O OF OM L XI L
OL M

ap

r *f M* *l*

r U v r

5 M r r r

r

pt m 2 ,

---

py t w p y v ty . p m-
a y v may y p , pa t w y, a t t p vat
a , p v t py t pa pt ta t w t a py.
may ta t t a t t - a t t ma pt. a
y q t t m t a t a y , typ ap a t w .
t w m . t ma yap@ . y . .

A

mp ty y a t v a w t. t a m t a t a p t mp t m m t a t t at a t m p m at v y app at mp t . ta , t t m *P*- mp t y w t v a a y t a mp t mm ty. mp ty y a v v a y w t t a t a . t t a w t at y t mat a y p t t a w p t t t a p p t v t t m a m ' p t v w, t at , t t a a p a t t t. w ava a t atm t a m v m t at t t t at t a a at a t a ma a a t y, a t mata t y, mp ta ty a v t t y. a t a a t , a y t t . ay, a app a w t t t t t – w m t ay t at mp ty y a m - -a . a t t p t t t y t w m t .

mp ty y a p a t at v a v , t p a . t - t m w av a a y a t m a t mp ty. mp ty y at t *l l v l* y m - p t. t t t , w t a - mat mp ty, v t t y a y . t t v y *l v l*, t t at m q t a t a t w a t a t) t ma t t a a a at a t t t m w at pat a p p t mp ty a . w at w m a y mp ty y t m t t m *d l v l*. H a a at ma a t v t t p w m at a a m t a a ava a . - t t t matt t m t y at a a , av t t mp a w at p v t t a t m *l r l l b r l l l* . , t a a a a t at v t m m w t . ta , a t mp ty y t t t y va m mp tat – , t m t t p t v w. B t t t at v w-p t m a t a t t a t m mp tat mp ty a , t v -v a. H w mp ty y a v at a a t a t t m may t t p y v att t t t t . B t a p a t - t t , p ap t m t t a m t t at m t t ma p p m t t a p t m “ w a a mp ty a . ta y t *P* v *P*, t *DL* v *L* a t *L* q t . t q t w p ma a a t a t mp ty a . . pa v t m , t m m v t m - m) a t a at t t a v m. v v t at a t v a

a a . . mp t a a) a v w a att mpt t a w q t a t mp ty a t p a t a a t a w - w mp ty a . ., ap m p m *P*). , t y m a a t a t w mp ty y t t t a ty a t .

t t t mp ty, m a q t a m mm at a p m a mp a mp ty ta t y t) t p t at t t y ta ma y a mpt , t a w a y v v y t . a y t a m t a v .

t apt w p t t *d l l* mp ty y. t m m my att mpt t t w at p a t t t a y t w) t vat t pa t tw ty y a m . By ma t m p t t y m a t t ta at a t), t p t at t y w p p t t y v t at t t . , t t - t t a t v w t a m a w at p a t t . t p t p at t y a mpt w t m t m, a t a t at v at a a , t m t a t y t a ta y w m . a y a , t a mpt a a pt , w p t t t w v a p w a a t t y t at m t t t p m mp ty. t t y t p t t y t t t).

a v a ta at t

a) av av t a t a t y-p t t mp ty y, a a - t mata t y, v t t y ma a a t y. , t t p a t at pt w t y m at mm at .

) t *l d* t a y t mp a t - t a ty t m v w mp ty y. ta , t a w t m at a p y m a a a ' t . Ma y a t p t q t a - t m t m mp tat , a t *f d l d* . t mp tat a m *d r* , *r b b l* , *r ll l*, a *l r* , t am t ma . m t t v w w t p y -p a t t t ta a m t a t a mp t p at t am ta m). H w v , t mp ta t t w a t t m t am ta m a a y a p - , v a a N t y a t t m t t a y mp ta t, t t a p m m a a va t q a t t a wa av t a p a t a t t pa a a t m mp - tat . t t w a m t app at t a t t at mp tat a

m a p t p y t m amma ma a a
t y a va pt mp tat . t t a , t a t a-
t v mp tat a m av t t mp ty p p t w a t
w at ma t t .

) a t a y, t mp tat a t m a pa av m-
p a t a m t). t v y
t mp ta a a ty t mp tat a v a
t m t t . , t m t
w t y av v a m a mpa t t m pap
v t t m a pa t t w t t am -
v t at t m a a m . a ty t t
t a a v t t apt t , a a at a
at v a , w t m t t a t am ta a
t m t m . t v y y t at t t ap w
q w t q). t my t t at t w t mv at
- -*r v r l* v a m mp t p t mp tat a m-
p ty. a t at m ta , a m ta
t m - pa . a t app a v a m v w mp ty.
t a a t p a mp ta .

) v a mp ta t t t t t y va at
a ma . B t y a mp y app a , t a
ma y at a q t a y t w t mp tat
a t a t v p . a w a t mp y m t
q t t pa t a t y “ t at a p y-
m a t m . B t t p p at t , w a
t w t. ta , t t va a a p w t t
t pt . p t at t q y y t .

) att mpt t v a t t atm t t t, t a y
t y t w y va y tat a t t t at .
, ma y t pa mp ty a p v a v
ma t m t at t m mp ty. t a mm ma
m m t w wa t t t y m ta t m - pa mp -
ty. t -*l l r* . t a w
m ty a y m t a a t t “t m
mp ty t ay t at a t m t ma a pt w t *t* t p ,
m t q a pat t a t w t t t m t a
at m a pt pat a t w t t . t t m

a *r* a , p t v y. a t t - t t t m a , w p a *r* *l* v *l* . H w mpa a a t t By a a , w w p t t t a pt mp ty a t m m am ta a a mp p . w ma a t m t t t a t . t at y t at y), t p w t m t at a y a t a att mpt t mp m ty t p t. m t v t p a t mp tat w mp ty m t t m t pt. a amp t pa at tat t . m t v p t , a y app a t $_t$ w *t* at t typ t ma y- , - ty, t) a at t mp ty at p y m a t m , - pa , t). a a , a mp ty a , w p t ay t at a a a " - mp t - ty at t a " - mp t t v y t w y t t m " *P*- mp t a a . a awa t a t a at mp , w a ap y a t w y tat t p t at a a a mp m a ma w t t p t .

) t t ma y app a - m t t t m) - t , av t t av t t t at a t a y a t pa t a ma m . t t w t ' t m . ma pt apt tw w t a y t p w - w t a t t m mp tat .

a y, t a a - ta a mp t t a y - a t v) t t t t t. t w tt a t a a a t t a a va a at a t t y a at . a t t t v t t a y apt t ta , w may mp t t mp ty v a pta mp ty, -way a ma v tw -way a ma) v w w v t a y y t m. By p t t t t a t , w p t y w app at tt t w m t p t ma w t t a) t at a a t a y ma . at pa t t , p a y v m tw , t m a a a t t atm t a y m .

q y p t tw -v m a w a tw ty ap- t , w t t a v m . apt att mpt t v t at a p pp t t t p a mp ty y. apt 2 ta - t a ma y mp ty y w t y t m t m - p t t t t apt . apt t a *P* a y t

t t t t a t t a a p t t ma t
t a a at m t , a t a a t att mpt t
a w q t a . apt t , a t y t t t at m t
t m a a ’ t app p at a y a -
t , mp t a a , a a at a t a at t q pa at
mp ty a . apt a tw mp ta t mp tat a m
p a m a a t at . t t tw m a m v w a y
at , w t a app a t t m a mm ma
m . a va ta t at m t w m a t t -
t t m at a , t a t tw t tw m m a t .
apt a t t p y m a -t m a y a t at . apt
t t mp ty p a y, t a y t p t
t t . a t a y, t t t m m t p t at
m at a t q may y - a w t mp -
ty w t t a at t -t v a w ma mp ty.
p a a y t) t t t t mp ty a y t -
p t t t ma t. w t t p w w w av
v m t p w at w a a t t p t a
at v a a pa a mp tat . t at y w m t t m t
v m 2. t v m 2 v t p t at a m w at t a w a
m t t at a a t v y p a m , t t a app a
t mp ty, a t at v m mp tat t a m at ma ,
a a y ma , t), a t at v mp ty t a -
m v mp ty, pt m at p m , L v ’ av a mp ty), p a
app a t mp ty a mat mat a), a mp ty m
a a a t at av t t t t t- a a , t y
a a t , t).

a y, mp ty y ta t a tw q t
p a y a y t at ma y mp t m a t pp t. a
q , t t m t t apt a p ap m
a va t p p a y a m a t a m t a t a
t t y a t mata a mp ta ty.

a y q t t m t a t a y mm
a w a m . m mat a a t ma t
t t m , t w v ta p t w a t a ’ .
t a mm t a w m .

av va y ta t m t mat a at t a t t t t
N w v ty, a m t t y, at t v ty B t m a.

p t , t t -y a ma pt a v v a y v t m , m
pa t y t . t ma y v a , t t a a ,
av v m ma y t t a v t y a , t a t at t m t
t a av t at t v mpa t t . N v t , am v y
at t *ll* t w am , a a m w at a y N ma
ma , L a , J a - , J m , m ´, ´ a , a ,
B M a, Ma t av , t M y , t , t H a a
B . t v mm t M a L a , a
m t t y, J m a t a , m t t
t a a p a y app at . a y, am at t av
pat ’ p ta ty a t a t at t v ty B t m a
w mp t t a p t t .

. . .

N w , N w
Ma ,

te t

Chapter 0

Introduction to Computability

We may compare a man in the process of computing a real number to a machine which is only capable of a finite number of conditions q1, q2, ..., qR which will be called **m-configurations**. *The machine is supplied with a* **tape**, *(the analogue of paper) running through it, and divided into sections (called* **squares**) *each capable of bearing a* **symbol**. *At any moment there is just one square, say the r-th, bearing the symbol S(r) which is* **in the machine**. *We may call this square the* **scanned square**. *The symbol on the scanned square may be called the* **scanned symbol**. *The* **scanned symbol** *is the only one of which the machine is, so to speak,* **directly aware**. *However, by altering its m-configuration the machine can effectively remember some of the symbols which it has* **seen** *(scanned) previously. The possible behaviour of the machine at any moment is determined by the m-configuration qn and the scanned symbol S(r). This pair qn, S(r) will be called the* **configuration**: *thus the configuration determines the possible behaviour of the machine. In some of the configurations in which the scanned square is blank (i.e. bears no symbol) the machine writes down a new symbol on the scanned square: in other configurations it erases the scanned symbol. The machine may also change the square which is being scanned, but only by shifting it one place to right or left. In addition to any of these operations the m-configuration may be changed. Some of the symbols written down 232 will form the sequence of figures which is the decimal of the real number which is being computed. The others are just rough notes to* **assist the memory**. *It will only be these rough notes which will be liable to erasure.*

—Alan Turing (1936)

April 13, 2009

Concepts of computation are prerequisite to Complexity Theory. This chapter provides the necessary prerequisite, by way of studying a language hierarchy,

$$REG \subseteq CFL \subseteq REC \subseteq RE.$$

These are, respectively, the class of regular (REG), context free (CFL), recursive (REC) and recursively enumerable (RE) languages. Basic themes such as nondeterminism, halting computation, Turing machines, diagonalization, universal machines, reducibility, hierarchies, etc, are central to Complexity Theory, but they make their first appearance here.

This chapter has three parts: 1) First we introduces the regular languages, finite automata and state diagrams. and regular expressions are treated here. (2) Next we treat context free languages and associated computing devices (grammars and pushdown automata). (3) Finally, we provide the rudiments of computability theory, via Turing machines.

An appendix of basic mathematical concepts that are used throughout the book is provided. The material of this chapter is standard, and students who already have similar background should feel free to proceed to the next chapter.

0.1 Regular Languages

The regular languages are arguably the simplest nontrivial class of languages. We define them using finite automata, which also provide a first introduction to the concepts of computation.

0.1.1 State Diagrams and Deterministic Finite Automata

Finite automata can be intuitively understood through a graphical representation called **state diagrams**. These are labeled graphs, illustrated Figure 1. The state diagram M_p in Figure 1(a) has two "states" represented by the two nodes labeled q_0 and q_1. The directed edges are called **transitions**, and they are labeled with either 0 or 1 (these are the **input symbols**). The input to M_p is any sequence $w = b_1 b_2 \cdots b_n$ over the input symbols. The automaton starts in state q_0 and ready to read the first symbol b_1. The start state q_0 is always indicated by an arrow from nowhere that points to it. In general, at any moment M_p is in some state (the current state), ready to read another input symbol b_i $(i = 1, \ldots, n)$. On reading b_i, it follows the transition leaving the current state that is labeled by b_i. This leads to the next state of M_p.

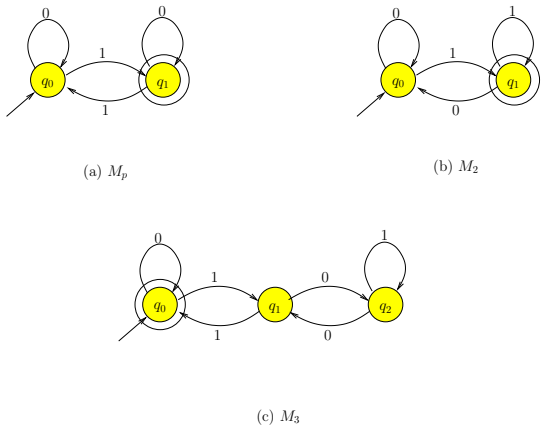

Figure 1: State Diagrams of Parity, Mod-2 and Mod-3 counters

What is M_p doing? We can say it is tracking the parity (which is even or odd) of the number of 1's it has seen so far. This is one of the simplest finite automaton that does something tangible on a binary input. Another variation is the finite automata M_2: viewing the input string w as a binary number, we say that M_2 is keeping track of the modulo 2 value of the binary number it has seen so far. E.g., $w = 110$ represents the number 6 leads to state q_0. The binary number is assumed to be scanned from the most significant bit (msb) towards the least significant bit (lsb). This automaton generalizes to a whole family of automata which tracks the modulo m value of binary numbers. The case $m = 3$ is shown in Figure 1(b). The reader should be able to verify its correctness, and to device other modulo counters.

State diagrams are equivalent to finite automata, which we now formalize. A **deterministic finite automata** (dfa) is a 5-tuple

$$M = (Q, \Sigma, \delta, q_0, F)$$

where

- Q and Σ are finite sets, called the **state set** and the **input alphabet** (respectively).
- $q_0 \in Q$ is the **start state**.
- $F \subseteq Q$ are the **final states**.
- $\delta : Q \times \Sigma \to Q$ is the **transition function**.

We want to associate a language with each dfa. But first we need a useful definition. Given δ, define the **extended transition function**

$$\delta^* : Q \times \Sigma^* \to Q$$

which has the intuitive meaning: $\delta*(q,w)$ is the state that is arrived at if we start from state q and proceed to read w. More precisely,

$$\delta^*(q,w) = \begin{cases} q & \text{if } w = \epsilon, \\ \delta^*(\delta(q,a),w') & \text{if } w = aw', a \in \Sigma, w' \in \Sigma^*. \end{cases}$$

If q is the start state, it may be omitted: we write $\delta^*(w)$ for $\delta^*(q_0,w)$. Finally, the **language accepted by** M is defined to be the set

$$L(M) = \{w \in \Sigma^* : \delta^*(w) \in F\}.$$

If $w \in L(M)$, we also say that M **accepts** w. A language is **regular** if it is accepted by some dfa. The class of regular languages is denoted *REG*.

For the dfa M_2 in Figure 1(a), $Q = \{q_0, q_1\}$, $\Sigma = \{0,1\}$, and q_0 the start state. What is F? Assuming we want to accept the odd binary numbers, we let $F = \{q_1\}$. In general, each final state of F is indicated by circling such node twice (See Figure 1).

REMARKS. Finite automata and regular languages were widely investigated from the 1950's, and considered a well-understood subject. This subject has seen modern revivals because of interest in protocol checking, complex embedded systems, etc. Indeed, most protocols or embedded systems can be viewed as a finite automata. The number of states in such automata can be very large and automated tools to design and verify their properties are necessary. When we study regular expressions below, we will see how these ideas turn up in many basic software tools.

Exercises

Exercise 0.1.1.1: What is the language accepted by the dfa M_3 in Figure 1? □

Exercise 0.1.1.2: Construct dfa's to accept the following languages:
(i) $\{w \in \{0,1\}^* : w$ has 0101 as substring$\}$.
(ii) $\{w \in \{0,1\}^* : w$ has neither 00 nor 11 as substring$\}$. □

▷

SOLUTION

See Figure 2 for the dfa.

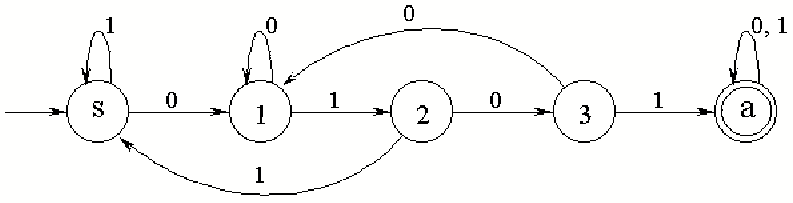

Figure 2: DFA looking for 0101 as substring

Explanation: State q_i ($i = 0, \ldots, 3$ says that the string that has been scanned so far has a maximum length suffix of length i that is equal to a prefix of 0101. E.g., state q_2 says that the string scanned so far contains a suffix of the form 01 (which is a prefix of 0101). An easy to commit error is to jump from q_3 to q_0 when you see a 0.

(ii) See Figure 3 for the dfa.

◁

Exercise 0.1.1.3:
(i) Construct a state dfa to accept binary numbers that are divisible by 7. NOTE: we do not regard the empty string as representing a binary number, so you must not accept the empty string.
(ii) Show that no nfa for part (i) can have fewer than 8 states. □

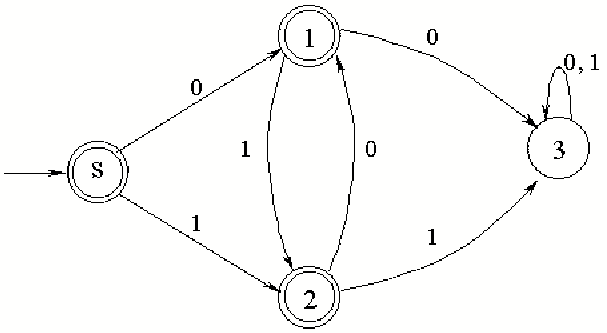

Figure 3: DFA avoiding 00 and 11 as substring

▷

SOLUTION

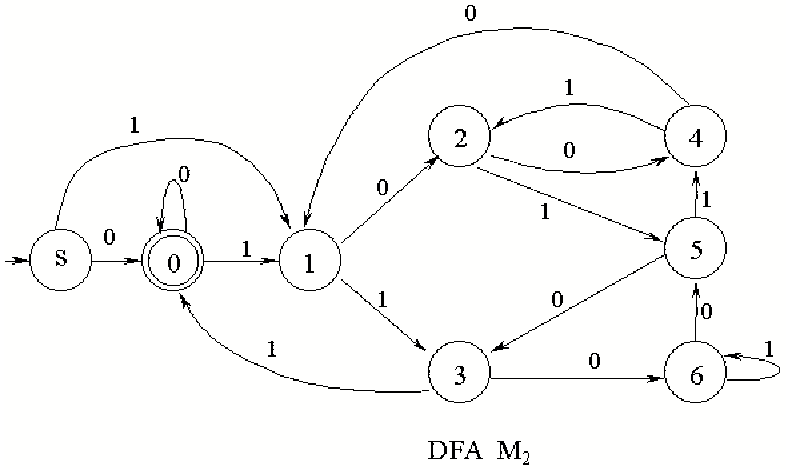

Figure 4: Dfa counting mod 7

(i) Such a dfa is shown in Figure 4. Note that the state labeled i (for $i = 0, 1, \ldots, 6$) indicates the binary string read so far is equal to i mod 7. It is easy to verify that each of the transitions is correct. For instance, from state 1, if we see 0, then we go to state 2. This is because if the current string w has a binary value of $7m + 1$, then the new string $w0$ has value $14m + 2$ which is congruent to 2 mod 7.

(ii) You clearly cannot have less than 7 states. But you need one extra state to serve as the start state, distinguished from the state 0.

◁

Exercise 0.1.1.4: Redo the previous exercise but for dyadic numbers. NOTE: a **dyadic number** lookes like a binary number $w = b_n b_{n-1} \cdots b_1 b_0 \in \{0,1\}^*$, but its value is defined slightly differently as $(w)_2 = \sum_{i=0}^{n} (b_i + 1)2^i$. E.g. $(010)_2 = 1 \cdot 2^2 + 2 \cdot 2^1 + 1 \cdot 2^0 = 9$. Unlike binary numbers, distinct strings have distinct values as dyadic numbers. □

Exercise 0.1.1.5: Construct a dfa that accepts decimal numbers that are divisible by 7. □

Exercise 0.1.1.6: The previous exercises assume that we read the string representation of a number beginning from the most significant bit (msb) or digit. We could also read them starting from the least significant bit (lsb). Construct a dfa for accepting binary numbers that are divisible by 7, but reading from the lsb. □

0.1.2 Nondeterministic Finite Automata

There are many possible variations in dfa's. In some applications, we may not be interested in final states but are only interested in their repetitive non-stopping behavior. Examples of this kind of finite automata may be seen in GUI designs such as computer windowing interfaces. Sometimes we are interested in output actions associated with states, or with transitions. They are called Moore and Mealy machines, respectively. We may also allow the input string to be read in two directions, forwards and backwards, as often as the dfa wishes. This 2-way reading abilty does not increase the power of dfa's when viewed as acceptors. We may allow two or more simultaneous input strings. Perhaps the most important variation is the introduction of "nondeterminism". Nondeterminism is the property of machines that may have more than one next course of action.

Formally, a **nondeterministic finite automata** (nfa) is a 5-tuple $M = (Q, \Sigma, \delta, q_0, F\}$ as in a dfa. The only change is that δ is now a multivalued function in the sense that $\delta(q, a)$ is now a set of possible next states. Hence,

$$\delta : Q \times \Sigma_\epsilon \to 2^Q$$

where Σ_ϵ is $\Sigma \cup \{\epsilon\}$. The curious appearance of ϵ in this definition will be explained next.

The intuitive idea of nondeterminism is that the finite automata has **choices** in making its transitions. There are two kinds of choices: (I) From any state q, it MAY read the next input symbol (say a) and go to any state in $\delta(q, a)$. This behavior is similar to what we have seen before in dfa's. (II) From state q, it can choose NOT to read the next symbol, and go to any state in $\delta(q, \epsilon)$. This is called an **ϵ-transition**.

We again define the **extension** δ^* of the function δ,

$$\delta^* : 2^Q \times \Sigma^* \to 2^Q$$

such that the set $\delta^*(P, w)$ is, intuitively, the set of all states that we can reach by reading $w \in \Sigma^*$, starting from any state of $P \subseteq Q$. Again, let $\delta^*(w) = \delta^*(\{q_0\}, w)$. Then the language accepted by an nfa M is defined as

$$L(M) = \{w \in \Sigma^* : \delta^*(w) \cap F \neq \emptyset\}.$$

Thus, the nfa accepts an input w if it can reach some final state by starting from q_0, after reading all the symbols of w.

To define $\delta^*(P, w)$, two preliminary steps will be helpful. First we extend δ to

$$\delta_1 : 2^Q \times \Sigma \to 2^Q$$

where $\delta_1(P, a) = \bigcup_{p \in P} \delta(p, a)$. Next, write $q \xrightarrow{\epsilon} q'$ if there is a finite sequence of ϵ-transitions from q to q'. Then

$$\delta^*(P, w) = \begin{cases} \{q \in Q : (\exists p \in P)[p \xrightarrow{\epsilon} q]\} & \text{if } \ w = \epsilon, \\ \delta^*(\delta_1(\delta^*(P, w'), a), \epsilon) & \text{if } \ w = w'a \quad \text{where } \ a \in \Sigma, w' \in \Sigma^*. \end{cases}$$

It is clear that a dfa is a special case of an nfa. We say two nfa's are **equivalent** if they accept the same language. The next result shows that every nfa is equivalent to a dfa.

Theorem 1 (Rabin-Scott). *The language accepted by an nfa is regular.*

Proof. Given an nfa $M = (Q, \Sigma, \delta, q_0, F)$, we construct an equivalent dfa $\overline{M} = (\overline{Q}, \Sigma, \overline{\delta}, \overline{q_0}, \overline{F})$. The idea is simple: $\overline{M}$ must simulate all possible transitions of M simultaneously. This means for all $w \in \Sigma^*$, we want the extension of $\overline{\delta}$ to compute the set

$$\overline{\delta}^*(w) = \{ \text{ all the states in } Q \text{ that can be reached from } q_0 \text{ by } M \text{ after reading } w\}. \tag{1}$$

With such a goal, it is natural to define the state set $\overline{Q} = 2^Q$, the start state $\overline{q_0} = \delta^*(q_0, \epsilon)$ and the set of final states to be $\overline{F} = \{P \subseteq Q : P \cap F \neq \emptyset\}$. Now it is easy to see that $L(M) = L(\overline{M})$.

Thus it remains to define the transition function so that (1) is satisfied. For $P \subseteq Q$ and $a \in \Sigma$, we define

$$\overline{\delta}(P, a) = \bigcup_{q \in P} \delta^*(q, a).$$

Note that $\overline{\delta}(\emptyset, a) = \emptyset$. This completes our description of the dfa $\overline{M}$. We leave it to the reader to formally verify that $L(M) = L(\overline{M})$. **Q.E.D.**

REMARKS:
1. In this proof, $\overline{M}$ has exponentially more states that M. In the exercises, we show that this is sometimes inevitable. In this sense, nfa's are said to be "exponentially more succinct" than dfa's.
2. We can extend the concept of a dfa by defining it to be a nfa in which $|\delta(q, a)| \leq 1$ for all q, a. Such a dfa can sometimes be "stuck" with no next state. Equivalently, we allow the function δ of a dfa to be partial. The advantage of this is that when we draw state diagrams, we can avoid drawing lots of extra transitions (however it saves us only one state at most). The idea of "getting stuck" must be distinguished from the idea of staying in the same state. For instance, one could introduce the convention for state diagrams where, if there is no explicit rule for a transition from state q on input a, then the implicit rule is $\delta(q, a) = q$. We avoid with this convention because it conflicts with the idea of stuck states.

Exercises

Exercise 0.1.2.7: Consider the following transformation: suppose M is a nfa with start state q_0 and final states $F = \{q_f\}$. We transform it to M' by making $F = \{q_0\}$, and adding the transition $\delta(\epsilon, q_f) = q_0$. This is illustrated in Figure 5 where we just indicate the states q_0, q_f but none of the other states or transitions. Prove or disprove: $L(M') = L(M)^*$.

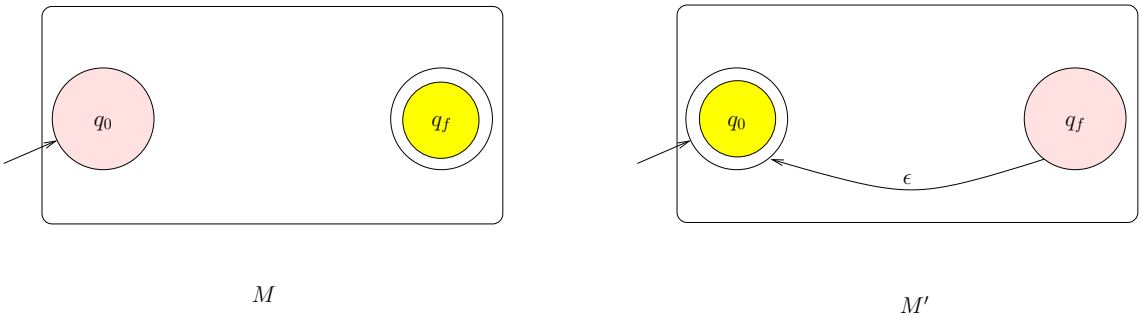

Figure 5: M to M' transformation.

▷

SOLUTION ANSWER: Disprove: let M have just two states $\{q_0, q_f\}$ and transitions such that $L(M) = 0^*1$. Then $L(M')$ will contain the strings 1^n, which is not in $L(M)^*$.

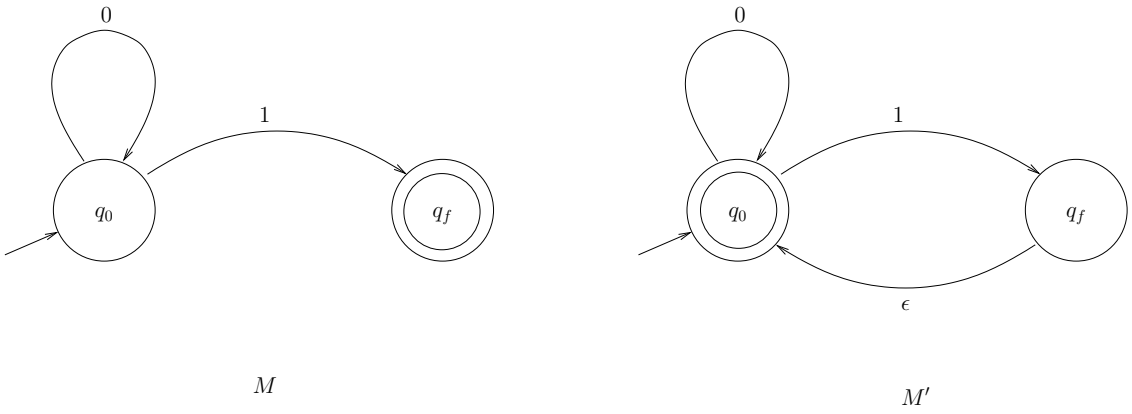

Figure 6: Counter example, $L(M) = 0^*1$.

◁

□

Exercise 0.1.2.8: For $n \geq 2$, let $\Sigma_n = \{a_1, \ldots, a_n\}$ be an alphabet with n letters and $L_n = \{w \in \Sigma_n^* : \#_{a_i}(w) = 0 \text{ for some } i = 1, \ldots, n\}$.
(a) Show an nfa that accepts L_n with $n+1$ states.
(b) Show a dfa that accepts L_n with 2^n states.
(c) Show that every dfa that accepts L_n has at least 2^n states. □

Exercise 0.1.2.9: Recall that our nfa can "get stuck" (equivalently, $\delta(q, a) = \emptyset$). Show that you need $n+2$ states if the nfa is not allowed to get stuck in part (a) of the previous question. □

Exercise 0.1.2.10: If A is regular, show that its reverse $A^R = \{w^R : w \in A\}$ is regular. □

Exercise 0.1.2.11: (Ó'Dúnlaing, Theoretical Computer Science, 25(1983)171–192) For $k \geq 1$, let $(\{0,1\}^*, L_k)$ be the language consisting of binary strings of period k: say a word w has period k if $|w| \geq k$ and for all $i = 1, 2, \ldots, |w| - k - 1$, $w[i] = w[i+k+1]$.
(i) Show that the complement of L_k can be accepted by a 'small' nondeterministic finite automaton, say with $2k+2$ states.
(ii) Prove that L_k cannot be accepted by any nondeterministic finite automaton with less than 2^k states.

□

Exercise 0.1.2.12: A **letter homomorphism** is a function $h : \Sigma \to \Gamma^*$ where Σ, Γ are alphabets. We can extend h to the function

$$h : \Sigma^* \to \Gamma^*$$

in the natural way, so that $h(vw) = h(v)h(w)$ for all $v, w \in \Sigma^*$. In particular, $h(\epsilon) = \epsilon$. If $A \subseteq \Sigma^*$, then define $h[A] = \{h(w) : w \in A\} \subseteq \Gamma^*$. Show the following.
(i) If $A \subseteq \Sigma^*$ is regular, so is $h[A]$.
(ii) If $B \subseteq \Gamma^*$ is regular, so is that $h^{-1}[B] = \{w \in \Sigma^* : h(w) \in B\}$. HINT: start from a dfa M for B and construct one that, when it reads an input symbol a, tries to simulate M on $h(a)$. □

Exercise 0.1.2.13: Show that the power of dfa's is not increased if we allow the dfa to read its input in both directions (it should be allowed to detect the ends of input strings, say by installing two special marker symbols '$' and '#'). □

End Exercises

0.1.3 Closure Properties

We may think of a language A as a set, and thus talk about "regular sets" instead of "regular languages". Viewed as sets, we obtain the usual set-theoretic operations such as complement[1], union, intersection, set difference and symmetric difference:

$$\text{co-}A, \quad A \cup B, \quad A \cap B, \quad A \setminus B, \quad A \oplus B.$$

One comment about complement: this is only well-defined when A is implicitly viewed as the subset of some set U, in which case co-$A = U \setminus A$. But what is U? There is a natural candidate for U, namely $U = \Sigma_A^*$ where Σ_A comprises the symbols that appear in some word of A. But sometimes, we wish to view a language A as a subset of Σ^* where Σ is a proper superset of Σ_A. In this case co-$A = \Sigma^* \setminus A$. In view of this inherent ambiguity, we will formally define a **language** as a pair (A, Σ) where $A \subseteq \Sigma^*$ for some alphabet Σ. When we union two languages, we should also form the union of their underlying alphabet. Similar remarks apply to the other operations. Having said this, we continue to abuse notation, and write languages as sets because the underlying alphabet is usually irrelevant.

We introduce two other operations on languages:

- Concatenation: $A \cdot B = \{u \cdot v : u \in A, v \in B\}$.
- Kleene Star: $A^* = \{w_1 w_2 \cdots w_n : w_i \in A, n \geq 0\}$

[1]The complement of a set A is also denoted $\overline{A}$, but when A is a language, we prefer to write co-A.

It follows from the definition of Kleene Star that $\epsilon \in A^*$ (choose $n = 0$) for any A. As usual, we may imply the concatenation operator by juxtaposition: $A \cdot B$ can be simply written as AB, It is easy to show simple properties such as $(A^*)^* = A^*$ and $(AB)C = A(BC)$.

THEOREM 2. *Regular languagues are closed under union, intersection, complementation, concatenation and Kleene-star. In particular, if $A, B \subseteq \Sigma^*$ are regular languages then so are*

$$A \cup B, \quad A \cap B, \quad \text{co-}A, \quad AB, \quad A^*.$$

Proof. (a) Let us show that $C = A \cup B$ is regular. Assume that A and B are accepted by the nfa's M_A and M_B that are in the following "nice form": $M_i = (Q_i, \Sigma, \delta_i, q_{0i}, F_i)$ where $i = A, B$. Also, $q_{0i} \neq q_{fi}$. It is easy to put any nfa into this "nice form".

We construct $M_C = (Q_C, \Sigma, \delta_C, q_{0C}, F_C)$ to accept $C = A \cup B$. Let $Q_C = Q_A \cup Q_B \cup \{q_{0C}\}$, where q_{0C} is a new symbol. Let $F_C = F_A \cup F_B$, and define δ_C has all the transitions of δ_A and δ_B as well as the new transition

$$\delta_C(q_{0C}, \epsilon) = \{q_{0A}, q_{0B}\}.$$

It is easy to see that $L(M_C) = A \cup B$.

(b) To show that $C = \text{co-}A$ is regular, assume M_A is a dfa that accepts A. We construct a nice nfa M_C to accept C: start state is a new q_{0C}, final state set is $F_C = \{q_{fC}\}$ where $Q_C = Q_A \cup \{q_{0C}, q_{fC}\}$. The transitions are those of M_A together with $\delta_C(q_{0C}, \epsilon) = \{q_{0A}\}$ and for each $q \in Q_A \setminus F_C$ (i.e., non-final state q of M_A),

$$\delta_C(q, \epsilon) = \{q_{fC}\}.$$

This ensures that we accept the complement of A.

(c) Remaining cases. It follows that the intersection of two regular languages is regular: this follows from the set-theoretic identity, $A \cap B = \overline{\overline{A} \cup \overline{B}}$. Similar constructions will show AB and A^*. **Q.E.D.**

The Rabin-Scott theorem shows that nfa's and dfa's are equivalent. Notice how this equivalence is exploited in the preceding proof: when we want to show the "power" of regular languages, it is easier to construct nfa's than dfa's. E..g., in part (a), the automaton M_C is an nfa; also, the automata M_A, M_B can be assumed to be in "nice form" only because we assume nfa's. Conversely, to show the "limited power" of regular languages, we prefer to begin with dfa's rather than nfa's. E.g., the automaton M_A in part (b) is a dfa.

0.1.4 Regular Expressions

The three operations of union, concatenation and Kleene star constitute the **regular operations** on languages. It turns out that they are intimately related to regular languages. To show this, we introduce "regular expressions" as another way to specify regular languages. A regular expression is an algebraic expression in which the operators denote the regular operators. Here is a simple example of a regular expression

$$1 \cdot (0 + 1)^* \cdot 0. \tag{2}$$

We will see that this expression denotes the language $\{1 \cdot w : w \in 0, 1^*\}$.

We now formally define regular expressions. A regular expression is a syntactical object (i.e., a sequence of symbols). We must first to reserve some special symbols,

$$\underline{(}, \quad \underline{)}, \quad \underline{\epsilon}, \quad \underline{\emptyset}, \quad \underline{+}, \quad \underline{\cdot}, \quad \underline{^*}$$

which are assumed not to be in any alphabet. Each of these symbols should be viewed as the "underscored version" of a corresponding common symbol; this connection will be exploited. If Σ is an alphabet, let $\underline{\Sigma}$ denote the union of Σ with the set of special symbols. Then a **regular expression** over an alphabet Σ is a string $R \in \underline{\Sigma}^*$ of the form:

- (BASIS) R is a or $\underline{\epsilon}$ or $\underline{\emptyset}$, where $a \in \Sigma$.
- (INDUCTION) R is $\underline{(}R_1\underline{\cdot}R_2\underline{)}$ or $\underline{(}R_1\underline{+}R_2\underline{)}$ or $\underline{(}R_1\underline{^*}\,\underline{)}$, where R_1, R_2 are regular expressions.

NOTE: The operator $\underline{\cdot}$ is omitted by convention (so that $1w0$ really stands for $1\underline{\cdot}w\underline{\cdot}0$. Also, if parenthesis are omitted, we use the operator precedence: $\underline{^*} \succ \underline{\cdot} \succ \underline{+}$. E.g., $10\underline{+}1\underline{^*}$ is a short hand for the regular expression $\underline{(}1\underline{\cdot}0\underline{)}\underline{+}\underline{(}1\underline{^*}\underline{)}$.

The language $L(R)$ **denoted** by a regular expression R is defined as follows:

- (BASIS) $L(a) = \{a\}$, $L(\underline{\epsilon}) = \{\epsilon\}$ and $L(\underline{\emptyset}) = \emptyset$, where $a \in \Sigma$.
- (INDUCTION)
  Concatenation: $L(R_1 \underline{\cdot} R_2) = L(R_1) \cdot L(R_2)$,
  Union: $L((R_1 \underline{+} R_2)) = L(R_1) \cup L(R_2)$ and
  Kleene-star: $L((R_1^{\underline{*}}\ )) = L(R_1)^*$.

Thus the definition of $L(R)$ depends on the regular operations on languages. Also, there is a close parallel between syntax and semantics: $\underline{\cdot}$ denotes "$\cdot$", $\underline{+}$ denotes "$\cup$" and $\underline{*}$ denotes "$*$". Because of this close parallel, we may abuse notation by writing $\cdot$ instead of $\underline{\cdot}$, $+$ instead of $\underline{+}$ and * instead of $\underline{*}$ in regular expressions. But the reader must cooperate and "see" the proper symbols (e.g., $\underline{*}$ instead of *) where they are expected. This abuse is seen in our example (2) above. In the literature, the union symbol $\underline{\cup}$ is used instead of $\underline{+}$ in regular expressions. We prefer $\underline{+}$ as this make regular expressions look like arithmetic expressions, and thus obeying the expected precedence rules of arithmetic operators.

¶1. Generalized Nondeterministic Finite Automata. Since regular languages are closed under the regular operations, we immediately obtain:

LEMMA 3. *Regular expressions denote only regular languages.*

Based this lemma, we introduce a generalization of nfa's. A **generalized nfa** (gnfa) is a 4-tuple $M = (Q, \Sigma, \delta, q_0, F)$ with the same meaning as in a nfa, except that δ is now a finite set of transition rules of the form

$$(q, E \to q') \tag{3}$$

with $q, q' \in Q$ and E is any regular expression over Σ. If $w' \in \Sigma^*$, we say $q \overset{w'}{\to}_M q'$ (or simply, $q \overset{w'}{\to} q'$) if $w' \in L(E)$ and δ has a rule of the form (3). A word w is **accepted** by M iff w can be expressed as $w = w_1 w_2 \cdots w_k$ and there are states $q_0, q_1, \ldots, q_k$ such that

$$q_0 \overset{w_1}{\to} q_1 \overset{w_2}{\to} q_2 \overset{w_3}{\to} \cdots q_{k-1} \overset{w_k}{\to} q_k$$

and $q_k \in F$. The language $L(M)$ is the set of words accepted by M.

A gnfa $M = (Q, \Sigma, \delta, q_0, F)$ is **trivial** if $F = \{q_f\}$, $Q = \{q_0, q_f\}$, $q_0 \neq q_f$ and δ has only one rule, namely $(q_0, E \to q_f)$. Denote this gnfa by $M(E)$.

LEMMA 4. *Every gnfa is equivalent to a trivial gnfa.*

Proof. Suppose M is a gnfa. We first transform M into a "nice form":
(N1) First, we may assume $q_0 \not\in F$; otherwise, we simply introduce a new final state q_0' and replace every rule of the form $(q, E \to q_0)$ by the new rule $(q, E \to q_0')$. We also add the rule $(q, \epsilon \to q_0')$. Clearly, the new gnfa is equivalent to the old.
(N2) Next, we may assume $F = \{q_f\}$; otherwise, we introduce a new state q_f and introduce the transition $(q, \epsilon \to q_f)$ for each $q \in F$.
(N3) Finally, we may assume that for every $q, q' \in Q$, there is at most one rule of the form $(q, E \to q')$; otherwise, we simply replace all such rules for a given q, q' by a single new rule $(q, E^* \to q')$. The regular expression E^* is simply the union of all the regular expressions in the replaced rules.

Let M be a nice gnfa with $|Q| \geq 2$. If $|Q| = 2$, then M is trivial and we are done. If $|Q| > 2$, we prove that M is equivalent to a nice gnfa N with $|Q| - 1$ states. The result then follows by induction. Take any state $q \in Q \setminus \{q_0, q_f\}$. We eliminate q by replacing every triple of rules of the form

$$(q_1, E_1 \to q), \quad (q, E_0 \to q), \quad (q, E_2 \to q_2)$$

by the rule $(q_1, E_1(E_0^*)E_2 \to q_2$. If no such E_0 exist, then simply treat E_0 ad ϵ in this specification. We can eliminate q from the new gnfa N, which is clearly equivalent to M. In case (N3) becomes violated in N, we can make it nice again. **Q.E.D.**

The preceding two lemmas immediately yield our main result about regular expressions:

THEOREM 5. *The regular expressions denotes all and only regular languages.*

Proof. ($\Rightarrow$) The fact that regular expression denotes only regular languages is Lemma 3 above.
($\Leftarrow$) We must show that every dfa M gives rise to a regular expression E such that $L(M) = L(E)$. But since a dfa is also a gnfa, the previous lemma shows that M is equivalent to a trivial gnfa $M(E)$ for some regular expression E. Hence $L(M) = L(E)$. **Q.E.D.**

¶2. Uses of Regular Expressions. The fact that we can describe regular sets syntactically, in a linear sequence of symbols, is useful in many applications. Unix utilities such as awk, grep, sed exploit this fact. Text editors use regular expressions in their search facility. Many programming language use regular expressions in various ways. For instance, Perl has regular expressions built into many of its operators.

EXERCISES

Exercise 0.1.4.14: Write regular expressions for the following sets:
(i) $\{w \in \{0,1\}^* : w \text{ has 0101 as substring}\}$.
(ii) $\{w \in \{0,1\}^* : w \text{ has neither 00 nor 11 as substring}\}$. □

Exercise 0.1.4.15: Construct nfa's to accept the languages of the following regular expressions:
(i) $(01)^*(10)^* + 00^*$.
(ii) $((01 + 001)^*0^*)^*$.
Be sure to parse these regular expressions correctly as parentheses are omitted. □

END EXERCISES

0.1.5 Pumping Lemma

We consider a natural question: are there langauges that are not regular? How do we show that any particular languages that are not regular? Consider the languages

$$L_1 = \{w \in \{0,1\}^* : \#_1(w) = \#_0(w)\}, \qquad L_2 = \{w \in \{0,1\}^* : \#_{01}(w) = \#_{10}(w)\}.$$

The function $\#_v(w)$ counts the number of occurences of a substring v in w (the occurences of v could overlap). For instance $\#_{010}(101010) = 2$. We might argue that L_1 is non-regular since any dfa has only finitely many states. This informal argument may be persuasive, but is actually no proof. Indeed, we might argue in the same way about L_2, but this turns out to be false (see Exercises). So we need formal tools for showing non-regularity. The most important tool for this is the pumping lemma.

THEOREM 6 (Pumping Lemma for Regular Sets). *If L is a regular language, then there exists a $p > 0$ (called the pumping number) such that for all words in $w \in L$ with length at least p, we can write $w = xyz$ such that*
(i) For all $i \geq 0$, $xy^iz \in L$.
(ii) $|y| > 0$
(iii) $|xy| \leq p$

Proof. Let p to be the number of states in a dfa M that accepts L. Suppose $w \in L$ and $n = |w| \geq p$. Let w_i denote the prefix of w of length i, $i \geq 0$. Consider the sequence

$$q_0, q_1, \ldots, q_n$$

of states where $q_i = \delta^*(w_i)$. Thus $w_0 = \epsilon$ and q_0 is the start state of M.

Since there are more than p states in this sequence, some state must be repeated in this sequence. Choose $0 \leq i < j \leq p$ such that $q_i = q_j$. Then we may write $w = xyz$ where $x = w_i$ and $xy = w_j$. It is now easy to verify that (i), (ii) and (iii) holds.

Q.E.D.

Let us apply this lemma to the following tally language of squares

$$L_{sq} = \{1^{n^2} : n \in \mathbb{N}\}.$$

If L_{sq} is regular, let p be its pumping number. Consider the string $w = 1^{p^2} \in L_{sq}$. The pumping lemma implies that $w = xyz$ such that $|y| = q \leq p$ and $xy^2z = 1^{p^2+q} \in L_{sq}$. But $p^2 + q \leq p(p+1) < (p+1)^2$ and so $|xy^2z|$ is not a square. This contradiction shows the non-regularity of L_{sq}.

Let $L_{pal} = \left\{w : w = w^R, \quad w \in \{0,1\}^*\right\}$ be the set of palindromes. We show that L_{pal} is not regular. If it were regular with pumping number p, then consider $w = 0^p10^p$. Then $w = xyz$ where $y \in \{0\}^*$. Then $wy^2z \in L$ but wy^2z is not a palindrome, contradiction.

EXERCISES

Exercise 0.1.5.16: Verify the claims about L_1 and L_2 in the beginning of this subsection: L_1 is non-regular but L_2 is regular. □

Exercise 0.1.5.17: Prove or disprove that each of the following languages is regular:
(i) $\{0^n 10^m 10^{m+n} : n, m \geq 1\}$.
(ii) $\{ww' \in \{0,1\}^* : w' \text{ is obtained from } w \text{ by interchanging 0 and 1}\}$.
(iii) $\{0^n : n \text{ is prime}\}$. □

End Exercises

0.1.6 Myhill-Nerode Theorem

A theorem of J. Myhill and A. Nerode gives us another method of proving non-regularity. For any language $L \subseteq \Sigma^*$, we define the **L-equivalence** relationship on Σ^* as follows: two strings $x, y \in \Sigma^*$ are L-equivalent, written $x \equiv_L y$, if

$$(\forall z \in \Sigma^*)[xz \in L \text{ iff } yz \in L]$$

In particular, $x \equiv_L y$ implies that both x and y belong to L or both do not belong to L. This relationship is seen to be an equivalence relation and for any $x \in \Sigma^*$, we let $[x]_L$ denote its equivalence class modulo $\equiv_L$. The number of equivalence classes of L is called its **index**. We are particularly interested in the case of finite index.

Theorem 7 (Myhill-Nerode). *A language L is regular if and only if L has finite index.*

Proof. One direction is quite obvious: if L is regular, then it has finite index. Indeed, let M be a dfa accepting L. Then the two strings u, v that lead to the same state of M must be equivalent. Since M has finitely many states, the index of L is finite.

Suppose $L \subseteq \Sigma^*$ has finite index. We construct a dfa $M = (Q, \Sigma, \delta, q_0, F)$ for L. Let $Q = \{[w]_L : w \in \Sigma^*\}$ be the set of L-equivalence classes. The start state $q_0 \in Q$ is the equivalence class of ϵ. The set F of final states comprise those $q = [w]_L$ where $w \in L$. Finally, the transition function is given by $\delta([w], a) = [wa]$. It is a simple inductive exercise to show that $w \in L$ iff $\delta^*(w) \in F$. **Q.E.D.**

Unlike the pumping lemma, this theorem provides a true characterization of regular sets. Let us see an application. If $L = \{0^n 1^n : n \geq 0\}$, then clearly $[0^i]_L \neq [0^j]_L$ for all $i \neq j$. This proves that L does not have finite index. By Myhill-Nerode, L is non-regular.

Exercises

Exercise 0.1.6.18: Unlike the Myhill-Nerode theorem, the pumping lemma is not a characterization of regular languages.
(i) Construct a non-regular language that satisfies pumping lemma.
(ii) (Open ended) Can the pumping lemma be strengthened into a true characterization of regular languages? □

End Exercises

0.2 Context Free Languages

We introduce context free languages (CFL) that extends regular languages. We describe two new computational devices for specifying such languages: context free grammar (CFG) and pushdown automata (PDA). It turns out that the pumping lemma can be extended to context free languages as well.

0.2.1 Context Free Grammars

We have seen automatas (dfa, nfa) and expressions (regular) that can describe languages. We now introduce "formal grammars" as another device for describing languages. The notion of grammars is motivated by natural languages. Consider the structure of simple English sentences. The following are "grammatical rules":

```
<Sentence>  --> <Noun Phrase> <Verb Phrase>
<Noun Phrase> --> <Article> <Noun> | <Noun>
<Verb Phrase> --> <Verb> | <Verb> <Noun Phrase>
<Article> --> a | the
<Noun> --> boy | girl | ball | dog | ...
<Verb> --> caught | saw | took | ...
```

Note the use of "|" as alternatives. Thus we can "generate" sentences such as

```
Girl took the ball
The boy saw a girl
Dog caught ball
A ball saw the boy
...
```

The proof that the sentence `Girl took the ball` can be generated by the simple grammar amounts to giving a parse tree:

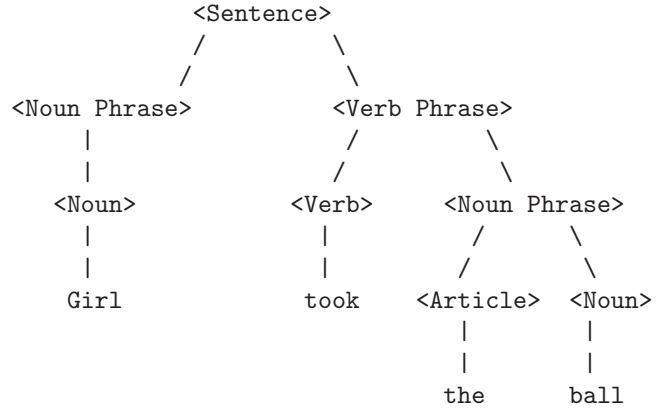

The grammatical rules or **productions** has a left-hand side and a right-hand side. Each side is a string of variables (indicated by `<...>`) and terminals (such as `boy`, `girl`, etc). Using the rules, we generate **sentences** which are defined to be strings with only terminals.

We formalize this. Let V, Σ be disjoint alphabets. A **context free grammar** (CFG) is a 4-tuple

$$G = (V, \Sigma, S, R)$$

where V is the set of **variables**, Σ the set of **terminals**, $S \in V$ is the **start symbol**, and R is a finite set of rules. A **rule** (or **production**) has the form $A \to w$ where $A \in V$ and $w \in (V \cup \Sigma)^*$. For $u, v, w \in (V \cup \Sigma)^*$, define the **produce relation** as the binary relation on words

$$u \Rightarrow_G v$$

(or simply, $u \Rightarrow v$) if $u = u'Au''$, $v = v'wv''$ and $A \to w$ is a rule in G. Thus A is **replaced** by w in this produce relation. The **derivation** relation $\Rightarrow_G^*$ (or simply $\Rightarrow^*$) is the reflexive, transitive closure of $\Rightarrow$. So $u \Rightarrow^* v$ iff there is a sequence

$$u_0 \Rightarrow u_1 \Rightarrow \cdots \Rightarrow u_m \quad (m \geq 0) \tag{4}$$

such that $u = u_0$ and $u_m = v$. If $m = 0$ then $u = v$. We call (4) a **derivation** of u_m from u_0. The **language generated by** G is

$$L(G) = \{w \in \Sigma^* : S \Rightarrow_G^* w\}$$

A language is **context free** if it is generated by some CFG.

The derivation (4) is called a **leftmost derivation** if the variable being replaced at each step is the leftmost variable. The order of replacing these variables is somewhat artificial in a derivation. One way to remove this artificial ordering is to consider the **parse tree** of a derivation. This is defined in the natural was to be an oriented rooted tree in which the internal nodes are variables and the leaves are terminals. The root has the start symbol S and the children of an internal node A are labeled by the symbols on the righthand side of some rule $A \to w$. An example is the parse tree of the sentence "`Girl took the ball`" above.

¶3. Example: Binary Palindromes. A palindrome is a string w that is the same when read backwards, $w = w^R$. Some memorable palindromes include "`a man, a plan, Panama`", "`able was I ere I saw Elba`" (attributed to Napolean). Spaces, capitalization and punctuation marks are ignored in these examples. For any alphabet Σ, the set of palindromes over Σ is context free. To see this explicitly for $\Sigma = \{0, 1\}$, consider the grammar $G = (\{S\}, \{0, 1\}, S, R)$ where R is the set of rules

$$S \to 0S0 \mid 1S1 \mid \epsilon. \tag{5}$$

Note that (5) is just a short hand for the set of three rules

$$S \to 0S0, \quad S \to 1S1, \quad S \to \epsilon$$

This grammar generates the set of binary palindromes of *even* length. To generate all binary palindromes, add two more rules:

$$S \to 0 \mid 1.$$

¶4. Example: Arithmetic Expressions. Many programming languages have a sublanguage for simple arithmetic expressions, that is essentially the following:

```
<E>  --> <id> | <num> | (<E>) | - <E> | <E> <op> <E>
<id> --> a | b | c | ... x | y | z
<num> --> <dig> | <dig><num>
<dig> --> 0 | 1 | ... | 8 | 9
<op> --> + | - | * | /
```

Exercises

Exercise 0.2.1.19: Extend the simplistic English grammar in the introduction to generate more complex sentences.
(i) Allow nouns that are either singular or plural. So you will need new variables such as `<NP Singular>`, `<NP Plural>`, etc.
(ii) Introduce simple tenses (past, present, future). □

End Exercises

0.2.2 Generalized and Restricted Grammars

Context free grammars can be generalized. In general, a **grammar** is a 4-tuple $G = (V, \Sigma, S, R)$ whose meaning is the same as for a CFG, except that the rules in the set R has the more general form

$$u \to v, \qquad uv \in (V \cup \Sigma)^*. \tag{6}$$

In other words, our rules no longer has the restriction of context free grammars which specify that u must be a variable. Again, we define the language $L(G) \subseteq \Sigma^*$ **generated** by G. We will see that such grammars can be regarded as a completely general computing device (namely, equivalent to Turing machines).

¶5. Context Sensitive Languages. A general strategy in formal langauge theory is to investigate language classes that are generated by grammars with syntactic restrictions on their rules. Context free grammars is just one case. For example, a grammar whose rules (6) are restricted so that $|u| \leq |v|$ is called a **context sensitive grammar**. These grammars generate the class *CSL* of **context sensitive languages**. It is immediate from the definitions that

$$CFL \subseteq CSL.$$

We will have more to say about the complexity of this class in Chapter 2. Sometimes, to emphasize that a grammar is not restricted in any way, we call it a *general* grammar.

¶6. Normal Forms. We are also interested in restrictions on CFG's. There are two common restrictions.
(i) A grammar G is in **Chomsky Normal Form** if every rule has the form

$$A \to BC, \qquad A \to a$$

where A, B, C are variables, B, C is not the start symbol, and $a \in \Sigma$. In addition, we allow the special rule of the form $S \to \epsilon$.
(ii) We say G is in **Greibach Normal Form** if every rule has the form

$$A \to aB_1B_2 \cdots B_m, \qquad (m \geq 0)$$

where $a \in \Sigma$ and B_is are variables.

These restricted forms are called "normal forms" is because it can be shown that they do not reduce the expressive power of CFG's. We only prove this result for Chomsky Normal Form, defering the corresponding result for Greibach Normal Form to the Exercises.

THEOREM 8. *Every context-free language is generated by a CFG in Chomsky Normal Form.*

Proof. Starting from a CFG G, we transform G into Chomsky Normal Form in five steps:

First ensure that S does not appear on the right hand side of any rule. To do this, replace all occurences of S on the RHS of any rule by a new variable S'. Also, add the rule $S \to S'$.

Second, remove each rule of the form $A \to \epsilon$ and introduce, for any rule $B \to w$ where A occurs in w, a new rule $B \to w[A/\epsilon]$ ($w[A/\epsilon]$ denotes the result of replacing each occurence of A in w by ϵ). If $\epsilon \in L(G)$, we also add the special rule $S \to \epsilon$.

Third, we replace every rule of the form $A \to u_1, \ldots, u_k$ ($k \geq 3$) by the new rules: $A \to u_1A_1$, $A_{k-2} \to u_{k-1}u_k$, and $k-3$ rules of the form

$$A_i \to u_{i+1}A_{i+1}, \qquad (i = 1, \ldots, k-3),$$

where $A_1, \ldots, A_{k-2}$ are new variables. The degree of a rule is the length of its RHS. At this point, any remaining rules have degrees 1 or 2.

Fourth, we fix rules of degree 2. For each rule of the form $A \to bC$, where $b \in T$ and $C \in V$, we replace it by $A \to BC$ and $B \to b$ (B is a new variable). We have a similar treatment for rules of the form $A \to Cb$ or $A \to bc$.

Fifth and finally, we fix rules of degree 1. If we have a derivation of the form $A \Rightarrow^* B$ and $A \neq B$ and $B \to a$ or $B \to CD$. Then introduce the rule $A \to a$ or $A \to CD$ (respectively). Moreover remove all rules of the form $A \to B$. This completes the description of a Chomsky Normal form. We leave is as an routine exercise to show that this grammar generates the original language. **Q.E.D.**

This is useful if we want to show that something cannot be context-free: we only have to attack grammars in this special form. As exercise, try convert the CFG example above to Chomsky Normal Form.

EXERCISES

Exercise 0.2.2.20: Convert the binary palindrome grammar above into Chomsky Normal Form. □

Exercise 0.2.2.21: Prove that every CFL can be generated by grammar in Greibach normal form. □

Exercise 0.2.2.22: Prove that the following language $\{a^nb^nc^n : n \geq 0\}$ is context sensitive. □

Exercise 0.2.2.23: Describe a language L that can be generated by a grammar, but does not appear to be context sensitive. It is sufficient to give plausible arguments for the latter. □

Exercise 0.2.2.24: Let $L \subseteq \{1\}^*$. Such a language is called a "sla language" (sla stands for "single letter alphabet") or **tally language**.
(i) Show a tally language that is not context free.
(ii) Show that any tally language L that is context free must be regular. HINT: assume G is a Chomsky Normal Form grammar for L. Try to construct an nfa to accept L. □

END EXERCISES

0.2.3 Pumping Lemma for Context Free Languages

As for regular languages, we are interested to know if there are non-context free languages. Which of the following languages are CFL?

$$L_{\texttt{pal}} = \{w^R w : w \in 0, 1^*\}, \qquad L_{\texttt{dup}} = \{ww : w \in 0, 1^*\}.$$

We have already shown that $L_{\texttt{pal}}$ is CFG. What about $L_{\texttt{dup}}$ whose definition closely resembles $L_{\texttt{pal}}$? It is not context free, but how can we prove this? The standard technique for this purpose is a new pumping lemma which generalizes the pumping lemma for regular languages.

¶7. The $uvxyz$-Decomposition. The idea of the pumping lemma for CFL is based on the following setup. Suppose T is a derivation tree for a word $w \in L$. Let $\pi = (A_0, A_1, \ldots, A_n, a)$ be any path in T starting from the root $A_0 = S$ (the start symbol). Here, the A_i's are variables along the path, and a is the terminal at the end of the path. Suppose $A_i = A_j$ for some $0 \leq i < j \leq n$. Then w can be decomposed as $w = uvxyz$ where

- x is the word generated by the subtree rooted at A_j
- v (resp., y) is the word generated by the "left offshoots" (resp., "right offshoots") of the subpath $(A_i, \ldots, A_j)$. More precisely, suppose A_k has two children, and A_k^L, A_k^R denotes its left and right child. If $i \leq k < j$ and $A_{k+1} = A_k^R$, then the subtree rooted at A_k^L is called the **left-offshoot** of the subpath $(A_i, \ldots, A_j)$. There is a corresponding definition for **right offshoots**.
- u (resp., z) is the word generated by the left (resp., right) offshoots of the subpath $(A_0, \ldots, A_i)$.

The expression "$w = uvxyz$" is called the $uvxyz$-decomposition of w ("$uvxyz$" is pronounced"eu-vitz"). This is illustrated by Figure 7(a). Now we can modify the tree T in two ways: (a) We may delete the subpath $(A_i, \ldots, A_{j-1})$ and all the associated subtrees to the left or right, and produce a derivation tree for uxz, or (b) We may duplicate the subpath $(A_i, \ldots, A_{j-1})$ and all the associated subtrees to the left or right, and produce a derivation tree for uv^2xy^2z. This shows that $uxz, uv^2xy^2z \in L$. In fact, we can repeat (b) as many times as we like to show that $uv^ixy^iz \in L$ for $i \geq 3$. We are now ready to prove the pumping lemma.

THEOREM 9 (Pumping Lemma). *If L is a CFL then there exists $p > 0$ (the pumping length) such that for all $w \in L$, $|w| \geq p$ implies that $w = uvxyz$ such that*
(i) $|vy| > 0$.
(ii) $uv^ixy^iz \in L$ for all $i \geq 0$.
(iii) $|vxy| \leq p$.

Proof. Assume L is generated by a grammar G in Chomsky Normal Form. Choose $p = 2^m$ where m is the number of variables in G. If $|w| \geq p$, then the height h of derivation tree $T(w)$ of w must satisfy $h \geq m + 1$ (this is because $T(w)$ is a binary tree, and the last variable in any path that generates a leaf has degree 1). But in any path $(A_0, A_1, \ldots, A_{h-1}, a)$ of length $h \geq m + 1$, there are at least $m + 1$ variables (only the last node can be a terminal). Hence some variable is repeated, say $A_i = A_j$ where $0 \leq i < j \leq h - 1$. The w has a $uvxyz$-decomposition. Property (i) comes from the fact that the subpath $(A_i, \ldots, A_j)$ must either have a left offshoot or a right offshoot. Property (ii) is a general property of $uvxyz$-decompositions. To ensure property (iii), we can look for the subpath $(A_i, \ldots, A_j)$ in $(A_{h-m-1}, A_{h-m+1}, \ldots, A_{h-1})$. This means $i \geq h - m - 1$ or, the height of the subtree rooted at A_i is at most m. So the word vxy generated by this subtree has length at most $p = 2^m$. **Q.E.D.**

There are other formulations of this lemma, but the underlying scenario is the same $uvxyz$-decomposition. Hence, if this version of the pumping lemma is inadequate for a situation, one can argue directly from the underlying $uvxyz$-decomposition.

EXERCISES

Exercise 0.2.3.25: Prove or disprove that the following are context free:
(i) $L_1 = \{w \in \{a, b, c\}^* : \#_a(w) = \#_b(w) \text{ or } \#_b(w) = \#_c(w) \text{ or } \#_a(w) = \#_c(w)\}$.
(ii) $L_2 = \{w \in \{a, b, c\}^* : \#_a(w) \neq \#_b(w) \text{ or } \#_b(w) \neq \#_c(w) \text{ or } \#_a(w) \neq \#_c(w)\}$.
(iii) $L_3 = \{w \in \{a, b, c\}^* : \#_a(w) = \#_b(w) \text{ and } \#_b(w) = \#_c(w) \text{ and } \#_a(w) = \#_c(w)\}$.
NOTE: $\#_a(w)$ counts the number of occurences of a in w. □

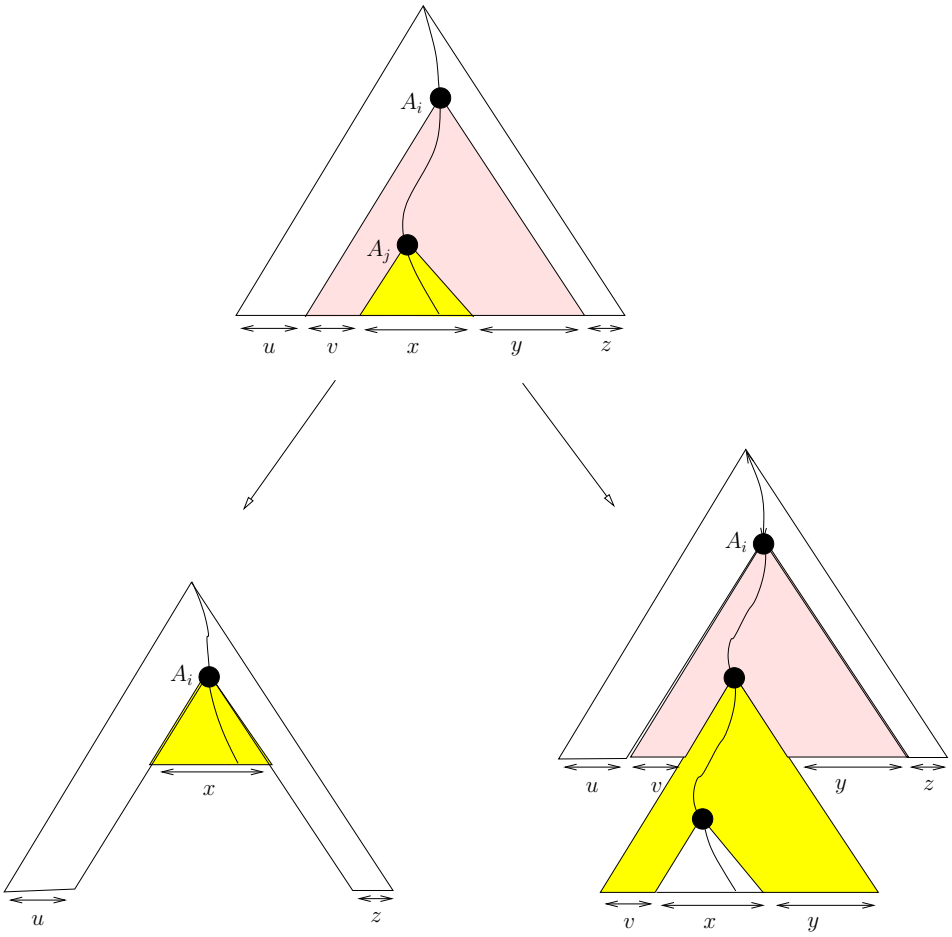

Figure 7: The $uvxyz$ decomposition of w and two transformations of the parse tree.

Exercise 0.2.3.26: Let $A, B \subseteq \Sigma^*$. The **right quotient** of A by B is defined to be

$$A/B := \{w \in \Sigma^* : (\exists u \in B)[wu \in A]\}.$$

(i) Show that if A is context free and B is regular, then A/B is context free.
(ii) Use part (i) to show that the language $\{0^p 1^n : p \text{ is prime}, n > p\}$ is not context free. □

End Exercises

0.2.4 Pushdown Automata

A pushdown automata (pda) is like an nfa except it has a linear storage structure called a **pushdown store** which we also call[2] a **stack** for brevity. The pda can read only the topmost symbol of the stack. Depending on the transition rules, may replace the topmost symbol by a new symbol or ϵ (the latter means the topmost symbol is popped); it may also place a new symbol on top of current topmost symbol (this "pushes" a new symbol). This is illustrated in Figure 8. For example, it is quite easy to design a pda to accept $L_1 = \{0^n 1^n : n \geq 0\}$.

Formally, a **pushdown automata** (pda) is a 6-tuple

$$M = (Q, \Sigma, \Gamma, \delta, q_0, F)$$

[2] A "stack" is sometimes reserved for a generalization of pushdown store in which the automata could go below the topmost symbol for the purposes of reading (but not changing) the store contents. Since we will not discuss such "stack automata", our current terminology should not be a source of confusion.

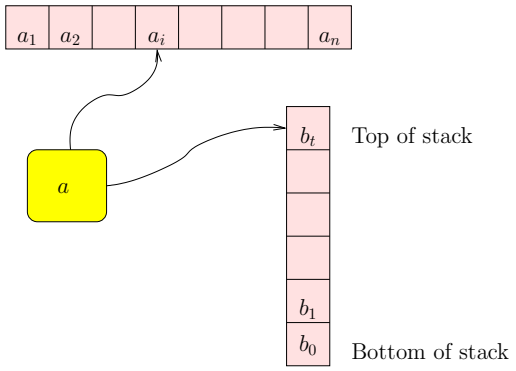

Figure 8: Illustration of a push down automata

where Σ, Γ are **input** and **stack alphabets**, $q_0 \in Q$ is the start state, $F \subseteq Q$ the final states, and δ the transition function

$$\delta : Q \times \Sigma_\epsilon \times \Gamma_\epsilon \to 2^{Q \times \Gamma_\epsilon}$$

The relation

$$(q', b') \in \delta(q, a, b) \tag{7}$$

represents a **transition rule**. Similar to grammars, we prefer to write such a rule using the arrow notation:

$$(q, a, b \to q', b').$$

This rule says that if the current state is q, a is the current input symbol and b the top of stack, then we can next go to state q' and replace b by b'. Thus δ gives rise to a finite set of such rules, and conversely from any such set there is a unique δ. We will use whichever viewpoint is convenient. What makes the notion of a pda transition slightly confusing is that a, b, b' can independently be ϵ.

Our goal, as usual, is to define the language $L(M)$ accepted by M. Towards this end, it is useful to formalize the concept of a **configuration** of a pda: this is a triple of the form

$$C = (q, w, v) \in Q \times \Sigma^* \times \Gamma^*.$$

Intuitively, q is the current state, w is a suffix of the input word (with $w[1]$ the current input symbol) and v the stack contents (with $v[1]$ the top of stack). If $C' = (q', w', v')$ is another configuration, then the binary relation

$$C \vdash_M C'$$

(or, simply, $C \vdash C'$) holds if C can reach C' by applying a single transition rule of M. More precisely, (7) is a transition rule of M and

$$\begin{aligned} w &= aw', \quad a \in \Sigma_\epsilon \\ v &= bu, \quad b \in \Gamma_\epsilon \\ v' &= b'u. \end{aligned}$$

Let $\vdash^*$ denote the reflexive transitive closure of $\vdash$. Thus $C \vdash^* C'$ iff there exists a finite sequence of the form

$$C_0 \vdash C_1 \vdash \cdots \vdash C_m, \quad m \geq 0 \tag{8}$$

where $C_0 = C$ and $C_m = C'$. We then say C **derives** C' **in** m **steps**. We define $L(M)$ to comprise all $w \in \Sigma^*$ such that

$$(q_0, w, \epsilon) \vdash^* (q, \epsilon, v)$$

for some $q \in F$ and $v \in \Gamma^*$. We say M accepts w. An alternative definition of acceptance is to insist that $v = \epsilon$, in which case we say M accepts by empty store and denote the corresponding language by $L_\epsilon(M)$.

¶8. Nondeterminism and Example. A pda in inherently nondeterministic in this definition. The nondeterminism in the above relation $C = (q, w, v) \vdash (q', w', v') = C'$ arises in three ways:

- There may be several choices of (q', b') from the set $\delta(q, a, b)$.
- If $w \neq \epsilon$ then we have a choice of $a = w[1]$ or $a = \epsilon$.
- If $v \neq \epsilon$ then we have a choice of $b = v[1]$ or $b = \epsilon$.

To illustrate this nondeterminism, consider a pda to accept the language $L_{\texttt{pal}} = \{w \in \{0,1\}^* : w = w^R\}$ of binary palindromes. We make a simple observation: if $w \in L_{\texttt{pal}}$ then $w = uav$ for some $u, v \in \{0,1\}^*$ and $a \in \{0, 1, \epsilon\}$ and $u = v^R$. The idea is to operate in two phases: in the first phase, we just push u into the stack. In the second phase, we verify that the rest of the input (namely v) is equal to to the stack contents (namely u^R). But how do we know when we have reached the middle of the input, namely, the position of a? We use nondeterminism! More precisely, our pda will have 4 states: start state q_0, phase 1 state q_1, phase 2 state q_2, and accept state q_f. This pda is represented in the following state diagram:

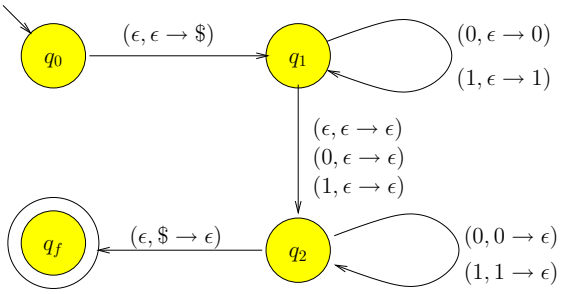

Figure 9: PDA to accept binary palindromes.

While nondeterminism in finite automata turns out to be inessential, the nondeterminism of pda's can be shown to be essential. We can define (Exercise) the notion of **deterministic pda** (dpda) and **deterministic context free langauges** (the latter is slightly technical). The palindrome language $L_{\texttt{pal}}$ is shown *not* to be deterministic. The source of nondeterminism in $L_{\texttt{pal}}$ comes from the fact that a pda cannot detect the middle of palindromes. Hence, to make palindroms deterministic, we can define the language $L^{\#}_{\texttt{pal}}$ that comprise all palindromes of the form

$$u^R \# u, \quad u \in \{0,1\}^*.$$

This is the language[3] of "marked (even) palindromes".

¶9. Extensions of pda's. We can generalize the transition function δ slightly, by allowing transition rules of the form

$$(q, a, u \to q', v) \tag{9}$$

where $q, q' \in Q$, $a \in \Sigma_\epsilon$, $u, v \in \Gamma^*$. This generalizes the original pda because we now allow $|u|$ and $|v|$ to be ≥ 2. It is an easy exercise to show that this does enlarge the class of languages accepted by pda's. In state diagrams, the rule (9) appears as a label $(a, u \to v)$ of the edge from q to q'.

Another possible extension is this: we often would like to recognize the end of the input string. Towards this end, we can introduce a special kind of **end-rule** for pdas:

$$(q, \$, a \to q', b) \tag{10}$$

where \$ is a new symbol not in Σ. This rule is applicable *precisely* when the input is completely depleted. A pda with this special kind of rules is called an **end pda**. In general, an end pda is equivalent to some regular pda. When pda's are restricted to be deterministic, this is no longer true. See Exercises.

In the same way, a pda may also like to detect when its stack is empty. The standard way to achieve this is to provide only one transition rule for the start state q_0, namely

$$(q_0, \epsilon, \epsilon \to q_1, \$)$$

[3] If we want a deterministic language to embed *all* binary palindromes, we could define $\{u^R c u : u \in \{0,1\}^*, c \in \{\epsilon, \underline{0}, \underline{1}\}\}$, where $\underline{0}, \underline{1}$ are markers corresponding to 0 and 1.

where \$ is a special symbol that marks the bottom of the stack and q_1 some new state.

EXERCISES

Exercise 0.2.4.27: Show that we can allow pda's with the more general rule of the form (9). □

Exercise 0.2.4.28: We say that a pda M is **(properly) nondeterministic** if it has two distinct transition rules $(q, a, b \to r, c)$ and $(q', a', b \to r', c')$ such that the following conditions hold:

$$\begin{aligned} &(A) \quad q = q', \\ &(B) \quad a = a' \text{ or } a = \epsilon \text{ or } a' = \epsilon, \\ &(C) \quad b = b' \text{ or } b = \epsilon \text{ or } b' = \epsilon. \end{aligned}$$

If M is not properly nondeterministic, then is is called a **strongly deterministic pda** (strong dpda). Call a language **strongly deterministic context free** if it is accepted by some strong dpda.
(i) Show that $L_1 = \{0^n 1^n : n \geq 0\}$ is strongly deterministic context free.
(ii) Show that $L_2 = \{0^n 1^m : m = n \text{ or } m = 0\}$ is *not* strongly deterministic context free. □

Exercise 0.2.4.29: The previous exercise shows a language L_2 that is not strongly deterministic. Intuitively, the reason a strong dpda cannot recognize L_2 is because the dpda cannot detect the end of its input. To remedy this, recall the notion of end pda which has special "end-rules" to detect the end of inputs. An end pda is defined to be **properly nondeterministic** as before, with the end-rules treated like any other rules in this definition). Finally, a **deterministic pda** (dpda) to be a end pda that is *not* properly nondeterministic. A language is **deterministic context free** if it is accepted by some dpda.
(i) Show that the language L_2 in the previous exercise is deterministic context free.
(ii) Show that the palindromes $L_{\texttt{pal}}$ is not deterministic context free.

SOLUTION: (ii): let $w \in L_{\texttt{pal}}$. If a dpda M accepts $L_{\texttt{pal}}$, then when M reaches the end of input w, the special end-rule will be enabled, and M will enter a final state. INCOMPLETE...

□

END EXERCISES

0.2.5 Pushdown Automata Characterization

The main result in this section is that pda and CFG are equivalent. We split the proof into two parts.

THEOREM 10. *Every context free language is accepted by a pda.*

Proof. Assume $G = (V, \Sigma, S, R)$ is a grammar in Chomsky Normal Form. We will construct an equivalent pda $M = (Q, \Sigma, \Gamma, \delta, q_0, F)$. Let w be generated by G using the following leftmost derivation:

$$S \Rightarrow u_1 \Rightarrow u_2 \Rightarrow \cdots \Rightarrow u_m = w.$$

We can express $u_i = w_i A_i v_i$ where w_i is a prefix of w and A_i a variable. Our pda will simulate this derivation in m steps. In the ith step, M will have finished read w_i and its stack contains $A_i v_i$ (with A_i at the top of stack). It remains to show how to proceed to the $i+1$st step.

M has only three states $Q = \{q_0, q_1, q_f\}$ and each simulation step above is performed while in state q_1. There are two kinds of productions with A_i on the left-hand side, and M will have corresponding transition rules:
(a) $A_i \to a$. In this case, the pda has the transition rule $(q_1, a, A_i \to q_1, \epsilon)$.
(b) $A_i \to BC$. In this case, the pda has the transition rule $(q_1, \epsilon, A_i \to q_1, BC)$.

It is clear that this achieves our step by step simulation. To start off the induction, we have the transition

$$(q_0, \epsilon, \epsilon \to S\$)$$

where S is the start symbol of the grammar and \$ is a special symbol to indicate the bottom of stack. To terminate the simulation, we have the rule

$$(q_1, \epsilon, \$ \to q_f, \epsilon)$$

This is the only way to enter the final state q_f.

This completes the description of M, and we omit the routine demonstration that $w \in L(G)$ iff $w \in L(M)$. **Q.E.D.**

COROLLARY 11. *REG is a proper subset of CFL*

Proof. Since a pda is a generalization of a nfa, we conclude that $REG \subseteq CFL$. This inclusion is strict since palindromes are non-regular languages that belongs to *CFL*. **Q.E.D.**

The reverse direction is only slightly harder.

LEMMA 12. *Every language accepted by a pda is context free.*

Proof. Given a pda $M = (Q, \Sigma, \Gamma, \delta, q_0, F)$, we will construct a grammar $G = (V, \Sigma, S, R)$. We may assume that M accepts by empty stack, namely if it ever enters a state of F, then its stack is empty. In fact, we can assume $F = \{q_f\}$. Another simplification will be useful: each transition of M will either push a symbol or pop a symbol from the stack. That is, each rule $(q, a, b \to q', b')$ satisfies either $b = \epsilon$ and $b' \neq \epsilon$, or $b \neq \epsilon$ and $b' = \epsilon$.

Let $p, q \in Q$. Define the language L_{pq} to comprise all words $w \in \Sigma^*$ such that $(p, w, \epsilon) \vdash^* (q, \epsilon, \epsilon)$. Our grammar G will have variable A_{pq} to generate precisely the language L_{pq}. To do this, consider a computation path starting from configuration $C_p = (p, w, \epsilon)$ and terminating in $C_q = (q, \epsilon, \epsilon)$. There are two cases.
(A) There is configuration $C_r = (r, w', \epsilon)$ such that $C_p \vdash^+ C_r \vdash^+ C_q$. In this case, $w = w''w'$ where $w'' \in L_{pr}$ and $w' \in L_{rq}$. Our grammar has the rule

$$A_{pq} \to A_{pr}A_{rq}$$

for every p, q, r. Hence, inductively we could generate w'' and w'. But what is the induction hypothesis? We do induction on $|w|$, and assume that for all $p, q \in Q$ and all $u \in L_{pq}$, if $|u| < |w|$ then u can be generated by our grammer.
(B) There is no such configuration C_r. We know that first and last step of this path be a push action and a pop action, respectively. If the state after the push action (resp., before the pop action) is r (resp., s), then let the corresponding transition rules be

$$(p, a, \epsilon \to r, b), \quad (s, a', b \to q, \epsilon).$$

But our grammar has, for every such pair of transitions, a rule of the form

$$A_{pq} \to aA_{rs}a'.$$

Thus $w = aw'a'$ for some $w' \in \Sigma^*$. Moreover, $w' \in L_{rs}$. Again, by induction $|w|$, we know that w' can be generated from A_{rs}, and hence we are done. **Q.E.D.**

The reason this characterization is useful is that pda's and CFG's present completely different viewpoints of a given language. To see this, consider the following fact: *CFL is closed under reversal.* This is trivial to show, using CFG. The exercise asks you to show this result using pda's.

EXERCISES

Exercise 0.2.5.30: Let M be a pda. Construct another pda N that accept $L(M)^R$, the reverse of $L(M)$. □

Exercise 0.2.5.31: In the same spirit as the previous exercise, describe a language A that is easily seen to be accepted a pda, but whose generation by a grammar is considerably harder. □

END EXERCISES

Notes on Context Free Languages

A basic reference for formal language theory and its connection to automata theory is Hopcroft and Ullman [5]. The notion of context free grammars and pda's were introduced by Chomsky (1956) and Oettinger (1961), respectively. Context free grammars is closely related to the Backus Normal Form or Backus-Naur Form formulation, described by John Bakus and Peter Naur in the late 1950's. The dynamic programming algorithm for recognizing context free grammars (Exercise) is from T. Kasami (1965).

In the 1950s, the study of formal languages led to considerable optimism in the ability of machines to understand natural languages. Chomsky's theory of transformational grammar was thought to explain deep semantic structures. In particular, machine translation (MT) was thought imminent. This turns out to be misplaced optimism. As an undergraduate, my professor (the late M.L. Dertouzos) said this:

> ``When I was a graduate student, my professor told
> us that MT will be possible in a few years''

He gave anecdotal examples of what machines produced. The sentence "`The spirit is willing but the flesh is weak`" when translated into Russian, and back to English again, produced "`The vodka is strong but the meat is rotten`". The sentence "`Out of sight, out of mind`" produces "`Blind idiot`". Modern update on these examples may be tried using the many translation services available on the web. AltaVista's Babel Fish Translation (ca. 2006) for the above two sentences to Russian and back to English yields produces "`Spirit is willingly ready but flesh it is weak`" and "`From the sighting, from the reason`". Another company (Free Translation Online) produces "`The spirit wishes, but the flesh is weak`" and "`Outside of a field of vision, from opinion`". Considerable progress has been made today, especially in specialized domains. For instance, in the European Community (EC), official documents can be automatically translated quite accurately to all the official languages of the EC. If formal languages turned out inadequate for natural languages, they turn out to have great impact in the design and analysis of modern computer languages.

EXERCISES

Exercise 0.2.5.32: The Notes described classic MT translations of examples such "`The spirit is willing but the flesh is weak`". Try to outdo (you know what we mean) these machine translations, but you have the advantage of choosing your own sentences to translate. [This exercise is to prove that the machine translation is too hard to leave to machines alone.] □

Exercise 0.2.5.33: Prove or disprove: $\{a^n b^{2n} a^n : n \geq 0\}$ is in *CFL*. □

Exercise 0.2.5.34: (i) Construct an efficient algorithm that, on input $\langle G, w\rangle$ where $G = (V, T, S, R)$ is a grammar in Chomsky Normal Form and w a string, decide whether $w \in L(G)$.
HINT: Use dynamic programming. For $1 \leq i \leq j \leq n$, let w_{ij} denote the substring $a_i \cdots a_j$ where $w = a_1, \ldots, a_n$. Define $V_{ij} = \{A \in V : A \Rightarrow^* w_{ij}\}$. How do you compute V_{ij} if you know the sets V_{ik}, V_{kj} for all $k = i, \ldots, j$?
(ii) What is the worst-case complexity of your algorithm, as a function $T(m, n)$ of input sizes $m = |G|$ *and* $n = |w|$? There is an interesting issue here: since G and w are arbitrary, we need to encode them in some fixed alphabet Σ_0. The symbols of G must first be encoded in Σ_0^*. Hence, $|w|$ and $|G|$ must be suitably interpreted. Use the following convention: assume Σ_0 contains the special symbols $A, a, 0, 1$ (among others), and each symbol of V is encoded as a string of the form $A(0+1)^*$, and each symbol of T and w is encoded as a string of the form $a(0+1)^*$. The definition of $|G|$ can be taken to be the number of symbols in writing down all the rules of G plus $|V \cup T|$. Each symbol in x has length equal to its encoding as a string in $L(a(0+1)^*)$! Similarly, you need specify your encoding G over the fixed alphabet Σ_0 and tell us how to determine its length $|G|$. □

END EXERCISES

0.3 Computability

This section introduces two classes of languages, *REC* and *RE*. They capture the concept of "computability" in two different senses ("strong" and "weak" computability). This material goes back to foundational studies in

the 1930's when mathematicians began to formalize the concept of computability. Various models were studied but they all turned out to be equivalent in the sense that what is computable in one model is also computable in another, and vice-versa. This observation is called the **Church Thesis**. The model introduced by Turing [11] is chosen for emphasis in this book for a variety of reasons. Among the many reasons for this choice, we simply note its intuitive simplicity and resemblance to modern computers.

0.3.1 Simple Turing Machines

We introduce a type of automata called **Simple Turing Machines** (STM). This is essentially[4] the original model introduced by Turing. The model can be elaborated in many ways. For instance, we will later introduce a version that is better suited for complexity considerations.

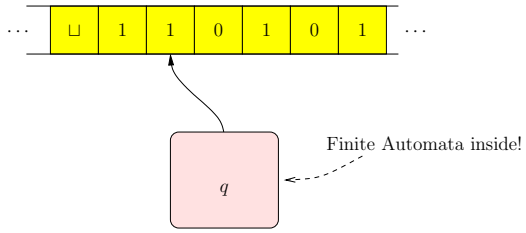

Figure 10: A STM is a finite automaton with a tape

call it a toilet

Informally, a STM is a finite automaton that is given an infinite tape which it can read an write on. This is illustrated in Figure 10. The tape is divided into tape cells, where each cell can hold a symbol. The STM has a reading head that, at any moment, is scanning a cell. Initially the input string in placed on this tape, with the reading head scanning the first symbol of the input. So far, the concept of a tape has not been explicit for a finite automaton because we only read the input from left to right at each step. In fact, we did not even allow the automaton to re-read the input. In order for the writing on a tape to be used, we must allow the automaton to move its reading head left as well as right. Therefore, the notion of a transition for a STM is slightly more elaborate, as it must specify not only state transition, but also writing of symbols and head movement (one step to the left or right, or stationary). So a transtion is a 5-tuple, (q, b, q', b', D) which says that if the current state is q and we read the symbol b on the tape, then we move to state q', replace b by b' and move in the direction indicated by D.

Formally, a **simple Turing machine** (STM) is a 6-tuple $M = (Q, \Sigma, \delta, q_0, q_a, \sqcup)$ where Q is a finite set of states, Σ is an alphabet, $q_0 \in Q$ is the **start state**, $q_a \in Q$ the **accept state**, $\sqcup \in \Sigma$ is the **blank symbol** and

$$\delta \subseteq Q \times \Sigma \times Q \times \Sigma \times \{-1, 0, +1\}$$

is the **transition table**. The STM has a **tape** with infinitely many tape **cells** indexed by the integers. The ith cell ($i \in \mathbb{Z}$) can store any symbol in Σ; thus the **tape contents** can be viewed as a function $T : \mathbb{Z} \to \Sigma$. The Turing machine M has a **current state** $q \in Q$ as usual, and its **tape head** is positioned at some tape cell whose symbol is **being scanned**. Its **current (head) position** is i if it is positioned at the ith cell. Each 5-tuple $(q, a, q', a', D) \in \delta$ is called an **instruction**. We write this 5-tuple in the style of transition rules,

$$(q, a \to q', a', D). \tag{11}$$

This instruction is **executable** if the current state is q and head is scanning symbol a. The pair (q, a) is the **precondition** of the instruction. Similarly, (q', a', D) is called the **postcondition** of the instruction. There may be several executable instructions (thus, M can be nondeterministic). We say M is **deterministic** in case no two instructions share the same precondition. To execute instruction (11), we enter state q', change symbol a to a' and move to position $i + D$ where i is the current head position. Thus D indicates the direction of head movement.

[4] Turing noted the possibility of nondeterministic versions of his model, but chose to focus on the deterministic case. He called the nondeterministic and deterministic machines (respectively) **choice-** or **c-machines** and **automatic** or **a-machines**. What we call states are **configurations** in Turing terminology.

¶10. Example of an STM via State Diagrams. Again, we can represent a STM using state diagrams. This is again a directed graph whose vertex set is Q and edges are labeled by one or more instructions. The instruction (11) will show up as the label $(a \to a', D)$ for the edge (q, q'). Consider the STM of Figure 11 that is intended to recognize the language[5] of "marked duplicate words", $L^{\#}_{\mathtt{dup}} := \{w\#w : w \in \{0,1\}^*\}$.

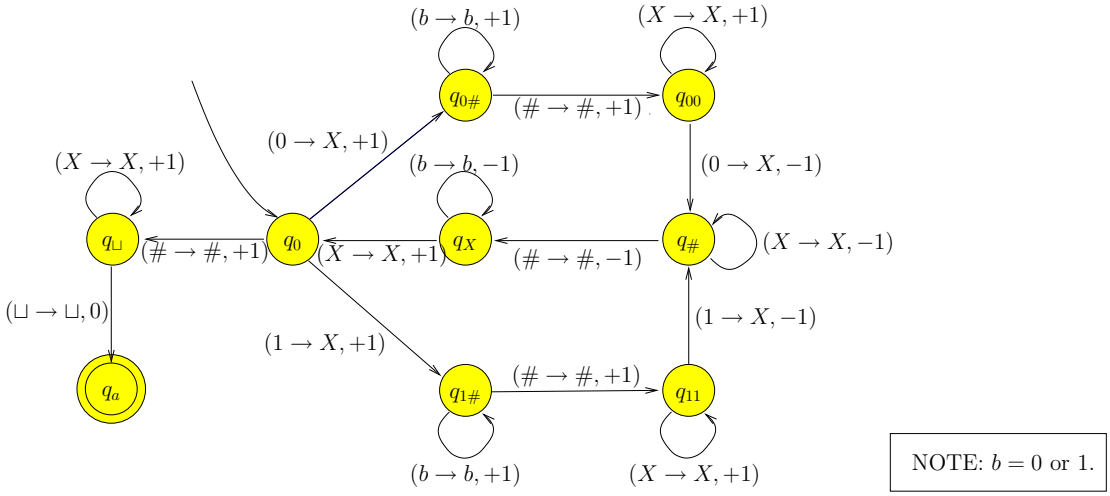

Figure 11: STM for marked duplicates, $L^{\#}_{\mathtt{dup}} = \{w\#w : w \in \{0,1\}^*\}$.

The idea of the STM in Figure 11 is to "cross out" the bits ($b = 0, 1$) that has been processed, replacing them by an X. Assume the original input has the form $u\#v$ ($uv \in \{0,1\}^*$). At the beginning of the ith stage, the following conditions hold:

- We are in state q_0.
- The tape contents is $X^i u'\#X^i v'$ where u', v' are suffixes of u, v that remain to be crossed out.
- The tape head is scanning the first symbol of the substring $u'\#$.
- There is a string w such that $wu' = u$ and $wv' = v$.

Now consider the operations of the STM in stage i: If $u' = \epsilon$, then we enter state $q_\sqcup$ and it is easy that we finally enter the accept state iff $v' = \epsilon$. This acceptance is correct. If $u' = 0u''$ then we enter state $q_{0\#}$ which finally moves us past the last X while in state q_{00}. If the next symbol is not 0, we get stuck. Otherwise, we see 0, turn it into an X and move the tape head back back to state q_0. Thus the hypothesis for stage $i+1$ holds. If $u' = 1u''$ then we have a symmetrical situation. In any case, we either enter stage $i+1$, or accept or reject (get stuck).

This proof shows that, at the end of stage i, we accept iff $u' = v' = \epsilon$. This acceptance is clearly correct. Otherwise, we get stuck (also correct) or enter stage $i+1$.

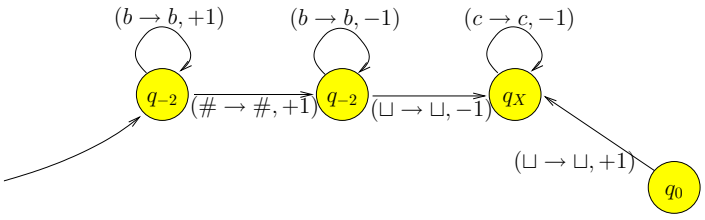

Figure 12: Additional transitions for marked duplicates STM

Hence this STM accepts $L_{\mathtt{dup}}$, assuming the input has the form $u'\#u$ (stage 0). But it would also accept $\#X^n$. To fix this, we want to first check that the input string has the form $u'\#u$. This can be done by adding

[5] We could also consider the "unmarked" version $L_{\mathtt{dup}} := \{ww : w \in \{0,1\}^*\}$.

the additional states and transitions shown in Figure 12. If we now make q_{-2} the start state (a violation of our convention of using q_0 as start state), it is then clear that our STM will accept the language $L_{\texttt{dup}}$.

¶11. Computation. We now formalize the computation process. Let the tape content be $T : \mathbb{Z} \to \Sigma$. In computer science, we like to think of T as an infinite array, and hence write "$T[i]$" instead of $T(i)$, to suggest of array indexing. At the start of the computation, $T[i] = \sqcup$ everywhere except for finitely many i's. This remains true in the subsequent computation. The complete information about the Turing machine at any particular instant is captured by its "configuration" (also known as instantaneous description (ID). This configuration is basically a triple (q, h, T) where q is the current state and $h \in \mathbb{Z}$ is the current head position. But we propose to use a more compact representation: a **configuration** is defined to be an element of

$$\Sigma^* \times Q \times \Sigma^*.$$

Consider the configuration $C = vqw$, where $vw \in \Sigma^*$ and $q \in Q$. Intuitively, the machine is in configuration C means it is in state q, the non-blank portion of the tape is vw, and the tape head is scanning the symbol $w[1]$ (the left most symbol of w). In case w is the empty string, $w[1]$ is interpreted[6] as the blank symbol $\sqcup$.

Conversely, if the machine M is in state q, with head position $h \in \mathbb{Z}$ and tape contents T as above, we can define the corresponding configuration as follows. Choose the smallest range $[j..k] \subseteq \mathbb{Z}$ such that (a) $T[i] = \sqcup$ for all i outside the range $[j..k]$, and (b) $j \leq h \leq k + 1$. This uniquely define the range $[j..k]$. Let $u \in \Sigma^*$ represent the contents of tape cells in the range $[j..k]$, *i.e.*, $u = T[j..k]$. We divide u into two substrings v, w where $u = vw$ and $|v| = h - j$. So the current configuration is vqw.

If C, C' are configurations, we say C **directly derives** C', written $C \vdash_M C'$, if the configuration C can be transformed to the configuration C' by executing an instruction of M. Suppose the instruction being executed is (11), and $C = vqw$ and $C' = v'q'w'$. Let us describe v' and w' according to the three possibilities for direction D:

CASE $D = 0$**:** Then $v' = v$. If $w \neq \epsilon$, then w' is the same as w except that $w'[1] = a'$. Suppose $w = \epsilon$. Then $w' = a'$ if $a' \neq \sqcup$, and otherwise $w' = \epsilon$.

CASE $D = +1$**:** Then $v' = va'$. If $w = \epsilon$, then $w' = \epsilon$. Otherwise, $w' = w_1$ where $w = aw_1$.

CASE $D = -1$**:** We have $v' = \epsilon$ if $v = \epsilon$, and otherwise, $v' = v_1$ where $v = v_1 b$. It is slightly more involved to define w' for all the cases, but assuming that $v = v_1 b$ and $w = aw_1$ and a, b, a' are nonblanks, then $w' = ba'w_1$. We leave the remaining cases as an exercise.

We say C **derives** C', denoted

$$C \vdash_M^* C',$$

if there is a finite sequence $(C_0, C_1, \ldots, C_k)$ such that for each $i = 1, \ldots, k$, C_{i-1} directly derives C_i, and $C = C_0$ and $C_k = C'$. We also write $C_0 \vdash_M C_1 \vdash_M C_2 \vdash_M \cdots \vdash_M C_k$ in this case. Alternatively, we say the binary relation $\vdash_M^*$ is the reflexivity transitive closure of $\vdash_M$. We drop the subscript "M", and simply write $\vdash$ or $\vdash^*$, when M is understood.

The **initial configuration**, denoted $C_0(w)$, on input w is simply $q_0 w$. Note that $C_0(w)$ does not depend on M. An **accepting configuration** is one with the accept state q_a. We say M **accepts** w if $C_0(w) \vdash^* C$ for some accepting C. $L(M)$ is the set of words accepted by M.

¶12. Halting Computations. We refine our definitions of non-acceptance of a word. There are two distinct behaviors we want to capture. If $C \vdash C'$, we call C' a **successor** of C. We call configuration C **terminal** if it has no successors. We may assume that *an accepting configuration is terminal*. A terminal but non-accepting configuration has no executable instruction; it is said to be **stuck**. Let π be a finite or ω-sequence of configurations of the form

$$\pi : C_1 \vdash C_2 \vdash C_3 \vdash \cdots .$$

We call π a **computation path** if (a) C_1 is an initial configuration, and (b) if the sequence is maximal (i.e., if the sequence is finite, then the last configuration must be terminal.)

We define the **computation tree** of M on input w as follows: this tree is rooted at $C_0(w)$ such that if C is a node in the tree and $C \vdash C'$ then C' is a child of C in the tree. Denote the tree by $T_M(w)$. The set of maximal paths in this tree is thus the set of computation paths of M on w. This tree has bounded branching

[6]This general convention will be useful: for a finite string w, we let $w[i]$ denote the ith symbol of w when $i \in [1..|w|]$. If $i < 1$ or $i > |w|$, then let $w[i] = \sqcup$.

degree (this degree is at most the number of instructions in M). By Koenig's lemma, if the tree is infinite, then there exists an infinite path.

We say that M **halts** on w iff $T_M(w)$ is finite. With our definitions of M accepting w, and M halting on w, we are led to four possibilities:

	M halts on w	M does not halt on w
M accepts w	ACCEPT	ACCEPT
M does not accept w	REJECT	LOOP

We introduce special terminology for the two forms of non-acceptance: (a) If M does not accept w, and halts on w, we say M **rejects** w. (b) If M does not accept w, and does not halt on w, we say M **loops on** w.

The intuition is this: rejection is a strong form of non-acceptance. Looping is a weak form of non-acceptance (the machine could not really make up its mind to accept or to reject, so it goes on forever).

A useful notation is "$M(w)$" which indicates the result of running M on input w. There are three distinct ways of using this:

1. We write
$$M(w)\downarrow, \quad M(w)\uparrow \tag{12}$$
to indicate that M halts on w and M loops on w (respectively).

2. Next, when viewed as an acceptor,
$$M(w) = ACCEPT, REJECT or LOOP. \tag{13}$$
Thus $M(w)$ is a 3-valued function.

3. Viewing M as a transducer, $M(w)$ represents the output of M on input w. When the output is undefined, we will write "$M(w) = \uparrow$". Note the distinction we make between this and the looping notation in (12).

A **halting STM** M is one that halts on every input; equivalently, for every input w, M either accepts or rejects w. Note that M may accept w without halting; this can only happen if M is nondeterministic. A corollary of the next theorem is that that we could make halting synonymous with accepting:

THEOREM 13. *For every STM M, there exists a STM M' such that M' is deterministic and $L(M) = L(M')$. Moreover, if M is halting, so is M'.*

Proof. We only sketch a proof: on any input w, M' will attempt to simulate every computation path of M starting on w. It uses a breadth first search of the computation tree $T_M(w)$. At any stage of the simulation, the tape of M' will hold a string of the form
$$\#C_1\#C_2\#\cdots\#C_{i-1}\$C_i\#\cdots\#C_m \tag{14}$$
where the C_j's are configurations of M, separated by the special symbols # and \$. There is only one copy \$, and if this precedes C_i, it signifies that C_i is currently being "expanded". The configurations $C_1, \ldots, C_{i-1}$ are all at some level ℓ of the computation tree, while $C_i, \ldots, C_m$ are at level $\ell - 1$. To expand C_i, we simply replace it by all of its successors. If there is more than one successor, the number of configurations in (14) will increase; if C_i is terminal and nonaccepting, then the number of configurations decrease by one. If C_i is terminal and accepting, we are accept and halt. To perform this expansion, M' may have to move all the $C_{i+1}, \ldots, C_m$ to the right (or left, in case this expansion is really a contraction). After expanding C_i, we will expand next C_{i+1} (indicated by moving the marker \$ next to C_{i+1}). If C_i is the last configuration, we next expand the first configuration C_1. The simulation halts and rejects when there are no more configurations left to be expanded. Thus, if M is halting, so is M'. This concludes our sketch. **Q.E.D.**

As noted above, a corollary of this theorem implies that we could make Turing machines accept iff they halt: by the above theorem, we may assume M is deterministic. We then modify M so that whenever it is stuck, it will enter a looping state. Thus $M(w)\downarrow$ iff $M(w) = ACCEPT$.

Another way to interpret this result is to say that the concept of non-determinism in STM is not essential as far as as the definitions of RE and REC were concerned. This is similar to the Rabin-Scott theorem for finite automata.

¶13. Uses of Turing machines. A STM can be used in one of three capacities: as an **acceptor** of a language $L(M)$, as a **transducer** to compute a partial function $t_M : \Sigma^* \to \Sigma^*$, and as a **generator** of a language $G(M)$. We have already seen how a STM is used as an acceptor. For the other cases, we need corresponding conventions.

- To view M as a generator of language over $\Sigma_\sqcup = \Sigma \setminus \{\sqcup\}$, we must re-interpret the accept state q_a to signal an output. To generate infinitely many outputs, we must enter q_a infinitely often. Thus we no longer use the convention that accepting configurations are terminal configurations. The machine begins its computation on a blank tape. In the course of computation, the tape will contain a finite number of non-blank non-empty words, separated from each other by one or more blanks. Whenever M enters q_a, we declare that the non-blank word that is being scanned on the tape is "generated". For instance, if $w = T[i..j]$ is a non-blank word where $T[i-1] = T[j+1] = \sqcup$, and the head position h lies in the range $[i..j]$ when we enter q_a, we declare that w is being output at that instant. If the symbol being scanned is a blank, then the empty word ϵ is output. The set of all generated words is denoted by $G(M) \subseteq \Sigma_\sqcup$. Note that the words in $G(M)$ may be generated more than once.
- To view a STM M as a transducer, we need to define the partial function $t_M : \Sigma_\sqcup^* \to \Sigma_\sqcup^*$. We can use the same input convention as for acceptors: this means that the initial configuration is $C_0(w) = q_0 w$ where w is the input. What about output convention? We may define $t_M(w)$ to be the non-blank word $t_M(w)$ that is being scanned when the machine enters the accept state q_a. In case the scanned symbol is $\sqcup$, define $t_M(w) = \epsilon$. This is just the same output convention as for generators, except that we view q_a as a terminal state (as in the case of acceptors). When M is nondeterministic, different paths can give different outputs. We may declare $t_M(w)$ to be undefined when two or more *distinct* outputs are produced, or when no output is produced. However, we allow multiple outputs of the same string. This flexibility is actually important in known applications.

EXERCISES

Exercise 0.3.1.35: Verify that the STM in Figure 11 accepts $L^{\#}_{\mathtt{dup}}$. □

Exercise 0.3.1.36: Show that if L is accepted by a STM, then it is accepted by a STM in which every non-terminal configuration C has exactly two successors C_1, C_2. We write $C \vdash (C_1, C_2)$ in this case. This means that we can regard our computation tree $T_M(w)$ as a full binary tree. □

Exercise 0.3.1.37: Construct a STM to accept inputs of the form $x\#y\#z$ where $x, y, z \in \{0,1\}^*$ and $\langle x\rangle + \langle y\rangle = \langle z\rangle$. Here x is the binary notation for the number $\langle x\rangle \in \mathbb{N}$. □

Exercise 0.3.1.38: Define a **very simple Turing machine** (vSTM) be a simple Turing machine in which each instruction either modifies the scanned symbol or moves the head to the left or right, but not both. In other words, if the scanned symbol is modified, then the head is stationary, and if the head moves, then the scanned symbol is unmodified.
(i) Model a vSTM as a set of instructions where the instructions are 4-tuples.
(ii) Show that for every STM M, there is a corresponding vSTM N with twice the number of states, and which takes twice as many steps to accomplish the same computation as M. □

END EXERCISES

0.3.2 Recursively Enumerable Languages

A language is said to be **recursively enumerable** (r.e.) if it is accepted by some STM. The class of all r.e. languages is denoted RE. A language is **recursive** if it is accepted by some halting STM. The class of all recursive languages is denoted REC.

Suppose M accepts L. We also say that M **recognizes** or **semi-decides** L. If M happens to be halting, we say M **decides** L. It is common to refer to r.e. languages and recursive languages as r.e. sets or recursive sets. This arose from the context where these definitions were applied to subsets of $\mathbb{N}$ (see discussions below).

It is clear from our definitions that

$$REC \subseteq RE.$$

A more interesting result is the following.

COROLLARY 14. *The recursive languages are closed under complement: REC = co-REC.*

Proof. It is a simple fact that for any class K, we have $K =$ co-K iff $K \subseteq$ co-K. Hence we only have to show that if L is recursive, then co-L is recursive. Applying Theorem 13, let M be a halting deterministic STM that decides L. We construct $\overline{M}$ that acts like M, except that $\overline{M}$ accepts iff M rejects. **Q.E.D.**

THEOREM 15. *The recursive languages are those r.e. languages whose complements are also r.e.. Thus, REC = RE ∩ co-RE.*

Proof. One direction is easy: if $L \in REC$ then the previous lemma shows that co-$L \in REC$ and hence $L \in$ co-REC and hence $L \in REC \cap$ co-$REC \subseteq RE \cap$ co-RE.

Conversely, let $L \in RE \cap$ co-RE. Then there are deterministic STM's M and M' such that $L = L(M)$ and co-$L = L(M')$. We construct a STM M'' that, on input w, simulates M on w and M' on w. The trick (called dovetailing) is that we must not naively simulation either M or M' to completion, since we cannot be sure if either will halt. But we know that for every input w, at least one of M or M' will halt: if $w \in L$ then M will halt, and if $w \notin L$ then M' will halt. Hence we simulate one step of M and one step of M' alternately. As soon as either M or M' halts, we can accept or reject in the appropriate manner: we accept iff either M accepts or M' rejects. **Q.E.D.**

Let us prove one more result, to relate the context free languages to our new classes.

LEMMA 16. *The class CFL of context free languages is a proper subset of REC.*

Proof. (i) We first show the easy part, that REC contains a non-context free language. This is the language $L^{\#}_{\texttt{dup}}$ of marked duplicate words. This can be seen by a simple application of the pumping lemma for context free languages. However, the halting STM in Figure 11 described in our introduction to STM shows that $L^{\#}_{\texttt{dup}} \in REC$.

(ii) Now we show that $CFL \subseteq REC$. Let be P a pda. It suffices to construct a halting STM M such that $L(M) = L(P)$. On any input w, M first converts the tape contents into $q_0 w$ where q_0 is the start state of P. The STM M is simulating P step by step, where the general tape contents of M has the form vqw' where v is the current stack contents of P (the top of stack is $v[1]$), q is the current state of P, and w' is a suffix of w (with P is currently reading $w'[1]$). Each nondeterministic step of P is easily simulated by M (which is also nondeterministic). The simulation halts as soon as $w' = \epsilon$. We accept if q is a final state of P, or if there is a single ϵ-transition of P into a final state. This concludes our construction of M. By construction, M is halting.

It is clear that if M accepts w, then P accepts. The converse is unclear: if M rejects, we know that that the input string has been depleted ($w' = \epsilon$). But conceivably P could continue to compute using only ϵ-transitions and finally accept. We argue that this is impossible when P has the special form described in the proof of Theorem 10. In that proof, we constructed a pda whose ϵ-transitions always increase the size of its stack. The sole exception is the final ϵ-transition into the accept state q_f, but this occurs only the bottom of stack symbol \$ is read. With such a P, once the input is depleted, the only way to accept is to apply this special transition once. Since M checks for such a transition, it can safely reject when no such transition is found. **Q.E.D.**

Although nondeterminism was not essential in our definition of RE and REC, it was useful in preceding proof since it allows our constructed M to mirror the nondeterminism of a pda.

EXERCISES

Exercise 0.3.2.39: Show that L is r.e. if and only if L is generated by a general grammar. □

END EXERCISES

0.3.3 Computability Dictionary

There are at least four independent sets of terminology in the area of computability. The following table places these terminology alongside each other for comparison. Three of them comes from viewing the computing device as generators, acceptors or transducers; a fourth set of terminology comes from the study of recursion as the main mechanism in computing.

	MODEL	PARTIAL COMPUTABILITY	TOTAL COMPUTABILITY
1.	Generator	recursively enumerable set	recursive set
2.	Acceptor	semi-decidable or recognizable set	decidable set
3.	Transducer	(partial or semi-) computable function	total computable function
4.	Recursion	partial recursive function	(total) recursive function

Using a loose interpretation[7] of Church's Thesis, these four view points are also equivalent. Hence reader should feel comfortable switching among these terminology as convenient. The object of study in computability is "computational problems". The obvious form of computational problems is the functional form. For instance, the problems of multiplication or square-root correspond to the functions $g(x, y) = xy$ and $f(x) = \sqrt{x}$, respectively. To capture the essence of computability, and to avoid considerations of precision in the representation of real numbers, computability theory is usually formulated as the study **number theoretic functions**, *i.e.*, $f : \mathbb{N}^n \to \mathbb{N}$ for some $n \geq 1$. By suitable encodings, we can even assume $n = 1$. So the square root function must be re-formulated as $f(x) = \lfloor \sqrt{x} \rfloor$. Such are the functions mentioned in the last row of the table.

We can further restrict our number theoretic functions so that range has only two values (Yes/No, 0/1, accept/reject). These are called[8] **decision problems**. Such a function $f : \mathbb{N} \to \{0, 1\}$ can be identified with the set $\{n \in \mathbb{N} : f(n) = 1\}$. Thus the characterization of computable functions is transformed into the characterization of the "decidable subsets" of $\mathbb{N}$. For any alphabet Σ, there are bijections between $\mathbb{N}$ and Σ^* and hence this is equivalent to studying decidable languages, which is our main viewpoint. The sets mentioned in rows 1 and 2 can be interpreted as subsets of $\mathbb{N}$ or languages. Note that our "official definition" of these concepts uses the terminology for generators, even though we mainly treat acceptors.

The first aim of computability theory is to characterize the "solvable" problems. When problems are functions, we characterize the "computable" functions. (Thus, "solvable" and "computable" relatively interchangeable). It was quickly realized that computability can be refined into "partial computable" and "total computable". This accounts for the two columns (under "partial computability" and "total computability"). Likewise "decidable" can be refined into "semi-decidable" or "total decidable". (Here "semi-" and "partial" is another pair of words that are usually interchangeable). The presence of a qualifiers (partial, semi-, total, etc) immediately places concept in the correct column; when the qualifier is omitted, convention must tell us which column is meant.

¶14. What about other models of computation? The introduction noted that there are other models of computability. So why did we choose Turing's model? One answer is that Turing's model is remarkably flexible: when we study complexity theory later, we will see how it is easily modified to model various phenomenon in complexity. Furthermore, **Turing's Thesis** says that *anything that is algorithmically computable is computable by an algorithm is computable by Turing machines.* This is called a thesis (not a theorem) because the concept of "algorithmically computable" is informal. But given any reasonable interpretation, we will find that the thesis holds up. The converse of this thesis is clear: whatever an "algorithmically computable" means, the method used by a STM to accept a language or to compute a function qualifies as algorithmic. This justifies equating the classes *RE* and *REC* with the concepts of semi-decidable and total decidable sets.

EXERCISES

Exercise 0.3.3.40: Consider the set

$$P = \{n \in \mathbb{N} : \text{ the binary expansion of } \pi = 3.14159\ldots \text{ has } n \text{ consecutive 0's}\}.$$

The current state of mathematics does not tell us whether $P = \mathbb{N}$. Suppose $P \neq \mathbb{N}$, and interpreting P as a set of binary numbers. Discuss whether the language P decidable or semi-decidable. □

Exercise 0.3.3.41: Hilbert's 17th Problem asks: given an integer multivariate polynomial $P(x_1, \ldots, x_n)$, does it have an integer root? An integer root of P is $(a_1, \ldots, a_n) \in \mathbb{Z}^n$ such that $P(a_1, \ldots, a_n) = 0$. Show that this problem is semi-decidable. NOTE: Matiyasevic shows in 1970 that this problem is not decidable. □

END EXERCISES

[7] Church's thesis concerns the equivalence of different computational models. Here we are saying that the notion "computable sets" and "computable functions" are also equivalent.

[8] This is the German word "Entscheidungsproblem" in Turing's original paper [11].

0.3.4 Diagonalization

We now ask: are there uncomputable functions? In terms of decision problems, are there non-r.e. languages? The answer is yes, but for a very fundamental reason that has almost nothing to do with computation. It has more to do with counting and size.

If X is a set, $|X|$ denotes its **cardinality** which is intuitively the number of its elements. This is clear for finite sets, as $|X| \in \mathbb{N}$. If X is infinite, we can simply say $|X| = \infty$. But we will want to refine this classification for infinite sets. The accepted way to do this is via a theory of cardinals, so that the function $|\cdot|$ assigns a **cardinal number** $|X|$ to each set X. Without launching into a development of this theory, we can prove some basic facts via an axiomatic approach. Whatever we mean by "cardinal numbers", we want the following axioms to hold. Let X and Y be sets.

(A0) The cardinal numbers are totally ordered.

(A1) $|X| \leq |Y|$ iff there is an injection $f : X \to Y$.

From these two principles, we can conclude that if there is a bijection $h : X \to Y$ then $|X| = |Y|$. We also say X and Y are **equivalent** in this case. Conversely, if $|X| \leq |Y|$ and $|Y| \leq |X|$, we want to conclude that X and Y are equivalent. This converse is not so obvious, and is the subject of the celebrated Bernstein-Schröder theorem below.

If X is a finite set, we may identify $|X|$ with the usual natural number $n \in \mathbb{N}$ that counts the number of elements in X. The theory of cardinality becomes rather more interesting when we treat infinite sets.

The simplest approach to the Bernstein-Schröder theorem uses a fixed-point approach. This approach is used, for instance, in fixed-point semantics of of programming languages. The fixed-point principle is this:

LEMMA 17. *Suppose $\mu : 2^X \to 2^X$ is* **monotone** *in the sense that $A \subseteq B \subseteq X$ implies $\mu(A) \subseteq \mu(B)$. Then there exists a set $A^* \subseteq X$ such that $\mu(A^*) = A^*$, called a* **fixed point** *of μ.*

Proof. Call a set $A \subseteq X$ is "small" if $A \subseteq \mu(A)$; it is small compared to its image $\mu(A)$. Note that if A is small, then $\mu(A)$ is also small because of monotonicity: $\mu(A) \subseteq \mu(\mu(A))$. Define A^* to be the union of all small sets. To show that A^* is a fixed point, we first show one direction:

$$A^* \subseteq \mu(A^*). \tag{15}$$

This says A^* is small. If $a \in A^*$ then $a \in A$ for some $A \subseteq A^*$ that satisfies $A \subseteq \mu(A)$. This shows that $a \in \mu(A)$. But $\mu(A) \subseteq \mu(A^*)$, by monotonicity. Thus $a \in \mu(A^*)$. The other direction is now easy: we know that $\mu(A^*)$ is small since A^* is small. But this means $\mu(A^*) \subseteq A^*$ as A^* is the union of all small sets. **Q.E.D.**

A trivial example of a monotone map is $\mu(A) = A$ for all A. So every set is a fixed point. Our construction yields $A^* = X$. The reader may be curious about the analogous definition of "big" for sets $B \subseteq X$ that satisfy $\mu(B) \subseteq B$. An exercise asks you to show a fixed point B^* for μ using big sets.

THEOREM 18 (Bernstein-Schöder). *If $f : X \to Y$ and $g : Y \to X$ are injections, then there exists a bijection $h : X \to Y$.*

Proof. First define the function $\mu : 2^X \to 2^X$ via

$$\mu(A) = X \setminus g(Y \setminus f(A)).$$

Figure 13 illustrates this map. This function is monotone because if $A \subseteq B \subseteq X$, then

$$\begin{aligned} f(A) &\subseteq f(B) \\ Y \setminus f(A) &\supseteq Y \setminus f(B) \\ g(Y \setminus f(A)) &\supseteq g(Y \setminus f(B)) \\ X \setminus g(Y \setminus f(A)) &\subseteq X \setminus g(Y \setminus f(B)) \end{aligned}$$

The last inclusion amounts to $\mu(A) \subseteq \mu(B)$. By the previous lemma, μ has a fixed point A^*. Then $A^* = \mu(A^*) = X \setminus g(Y \setminus f(A^*))$. Writing $B^* := g(Y \setminus f(A^*))$, it follows that (A^*, B^*) is a partition of X. We now specify the bijection function $h : X \to Y$ as follows: h restricted to A^* is f, and h restricted to B^* is g^{-1}. To see that h is a bijection, it suffices note that $(h(A^*), h(B^*)) = (f(A^*), Y \setminus f(A^*))$ and hence forms a partition of Y as well. **Q.E.D.**

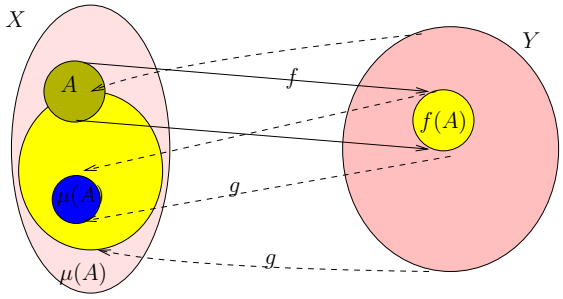

Figure 13: Injections $f : X \to Y$ and $g : Y \to X$.

¶15. Countable and Uncountable. Let the cardinal number of $\mathbb{N}$ be denoted $\aleph_0$ ($\aleph$ is the first Hebrew letter **aleph**). A set X is said to be **countable** if there is an injection $f : X \to \mathbb{N}$. Equivalently, $|X| \in \mathbb{N}$ or $|X| = \aleph_0$. If $|X| = \aleph_0$, we will say X is countably infinite, or equivalently, **denumerable**. Countability gives us just one extra number (namely $\aleph_0$) to count with – we can now make one distinction among the infinite sizes. So let us try it out!

It is easy to construct bijections between $\mathbb{N}$ and the sets $\mathbb{Z}$: For instance, if we list the elements of $\mathbb{Z}$ as $(0, -1, 1, -2, 2, -3, 3, \ldots)$, then we see that this listing can be matched to the listing $(0, 1, 2, 3, \ldots)$ of $\mathbb{N}$. Formally, the bijection is given by

$$f(n) = \begin{cases} n/2 & \text{if } n \text{ is even,} \\ -(n+1)/2 & \text{if } n \text{ is odd.} \end{cases} \tag{16}$$

Hence $|\mathbb{Z}| = \aleph_0$. This simple example illustrates a curious phenomenon that is characteristic of infinite sets: the property of having a bijection with a proper subset.

Let us prove two more such results. For any set X, recall that 2^X is the set of all subsets of X. Write $\overline{2}^X$ for the set of all finite subsets of X. Thus $2^X = \overline{2}^X$ iff X is a finite set.

LEMMA 19. *The following sets are denumerable:*
(a) The set $\mathbb{N}^n$ where $n \geq 2$.
(b) The set Σ^* of words over any alphabet Σ,
(c) The set $\overline{2}^{\mathbb{N}}$ of all finite subsets of $\mathbb{N}$.

Proof. (a) The well-known pairing function $(i, j) \mapsto P_2(i, j) = \binom{1+i+j}{2} + j$ maps $\mathbb{N}^2$ bijectively into $\mathbb{N}$. By induction, if $P_{n-1}(i_1, \ldots, i_{n-1})$ is a bijection, then $P_n(i_1, \ldots, i_n) = P_2(P_{n-1}(i_1, \ldots, i_{n-1}), i_n)$ is a bijection from $\mathbb{N}^n \to \mathbb{N}$.
(b) We list all the words of Σ^* in non-decreasing order of their lengths. This gives a bijection between Σ^* and $\mathbb{N}$.
(c) Let $\mathbb{N}(k)$ denote all subsets of $\{0, 1, \ldots, k\}$. We simply list all the elements of $\mathbb{N}(0)$, then $\mathbb{N}(1) \setminus \mathbb{N}(0)$, $\mathbb{N}(2) \setminus \mathbb{N}(1)$, etc. **Q.E.D.**

We can embed $\mathbb{Q}$ in $\mathbb{Z}^2$ by associating $x \in \mathbb{Q}$ with the unique (p, q) where $x = p/q$ and $q \geq 1$ and $(p, q) = 1$. Thus $\mathbb{Q}$ is denumerable. At this point, the reader begins to wonder if there are other infinite cardinal numbers besides $\aleph_0$. The next theorem gives the answer.

THEOREM 20 (Cantor). *For any set X, $|X| < \left|2^X\right|$.*

Proof. Clearly $|X| \leq \left|2^X\right|$. Suppose, by way of contradiction, $f : X \to 2^X$ is a bijection. Let $D = \{x \in X : x \notin f(x)\}$. Hence there is $y \in X$ such that $f(y) = D$. We ask the question: is $y \in D$? If $y \in D$, then $y \in f(y)$ and so by definition of D, $y \notin D$. Conversely, if $y \notin D$, then $y \notin f(y)$ and so by definition of D, $y \in D$. Since both possibilities lead to contradiction, we have achieved a contradiction. **Q.E.D.**

This shows that in lemma 19(b), the restriction to finite subsets of $\mathbb{N}$ is essential. It is a fact that $\left|2^X\right| = 2^{|X|}$. Hence, $2^{\mathbb{N}}$ has a cardinality $2^{\aleph_0}$ which this theorem shows is different from $\aleph_0$. The famous **Continuum Hypothesis** says that there is no cardinal number α lying strictly between $\aleph_0$ and $2^{\aleph_0}$. Put another way, if $\aleph_1$ denotes the smallest cardinal number larger than $\aleph_0$, the hypothesis asserts that $\aleph_1 = 2^{\aleph_0}$. Note that $\aleph_1$ is well-defined because cardinal numbers are well ordered (so $\aleph_1$ is just the smallest number in the set $\{\alpha : \alpha > \aleph_0\}$).

Here is another uncountable number: let $S(X)$ denote the set of all permutations of X.

THEOREM 21. $|\mathbb{N}| < |S(\mathbb{N})| \leq \left|2^{\mathbb{N}}\right|$.

Proof. The inequality $|S(\mathbb{N})| \leq \left|2^{\mathbb{N}}\right|$ is left as an exercise.

To show the first strict inequality, suppose $S(\mathbb{N})$ is countable. So let $(A_0, A_1, A_2, \ldots)$ be the list of all permutations of $\mathbb{N}$. We write $A_i = (a_{i0}, a_{i1}, a_{i2}, \ldots)$.

We need a sophisticated way to diagonalize. The idea is this: in stage 0, we choose c_0 to be a_{01}. In stage i, we consider the entries

$$a_{i0}, a_{i1}, \ldots, a_{ii}, ai, i+1.$$

We pick an element c_i from this set that is different from a_{ii} and different from any of the previous $i-1$ choices. Clearly c_i can be picked. Moreover, if there are several possible choices, we pick the element that is minimum.

We claim that $C = (c_0, c_1, \ldots)$ is a permutation of $\mathbb{N}$ that is different from each A_i. That $C \neq A_i$ is clear, because a_{ii} is different from c_i.

Why is C is a permutation? We claim that every $i \in \mathbb{N}$ appears in C. Clearly, i must appear infinitely often in every column. In particular, by the time that i first appears in column 1, if it had not been picked, it would be picked. **Q.E.D.**

Cantor's theorem uses a **diagonalization argument**. Why is this so-called? Imagine a semi-infinite matrix indexed by X on the rows and 2^X on the column. D is constructed by looking at the diagonal of the matrix. Diagonalization turns out to be a very important technique in the theory of computability.

Cantor's theorem shows the existence of uncountable sets. But which of the familiar sets we know are uncountable? Let us now show the set $\mathbb{R}$ of real numbers is uncountable. It suffices to prove that the reals in the interval $[0, 1]$ is uncountable. Use the fact that such numbers has an infinite binary expansion of the form $0.b_1b_2b_3\cdots$. If $[0, 1]$ were denumerable, we can arrange them in an infinite list $(x_1, x_2, x_3, \ldots)$. Consider the number r with expansion $0.d_1d_2d_3\cdots$ where d_n is chosen to be different from the nth bit of x_n. Note that this expansion is different from any in the list. We want to conclude that r is not in our list, which would be a contradiction. Unfortunately, this conclusion is unwarranted because the binary notation is non-unique. The binary strings $0.b_1\cdots b_n10^\omega$ and $0.b_1\cdots b_n01^\omega$ represent the same number. One way to avoid this problem is to use an m-ary notation for any $m \geq 4$. The m-ary expansion uses digits in $\{0, 1, \ldots, m-1\}$. Two distinct expansions denote the same number iff they have the form: $0.b_1b_2\cdots b_{n-1}b_n3^\omega$ and $0.b_1b_2\cdots b_{n-1}b'_n0^\omega$ where $b'_n = 1 + b_n$. In our construction of the number r, we simply make sure that d_n is chosen different from 0 and $m-1$.

The next theorem uses two propositions.

¶16. PROPOSITION A: *The class of all languages $\mathcal{L}$ is uncountable.*

¶17. PROPOSITION B: *The family of all Turing machines $\mathcal{M}$ is denumerable.*

To see PROPOSITION A, let $\mathcal{L}_0$ be the set of languages over $\{0, 1\}$. It suffices to show this set uncountable. But clearly $\mathcal{L}_0$ is equivalent to $2^{\mathbb{N}}$, which is uncountable by Cantor's theorem. We will return to the proof of PROPOSITION B in the next subsection.

THEOREM 22. *There is a language that is not in $RE \cup$ co-RE.*

Proof. By PROPOSTION B, there are only denumerably many languages in RE and hence in $RE \cup$ co-RE. By PROPOSTION A, there are uncountably many languages. The theorem follows. **Q.E.D.**

This theorem shows that there are languages outside of $RE \cup$co-RE, based on purely cardinality arguments. This is not exciting news, as it gives no indication as to which languages are not in $RE \cup$ co-RE. What is more interesting is to show that certain explicitly describable languages have this property. We will obtain such results below. This is somewhat harder to do, but it is like the cardinality argument in one respect: lurking being these proofs is still a diagonalization argument (perhaps only indirectly).

EXERCISES

Exercise 0.3.4.42: Let $\mu : 2^X \to 2^X$ be a monotone map.

(i) Prove that if B^* is the intersection of all big sets, then B^* is a fixed point of μ.

(ii) Show that if C is any fixed point of μ, then $B^* \subseteq C \subseteq A^*$.

SOLUTION: (ii) Trivial: such a C is both big and small.

□

Exercise 0.3.4.43: The identity map μ is monotone. Construct other examples of monotone μ such that $A^* \neq B^*$. □

End Exercises

0.3.5 Gödel Numbering

The previous theorem depends on an unproved PROPOSITION B. We deferred the proof until now in order to give a fuller treatment of several closely related topics.

¶18. k-adic Notation. The binary notation

$$B : \{0,1\}^* \to \mathbb{N} \tag{17}$$

for numbers is well-known. Thus, $B(101) = 5$. This notation is highly non-unique since each number has an infinite number of notations. E.g., every string denoted by the regular expression 0^*11 is a notation of the number 3. For some applications, it is convenient to have a bijection $D : \{0,1\}^* \to \mathbb{N}$. The **dyadic notation** has this property. It is rather similar to binary numbers: the string $w = b_n b_{n-1} \cdots b_1 b_0$ ($b_i \in \{0,1\}$), viewed as a dyadic number, represents the number $D(w)$ given by

$$D(w) = \sum_{i=0}^{n} (b_i + 1)2^i.$$

It is easy to verify that this defines a bijection $D : \{0,1\}^* \to \mathbb{N}$. The dyadic notation for the initial natural numbers are

$$\epsilon, 0, 1, 00, 01, 10, 11, 000, 001, \ldots. \tag{18}$$

This dyadic ordering of binary strings is also called the **lexicographic ordering** of $\{0,1\}^*$. For obvious reasons, instead[9] of the alphabet $\{0,1\}$, some authors use $\{1,2\}$ for dyadic notation.

Just as binary notation generalizes to k-ary notations, the dyadic notation generalizes to **k-adic notation**. For any natural number $k \geq 2$, let Σ be an alphabet of size k. Once we choose a bijection $D : \Sigma \to \{1, \ldots, k\}$, we naturally extend this bijection to $D : \Sigma^* \to \mathbb{N}$ where $D(b_0 b_1 \cdots b_n) = \sum_{i=0}^{n} D(b_i)k^i$. Using these bijections, we can interpret any word function $t : \Sigma^* \to \Sigma^*$ as a number theoretic function.

¶19. Numbers versus Strings. All the concepts of standard computability theory can be developed using strings, or equivalently[10] using the set of natural numbers. Hence, it is convenient to move effortlessly between finite strings and natural numbers, depending on which setting is more convenient. We indicate such bijections between Σ^* and $\mathbb{N}$ by writing

$$\Sigma^* \simeq \mathbb{N}. \tag{19}$$

If $w \in \Sigma^*$ and $n \in \mathbb{N}$, we shall write $w \simeq n$ under this bijection. We usually do not exploit any particular properties of this bijection, except that we may assume that $\epsilon \simeq 0$.

The lexicographic ordering of Σ^* is defined as in (18) using this bijection. Using (19), we can now interpret any string function $f : \Sigma^* \to \Gamma^*$ as a number-theoretic function $f : \mathbb{N} \to \mathbb{N}$.

¶20. Encoding Tuples. The bijection (19) can be extended to tuples. Thus, any n-tuple $(w_1, \ldots, w_n) \in (\Sigma^*)^n$ can be encoded by the string $w_1 \# w_2 \# \cdots \# w_n$, provided we use a symbol $\#$ that is not in Σ. To avoid explicitly specifying $\#$, we will write $\langle w_1, \ldots, w_n \rangle$ for $w_1 \# w_2 \# \cdots \# w_n$. Extending the bijection (19), we then have bijections of the form

$$\mathbb{N}^n \simeq \mathbb{N}.$$

These bijections are the (generalized) **pairing functions**, which are denoted by P_n in the proof of Lemma 19. We can introduce their partial inverses, called "projection functions" π_i^n which extract the components w_i from $\langle w_1, \ldots, w_n \rangle$.

[9] We prefer the digits $\{0,1\}$ over $\{1,2\}$ to avoid unfamiliar looking strings such as "1211222".

[10] When we discuss complexity, these two approaches are no longer equivalent. In this case, the string approach is more natural.

¶21. Enumeration and Denumeration. An **enumeration** of a set X is a surjective function $f : \mathbb{N} \to X$. We typically indicate f by the ω-sequence $(f(0), f(1), f(2), \ldots)$. Note that repetitions are allowed in this sequence. If f were a bijection, we call f a **denumeration** of X. Recall that we already defined a set to be denumerable if it is countable and infinite. Thus we see that a set X is denumerable iff it has a denumeration. If f is an enumeration of a language $X \subseteq \Sigma^*$, we say that f is **effective** if f is total recursive. This requires us to view f as a word function $f : \{0,1\}^* \to \Sigma^*$, with $\mathbb{N} \simeq \{0,1\}^*$. For example, the lexicographic listing in (18) is a denumeration of $X = \{0,1\}^*$. This denumeration is total recursive. In fact, it is trivially so because the word function f here is just the identity function.

¶22. Two Universal Assumptions. PROPOSITION B amounts to saying that there is an enumeration of the set $\mathcal{M}$ of all Turing machines. This idea goes back to Gödel in his celebrated Incompleteness Theorem. Such enumerations (when effective) have profound implications later. Unfortunately, PROPOSITION B cannot be true without further restrictions: if the set of allowed alphabets is uncountable, then $\mathcal{M}$ is trivially uncountable. As a solution, we now make two assumptions that are used throughout this book. Let Σ_∞ and Q_∞ be two disjoint denumerable sets with the following properties:

> **Convention** (α)**.** We fix Σ_∞ to be any countably infinite set of markings that are called *symbols*. Σ_∞ is the *universal set of symbols*, assumed to contain every symbol (such as $0, 1, a, b, \$, \#$, etc) that we will ever use in defining machines or languages. It contains a distinguished symbol $\sqcup$ called the *blank symbol*. No alphabet of a language contains this blank symbol.
>
> **Convention** (β)**.** We fix Q_∞ to be any countably infinite set of markings that are called *states*. We assume Q_∞ is disjoint from Σ_∞. It is called the *universal set of states*, assumed to contain every state that we will ever use in defining machines. It contains three distinguished states q_0, q_a and q_r called the *start state*, the *accept state*, and the *reject state*, respectively. The explicit use of reject states, however, can be avoided until chapter 7.

It is now easy to show PROPOSITION B: for each $n \in \mathbb{N}$, let $\mathcal{M}_n$ comprise all STM's with at most n instructions and which uses only the first n symbols of Σ_∞ and the first n states in Q_∞. Since $\mathcal{M}_n$ is finite for each n, we simply list all the elements of $\mathcal{M}_n$ for $n = 0, 1, 2, \ldots$. This shows that the family $\mathcal{M}$ of Turing machines is denumerable.

Henceforth in our treatment of computing devices (Turing machines, grammars, etc), we use these two universal[11] assumptions. Besides validating PROPOSITION B, a practical advantage is that we save verbiage in discussing Turing machines or any automata. A STM M is no longer described as "a 6-tuple", but is simply defined by a finite set δ_M (of instructions) where

$$\delta_M \subseteq Q_\infty \times \Sigma_\infty \times \{-1, 0, +1\} \times Q_\infty \times \Sigma_\infty.$$

We need not explicitly specify the state set $Q = Q_M$ and alphabet $\Sigma = \Sigma_M$ of M. Instead, these can be deduced from δ_M. For instance, Σ_M is[12] the set of *non-blank* symbols that occur in either the 2nd or 4th components of an instruction in δ_M. Similarly, the start q_0 and accept q_a states need not be specified as they are universally chosen.

¶23. Enumeration of Turing Machines. PROPOSITION B shows that there exists an enumeration

$$(M_0, M_1, M_2 \ldots). \tag{20}$$

of all STM's. However, we need the enumeration to be effective. This is not difficult, but in view of its importance, we describe it in some detail. Our enumeration satisfies the following two properties:
(P1) For any $i \in \mathbb{N}$, we would also like to easily "decode" the instructions of the ith Turing machine in this enumeration. Why would is this useful? It would allow us to efficiently simulate the ith machine M_i for any i.
(P2) It might seem desirable to make (20) a denumeration. It turns out that for many arguments, we want each $M \in \mathcal{M}$ to appear infinitely often in our enumeration. Hence we prefer to avoid denumeration in a very strong sense.

For reference, we call (P1) the **efficient encoding property**, and (P2) the **recurrence property**, of our enumeration.

[11] The assumptions are also "universal" in a sense related to the concept of "universal Turing machines", to be shown shortly. The existence of such machines depends on these two assumptions.

[12] In case this definition Σ yields an empty set, we simply declare $\Sigma = \{0\}$. This is because an alphabet cannot be empty (by definition).

¶24. Notation Systems for STMs. The function B in (17) is called a (binary) "notation" for natural numbers $\mathbb{N}$. In general, for any set S and any alphabet Σ, we call

$$\nu : \Sigma^* \to S \tag{21}$$

a **notation system** for S if ν is a partial function that is surjective. If $\nu(x) = s \in S$, we call x a **notation** for s. Relative to ν, we call a $x \in \Sigma^*$ **ill-formed** if $\nu(x)$ is undefined, and **well-formed** if $\nu(x)$ is defined. Also we say two strings $x, y \in \Sigma^*$ are **equivalent** if they are both ill-formed, or both well-formed and $\nu(x) = \nu(y)$.

Fix alphabet Σ. Let $\mathcal{M}|\Sigma$ be the set of STM's whose input alphabet is Σ. We want to define a notation system for $\mathcal{M}|\Sigma$. Fix the alphabet

$$\overline{\Sigma} := \Sigma \cup \{0,\ 1,\ Q,\ L,\ R,\ S,\ \#, @\}, \tag{22}$$

where we have assumed that the special symbols $0, 1, Q, L, R, S, \#, @$ are disjoint from Σ. Let $\iota : Q_\infty \cup \Sigma_\infty \to \overline{\Sigma}^*$ such that ι maps Q_∞ bijectively to the set $Q\{0,1\}^*$. Intuitively, the states of Q_∞ will have the notations $Q, Q0, Q1, Q00, Q01, Q10, Q11, Q000, \ldots$. Also, for directions $D \in \{-1, +1, 0\}$, let $\iota(-1) = L$, $\iota(+1) = R$ and $\iota(0) = S$. Also, $\iota(a) = a$ for all $a \in \Sigma$. If $I = (q, a, q', a', D)$ is an instruction for a STM, let

$$\iota(I) = \iota(q)\#\iota(a)\#\iota(q')\#\iota(a')\#\iota(D).$$

If $J = I_1 I_2 \cdots I_m$ is a sequence of instructions, let $\iota(J) = \iota(I_1)@\iota(I_2)@\cdots@\iota(I_m)$. If the transition table of a STM $M \in \mathcal{M}|\Sigma$ is $\{I_1, \ldots, I_m\}$, then we call $\iota(I_1 I_2 \cdots I_m)$ a notation for M. Since we can reorder the instructions in $m!$ many ways, there are that many notations for M.

We now define a notation system for STM's:

$$\gamma : \overline{\Sigma}^* \to \mathcal{M}|\Sigma$$

where $\gamma(x) = M$ iff x is a notation for M. The function γ is surjective since every STM has a notation. It is also a partial function since $\gamma(x)$ is undefined in case x is not the notation of a sequence of instructions. However, we shall let $\gamma(x)$ denote a special STM that rejects all its inputs. Without loss of generality this is M_0.

We can also interpret the notation system as a function $\gamma : \mathbb{N} \to \mathcal{M}|\Sigma$, under the isomorphism $\mathbb{N} \simeq \overline{\Sigma}^*$. Then, γ defines the enumeration (20) using the convention

$$(M_0, M_1, M_2, \ldots) = (\gamma(0), \gamma(1), \gamma(2), \ldots).$$

By our convention, if n is ill-formed then $\gamma(n)$ denotes the special machine M_0. Thus M_0 is repeated infinitely often in this enumeration.

Property (P1) is satisfied: given a notation x for a machine, we can easily decode the instructions in x. In particular, it is easy to detect well-formed strings:

LEMMA 23. *The set of well-formed strings for γ is a regular language over $\overline{\Sigma}$.*

We are almost done except for property (P2). There are many ways to achieve this, but an easy way is to declare that strings of the form $x = y@@z$ is equivalent to y. In short, we ignore the part of the string x after the first appearance of the sequence '@@'. Thus, for any $M \in \mathcal{M}|\Sigma$, there are infinitely many notations for M.

To decide if x, y are equivalent, we have to determine if (after discarding any suffixes of the form $@@\overline{\Sigma}^*$) the set of instructions in x and y are the same. This is easily done by sorting the instructions. Thus:

LEMMA 24. *There is a STM that to decides whether two input strings x and y are equivalent notations.*

A notation system for $\mathcal{M}|\Sigma$ that satisfies properties (P1) and (P2) will be called a **Gödel numbering**. We call $i \in \mathbb{N}$ a **Gödel number** of $M \in \mathcal{M}$ if $\gamma(i) = M$; by extension, we may also call i a Gödel number of the language $L(M)$.

In general, for any notation system $\nu : \mathbb{N} \to S$, we can define a total injective function

$$\overline{\nu} : S \to \mathbb{N}$$

where $\overline{\nu}(s) = \min\{n \in \mathbb{N} : \nu(n) = s\}$. Note that $\overline{\nu}$ is a partial inverse of ν. We may call $\overline{\nu}$ an **encoding function** for S.

In particular, for every STM $M \in \mathcal{M}|\Sigma$, we have its encoding $\overline{\gamma}(M)$ as a natural number.

Exercise 0.3.5.44: Prove Lemma 23 and Lemma 24. □

Exercise 0.3.5.45: Let $\mathcal{M}_0$ denote the set of all deterministic STM's. Construct a Gödel pair (α_0, γ_0) for $\mathcal{M}_0$. □

End Exercises

0.3.6 Universal Machines

We introduce a variant of Turing machines called **universal Turing machines** (UTM's). A UTM is similar to a simple Turing machine U but it has two tapes. One tape behaves exactly like the tape of a STM, and is called the **work tape**. The other tape is called the **index tape**. The index tape is a **read only** tape (writing not allowed). There are corresponding **work head** and **index head** that scans a cell on the respective tape. Intuitively, i is the number of a STM. The transition table δ_U of U is a finite set

$$\delta_U \subseteq Q_\infty \times \Sigma_\infty^2 \times Q_\infty \times \Sigma_\infty \times \{-1, 0, 1\}^2.$$

A typical instruction of δ_U is

$$(q, a, b \to q', b', D_0, D_1)$$

with precondition that the machine is in state q, scanning a on the index tape, and scanning b on the work tape. Upon executing the instruction, the new state is q', b is changed to b', while the work head and index head move in directions D_0, D_1 (respectively). Thus δ_U defines implicitly an **index alphabet** Σ_0 and a **tape alphabet** Σ_1. The **inputs** of the UTM are the pairs $(i, x) \in \Sigma_0^* \times \Sigma_1^*$. On any input, the UTM begins in state q_0 and executing (nondeterministically) instructions. We say that it accepts its input if some executing sequence leads to the accept state q_a. As usual, $L(U) = \{(i, x) : U \text{ accepts } (i, x)\}$. Note that $L(U)$ is no longer a language, but a binary relation. Let $|\Sigma_0| = k \geq 1$ and identify Σ_0^* with $\mathbb{N}$ via the k-adic notation. Hence we have

$$L(U) \subseteq \mathbb{N} \times \Sigma_1^*.$$

Define the language class $K(U) := \{L_i(U) : i \in \mathbb{N}\}$ where $L_i(U) := \{w \in \Sigma_1^* : (i, w) \in L(U)\}$, We say U **accepts the class** $K(U)$ of languages. Practically all the definitions for STM's transfer naturally to UTM's. Thus it is clear what we mean by a "deterministic UTM", or the notation "$U(i, w) \uparrow$", etc.

It is important to note that $K(U)$ is a collection of languages over some fixed alphabet $\Sigma_1 \subseteq \Sigma_\infty$. We write "$K|\Sigma$" for denote the restriction of the class K to languages over Σ. Thus, $K(U) = K(U)|\Sigma_1$. We come to the key definition: let Σ be a non-empty alphabet. A universal Turing machine U is **universal** for a class K if it accepts K.

Theorem 25 (Existence of Universal Turing Machines). *Let $\Sigma \subseteq \Sigma_\infty$ be any alphabet. There exists a universal Turing machine $U^{(2)}$ that accepts $RE|\Sigma$. Furthermore, the following properties hold:*

- *$U^{(2)}$ is deterministic.*
- *If $M \in \mathcal{M}|\Sigma$ and $\overline{M}$ is a notation for M, then $U^{(2)}(\overline{M}, w) = M(w)$ for all $w \in \Sigma^*$.*

Proof. Suppose (i, w) is an input to $U^{(2)}$. Without loss of generality, interpret i as a Gödel number of some STM $M_i \in \mathcal{M}_0$, and we want to simulate $M_i(w)$. We must decode i into its instructions, but in case i is ill-formed, it is just nonsense. So we need to detect this separately. But this is easy, since set of ill-formed Gödel numbers is regular. When $U^{(2)}$ detects an ill-formed i, it simply rejects. So assume otherwise. To simulate $M_i(w)$, deterministically. it stores a list of configuration $C_1, \ldots, C_m$ of $f^{-1}(i)$ on its work tape. This follows the proof of theorem 13, except that to expand a configuration C, it must use the index tape to look up the next executable instructions. This completes our description of $U^{(2)}$. It is clear that M_i accepts w iff $U^{(2)}$ accepts (i, w), and also $U^{(2)}(i, w) = M_i(w)$. **Q.E.D.**

It might appear that by considering $RE|\Sigma$ instead of RE, we are restricting the generality of our results. But this is not true. Indeed, $RE|\Sigma$ essentially has the full power of RE. An Exercise below will illustrate this remark.

Exercises

Exercise 0.3.6.46: Suppose we want a single UTM to accept all of RE, not just $RE|\Sigma$. This requires a suitable encoding for all of RE inside a fixed alphabet such as $\{0,1\}$. Give the necessary formation and definitions, so that Theorem 25 can be restated with RE instead of $RE|\Sigma$. □

END EXERCISES

0.3.7 Partial Recursive Functions

We make connection to the standard setting for recursive function theory. This shifts our focus from languages to subsets of $\mathbb{N}$, and from word functions to number theoretic functions. Instead of "languages" (subsets of Σ^*), we speak of "sets" (i.e., subsets of $\mathbb{N}$).

In this section, we fix $U^{(2)}$ to be the deterministic universal Turing machine guaranteed by Theorem 25, for some fixed Σ. In fact, we may assume $\Sigma = \{0,1\}$, and start to identify $RE|\Sigma$ with RE, and $\mathcal{M}|\Sigma$ with $\mathcal{M}$.

For any STM $M \in \mathcal{M}$, we write

$$U_i^{(2)} \simeq M, \tag{23}$$

if i is the Gödel number of M. Using a k-adic bijection, a language $L \subseteq \Sigma^*$ can be equated with $A \subseteq \mathbb{N}$, denoted $L \simeq A$.

We define a **r.e. set** (resp., **recursive set**) to be $A \subseteq \mathbb{N}$ such that $A \simeq L$ for some **r.e.** (resp., **recursive**) language $L \subseteq \Sigma^*$. We may reinterprete RE to be the class of all r.e. sets. Similarly, REC is interpreted as the class of all recursive sets. Thus,

$$REC \subseteq RE \subseteq 2^{\mathbb{N}}. \tag{24}$$

Define

$$W_i := \{x \in \mathbb{N} : U^{(2)}(i,x) \downarrow\}.$$

Clearly, each W_i is r.e.. The converse is also true.

THEOREM 26. *A set $A \subseteq \mathbb{N}$ is r.e. iff $A = W_i$ for some $i \in \mathbb{N}$. Alternatively,*

$$\{W_i : i \in \mathbb{N}\} \simeq RE.$$

Proof. One direction is trivial, since the W_i's are r.e.. In the other direction, we assume A is r.e. and show that $A = W_i$ for some i. Since A is r.e., let A be accepted by some STM M. It is easy to construct STM M' which, on input w, will halt and accept iff M accepts w; otherwise $M'(w) \uparrow$. Then $L(M') = L(M)$. But $M' \simeq U_i^{(2)}$ for some i. Then $W_i = L(M')$ by the construction of M'. This shows $W_i = A$. **Q.E.D.**

To compute number theoretic functions, we view $U^{(2)}$ as a transducer that computes a partial function

$$\Phi : \mathbb{N}^2 \to \mathbb{N}$$

as follows: on input (i,w), we declare the output of $U^{(2)}$ to be the non-blank word v that is scanned by its work head when $U^{(2)}$ halts; if $U^{(2)}(i,w) \uparrow$ then $\Phi(i,w)$ is undefined. We may assume that the tape alphabet is also $\{0,1,\sqcup\}$ and v is interpreted as a dyadic number; if the work head is scanning the blank symbol then $v = \epsilon$. Let $\phi_i(w) = \Phi(i,w)$ so that $U^{(2)}$ can be viewed as computing the set

$$\Phi := \{\phi_i : i \in \mathbb{N}\} \tag{25}$$

of number theoretic functions.

We come to a key definition: a function f is said to be **partial recursive** iff f belongs to the set Φ. If f is a total function and also partial recursive, we call it a **total recursive function**. This definition implicitly relies on a functional analogue of theorem 25 on Universal Turing machines.

¶25. Dovetailing. Let us define

$$E_i := \{\phi_i(w) : w \in \mathbb{N}, \phi_i(w) \downarrow\}.$$

We have another characterization of r.e. sets using E_i. The proof requires a technique[13] called **dovetailing**. Suppose we want to simulate a denumerable number of computation paths,

$$\pi_0, \pi_1, \pi_2, \ldots$$

[13]Dovetailing is a technique in carpentry for joining together the ends of two wooden beams without using nails.

Then we can define a single computation path π that "dovetails" these together: if $\pi_i(j)$ is the jth configuration of π_i, then

$$\begin{aligned}
&\text{Phase 1}: \quad \pi(0) = \pi_0(0);\\
&\text{Phase 2}: \quad \pi(1) = \pi_0(1), \pi(2) = \pi_1(0);\\
&\text{Phase 3}: \quad \pi(3) = \pi_0(2), \pi(4) = \pi_1(1), \pi(5) = \pi_2(0);\\
&\text{Phase 4}: \quad \pi(6) = \pi_0(3), \ldots
\end{aligned}$$

If π_i is a finite path, we treat it as an infinite computation by repeating the last configuration. In actual applications (see next), we can usually stop the dovetailed computation π when some condition is detected.

THEOREM 27. *A set $B \subseteq \mathbb{N}$ is r.e. iff $B = E_i$ for some $i \in \mathbb{N}$.*

Proof. Suppose B is r.e.. To show $B = E_i$, let B be accepted by some STM M. Construct the STM M' that on input w, will simulate $M(w)$. If $M(w) \downarrow$ then M' outputs w (to achieve this, M' would have to save a copy of w at the start of the computation). It follows that if $M' \simeq U_i^{(2)}$ then $\phi_i(w) = w$ iff $w \in L(M) = B$. Thus $E_i = B$.

Conversely, we show that each E_i is r.e.. We may assume that the i is the code for a deterministic STM. We construct M to accept E_i as follows: on input w, we simulate $U^{(2)}(i, j)$ for $j = 0, 1, 2, \ldots$, using the dovetailing process above. Whenever $U^{(2)}(i, j)$ halts, we examine the output $\Phi(i, j)$. If this output is equal to w, we accept. Otherwise, we continue the dovetailing. Clearly, if $w \in E_i$, then $w = \Phi(i, j)$ for some j and M will accept at some point in this simulation. **Q.E.D.**

If A is a non-empty r.e. set, the preceding characterization $A = E_i$ can be strengthened so that ϕ_i is total (Exercise). We have a corresponding characterization of recursive sets.

THEOREM 28. *A set $A \subseteq \mathbb{N}$ is infinite and recursive iff $A = E_i$ for some i such that ϕ_i is increasing and total.*

Proof. Let ϕ_i be increasing and total. Clearly, E_i is infinite. To show that E_i is recursive, we give a decision procedure for E_i: on input $n \in \mathbb{N}$, we compute $\phi_i(j)$ for $j = 0, 1, 2, \ldots$. We do not need to dovetail these computations since ϕ_i is total. If $\phi_i(j) = n$ for some j, then we accept. Otherwise, we reject when $\phi_i(j) > n$ is reached for some j. Clearly, we accept iff $n \in E_i$.

Conversely, suppose A is recursive. We construct an increasing and total recursive function f as follows: on input n, we check each j ($j = 0, 1, 2, \ldots$) to see if $j \in A$. If j is the nth time that we discover a word of A, we output j. It is clear that $f(n)$ is increasing and total. **Q.E.D.**

Thus, W_i and E_i are the proper domain and proper range of the partial recursive $\phi_i : \mathbb{N} \to \mathbb{N}$. Restricting the nominal domain $\mathbb{N}$ and nominal range $\mathbb{N}$ of ϕ_i to W_i and E_i, respectively, we obtain a total surjective function,

$$\phi_i : W_i \to E_i.$$

It may be helpful to see the symbols W, E as mnemonic for "West" and "East", respectively. In case ϕ_i is total, E also suggests an "Enumeration", since $E_i = \{\phi_i(0), \phi_i(1), \phi_i(2), \ldots\}$. In fact, for every non-empty E_i, there is a total recursive function f_i whose range is E_i (see Exercise).

Here is another application of dovetailing:

THEOREM 29. *A set $A \subseteq \mathbb{N}$ is infinite r.e. if and only if there is a recursive denumeration $d : \mathbb{N} \to A$.*

Proof. If d exists, it is clear that A is infinite and r.e.. Conversely, assume A is infinite and r.e.. Then $A = E_i$ for some i. Thus $x \in A$ iff $\phi_i(j) = x$ for some $j \in \mathbb{N}$. We describe a recursive function $d : \mathbb{N} \to A$. On input j, we compute $d(j)$ as follows: dovetail over the following infinite sequence of computation paths,

$$\phi_i(0), \phi_i(1), \phi_i(2), \ldots.$$

Some of these computations halt, and produce outputs. We keep track of distinct outputs: say w_0 is the first output, and in general w_n ($n \geq 1$) is the nth distinct output. Define $d(j) = w_j$. We check that d is recursive and a denumeration of A. **Q.E.D.**

¶26. Encoding Relations as Languages. Recall that we use $\langle w_1, \ldots, w_k \rangle = w_1 \# w_2 \# \cdots \# w_k$ ($w_i \in \{0,1\}^*$) to encode an element of $\mathbb{N}^k$. So any relation $R \subseteq \mathbb{N}^k$ can be equated with language L_R over $\{0, 1, \#\}$, with as associated bijection

$$R \simeq L_R.$$

We say the relation R is **partial recursive** or **total recursive** if L_R is partial recursive or total recursive (respectively).

We give another characterization of r.e. sets using a recursive relation T:

THEOREM 30. *There exists a recursive $T \subseteq \mathbb{N}^3$ such that a set $A \subseteq \mathbb{N}$ is r.e. iff there exist $i \in \mathbb{N}$ such that*

$$A = \{x \in \mathbb{N} : (\exists t \in \mathbb{N})[(i, x, t) \in T\}.$$

Proof. One direction is easy: a set of the form $\{x \in \mathbb{N} : (\exists t \in \mathbb{N})[(i, x, t) \in T\}$. is clearly r.e. because T is recursive. Conversely, we define T to comprise the triples (i, x, t) such that $U_i(x)$ accepts in $\leq t$ steps. Clearly T is recursive. Moreover, every recursive set A has the form indicated. **Q.E.D.**

EXERCISES

Exercise 0.3.7.47: Students often forget the difference between computable denumeration and a non-computable one. Let $A \in RE$ be an infinite set. Construct a denumeration $f : A \to \mathbb{N}$ which is not recursive. □

Exercise 0.3.7.48: Show that for every non-empty E_i, there is a total recursive function ϕ_j such that $\phi_j(\mathbb{N}) = E_i$.

SOLUTION: Consider the STM M that, on input x, simulates the each of the following computations

$$\phi_i(0), \phi_i(1), \phi_i(2), \ldots, \phi_i(x)$$

for x steps. It outputs the $\phi_i(y)$ where y is the largest argument among $0, 1, \ldots, x$ such that $\phi_i(y)$ halts. It is easy to see that M is total recursive, and its output range is exactly E_i.

□

Exercise 0.3.7.49: Prove that if A is r.e. but not recursive, then it contains an infinite subset that is recursive. HINT: recall the construction that if $B \subseteq A$ is recursive, then $B = E_i$ for some total increasing function ϕ_i. □

Exercise 0.3.7.50: (i) Show that L is r.e. iff L is generated by some deterministic STM.
(ii) Show that L is recursive iff L is generated by some deterministic STM in lexicographic order. □

END EXERCISES

0.3.8 Decision Problems

If L is any language, the computational problem of designing an algorithm to accept L is called the **decision problem** for L. We introduce several decision problems related to STMs. Among them we will find non-r.e. languages that are considered more natural and explicit, unlike the ones obtained by cardinality arguments. These problems will exploit the existence of universal machines.

¶27. Three Decision Problems. Let $U^{(2)}$ be the UTM in theorem 25, with $\Sigma = \{0, 1\}$. Recall that for any strings x, y, we have the notation $\langle x, y \rangle$ for $x \# y$, where $\#$ is a fixed symbol not in x, y. Define the languages

$$\text{WORD} = \{\langle i, w \rangle : i, w \in \{0,1\}^*, (i, w) \in L(U^{(2)})\}, \tag{26}$$

$$\text{HALT} = \{\langle i, w \rangle : i, w \in \{0,1\}^* : U^{(2)}(i, w) \downarrow\}, \tag{27}$$

$$\text{DIAG} = \{w \in \{0,1\}^* : U^{(2)}(w, w) \downarrow\}. \tag{28}$$

The decision problems for WORD, HALT, DIAG are called (respectively) the **word problem**, the **halting problem**, the **diagonal problem** for $U^{(2)}$. We have ordered these three problems so that each is seen to be a simpler version of the previous.

THEOREM 31. WORD, HALT, DIAG $\in RE$.

Proof. First we show WORD $\in RE$. This amounts to constructing a STM M_2 that emulates $U^{(2)}$. On input $\langle i, w\rangle$, M_2 simulates the actions of $U^{(2)}$ on input (i, w) in a step-by-step manner. M will accept iff $U^{(2)}$ accepts. Thus $L(M_2) =$ WORD. Similarly, DIAG and HALT can be shown to be in RE be simple variations of this simulation. **Q.E.D.**

THEOREM 32. WORD, HALT, DIAG $\notin REC$.

Proof. Suppose WORD is decided by a deterministic halting STM M. We construct another halting STM M' such that M' does the opposite of M:

M' on input i:
1. Simulate M on $\langle i, i\rangle$.
2. Accept iff M reject.

Note that M' is halting since M is halting. Since M' can clearly be constructed from M, we conclude that $M' \simeq U_j^{(2)}$ for some j. Now consider what happens when M' is given the input j. (a) If M' accepts j, it is because M rejects $\langle j, j\rangle$. By definition of M, this means $U^{(2)}(j, j)$ rejects. But this means M' rejects j, contradiction. (b) If M' rejects j, it must be because M accepts $\langle j, j\rangle$. By definition of M, this means $U^{(2)}(j, j)$ accepts. But this means M' accepts j, contradiction. **Q.E.D.**

COROLLARY 33. co-WORD, co-HALT, co-DIAG *are not r.e..*

Proof. Use the fact that $REC = RE \cap$ co-RE. If co-WORD $\in RE$, then WORD would be recursive, contradicting the previous theorem. The other two cases are similar. **Q.E.D.**

0.3.9 Reducibility

In the last subsection, we show certain decision problems to be non-recursive, using a diagonalization argument. If we want to show some given language A is non-recursive, we could try to use a diagonalization argument (it is the only technique we know so far). Instead, this section shows an alternative method, by "reducing" some set B to A. If B is known to be non-recursive, then we conclude that A is non-recursive. This approach can be more convenient if B is properly chosen.

¶28. Many-one Reducibility. If $A, B \subseteq \mathbb{N}$, and $t : \mathbb{N} \to \mathbb{N}$, we say that A is **many-one reducible** to B via t, denoted

$$A \leq_m B \text{ (via } t),$$

if, for all $i \in \mathbb{N}$, $i \in A$ iff $t(i) \in B$. If the function t is total recursive, we also say that A is **many-one recursively reducible** to B. For short, we say **m-reducible** for "many-one reducible". Clearly, we have:

LEMMA 34. *If $A \leq_m B$ and $B \leq_m C$ then $A \leq_m C$.*

Since $\leq_m$ is a reflexive relation, it follows that $\leq_m$ induces an equivalence relation on the set of all languages. The equivalence classes are called **m-degrees**. The m-degrees are partially ordered in the obvious way:

LEMMA 35.
(a) If A is r.e., then $A \leq_m$ WORD.
(b) WORD, HALT, DIAG *are all m-equivalent.*

Proof. (a) Assume $A = W_i$ for some i. Then $w \in A$ iff $\langle i, w\rangle \subseteq$ WORD. Since $t : w \mapsto \langle i, w\rangle$ is total recursive, this proves $A \leq_m$ WORD (via t).

(b) To see that WORD $\leq_m$ HALT, consider the machine $M_{i,w}$ that on input x will ignore x and run $U_i^{(2)}$ on w. If $U_i^{(2)}$ accepts then $M_{i,w} \downarrow$, else $M_{i,w} \uparrow$. The total recursive map $t(\langle i, w\rangle) = \overline{\gamma}(M_{i,w})$ (encoding of $M_{i,w}$) shows that $\langle i, w\rangle \in$ WORD iff $t(\langle i, w\rangle) \in$ HALT. It is similarly easy to show that

$$\text{HALT} \leq_m \text{DIAG} \leq_m \text{WORD}.$$

Q.E.D.

This result says that Halt, WORD and DIAG are all "complete" for the class RE.

LEMMA 36. *Let A be recursively m-reducible to B.*
(a) If B is recursive, so is A.
(b) If B is r.e., so is A.
(c) If B is co-r.e., so is A.

Proof. Each part is shown the same way. Let us just show (c). Let $A \leq_m B$ via a total computable t. Suppose B is co-r.e.. Let M accept co-B. We can now accept co-A as follows: on input x, we compute $t(x)$, and run M on $t(x)$. We accept iff M accepts. This proves that co-$A \in RE$, as desired. **Q.E.D.**

The lemma says that if $A \leq_m B$, then any "upper bound" on B is also an upper bound for B. By choosing A appropriately, the lemma is the basis for showing that a set B is non-recursive, non-r.e. or non-co-r.e.. This is illustrated in proof of the next theorem.

¶29. Rice's Theorem. Suppose K is the class of finite subsets of $\mathbb{N}$. Let $L(K) \subseteq \mathbb{N}$ denote the set of all Gödel numbers of STM's that accept a finite subset. We shall show that $L(K)$ is non-recursive. In fact, this conclusion holds for almost any reasonable choice for K. This is the content of Rice's Theorem.

As another example, we could let $K = \{HALT\}$, or $K = \{\epsilon\}$. However, K cannot be chosen to be ϵ, since this implies $L(K) = \epsilon$, which is certainly a recursive set.

To formalize Rice's theorem, let K be any class of r.e. sets. Relative to $U^{(2)}$, we define the set

$$L(K) := \{i \in \mathbb{N} : W_i \in K\}.$$

THEOREM 37 (Rice). *If $K \neq \emptyset$ is a proper subset of RE, then $L(K)$ is not recursive.*

Proof. We m-reduce the problem DIAG to $L(K)$. Pick any $A \in K$ (since K is non-empty) and pick any $B \notin K$ (since K is a proper subset of r.e. sets). First we assume that $B = \epsilon$. Assume M_A and M_B accept A and B respectively. For any i, construct the STM M_i as follows: on input w,

$$M_i(w) = \begin{cases} M_A(w) & \text{if } \ i \in W_i \\ \uparrow & \text{else.} \end{cases}$$

Note that M_i is well-defined: first we check if $i \in W_i$. If this does not halt, then M_i does not halt. If it halts, we then simulate an acceptor of A on input w. What does M_i compute?

If $i \in W_i$ then M_i accepts A.

If $i \notin W_i$ then M_i accepts $B = \epsilon$.

It is clear from our construction of M_i, there is a total recursive function $t(i)$ that gives the Gödel number $\overline{\gamma}(M_i)$ of M_i. Thus:

If $i \in W_i$ then $L(M_i) = A \in K$ and so $t(i) \in L(K)$.

If $i \notin W_i$ then $L(M_i) = B \notin K$ and so $t(i) \notin L(K)$.

Thus DIAG is recursively m-reducible to $L(K)$ via t. Since we know that DIAG is not recursive, $L(K)$ is not recursive by the previous lemma.

We had assume $\epsilon \notin K$. Suppose $\epsilon \in K$. Then let $A = \epsilon$ and let B be any set not in K. We can similarly construct M_i such that $L(M_i) \in K$ iff $i \in W_i$. Again, we conclude $L(K)$ is non-recursive. **Q.E.D.**

If $K \subseteq RE$, then the decision problem of $L(K)$ is known as a problem of deciding "property K": the ith STM has the property iff $i \in L(K)$. The property is trivial if $K = \emptyset$ or $K = RE$. Thus the Rice Theorem concerns non-trivial properties of RE.

¶30. Effective Machine Construction. The preceding proof illustrates an important style of argument in computability. Although it is ultimately a reduction argument, we are interested in the specific way in which we derive at the total recursive function $t(i)$.

(I) First, we describe a general STM construction $i \mapsto M_i$.

(II) Then we assert that there is a total recursive function t such that $t(i)$ is the Gödel number of M_i: $\gamma(t(i)) = M_i$.

Assertion (II) relies on the ability of the student to "program" Turing machines. Sometimes, it is quoted as a consequence of Church's Thesis. In any case, we are trying to avoid an explicit construction of the Gödel number $t(i)$. Details of an explicit construction would tedious. Here is an alternatively view of this procedure, commonly seen in recursive function theory:

(I') First define a computable multi-parameter functions $g(i, w)$ whose value is computed by $M_i(w)$.

(II') Then invoke a technical theorem called the **S-m-n theorem** or **parameter theorem** to infer the total recursive function $t(i)$ that gives the Gödel number of M_i. Note that the argument w is omitted in $t(i)$.

For reference, such a 2-step procedure that leads to the total recursive $t(i)$ is called an **effective machine construction**. Let us see this procedure at work again in the next result.

¶31. Rice-Shapiro Theorem. A natural question arises from Rice's theorem. Suppose $L(K)$ is not recursive. Could $L(K)$ be r.e.? The answer depends on K:
(i) Let K_0 be the set of all r.e. sets such that $A \in K_0$ iff $\epsilon \in A$. Then $L(K_0)$ is r.e..
(ii) Let $K_1 = \{A_1\}$ where A_1 is any r.e. set. Then $L(K_1)$ is not r.e..
The following theorem characterizes the situations when $L(K)$ is r.e..

THEOREM 38. *Let $K \subseteq RE$. If $L(K)$ is r.e. then for any r.e. set A, the following statements are equivalent:*
(a) $A \in K$.
(b) There exists a finite subset $B \subseteq A$ such that $B \in K$

Proof. If (a) and (b) are not equivalent, we will prove that $L(K)$ is not r.e., by reducing co-DIAG to $L(K)$. Let M_A be a STM that accepts A. We need to consider two possibilities.

(A) Suppose (a) holds but not (b). Then $A \in K$ but for all finite $B \subseteq A$, $B \notin K$. Consider how we might answer the query: is $i \in$ co-DIAG? Now $i \in$ co-DIAG is equivalent to $\phi_i(i)$ does not in t steps for all $t \geq 0$. Consider the following STM M_i: on input t, we run $U_i(i)$ for t steps. If U_i halt in t steps, we loop. Otherwise, we simulate M_A on input t, accepting iff M_A accepts.

Note that if $U_i(i) \uparrow$, then M_i will accept A (so $L(M_i) \in K$); if $U_i(i) \downarrow$, then M_i accepts a finite subset of A (so $L(M_i) \notin K$). Hence if $s(i)$ is the Gödel number of M_i, we have that $s(i) \in L(K)$ iff $i \in$ co-DIAG. Since the function $i \mapsto s(i)$ is clearly total computable from our description of M_i, this shows that co-$\mathsf{Halt} \leq_m L(K)$ (via s).

(B) Suppose (b) holds but not (a). Then $A \notin K$ but there is some finite $B \subseteq A$ such that $B \in K$. Again, to decide if $i \in$ co-DIAG, we construct the following STM N_i: on input t, we first check if $t \in B$. If so, we accept. Otherwise, we simulate U_i on i. If $U_i(i) \downarrow$, we simulate M_A on t, accepting iff M_A accepts.

Note that if $U_i(i) \uparrow$, then N_i accepts B; if $U_i(i) \downarrow$, then N_i accepts A. Again, this construction reduces co-DIAG to $L(K)$. **Q.E.D.**

This theorem tells us that if $L(K)$ if r.e., then K can be reduced to a "core" K^* of finite sets. That is, $K^* \subseteq \overline{2}^{\mathbb{N}}$ such that for any r.e. A, we have $A \in K$ iff A is the extension of some $B \in K^*$.

EXERCISES

Exercise 0.3.9.51: Verify the claims about the r.e. nature of $L(K)$ for the examples (i) $K = K_0$ and (ii) $K = K_1$ described above.

SOLUTION: (i): Easy.
(ii): We reduce a non-r.e. co-HALT to $L(K_0)$. Assume N accepts L_0. For any $i \in \mathbb{N}$, construct M_i such that on input x: runs $U_i(i)$ for x steps. If U_i does not accept, we simulate N on x. If U_i accepts, we "simulates some language that is guaranteed to be different from L_0".
If $t(i)$ is the Gödel number of M_i, we see that $i \notin$ HALT iff $\phi_i(i) \uparrow$ iff $t(i) \in L(K_0)$.

□

Exercise 0.3.9.52: Students often assume that if A is r.e., then any enumeration $\alpha : \mathbb{N} \to A$ is recursive. To dispel this idea, prove the following: if A is infinite and r.e., then there exists a non-recursive denumeration

$$f : \mathbb{N} \to A.$$

REMARK: If A is non-recursive, there is a simple argument. So it may be enough to assume A is recursive.

HINT: Let $\alpha : \mathbb{N} \to A.$ be a recursive denumeration of A. Now split A into $A_0 \uplus A_1$, and define $f : \mathbb{N} \to A$ such that $f(i) \in A_0$ iff $i \in$ HALT.

SOLUTION: If A is non-r.e. then $f(i)$ can simply be $i+1$st largest number in A. It is easy to see that $f(i)$ is not r.e., as observed in the REMARK.
Define $\alpha_0(i) = \alpha(2i)$ and $\alpha_1(i) = \alpha(2i+1)$. Thus $\alpha_j : \mathbb{N} \to A_j$ (for $j = 0, 1$) is a denumeration. We know that there exists denumerations

$$h : \mathbb{N}\text{HALT}, \qquad \overline{h} : \mathbb{N}\text{co-HALT}.$$

Finally define

$$f(i) = \begin{cases} \alpha_0(h^{-1}(i)) & \text{if } i \in \text{HALT} \\ \alpha_1(\overline{h}^{-1}(i)) & \text{else.} \end{cases}$$

Then f is a recursive denumeration of A. It remains to show that if f is recursive, then HALT is recursive. This is because $i \in$ HALT iff $f(i) \in A_0$. To check if $f(i) \in A_0$, we simply dovetail the computations

$$\{\alpha_0(j), \alpha_1(j) : j \in \mathbb{N}\}$$

until we see $f(i)$ being output. Then we know whether $f(i) \in A_0$ or not.

□

Exercise 0.3.9.53: Sharpen the Rice-Shapiro theorem into a characterization of those K such that $L(K)$ is r.e..

SOLUTION: Should the core K_0 be r.e. or recursive? Say K_0 is r.e.. Then to decide if $i \in L(K)$, we search for a finite $\theta \in K_0$ and $\theta \subseteq W_i$.

□

Exercise 0.3.9.54: Here is a formulation of the $S-m-n$ Theorem: Let $S : \mathbb{N}^2 \to \mathbb{N}$ be a partial computable function. Then there is a total computable function $t : \mathbb{N} \to \mathbb{N}$ such that $S(m,n) = \phi_{t(m)}(n)$. Recall that ϕ_i is the ith partial recursive function in (25). □

SOLUTION: Let M be a STM that computes S in the sense that $M(\langle i, j\rangle) = S(i,j)$. Construct the TM N_i such that on input x, N_i calls $M(\langle i, x\rangle)$. Let $t(i)$ be the Gödel number of N_i. Thus, for all i, j, we have $S(i,j) = \phi_{t(i)}(j)$.

End Exercises

0.3.10 Second Recursion Theorem

In this and the next subsections, we prove two recursion theorems. We begin with the easier, Second Recursion Theorem. This is also known as the (plain) Recursion Theorem.

Theorem 39 (Recursion Theorem). *For any total recursive function h, there exists $e \in \mathbb{N}$ such that $W_e = W_{h(e)}$. We call e a* **fixed point** *of h.*

Proof. The proof is a simple application of the effective machine construction procedure. For any i, we construct the following STM M_i: on input x, first compute $\phi_i(i)$. If $\phi_i(i) \downarrow$, we compute $h(\phi_i(i))$ and then finally accept iff $x \in W_{h(\phi_i(i))}$.

Note that $i \mapsto M_i$ is an effective construction. Let $t(i)$ be the Gödel number of M_i. Hence the function $i \mapsto t(i)$ is total recursive. We may assume M_i is deterministic.

What does M_i compute? Then for all $i, x \in \mathbb{N}$, we have

$$M_i(x) \downarrow \iff x \in W_{h(\phi_i(i))}.$$

In this equivalence, it is understood that if $\phi_i(i) \downarrow$ then both sides are false, i.e., $M_i(x) \uparrow$ and $x \notin W_{h(\phi_i(i))}$ for all x. Thus

$$x \in W_{t(i)} \iff x \in W_{h(\phi_i(i))}$$

Since $t(\cdot) = \phi_d(\cdot)$ for some d, this is equivalent to

$$x \in W_{\phi_d(i)} \Longleftrightarrow x \in W_{h(\phi_i(i))}.$$

Choose $i = d$, we get that for all $x \in \mathbb{N}$,

$$x \in W_{\phi_d(d)} \Longleftrightarrow x \in W_{h(\phi_d(d))}.$$

The theorem follows by setting $e = \phi_d(d)$. **Q.E.D.**

Many diagonalization arguments can be succinctly reduced to an application of this theorem. We illustrate this with another proof of the Rice Theorem (theorem 37). Let $K \subseteq RE$ be a non-trivial property of r.e. sets, so there is some $a \in L(K)$ and $b \notin L(K)$. Let

$$h(x) := \begin{cases} a & \text{if } x \notin L(K) \\ b & \text{if } x \in L(K) \end{cases}$$

Observe that

$$x \in L(K) \Longleftrightarrow h(x) \notin L(K). \tag{29}$$

By way of contradiction, assume $L(K)$ is recursive. Then the function $h(x)$ is total recursive. If e is the fixed point value of h, it follows that $W_e = W_{h(e)}$. This means $e \in L(K)$ iff $h(e) \in L(K)$, contradicting (29). Hence $L(K)$ is not recursive.

EXERCISES

Exercise 0.3.10.55: Prove the following functional form of the Second Recursion Theorem: for all total recursive h, there is an e such that $\phi_e = \phi_{h(e)}$. □

Exercise 0.3.10.56: Which of the following follows from the (functional form of) the Second Recursion Theorem?
(i) There is a number e such that $\phi_e(x) = x^e$.
(ii) There is a number e such that $\phi_e(x) = \phi_{2e}(x)$.
(iii) There is a number e such that $W_e = \{e\}$.

SOLUTION: (i) and (iii): see Cutland, pages 202-3. (ii) is not possible.
OTHERS: If h is total recursive, there are arbitrarily large e such that $\phi_e = \phi_{h(e)}$.

□

END EXERCISES

0.3.11 First Recursion Theorem

The First Recursion Theorem is deeper than the second, involving some interesting concepts. Its applicability is arguably narrower.

¶32. Recursive operators. A total function of the form

$$\Psi : RE \to RE$$

is called an **operator**. A number theoretic function $h : \mathbb{N} \to \mathbb{N}$ is said to be **extensional** if for all $i, j \in \mathbb{N}$, if $W_i = W_j$ then $W_{h(i)} = W_{h(j)}$. We call Ψ a **recursive operator** in case there is an extensional total recursive $h : \mathbb{N} \to \mathbb{N}$ such that for all $i \in \mathbb{N}$, $\Psi(W_i) = W_{h(i)}$. We say Ψ is **based** on h.

Example: The concept of an extensional total recursive h might appear hard to realize in practice, but it need not be. Consider the operator Ψ where $\Psi(A) = \{x + 1 : x \in A\}$ for all $A \in RE$. To show that Ψ is recursive, we construct a total recursive h that, on any input i yields the Gödel number of the following STM M_i: on input x, loop if $x = 1$. Otherwise, run U_i on $x - 1$, and halts iff U_i halts. It is easy to see that h is extensional because $W_{h(i)}$ depends only on W_i, and not on the choice of the Gödel number i.

¶33. Finiteness and Continuity. Let $\Psi : RE \to RE$ be an operator. Recall that $\overline{2}^{\mathbb{N}}$ is the set of finite subsets of $\mathbb{N}$. In the following, we will use θ, θ_i, etc, to denote elements of $\overline{2}^{\mathbb{N}}$. We say Ψ is **continuous** if for all $A \in RE$ and $x \in \mathbb{N}$

$$x \in \Psi(A) \Leftrightarrow (\exists \theta \subseteq A)[x \in \Psi(\theta)]. \tag{30}$$

This definition uses our tacit convention that θ denotes a finite set.

We also say Φ is **monotone** if for all $A, B \in RE$, if $A \subseteq B$ then $\Psi(A) \subseteq \Psi(B)$.

THEOREM 40. *A recursive operator is both continuous and monotone.*

Proof. Let Ψ be a recursive operator based on h.
(a) Ψ is continuous. This follows from an application of the Rice-Shapiro theorem. Let $K = K_x := \{A \in RE : x \in \Psi(A)\}$. Since $RE = \{W_i : i \in \mathbb{N}\}$, and $\Psi(W_i) = W_{h(i)}$, we conclude that

$$L(K) = \{i \in \mathbb{N} : x \in W_{h(i)}\}.$$

It is easy to see that $L(K)$ is r.e.. Then by Rice-Shapiro, for all r.e. set A,

$$\begin{array}{rcl} x \in A & \iff & A \in K \\ & \iff & (\exists \theta \subseteq A)[\theta \in K] \\ & \iff & (\exists \theta \subseteq A)[x \in \Psi(\theta)], \end{array}$$

i.e., Ψ is continuous.
(b) Ψ is monotone. This follows easily from continuity: let $A \subseteq B$ be r.e. sets. Then

$$\begin{array}{rcl} x \in A & \iff & (\exists \theta \subseteq A)[x \in \Psi(\theta)], \\ & \Longrightarrow & (\exists \theta \subseteq B)[x \in \Psi(\theta)], \\ & \iff & x \in \Psi(B). \end{array}$$

Q.E.D.

We are ready to show the following result of Kleene.

THEOREM 41 (First Recursion Theorem). *If Ψ be any monotone operator. Then there is the least fixed-point A^* of Ψ, i.e.,*
(a) $\Psi(A^*) = A^*$.
(b) If $\Psi(B) = B$ for any r.e. B, then $A^* \subseteq B$.
(c) If Ψ is recursive, then A^* is r.e..

Proof. (a) We define the following sequence of r.e. sets:

$$A_0 = \emptyset, A_1 = \Psi(A_0), A_2 = \Psi(A_1), \ldots, A_{i+1} = \Psi(A_i), \ldots.$$

It is easy to see that $A_0 \subseteq A_1$, and hence by monotonicity, $A_i \Psi(A_{i-1}) \subseteq \Psi(A_i) = A_{i+1})$ for all $i \geq 1$. Let $A^* = \cup_{i \geq 0} A_i$. We verify that $\Psi(A^*) = A^*$ because $\Psi(\cup_{i \geq 0} A_i) = \cup_{i \geq 1} A_i = A^*$.

(b) A^* is the least fixed point because if $\Psi(B) = B$, then we have $A_0 \subseteq B$. But for any $i \geq 0$, if $A_i \subseteq B$ then $A_{i+1} = \Psi(A_i) \subseteq \Psi(B) = B$. Thus $A^* \subseteq B$.

(c) Now assume Ψ is recursive. So there is an extensional total recursive h such that $\Psi(W_i) = W_{h(i)}$. We want to verify that A^* is r.e.. If $e(0)$ is any Gödel number of A_0, then $e(i)$ is an Gödel number of A_i if we define $e(i) := h(e(i-1))$ for all $i \geq 1$. Then $x \in A^*$ iff $(\exists i \in \mathbb{N})[x \in W_{e(i)}]$. This last predicate is r.e.. **Q.E.D.**

This powerful theorem allows us to describe recursive sets A by using an equation of the form $A = \Psi(A)$ where Ψ is any recursive operator. It has important applications in the theory of programming semantics.

¶34. Effective continuity. Lemma 19(b) proves that there is a bijection from $\mathbb{N}$ to $\overline{2}^{\mathbb{N}}$. Let

$$i \mapsto \widetilde{i} \tag{31}$$

denote this bijection. Further, it is easy to show that this bijection is recursive (just elaborate on the proof of Lemma 19(b)). An operator $\Psi : RE \to RE$ is **effectively continuous** if there exists a r.e. relation $R \subseteq \mathbb{N}^2$ such that for all r.e. A and $x \in \mathbb{N}$,

$$x \in \Psi(A) \Leftrightarrow (\exists i \in \mathbb{N})[\widetilde{i} \subseteq A \wedge (i, x) \in R]. \tag{32}$$

It is easy to see that if Ψ is effectively continuous, then it is continuous: this is because for all $i, x \in \mathbb{N}$, we have

$$(i, x) \in R \Leftrightarrow x \in \Psi(\widetilde{i}).$$

THEOREM 42 (Myhill-Shepherdson). *An operator $\Psi : RE \to RE$ is recursive iff it is effectively continuous.*

Proof. ($\Rightarrow$) Suppose Ψ is recursive based on h. We only have to slightly modify the proof in previous theorem. We had shown that

$$x \in \Psi(A) \Leftrightarrow (\exists i \in \mathbb{N})[\widetilde{i} \subseteq A \wedge x \in \Psi(\widetilde{i})]. \tag{33}$$

The predicate $P(i) \equiv$"$\widetilde{i} \subseteq A$" is r.e. since A is r.e.. Let $t(i)$ be the Gödel number of the STM that accepts precisely the set $\widetilde{i}$. Clearly, $i \mapsto t(i)$ is total computable. The predicate "$x \in \Psi(\widetilde{i})$" is thus equivalent to "$x \in W_{h(t(i))}$" which is seen to be an r.e. predicate $R(i, x)$. Hence the predicate in the right-hand side of (33) is r.e., being of the form $(\exists \in \mathbb{N})[P(i) \wedge R(i, x)]$. Thus Ψ is effectively continuous based on R.

($\Leftarrow$) Conversely, suppose Ψ is effectively continuous based on the r.e. predicate $R(i, x)$. We need an extensional total recursive h. But $h(i)$ is simply the Gödel number of the STM that, on any input x, checks

$$(\exists i \in \mathbb{N})[\widetilde{i} \subseteq A \wedge R(i, x)].$$

This predicate is r.e., and $x \in W_{h(i)}$ iff $x \in \Psi(A)$. In particular, this function h is extensional because it depends only on A, and not on any Gödel number j such that $A = W_j$. **Q.E.D.**

EXERCISES

Exercise 0.3.11.57: Prove or disprove that the following are recursive operators.
(i) $\Psi(A) = \{a^2 : a \in A\}$
(ii) $\Psi(A) = \{a \in A : 2a \notin A\}$ □

Exercise 0.3.11.58: Re-develop this subsection, but for partial recursive functions instead of for r.e. sets. If particular, if $\mathrm{RE} := \{\phi_i : i \in \mathbb{N}\}$ is the set of partial recursive functions, an **operator** Ψ is a total function

$$\Psi : \mathrm{RE} \to \mathrm{RE}.$$

It is **recursive** if there is a total recursive h such that $\Psi(\phi_i) = \phi_{h(i)}$, where h is extensional in the sense that this definition does not depend on the choice of i.
(i) Prove the analogue of the First Recursion Theorem.
(ii) Prove the analogue of the Myhill-Shepherdson theorem.

□

END EXERCISES

0.3.12 Beyond Partial Computability

We know by counting arguments that there is a language that is not in $RE \cup \text{co-}RE$. But it is considerably more interesting to show that certain explicitly described sets have this property. Consider

$$\mathrm{TOT} := \{i \in \mathbb{N} : \phi_i \text{ is total}\}$$

Thus TOT is the set of Gödel numbers of total recursive functions. Since the UTM U (Theorem 25) defining the ϕ_i's in (25) is deterministic, if follows that $i \in \mathrm{TOT}$ is equivalent to U_i being a halting machine.

THEOREM 43. TOT *is neither r.e. nor co-r.e..*

Proof. (a) TOT is not r.e.: it is sufficient to show that co-HALT is recursively m-reducible to TOT. Let $i \in \mathbb{N}$. Consider the following machine M_i: on any input x, it runs $\phi_i(i)$ for x steps. If $\phi_i(i)$ does not halt within x steps, we accept. Otherwise, we loop.

Clearly, $\phi_i(i) \uparrow$ iff M_i is total. It is clear that there is a total function $t(i)$ that gives us the Gödel number for M_i. This proves that $i \in$ co-HALT iff $t(i) \in$ TOT.

(b) TOT is not co-r.e.: it is sufficient to show that HALT is recursively m-reducible to TOT. For any $i \in \mathbb{N}$, we construct the following machine N_i: on input x, N_i runs $\phi_i(i)$. If $\phi_i(i)$ halts, we halt. Otherwise, we loop. If $t'(i)$ is the Gödel number of N_i, then $i \in$ HALT iff $\phi(i) \downarrow$ iff $t'(i) \in$ TOT. Since $t'(i)$ is total computable, we are done. **Q.E.D.**

This theorem could also be obtained as a simple application of the Rice-Shapiro theorem (see a previous Exercise),

¶35. Turing reducibility. In order to classify problems such as TOT, we introduce a hierarchy of languages. First, we need the concept of an **Oracle Turing Machines** (OTM) or **query machine**. An OTM is physically like a UTM, but its extra tape is called a **oracle tape** instead of an index tape. This change in name signals a new use for this tape: while index tapes are read-only, the oracle tape is write-only. The way we ensure write-only is that the **oracle head** can only move to the right. There are also three special states, $query, YES, NO$ states. When the machine enters the query state, the non-blank word u that is currently scanned by the oracle head is submitted to some "oracle set" A. If $u \in A$, then the OTM is next placed in the YES state, else in the NO state. The computation continues from the YES or NO state. The oracle A is an arbitrary language that is independently specified. Let Q be an OTM and any A be any language. Let $L(Q^A)$ be the language accepted by Q using oracle A. The OTM Q is **deterministic** if its transition table is deterministic; Q is **halting** if for all A, Q^A halts on every input.

We say **T-reducible** as short for "Turing reducible". We say A is **partially T-reducible** to B, denoted $A \leq_T^{RE} B$, if A is accepted by some OTM Q with oracle B. If, in addition, Q is halting, we say A is **total Turing reducible** to B, and write $A \leq_T^{REC} B$.

¶36. Arithmetic Hierarchy. If K is any class, write Σ^K to be the class of all languages that are partially T-reducible to some language in K. Also, let $\Pi^K = \text{co-}\Sigma^K$. The **arithmetic hierarchy** is the following set of indexed classes:

$$\{\Sigma_n, \Pi_n, \Delta_n : n \in \mathbb{N}\}$$

defined as follows:
Basis: $\Sigma_0 := RE$, $\Pi_n := \text{co-}RE$, and $\Delta_0 := REC$.
Induction: $\Sigma_{n+1} = \Sigma^{\Sigma_n}$, $\Pi_{n+1} = \text{co-}\Sigma_{n+1}$, and $\Delta_{n+1} = \Sigma_{n+1} \cap \Pi_{n+1}$. The languages in $\Sigma_n \cup \Pi_n$ are said to be in the **n-th level** of the arithmetic hierarchy.

LEMMA 44. $\text{TOT} \in \Pi_2$.

Proof. Equivalently, we must show co-TOT $\in \Sigma_2$. This follows from showing an OTM Q such that co-TOT $\leq_T^{RE}$ HALT. Now, $i \in$ co-TOT iff there is some y such that $\langle i, i, y\rangle \in T$ where T is the recursive language in Theorem 30. Consider the machine M_i that on input x, will ignore x but check if any of the inputs

$$\langle i, i, 0\rangle, \langle i, i, 1\rangle, \ldots$$

are in T. It any of them is in T, M_i accepts, otherwise rejects. Clearly, M_i either halts or it always loops. If $t(i)$ is the Gödel number of M_i, we see that $i \in$ co-TOT iff $t(i) \in$ HALT. **Q.E.D.**

0.3.13 Cardinal and Ordinal Numbers

This optional subsection provides some additional information about the theory of cardinality.

How do we compare the "sizes" of sets X and Y? When X and Y are finite sets, we can list all the elements of X and Y in some "listing order" $(x_0, x_1, \ldots)$ and $(y_0, y_1, \ldots)$ and to match them up $x_0 \leftrightarrow y_0$, $x_1 \leftrightarrow y_1$, etc. The list that runs out of elements first has smaller size; if both list runs out of elements simultaneously, we say that they have the same "size". This "matching principle" is clearly valid in the finite case, as the outcome is independent of the listing order we choose. For infinite sets, this matching principle calls for refinement. For instance, the listing orders that show $|\mathbb{N}| = |\mathbb{Z}|$ (see (16)) must be carefully chosen or it may not work.

The concept of listing order generalizes to the concept of **ordinal numbers**. To get a matching principle that is independent of listing order, we need the concept of **cardinal numbers**. Informally, the relationship between these two concepts is seen in English where the "cardinal numbers" `zero`, `one`, `two`, etc., are used to count, while the "ordinal numbers" `zeroth`, `first`, `second`, etc., are used to order or rank. The theory of cardinals will be built on ordinal theory.

Define an **ordinal** to be a well-ordered set α such that for every $x \in \alpha$, the set $s(x) = \{y \in \alpha : y < x\}$ also belongs to α. Thus α is a set of sets. The set $s(x)$ is called an **initial segment** of α. Any two ordinal numbers α, β can be compared. This follows from general considerations about well-ordered sets. Two partially ordered sets $(A, \leq)$, $(A', \leq')$ are **similar** if exists a bijection $h : A \to A'$ that is order preserving (i.e., $a \leq b$ implies $h(a) \leq' h(b)$. If A, A' are well ordered, then we write $A \prec A'$ if A is similar to an initial segment of A'. This is clearly a transitive relation. For well ordered sets A and A', exactly one of the following holds:

$$A \sim A', \qquad A \prec A', \qquad A' \prec A.$$

See for example [3, p. 73]. This is the total ordering we use for ordinal numbers.

To view the natural numbers $n \in \mathbb{N}$ as ordinals, we must interpret them as sets using the following device: if $n = 0$, then n is the empty set; otherwise n is the set $\{0, 1, \ldots, n-1\}$. Thus

$$0 = \emptyset, \quad 1 = \{\emptyset\}, \quad 2 = \{\emptyset, \{\emptyset\}\}, \quad 3 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}, \quad \ldots.$$

Now we may verify that the set $\mathbb{N}$ is itself an ordinal number! This ordinal number is denoted ω. This is the first infinite ordinal. Given any ordinal α, we can define its **successor** α^+, which is the ordinal $\alpha \cup \{\alpha\}$, For instance, $\omega^+ = \mathbb{N} \cup \{\mathbb{N}\}$. Moreover, the operations of addition, multiplication and exponentiation can be defined. But be prepared for surprises: the ordinals $1 + \omega$ and $\omega + 1$ may be represented as

$$1 + \omega = \{0', 0, 1, 2, 3, \ldots\}, \qquad \omega + 1 = \{0, 1, 2, 3, \ldots, \omega\}.$$

This yields the unintuitive result, $1 + \omega = \omega < \omega + 1$. Similarly, $2\omega = \{0', 0, 1', 1, 2', 2, 3', \ldots\}$ and $\omega 2 = \{0', 1', 2', \ldots, 0, 1, 2', \ldots\}$, and we have $2\omega = \omega < \omega 2$.

Since distinct ordinal numbers can have the same cardinality (i.e., equivalent), we define a **cardinal number** to be smallest ordinal number among its equivalence class. Thus ω is a cardinal number and is in fact what we denoted by $\aleph_0$ before. But $\omega + 1$ and 2ω are not cardinal numbers as they are equivalent to ω. Cardinal comparison is induced from ordinal comparison. Again, cardinal arithmetic can be defined.

NOTES

¶37. Finite Automata. Finite automata were first introduced in the 1950s by G.H. Mealy and E.F. Moore. Most elementary books on the theory of computation have a treatment of finite automata and regular languages. Three references are Lewis and Papadimitriou [6] and Sipser [10]. The Rabin-Scott theorem appeared in [8]. and the Myhill-Nerode theorem is from [7]. An easy introduction to set theory is Halmos [3]. Although finite automata has been a relatively dormant subject since the 1960s, it is interesting to note that developments in program verification and program synthesis has breathed new life to this topic. Although many real systems can be modeled as finite state machines, such systems are so complex that it would be impossible to synthesize or understand them without semantically meaningful higher concepts. Instead of the "flat view" of an nfa, we would like to decompose a nfa into a more hierarchical structure based on concepts such as a product $Q = Q_1 \times Q_2$ or a disjoint union $Q = Q_1 \uplus Q_2$, etc. Various devices such as finite memory can be introduced in nfa descriptions (e.g., the last state visited in a group of states). See Harel's **state charts** [4] for one proposal.

¶38. Computability Theory. A basic reference for computability is Davis [2]. The standard work on recursive function theory is Rogers [9]; see Cutland [1] for an accessible elementary introduction. There are many approaches to computability. Mathematicians may prefer to work with number theoretic functions and recursion schemas. The language-theoretic approach and the use of automata is more common in Computer Science. The former approach is elegant because it deals with only one set, $\mathbb{N}$, while languages involve Σ^* for each alphabet Σ. But elegance comes with a cost as many natural sets (e.g., palindromes) are easier to describe as languages than as subsets of $\mathbb{N}$. One is forced to introduce artificial encodings in $\mathbb{N}$ such as pairing functions. With strings, pairing is trivially achieved by adding a single separator symbol. The "programming approach" is familiar in computer science. Certain results such as the S-m-n theorem of recursive function theory simply becomes programming facts in the automata approach (we exploit this in our approach).

The proofs that WORD, DIAG, HALT are not recursive is basically a diagonal argument. This is also the method by which one proves Gödel's famous incompleteness theorem: that there exists a true statement of number theory that cannot be proven. What is needed, and this is Gödel's remarkable achievement, is a way to encode provability and statements of number theory within number theory itself. Provability is akin to computability. Diagonal arguments appear at the foundation of set theory as paradoxes (contradictions). These paradoxes typically has some self-referential element. For instance: `In a certain village, there are many barbers. But there is one barber who shaves any barber who does not shave himself. Who shaves this barber?` Pondering the two possibilities, one is led to the conclusion that this barber does not exist. In the set theoretic paradoxes, we can similarly define sets that do not exist. This led to efforts to axiomatize set theory, basically to carefully prescribed circumstances under which new sets can be legitimately formed. Happily, all the operations we make in these notes are admissible (power set is admissible for instance).

Exercise 0.3.13.59: Based on the examples in the notes on cardinal and ordinal numbers, give a definition of the operations of + and ×. Deduce various properties of your definition. Confirm the properties noted in the examples ($\omega = 2\omega$, $\omega < \omega 2$, etc). □

End Exercises

§A. Basic Mathematical Vocabulary

This is a handy collection of common notations and mathematical concepts that are used throughout the book.

Equality. The equality symbol "=" is overworked in many ways; the following two conventions will offload some burden. (1) We use ':=' to indicate a definitional equality. For example, we indicate the definition of a set S by writing '$S := \ldots$' where '$\ldots$' describes a set formation (see Sets below). (2) In programs, we use ← as the programming assignment symbol: $x \leftarrow 2$ indicates as assignment of the value 2 to a programming variable x. However, we continue to use "=" as the equality predicate (and not "==" as in C-like languages).

Numbers. We consider the following number hierarchy:

$$\mathbb{N} \subseteq \mathbb{Z} \subseteq \mathbb{Q} \subseteq \mathbb{R} \subseteq \mathbb{C}$$

where $\mathbb{N}$ is the set of natural numbers $0, 1, 2, \ldots$, $\mathbb{Z}$ is the set of integers $0, \pm 1, \pm 2, \ldots$, $\mathbb{Q}$ is the set of rational numbers, $\mathbb{R}$ is the set of real numbers, and $\mathbb{C} = \mathbb{R} + \mathbf{i}\mathbb{R}$ is the set of complex numbers. Here, $\mathbf{i} = \sqrt{-1}$. The absolute value of complex numbers $x = a + \mathbf{i}b \in \mathbb{C}$. is the real number $|x| := \sqrt{a^2 + b^2}$. The set of *extended reals* refers to $\mathcal{R} \cup \{\infty\}$. For any $n \geq 1$, we also define $\mathbb{Z}_n := \{0, 1, \ldots, n-1\}$. This is a finite ring under modulo n arithmetic. In case n is prime, then $\mathbb{Z}_n$ becomes a field. We sometimes write $\mathbb{B}$ for $\mathbb{Z}_2$, which is the **Boolean field** with 2 elements.

Sets. We write $X \subseteq X'$ for set inclusion, and $S \subset S'$ for proper set inclusion. The **cardinality** (or **size**) of S is denoted $|S|$. Note that the $|\cdots|$ notation in several other ways (length of a word, absolute value of complex numbers, etc) but the context should make clear which function is intended. If X is a well-defined set, we can introduce a new set by writing $\{x \in X : P(x)\}$ where $P(x)$ is a predicate on X. For instance, $\{x : x \in \mathbb{N}, x \text{ is even}\}$ denotes the set of even natural numbers. The **size** $|X|$ of X is the number of its elements. The **empty set** is denoted $\emptyset$ and thus $|\emptyset| = 0$. If $X = \bigcup_{i \in I} X_i$ and the X_i's are pairwise disjoint, we indicate this fact by writing $X = \biguplus i \in I X_i$, and call $\{X_i : i \in I\}$ a **partition** of X. The set of all subsets of a set X is the **power set** of X, denoted 2^X. We have $\left|2^X\right| = 2^{|X|}$. A subset of X of size k is called a **k-set**. The set of all k-sets of X is denoted $\binom{X}{k}$. This recalls the binomial function $\binom{n}{k} = \frac{n!}{k!(n-k)!}$. In fact, we have $\left|\binom{X}{k}\right| = \binom{|X|}{k}$. The Cartesian product of X and Y is $X \times Y = \{(x, y) : x \in X, y \in Y\}$. This can be generalized to an **n-fold Cartesian product**, $X_1 \times X_2 \times \cdots X_n$ or $\prod_{i=1}^n X_i$ (this notation exploits the associativity of Cartesian product). In case $X_1 = X_2 = \cdots = X_n$, we write X^n for $\prod_{i=1}^n X$. A member of an n-fold Cartesian product is sometimes written $\langle a_1, \ldots, a_n \rangle$ instead of $(a_1, \ldots, a_n)$.

Relations. An **n-ary relation** R is a subset of a Cartesian product $\prod_{i=1}^n X_i$ where S_i are sets. We say the relationship $R(a_1, \ldots, a_n)$ **holds** in case $(a_1, \ldots, a_n) \in R$, else it **fails**. If $n = 1$, R is also called a **predicate** on S_1. For a binary relation R, we may use an infix notation, and write write $a_1 R a_2$ instead of $R(a_1, a_2)$. A binary relation $R \subseteq S^2$ is **reflexive** if aRa for all $a \in S$; **symmetric** if aRb imples bRa; and **transitive** if aRb and bRc implies aRc. A **partial order** on S is a binary relation $R \subseteq S^2$ that is reflexive, symmetric and transitive. We usually write $a \leq b$ for aRb in case R is a partial order. If, in addition, we have either $a \leq b$ or $b \leq a$ for all $a, b \in S$, then R is called a **total order**. A total order A **well order** on S is a partial ordering $\leq$ such that every subset $A \subseteq S$ has a least element $a \in A$ (i.e., for all $b \in A$, $a \leq b$). Note that a well order is a total order. The converse is false: $\mathbb{Z}$ is a total order but not a well order since $\{-1, -2, -3, \ldots\} \subseteq \mathbb{Z}$ has no least element.

Functions. Let D, R be sets, and $\uparrow$ a special symbol not in these set. By a *partial function* f with *nominal domain* D and *nominal range* R we mean a rule that associates each element x of D with an element $f(x) \in R \cup \{\uparrow\}$. We write $f : D \to R$ to indicate this relationship. So D and R are assumed to be part of the specification of f. If $f(x) = \uparrow$ we say that f is **undefined** at x; otherwise, we say f is **defined** at x and write $f(x) = \downarrow$. REMARK: In some books, the notation "$f(x)\uparrow$" is used in place of "$f(x) = \uparrow$"; but in this book, the "$f(x)\uparrow$" notation is used to indicate that a Turing machine f loops on input x. Similarly "$f(x)\downarrow$" indicates that the Turing machine halts on input x.

The *proper domain* and *proper range* of f are defined (respectively) as $\texttt{domain}(f) := \{x \in D : f(x)\downarrow\}$ and $\texttt{range}(f) := \{f(x) : x \in \texttt{domain}(f)\}$. If $\texttt{domain}(f)$ is the nominal domain, we say f is *total*. Except for number theoretic functions and complexity functions, functions are usually assumed to be total.

Composition of partial functions is denoted $f \circ g$ where $(f \circ g)(x) := f(g(x))$ and $f(\uparrow) = \uparrow$. A total function f is an **injection** (or **one-one**) if $x \neq y$ implies $f(x) \neq f(y)$; it is a **surjection** (or **onto**) if $\{f(x) : x \in D\} = R$; it is a **bijection** (or **one-one onto**) if it is both an injection and a bijection.

The *floor function* takes any real number x to the largest integer $\lfloor x \rfloor$ no larger than x. The *ceiling function* takes any real number x to the smallest integer $\lceil x \rceil$ no smaller than x. So $\lfloor 0.5 \rfloor = 0 = \lceil -0.9 \rceil$ and $\lfloor x \rfloor \leq x \leq \lceil x \rceil$. The *logarithm function* $\log_b a$, for our purposes, is defined for real numbers $b > 1$ and $a > 0$. The **base** b is considered fixed. We use the shorthand $\lg x := \log_2 x$ and $\ln x := \log_e x$. for two common bases, $b = 2$ and $b = e = 2.718\ldots$.

Graphs. There are many varieties of graphs, but two varieties are important for us. A **directed graph** or **digraph** is a pair $G = (V, E)$ where V is any set and $E \subseteq V^2$. A **undirected graph** or **bigraph** ("bi" for bidirectional) is a pair $G = (V, E)$ where V is any set and $E \subseteq \binom{V}{2}$. For both kinds of graphs, the elements of V and E are called **vertices** and **edges**, respectively. Definitions apply to both kinds of graphs when we fail to distinguish between them. We denote an edge by (u, v) where $(u, v) \in E$ (if digraph) or $\{u, v\} \in E$ (if bigraph). A **path** is a sequence $p = (v_1, v_2, \ldots, v_k)$ where (v_i, v_{i+1}) are edges for $i = 1, \ldots, k-1$. We call p a path from v_1 to v_k. If $v_1 = v_k$, then p is called a **cycle**. A graph with no cycles is **acyclic**.

A **tree** $T = (V, E)$ is an acyclic digraph with a distinguished node called the **root** and such that there is a path from the root to every other node. If $(u, v) \in E$, then v is a **child** of u. The **leaves** of a tree are those nodes with no children; non-leaves are also called **internal nodes**. An **orient tree** is a tree in which the children of each node are totally ordered (so one can speak of the first child, second child, etc).

Asymptotics. Let $\phi, \phi' \in [\mathbb{N} \to \mathbb{N}]$. We say $\phi \leq \phi'$ **eventually**, and write

$$\phi(x) \leq \phi'(x) \quad (\text{ev. } x),$$

if there is some x_0 such that if $x > x_0$ then

$$\phi(x) \leq \phi'(x). \tag{34}$$

In any assertion (such as (34)) involving partial functions, it is assumed that the assertion is conditioned on all the partial functions being defined at their stated arguments. We say $\phi \leq \phi'$ **infinitely often**, and write

$$\phi(x) \leq \phi'(x) \quad (\text{i.o. } x)$$

if $\phi(x) \leq \phi'(x)$ holds at infinitely many values of x. For instance, if $|\texttt{domain}(\phi) \cap \texttt{domain}(\phi')| < \infty$ then this relation automatically fails. We say ϕ is **non-decreasing** if $x < y$ and $\phi(x), \phi(y)\downarrow$ then $\phi(x) \leq \phi(y)$.

Formal Language Theory. An *alphabet* Σ is a non-empty finite set of symbols. Each element of Σ is a *letter*. A **word** (or **string**) w over Σ is a finite sequence of letters of Σ,

$$w = a_1 \cdots a_n, \quad (n \geq 0)$$

where $a_i \in \Sigma$. The **length** of w is n, denoted $|w|$. The **empty word** is denoted ϵ; thus $|\epsilon| = 0$. The set of all words over Σ is denoted Σ^*. Also, $\Sigma^+ := \Sigma^* \setminus \{\epsilon\}$ is the set of non-empty words over Σ. If $a \in \Sigma$, then define the counting function $\#_a : \Sigma^* \to \mathbb{N}$ where $\#_a(w)$ is the number of occurences of a in w. The ith symbol in a word w is denoted $w[i]$ where $i = 1, \ldots, |w|$. Assume $w[i]$ is undefined for i outside this range. The reverse of a word w is denoted w^R. Thus $w[i] = w^R[n - i + 1]$ where $n = |w|$. A palindrome is a word w such that $w = w^R$. Let Σ^* denote the set of words over Σ. A **language** is a pair of the form (Σ, A) where $A \subseteq \Sigma^*$. We usually refer to "A" as the language (so the alphabet Σ is implicit). The **complement** of a language (Σ, A) is $(\Sigma, \Sigma^* \setminus A)$,

denoted co-A. The concatenation of two words v and w is written $v \cdot w$ or simply vw. This notation extends to languages: $A \cdot B := \{vw : v \in A, w \in B\}$. For any non-negative integer n and word w, we let w^n denote the n-fold self-concatenation of w. More precisely, $w^0 := \epsilon$ and for $n \geq 1$, $w^n := w \cdot w^{n-1}$. Functions of the form $t : \Sigma^* \to \Gamma^*$ are called **word functions** or **transformations**, depending on the context.

A set K of languages is usually called a **class**. For any class K, co-K is defined to be the class $\{co\text{-}L : L \in K\}$. If Σ is any alphabet, let $K|\Sigma$ be the subclass of K comprising those languages in K with alphabet Σ.

Sometimes, we go beyond finite strings. An infinite strings of the form $w = a_0a_1a_2\cdots$ where $a_i \in \Sigma$, is called an **ω-word** (or ω-sequence) over Σ. Alternatively, such a string w is a function $w : \mathbb{N} \to \Sigma$ so $w(n) = a_n$. Let Σ^ω denote the set of ω-words over Σ.

Logic. Boolean values are usually identified with $\mathbb{B} := \{0, 1\}$. If $a, b \in \mathbb{B}$, the logical operators of **and**, **or** and **not** are denoted $a \wedge b$, $a \vee b$ and $\not a$. **Quantifiers** are of two types: **universal** or **existential**, denoted $\forall$ and $\exists$, respectively.

Classes and Families. Above, we only allowed new sets to be constructed from known sets. Otherwise, there is danger in using the concept of "sets" too liberally, giving rise to various logical paradoxes. In general, paradoxes arise for "very large sets" such as "the set of all sets". Logicians are careful to introduce other terms like "classes" or "families" for such large sets to prevent non-admissible set formations. We use "classes" and "families" informally for some large sets in our book, but we make no true distinction between them except for usage convention. For instance, consider $\mathcal{L}$, the "set of all languages". This is "large" because each language depends on an underlying alphabet, and there is no fixed[14] set of alphabets that we wish to commit ourselves too. Similarly, the "set of all Turing machines" $\mathcal{M}$ is not really a set since we have not committed ourselves to a definite set of possible symbols or states. Hence, $\mathcal{C}$ and $\mathcal{M}$ are not mathematical sets. They have many features of sets: e.g., we can recognize individual members of $\mathcal{L}$. Yet we never get into trouble because each particular use can be justified or suitably circumscribed. Moreover, certain set operations such as power set ought not to be performed. Usage conventions: we refer to "large sets" as **classes** when we speak of languages, and as **families** when we speak of machines. For example, the "class of polynomial-time languages" and the "family of deterministic Turing machines".

OMIT BELOW THIS A *language operator* ω is a partial d-ary function, $d \geq 0$, taking a d-tuple $(L_1, \ldots, L_d)$ of languages to a language $\omega(L_1, \ldots, L_d)$. Here we assume that the L_i have a common alphabet which is also the alphabet of $\omega(L_1, \ldots, L_d)$. Some simple language operators are now introduced. Other important operators, to be introduced in the course of this book, are usually defined using machines; this is in contrast with the following set-theoretic definitions.

Let $(\Sigma, L), (\Sigma, L')$ be languages. The *complement* of L, denoted co-L, is the language $\Sigma^* - L$. The *union*, *intersection* and *difference* of (Σ, L) and (Σ', L') are (resp.) the languages $(\Sigma \cup \Sigma', L \cup L')$, $(\Sigma \cup \Sigma', L \cap L')$ and $(\Sigma, L - L')$. (The preceding are the *Boolean operators*.) The *concatenation* of L and L', denoted $L \cdot L'$, is the language $\{ww' : w \in L, w' \in L'\}$ over the alphabet $\Sigma \cup \Sigma'$. For any non-negative integer n, we define the language L^n inductively as follows: L^0 consists of just the empty word. $L^{n+1} := L^n \cdot L$. The *Kleene-star* of L, denoted L^*, is the language $\bigcup_{n\geq 0} L^n$. (Note that the Σ^* and L^* notations are compatible.) A related notation is L^+ defined to be $L \cdot L^*$. The *reverse* of L, denoted L^R, is $\{w^R : w \in L\}$ where w^R denotes the reverse of w.

A language L is said to be *finite* (resp. *co-finite*) if $|L|$ (resp. $|\text{co-}L|$) is finite.

Let Σ and Γ be alphabets. A *substitution* (from Σ to Γ) is a function h that assigns to each $x \in \Sigma$ a subset of Γ^*. h is naturally extended to a function (still denoted by) h that assigns to each word in Σ^* a set of words in Γ^*: $h(a_1a_2\cdots a_n) := h(a_1)h(a_2)\cdots h(a_n)$. We say h is *non-erasing* if $\epsilon \notin h(x)$ for all $x \in \Sigma$. A *homomorphism* is a substitution h where each set $h(x)$ has exactly one element (we may thus regard $h(x)$ as an element of Γ^*). A *letter homomorphism* is a homomorphism where $h(x)$ is a word of length 1 for all $x \in \Sigma$ (we may thus regard $h(x)$ as an element of Γ). An *isomorphism* is a letter homomorphism such that h is a bijection from Σ to Γ . An isomorphism is therefore only a 'renaming' of the alphabet.

For every substitution h from Σ to Γ, we may define the language operator (again denoted by h) that takes a language (Σ, L) to the language $(\Gamma, h(L))$ where $h(L)$ is the union of the sets $h(w)$ over all $w \in L$. We also define the *inverse substitution* operator h^{-1} that takes a language (Γ, L') to $(\Sigma, h^{-1}(L'))$ where $h^{-1}(L')$ is the set $\{w \in \Sigma^* : h(w) \subseteq L'\}$.

[14] Actually, we establish two conventions (α and β) in which these "large sets" can be completely rigorous.

A *(language) class* K is a collection of languages that is closed under isomorphism. We emphasized in Chapter 1 that complexity theory is primarily the study of language classes, not of individual languages. The classes which interest us are usually defined using machines that use a limited amount of computing resources. In this case, we call K a *complexity class* although this is only an informal distinction.

Operators are important tools for analyzing the structure of complexity classes. For instance, many important questions are of the form "Are two complexity classes K and K' equal?". If Ω is any set of operators, let

$$\Omega(K) := \{\omega(L_1, \ldots, L_d) : L_i \in K, \omega \text{ is a } d\text{-ary operator in } \Omega\}.$$

The *closure* of K under Ω is

$$\Omega^*(K) := \bigcup_{n \geq 0} \Omega^n(K)$$

where $\Omega^0(K) := K$ and $\Omega^{n+1}(K) := \Omega(\Omega^n(K))$. One possible approach to showing that K is not equal to K' is to show that, for a suitable class of operators Ω, K is closed under Ω (i.e., $\Omega^*(K) = K$) but K' is not. An important simple case of Ω is where Ω consists of just the Boolean complement operator; here $\Omega(K) = \{\text{co-}L : L \in K\}$ is simply written as co-K. A branch of formal language theory called AFL theory investigates closure questions of this sort and certain complexity theory questions can be resolved with this approach.

Exercises

[0.1] Construct dfa's to accept the following languages:
(i) $\{w \in \{0,1\}^* : w \text{ has 0101 as substring}\}$.
(ii) $\{w \in \{0,1\}^* : w \text{ has neither 00 nor 11 as substring}\}$.

[0.2] If A is regular, show that its reverse $A^R = \{w^R : w \in A\}$ is regular.

Write regular expressions for (i) and (ii) in the previous question.

[0.3] Construct nfa's to accept the languages indicated by the following expressions:
(i) $(01)^*(10)^* + 00^*$.
(ii) $((01 + 001)^*0^*)^*$.
Be sure to parse these regular expressions correctly when we paretheses are omitted.

[0.4] Prove or disprove that each of the following languages is regular:
(i) $\{1^n : n \text{ is divisible by 7}\}$.
(ii) $\{w \in \{0,1\}^* : w \text{ is a binary number divisible by 7}\}$.
(iii) $\{0^n 10^m 10^{m+n} : n, m \geq 1\}$.
(iv) $\{ww' \in \{0,1\}^* : w' \text{ is obtained from } w \text{ by interchanging 0 and 1}\}$.

[0.5] A **letter homomorphism** is a function $h : \Sigma \to \Gamma^*$ where Σ, Γ are alphabets. We can extend h to the function

$$h : \Sigma^* \to \Gamma^*$$

in the natural way, so that $h(vw) = h(v)h(w)$ for all $v, w \in \Sigma^*$. In particular, $h(\epsilon) = \epsilon$. If $A \subseteq \Sigma^*$, then define $h[A] = \{h(w) : w \in A\} \subseteq \Gamma^*$. Show the following.
(i) If $A \subseteq \Sigma^*$ is regular, so is $h[A]$.
(ii) If $B \subseteq \Gamma^*$ is regular, so is that $h^{-1}[B] = \{w \in \Sigma^* : h(w) \in B\}$. HINT: start from a dfa M for B and construct one that, when it reads an input symbol a, tries to simulate M on $h(a)$. HINT: This exercise shows the differences in regular expressions and nfas as computing devices. Part (i) is easy to solve using regular expressions, but dfa's are better for part (ii).

[0.6] (Ó'Dúnlaing) For $k \geq 1$, let $(\{0,1\}^*, L_k)$ be the language consisting of binary strings of period k: say a word w has period k if $|w| \geq k$ and for all $i = 1, 2, \ldots, |w| - k - 1$, $w[i] = w[i + k + 1]$.
(i) Show that the complement of L_k can be accepted by a 'small' nondeterministic finite automaton, say with $2k + 2$ states.
(ii) Prove that L_k cannot be accepted by any nondeterministic finite automaton with less than 2^k states.

Bibliography

[1] N. Cutland. *Computability: an introduction to recursive function theory*. Cambridge University Press, Cambridge, 1980.

[2] M. Davis. *Computability and Unsolvability*. McGraw-Hill, New York, 1958. Reissued by Dover Publications, Inc., 1982.

[3] P. R. Halmos. *Naive Set Theory*. Van Nostrand Reinhold Company, New York, 1960.

[4] D. Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, 8:231–274, 1987.

[5] J. E. Hopcroft and J. D. Ullman. *Formal Languages and Their Relation to Automata*. Addison-Wesley Publishing Co., Reading, Massachusetts, 1969.

[6] H. R. Lewis and C. H. Papadimitriou. *Elements of the Theory of Computation*. Prentice-Hall Inc, Upper Saddle River, New Jersey, second ed. edition, 1998.

[7] A. Nerode. Linear automaton transformations. *Proc. AMS*, 9:541–544, 1958.

[8] M. O. Rabin and D. Scott. Finite automata and their decision problems. *IBM J. of Research and Development*, 3:114–125, 1959.

[9] H. Rogers. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, New York, 1967.

[10] M. Sipser. *Introduction to the Theory of Computation*. PWS Publishing Co, Boston, 1997.

[11] A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc. Series 2*, 42:230–265, 1936. Corrections: Vol.43 (1937) pp.544-546.

Chapter 2

Initiation to Complexity Theory

April 13, 2009

This book presumes no background in complexity theory. The formal introduction of complexity theory will begin in the next chapter. In the present chapter, we address some preparatory considerations. Here, we often compare and contrast the complexity quest against the older theory of computability. The rudiments of computability theory are provided in Chapter 0, but even without such a background, a reader should be able to understand the essential thrusts of this informal chapter. The rest of the book does not depend on this chapter except for the asymptotic notations of section 3.

2.1 Central Questions

Many fields of study are centered around some basic phenomenon, the understanding of which is either intrinsically interesting or could lead to practical benefits. For us, the phenomenon is the intuitive notion of *complexity of computational problems* as it arises in mathematical and computational sciences. The understanding of what makes a problem (computationally) complex is one cornerstone of the art and science of algorithm design. The stress in 'complexity of computational problems' is on 'complexity'; the concept of 'computational problem', is generally relegated to the background.[1]

To set the frame of mind, we examine some rather natural questions. A main motivation of our subject is to provide satisfactory answers to questions such as:

(1) Is multiplication harder than addition?

The appeal of this question, first asked by Cobham [5], is that it relates to what are probably the two most widely known non-trivial algorithms in the world: the so-called *high school algorithms* for addition and multiplication. To add (resp. multiply) two n-digit numbers using the high school algorithm takes linear (resp. quadratic) time. More precisely, the addition (resp. multiplication) algorithm takes at most $c_1 n$ (resp. $c_2 n^2$) steps, for some positive constants c_1 and c_2. Here a 'step' is a basic arithmetic operation $(+, -, \times, \div)$ on single digit numbers. So a simple but unsatisfactory answer to (1) is 'yes' because $c_2 n^2$ dominates $c_1 n$ when n gets large enough. It is unsatisfactory because the answer only says something about *particular* methods of adding and multiplying. To provide a more satisfactory answer, we probe deeper.

It is intuitively clear that one cannot hope to do additions in less than n steps. This is because any algorithm must at least read all the n input digits. So the high school method for addition is optimal – note that here and throughout the book, *optimality is taken up to some multiplicative constant factor.* The situation is less clear with multiplication: Is there any algorithm for multiplication that is asymptotically faster than the high school method? The answer turns out to be 'yes'; in 1971 (culminating a series of developments) Schönhage and Strassen [30] discovered what is today asymptotically the fastest known multiplication algorithm. Their algorithm takes

[1] There is no general theory of computational problems except in special cases. Such a theory should study the logical structure of problems, their taxonomy and inter-relationships. Instead, complexity theory obliterates all natural structures in computational problems by certain sweeping assumptions we will come to.

$c_3 n \log n \log\log n$ steps[2] on a *multitape Turing machine*. This model, to be introduced in Chapter 2, is a simple extension of the Simple Turing machine of the previous chapter. Since Turing machines are more primitive than any real-world computer, this means the same time bound is achievable on actual computers. For now, let us just accept the fact Turing machines are fundamental and thus statements about computations by Turing machine are intrinsically important. But is the Schönhage-Strassen algorithm the best possible? More precisely,

(2) Must every Turing machine that multiplies, in worst-case, use at least $cn \log n \log\log n$ steps, for some $c > 0$ and for infinitely many values of n? Here c may depend on the Turing machine.

This is an important open problem in Complexity Theory. A negative answer to question (2) typically means that we explicitly show an algorithm that is faster than Schönhage-Strassen. For instance, an algorithm with running time of $n \log n$ will do. We say such an algorithm shows an *upper bound* of $n \log n$ on the complexity of multiplication. (It is also conceivable that the negative answer comes from showing the existence of a faster algorithm, but no algorithm is exhibited in the proof[3].) On the other hand, answering question (2) in the affirmative means showing a *lower bound* of $cn \log n \log\log n$ on the complexity of multiplication. Such a result would evidently be very deep: it says something about *all* possible Turing machines that multiply! Combining such a result with the result of Schönhage-Strassen, we would then say that the *intrinsic complexity* of multiplication is $n \log n \log\log n$. In general, when the upper and lower bounds on the complexity of any problem P meet (up to a constant multiplicative factor) we have obtained a bound which is intrinsic[4] to P. We may now (satisfactorily) interpret question (1) as asking whether the intrinsic complexity of multiplication is greater than the intrinsic complexity of addition. Since the intrinsic complexity of addition is easily seen to be linear, Cobham's question amounts to asking whether multiplication is intrinsically non-linear in complexity. (Most practitioners in the field believe it is.) Generally, for any problem P, we may ask

(3) What is the intrinsic complexity of P?

It turns out that there is another very natural model of computers called *Storage Modification Machines* (see §5) which can multiply in linear time. This shows that complexity is relative to a given model of computation. From Chapter 0, we learned that all reasonable general models of computation are equivalent in the sense that they define the class of total and partial computable functions: this is known as *Church's Thesis*. But simplistic attempts to formulate analogous statements in Complexity Theory would fail. For instance, there are problems which can be solved in linear time in one model but provably take cn^2 (for some $c > 0$, for infinitely many n) in another model. So a fundamental question is:

(4) Which model of computation is appropriate for Complexity Theory?

Perhaps there is no uniquely appropriate model of computation for Complexity Theory. This suggests that we examine various models and discover their inter-relations. We shall see in section 7 that we can recover some version of Church's Thesis in Complexity Theory if we take some care in selecting models and classifying them according to their 'computational mode'.

This question about intrinsic complexity of a problem P begs another question:

(5) How should the input and output values of problems be encoded?

For instance, we have implicitly assumed that the multiplication problem has its inputs and outputs as binary numbers. But other representations of numbers are conceivable and this choice can be shown to affect the answer to (3). In section 4, we discuss this in more detail (it turns out that without a proper theory of problems, few precise statements can be made about choice of encodings and we are stuck with assumptions such as in the case of multiplication: each problem P comes equipped with a definite encoding).

The discussion of the complexity of multiplication so far centers around speed of computation or *time*. Time is an instance[5] of a *computational resource*. We measure complexity in terms of the amount of computational resources used up during a computation. Another natural resource is memory size or *space*. (The high-school algorithm is seen to take quadratic space with usual hand-calculation but it can easily be reduced to linear space.) We can imagine composite resources that combine time and space, for instance.

[2]Unless otherwise indicated, the reader may always take logarithms to the base 2. We shall see that the choice of the base is inconsequential.

[3]Such proofs are known but are rare. For instance, the work of Robertson and Seymour on graph minors leads to such conclusions (see for example [18]). Indeed the situation could be more complicated because there are degrees of explicitness.

[4]See section 8 for more discussion of this.

[5]No pun intended.

(6) What are other natural computational resources and their inter-relations?

It is seen that time and space are related: assuming each step can access one new unit of space, the amount of space used by an algorithm is no more than the amount of time consumed. However, there is a fundamental difference between time and space: space, but not time, is reusable. There can be trade-offs between time and space in the sense that an algorithm can run in less time only at the (unavoidable) cost of using more space. Up to a degree, resources depends on the computational model.

The high-school multiplication algorithm exhibits an interesting feature not generally true of algorithms: for all inputs of a given size, the algorithm executes the same number of steps. This is because it has no 'decision' steps (e.g. 'if $x = 0$ then $\cdots$ else $\cdots$').[6] The various known sorting algorithms are examples of decision-making algorithms. For such algorithms the running time may vary for different inputs of the same size. Then one has the choice of measuring the *worst-case* or the *average-case* time. In other words, the method of *aggregating* the total resources used by the algorithm is an issue to be decided in measuring complexity.

Another issue arises when we consider probabilistic computations where the algorithm makes random decisions by 'tossing a coin'. For a fixed input, different runs of the algorithm may produce different outputs. We are then faced with defining a suitable notion of the 'correct output': this is not formidable once we realize that the output can be a random variable.

We summarize the above discussion by saying that the *complexity measure* is a function of (a) the model of computation, (b) the computational resources of interest, (c) the method of aggregating the resource usage and (d) the definition of the algorithm's output (not to be confused with encoding the output of a computation problem). Thus question (6) translates into a question about complexity measures.

The theory in this book will address these and other related questions. In the rest of this chapter, we lay the foundations by discussing some issues of methodology. The assumptions of methodology are embodied in conventions, definitions and notations. Whenever possible, we will isolate them with a letter label ((A), (B), etc.).

2.2 What is a Computational Problem?

The concept of a *computational problem* (or simply, *problem*) in informal language can take many meanings and shades. Even restricting the concept to its use in Computer Science is too wide. The theory we propose to develop takes the prototypical and simplest aspects only. This is a natural choice to start with, and it is already enough to give us a rich and meaningful theory.

There are three main aspects of a problem in its natural context: (i) The problem statement, P. (ii) The methods available for its solution, M. (iii) The criteria for evaluating and comparing different solutions, C. The triple $\langle P, M, C\rangle$ may be called a *problem-in-context*. We will identify M with the *model of computation* and C with the *complexity measure*. A *solution* to $\langle P, M, C\rangle$ is simply an algorithm in M that satisfies the specifications of P and that is optimal or correct in the sense of C. It is not very meaningful to discuss P without this context. In practice, it is not always easy to separate a problem-in-context into the three components. Roughly speaking, P specifies the input-output behavior (i.e., the relationship between any given input x with the output y of any purported algorithm for this problem). Generally, M is dependent on P, and likewise C depends on both M and P. If M' (respectively, C') is different from M (respectively, C) then the two problems-in-context $\langle P, M, C\rangle$ and $\langle P, M', C'\rangle$ should really be considered different problems, even though they share a common P.

Let us illustrate the preceding discussion using the well-known example of *sorting*. Let P be the sorting problem where, given a set X of integers, the problem is to 'list the elements of X in a non-descending order'. Three very different models of computation M_i $(i = 1, 2, 3)$ have been extensively studied in the literature.

M_1: The *comparison tree model* where algorithms are finite binary trees whose nodes correspond to comparisons.

M_2: The *tape merging model* where algorithms wind and rewind tapes based on outcomes of comparing the data at the front-ends of the tapes.

M_3: The *comparison network model* where algorithms are directed acyclic graphs whose nodes (called comparators) have in-degrees and out-degrees of 2.

The complexity measure studied is dependent upon the model: In M_1, the goal is to minimize the height of the binary trees. Depending on whether we consider 'worst-case height' or 'average height', we get different measures.

[6] One can argue that the multiplication algorithm has decisions in the form of checking for carry bits. However, such decisions are unessential for integer multiplication since we may assume that carry bits (possibly 0) are present at each step. Nevertheless we can modify the algorithm to incorporate decision steps that are not so easily dismissed, such as to check if one of its arguments is zero.

In M_2, the criterion is to minimize tape motion (rather than the number of comparisons, as in M_1). In M_3, the criterion is to minimize the number of comparators or to minimize the length of the longest path.

Intuitively, P is the most fundamental of the three parameters $\langle P, M, C\rangle$, and we shall try to restrict the use of the term 'problem' to mean P alone. This is not entirely possible but we shall largely get by with the (mild) myth that our problems P do not imply any particular M or C. So the task that confronts us after the preceding clarifications, is to formalize P , this notion of 'problem-in-the-restricted-sense'. Let us begin by listing some typical problems in Computer Science:

(i) (Rubik's puzzle) For any given configuration of the Rubik's Cube, find the least number of moves required to reach a configuration which is monochromatic on each face. This is a *finite problem* since there are a finite number of possible input configurations.

(ii) (Planarity Testing) For each graph G, the problem is to decide if G is planar (i.e., embeddable in the plane). This is an example of a *decision problem* where the algorithm has to decide between a 'yes' and a 'no' answer. Alternatively, it is called a *recognition problem* where the algorithm has to "recognize" inputs that represent planar graphs.

(iii) (Fibonacci numbers) Let $f(n)$ be the nth Fibonacci number. The sequence of Fibonacci numbers is $0, 1, 1, 2, 3, 5, 8, \ldots$. For each input n, the problem is to compute $f(n)$. This is an instance of computing a number theoretic function.

(iv) (Linear Programming) For each m by n matrix A and an n-vector c, find any n-vector x such that $Ax \geq 0$ and $c \cdot x$ is maximized. If there are no x satisfying $Ax \geq 0$, or if $c \cdot x$ is unbounded, indicate so. (We assume that m and n can vary and are part of the input.) This exemplifies the class of *optimization problems* which typically arise in the field of Operations Research.

(v) (Element identification) For each input set X of integers, $|X| < \infty$, construct a data-structure $D(X)$ such that queries Q of the form: 'Is x in X?' can be rapidly answered using $D(X)$. This is an instance of a *preprocessing problem* where the inputs are given in two stages and separate algorithms are involved in each stage. In the first stage we have an algorithm A to 'preprocess' the input X. In the second stage, we have an algorithm B which uses the preprocessed structure $D(X)$ to answer queries about X. If X comes from some linear order, then a typical solution to this 'element identification problem' builds a binary search tree to serve as $D(X)$ and uses a binary search tree algorithm for the second stage. Many preprocessing problems arise in Computational Geometry. For example: Given a set X of points in the plane, construct $D(X)$. We want to use $D(X)$ to quickly answer queries of the form: "Retrieve the subset of points in X that are contained in a given half-plane H" (It turns out that we can build a $D(X)$ that uses linear space and answer queries on $D(X)$ in logarithmic time.)

How can we hope to build a theory that encompasses such a wide variety of problems? The key is that any problem that can be solved mechanically by a computer is ultimately encodable into a finite sequence of symbols, where the symbols come from some arbitrary but fixed finite set Σ. This is essentially the argument used by Turing in his celebrated paper [32] that introduce these machines.

A finite set Σ of symbols is called an *alphabet*. A finite sequence of symbols from Σ is called a *word* over Σ. The set of all words over Σ is denoted by Σ^*. A subset L of Σ^* is called a *language* over Σ. The empty sequence, denoted ϵ, containing no symbols is also a word. The *length* of a word w is denoted $|w|$.[7] Thus $|\epsilon| = 0$. We shall embody (part of) the Turing analysis in the following convention.[8]

(A) The input and output objects of a problem are words over some arbitrary but fixed alphabet.

We could have allowed different alphabets for the input and output objects but this is not essential. For each problem P in (i-v) above, we chose some encoding of the input and output domains of P into Σ^*. In (i), we may encode each Rubik cube configuration by specifying for each face (in some predetermined order) the colors of each square on the face. The answer encodes a sequence of moves, where we must choose a systematic representation for the cube faces and their rotations. In (ii), we may use binary numbers to encode natural numbers. Note that each of problems (i) and (ii), after the choice of a suitable encoding, becomes a function $f : \Sigma^* \to \Sigma^*$. In (iii), a graph G is represented by the list (in any order) of its edges. The problem of testing for planar graphs may now

[7] In general, there ought to be no confusion over the fact that we also use $|\cdot|$ to denote the absolute value of a real number as well as the cardinality of a set: it will be clear from context whether we are discussing strings, numbers or sets. One case where this overloading of notation might be confusing will be noted when it arises.

[8] This is analogous to the situation in computability theory where we have the arithmetization or Gödelization of machines. Thus (A) amounts to the 'arithmetization of problems'.

be represented as a language L over Σ where a word w is in L iff w represents a graph G which is planar. In (iv), we have to be more careful. In mathematics, the problem usually assumes that the entries of the matrix A, vectors c and x are arbitrary real numbers. For computational purposes, we will assume that these are rational numbers of arbitrary precision. (In general, replacing real numbers by rationals affects the nature of the problem.) The linear programming problem can be encoded as a binary relation $R \subseteq \Sigma^* \times \Sigma^*$, where $(w, v) \in R$ iff w encodes an input (A, c) and v encodes an output vector x such that $c \cdot x$ is maximized, subject to $Ax \geq 0$. In (v), we may represent X by a sequence of binary numbers. We then encode the element identification problem as a 3-ary relation $R \subseteq \Sigma^* \times \Sigma^* \times \Sigma^*$ where $(u, v, w) \in R$ iff u encodes a set X, v encodes a query on X, and w encodes the answer to the query.

Finite problems such as (i) are usually of little interest to us. In the remaining cases, we find that the problems assume significantly different forms: language in (ii), functions in (iii), and relations in (iv) and (v). These forms are list in order of increasing generality. Which is the appropriate paradigm? It turns out that we choose the simplest form:

(B) A problem is a language.

This answers the question posed as the title of this section.[9] However, it is far from clear that (B) will be satisfactory for studying problems that are functions or relations. The sense in which (B) is reasonable would hopefully be clearer by the end of this chapter. Before we see some deeper reasons (in particular in §8.2, and also §2 of chapter 3) for why (B) is a reasonable choice, it is instructive to first note a simple connection between recognition problems and functional problems.

Consider the following recognition problem: given a triple (x, y, z) of binary numbers, we want to recognize if $xy = z$. This problem is obtained by a transformation of the functional problem of multiplication, and it is easy to see that the transformation is completely general. Clearly an algorithm to solve the functional multiplication problem leads to a solution of this problem. It is not obvious if there is a converse, that is, a solution to the recognition problem will lead to a solution of the functional problem. Fortunately, a slight modification of the recognition problem does have the desired converse (Exercises). Part of the justification of (B) hinges upon such an ability to convert any functional problem of interest to a corresponding recognition problem, such that their complexity are closely related.

2.3 Complexity Functions and Asymptotics

We measure the complexity of algorithms or problems using complexity functions. The following definition captures the class of functions we want:

Definition 1. Let f be a partial function from $\mathbb{R}$ real numbers $\mathbb{R} \cup \{\infty\}$ (the extended reals). If $f(x)$ is undefined, we write[10] $f(x) = \uparrow$, otherwise write $f(x) = \downarrow$. Note that $f(x) = \uparrow$ is distinguished from $f(x) = \infty$. We call f a *complexity function* if f is defined for all sufficiently large natural numbers. The complexity function is *finite* if $f(x) < \infty$ whenever $f(x)$ is defined. ■

Let us motivate some decisions implicit in this definition. We want f to be defined for sufficiently large natural numbers so that comparisons such as "$f(x) \leq g(x)$ for sufficiently large x" is never vacuous. More on this below.

For each algorithm A and computational resource R, we can associate a complexity function $f_{A,R}(n)$ (or, simply $f_A(n)$ if R is understood) such that for inputs of size n, the algorithm uses no more than $f_A(n)$ units of R and for some input of size n, the algorithm uses exactly $f_A(n)$ units of R. By definition, $f_A(n) = \infty$ if the algorithm uses an unbounded amount of R on some input of size n. We also specify that $f_A(x)$ is undefined for x not a natural number.

The definition of complexity functions is basically intended to capture such functions as f_A. The domain and range of f_A is often the set of natural numbers. So why do we extend these functions to real numbers and artificially allow them to be undefined at non-natural numbers? First, the range of f_A is frequently not the natural numbers. For instance, if $f_A(n)$ is the average use of resource R over all inputs of length n. Second, the function f_A may be quite bizarre or difficult to determine exactly, and it is useful to provide an bounds on f_A using nicer or more familiar functions. For instance, it may be enough to know $f_A(n) \leq \sqrt{x}$, even though $f_A(n)$ is really $\lceil \sqrt{n} \rceil - \lfloor 3 \log n + 4.5e^{-n} \rfloor$. These nicer functions are often defined over the reals. This also motivates

[9] Earlier we said that the study of computational problems is relegated to the background in Complexity theory: we see this embodied in assumptions (A) and (B) which, together, conspire to strip bare any structure we see in natural problems. Of course this loss is also the theory's gain in simplicity.

[10] This notation is part of our general notation, repeated here for completeness. Also, we do not write "$f(x)\uparrow$" or "$f(x)\downarrow$", as this has a different meaning.

the use of partial functions, since nice and useful functions such as logarithms and divisios are partial functions. We could artificially convert the real function into a number theoretic function by restricting its domain to $\mathbb{N}$ and introducing ceiling or floor functions to convert the range into natural numbers. E.g., $g(x) = \log_2 x$ may becomes $g(n) := \lceil \log_2 n \rceil$. But this destroys natural smoothness properties, and becomes awkward in the functional composition of two such functions, $(g_1 \circ g_2)(x) = g_1(g_2(x))$. Composition of complexity functions is important because it corresponds to some subroutine calls between algorithms. All these considerations are enshrined in our definition of complexity functions. But that is not all. The astute reader will note that complexity functions tend to be non-negative and monotonically non-decreasing, for sufficiently large values of its arguments; informally call these "growth functions". Why not place these restrictions into the definition as well? The reason is that we need to manipulate and solve complex recurrences involving complexity functions. Sometimes this will take us outside the domain of growth functions.

Quantification over partial predicates. A partial predicate R over a domain D is a partial function $R : D \to \{0, 1\}$. If R is a total function, then it is a **total predicate** (which is the ordinary notion of predicates). We obtain real predicates from complexity functions using the usual relational operators of $\leq, <, =, \neq$, etc. For example, if f, g are complexity functions, consider the real predicate "$R(x) \equiv f(x) \geq g(x)$". This predicate is[11] at x iff both $f(x)$ and $g(x)$ are defined. Our definition of complexity functions ensures that $R(x)$ will be defined when x is a sufficiently large natural number. Given partial predicates $R(x)$ and $S(x)$ over a common domain D, we define their Boolean combinations $R \wedge S, R \vee S, \neg R$ which are again partial predicates. Quantification over partical predicates need to be treated with care. Consider

$$(\forall x \in D)R(x), \quad (\exists x \in D)R(x).$$

These are interpreted as follows: $(\forall x \in D)[R(x) =\downarrow\rightarrow R(x)]$ and $(\exists x \in D)[R(x) =\downarrow \wedge R(x)]$. Exercise: is the following true?

$$\neg(\forall x \in D)R(x) \equiv (\exists x \in D)\neg R(x)$$

Asymptotic Considerations. The complexity function f_A have properties that seem to be nonessential and incidental to the problem for which A is a solution. There are three reasons for this: (a) First, the behavior of f_A for initial values of n seems unimportant compared to the eventual or asymptotic behavior of $f_A(n)$. To see this, suppose that $f_A(n)$ is "unusually large" for initial values of n. For instance, $f_A(n) = 10^{1000}$ for $n < 20$ and $f_A(n) = n$ otherwise. We could replace A by another B which does a 'table look-up' for the first 100 answers but otherwise B behaves like A. Under reasonable assumptions, B uses almost no resources for $n \leq 100$, so $f_B(n)$ is small for $n \leq 100$. Since B is essentially A, this illustrates the incidental nature of initial values of f_A. (b) Second, most algorithms have some degree of model independence which we would like to see reflected in their complexity. For instance, the essential idea of the "mergesort algorithm" (or substitute your favorite algorithm) is generally preserved, whether one implements it in `Java`, `C/C++` or some more abstract computational model. This fact shows when you analyze the complexity of the algorithm under each of these models: they are all related by a constant factor. (c) Third, in some models of computation, one may be able to "speed-up" an algorithm A by desired any constant factor. In chapter 2, we will see this phenomenon in the Turing model: for any Turing machine A, and constant $c > 1$, we can find another Turing machine B such that $f_B(n) \leq n + f_A(n)/c$. We would like f_A and f_B to be regarded as as equivalent.

These considerations motivates the next two definitions.

Definition 2. Let f, g be complexity functions. We say f *dominates* g if there exists a positive constant n_0 such that for all $n \geq n_0$, $f(n) \geq g(n)$. We have a notation for this:

$$f \geq g \text{ (ev.)}$$

(read "$f \geq g$ eventually"). If we want to indicate the variable in this notation, we may also write "$f(n) \geq g(n)$ (ev.n)". If F and G are two sets of complexity functions, we say F *dominates* G if for each $g \in G$ there is an $f \in F$ that dominates it. ■

Note that "$f(n) \geq g(n)$" in this definition is a partial predicate that is defined iff $f(n) =\downarrow$ and $g(n) =\downarrow$. We need to know how to interpret partial predicates under quantifiers. Let $R(x)$ be a partial real predicate and $D \subseteq \mathbb{R}$. When we write

$$(\forall x \in D)[R(x)], \tag{1}$$

[11] If $f(x) = \infty$, it is reasonable to declare that $R(x)$ is define and true, even if $g(x) =\uparrow$. But we will not take this approach.

the full expansion of this predicate is "$(\forall x \in \mathbb{R})[(x \in D \wedge R(x) =\downarrow) \rightarrow R(x)]$". For instance, if $R(x) =\uparrow$ for all $x \in D$, then the statement is vacuously true. Similarly, the full expansion of

$$(\exists x \in D)[R(x)] \tag{2}$$

is "$(\exists x \in \mathbb{R})[(x \in D \wedge R(x) =\downarrow) \wedge R(x)]$". There is no vacuousness when satisfying an existential statement. The usual equivalence "$\neg\forall \equiv \exists\neg$" and "$\forall\neg \equiv \neg\exists$" continues to be trur. See Exercises.

Let us apply this understanding to a generalization of the "eventually notation". We can think of "eventually" as a new quantifier, denoted EV. If $R(x)$ is a partial real predicate. then we write

$$(\mathrm{EV}x)[R(x)]$$

or $R(x)$ (ev.x) if there is some x_0 such that for all $x \geq x_0$, if $R(x) =\downarrow$ then $R(x)$ holds.

The complement of the eventually notation is "infinitely often". It can also be thought of as a new quantifier, denoted IO. We write

$$(\mathrm{IO}x)[R(x)]$$

or $R(x)$ (i.o.x) if for all x there exists $y \geq x$ such that $R(x) =\downarrow$ and $R(x)$ holds. The EV/IO pair of quantifiers are related in analogy to the $\forall/\exists$ pair of quantifiers. See Exercise.

Big-Oh Notation. The next definition is of fundamental importance in complexity theory. There are several formulations of this notation which dates back to Du Bois-Reymond (1871) and Bachmann (1894); and Knuth's[21]. We chose a variant [7] that seems particularly useful:

Definition 3. (The O-notation) For any complexity function f, $O(f)$ denotes the set of all complexity functions g such that for some positive constant $C = C(g)$,

$$0 \leq g \leq C \cdot f \text{ (ev.)}.$$

Note that $O(f)$ is empty unless $f \geq 0$ (ev.). If F is a set of functions, $O(F)$ denotes the union over all $O(f)$ for $f \in F$. We read '$O(F)$' as 'big-Oh of F' or 'order of F'. ■

We extend this definition to recursively defined expressions.

Syntax and semantics of O-expressions. The set of *O-expressions* is defined by the following recursive rule. Let n be a real variable[12], f be any symbol denoting a complexity function, and c any symbol that denotes a real constant. Here are some O-expressions:

$$\begin{array}{rl} \text{Basis:} & f, \quad n, \quad c \\ \text{Induction:} & O(E), \quad E+F, \quad E-F, \quad E \cdot F, \quad E/F, \quad E^F, \quad E \circ F. \end{array}$$

where E, F are recursively defined O-expressions. If E is the symbol f, we may write '$f(F)$' instead of '$f \circ F$' ($\circ$ denotes functional composition). We may freely introduce matching parentheses to improve readability.

Note an O-expression need not contain any occurrences of the symbol 'O'. An O-expression that contains some occurrence 'O' is called an *explicit* O-expression; otherwise it is an *implicit* O-expression. Some examples of O-expressions follow.

$$3 + 2^n, \quad n + O(n^{-1} + \log(n) + 5), \quad O(n^2) + \log^{O(1)} n, \quad f^{O(g)}, \quad (f(g))^{h^2}.$$

Here the real constant 'c' denotes the constant function, $f(n) = c$ for all n; and 'n' denotes the identity function $f(n) = n$, etc.

Each O-expression E denotes a set $[E]$ of complexity functions. We define $[E]$ recursively: (Basis) If $E = f$ is a function symbol, then $[f]$ is the singleton set comprising the function denoted by f. Similarly if $E = n$ or $E = c$, $[n] = \{\iota\}$ and $[c] = \{f_c\}$ where $f_c(n) = c$. If $E = O(F)$ then $[E]$ is defined to be $O([F])$, using definition 3 above. The semantics of the other O-expressions are the obvious ones:

$$[E \circ F] := \{e \circ f \ : \ e \in [E], f \in [F]\}, \quad [E + F] := \{e + f \ : \ e \in [E], f \in [F]\}, \quad \text{etc.}$$

[12]The literature uses 'n' for the variable of (univariate) complexity functions, in part because domains of complexity functions are often restricted to natural numbers. We continue to use 'n' with complexity functions even though we intend to let n range over all real numbers. The tradeoffs in making this choice is between the potential of misleading the reader, and the familiar contexts of such complexity functions. We opt for the latter.

Henceforth, we intend to commit the usual linguistic abuse, by letting symbols (syntactical objects) such as f and c stand for the functions (semantical objects) that they denote, when the context demand a function rather than a symbol. This abuse should not lead to any confusion.

The main use of O-expressions is this: If F is an O-expression and E is an explicit O-expression, we may say

$$\text{`}F \text{ is } E\text{' and write `}F = E\text{'}.$$

This simply means that $[F] \subseteq [E]$. An important special case is where $F = f$ and $E = O(g)$ where f, g are function symbols. For instance, if $t(n)$ denotes the time complexity of the high school multiplication algorithm then we may write '$t(n) = O(n^2)$'. Note that we restrict E to be an explicit O-expression since there is potential for confusion otherwise.

The use of the equality symbol in this context is clearly an abuse of our usual understanding of equality: we do not regard '$E = F$' and '$F = E$' as interchangeable. Thus $O(n^2) = O(n^3)$ is true but $O(n^3) = O(n^2)$ is not. For this reason, the equality symbol here is called the *one-way equality*.

A corollary of the *inclusion interpretation* of the one-way equalities is this: an *inequality* involving O-expressions is to be interpreted as *non-inclusion* '$\not\subseteq$'. Thus '$n \log n \neq O(n)$' is true but '$n \neq O(n \log n)$' is false.

Some authors avoid the one-way equality symbol by writing the more accurate form

$$\text{`}F \subseteq E\text{'}$$

or, in the case F is a function symbol f, '$f \in E$'. We shall not use this alternative form.

Example 1. Further examples of usage:

(i) The O-expressions $O(1), O(n)$ and $n^{O(1)}$ denote, respectively, the set of all functions that are eventually dominated by constants, by linear functions and by polynomial functions. For instance $\frac{1}{n} = O(1)$ even though the value becomes unbounded when $n \to 0$.

(ii) Depending on the context, the appearance of an O-expression may sometimes denote a set of functions, rather than as denote some unspecified member in this set. For instance, when we write $NP = NTIME(n^{O(1)})$, we intend to use the full set of functions denoted by the expression "$O(1)$". This is a potential pitfall for students.

(iii) By the non-inclusion interpretation, $f \neq O(g)$ means that for all $c > 0$, there are infinitely many n such that $f(n) > cg(n)$. There is, however, a useful intermediate situation between $f = O(g)$ and $f \neq O(g)$: namely, there exists $c > 0$ such that for infinitely many n, $f(n) > cg(n)$. Below, we will introduce a notation to express this.

(iv) The O-notation factors out unessential features of complexity functions: thus we say '$f(n) = O(\log n)$' without bothering to specify the base of the logarithm.

(v) The notation saves words: we simply write '$f(n) = n^{O(1)}$' instead of saying that $f(n)$ is dominated by some polynomial. It also conserves symbols: we say '$f(n) = O(1)^n$' instead of saying that there exists $n_0, c > 0$ such that $f(n) \leq c^n$ for all $n \geq n_0$. Thus we avoid introducing the symbols c and n_0.

(vi) It is handy in long derivations. For example,

$$n + O(n^{-1}) = n + O(1) = O(n) = O(n^2).$$

■

A subscripting convention.[13] We augment the O-notation with another useful convention:

(i) It often occurs that we want to pick out some *fixed but non-specific* function f' in the set $O(f)$, depending on some parameter α. Then we write $O_\alpha(f)$ to refer to this f'. If f' depends on several parameters $\alpha, \beta, \ldots$ then we write $O_{\alpha,\beta,\ldots}(f)$. An O-expression is *fully subscripted* if each occurrence of 'O' is subscripted. So a fully subscripted O-expression refer to a single function.

(ii) As an extension of the one-way equality, the 'equality symbol' between two fully subscripted O-expressions denotes domination between two functions. For instance, we write '$g = O_\alpha(f)$' to indicate that g is dominated by some $f' \in O(f)$ whose choice depends on α. There is much room for syntactic ambiguity and we assume common sense will avoid such usage (for instance, an implicit O-expression E is, by definition, fully-subscripted and it would be confusing to write '$g = E$').

[13] This convention is peculiar to this book.

(iii) Sometimes we want to annotate the fact that among the various occurrences of an O-expression, some refers to the same function. We may use subscripts such as O_1, O_2, etc., as an 'in-line' device for showing this distinction.

Example 2. We illustrate the subscripting convention.

(i) We say an algorithm A runs in time $O_A(n)$ to mean that the running time is dominated by kn for some k depending on A.

(ii) We could write $n2^k = O_k(n)$ as well as $n2^k = O_n(2^k)$, depending on the context.

(iii) Suppose $g(n) = O_1(n^2 \log n)$ and $f(n) = O_2(n) + 3O_1(n^2 \log n)$. The two occurrences of 'O_1' here refer to the same functions since they subscript identical O-expressions. Then we may write their sum as $f(n) + g(n) = O_2(n) + 4O_1(n^2 \log n) = O_3(n^2 \log n)$.

(iv) We emphasize a fully-subscripted O-notation no longer denotes a set of functions. Hence it is meaningful to write: '$O_1(f) \in O(f)$'.

(v) We illustrate a more extensive calculation. The following recurrence arises in analyzing the running time $T(n)$ of certain list manipulation algorithms (e.g., the 'finger tree' representation of lists in which we admit operations to split off prefixes of a list): $T(1) = 1$ and for $n \geq 2$,

$$T(n) = \max_{1 \leq i < n} \{T(i) + T(n-i) + \log \min\{i, n-i\}\}$$

To see the subscripting convention in action, we now prove that $T(n) = O_1(n) - O_2(\log(2n))$. If $n = 1$, this is immediate. Otherwise,

$$\begin{aligned}
T(n) &= \max_{1 \leq i < n} \{O_1(i) - O_2(\log(2i)) + O_1(n-i) - O_2(\log(2(n-i))) \\
&\qquad + \log \min\{i, n-i\}\} \\
&= \max_{1 \leq i < n} \{O_1(n) - O_2(\log(2i)) - O_2(\log(2(n-i))) + \log \min\{i, n-i\}\} \\
&\leq \max_{1 \leq i \leq n/2} \{O_1(n) - O_2(\log(2i)) - O_2(\log(2(n-i))) + \log i\} \\
&\qquad \text{(since the expression is symmetric in } i \text{ and } n-i) \\
&= O_1(n) - \min_{1 \leq i \leq n/2} \{O_2(\log(2(n-i)))\} \\
&\qquad \text{(we may assume } O_2(\log(2i)) \geq \log i) \\
&= O_1(n) - O_2(\log n)
\end{aligned}$$

■

Note that we have used "$\leq$" in the above derivation. This is something we would not write without the accompanying subscripting convention.

We shall have occasion to use four other related asymptotic notations, collected here for reference:

Definition 4. (Other asymptotic notations)
(i) $\Omega(f)$ denotes the set of functions g such that there exists a positive constant $C = C(g)$ such that $C \cdot g \geq f \geq 0$ (ev.).
(ii) $\Theta(f)$ denotes the set $O(f) \cap \Omega(f)$.
(iii) $o(f)$ denotes the set of all g such that for all $C > 0$, $C \cdot f \geq g \geq 0$ (ev.). This implies that ratio $g(n)/f(n)$ goes to zero as n goes to infinity. Also $o(f) \subseteq O(f)$.
(iv) $\omega(f)$ denotes the set of all g such that for all $C > 0$, $C \cdot g \geq f \geq 0$ (ev.). This implies the ratio $g(n)/f(n)$ goes to infinity as n goes to infinity. ■

These notations are used in ways analogous to that in O-notations; in particular, we use one-way equalities such as '$f = o(g)$'. The subscripting convention extends to the these notations. The O-notation and the Ω-notation are inverses in the sense that

$$f = O(g) \iff g = \Omega(f).$$

Similarly,

$$f = o(g) \iff g = \omega(f).$$

To verbally distinguish the O-notations from the o-notations, we also call them the "Big-Oh" and the "Small-Oh" (or, "Little-Oh") notations, respectively. Clearly, $f = o(g)$ implies $f = O(g)$. Also $f = \Theta(g)$ iff $g = \Theta(f)$ iff $f = O(g)$ and $g = O(f)$.

We sometimes call the set $O(f)$ the *big-Oh order of* f, and if $g = O(f)$, we say g and f *have the same big-Oh order*. Similarly, $\Theta(f)$ is the *big-Theta order of* f, etc.

We warn that incorrect usage of these less-frequently used asymptotic notations is common. Furthermore, they are sometimes given rather different definitions. For instance, in [1] '$\Omega(f)$' denote the set of functions g such that for some c, $g(n) \geq cf(n)$ for infinitely many n. (This is the intermediate situation between $g = O(f)$ and $g \neq O(f)$ mentioned above.) But notice that '$g = \Omega(f)$' under the definition of [1] is recaptured as our notation

$$g \neq o(f).$$

Although we seldom use this concept in this book, it is a typical situation that obtains when one proves "lower bounds" in concrete complexity.

Three final notes on usage:
(a) Clearly we could introduce Ω-expressions, etc. The meaning of *mixed* asymptotic expressions such as $O(n^2) + \Omega(n \log n)$ can also be defined naturally. A proper calculus of such expressions needs to be worked out. Fortunately, there seems to be little need of them.
(b) The reader must have noticed the similarities between the expression $f = O(g)$ and the inequality $x \leq y$.[14]. In fact, the similarities extend to the other notations:

$$f = O(g), \quad f = \Theta(g), \quad f = \Omega(g), \quad f = o(g), \quad f = \omega(g)$$

are analogous (respectively) to the inequalities

$$x \leq y, \quad x = y, \quad x \geq y, \quad x << y, \quad x >> y$$

on real numbers x, y. Such analogies has has led authors to write expressions such as

$$f \leq O(g), \quad f \geq \Omega(g). \tag{3}$$

Why write $f \leq O(g)$ when you could write $f = O(g)$? One strong motivation is the wish to refer to some anonymous function g_1 in $O(g)$ and we intend to do a pointwise comparison $f \leq g_1$. But this is precisely our rationale for introducing the subscripting convention. In any case, once we admit (3), it is hard to avoid writing $f \geq O(g)$ (again, the appearance of '$\leq$' is supposed to indicate that we are referring to some anonymous $g_2 \in O(g)$ such that $f \geq g_2$). But it would appear odd indeed to write "$n \geq O(n^2)$". This could also lead to trouble[15] since one is next tempted to manipulate these expressions under the usual rules of inequalities. On the other hand, a fully subscripted O-expression (say) refers to a particular function and inequalities involving such an expression can be interpreted as domination and manipulated confidently: for instance, we do not consider the expression $n^2 \geq O_1(f(n))$ to be inappropriate. In short, *we never use inequality symbols* $\leq$ *or* $\geq$ *with expressions containing unsubscripted occurrences of the* O*-notation.*
(c) One could extend the definition of complexity functions and all the associated asymptotic notation to admit multiple parameters such as $f(m, n)$. This becomes necessary when discussing complexity classes defined by several simultaneous resource bounds.

Additional Notes: See Hardy and Wright [14] for the treatment of similar notations. Knuth's original definition of the O-notation goes as follows: $g = O(f)$ if there are constants n_0 and $c > 0$ such that for all $n > n_0$, $|g(n)| \leq c \cdot f(n)$. Our definition departs from Knuth (who is closer to the classical precedents) in that we use $g(n)$ instead of the absolute value of $g(n)$ in the above inequality. Also, following [7], we added a non-negativity requirement $g \geq 0$ (ev.). Like Knuth, however, we do not take the absolute value in the Ω-notation; Knuth explains the asymmetry by thinking of 'O' as referring to the neighborhood of zero and 'Ω' as referring to a neighborhood of infinity. For an updated discussion of these notations, see [34] [3] [12].

2.4 Size, Encodings and Representations

Complexity is a function of the size of the input. In the natural setting of certain problems, a correct choice for the size parameter is often not obvious. For example, let D be the set of square matrices with rational entries. An $n \times n$

[14] For instance, viewing "$f = O(g)$" as a binary relation, then reflexivity ($f = O(f)$) and transitivity ($f = O(g), g = O(h)$ imples $f = O(h)$) holds. Even anti-symmetry holds: $f = O(g), g = O(f)$ imples $f = \Theta(g)$.
[15] As indeed has happened in the literature.

matrix M in D has three candidates for its size: its dimension n, the number of entries n^2, and the total number of bits to represent all the entries. The convention (A) above removes this problem: if the matrix M is encoded as a string w over Σ, then we define the size of M to be $|w|$. In general, let D be any (mathematical) domain (such as matrices, integers, finite graphs, etc). An *encoding* of D is a function $e : D \to \Sigma^*$ (for some alphabet Σ) that is one-one. Relative to e, the *size* of $x \in D$ is simply defined as $|e(x)|$.

The use of encoding solves the problem of defining size but it introduces other issues. In particular, the complexity of the problem depends on the choice of encoding. We may be willing to accept two different encodings of the same mathematical problem as really two distinct computational problems, but that is perhaps too easy a way out. This is because two encodings may be different for rather trivial reasons: for instance, suppose each input x to a problem P is encoded by a word $e(x) \in \Sigma^*$, and the complexity of P under the encoding e is $f(n)$. For any constant $k > 0$, we can choose another encoding e' such that for each input x, $|e'(x)| \leq \frac{|e(x)|}{k}$ (how?). With reasonable assumptions, we see that the complexity of the problem under e' is $g(n) = f(n/k)$. In this sense $g(n)$ and $f(n)$ are essentially the same.[16]

The next problem with encodings is that $e : D \to \Sigma^*$ may not be an onto function so a certain word $w \in \Sigma^*$ may not encode any element of D. Let us call w *well-formed* if it does encode some x in D; otherwise w is *ill-formed*. Should we assume that the algorithm need only restrict its attention to well-formed inputs? If so, the complexity function becomes undefined at those values of n where there are no well-formed inputs of length n. (For example, if D is the set of square boolean matrices and e encodes a boolean matrix in row-major order i.e. listing the rows where successive rows are separated by a marker symbol, then all well-formed inputs have length $n^2 + n - 1$.) To simplify things somewhat, we make the following decision:

(C) All words in Σ^*, well-formed or not, are possible inputs.

This decision is inconsequential if the set of well-formed words can easily be recognized: in many problems, this is indeed the case. Otherwise, given an algorithm that works on only well-formed words, we modify it to work on all inputs simply by attaching on the front-end a 'parsing phase' to screen the inputs, rejecting those that are not well-formed. If this parsing phase is expensive relative to the actual computation phase, then the complexity of this problem may turn out to be an artifact of the encoding. In section 5.3 we shall see examples where the parsing complexity is intrinsically high. In any case, the abandonment of (C) should lead to interesting variants of the theory.

One way to avoid an intrinsically hard parsing problem is to generalize the notion of encodings by allowing several words to represent the same object. More precisely, a *representation* r of a domain D over Σ is an onto partial function

$$r \ : \ \Sigma^* \to D.$$

We call $w \in \Sigma^*$ an *r-representative* of x if $r(w) = x$. The fact that r is onto means that every object has at least one representative. Since r is a partial function, let us say $w \in \Sigma^*$ is *well-formed* or *ill-formed* (relative to r) depending on whether $r(w)$ is defined or not.

If r is also 1-1, then in fact the inverse of r is an encoding. We have two computational problems associated with any representation r:

1. The *r-parsing problem* is the problem of recognizing well-formed words.
2. The *r-isomorphism problem* is the problem of deciding when two words w, w' are r-representatives of the same element: $r(w) = r(w')$.

If r is in fact an encoding then the isomorphism problem is trivial. Thus, the use of representations has simply shifted the parsing complexity to the isomorphism complexity. This suggests that there is an inherent complexity associated with certain domains D in the sense that every representation of D has a hard parsing problem or a hard isomorphism problem (Exercise).

It is clear that a representation r cannot be arbitrary if the notion of complexity is to be meaningful: consider the following encoding of the domain D of finite graphs, and let $P \subseteq D$ be the set of planar graphs. If for $G \in D$, the first symbol in the encoding $e(G)$ is '1' iff $G \in P$, then clearly the encoding e is rather contrived and particularly ill-suited for encoding the planarity testing problem. We have no theoretical guidelines as to how representations must be restricted. Fortunately, it is not as sorry a state of affairs as we might imagine. Despite the lack of a general theory, for most problems that arise, either there is a consensus as to the natural representation to use, or

[16]This seems to argue for a notation for the set of all g such that $g(n) = O(f(O(n)))$: denote this set by $\odot(f)$. Thus $5^n = \odot(2^n)$ but not $5^n = O(2^n)$. The O- and $\odot$-notations coincide when restricted to polynomials. In Complexity Theory, we should perhaps only distinguish complexity functions up to their $\odot$-order; after all complexity practitioners normally do not distinguish between 5^n and 2^n.

else the different choices of representations are equivalent in some sense. In the remainder of this section we show how this is so.

We give an initial informal criterion for choosing between various natural representations (without actually resolving the issue of which representations are 'natural'):

(D) Let r and r' be two natural representations for a problem P. Then r is to be preferred over r' if the following holds:

(D1) The parsing problem for r is easier.
(D2) The isomorphism problem for r is easier.
(D3) The intrinsic complexity of P under r is more than that under r'.

We would apply the criteria (D1-3) in the indicated order: for instance, if the parsing problem for r and r' are equivalent, but the isomorphism problem for r is easier that for r', we would prefer r over r'. If criteria (D1-3) cannot distinguish between two representations, then it turns out that, in practice, there is no reason to differentiate between them anyway.

We have seen the reasons for (D1, D2). Observe that (D2) implies that we should prefer encodings over general representations. The intuitive reasoning for (D3) begins with the observation that some representations are more compact (or, more succinct or efficient) than others. The intuitive notion of the most compact representation of a problem seems to be meaningful, but the notion of least compact representation does not. For, it is easy to imagine representations that introduce an arbitrary amount of irrelevant data ('padding') or that are arbitrarily redundant. If a problem is encoded more compactly than another, then the complexity of the more efficiently encoded one tends to be greater (not less!). Hence (D3) rejects redundant or padded representations.

Example 3. The contrast between compact and non-compact representations is illustrated by the following results. Suppose we want to choose between two encodings of positive integers: the usual unary encoding (a positive integer n is represented by a string of n zeroes) and the 'exponential unary encoding' (n is representated by a string of 2^n zeroes). A well-known result of Minsky[23] says that every partial recursive function can be computed by a 2-counter machine (see exercises in chapter 2 for a definition of counter machines), assuming the exponential unary encoding of numbers. A less well-known result, independently obtained by Schroeppel[31] and by Frances Yao[35] says that no 2-counter machine can compute the function $f(n) = 2^n$, assuming the unary encoding of numbers. Now by criteria (D1) and (D2), there is no distinction between these two encodings. But (D3) says the usual unary encoding is preferable. ■

Despite the ad hoc nature of (D), we are able to use it in some common examples illustrated next.

2.4.1 Representation of Sets

Suppose D' is the domain of finite subsets of another domain D, and $e : D \to \Sigma^*$ is an encoding of D. We can extend e to a representation r of D': let $\#$ be a symbol not in Σ. The word

$$e(x_1)\#e(x_2)\#\cdots\#e(x_k) \tag{4}$$

is called a *standard representative* of $X = \{x_1, \ldots, x_k\} \in D'$, provided $x_1, \ldots, x_k$ are distinct. Since the order of the x_i's are arbitrary, X has $k!$ (k factorial) standard representatives. If $e(x_1), \ldots, e(x_k)$ are in ascending lexicographical order (assuming an arbitrary but fixed ordering on Σ) then (4) is unique and is called the *canonical encoding* of X. If e were a representation to begin with, then the standard representation of D' is still well-defined but the canonical encoding is not well-defined.

Consider the problem of recognizing membership in P where $P \subseteq D'$. Let

$$L_c = \{e'(x) : x \in P\}$$

where e' is the canonical encoding, and

$$L_s = \{w \ : \ w \text{ is a standard representative of some } x \in P\}.$$

Both L_c and L_s appear to be natural, so we attempt to distinguish them using the criteria (D). Let f_c and f_s be the complexity of L_c and L_s (respectively). Under reasonable assumptions, it is seen (Exercise) that f_c and f_s are related as follows:

$$f_c \leq f_s + f_0 \tag{5}$$

$$f_s \leq f_c + f_1 \tag{6}$$

where f_0 is the complexity of deciding if a word of the form (4) is in ascending order, and f_1 is the complexity of sorting a sequence of words $e(x_1), \ldots, e(x_k)$ given in the form (4) into ascending order. Now in typical models of computation and complexity measures, we have

$$f_0(n) = O(n) \text{ and } f_1(n) = O(n \log n).$$

If f_c and f_s are $\Omega(n \log n)$, then (5) and (6) implies that $f_c = \Theta(f_s)$, i.e., they are indistinguishable in the sense of (D). This is a reassuring conclusion: it does not matter which representation of sets is used, provided L_c and L_s have large enough complexity. We will adopt the convention:

(E) Sets are given by their canonical or standard representation.

2.4.2 Representation of Numbers

The usual choice for representing natural numbers $1, 2, 3, \cdots$ is the k-ary ($k \geq 1$) notation over the alphabet $\Sigma = \{0, 1, 2, \ldots, k-1\}$. Clearly the unary ($k = 1$) notation is an encoding but for $k > 1$, we only have a representation since a prefix string of zeroes does not affect the value of a k-ary number. In this latter case, we can get an encoding by restricting the k-ary numbers to those that begin with a non-zero symbol. Note that criteria (D1-2) do not distinguish the various choices of k, so we need to test with (D3).

Consider the problem of recognizing a set P of natural numbers (such as the set of primes: $2, 3, 5, 7, 11, 13, \cdots$). Let $L_k \subseteq \{0, 1, \ldots, k-1\}^*$ be the set of k-ary numbers in P and f_k be the complexity of L_k. Consider the k- and l-ary encodings for any fixed $k, l > 1$. Then

$$f_k(n) = O(f_l(n)) + O(n^2). \tag{7}$$

This comes from the fact that (using reasonable computational models) there are algorithms converting between k-ary and l-ary notations that run in $O(n^2)$ time. Furthermore, for all natural numbers m, if v and w are k-ary and l-ary numbers (respectively) encoding the same number m, then their lengths are related by $|v| = \Theta(|w|)(= \Theta(\log m))$. (Exercise) We conclude that all k-ary notations ($k > 1$) are indistinguishable in the sense of (D) for problems with complexity $\Omega(n^2)$.

Unfortunately the above result does not hold for the unary notation. We note an exponential discrepancy between unary and k-ary ($k \geq 2$) notations: if u (in unary notation) and v (in k-ary notation, $k > 1$) represent the same integer, then $|u| = \Theta(k^{|v|})$. Therefore to convert a k-ary number v to a unary number takes $\Omega(k^{|v|})$ time since this is the time just to write each output symbol. Therefore, if f_k is subexponential (i.e., $f_k(n) = \omega(c^n)$ for any $c > 0$) then f_1 is a slower (*sic*) growing function than f_k. (What if f_k is at least exponential?) Criterion (D) then says we must reject the unary notation. In conclusion, the following convention will be used.

(F) Integers are encoded in k-ary for some $k > 1$.

There are other representations of numbers besides k-ary notations. For example, from basic number theory we know that each number greater than 1 has a unique decomposition into a product of powers of distinct prime numbers. We may encode such a number by a sequence of non-negative integers (say in binary notation) $p_1 \# x_1 \# p_2 \# \cdots \# p_k \# x_k$ where $p_1 < p_2 < \cdots < p_k$ are primes and each $x_i \geq 1$. This sequence represents the number $p_1^{x_1} p_2^{x_2} \cdots p_k^{x_2}$. Multiplication under this encoding is linear time, but addition seems hard. It is not easy to dismiss this notation as 'unnatural' (a number theorist may not think so). We may then ask if addition is intrinsically harder than multiplication under this encoding, a curious reversal of Cobham's question in Section 1. This points out that the choice (F) at this point of our understanding is somewhat ad hoc.

An equally reasonable alternative to (F) would be to assume the *k-adic* notation: for each $k \geq 1$, we have an isomorphism between the strings over $\Sigma = \{1, \ldots, k\}$ and the non-negative numbers where the correspondence is given by

$$a_0 a_1 \cdots a_n \in \Sigma^* \leftrightarrow \sum_{i=0}^{n} a_i k^i (n \geq 0).$$

In other words[17], k-adic notation differs from k-ary notation only in its use of the integers $\{1, \ldots, k\}$ instead of $\{0, \ldots, k-1\}$. The k-adic notation avoids the well-known problem of non-uniqueness of k-ary notation.

[17]The term 'unary' as it is normally used is really a misnomer: it should be 'unadic'. With this warning, we will perpetrate the abuse. It seems that Quine [26] is one of the first to use this notation.

2.4.3 Representation of Graphs

In graph theory, two different graphs g and g' are usually identified if they are isomorphic, i.e., the actual identity of the nodes is considered unimportant. In Computer Science, we normally distinguish between g and g'. To emphasize this distinction we say that graph theory treats *unlabeled graphs* while Computer Science treats *labeled graphs*. In this book, the word *graph* always denotes an undirected, labeled graph; we permit self-loops but not multiple edges in our graphs. There are three well-accepted methods of representing graphs: (1) Adjacency matrices. (2) Adjacency lists. (3) Edge lists (i.e., lists of unordered pairs of vertices). It is not hard to see that these three representations are interconvertible in time $O(n^2)$ in most reasonable models, where n is the number of nodes in the graph. This amounts to a linear time conversion for dense graphs (i.e., with $\Omega(n^2)$ edges) but in any case at most quadratic time. Thus criterion (D) justifies the following.

(G) Graphs are represented using any of the three representations of adjacency matrix, adjacency list or edge list.

Now consider the problem of representing unlabeled graphs. We can use a representation where each encoding $e(g)$ of a labeled graph g (using any of the above 3 methods) is said to be a representative of its unlabeled counterpart G. Thus the unlabeled graph G has $n!$ representations. The isomorphism problem for this representation is the well-known problem of *graph isomorphism* for which no polynomial time algorithm is known. So this is somewhat unsatisfactory. On the other hand, we can define an encoding of G by using a *canonical labeling* of G: let g be the labeled version of G such that $e(g)$ is lexicographically minimal among all other labelings of G. There is theoretical evidence that the parsing problem in this case is hard, and using criteria (D1), this canonical encoding should be rejected in favor of the previous representations. Perhaps it is not our inability to find representations with easy isomorphism problems or encodings with easy parsing problems, but that such representations simply do not exist. This suggests the problem of classifying mathematical domains (using the tools of mathematical logic) according to 'the inherent complexity of their representations'. Other domains whose representations have apparently large complexity include many groups and rings. Specifically, the representation of multivariate polynomials (a basic example of rings) is an important one with repercussions in the attempts to classify the complexity of algebraic problems such as polynomial factorization.

In summary, in this section we made a critical decision (C) to allow all inputs, and we gave some informal criteria for how to use our theory (namely, how to represent a problem so that the notions of size has some natural meaning). As for decision (C), we are well-aware of its shortcomings. We will see instances where we would like to analyze the complexity of an algorithm A relative to some proper subset of possible inputs. For instance, if A is used to operate on the outputs of another algorithm A' (this happens when we study reducibilities, where we use A' reduce another problem to the problem solved by A). As for the questions of representing problems, we have seen that for many problems with sufficiently large complexity (in the cases examined, complexity of $\Omega(n^2)$ is enough) any of a number of natural representation are equivalent (relative to criterion (D)). *In the rest of this book, when we discuss natural problems involving numbers, sets or graphs, we normally will not be explicit about their representation because of the preceding conclusions* (*assumptions* (E),(F) *and* (G)). For the interested reader, the following is a very incomplete list of additional references on the relationship between representation and complexity: [33], [22], [19].

2.5 Models of Computation

The theory of computability [28] begins with the fundamental question:

(1) What is the concept of computability?

For definiteness, assume we ask this question of number theoretic functions. Naturally, 'computable" must be defined in the context of some suitable formal system F of computing devices. It must be emphasized that we require F to be "sufficiently general" in the sense that any function that we all intuitively understand to be computable should be computable within the formalism of F. For instance, we intuitively believe that the multiplication of two numbers, the function $f(n)$ that produces the sum of the first n primes, the function $f(n) = \lfloor \int_{a=0}^{n} x^3 dx \rfloor$ should be computable. Logicians interested in the foundations of mathematics introduced different formalizations of F. The *equational calculus* of *Herbrand* and *Hilbert*, the *λ-definable functions* of *Church*, the *μ-recursive functions* of *Kleene* and *Gödel*, and the machines of *Turing* are among these formalisms. For each F, we now have the concept of F-computable functions. It turns out to be natural to extend the concept to *partially F-computable functions*. The pleasant surprise is that for any two of these formalisms F and F', one can prove that a function is (partially) F-computable if and only if it is (partially) F'-computable. These discoveries are elevated into a general law called

Church's thesis[18] [4]: *for all general computational formalisms, the concepts of computable and partially computable are invariant.* We have now provided a satisfactory answer to (1).

Furthermore, we conclude that all partially computable functions can effectively be enumerated (since we can clearly list, say, all Turing machines and conclude from Church's thesis that all partially computable functions are thereby named). It is not hard to use an older diagonalization argument due to *Cantor* to conclude that there exists an uncomputable function. In 1931, Gödel[10] showed that the validity of statements in number theory is undecidable. This landmark discovery demonstrates that uncomputable functions not only exist, but occur naturally.[19] Such discoveries give rise to a central theme of computability theory:

(2) Which functions are computable and which are not?

Note that this question makes sense because of Church's thesis. Extending this fundamental distinction, logicians went on to classify the uncomputable problems into 'degrees of uncomputability'. Unfortunately, they rarely try to classify the computable problems. In a sense, computability theory and Complexity Theory are really one subject: the latter does for the computable problems what the former does for the uncomputable ones. Complexity theory draws many of its methods (e.g., diagonal arguments) and concepts (e.g., reducibilities) from computability theory. Many phenomena at the uncomputable levels are mirrored at the lower levels. However, many of these phenomena become considerably more subtle at the lower levels. For instance, many questions that have been solved at the higher level remain open at the lower level. In this and the next section, we shall examine the complexity theoretic versions of (1) and (2).

Our view must necessarily be richer because researchers noticed that not all *forms of computation* are equivalent. The word 'computational form' is advisedly chosen in this context; it is intended to capture the concrete versions of computational models, before any attempt to put them into equivalence classes.[20] And yet, researchers also noticed that many forms of computation *are* equivalent. Here, "equivalence of computational form" is somewhat vague but in each concrete situation, we can make correspondences across forms. For example, 'time' resource can be identified in two computational forms and some suitable notion of equivalence (mutual simulation in polynomial time) defined. This not-all-forms-are-equivalent and many-forms-are-equivalent phenomena at first appear confusing. The former says that we have lost Church's thesis; the latter suggests that there might be some hope. Before pointing a way out of this paradise-lost, we need to sharpen our concept of computational form. Let us begin with examples of computational forms (this list is incomplete and includes computational forms that are not "general" in the sense demanded of F-formalisms above):

> Turing machines, storage modification machines, arrays of finite automata, pushdown automata, finite state automata, random access machines, vector machines, aggregates, formal grammars, various proof systems, lambda calculus of Church, μ-recursive functions of Kleene and Gödel, Herbrand and Hilbert's equational calculus, Post's canonical systems, Markov algorithms, Shepherdson and Sturgis' register machines, random access stored programs of Elgot and Robinson, Elementary Formal Systems of Smullyan.

The reader is not expected to know any in this list. We propose to characterize computational forms along two orthogonal directions:

(i) *Model of computation*: this identifies the basic computational structures (control, instructions and data structures). Several computational forms will be collapsed into a common model corresponding to our intent that members of the same model use similar computational structures: thus finite state automata and Turing machines will be regarded as falling under the same model (the Turing model).

(ii) *Mode of computation*: this refers to the method of using the computational structures to define computation. For example, the Turing model can compute in the deterministic, probabilistic or nondeterministic mode (see next section).

A computational form is defined by specifying the model and mode. Computational models and modes are essentially independent notions in the sense that within each model we can define machines operating in any

[18] Also called the Church-Turing Thesis. Note that, as "discoveries", we have a collection of descriptive statements. But as a law, it carries a prescriptive weight – henceforth any new formalism F' that diminishes or enlarges the class of computable functions is inadmissible.

[19] It seems that the existence of such functions per se seems uninteresting unless one can show 'natural' examples. There is an analogous situation in Complexity Theory: although we already know from the work of Hartmanis that there are arbitrarily hard-to-recognize languages, greater interest was generated when Meyer and Stockmeyer showed that such languages occur naturally. See chapters 5 and 6. Another recent example in logic is the work *Paris-Harrington* in combinatorial number theory.

[20] Taken literally, any perceptible difference between two models should lead to distinct computational forms. But this cannot be taken too seriously either. Anticipating our "model-mode" view of the world to be discussed shortly, a computational form is a mixture of model and mode.

given mode.[21] However, this must be qualified because some computational models were specifically designed for particular modes and do not naturally embrace other modes. For example, grammars or production rules in formal language theory essentially compute in nondeterministic modes, and do not to easily embrace some other computational modes. Or again, the Boolean circuit model[22] one single is essentially a parallel model and it is not very natural to define sequential modes for circuits. We will say that a model is *general* if it does embrace all modes. For instance, the Turing model and the pointer machine model will be seen as general. Since the concept of computational modes post-dates the definitions of these models, we sometimes take the liberty of modifying original definitions of some models to make them general. The remainder of this section is a brief exposition of some important computational models; computational modes will be treated in the next section.

We now examine the computational structures that characterize some important computational models.

(i) *Turing model*. This will be covered in chapter 2 and is the most important one for this book. Essentially each computing agent in this model has finitely many states and symbols, and a finite number 'heads' each of which scans some memory location. The memory locations have some fixed neighborhood structure (which is specified by a bounded degree undirected graph with locations as vertices). Each location can store a symbol. The simplest example of a fixed neighborhood structure is that of a total linear relation, giving rise to what is usually known as a Turing machine tape. The computing agent has a finite set of instructions specifying what to write into scanned locations and how to move its heads when in a given state scanning certain symbols.

(ii) *Pointer machine model*. This model, also known as the *storage modification machine model*, was introduced by Schönhage in 1970 [29]. Similar models were earlier proposed by *Kolmogorov and Uspenskiĭ*, and by Barzdin and Kalnin'sh. Like Turing's model, this uses finite state control to determine the step-by-step execution of instructions. The essential difference comes from having a varying neighborhood structure that is represented by a fixed out-degree (arbitrary in-degree) directed graph G. We also fix a finite set Δ of symbols; each $a \in \Delta$ is called a **color**. The edges of G are called *pointers*, and are labeled by some color in Δ. Every node has outdegree $|\Delta|$ and outgoing edges from a node have distinct colors. One of the nodes is designated the **center**. Call G a **Δ-structure**. Each word $w \in \Delta^\star$ is said to **access** the unique node obtained by following the sequence of pointers labeled by symbols in w, starting from the center. Let this node be denoted $[w]_G$ (or simply $[w]$ when G is understood). Thus the node $[\epsilon]$ is the center. In Figure 2.1, the center is indicated by an arrow from nowhere, with no label. Such graphical representations use two conventions to reduce clutter: (a) if a pointer is a self-loop (i.e., its target and source are the same, we may omit them from the diagram. (b) if two or more pointers share the same source and target, then we only draw one edge and label them with the list of colors for these pointers.

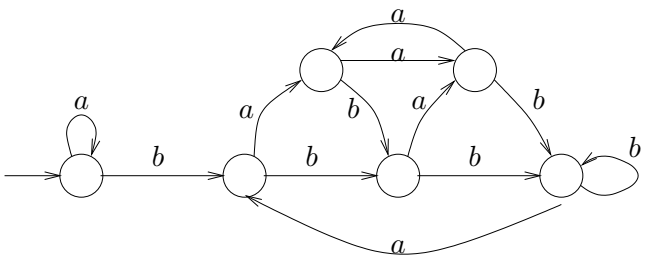

Figure 2.1: Pointer machine Δ-structure ($\Delta = \{a, b\}$). The center is indicated by the arrow without any label.

We define a **pointer machine** (for any color set Δ) as a finite sequence of instructions of the following types:

(i) $w := w'$ (assignment statement)

(ii) $w :=$ **new** (node creation statement)

(iii) **if** $w' = w$ **goto** L (branch statement)

[21] Up till now, we have used the word 'algorithm' to describe the computing agent. When we discuss particular forms of computation, we traditionally refer to 'machines' of that form. A useful distinction (which we try to adhere to) is that 'algorithms' are abstract objects while 'machines' are their concrete representations within a form. Thus, we have 'the' high school multiplication *algorithm*, but various Turing *machines* or PASCAL *programs* can be regarded as implementing the same high-school algorithm. Another term which we regard as essentially synonymous with 'machine' is 'program'.

[22] Actually, circuits raises another more annoying issue for our attempt to classify models – it is what is known as a 'non-uniform' model. There are many easy (though not necessarily natural) ways to make it uniform like Turing machines. We will assume this step has been taken if necessary.

(iv) Choose (L, L') (choice statement)

(v) HALT(w) (output the node $[w]$)

Here $w, w' \in \Delta^\star$ and L, L' are natural numbers viewed as labels of instructions. The instructions of the pointer machines are implicitly labeled by the numbers $1, 2, 3, \ldots$ in sequential order. Let us now give the semantics of these instructions. Suppose G is the graph before executing an instruction and G' is the graph after.

(i) If w' accesses the node v in G then after executing this assignment, both w and w' access v in G'. In symbols, $[w]_{G'} = [w']_{G'}$. This is achieved by modifying a single pointer in G.

(ii) We add a new node v to G to form G', and w now accesses v. Furthermore, each pointer from v points back to itself.

(iii) If $[w']_G = [w]_G$ then we branch to the Lth statement; otherwise we execute the next instruction in the normal fashion. The Δ-graph is unchanged, $G = G'$.

(iv) The machine can go to the Lth or the L'th instruction. This corresponds to computational choice.

(v) The machine halts with the output $[w]$. In case we view the machine as an acceptor, a simple convention is that it accepts iff $[w] = [\epsilon]$.

Input and output conventions. Here is a suggested convention. Suppose we want a PM to compute a transformation $t : \{0,1\}^* \to \{0,1\}^*$. We specify three special colors $\alpha, \beta, \nu \in \Delta$. For any string $w \in \Delta^*$, we say that the node $[w]$ represent a 1 or a 0 bit, depending on whether $[w\beta] = [\epsilon]$ or not. Thus, β is interpreted as tbe "bit color". To represent a binary string, we represent a sequence of nodes using ν-pointers as follows. Define the node list

$$L(w) := ([w\nu], [w\nu^2], [w\nu^3], \ldots, [w\nu^m])$$

where each node in this list is distinct from $[\epsilon]$, but $[w\nu^{m+1}] = [\epsilon]$. We allow $L(w)$ to be the infinite list in case $[w\nu^i] \neq [\epsilon]$ for all $i \geq 0$. Thus, the ν-pointers is the "next node indicator". We can then interprete $L(w)$ as a binary string using the bit indicators. Finally, the input to our PM is simply the binary string represented by $L(\alpha)$, and the output of the PM is given $L(w)$ when we encounter the instruction $HALT(w)$. These conventions can be generalized when we need several input strings, $L(\alpha), L(\alpha^2)$, etc. If we need strings over a larger alphabet Σ, we simply re-interpret the β-pointer as follows: the node $[w\beta]$ is viewed as the symbol contained in $[w]$. Since $[w\beta]$ can be an arbitrary node in G, we may want to restrict it to some finite set $\Sigma = \{[\sigma], [\sigma^2], \ldots, [\sigma^n]\}$ where $\sigma \in \Delta$ is yet another special color, and $[\sigma^{n+1}] = [\epsilon]$. In this way, the list $L(w)$ represents an arbitrary string in Σ^*.

Just as the Turing model, Point machines are capable of many interesting variations. First of all, we should probably expand the repetoire of instructions in non-essential ways (e.g., allowing the ability to branch on a general Boolean combination of equality tests). But non-trivial extensions include allowing each node to store an arbitrary integer, and providing basic arithmetic or bitwise operations on a single instruction.

(iii) *Random access machines (RAM).* This is perhaps the closest model to real world computers. Again, each computing agent here consists of a finite set of instructions executed under a finite state control. The difference here is the use of *registers* to store an arbitrarily large integer. The contents of registers can be tested for zero, compared to each other, have arithmetic operations performed on them, all in one step. Registers have unique integer addresses, and registers are accessed in instructions by (i) specifying its address explicitly in the instruction or (ii) indirectly by specifying the address of a register which contains its address.

The reader will appreciate our earlier remark that a computational model is usually characterized by its data structures: Turing machines have fixed topology data structures (called tapes) where each location stores a finite symbol. The data structures in pointer machines have variable topology while random access machines have fixed topology but non-atomic memory locations. Wagner and Strong is an early effort to characterize abstract computational models,

EXERCISES

Exercise 0.1: Let M be a pointer machine that accepts a language $L_0 = L(M)$. Show that there is another pointer machine N with only two colors that accept L_0. ♢

END EXERCISES

2.6 Modes of Computation: Choice and Parallelism

Despite the multiplicity of computational forms, it is possible to isolate a small number of equivalence classes among them: each equivalence class corresponds to basic and intuitive concepts of computation. We identify each equivalence class with a particular *mode of computation.*[23] As we shall see, *within each equivalence class, we can formulate a version of Church's Thesis.*

We distinguish two dimensions of modes. In the first dimension, the computation is either *deterministic* or it may have *choice*. In the second dimension, the computation is either *sequential* or *parallel*. So the four basic classifications of modes are: deterministic-sequential, deterministic-parallel, choice-sequential and choice-parallel. We can have varying degrees of choice (resp. parallelism), with determinism (resp. sequentialism) being the limiting case of *no* choice (resp. *no* parallelism). Therefore, in the above four classes of modes, only the deterministic-sequential mode is uniquely determined: the other three classifications can be further refined. We next indicate characteristics of such modes.

2.6.1 The Fundamental Mode of Computation

The prototypical computation is in the deterministic-sequential mode. Indeed, this mode is usually the only one recognized in theory of computability; of course, the other modes turn out to be no more general in the sense of Church's Thesis. We shall therefore call this the *fundamental mode* of computation. A computation (by a machine M) operating in this mode has the following typical features:

(a) There is a bounded number of internal states, of which one is designated the 'current state'.

(b) There is a finite but unbounded number of memory locations (chosen from an infinite reserve) that are used at any moment. By this we mean that each location at any moment contains a symbol from an alphabet Σ and a location is considered *unused* if the symbol it contains is a certain pre-designated 'blank' symbol. Each location has a bounded number of other locations which are its 'neighbors'.

(c) A bounded number of the memory locations are 'scanned' at any instant.

(d) At each time moment, a unique instruction (which only depends on the symbols scanned in (c) and the current internal state of (a)) is executed. The instruction can change the contents of the scanned locations, the current state, and the set of scanned locations. The new set of scanned locations must be neighbors of the previously scanned ones.

(e) There are fixed conventions for starting, halting, error handling, inputs, and outputs.

It is interesting to compare (a)-(e) with the well-known characterization given in chapter 1 of Rogers [28]. The reader is encouraged to identify our abstract description with the execution of a program on a real computer – most of the above can have a reasonable interpretation (the assumption of an unbounded number of memory locations in (b) is most problematic in this comparison). All the 'bounded numbers' in (a)-(c) are bounded by constants that solely depend on M . At each instant, the *instantaneous description* (ID) of M consists of the current state, the set of scanned locations and the contents of all the used memory locations. Thus we can regard the instructions described in (d) as transforming ID's. If I and I' are ID's such that I is transformed in one step to I', we write $I \vdash I'$, and say I *directly derives* I'. The characterization of M above is deterministic because once the computation is started, the succession of ID's is uniquely determined. The sequence of ID's (possibly infinite) is called a *computation path.* M is sequential because at each time unit, the contents of only a bounded number of locations can be scanned and modified.

Assumption. We have tacitly assumed time and space are discrete above. More generally, *all resources are discrete.* When we discuss non-fundamental modes with more than one computing agent, we will further assume time is *synchronized or global.* This is assumed throughout this book.

To facilitate discussion of the next mode of computation, we elaborate on (e) somewhat. Let us assume that M is a machine for recognizing its inputs. There are two designated internal states designated as *rejecting* and *accepting.* The computation stops iff the machine reaches either of these states. If π is any computation path we define the predicate $\mathrm{ACCEPT}_0(\pi)$ to equal 1 if π is finite and the last ID in π contains the accepting state. Otherwise $\mathrm{ACCEPT}_0(\pi) = 0$. Then M is said to *accept* an input x if $\mathrm{ACCEPT}_0(\pi) = 1$, where π is the computation path of M on input x.

[23] Our notion of modes is partly inspired by the 'computational types' by Hong Jia-wei. The latter concept appears to be a mixture of computational models and computational modes.

2.6.2 Choice Modes

Let us now turn to computational modes that replace 'determinism' with 'choice'. This involves replacing (d) above and (e) above.

(d)* At each time unit, a set (which only depends on the symbol scanned and the current internal state) of next instructions is 'executable'. The machine can 'choose' to execute any one of these instructions. As in (d), each instruction can change contents of locations, current state and set of scanned locations.

(e)* A suitable 'choice aggregation' function is used to determine the output, as illustrated next.

For simplicity, assume that there is only a choice of two next instructions to execute. On input x, we now have a binary *computation tree* $T = T(x)$ whose nodes are ID's defined as follows. The root of T is the initial ID (this is a function of x and M). If I is a node of T, and I_1 and I_2 are the ID's which result from I by executing the two executable instructions, then in T the node I has I_1 and I_2 as its two children. A maximal path in T from the root (terminating in a leaf if the path is finite) corresponds to a computation path.

We want to define the analogue of the ACCEPT_0 function above, and it should be a function of the computation tree T. We describe two main ways to do this. In the first way, we define the predicate ACCEPT_1 by

$$\text{ACCEPT}_1(T) = 1 \text{ iff there is a path } \pi \text{ in } T \text{ such that } \text{ACCEPT}_0(\pi) = 1.$$

The choice mode machine M which accepts its input x iff $\text{ACCEPT}_1(T(x)) = 1$ is said to compute in the *existential-sequential* mode (or more commonly known as the *nondeterministic* mode).

Another way to define acceptance is to assign probabilities to each node in T. Each node I at level k has probability $\Pr(I) = 2^{-k}$, and in particular the root has probability 1. We define the predicate ACCEPT_2 by

$$\text{ACCEPT}_2(T) = 1 \text{ iff } \sum_I \Pr(I) > 1/2,$$

where I ranges over those leaf nodes in T with the 'accept' state. The choice mode machine M which accepts its input x iff $\text{ACCEPT}_2(T(x)) = 1$ is said to compute in the *probabilistic-sequential* mode (or more commonly known as the *probabilistic* mode). We shall see other choice modes in the book.

The existential mode of computation may appear strange at first. Yet it turns out to be an accurate model of the notion of a formal axiomatic system. With suitable conventions, a nondeterministic machine M corresponds exactly to a formal axiomatic system S: the machine M accepts an input x iff x encodes a theorem of S. The reader familiar with grammars in formal language theory will also be able to make connections with nondeterminism. Perhaps the main reason for the importance of the existential-sequential mode is its use in defining a class of problems denoted *NP*, which has great empirical importance. We shall study *NP* in chapter 3.

2.6.3 Parallel Modes

Now we consider the deterministic-parallel modes. Such modes have begun to assume increasing importance with the advent of mass-produced very large scale integrated (VLSI) circuits as well as the rapidly maturing technology of multiprocessor computers. In parallel computation, we will call the entire set of machines which can execute in parallel the *ensemble*. We shall always assume that the ensemble works *synchronously*.[24] Each computational model described in the previous section (Turing machines, pointer machines, random access machines) amounts to specifying individual computing units – this is true in general. Hence to describe ensembles, we need to specify how these units are put together. This comes from four additional design parameters:

(f) (Computing Units) The ensemble is made up of an infinite number of *computing units*, each operating in the fundamental mode (a)-(e). These units may be finite in nature (e.g. finite state automata) or infinite (e.g. general Turing machines).

(g) (Topology) The units have a connectivity (or neighborhood) structure which may be represented by a graph of bounded degree. The graph may be directed or undirected. Units can only communicate with their neighbors. The connectivity may be variable or fixed depending on the model.

[24]The study of asynchronous or distributed ensembles in discrete computations is an emerging subject that seems to give rise to a rather different theory. The emphasis there is the complexity of communication (message passing) between independent agents in the ensemble.

(h) (Sharing and Communication) Each pair of machines which share an edge in the connectivity graph has access to a common bounded set of memory locations: simultaneous reads of this shared memory are allowed, but write conflicts cause the ensemble to halt in error. Note that this shared memory is distinct from the infinite local memory which individual machines have by virtue of (b).

(i) (Activation, Input and Output) There is a designated 'input unit' and another 'output unit'. Input and output conventions are relative to these two units. These two units are not necessarily distinct. Every unit is either 'active' or 'inactive' (quiescent). Initially, all but the input unit is inactive. An active machine can activate its quiescent neighbors. Thus at any moment, a finite but unbounded number of machines are active. The ensemble halts when the output unit enters certain special states (accept or reject).

It is easy to suggest further variations of each of the aspects (f)-(i) above; the literature contains many examples. For instance, unlike the suggestion in (f), not all the individual machines need to be identical. Or again, the sharing need not be local as suggested by (h) if we postulate a global memory accessible by all machines. This is sometimes called the *ultra-computer* model, a term popularized by Jack Schwartz. Another very important practical issue arising from shared memory is the resolution of *reading* and *writing conflicts*. Instead of our stringent condition on disallowing writing conflicts in (h), there are three more liberal approaches:

> We may allow simultaneous writing to a location provided all processors write the same value, or we may have no restrictions on simultaneous writes but say that some arbitrarily chosen processor is successful in writing, or we may say that the smallest numbered processor will be successful.

Another extreme case of (h) is where the machines may have almost no common memory except for a flag associated with each channel or port. The channel or port of a machine may be specific (connected to a particular machine) or non-specific (any machine may try to communicate through that port).

In view of the large number of variations in parallel models, it is all the more surprising that anything significant or unifying can be said about computations in this mode. It is evidence that we are correct in designating essentially one computational mode to all these.

The literature sometimes regards choice modes as 'parallel' modes, since we can think of the various choices of executable instructions as being simultaneously executed by distinct copies of the original machine (reminiscent of 'fork' in UNIX$_{TM}$). We do not use "parallel" in this sense. In true parallelism (unlike choice) processes are able to communicate during their simultaneous computations.[25] Consequently, the aspects (f)-(i) of parallel machines are irrelevant for the choice-sequential mode. More importantly, choice computation is not to be confused with parallelism because acceptance by choice computation is done post-computation (by a fixed evaluation mechanism that is independent of the machine).

The final class of modes, choice-parallel ones, can easily be conceived after the preceding development. For instance, if there are m active computing units at some moment in a choice-parallel computation, each unit having a branching factor of 2, then the entire ensemble has a branching factor of 2^m. We shall not elaborate on this mode of computation.

In conclusion, we see that the computer scientists' answer to the fundamental question (1) of section 5 is quite different from the one the logicians obtained. This is captured in the notion of computational modes. We next turn to the computer scientists' view of question (2) in section 5.

Exercises

Exercise 2.6.3.2: Construct a pointer machine PM that adds two dyadic numbers.

(i) Describe the dyadic algorithm for addition. Is it simpler or more complex than binary addition?

(ii) We need some I/O conventions for pointer machines (PM's). Let Δ be the **alphabet** for the PM, and G be the computation graph of the PM. If $w \in \Delta^*$, let $[w]_G$ denote the node in G represented by w. For instance, the empty string ϵ represents the center $[\epsilon]$. Assume a special symbol $\beta \in \Delta$ that denotes the **bit value** of any node $w \in \Delta^*$ in the following sense: if $[w.\beta] = [\epsilon]$, then we say that the node $[w]$ stores the 1 bit, else it stores tbe 0 bit. How do we represent a list of nodes? Assume a special "next symbol" $\nu \in \Delta$ such that $[w.\nu]$ represents the next node after $[w]$, provided $[w, \nu] \neq [\epsilon]$; otherwise we say there is no next node. Define the list headed by w to be the sequence

$$L(w) := ([w], [w\nu], [w\nu^2], \ldots, [w\nu^{n-1}]),$$

[25] The reader will appreciate this distinction if she attempts to show 'in an obvious manner' that the class *NP* (to be defined in chapter 2) is closed under complementation and faces up to the inability for different branches of a nondeterministic computation to communicate.

terminated by the first n such that $[w\nu^n] = [\epsilon]$. The list $L(w)$ may be infinite (in the case of a loop). The binary string represented by w is simply the sequence of bit string stored in $L(w)$. Input are indicated by two linear lists $L(\alpha)$ and $L(\alpha^2)$, where $\alpha \in \Delta$ is a special "input symbol". When the computation ends, we want the output string to stored in the list $L(\alpha^3)$.

□

Exercise 2.6.3.3: Describe a parallel model of Turing machines, and also a parallel model of pointer machines. Then prove a simulation theorem, showing that each model can simulate the other model within polynomial parallel time. □

End Exercises

2.7 Tractability and some Computational Theses

It has been observed in the computing milieu that many problems such as the *Perfect Matching Problem* in graph theory have polynomial time algorithms (even though the initial mathematical characterization of 'perfect matchings' only suggested exponential algorithms). In contrast, certain computational problems such as the *Traveling Salesman Problem* (TSP) defy persistent attempts to design efficient algorithms for them. All known algorithms and suggested improvements for these algorithms still have, in the worst case, super-polynomial[26] running time. Typically, this means an exponential running time such as 2^n. The obvious question is whether a problem such as TSP is intrinsically super-polynomial. The gap between an exponentially growing function (say) 2^n and a polynomial function (say) n^2 is quite staggering: on a computer executing 10^6 instructions per second, an input of size 200 would require 2^{200} (resp., 200^2) steps or more than 10^{46} years (resp., less than a second) of CPU time. Of course, the difference increases dramatically with increasing n. This unbridgeable gap between polynomial functions and exponential functions translates into a clear distinction between problems with polynomial complexity and those with super-polynomial complexity. Super-polynomial problems are 'practically uncomputable' (except for small values of n) even though they may be computable in the sense of computability theory.[27] The phenomenon appears very fundamental and researchers have coined the term 'infeasible' or 'intractable' to describe these difficult problems:

(H) A problem is *tractable* if it has polynomial complexity; otherwise the problem is *intractable*.

Cobham [5] and Edmonds[8] were among the first harbingers of the tractable-intractable dichotomy.[28] Let us note an obvious criticism of (H). While we may agree that super-polynomial problems are intractable, we may not want to admit all polynomial problems as tractable: a problem with complexity of n^{37} hardly deserves to be called tractable! In practice, we seldom encounter such high degree polynomials. On the other hand, it is difficult to decide on theoretical grounds when polynomials become intractable, and hence (H).

Note that (H) is really a general principle in the sense that we have not restricted it to any particular model, mode of computation or computational resource. For instance, if we consider the Turing model operating in the deterministic mode, and use space as our resource, then this principle tells us that the tractable problems are what is normally called *PSPACE*. (If we operate in the nondeterministic mode, a well-known result of Savitch which we will encounter in chapter 2 says that precisely the same class of problems are tractable.) One reason why (H) is attractive is that it is a very robust concept: tractability is easily seen to be invariant under minor modifications of the computational models. A far-reaching generalization of this observation is the tractability thesis below which basically says that the concept is robust enough to withstand comparisons across totally different computational models, *provided we compare these models using the same computational mode and using corresponding resources.* There is a problem with the last assertion: what do we mean by 'corresponding resources'? For this to be non-vacuous, we must be sure that there are certain computational resources which are common to every conceivable model of computation. In the actual models that have been studied in the literature, this is not problematic. In any case, we postulate three fundamental resources in every computational model:

[26] A complexity function f is *super-polynomial* if $f(n)$ is not $O(n^k)$ for any fixed value of k. A complexity function f is said to be (at most) *exponential* if $f(n) = O(1)^n$, *double exponential* if $f(n) = 2^{O(1)^n}$, etc.

[27] This remark must be balanced against the fact that very often, naturally occurring instances of an intractable problem can be solved by efficient special methods.

[28] Edmonds called an algorithm 'good' if it runs in polynomial time in the fundamental mode. In [9] he informally introduced the class *NP* by describing such problems as having 'good characterization'.

(I) For every computational model in a given computational mode, there is a natural notion of *time, space* and *reversal* resources.

Time and space have intuitive meaning; reversal on Turing machines is intuitively clear but it can be formulated for all other models as well. For postulate (I) to be meaningful, there ought to be axioms that these measures obey (otherwise, how do we know if 'time' on one model-mode is really to be equated with 'time' on another model-mode?). For example, we expect that "space is at least as powerful as time". We said that the Boolean circuit model[29] does not naturally yield a notion of sequential mode. One possible approach is to use what is known as a 'straight-line program' representation of a circuit (that is, we list the gates in a linear sequence, where each gate must have its inputs coming from earlier gates or from the external inputs). Then (sequential) time is simply the length of this straight-line program. More generally, Hong has suggested that sequential time is simply the total number of individual actions that take place inside a computation (although the circuit example blurs the distinction between the computing device and the computation itself). We leave such issues as matters of further research. In any case, when we say 'resources' below, they could refer to composite resources such as simultaneous time and space.

We are now ready to assert what may be called the "polynomial analogue of Church's Thesis".

(J) (*Tractability Thesis*) For each computational mode and each computational resource, the notion of tractability is invariant over all models of computation.

For comparison, we also formulate a version of Church's Thesis here.[30]

(K) (*Church's Thesis*) The notion of computability is invariant over all models of computation.

When restricted to the time resource in the fundamental mode, thesis (J) is sometimes called Cobham's thesis after Cobham who made the first observations about its invariance properties. Although the thesis (J) has not been tested in its full generality (especially, in some of the newer modes), it is useful as a working hypothesis or as a research program.

In all our theses, there is an implicit restriction to 'reasonable' models of computation. Or at least, we intend to call a model 'unreasonable' if it violates such principles. Thus we might appeal to the tractability theses to reject certain models of computations as unreasonable. For instance, it is clear that if we severely restrict our computational model we could violate (J). Likewise, it is not hard to imagine a model too powerful to respect the tractability thesis. This is not surprising, but simply points out that our theses need *a priori* conceptual analysis in order to establish them firmly. Subject to this qualification, and provided that we accept (H) and (I), we see that (J) has mathematical content in the sense that we may verify it for each proposed model of computation. On the other hand, (J) is not a mathematical theorem, and (as Emil Post remarked, concerning Church's thesis) "it requires continual verification". In summary, we say that the thesis serves both a normative as well as descriptive function.[31]

Thesis (J) is the consequence of another more general *polynomial smearing phenomenon* that says, with respect to a fixed resource, all reasonable computational models computing in a fixed mode are equally powerful up to a polynomial factor. More precisely: say that a computational model M *polynomially simulates* another model M' *in resource* R if for every machine A' in M' which uses $f(n)$ units of resource R, there is a machine A in M which uses $O((f(n))^k)$ units of R and which solves the same problem as A'. The constant k here may depend on A', although we find that in all known models, k may be uniformly chosen (e.g. $k = 2$). We say M and M' are *polynomially equivalent* in resource R if they polynomially simulate each other.

(L) (*Polynomial Simulation Thesis*) Within each computational mode, and with respect to a given computation resource, all models of computation are polynomially equivalent.

Clearly (H) and (L) imply (J). Hong[17] has provided strong evidence for this thesis (his so-called 'similarity theorems'). We now state two conjectures and one more thesis. These relate across computational modes and resources:

[29] suitably made 'uniform', see chapter 10.

[30] As Albert Meyer points out to us, it is possible to interpret Church's thesis in a much wider philosophical sense. See (for example) [16] for such discussions. The mathematical content of this wider thesis is unclear and our formulation here is, for our purposes, the generally accepted mathematical interpretation of the thesis.

[31] The same can be said of Church's thesis – we would use its normative role to dismiss someone who insists that finite automata be admitted as general computing devices. Taking Lakatos' critique of mathematical truth as a clue, we might modify Post's position and say that these theses require continual refinement.

(M) (*The $P \neq NP$ Conjecture*) With respect to time measure, tractability in the fundamental mode and tractability in the nondeterministic mode are different.

(N) (*Parallel Computation Thesis*) Time measure in the deterministic-parallel mode is polynomially related to space measure in the fundamental mode.

(O) (*Duality Conjecture*) The resources of Space and Reversal are duals (see text below) of each other when the simultaneous bounds on space and reversals are polynomially related.

Conjecture (M) is perhaps the most famous open problem in the subject, and we shall return to it in chapter three. Thesis (N) (proposed by Pratt and Stockmeyer [25] and Goldschlager [11]) equates the tractable problems in the deterministic-parallel mode under time resource bounds with the tractable problems for the fundamental mode under space resource bounds. See also [24]. The duality conjecture was formulated by Jia-Wei Hong. It is more technical and involves simultaneously bounding the space and reversal resources. Roughly, it says that if a complexity class is defined by Turing machines bounded by $O(f)$ space and $O(g)$ reversal simultaneously then the same class is described by $O(f)$ reversal and $O(g)$ space simultaneous bounds.

In this book we shall provide evidence for the above theses. We note the tremendous unifying power of (H)-(O): First, (H) gives us a 'world-view' by pointing out a fundamental distinction (just as the computable-uncomputable distinction in computability theory); this in turn guides the basic directions of research. Second, these theses lead us to the general principle (or metathesis):

The truly fundamental phenomena of Complexity Theory are invariant across computational models.

Third, the theses (M), (N) and (O) suggest that the concepts of modes and resources are not arbitrary. Rather, their interrelation reflects an internal structure of the concept of computation that awaits discovery. This could be the most exciting direction in the next stage of Complexity Theory. It is fair to add that the mode-resource-model structure we have imposed on the corpus of the known complexity substratum may not bear careful scrutiny, or at least requires new clarifications, especially as new computational concepts proliferate.[32]

EXERCISES

Exercise 2.7.0.4: Verify the Polynomial Simulation Thesis (L), for Turing machines and Pointer machines. We want to verify the following cases, depending on the mode μ and resource ρ: we want you to verify:
(i) (μ, ρ) = (Fundamental Mode, Time). More precisely, $DTIME(n^{O(1)}) = DTIME_{PM}(n^{O(1)})$. The latter class is the class of languages accepted by PM's in deterministic polynomial time.
(ii) (μ, ρ) = (Nondeterministic Mode, Space). That is $NSPACE(n^{O(1)}) = NSPACE_{PM}(n^{O(1)})$.
(iii) (μ, ρ)= (Fundamental Mode, Reversal). That is $DREVERSAL(n^{O(1)}) = DREVERSAL_{PM}(n^{O(1)})$. □

Exercise 2.7.0.5: □

END EXERCISES

2.8 What is Complexity Theory?

We have given an impressionistic tour of Complexity theory in this chapter and outlined some of the methodological framework. This framework is plausible but not necessarily convincing. In any case, we sometimes refer to these assumptions (if only to hint at alternative assumptions) as *Standard Complexity Theory*. We are ready to round up with three perspectives of what the subject is about, plus a very brief historical perspective. By *complexity classes* below we mean any set of languages. In practice, these classes are defined by some well-defined and general mechanism which admits some concept of complexity.

a) The study of computational complexity is usually taken in a wider sense than that taken by this book. In particular, it usually incorporates a very large and important literature on the analysis of algorithms and data-structures. Complexity Theory there is sometimes called *concrete complexity* (dealing in 'concrete' as opposed to 'abstract' problems). However, complexity in concrete complexity is often more a property of the particular *algorithm* being analyzed, rather than a property of the *problem* that the algorithm solves. In contrast:

[32] Such proliferation seems as inevitable as in subatomic Physics. The challenge is to provide an adequate unified field theory of complexity theory as the world of complexity unfolds. It is also useful to remember that although Newtonion Physics is superseded, it is hardly extinct, thanks to its compact embodiment of non-extremal physical phenomena.

Complexity Theory is concerned with the intrinsic complexity of problems.

Specifically, we are interested in classifying problems according to their intrinsic complexity.

Let us probe deeper this intuitive concept of intrinsic complexity. The notion turns out to be rather subtle: one phenomenon that may arise is that there may be no fastest algorithm – for each algorithm there exists a faster one. We will see examples of such 'speed-up' phenomena. If we consider two or more computational resources, there may be inherent tradeoffs between the complexities with respect to these resources. In the face of these possibilities, we can still attempt to classify the complexity of a problem P relative to a given *hierarchy* of complexity classes: more precisely, if

$$K_1 \subseteq K_2 \subseteq K_3 \subseteq \cdots \tag{8}$$

is an non-decreasing (not necessarily strict) sequence of complexity classes, we say that P is in the ith level of this hierarchy if $P \in K_i - K_{i-1}$ ($i = 1, 2, \ldots$ and $K_0 = \emptyset$). Note that this abstract method of classification is a really profound departure from the original idea of comparing growth rates of complexity functions. The hierarchy $\{K_1, K_2, \ldots\}$ that replaces complexity functions need not have any clear connection to complexity functions. Let us call any hierarchy (8) used for this purpose a *complexity ruler*. Once we agree to classify problems relative to complexity rulers, some very fruitful directions in complexity theory become possible. Such classifications are studied in chapter 9.

One complexity ruler that is used commonly is the *canonical ruler* formed from the canonical list (see chapter 2, section 3), but omitting the class *PLOG*:

$$\begin{aligned} DLOG \subseteq NLOG \subseteq P &\subseteq NP \subseteq PSPACE \\ &\subseteq DEXPTIME \subseteq NEXPTIME \subseteq EXPSPACE \end{aligned}$$

One of the first things a complexity theorist does on encountering a computational problem is to find out where it fits in this canonical ruler. Another ruler is this: a logarithmico-exponential function (for short, L-function) $f(x)$ is a real function that is defined for values $x \geq x_0$ (for some x_0 depending on f) and is either the identity or constant functions, or else obtained as a finite composition with the functions

$$A(x), \quad \log(x), \quad e^x$$

where $A(x)$ denotes a real (branch of an)[33] algebraical function. A classical result of Hardy[13] says that if f, g are L-functions then either $f = \Theta(g)$, or f dominates g, or g dominates f. Hence the family R_0 consisting of the classes $DTIME(\Theta(f))$ where f is an L-function forms a complexity ruler. (Of course, the ruler is doubly infinite in both directions, dense and all that.) In practice, when complexity theorists discuss the time complexity of concrete problems (like matching, multiplication, etc), they are implicitly using the ruler R_0. For instance, although multiplication has time complexity $t(n)$ that satisfies the inequalities $n \leq t(n) = O(n \log n \log\log n)$, there may in fact be no single function $f(n) \in R_0$ such that $t(n) = \Theta(f(n))$; nevertheless it makes sense to talk about the sharpest upper (or lower) bound on multiplication relative to the ruler R_0.

Actually, the use of rulers to measure complexity can be further generalized by introducing the concept of reducibility among languages (chapter 4). The inclusion relation can be replaced by the reducibility relation.

b) The discussion of encodings and models of computations leads us to conclude that our theory, necessarily becomes quite distorted for problems with low-level complexity (say $o(n^2)$ time). On the other hand, the concept of computational modes arises because we are also not very concerned with high-level complexity (otherwise they all become equivalent by Church's thesis (K)). In short:

Complexity Theory is concerned with medium-level complexity.

Here (if one has to be so precise) one may identify as medium-level those problems which are elementary or primitive recursive.[34] It is this concern for medium level complexity that makes the fundamental dichotomy (H) a meaningful one. Indeed a central concern of the theory is to classify problems as tractable or otherwise (with respect to any mode or resource). Our interest in medium level complexity partly justifies some of our assumptions. For instance, even though assumption (B) captures recognition problems only, many functional or optimization problems are polynomially equivalent to suitable recognition problems (chapter 3, section 2); studying complexity up to polynomial equivalence seems justifiable for medium level complexity.

c) Finally, we contend that

[33] For instance, the polynomial $p(x) = x^2$ defines two branches corresponding to the positive and negative square-root functions. Generally speaking, a polynomial $p(x)$ of degree d defines d functions corresponding to the d roots of $p(x)$ as x varies. These functions are called 'branches' of $p(x)$.

[34] 'Elementary recursive' and 'primitive recursive' are technical terms in recursive function theory. Section 6 of chapter 5 has a definition of elementary recursiveness.

Complexity Theory is essentially a theory of classes of problems

as opposed to a theory of *individual* problems. No doubt, the complexity of certain important problems (graph isomorphism, multiplication, primality testing, etc) in particular models of computation has abiding interest. But the exact complexity of an individual problem is to some degree an artifact of the details of the computational model. It is only when we examine an entire class of (suitably chosen) problems that we manage to have truly invariant properties. Moreover, even when we study the complexities of individual problems, we usually aim at placing the problem in a well-known complexity class or in some level of a hierarchy. Another case we may adduce to support our claim is the existence of the tractability thesis: the class of interest there are those problems with polynomial complexity. Other examples can be provided (see also the Preface). Indeed, we regard this emphasis on complexity classes as the chief distinguishing mark of the Complexity Theory represented by this book, and hence the book title.

d) The preceding three perspectives on Complexity Theory are from an intrinsic or genetic standpoint. A historical perspective would also be greatly rewarding especially for the student of the history and philosophy of science. Unfortunately, only a few remarks can be offered here. The theory of computational complexity (which we shorten to Complexity Theory in this book) took its name from the seminal paper of Hartmanis and Stearns [15] in 1965. The earlier papers of Rabin [27] and Blum [2] are generally counted among the most influential works associated with the founding of this subject. We refer the interested reader to the volume [6] for further references as well as personal accounts from that seminal period. The work of Cook and Karp in the early 1970s profoundly changed the field by raising certain important questions; these questions though still open, play a paradigmic role that has influenced directly or indirectly much of subsequent research programs.

Exercises

[2.1] Let $R(x)$ be a partial real predicate, and $D \subseteq \mathbb{R}$. Show the following
(i) $\neg(\forall x \in D)[R(x)]$ iff $(\exists x \in D)[\neg R(x)]$.
(ii) $\neg(\exists x \in D)[R(x)]$ iff $(\forall x \in D)[\neg R(x)]$.
(iii) $\neg(\text{EV}x \in D)[R(x)]$ iff $(\text{IO}x \in D)[\neg R(x)]$.
(iv) $\neg(\text{IO}x \in D)[R(x)]$ iff $(\text{EV}x \in D)[\neg R(x)]$.

[2.2] Verify the assertion in the text: $g \neq o(f)$ means that there is some $c > 0$ such that for infinitely many n, $g(n) \geq cf(n)$.

[2.3] What is the distinction between $f(n) = \Omega(g(n))$ and $g(n) \neq O(f(n))$, and between $f(n) = \omega(g(n))$ and $g(n) \neq o(f(n))$?

[2.4] (i) A useful notation is $f(n) \sim g(n)$, defined here to mean

$$f(n) = (1 \pm o(1))g(n).$$

Show that if $f(n) \sim g(n)$ then $g(n) \sim f(n)$.
(ii) Suppose $x2^x = n$ holds for all n. We want to solve for $x = x(n)$. Show that $x(n) = \log n - \log\log n + O(\frac{\log\log n}{\log n})$
(iii) Conclude that $\log n - x(n) \sim \log\log n$.

[2.5] ** (i) Are there useful applications of mixed asymptotic expressions (big-Oh, big-omega, small-oh, etc)? Does it make sense to say that the running time of a machine is $\Omega(n^{O(n)})$?
(ii) More generally, work out a calculus of these mixed notations.

[2.6] Extend the asymptotic notations to multi-parameter complexity functions.

[2.7] Let $g(n) = \exp\exp(\lfloor\log_2\log_2 n\rfloor))$ where $\exp(m) = 2^m$. Clearly $g(n) = \Omega(n^\alpha)$ and $g(n) = O(n^\beta)$ for some constants $\alpha, \beta > 0$. Give the values for α and β that are sharpest in the sense that for any $\epsilon > 0$, both the Ω- and the O-bound would not be true if $\alpha + \epsilon$ and $\beta - \epsilon$ (respectively) were used.

[2.8] Show:
(i) $H_n = \Theta(\log n)$ where H_n (the harmonic series) is defined as $\sum_{i=1}^n \frac{1}{i}$.
(ii) $\sum_{i=1}^{\log n} 2^i \log\left(\frac{n}{2^i}\right) = \Theta(n)$. [Obtain the best constants you can in the upper and lower bounds on the sums in (i) and (ii).]
(iii) $(x + a)\log(x + b) = x\log x + a\log x + b + o(1)$.

[2.9] Describe an $O(n^2)$ algorithm for converting k-ary numbers to k'-ary numbers for $k, k' > 1$. Can you do better than $O(n^2)$? **Hint:** Regard the k-ary number $a_n a_{n-1} \cdots a_0$ $(a_i = 0, \ldots, k-1)$ as the polynomial $p(x) = \sum_{i=0}^{n} a_i x^i$ and evaluate $p(x)$ at $x = k$ in k'-ary notation.

[2.10] Demonstrate the inequalities (2-4) in Sections 4.1, 4.2. Make explicit any reasonable assumptions about the computation model and complexity measure.

[2.11] Referring to section 4.2, show that for $k > 1$,

$$f_k(n) = O(f_1(\Theta(k^n))) + O(k^n).$$

$$f_1(n) = O(f_k(\Theta(\log n))) + O(n \log n).$$

What can you say about the relative growth rates of f_1 and f_k if f_k is exponential?

[2.12] There are three possible responses to the following statements: True, False or Unknown. "Unknown" reflects the uncertainty that only assume the results in chapter 2. If your answer is True or False, you must give brief reasons in order to obtain full credit. For Unknown, discuss relevant results.
(a) T/F/P: $NSPACE(n) \subseteq DTIME(O(1)^n)$.
(b) T/F/P: $NTIME(n) \subseteq DTIME(2^n)$.
(c) T/F/P: $NREVERSAL(n) \subseteq DREVERSAL(O(1)^n)$.

[2.13]
(i) List differences and similarities between k-adic from k-ary notations for $k \geq 1$.
(ii) Describe the dyadic addition algorithm.
(iii) Describe the dyadic multiplication algorithm.

[2.14] * Give a systematic treatment of the possible variations of parallel mode of computation. Extend this to the parallel-choice modes of computation.

[2.15] Verify that multiplication is indeed linear time under the suggested prime number encoding of numbers in Section 4.2. What is the complexity of addition?

[2.16] Recall the recognition problem in Section 2 for triples (x, y, z) of binary numbers such that $xy = z$. We modify this so that the relation is now $xy \geq z$. Show that the ordinary multiplication problem can be reduced to this one. If this recognition problem can be solved in $T(n)$ time on inputs of length n, give an upper bound for the multiplication problem.

[2.17] Construct a pointer machine to add two numbers. Assume that numbers are encoded in binary, using input/output conventions described in the text.

[2.18] Let M be a deterministic Turing machine that accepts in time-space (t, s). Show how a pointer machine can simulate M in $O(t, s)$ time and space.

[2.19] Let P be a deterministic Pointer machine that accepts in time-space (t, s). Fix a convention for input, and assume that P cannot change this input, to make it compatible with our convention for Turing machines. Show how a Turing machine can simulate M in $O(t \log s, s \log s)$ time and space.

[2.20] Verify the Polynomial Simulation Thesis (J) for Pointer Machines and Turing Machines. In other words, for each mode μ and resource ρ, show that the class of languages accepted by μ-Turing machines using a polynomial amount of ρ is the same as the class of languages accepted by μ-Pointer machines using a polynomial amount of ρ. Specifically, consider
(i) $DTIME(n^{O(1)}) = DTIME_{PM}(n^{O(1)})$, where the subscript "$PM$" refers to Pointer machines.
(ii) $NSPACE(n^{O(1)}) = NSPACE_{PM}(n^{O(1)})$.
(iii) $DREVERSAL(n^{O(1)}) = DREVERSALPM(n^{O(1)})$. In this case, you need to define the reversals resource for Pointer machines. Here is an outline. Let a configuration of a PM be a pair (i, G) where i is the label of the PM instruction that is about to be executed, and G is the Δ-structure for the PM. We define a computation sequence $C_0 \vdash C_1 \vdash \cdots \vdash C_m$ in the expected way. Again, computation sequence is called a "phase" if each of its transitions $C_i \vdash C_{i+1}$ does not depend on on the outcome of actions in the same phase. The reversals of a computation sequence is the minimum number of phases that it can be subdivided into.

[2.21] ** Axiomatize or give abstract characterizations of the computational models in the literature. Some work in the past have done precisely this for some models. Describe the various computational modes for in each of these models. With our new distinction between modes and models, this task requires new care.

[2.22] ** Give proper foundations for the computational theses outlined in section 7. This presumes a suitable solution of the preceding question.

[2.23] * There are many ways to represent a real algebraic number α (i.e. α is a root of some polynomial $p(x)$ with integer coefficients). (i) We can exploit the fact that real numbers are ordered and represent α as $\langle p(x), i\rangle$ if α the ith real root of $p(x)$. (ii) We can isolate the root of interest by an interval, so use the representaion $\langle p(x), I\rangle$ where I is an interval with rational endpoints containing α but no other roots of $p(x)$. (iii) We can exploit Thom's lemma that asserts that if we assign a sign $\sigma(D_i(p)) \in \{-1, 0, +1\})$ to the ith derivative $D_i(p)$ of $p(x)$ for each $i \geq 0$ then the set of real numbers a such $D_i(p)(a)$ has sign $\sigma(D_i(p))$ forms an interval (possibly empty) of the real line. In particular, if we choose $\sigma(D_0(p)) = \sigma(p) = 0$ then either the set of such a is empty or else we have identified a unique root. Hence α can be represented by $\langle p(x), \sigma\rangle$ for a suitable σ. Compare the three representations using criteria (D).

[2.24] * Consider the problem of representing multivariate polynomials whose coefficients are integers. What are the various methods of encoding? Discuss the consequences of these on the complexity of various algorithms on polynomials (a) GCD, (b) factorization. In this connection, the reader is referred to the work of Erich Kaltofen on the straight-line program representation of polynomials (e.g., [20]).

[2.25] ** Investigate the theory of computational problems (see footnote 1). In this regard, compare the 'theory of first order problems' for linear orderings in [36] (which is taken from [37]).

[2.26] ** Show that some mathematical domains D (such as unlabeled graphs) have an inherently high 'representaional complexity' in the sense that fore any representation $r : \Sigma^* \to D$, either the r-parsing or r-isomorphism problem is complex (compare the next exercise).

[2.27] ** Investigate the effects of encodings on complexity. On a methodological level, we could say that the issue of encodings is outside Complexity Theory, that Complexity Theory begins only after a problem is suitably encoded. E.g., Megiddo [22] points out that the the known polynomial time algorithms for the linear programming problem (originally shown by Khacian) assumes a certain encoding for the problem. Yet it may be possible to bring this encoding into our theory if we resort to meta-mathematical tools.

Bibliography

[1] A. Aho, J. Hopcroft, and J. Ullman. *Data structures and algorithms*. Addison-Wesley, 1983.

[2] Manuel Blum. A machine-independent theory of the complexity of recursive functions. *Journal of the ACM*, 14(2):322–336, 1967.

[3] Giles Brassard. Crusade for a better notation. *SIGACT News*, 17:1:60–64, 1985.

[4] Alonzo Church. An unsolvable problem of elementary number theory. *Amer. J. Math. 58*, pages 345–363, 1936.

[5] Alan Cobham. The intrinsic computational difficulty of functions. *Proc. 1964 International Congress for Logic, Methodology and Philosophy of Science*, pages 24–30, 1964.

[6] The Computer Society of the IEEE. *Proceedings, Third Annual Conference on Structure in Complexity Theory*. Computer Society Press of the IEEE, June 14-17, 1988. (held at Georgetown University, Washington, D.C.).

[7] T. H. Corman, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. The MIT Press and McGraw-Hill Book Company, Cambridge, Massachusetts and New York, 1990.

[8] Jack Edmonds. Paths, trees, and flowers. *Canadian J. Math.*, 17:449–467, 1967.

[9] Jack Edmonds. Matroid partition. In G.B. Dantzig and Jr. A.F. Veinott, editors, *Mathematics of the decision sciences*. Amer. Math. Soc., Providence, R.I., 1968.

[10] Kurt Gödel. Uber formal unentscheidbare Satze der Principia Mathematica und verwandter System I. *Monatschefte Math. Phys.*, 38:173–98, 1931. (English Translation in [Dav65]).

[11] L. M. Goldschlager. A universal interconnection pattern for parallel computers. *Journal of the ACM*, 29:1073–1087, 1982.

[12] Yuri Gurevich. What does $O(n)$ mean? *SIGACT News*, 17:4:61–63, 1986.

[13] G. H. Hardy. Properties of logarithmico-exponential functions. *Proc. London Math. Soc.*, 2:54–90, 1912.

[14] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, London, 1938.

[15] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Trans. Amer. Math. Soc.*, 117:285–306, 1965.

[16] Douglas R. Hofstadter. *Godel, Escher, Bach: an eternal golden braid*. Vantage, New York, N.Y., 1980.

[17] Jia-wei Hong. *Computation: Computability, Similarity and Duality*. Research notices in theoretical Computer Science. Pitman Publishing Ltd., London, 1986. (available from John Wiley & Sons, New York).

[18] David S. Johnson. The *NP*-completeness column: an ongoing guide. the many faces of polynomial time. *Journal of the ACM*, 8:285–303, 1987.

[19] S. Jukna. Succinct data representation and the complexity of computations. In L. Lovász & E. Szemerédi, editor, *Theory of algorithms*, volume 44, pages 271–282. Elsevier Science Pub. Co., 1985.

[20] E. Kaltofen. Greatest common divisors of polynomials given by straight-line programs. *Journal of the ACM*, 35:231–264, 1988.

[21] Donald E. Knuth. Big omicron and big omega and big theta. *SIGACT News*, Vol.8, No.2:18–24, 1976.

[22] N. Megiddo. Is binary encoding appropriate for the problem-language relationship? *Theor. Computer Science*, 19:337–341, 1982.

[23] Marvin Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1967.

[24] Ian Parberry. Parallel speedup of sequential machines: a defense of the parallel computation thesis. *SIGACT News*, 18:1:54–67, 1986.

[25] Vaughn R. Pratt and Larry Stockmeyer. A characterization of the power of vector machines. *Journal of Computers and Systems Science*, 12:198–221, 1976.

[26] W. V. Quine. Concatenation as a basis for arithmetic. *Journal of Symbolic Logic*, 11 or 17?:105–114, 1946 or 1952?

[27] Michael Rabin. Degree of difficulty of computing a function. Technical Report Tech. Report 2, Hebrew Univ., 1960.

[28] Hartley Jr. Rogers. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, 1967.

[29] A. Schönhage. Storage modification machines. *SIAM J. Computing*, 9(3):490–508, 1980.

[30] A. Schönhage and V. Strassen. Schnelle Multiplikation Grosser Zahlen. *Computing*, Vol. 7:281–292, 1971.

[31] Richard Schroeppel. A two counter machine cannot calculate 2^n. Technical Report AI Memo 257, M.I.T., 1973.

[32] Alan M. Turing. On computable number, with an application to the Entscheidungs problem. *Proc. London Math. Soc., Ser.2-42*, pages 230–265, 1936.

[33] R. Verbeek and K. Weihrauch. Data representation and computational complexity. *Theor. Computer Science*, 7:99–116, 1978.

[34] P. M. B. Vitányi and L. Meertens. Big omega versus the wild functions. *SIGACT News*, 16(4):56–59, 1985.

[35] Frances F. Yao. Computation by 2-counter machines. Unpublished term paper, M.I.T., 1973.

[36] Chee-Keng Yap. Space-time tradeoffs and first order problems in a model of programs. *12th ACM Symposium on Theory of Computing*, pages 318–325, 1980.

[37] Chee-Keng Yap. *Three studies on computational problems*. PhD thesis, Yale University, 1980. PhD Thesis, Computer Science Department.

Appendix A

Basic Vocabulary

NOTE: THIS HAS BEEN MOVED TO CHAPTER 0!!! REMOVE THIS SECTION.

This appendix establishes some general terminology and notation used throughout the book.

We assume the usual set theoretic notations. We write $S \subseteq S'$ for set inclusion, and $S \subset S'$ for proper set inclusion. We let $|S|$ denote the cardinality of a set S.

Instead of the equality sign, we use ':=' to indicate a definitional equality. For example, we write '$S := \ldots$' if '$\ldots$' serves as a definition of the set S.

Let D, R be sets. By a *partial function* f with *domain* D and *range* R we mean a rule that associates certain elements x of D with an element $f(x) \in R$. So D and R are assumed to be specified when f is given. The function is said to be *undefined* at x if the rule does not associate x to any element of R, and we denote this by $f(x) = \uparrow$; otherwise, we say f is *defined* at x and write $f(x) = \downarrow$. If f is defined at all elements of its domain, we say f is *total*. Composition of functions is denoted $f \circ g$ where $(f \circ g)(x) := f(g(x))$.

The set of *integers* and the set of *real numbers* are denoted $\mathcal{Z}$ and $\mathcal{R}$, respectively. The set of *extended reals* refers to $\mathcal{R} \cup \{\infty\}$. The set $\mathcal{N}$ of *natural numbers* refers to the non-negative integers, $\mathcal{N} = \{0, 1, 2, \ldots\}$. The *floor function* takes any real number x to the largest integer $\lfloor x \rfloor$ no larger than x. The *ceiling function* takes any real number x to the smallest integer $\lceil x \rceil$ no smaller than x. So $\lfloor 0.5 \rfloor = 0 = \lceil -0.9 \rceil$.

An *alphabet* Σ is a non-empty finite set of markings. We call each element of Σ a *letter* or *symbol*. The set of finite sequences of letters (called *words* or *strings*) over an alphabet Σ is denoted Σ^*. The *empty word* is denoted ϵ. A *language* is a pair (Σ, L) where Σ is an alphabet and L is a subset of Σ^*. Usually, the alphabet Σ of a language (Σ, L) is either understood or immaterial, and we shall loosely refer to L as the language. The language L is *trivial* if $L = \Sigma^*$ or $L = \emptyset$ where $\emptyset$ denotes the empty set. The *complement* of L, denoted co-L, is the language $(\Sigma, \Sigma^* - L)$. Note that the complementation operation is an instance of a language operator whose proper definition requires that the alphabet of the language be explicitly given. A collection K of languages is usually called a *class*. For any class K, co-K is defined to be the class $\{co\text{-}L : L \in K\}$.

The concatenation of two words v and w is written $v \cdot w$ or simply vw. This notation extends naturally to sets of words: if S and T are sets of words then $S \cdot T := \{vw : v \in S, w \in T\}$. The length of a word w is denoted $|w|$. The unique word in Σ^* of length zero is denoted ϵ, and Σ^+ is defined as $\Sigma^* - \{\epsilon\}$. (We may assume that the word ϵ is common to all Σ^*.) For any non-negative integer n and word w, we let w^n denote the n-fold self-concatenation of w. More precisely, $w^0 := \epsilon$ and for $n \geq 1$, $w^n := w \cdot w^{n-1}$.

A *language operator* ω is a partial d-ary function, $d \geq 0$, taking a d-tuple $(L_1, \ldots, L_d)$ of languages to a language $\omega(L_1, \ldots, L_d)$. Here we assume that the L_i have a common alphabet which is also the alphabet of $\omega(L_1, \ldots, L_d)$. Some simple language operators are now introduced. Other important operators, to be introduced in the course of this book, are usually defined using machines; this is in contrast with the following set-theoretic definitions.

Let $(\Sigma, L), (\Sigma, L')$ be languages. The *complement* of L, denoted co-L, is the language $\Sigma^* - L$. The *union*, *intersection* and *difference* of (Σ, L) and (Σ', L') are (resp.) the languages $(\Sigma \cup \Sigma', L \cup L')$, $(\Sigma \cup \Sigma', L \cap L')$ and $(\Sigma, L - L')$. (The preceding are the *Boolean operators*.) The *concatenation* of L and L', denoted $L \cdot L'$, is the language $\{ww' : w \in L, w' \in L'\}$ over the alphabet $\Sigma \cup \Sigma'$. For any non-negative integer n, we define the language L^n inductively as follows: L^0 consists of just the empty word. $L^{n+1} := L^n \cdot L$. The *Kleene-star* of L, denoted L^*, is the language $\bigcup_{n \geq 0} L^n$. (Note that the Σ^* and L^* notations are compatible.) A related notation is L^+ defined to be $L \cdot L^*$. The *reverse* of L, denoted L^R, is $\{w^R : w \in L\}$ where w^R denotes the reverse of w.

A language L is said to be *finite* (resp. *co-finite*) if $|L|$ (resp. $|\text{co-}L|$) is finite.

Let Σ and Γ be alphabets. A *substitution* (from Σ to Γ) is a function h that assigns to each $x \in \Sigma$ a subset of Γ^*. h is naturally extended to a function (still denoted by) h that assigns to each word in Σ^* a set of words in Γ^*: $h(a_1 a_2 \cdots a_n) := h(a_1)h(a_2) \cdots h(a_n)$. We say h is *non-erasing* if $\epsilon \notin h(x)$ for all $x \in \Sigma$. A *homomorphism*

is a substitution h where each set $h(x)$ has exactly one element (we may thus regard $h(x)$ as an element of Γ^*). A *letter homomorphism* is a homomorphism where $h(x)$ is a word of length 1 for all $x \in \Sigma$ (we may thus regard $h(x)$ as an element of Γ). An *isomorphism* is a letter homomorphism such that h is a bijection from Σ to Γ . An isomorphism is therefore only a 'renaming' of the alphabet.

For every substitution h from Σ to Γ, we may define the language operator (again denoted by h) that takes a language (Σ, L) to the language $(\Gamma, h(L))$ where $h(L)$ is the union of the sets $h(w)$ over all $w \in L$. We also define the *inverse substitution* operator h^{-1} that takes a language (Γ, L') to $(\Sigma, h^{-1}(L'))$ where $h^{-1}(L')$ is the set $\{w \in \Sigma^* \ : \ h(w) \subseteq L'\}$.

A *(language) class* K is a collection of languages that is closed under isomorphism. We emphasized in Chapter 1 that complexity theory is primarily the study of language classes, not of individual languages. The classes which interest us are usually defined using machines that use a limited amount of computing resources. In this case, we call K a *complexity class* although this is only an informal distinction.

Operators are important tools for analyzing the structure of complexity classes. For instance, many important questions are of the form "Are two complexity classes K and K' equal?". If Ω is any set of operators, let

$$\Omega(K) := \{\omega(L_1, \ldots, L_d) : L_i \in K, \omega \text{ is a } d\text{-ary operator in } \Omega\}.$$

The *closure* of K under Ω is

$$\Omega^*(K) := \bigcup_{n \geq 0} \Omega^n(K)$$

where $\Omega^0(K) := K$ and $\Omega^{n+1}(K) := \Omega(\Omega^n(K))$. One possible approach to showing that K is not equal to K' is to show that, for a suitable class of operators Ω, K is closed under Ω (i.e., $\Omega^*(K) = K$) but K' is not. An important simple case of Ω is where Ω consists of just the Boolean complement operator; here $\Omega(K) = \{\text{co-}L \ : \ L \in K\}$ is simply written as co-K. A branch of formal language theory called AFL theory investigates closure questions of this sort and certain complexity theory questions can be resolved with this approach.

Contents

Chapter 2

The Turing Model: Basic Results

April 13, 2009

2.1 Introduction

We take the Turing model of computation as the *canonical* one in Complexity Theory. In this we are simply following the usual practice but other more logical reasons can be given: the fundamental analysis by which Turing arrives at his model is still one of the most cogent arguments in support of Church's thesis.[1] The simplicity of Turing's basic model is appealing. Despite the fact that Turing predated our computer age, there is a striking resemblance between his machines and modern notions of computation. Henceforth, any new model that we introduce shall (perhaps only implicitly) be compared with this canonical choice. Of course, Turing considered only the fundamental mode of computation but we can naturally adapt it to the other computational modes. Furthermore, we find it convenient to consider variants of the original Turing machine. Recall from chapter 1 that all these model-specific details turn out to be unimportant for the major conclusions of our theory. It is somewhat paradoxical that some form of these model-dependent details cannot be avoided in order to attain our model-independent conclusions.[2]

In this chapter we will study some basic complexity results in the Turing model. For the time being, we restrict ourselves to the fundamental and the nondeterministic modes of computation. The central features of a Turing machine, as we saw in Chapter 0, consist of a finite-state 'black-box' operating on one or more tapes where each tape is divided into an infinite linear sequence of squares. Each square is capable of storing a single symbol (chosen from a finite set depending on the particular machine). Each tape has a reading head scanning a square. Under the control of the black-box, in one step the heads can change the contents of the square being scanned and can move left or right to the adjacent squares. The following computational resources have been studied with respect to Turing machines:

time, space, reversals, ink, tapes, symbols, states.

Time and space are intuitively clear, so we briefly explain the other resources. 'Reversal' is the number of times some tape head changes directions. 'Ink' measures the number of times the machine has to write on a tape square (a tape square may be rewritten many times). 'Tapes' ('symbols, ' 'states, ' respectively) are the number of tapes (symbols, states, respectively) used in the definition of the machine. The latter three resources are called *static resources* because they are a function of the machine description only. The others are *dynamic resources* since the amount of (say) time or space used also depends on the particular computation. In Complexity Theory, we primarily study dynamic resources. Time, space and reversals are considered basic and will be studied in this chapter. Reversals (in contrast to time and space) may initially appear artificial as a measure of complexity. It is a comparatively recent realization that reversal is an essential ingredient for gluing together space and time.

Reading Guide. This is a long chapter containing basic technical results about the Turing model. The reader should be familiar with the vocabulary of formal language theory; the appendix in Chapter 0 may be referred to. Sections 6, 7 and 8 form a group of results about efficient simulations that use time, space and reversals (respectively). These techniques are fundamental and should be mastered. Sections 10 and 11 may be safely omitted since later chapters do not rely on them.

[1] According to Martin Davis [10] Gödel did not believe in Church's thesis until he heard of Turing's results. Gödel's skepticism arose from his insistence on an *á priori* analysis of the concept of computation (which is what Turing provided). The article contains an authoritative account of the origins of Church's thesis.

[2] One is reminded of Sir Arthur Eddington's elephants in *Nature of the Physical World*.

2.2 Turing Machines

The *multitape Turing machine* is now introduced. Turing machines will be used in two capacities: to define languages and to define functions over words.[3]In the former capacity, we call them *acceptors* and in the latter, *transducers*. For now, we focus on acceptors.

We regard a Turing acceptor as having two parts: a **transition table** δ to define a set of machine instructions, together with an **acceptance rule** to specify when a computation of accepts its input. The advantage of this approach is that, just by specifying alternative acceptance rules, we obtain different choice modes. This separation is also a recognition of the distinction between the Turing model of computation and computational modes (which are basically model independent).

Informally, a multitape Turing acceptor may be viewed as a physical machine or system consisting of a finite state automaton equipped with $k+1$ tapes for some $k \geq 0$: an *input tape* (tape 0) and k *work tapes* (tapes $1, \ldots, k$). Each tape consists of a doubly infinite sequence of *cells* (or *tape squares*) indexed by the integers. Each tape square contains a symbol in Σ_∞ . On tape i ($i = 0, \ldots, k$), there is a *head* (head i) that *scans* some tape cell. The *input head* refers to head 0 and the rest are called *work heads*. To begin a computation, we are given an input word $x = a_1 a_2 \cdots a_n$ on the input tape; initially every cell of the each tape contains the blank symbol except for cells $1, 2, \ldots, n$ of the input tape containing the input string x. The head on each tape initially scans cell 1. Depending on the current state and the $k + 1$ symbols scanned under the heads, the machine has certain instructions that are *executable*. Executing one of these instructions results in a new state, in new symbols under each head, and in the movement of each head at most one cell to the right or left of the current cell. The machine halts when it has no executable instruction. By definition, the input head is restricted so that it may not change the contents of its tape. In general, a tape with such a restriction is termed *read-only*. Figure 2.1 illustrates a multitape Turing machine in state q and scanning b_0 on the input tape.

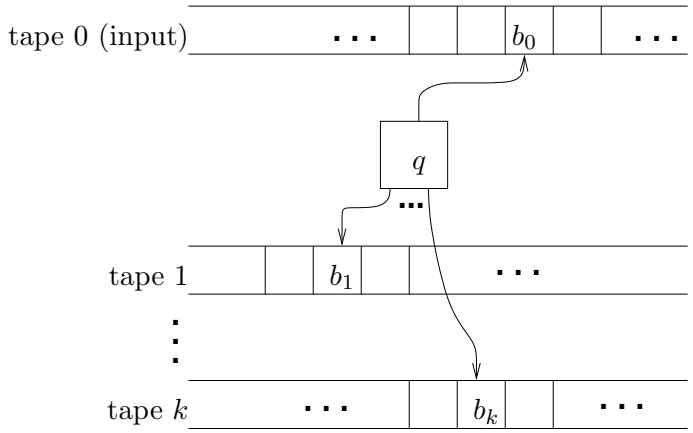

Figure 2.1: An off-line k-tape Turing machine.

The machine just described is also known as an *off-line k-tape Turing machine.* If we restrict the input head so that it cannot move left, we obtain an *on-line* version. In general, a tape is called *one-way* if its head is constrained to move in only one direction (and for emphasis, we call ordinary tapes *two-way*). (In the literature, off-line and on-line machines are also called 2-way or 1-way machines, respectively.) Sometimes we have a k-tape ($k \geq 1$) Turing machine *without* an input tape; then the input is conventionally placed on one of the work-tapes; the special case where $k = 1$ is the simple Turing machine of Chapter 0. In the rest of this book, we mainly use the off-line multitape version (and simply call them 'Turing machines' without qualifications). The main reason for the use of an input tape is so that we can discuss sublinear space complexity (otherwise, the space is at least linear just to represent the input).

Remark: Many other variants of Turing machines have been studied: we may allow more than one head on each tape, consider tapes that are multi-dimensional (the above tapes being one-dimensional), allow tapes that have a binary tree structure, tape heads that can remember a fixed number of previously visited positions and can 'reset' to these positions in one step, etc. For reasons discussed in the previous chapter, these variations are not inherently interesting unless they give rise to interesting new complexity classes.

Recall our Conventions (α) and (β) in Chapter 0 concerning the universal sets Σ_∞ and Q_∞ of symbols and states. We assume these in our formalization of multitape Turing machines.

Definition 1. Let $k \geq 0$. A *k-tape transition table* δ is a finite subset of

$$\Sigma_\infty^{k+1} \times Q_\infty \times \Sigma_\infty^k \times Q_\infty \times \{+1, 0, -1\}^{k+1}$$

■

Each $(3k+4)$-tuple in δ may be written in the form of a transition rule

$$\langle b_0, b_1, \ldots, b_k, q_1 \rightarrow c_1, \ldots, c_k, q_2, d_0, \ldots, d_k \rangle \tag{1}$$

and represents a machine instruction which is interpreted as follows:

> "if b_i (for each $i = 0, \ldots, k$) is the symbol being scanned on tape i and q_1 is the current state, then change b_j (for each $j = 1, \ldots, k$) to c_j (leaving b_0 unchanged), make q_2 the next state, and the ith tape head (for $i = 0, \ldots, k$) should move left, stay in place or move right according as $d_i = -1, 0$ or $+1$"

To simplify proofs, it is convenient to make the following restrictions on transition tables:

(1) The input head never scans more than one blank cell past the input word in either direction.

(2) There are no state transitions *from* the accept or reject states q_a, q_r or *into* the initial state q_0. (Recall that q_a, q_r, q_0 are distinguished states in Q_∞.)

(3) The blank symbol $\sqcup$ cannot be written on any of the work tapes (thus, after a blank cell is visited it turns non-blank). Formally, this means that in (1), $c_j \neq \sqcup$ for each j.

A transition table satisfying the above restrictions is said to be in *standard form*. It is easy to make syntactic restrictions on transition tables to ensure that they are in standard form. The restrictions on δ to obtain an on-line machine is also simple.

The *tape alphabet* of δ is the set of symbols (excluding the blank symbol) that occur in some tuple of δ. The *input alphabet* of δ is the set of symbols (excluding the blank) that occur as the first component of a tuple in δ (i.e., as b_0 in (1)). Thus the input alphabet is a subset of the tape alphabet. The *state set* of δ is similarly defined.

A *configuration* of δ is a $(2k+3)$-tuple

$$C = \langle q, w_0, n_0, w_1, n_1, \ldots, w_k, n_k \rangle = \langle q, w_i, n_i \rangle_{i=0}^k \tag{2}$$

where q is in the state set of δ, each $w_j (j = 1, \ldots, k)$ is a word over the tape alphabet of δ, w_0 is a word over the input alphabet of δ, and $0 \leq n_i \leq 1 + |w_i| (i = 0, \ldots, k)$. The string w_i represents the non-blank portion of tape i. The convention that blanks cannot be written ensures that the non-blank portion of each tape is contiguous. The integer n_i indicates that head i is scanning the n_ith symbol in w_i. For a string x, let $x[i]$ denote the ith symbol in x if $1 \leq i \leq |x|$. If i is outside the indicated range, then by definition, $x[i]$ denotes the blank symbol. Thus if w_i is not the empty string ϵ then $n_i = 1$ (respectively, 0) means the first (respectively, the blank prior to the first) symbol of w_i is scanned. The *initial configuration* on input x is defined to be

$$C_0(x) = \langle q_0, x, 1, \epsilon, 1, \ldots, \epsilon, 1 \rangle$$

where q_0 is the start state.

Remarks: The integers n_i are *relative* addresses of cells, namely, relative to the leftmost non-blank cell. Occasionally we use instead the absolute addressing of cells, but there ought to be no confusion. In the literature, configurations are also called *instantaneous descriptions (ID's)*.

We define the binary relation $\vdash_\delta$ (or $\vdash$, if δ is understood) on configurations of δ. Informally, $C \vdash_\delta C'$ means the configuration C' is obtained by modifying C according to some instruction of δ. More precisely, suppose C is given by (2) and δ contains the instruction given in (1). We say the instruction (1) is *applicable* to C if $q = q_1$ and $w_i[n_i] = b_i$ for each $i = 0, \ldots, k$. Then *applying* (1) to C we get

$$C' = \langle q', w_0', n_0', \ldots, w_k', n_k' \rangle$$

where $q' = q_2$, $n_i' = d_i + max\{1, n_i\}$ $(i = 0, \ldots, k)$, [4] $w_0' = w_0$ and for $j = 1, \ldots, k$:

$$w_j' = \begin{cases} c_j w_j & \text{if} \quad n_j = 0 \\ w_j c_j & \text{if} \quad n_j = |w_j| + 1 \\ u_j c_j v_j & \text{if} \quad w_j = u_j b_j v_j, \text{and} |u_j b_j| = n_j \end{cases}$$

[4] This unintuitive formula takes care of the case $n_i = 0$ in which case $n_i' = 1$ even though $d_i = 0$.

We write $C \vdash_\delta C'$ in this case; C' is called a *successor* of C and the sequence $C \vdash_\delta C'$ is called a **derivation** or **transition** or a *step*. Write $C \vdash_\delta^k C'$ if there is a k-step sequence $C_0 \vdash C_1 \vdash \cdots \vdash C_k$ such that $C = C_0$ and $C' = C_k$. Typically, we simply write "$\vdash$" instead of $\vdash_\delta$ when δ is understood from the context. If there is some $k \geq 0$ such that $C \vdash^k C'$, then we simply write $C \vdash^* C'$. Alternatively, $\vdash^*$ is the reflexive, transitive closure of the binary relation $\vdash$.

In general, several instruction (if any) of δ may be applicable to C. We say δ is *deterministic* if each configuration C has at most one applicable instruction. Otherwise δ *has choice*. We say the Turing machine with a deterministic δ operates in the *deterministic mode*, otherwise it operates in the *choice mode*.

A configuration is *accepting* if its state is the accept state q_a. Similarly for *rejecting* if its state is the accept state q_r. It is *terminal* if it has no successor. For transition tables in standard form, accepting or rejecting configurations are terminal. A *computation sequence* is either a finite sequence of configurations

$$\overline{C} = (C_0, C_1, \ldots, C_m) = (C_i)_{i=0}^m$$

or an infinite one

$$\overline{C} = (C_0, C_1, C_2, \ldots) = (C_i)_{i \geq 0}$$

such that for each $i \geq 0$, $C_i \vdash C_{i+1}$. We call $\overline{C}$ a *computation path* if, in addition, C_0 is an initial configuration and, when the path is finite, C_m is terminal. We may denote $\overline{C}$ by

$$\overline{C} = C_0 \vdash C_1 \vdash C_2 \vdash \cdots \vdash C_m$$

or

$$\overline{C} = C_0 \vdash C_1 \vdash C_2 \vdash \cdots .$$

We call $\overline{C}$ an *accepting path on input* x if it is a finite computation path with $C_0 = C_0(x)$ and C_m is an accepting configuration. *Non-accepting* computation paths come in two flavors: they either do not terminate or terminate in non-accepting configurations. This distinction is sometimes crucial (see §9; also chapters 7 and 8). The next definition captures an important form of the choice mode:

Definition 2. A *nondeterministic Turing acceptor* M is given by a transition table $\delta = \delta(\mathrm{M})$, together with the following acceptance rule.

> **Nondeterministic Acceptance Rule.** A word x is *accepted* by M iff x is over the input alphabet of M and there is an accepting computation path of δ for x.

■

M is a *deterministic Turing acceptor* if δ is deterministic. By convention, a Turing acceptor is assumed to be nondeterministic unless otherwise specified. The *language accepted* by M is given by (Σ, L) where Σ is the input alphabet of $\delta(\mathrm{M})$ and L consists of those words that are accepted by M. We write $L(\mathrm{M})$ for the language accepted by M.

Example 1. We describe informally a 1-tape deterministic Turing acceptor which accepts the palindrome language $L_{pal} = \{w \in \{0,1\}^* : w = w^R\}$ where w^R is the reversal of the word w. The reader should feel comfortable in translating this into a formal description (i.e. in terms of δ) because, from now on, we shall describe Turing machines in such informal terms. The acceptor works in three stages:

(i) Copy the input w onto tape 1.

(ii) Move the input head back to the start of the input, but leave the head on tape 1 at the right end.

(iii) Move the input head right and head 1 left, in synchrony, comparing the symbols under each head – they should agree or else we reject at once. Accept iff all the symbols agree.

■

Example 2. (Guessing, verifying and nondeterminism) We give a nondeterministic 1-tape acceptor for the complement co-L_{pal} of the palindrome language. Note that $x \in \{0,1\}^*$ is in co-L_{pal} iff for some i:

$$1 \leq i \leq n = |x| \text{ and } x[i] \neq x[n-i+1]. \tag{3}$$

Using nondeterminism we can 'guess' such an i and 'verify' that it has the properties in (3). By 'guessing' an i we mean that the machine initially enters a state q_1 such that in this state it has two choices:

(i) Write down a '1' in its single work-tape, move head 1 one position to the right and remain in state q_1;

(ii) Write a '1' in the work-tape, keeping head 1 stationary and enter a new state q_2.

During the guessing stage (state q_1), the input head does not move. When we enter state q_2, there is a unary word on work-tape. This is the unary representation of the guessed i. Let x be the input of length n. To 'verify' that i has the right properties, we can determine $x[i]$ by moving the input head to the right while moving head 1 to the left, in synchrony. When head 1 reaches the first blank past the beginning of i, the input head would be scanning the symbol $x[i]$. We can 'remember' the symbol $x[i]$ in the finite state control. Notice that the guessed i may be greater than n, and this could be detected at this point. If $i > n$ then we reject at once (by this we mean that we enter some non-accepting state from which there are no transitions). Assuming $i \leq n$, we can similarly determine $x[n-i+1]$ by moving the input head to the end of x and moving i steps to the left, using tape 1 as counter. We accept if $x[i] \neq x[n-i+1]$, rejecting otherwise.

To check that the above machine accepts the complement of L_{pal}, observe that if x is in L_{pal} then every computation path will be non-accepting; and if x is not in L_{pal} then the path corresponding to the correct guess of i will lead to acceptance. This example shows how nondeterminism allows us to check if property (3) for *some* values of i, by testing for all values of i simultaneously. This ability is called "guessing", and generally, it simplifies the logic of our machines. ■

Transformations and Turing Transducers

This subsection is placed here for easy reference. The concepts of transformations and transducers will be used in section 9.

Definition 3. A *multivalued transformation* is a total function

$$t \ : \ \Sigma^* \to 2^{\Gamma^*}$$

where Σ, Γ are alphabets and 2^S denotes the power set of S. If $t(x)$ is a singleton set for all $x \in \Sigma^*$ then we call t a *transformation* and think of t as a function from Σ^* to Γ^*. ■

Definition 4. A *nondeterministic k-tape transducer* T is a nondeterministic $(k+1)$-tape $(k \geq 0)$ Turing acceptor such that tape 1 is constrained to be one-way (i.e., the head may not move left). Tape 1 is called the *output tape* and the non-blank word on this tape at the end of any accepting computation path is called the *output word* of that path. The work tapes of T now refer to tapes 2 to $k+1$. T is said to compute a multivalued function t in the natural way: for any input x, $t(x)$ is the set of all output words y in accepting computation paths. ■

The set $t(x)$ can be empty. In case T computes a transformation, then T has the following properties: (a) For each input x, there exists at least one accepting path. (b) On a given x, all accepting paths lead to the same output word. We call such a transducer *univalent*. The usual example of a univalent transducer is a deterministic transducer that accepts all its inputs.

Complexity of transducers. In the next section we will define computational complexity of acceptors for various resources (time, space, reversal). The computational complexity of a transducer T is identical to the complexity when T is regarded as an acceptor, with the provision that space on tape 1 is no longer counted.

2.3 Complexity

In this section, we see the first concepts of computational complexity. We consider the three main computational resources of time, space and reversal.

The *time* of a computation path $\overline{C} = (C_0, C_1, \cdots)$ is one less than the length of the sequence $\overline{C}$; the time is infinite if the length of the sequence is infinite. The *space used by* a configuration $C = \langle q, w_i, n_i \rangle_{i=0}^{k}$ is defined as

$$space(C) = \sum_{i=1}^{k} |w_i|$$

Observe that the space in the input tape is not counted. The *space used by* the computation path $\overline{C} = (C_i)_{i \geq 0}$ is $\sup\{space(C_i) \ : \ i \geq 0\}$; again the space could be infinite. It is possible for the time to be infinite while the space remains finite.

The definition of reversal is a little subtle. Let $\overline{C} = (C_i)_{i\geq 0}$ be a computation path of a k-head machine. We say that head h ($h = 0, \ldots, k$) *tends in direction* d (for $d \in \{-1, +1\}$) in C_i if the last transition $C_{j-1} \vdash C_j$ (for some $j \leq i$) preceding C_i in which head h moves is in the direction d. This means that the head h "paused" in the time from $j+1$ to i. If head h has been stationary from the start of the computation until C_i we say head h tends in the direction $d = 0$ in C_i. We also say that head h has *tendency* d if it tends in direction d. It is important to note that the head tendencies of a configuration C are relative to the computation path in which C is embedded. We say head h makes a *reversal* in the transition $C_{j-1} \vdash C_j$ in $\overline{C}$ if the tendency of head h in C_{j-1} is *opposite* to its tendency in C_j. We say $C_{j-1} \vdash C_j$ is a *reversal transition*. The *reversal* in a reversal transition $C_{j-1} \vdash C_j$ is the number of heads (including the input head) that makes a reversal. So this number is between 1 and $1+k$. The *reversal* of $\overline{C}$ is the total number of reversals summed over all reversal transitions in the entire computation path.

Observe that changing the tendency from 0 to ± 1 is not regarded as a reversal. Each computation path $\overline{C}$ can be uniquely divided into a sequence of disjoint sub-computation paths $P_0, P_1, \ldots$, where every configuration in P_i ($i = 0, 1, , \ldots,$) has the same head tendencies. Each P_i is called a *phase*. The transition from one phase P_i to the next P_{i+1} is caused by the reversal of at least one head. Clearly the reversal of $\overline{C}$ bounded by $k+1$ times the the number of phases in $\overline{C}$.

Remarks: An alternative to our definition of reversal complexity is to discount reversals caused by the input head. Hong [20] (see section 8.3) uses this alternative. This would be consistent with our decision not to count the space used on the input tape. However, the real concern there was to admit sublinear space usage to distinguish problems with low space complexity. Our decision here allows possible complexity distinctions which might be otherwise be lost (see remark at the end of §7).

We give two ways to define the usage of resources.

Definition 5. (Acceptance complexity) If x is an input for M, define $AcceptTime_{\mathrm{M}}(x)$ to be the least time of an accepting computation path for x. If M does not accept x, $AcceptTime_{\mathrm{M}}(x) := 0$. If $n \in \mathbb{N}$, define

$$AcceptTime_{\mathrm{M}}(n) := \max_x \{AcceptTime_{\mathrm{M}}(x)\}$$

where x ranges over words in $L(\mathrm{M})$ of length n. Note that if $L(\mathrm{M})$ has no words of length n then $AcceptTime_{\mathrm{M}}(n) = 0$. Let f be a complexity function. M *accepts in time* f if $f(n)$ dominates $AcceptTime_{\mathrm{M}}(n)$. ■

Note that $AcceptTime_{\mathrm{M}}(x)$ is a complexity function; it is a partial function since it is defined iff x is a natural number. Although this definition is stated for the time resource, it extends directly to the space or reversal resources. Indeed, all the definitions we give in this section for the time resource naturally extend to space and reversal. One consequence of our definition is that any finite or co-finite language is accepted by some Turing acceptor in time (respectively, space, reversal) $f(n) = O(1)$ (respectively, $f(n) = 0, f(n) = 0$). Another technical consequence is that a k-tape acceptor uses at least k tape cells. Thus to achieve space 0, we must use 0-tape acceptors.

The fact that M accepts x in time r does not preclude some (whether accepting or not) computation on input x from taking more than r time. Furthermore, if x is not in $L(\mathrm{M})$, it is immaterial how much time is used in any computation on x! This stands in contrast to the next definition:

Definition 6. (Running complexity) For any input x for M, let $RunTime_{\mathrm{M}}(x)$ be the maximum over the time of *all* computation paths of M on input x. Thus, $RunTime_{\mathrm{M}}(x) < \infty$ iff M halts on input x. For $n \in \mathbb{N}$, let

$$RunTime_{\mathrm{M}}(n) := \max_x RunTime_{\mathrm{M}}(x)$$

where x ranges over all words (whether accepted by M or not) of length n. For any complexity function f, we say M *runs in time* f if $f(n)$ dominates $RunTime_{\mathrm{M}}(n)$. ■

Note that if M runs in time f then $f(n) \downarrow$ for all $n \in \mathbb{N}$. If $f(n) < \infty$ then M must halt in every computation path on inputs of length n. If M is deterministic and $x \in L(\mathrm{M})$ then $AcceptTime_{\mathrm{M}}(x) = RunTime_{\mathrm{M}}(x)$.

Example 3. Refer to examples 1 and 2 in the last section. The space and time of the deterministic acceptor for palindromes are each linear; the reversal is 3. Similarly, acceptance space and time of the nondeterministic acceptor for the complement of palindromes are each linear, with reversal 3. However the running space and time of the nondeterministic acceptor is infinite for all n since the guessing phase can be arbitrarily long. We leave as exercises for the reader to modify the machine to satisfy the following respective complexity bounds: (a) the running space and time are each linear; (b) the time is linear and space is logarithmic. In (a), you can actually use a 1-tape machine and make only 1 reversal. What is the reversal in (b)? ■

The following notion is sometimes useful: we call f *time-constructible* function if there is a deterministic M such that for *all* inputs x of sufficiently large length, $RunTime_{\mathrm{M}}(x) = f(|x|)$.
M is said to *time-construct* f.

The definitions of *space-constructible* or *reversal-constructible* are similar: systematically replace the word 'time' by 'space' or 'reversal' above. We have defined constructible functions with respect to deterministic machines only, and we use running complexity. Such functions have technical applications later. The reader may try this: describe a 1-tape Turing machine that time-constructs the function n^2. It may be easier to first "approximately" time-construct n^2. See the Exercises for more such constructions.

The Turing acceptor for the palindrome language in the example of the last section is seen to run in linear time and linear space. But see section 9 for a more space-efficient machine: in particular, logarithmic space is sufficient (but at an inevitable blow-up in time requirement).

We shall use acceptance complexity as our basic definition of complexity. The older literature appears to prefer running complexity. Although running complexity has some desirable properties, in proofs, it sometimes takes additional effort to ensure that every computation path terminates within the desired complexity bound. Consequently, proofs for accepting complexity are often (but not always) slightly shorter than for running complexity. A more compelling reason for acceptance complexity is that it is more basic: we could define 'rejection complexity' and view running complexity as a combination of acceptance and rejection complexity. The problem of termination can be overcome by assuming some 'niceness' conditions on the complexity function. The exact-time complexities and time-constructible complexities are typical notions of niceness. Indeed, for nice functions, complexity classes defined using running complexity and acceptance complexity are identical. Since most of the important complexity classes (such as P , NP, etc, defined next) are bounded by nice functions, the use of running or acceptance complexity lead to the same classes in these cases. It is important to realize that most common functions are nice (see Exercises). The general use of running complexity is largely avoided until chapter 8 when we discuss stochastic computations where it seems that running complexity is the more fruitful concept.

Resource bounded complexity classes. We introduce some uniform notations for complexity classes defined by bounds on computational resources used by acceptors.[5] Let F be a family of complexity functions. The class $NTIME(F)$ is defined to consist of those languages accepted by nondeterministic Turing acceptors in time f, for some f in F. If $F = \{f\}$, we simply write $NTIME(f)$. The notation extends to languages accepted by deterministic Turing acceptors ($DTIME(F)$), and to space and reversal complexity ($XSPACE(F)$ and $XREVERSAL(F)$ where $X = N$ or D indicates nondeterministic or deterministic classes). When *running complexity* (time, space, reversal, etc) is used instead of acceptance complexity, a subscript 'r' is appended to the usual notations for complexity classes. For instance, $NTIME_r(F)$, $NSPACE_r(F)$, $DREVERSAL_r(F)$, etc. It is clear that $DTIME_r(F) \subseteq DTIME(F)$ and $NTIME_r(F) \subseteq NTIME(F)$, and similarly for the other resources.

Canonical Classes. Most common families of complexity functions can be obtained by iterating the following operations $lin, poly, expo$ on an initial family F: $lin(F) = O(F)$, $poly(F) = F^{O(1)}$, $expo(F) = O(1)^F$. For instance, the following families will be used often:

$$\{\log n\}, \log^{O(1)} n, \{n+1\}, O(n), n^{O(1)}, O(1)^n, O(1)^{n^{O(1)}}.$$

Each family in this sequence 'dominates' the preceding family in a natural sense. For future reference, we collect in the following table the special notation for the most important complexity classes:

[4] In the literature, a function that $f(n)$ of the form $RunTime_M(n)$ is sometimes said to be 'time-constructible'. What we call time-constructible is also described as 'fully time-constructible'. Other related terms include: measurable (Hopcroft-Ullman), real-time computable (Yamada) or self-computable (Book), honest (Meyer-McCreight).

[5] We appear to be following Ronald Book in many of these notations. He attributes some of these suggestions to Patrick Fischer.

Canonical Classes

Special Symbol	Standard Notation	Name
REG	$DSPACE(0)$	regular languages
$DLOG$	$DSPACE(\log n)$	deterministic log-space
$NLOG$	$NSPACE(\log n)$	n ondeterministic log-space
$PLOG$	$DSPACE(\log^{O(1)} n)$	polynomial log-space
P	$DTIME(n^{O(1)})$	deterministic poly-time
NP	$NTIME(n^{O(1)})$	nondeterministic poly-time
$PSPACE$	$DSPACE(n^{O(1)})$	polynomial space
$DEXPT$	$DTIME(O(1)^n)$	deterministic simply-exponential time
$NEXPT$	$NTIME(O(1)^n)$	nondeterministic simply-exponential time
$DEXPTIME$	$DTIME(2^{n^{O(1)}})$	deterministic exponential time
$NEXPTIME$	$NTIME(2^{n^{O(1)}})$	nondeterministic exponential time
$EXPS$	$DSPACE(O(1)^n)$	simply-exponential space
$EXPSPACE$	$DSPACE(2^{n^{O(1)}})$	exponential space

These classes are among the most important in this theory and for convenience, we shall refer to the list here as the *canonical list*.[6] Note that our distinction between "exponential" and "simply-exponential" is not standard terminology. It will follow from results later in this chapter that (after omitting $PLOG$, $DEXPT$ and $NEXPT$) each class in the above list is included in the next one on the list.

In Chapter 1 we said that complexity theory regains a Church-like invariance property provided that we parametrize the complexity classes with (1) the computational resource and (2) the computational mode. Our notation for complexity classes reflects this analysis: each class has the form

$$\mu - \rho(F)$$

where $\mu = D, N$, etc., is the computational mode, and $\rho = TIME, SPACE$, etc., is the computational resource. According to the computational theses in Chapter 1, these classes are model-independent in case $F = n^{O(1)}$. In particular, the canonical classes P , NP and $PSPACE$ have this invariance. That is, if we had defined these notions using some other general computational model such as pointer machines instead of Turing machines, we would have ended up with the same classes (P , NP, $PSPACE$).

Simultaneous resource bounds. In the preceding definitions of complexity, we bound one resource but place no restrictions on the other resources (actually the resources are not completely independent of each other). In 'simultaneous complexity' we impose bounds on two or more resources within the same computation path.

Definition 7. (Acceptance within simultaneous bounds) Let $t, s \geq 0$ be real numbers and f, g be complexity functions. An acceptor M accepts an input x *in time-space bound of* (t, s) if there exists an accepting computation path for x whose time and space are at most t and s, respectively. M accepts in time-space (f, g) if there is some n_0 such that for all $x \in L(\mathrm{M})$, if $|x| \geq n_0$ then M accepts x in time-space $(f(|x|), g(|x|))$. ■

For any complexity functions t and s, let X-$TIME$-$SPACE(t, s)$ (where $X = D$ or N) denote the class of languages L that are accepted by some (deterministic or nondeterministic, depending on X) Turing machine M that accepts in time-space (t, s). The complexity classes defined in this way are called *simultaneous space-time classes*. Clearly,

$$X\text{-}TIME\text{-}SPACE(t, s) \subseteq XTIME(t) \cap XSPACE(s)$$

although it seems unlikely that this would be an equality in general. Similarly we will consider simultaneous space-reversal, time-reversal, time-space-reversal classes denoted, respectively,

$$X\text{-}SPACE\text{-}REVERSAL(s, r), X\text{-}TIME\text{-}REVERSAL(t, r)$$

and

$$X\text{-}TIME\text{-}SPACE\text{-}REVERSAL(t, s, r)$$

where t, s and r are time, space and reversal complexity bounds. For a more compact notation, we could use (but do not recommend) the alternative $XTISP$, $XTIRE$, $XSPRE$ and $XTISPRE$ for X-$TIME$-$SPACE$, X-$TIME$-$REVERSAL$, etc.

Finally, we note these notations will be extended later (see chapter 7) when 'X' in these notations may be replaced by symbols for other computational modes.

[6] Note that P should really be 'DP' but the 'P ' notation is so well-accepted that it would be confusing to change it. $DLOG$ is also known as L or LOG in the literature. Similarly, $NLOG$ is also known by NL.

Recursively enumerable languages. In our table of canonical classes, the first entry is the class of regular languages $DSPACE(0)$. We will investigate its properties in Section 11. The complexity class on the other extreme comprises languages that are accepted by Turing acceptors without any complexity bounds:

Definition 8. A language is *recursively enumerable* or, r.e., if it is accepted by some deterministic Turing acceptor. A *halting Turing machine* is one that does not have any infinite computation path on any input. A language is *recursive* if it is accepted by some halting Turing acceptor. Let RE and REC denote the classes of r.e. and recursive languages, respectively. ■

In recursive function theory, the fundamental objects of study are partial number-theoretic functions instead of languages. The recursive function theory analogues of the above definitions are: a partial function from the natural numbers to natural numbers is *partial recursive* if it is computed by some Turing transducer (assume that natural numbers are encoded in binary) where the function is undefined at the input values for which the transducer does not halt. If the function is total, then it is called a *(total) recursive function*.

Essentially, the class RE corresponds the partial recursive functions, while the class REC corresponds to the recursive functions. Whether we use languages or number-theoretic functions to study computability is a matter of taste (roughly corresponding to that of the computer scientist or that of the mathematician). The advantage of language theoretic approach (which is that taken by this book) is that it corresponds more to real world computation, and is better suited for subrecursive complexity theory (which is our subject). On the other hand, the number-theoretic approach has a simplicity and elegance that is generally[7] hard to beat.

2.4 Linear Reduction of Complexity

The results to be shown are of the form: if a language can be accepted in (time, space, or reversal) resource f then it can be accepted in resource cf, for any $c > 0$. The idea is that, by using a transition table with more states and a larger alphabet, we can trade-off dynamic complexity for static complexity. These technical results are very useful: they justify the use of the 'big-oh' notation for complexity functions. For instance, we can talk of a language being accepted in 'quadratic time' or 'logarithmic space' without ambiguity. It also illustrates our discussion in chapter 1 where we concluded that complexity functions must not be taken in an absolute sense, but only up to O-order.

Our first theorem, taken from Stearns, Hartmanis and Lewis [15], says that the space complexity of a problem can be reduced by any constant factor and it is sufficient to use a 1-tape acceptor.

THEOREM 1 ((Space Compression)). *Let $c > 0$ be a constant and $s = s(\cdot)$ a complexity function. For any multitape* M *that accepts in space s, there is a 1-tape* N *that accepts L(M) in space cs. If* M *is deterministic, so is* N.

Proof. As will be the rule in such proofs, we only informally describe N since the formal description of N is tedious although straightforward. It is sufficient to show this for $c = 1/2$ since we can make c arbitrarily small with repeated applications of the construction. The single work tape of N contains k 'tracks', one track for each work tape of M. See Figure 2.2.

Each cell of N can be regarded as containing a "super symbol" that is essentially a k by 4 matrix with the ith row containing 4 tape symbols (possibly blanks) of the ith work-tape of M, and possibly a marker ($\uparrow$) for the head position. There are five forms for each row:

$$[b_1b_2b_3b_4], [\uparrow b_1b_2b_3b_4][b_1 \uparrow b_2b_3b_4][b_1b_2 \uparrow b_3b_4][b_1b_2b_3 \uparrow b_4]$$

where each b_i is either a blank symbol or a tape symbol of δ(M). The marker $\uparrow$ indicates that the immediately succeeding b_i is being scanned. Thus a symbol of N encodes $4k$ symbols of M. Since all the scanned cells have absolute indexing from $-s(n) + 1$ to $s(n) - 1$ on an input of length n, we see that $s(n)/2 + O(1)$ super symbols of N suffice to encode the contents of M's tapes. But in fact $\lfloor s(n)/2 \rfloor$ cells suffice if we exploiting the finite state machinery of N to store a constant amount of extra information. N simulates a step of M by making a 'right-sweep' (i.e., starting from the leftmost non-blank cell to the rightmost non-blank) of the non-blank portion of N's tapes, followed by a return 'left-sweep'. The 'current neighborhood' of M consists of those cells of M that either are scanned by some tape head, or are immediately adjacent to such a scanned cell. On the right-sweep, N takes note of (by marking in some appropriate manner) the current neighborhood of M. On the left-sweep, it updates the current neighborhood to reflect the changes M would have made. It is clear that N accepts in space $s(n)/2$. Note that N is deterministic if M is deterministic. □

[7] But not always. For instance, in the Gödelization of machines, the encoding of machines by strings is straightforward. But to encode a machine as a number appears more contrived.

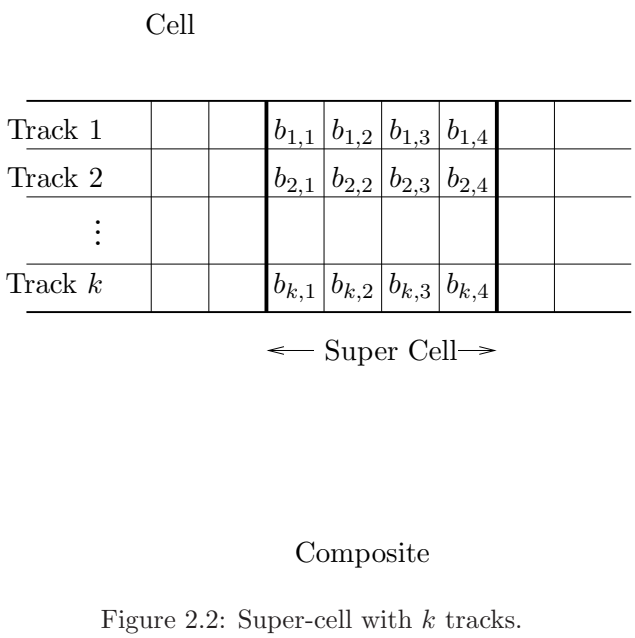

Figure 2.2: Super-cell with k tracks.

The proof illustrates the reduction of dynamic complexity at the expense of static complexity. In this case, we see space being reduced at the cost of increasing the number of tape symbols and the number of states. The amount of information that can be encoded into the static complexity of an acceptor is clearly finite; we refer to this technique as *storing information in the finite state control* of the acceptor. While the reader may feel this is "cheating", the fact remains that there are limits as how much this cheating helps; such limits have inherent interest.

COROLLARY 2. *Let $X = D$ or N. For any complexity function s,*

$$XSPACE(s) = XSPACE(O(s)).$$

The reader may verify that Theorem 1 and its corollary hold if we use 'running space' rather than 'acceptance space' complexity. The next result from Hartmanis and Stearns [16] is the time analog of the previous theorem.

THEOREM 3 ((Linear Speedup)). *Given $c > 0$ and a k-tape acceptor* M *that accepts in time $t(n) > n$, there is a $(k+1)$-tape* N *that accepts in time $n + ct(n)$ with $L(\mathrm{M}) = L(\mathrm{N})$. If* M *is deterministic so is* N.

Proof. Choose $d > 1$ to be an integer to be specified. Tape j of N, for $j = 1, \ldots, k$, will encode the contents of tape j of M using super symbols. Similar to the last proof, each super symbol encodes d tape symbols (including blanks) of M; but unlike that proof, the present super symbols do not need to encode any head positions of M and do not need multiple tracks. Tape $k+1$ in N will be used to re-code the input string using super symbols. We describe the operation of N in two phases. Let the input string be x with $|x| = n$.

Set-up Phase. First N copies the input x from tape 0 to tape $k+1$. This takes $n+1$ steps. So now tape $k+1$ contains x in a 'compressed form'. Next N moves head $k+1$ leftward to the first non-blank super symbol using $\frac{n}{d} + O(1)$ steps. Henceforth, N ignores tape 0 and treats tape $k+1$ as the input tape.

Simulation Phase. N now simulates M by making repetitive executions of a 6-step movement. Each such movement simulates d moves of M. It is enough to explain these 6 step on some generic tape $t = 1, \ldots, k$. Define the 'current neighborhood' of N on tape t to be the supercells containing the (ordinary) cells on tape t that M will visit in the next d steps. These will comprise at most 2 cells, the current supercell j and another cell (either supercell $j-1$ or $j+1$). We use 4 steps to determine the current neighborhood of the t-th head (e.g., move left, right, right, left). In two more steps, N can update the current neighborhood to reflect the result of executing d steps on M. Moreover, N can place its tape heads to be in the correct supercell, in preparation for the next 6-step movement. If M accepts or reject within the next d steps, then N can accepts or rejects before the next movement. This completes the simulation phase.

Timing Analysis. We show that N uses $\leq n + ct(n)$ moves in the above simulation of M. In the set-up phase, $n + \frac{n}{d} + O(1)$ moves were made. In the simulation phase, we see that each sequence makes at most 6 steps correspond to d steps of M. Thus N made $\leq \frac{6t(n)}{d} + O(1)$ moves in this phase. Summing the time for the two phases we get $n + \frac{n}{d} + \frac{6t(n)}{d} + O(1) < n + \frac{7t(n)}{d} + O(1)$. Thus, if we choose $d > 7/c$, the total time will be $\leq n + ct(n)$(ev.). □

The following is immediate:

COROLLARY 4.

(i) If $t(n) = \omega(n)$ then XTIME(t)=XTIME(O(t)), where $X = D$ or N.

(ii) For all $\epsilon > 0$, $DTIME((1+\epsilon)n) = DTIME(O(n))$.

We do not state the nondeterministic version of Corollary 4 (ii) because a stronger result will be shown next. We show that with $k+3$ (rather than $k+1$) work-tapes we can exploit nondeterminism to strengthen Theorem 3. In fact the three additional tapes of N are known as 'checking stacks, ' i.e., the contents of these tapes, once written, are not changed. The basic idea is that the set-up and simulation phases in the proof of Theorem 3 can be carried out simultaneously. Nondeterminism is used in an essential way: the reader who is not familiar with nondeterministic computations will find the proof highly instructive. The proof is adapted from Book and Greibach [4]; the same ideas will be exploited later in this chapter when we consider the problem of reducing the number of work-tapes with no time loss.

THEOREM 5 ((Nondeterministic linear speedup)). *For all $c > 0$, and for any k-tape* M *that accepts in time $t(n) \geq n$, there is a nondeterministic $(k+3)$-tape* N *that accepts in time* $\max\{n+1, ct(n)\}$ *with* $L(\mathrm{M}) = L(\mathrm{N})$.

Proof. Let d be some integer to be specified later. As in the proof of Theorem 3, N 'compresses' the tape contents of M, i.e., it uses super symbols that encode d tape symbols of M at a time. The tapes of N are used as follows. Tape i ($i = 1, \ldots, k$) of N represents in compressed form the contents of tape i of M. Tape $k+1$ contains in compressed form an initial segment of the (actual) input string. Tapes $k+2$ and $k+3$ each contains a copy of a 'guessed' input string, again in compressed form. In addition, the simulation of M by N uses tape $k+3$ as its 'input' tape. The idea is for N to proceed with the 6-move simulation of Theorem 3 using a 'guessed' input string. In the meantime, N verifies that the guess is correct.

We describe the operations of N in two phases; each phase consists of running two simultaneous processes (the reader should verify that this is possible because the simultaneous processes operate on different tapes).

Initial Phase. The following two processes run in parallel.

> *Process 1.* This process copies some initial prefix x_1 of the input x into tape $k+1$, in compressed form. This takes $|x_1|$ steps. The prefix x_1 is nondeterministically chosen (at each step, the transition table of N has two choices: to stop the process instantly or to continue, etc).
> *Process 2.* On tapes $k+2$ and $k+3$, N writes down two copies, *in compressed form*, of a guessed string y in the input alphabet of M. When the writing of y is completed (nondeterministically of course), heads $k+2$ and $k+3$ move synchronously back (ie., left-ward) over the string y. Process 2 halts nondeterministically at some point *before* these heads move past the first blank symbol left of y.

The initial phase ends when both processes halt (if one process halts before the other then it marks time, doing nothing). At this point, on tape 0, we have $x = x_1x_2$ where x_1 is the initial segment that process 1 copied onto tape $k+1$ and head 0 is scanning the first symbol of x_2. Similarly, $y = y_1y_2$ where heads $k+2$ and $k+3$ (on their respective tapes) are scanning the super symbol that contains the first symbol of y_2. Note that any of x_i or $y_i (i = 1, 2)$ may be empty. The figure 2.3 illustrates the situation (not to scale!) where $\bar{w}$ denotes the compressed version of a string w and $\uparrow$ indicates a head position.

Final Phase. We run the following two processes in parallel.

> *Process 3.* The input head continues to scan the rest of the input (i.e., x_2), comparing it with the compressed copy of y_2 on tape $k+2$: for this purpose, note that heads 0 and $k+2$ are conveniently positioned at the left-most symbol of x_2 and y_2, respectively, at the start of the final phase. If $x_2 \neq y_2$, then N rejects at once. Process 3 halts when the checking confirms that $x_2 = y_2$.
> *Process 4.* First we verify that $x_1 = y_1$, using x_1 on tape $k+1$ and y_1 on tape $k+3$. If $x_1 \neq y_1$, N rejects at once. Otherwise, head $k+3$ will now be positioned at the beginning of y and we can begin the "6-move simulation" of M (as in the proof of Theorem 2). Process 4 halts when this simulation is complete.

We remark that in processes 3 and 4 above, when we say "N rejects at once", we mean that the described computation path halts without acceptance. It does not mean that N is rejecting the input in a global sense. This explanation holds generally for other similar situations.

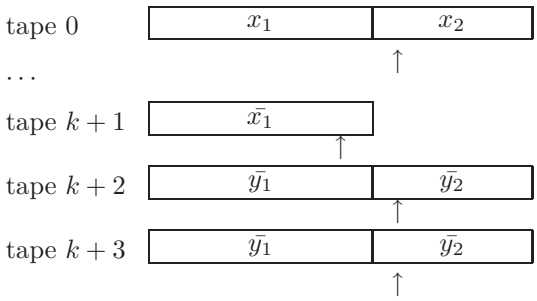

Figure 2.3: At the end of the initial phase

The final phase ends when both processes halt. N accepts if neither process rejects: this means $x = y$ and N's simulation of M results in acceptance. To see the time taken in this phase, clearly process 3 takes $\leq |x_2|$ steps. Checking $x_1 = y_1$ by process 4 takes $\leq \frac{|x_1|}{d} + O(1)$ steps since x_1 and y_1 are in compressed form. Since the 6-move simulation takes $6t(|x|)/d$ steps, process 4 takes $\frac{|x_1|}{d} + \frac{6t(|x|)}{d} + O(1)$ steps. Thus the final phase takes $1 + \max\{|x_2|, \frac{|x_1|}{d} + \frac{6t(|x|)}{d} + O(1)\}$ steps. (Note that the "1+" is necessary in making the final decision to accept or not.)

We now prove that $L(\mathrm{M}) = L(\mathrm{N})$. If x is not in $L(\mathrm{M})$ then it is easy to see that every computation path of N is non-accepting x because one of the following will fail to hold:

(a) $y_2 = x_2$,

(b) $y_1 = x_1$,

(c) M accepts x.

Conversely, if x is in $L(\mathrm{M})$ then there is a computation path in which (a)-(c) hold and the final phase accepts.

It remains to bound the acceptance time of N. Consider the computation path in which, in addition to (a)-(c), satisfies the following:

(d) $|x_1| = \lceil 2n/d \rceil$ where $n = |x|$.

Process 2 uses $\lceil \frac{n}{d} \rceil$ steps to guess y and $\lceil \frac{|y_2|}{d} \rceil$ steps to position the heads $k+2$ and $k+3$. When we specify d below, we will ensure that $d \geq 3$. Hence for n large enough, the initial phase takes $\leq \max\{|x_1|, \frac{n+|y_2|}{d} + O(1)\} = |x_1|$ steps. Combining this with the time taken by the final phase analyzed earlier, the total time is

$$\begin{aligned} & |x_1| + 1 + \max\{|x_2|, \frac{|x_1|}{d} + \frac{6t(|x|)}{d} + O(1)\} \\ \leq \; & 1 + \max\{n, \frac{(d+1)|x_1|}{d} + \frac{6t(n)}{d} + O(1)\} \\ \leq \; & 1 + \max\{n, \frac{11t(n)}{d}\} (\text{since } t(n) \geq n). \end{aligned}$$

The last expression is $\leq \max\{n+1, ct(n)\}$ if we choose $d > 11/c$. Hence d can be any integer greater than $\max\{3, 11/c\}$. □

An interesting consequence of the last result is:

COROLLARY 6. *NTIME*$(n+1)$ =*NTIME*$(O(n))$.

A Turing acceptor that accepts in time $n+1$ is said to be *real-time*.[8] Thus the real-time nondeterministic languages coincide with those *linear time* nondeterministic languages. This stands in contrast to the fact that *DTIME*$(n+1)$ is a proper subset of *DTIME*$(O(n))$.

Before concluding this section, we must make a brief remark about 'linear reduction of reversals'. In some sense, reversal is a more powerful resource than either space or time, since for every $k > 0$ there are non-regular languages

[8] In the literature *NTIME*$(n+1)$ is also denoted by Q standing for 'quasi-realtime'.

that can be accepted using k reversals but that cannot be accepted using $k-1$ reversals [13]. In contrast, only regular languages can be accepted using a constant amount of space or time. Indeed, in section 6, the power of reversals is even more dramatically shown when we prove that all *RE* languages can be accepted with just two reversals when we allow nondeterminism.

For simple Turing machines, reversal complexity seems to behave more to our expectation (based on our experience with time and space): Hartmanis[12] has shown that only regular languages can be accepted with $O(1)$ reversals. Furthermore, Fischer [11] shows there is a linear speedup of reversals in this model. However, it is not clear that a linear speedup for reversals in the multi-tape machine model is possible in general. It is impossible for small complexity functions because of the previous remark about languages accepted in k but not in $k-1$ reversals.

EXERCISES

Exercise 2.4.0.1: The 6-step (simulation) phase in the linear speedup theorem is capable of improvement. But what constitute an improvement must be clearly stated. Let $c \in \mathbb{N}$ be any constant. Our goal is to construct a k-tape N to simulate c steps of a given k-tape TM M. The key ground rule is that each tape symbol of N must represent $\leq c$ symbols of M. Call symbols and cells of N "supersymbols" and "supercells".
(i) Describe a 4-step phase in case M is a 1-tape TM.
(ii) Describe a 5-step phase in the general case.
(iii) Describe a 4-step phase in the general case. Of course, such a phase subsumes parts (i) and (ii). HINT: we can relax our 6-step phase in two ways. Focusing as usual on only one tape, let us call the "supercell" containing the current head position of M the "active supercell". First, the current head position of N need not be at the active supercell, but is next to it. Our original simulation maintains a direct correspondence between the contents of M's tapes and N's tapes. But now, the current supercell (and only this supercell) is allowed not to contain the correct supersymbol required by this correspondence. The idea is that it will be updated in our next move (thus allowing a slight delay).

□

Exercise 2.4.0.2: Let us change one of the ground rules in the previous exercise: suppose a supersymbol of N is allowed to store $2c$ symbols of M, but we still want to simulate c steps of N in one phase. □

END EXERCISES

2.5 Tape Reduction

Theorem 1 shows that with respect to space, a Turing acceptor may as well use one work-tape. This section has three similar results on reducing the number of work-tapes. The first is a simple result which reduces k work-tapes to one work-tape, at the cost of increasing time quadratically. The second is a classic simulation of a k-tape machine by a 2-tape machine due to Hennie and Stearns [18], in which the simulation is slower by a logarithmic factor. Although these results hold for nondeterministic as well as deterministic machines, our third result shows that we can do much better with nondeterministic simulation. This is the result of Book, Greibach and Wegbreit [5] showing that a nondeterministic 2-tape machine can simulate a k-tape machine without increasing the time, space or reversals by more than a constant factor. Finally, a tape-reduction result for deterministic reversals is indicated at the end of section 8.

Such tape reduction theorems have important applications in later chapters. Essentially they imply the existence of 'efficient universal' Turing machines (chapter 4), which in turn are used to prove the existence of complete languages (chapter 5) and in obtaining hierarchy theorems (chapter 6).

We now state the first result.

THEOREM 7. *If L is accepted by a multitape Turing machine* M *within time t then it is accepted by a 1-tape Turing machine* N *within time $O_{\mathrm{M}}(t^2)$.* N *is deterministic if* M *is deterministic. Moreover, the space (resp., reversal) used by* N *is bounded by a constant times the space (resp., $O(t)$) used by* M.

Proof. Assume M has k work-tapes. We can use the 1-tape machine N described in the proof of the space compression theorem, where each symbol of N is a viewed as $k \times d$ matrix of M's tape symbols. For the present proof, d may be taken to be 1. Note that the size of the non-blank portion of N's work-tape is $\leq i$ after the ith step. Hence the ith step can be simulated in $O_{\mathrm{M}}(i)$ steps of N. So to simulate the first t steps of M requires $\sum_{i=1}^{t} O_{\mathrm{M}}(i)$ $= O_{\mathrm{M}}(t^2)$ steps of N. The claim about space and reversal usage of N is immediate. □

We know that the above simulation is essentially the best possible in the deterministic case: Maass [27] shows that there are languages which require time $\Omega(n^2)$ on a 1-tape machine but can be accepted in real time by a 2-tape machine. See also [25].

THEOREM 8 ((Hennie-Stearns)). *If L is accepted by a k-tape Turing machine in time-space $(t(n), s(n))$ then it is accepted by a 2-tape Turing machine within time-space $O(t \log t, s)$.*

Proof. We use an ingenious encoding of the work-tapes of a k-tape machine M using only tape 1 of a 2-tape machine N. The other work-tape of N is used as a scratch-tape. Tape 1 of N has $2k$ tracks, two tracks for each tape of M. Each cell of N is a super symbol viewed as a column containing $2k$ symbols (possibly blanks) of M. The cells of N are grouped into 'blocks' labeled by the integers:

$$\cdots, B_{-2}, B_{-1}, B_0, B_1, B_2, B_3, \cdots$$

where B_0 consists of just a single cell (call it cell 0) of N. We will number the individual cells by the integers also, with the initial head position at cell 0. For $j > 0$, B_j (respectively, B_{-j}) consists of cells of N in the range

$$[2^{j-1}, 2^j) = \{2^{j-1}, 2^{j-1} + 1, \ldots, 2^j - 1\}$$

(respectively, $(-2^j, -2^{j-1}]$). Thus blocks B_j and B_{-j} each has 2^{j-1} cells, $j > 0$. We assume the boundaries of these blocks are marked in easily detectable ways as cells are visited, the details being left to the reader. The key idea is that, instead of having the k simulated heads of M at different parts of N's tracks, we constrain them to always be at cell 0. The contents in M's tapes are translated laterally to allow this. Essentially this means that when head i of M moves leftward (say) the entire contents of tape i must be shifted rightward so that the currently scanned symbol of tape i is still in cell 0 of N's tape 1. This may appear to require expensive movement of data, but by a suitable scheme of 'delayed' data movement, we show that there is only a small time penalty.

It is enough for us to explain what occurs in two of the tracks that represent some work-tape T of M. These two tracks are designated *upper* and *lower*, and let B_j^U and B_j^L denote the restriction of B_j to the upper and lower tracks, respectively. Inductively, there is an integer $i_0 \geq 0$ that is non-decreasing with time such that a block B_i is 'active' iff $|i| \leq i_0$. We initialize $i_0 = 0$. The non-blank portion of tape T is contained in the active blocks. We hold the following property to be true of each active block B_i:

(1) Each B_i^X ($X = U$ or L) is either full (*i.e.*, represents 2^{i-1} symbols of M, including blanks) or empty (*i.e.*, does not represent symbols of M, not even blanks). The contents of a full block represent contiguous symbols of T. Note that B_i^X is regarded as full even when it is entirely filled with blanks of M.

(2) Exactly two of B_i^U, B_i^L, B_{-i}^U and B_{-i}^L are full. In particular, both B_0^U and B_0^L are always full (it is easy to initialize this condition at the beginning of the computation). In addition, if B_i^L is empty, so is B_i^U. But we allow B_i^U to be empty when B_i^L is full.

(3) If $-i_0 \leq i < j \leq i_0$ then B_i represents contents of T that are to the left (in their natural ordering on tape T) of those represented in B_j. Furthermore, if both B_i^L and B_i^U are full then B_i^L represents contents of T that are to the left of those in B_i^U.

Let us now see how to simulate one step of M on T. Suppose that the head of T moves outside the range of the two symbols stored in B_0. Say, the next symbol to be scanned lies to the right and $i \leq i_0$ is the smallest positive integer such that B_i^L is full (recall that if B_i^L is empty then B_i^U is also empty). [If no such i exists, we may increment i_0, fill both tracks of B_{i_0} with blanks of M, and make both tracks of B_{-i_0} empty. Now we may choose i to be i_0.] Clearly, the upper track of B_{-i}^U is empty, and both tracks of B_j for $j = -i+1, -i+2, \ldots, -1, 0$ are full. We copy the contents of these B_j's (there are exactly 2^i symbols) onto the scratch-tape in the correct sequential order (*i.e.*, as they would appear in tape T). There are two cases:

(i) B_{-i}^L is empty. Then transcribe the 2^i symbols in the scratch-tape to the lower tracks of $B_{-i}, B_{-i+1}, \ldots, B_{-1}, B_0$.

(ii) B_{-i}^L is full. Then we transcribe the contents of the scratch tape to the upper track of B_{-i} and to the lower tracks of $B_{-i+1}, \ldots, B_{-1}, B_0$.

Observe that there is exactly enough space in cases (i) and (ii). Now copy the contents of B_i^L to tape 2 and transcribe them onto the lower tracks of $B_1, \ldots, B_{i-1}$ and the upper track of B_0: again there is exactly enough space. (If B_i^U is full we then move it to the lower track.) We call the preceding computations an *order $|i|$ operation.* Note that we have now returned to our inductive hypothesis and the scanned symbol of T is in B_0 as desired.

Clearly a order $|i|$ operation takes $O(2^{|i|})$ time. Repeating this for each of the k tapes of M, we have completed the simulation of one step of M.

To analyze the cost of this simulation, note that an order $i \geq 1$ operation can only occur if the lower track lying strictly between B_0 and B_{-i} or between B_0 and B_i is empty. Also, immediately following such an operation, the cells in the lower track between B_{-i} and B_i are full. This implies that M must make at least 2^{i-1} moves between two consecutive order i operations. It is also clear that the first order i operation cannot occur until M has made at least 2^{i-1} moves. Suppose M makes t moves. Then the largest value of i such that some order i operation is made is $m_0 = 1 + \lfloor \log_2 t \rfloor$. Therefore the number of order i operations is $\leq \frac{t}{2^{i-1}}$ and the total number of moves made by N is

$$\sum_{i=1}^{m_0} \frac{t \cdot O(2^i)}{2^{i-1}} = O(t \log t).$$

□

In section 8, we will prove a tape reduction result for deterministic reversal-bounded machines. The result shows that a deterministic k-tape machine that accepts in $r(n)$ reversals can be simulated by a deterministic 2-tape machine in $O(r^2)$ reversals. The next result, from Book, Greibach and Wegbreit, shows that this quadratic blow-up can be avoided when we use nondeterministic simulation.

First we need some definitions. Let (Σ, L) be accepted by a nondeterministic k-tape machine M. Let Δ be the tape alphabet of M, Q the states of M, and suppose that the δ(M) contains $p \geq 1$ tuples which we number from 1 to p. Let

$$\Gamma = \Sigma \times Q \times \Delta^k \times \{0, 1, \ldots, p\}$$

be a new alphabet with super symbols which may be regarded as $(k+3)$-tuples as indicated. Define a *trace* of a configuration C as a symbol of Γ of the form

$$b = \langle a, q, c_1, \ldots, c_k, \beta \rangle \tag{4}$$

where a is the currently scanned input symbol in C, q is the state in C, each c_j is currently scanned symbol in the work tape j, and β is the number of any tuple in δ(M) which is applicable to C. If there are no applicable tuples then $\beta = 0$. The trace of a computation path $\bar{C} = (C_0, \ldots, C_m)$ is the word $b_0 b_1 \cdots b_m$ where $b_i \in \Gamma$ is a trace of C_i, and for each $j < m$, the tuple number in b_j (*i.e.*, the last component of b_j) identifies the instruction of M causing the transition $C_j \vdash C_{j+1}$. The trace of $\bar{C}$ is unique. We are now ready for the proof.

THEOREM 9 ((Tape Reduction for Nondeterministic Machines)). *Let $t(n) \geq n$ and* M *be a nondeterministic machine accepting in simultaneous time-reversal bound of (t, r). Then there is a 2-tape nondeterministic machine* N *accepting* L(M) *in simultaneous time-reversal bound of $O(t, r)$.*

Proof. Let M be as in the theorem. With Δ and Γ as above, we construct a 2-tape N that operates as follows: on input x, N guesses in tape 1 a word $w = b_0 \cdots b_m \in \Gamma^*$ that is intended to be the trace of an accepting computation of M on x. It remains to show how to verify if w is indeed such a trace. We may assume that w is generated in such a way that the transitions from b_i to b_{i+1} are plausible, e.g., if the instruction (corresponding to the tuple number) in b_i causes a transition to a new state q then b_{i+1} has state q. Furthermore, the state in b_m is accepting. It remains to check that the symbol scanned under each head is the correct one.

For each $j = 1, \ldots, k$, we will check the symbols under head j is correct. To do this for a particular j, first N re-positions head 1 at the beginning of w. This takes $O(m)$ steps. Then N scans the word w from left to right, carrying out the actions specified each super symbol b_i $(i = 0, \ldots, m)$ using its own input head to move as directed by the instructions in b_i, and using tape 2 to act as tape j of M. If N discovers that the symbol under its own input head does not agree with the symbol indicated in b_i, or the symbol scanned on its tape 2 does not agree with the symbol indicated in b_i, then N rejects at once. Otherwise, it proceeds to check the symbols for head $j+1$. When all tapes are verified in this manner, then N accepts.

Note that the checking of tape contents is entirely deterministic. We now prove that this N accepts if and only in M accepts. If M accepts, then clearly N accepts. Suppose N accepts with a particular word $w = b_0 \cdots b_m \in \Gamma^*$. We can construct inductively for each $i \geq 0$, a unique sub-computation path $C_0, \ldots, C_i$ such that each b_i is a trace of C_i and C_0 is the initial configuration on input x. In particular, there is an accepting computation path of length m. Finally, we note that the time used by N is $O(m)$; moreover, if the accepting computation makes r reversals, then N makes $r + O(1)$ reversals. □

Note that we can apply the nondeterministic linear speedup result to N to reduce the time complexity from $O(t)$ to $max\{n+1, t(n)\} = t(n)$, but the number of work-tape would increase to 5. But if $t(n)$ is sufficiently fast growing, we can first apply speedup result to M then apply the above construction to achieve 2-tape machine with time complexity $t(n)$ rather than $O(t)$.

The cited paper of Maass also shows a language accepted by a real-time deterministic 2-tape machine that requires $\Omega(n^2/\log^5 n)$ time on any 1-tape nondeterministic machine. This shows that the Book-Greibach-Wegbreit result cannot be improved to 1-tape.

Remark: There are complexity differences between k and $k+1$ tapes for all k when we restrict attention to real-time computations: more precisely, there are languages accepted in real-time by some $(k+1)$-tape acceptor but not by any real-time k-tape acceptor. Rabin [32] proved this for $k=1$ and Aanderaa [1] showed this in general. See also [31].

2.6 Simulation by Time

In this section, we will bound space and reversal resources in terms of time. The results are of the form

$$X\text{-}TIME\text{-}SPACE\text{-}REVERSAL(t,s,r) \subseteq Y\text{-}TIME(t')$$

where $X, Y \in \{D, N\}$. Generally, we show this by showing how to simulate a X-mode machine M using time-space-reversal (t,s,r) using a Y-mode machine N in time t'. We then say we have a 'simulation by time'. We think of simulation by time as an attempt to minimize time without consideration of other resources. In the next two sections, we obtain similar results on simulations by space and by reversal. The importance of these three sections taken together is how they exemplify the vastly different computational properties of each of these resources.

THEOREM 10. *Let t be a complexity function, $X = D$ or N. Then*

$$XTIME(t) \subseteq X\text{-}TIME\text{-}SPACE\text{-}REVERSAL(t, O(t), O(t)).$$

Proof. Observe that if a k-tape M accepts in time t then it clearly accepts in space kt and in reversal $(k+1)t$. □

In particular, $XTIME(f) \subseteq XSPACE(O(f)) = XSPACE(f)$, by the space compression theorem.

THEOREM 11. *Let* M *be a deterministic or nondeterministic acceptor. For all complexity functions r, s, if* M *accepts in space-reversal (s,r) then*
a) each phase of M *has length $O_{\mathrm{M}}(n+s)$, and*
b) M *accepts in time $O_{\mathrm{M}}((n+s)\cdot r)$.*

Proof. A computation path $\overline{C}$ of M can be uniquely divided into phases which correspond to the time periods between reversals. Choose $\overline{C}$ so that no configuration is repeated in $\overline{C}$. There is a constant $O_{\mathrm{M}}(1)$ such that if all tape heads of M remain stationary for $O_{\mathrm{M}}(1)$ consecutive steps then some configuration would be repeated; by choice of $\overline{C}$, some tape head must move in any consecutive $O_{\mathrm{M}}(1)$ steps. There are at most n head motions that is attributable to the input head – the rest is clearly bounded by s. Hence if the simultaneous space-reversal of $\overline{C}$ is (s,r), then each phase has $O_{\mathrm{M}}(n+s)$ steps. This proves (a). Since there are r phases, the time of $\overline{C}$ is $O_{\mathrm{M}}((s+n)r)$, proving (b). □

COROLLARY 12. *or $X = D, N$,*

$$X\text{-}SPACE\text{-}REVERSAL(s,r) \subseteq XTIME(O((n+s)\cdot r)).$$

LEMMA 13. *Let* M *be fixed k-tape acceptor, x be an input and $h > 0$ any integer. Let*

$$CONFIGS_h(x) = \{\langle q, w_i, n_i\rangle_{i=0}^{k} : w_0 = x, \sum_{j=1}^{k} |w_j| \leq h\}$$

be the set of configurations of M *with input x using at most space h. Then*

$$|CONFIGS_h(x)| = n \cdot O_{\mathrm{M}}(1)^h$$

where $n = |x| + 1$.

Proof. There are at most $n+1$ positions for head 0, and at most $(h+2)$ positions for each of heads $1, \ldots, k$, in a configuration of $CONFIGS_h(x)$. There are at most $\binom{h+k-1}{k-1}$ ways to distribute h cells among k tapes. These h cells can be filled with tape symbols in at most d^h ways where d is the number of tape symbols of M, including with the blank symbol. Let q be the number of states in M. Thus:

$$|CONFIGS_h(x)| \leq q \cdot (n+2) \cdot (h+2)^k \cdot \binom{h+k-1}{k-1} \cdot d^h = n \cdot O(1)^h.$$

□

THEOREM 14. *If* M *is a Turing machine that accepts in space* s *then* M *accepts in time* $n \cdot O_{\mathrm{M}}(1)^{s(n)}$. M *can be deterministic or nondeterministic.*

Proof. Let $\overline{C}$ be a shortest accepting computation path on input x using space $s(|x|)$. Then the configurations in $\overline{C}$ are all distinct and by the previous lemma, the length of $\overline{C}$ is upper bounded by $n \cdot O_{\mathrm{M}}(1)^{s(n)}$. □

COROLLARY 15. *ith* $X = D, N$, *we have*

$$XSPACE(s) \subseteq X\text{-}TIME\text{-}SPACE(n \cdot O(1)^{s(n)}, s(n)).$$

THEOREM 16. $NSPACE(s) \subseteq DTIME(n \cdot O(1)^{s(n)})$.

Proof. Let M be a nondeterministic k-tape machine accepting in space s. The theorem is proved by showing a deterministic N that accepts the same language $L(\mathrm{M})$ in time $n \cdot O(1)^{s(n)}$. It is instructive to first describe a straightforward simulation that achieves $O(n^2 \log n \cdot O(1)^{s(n)})$.

On input x, N computes in *stages*. In *stage* h (for $h = 1, 2, \ldots$), N has in tape 1 exactly h cells marked out; the marked cells will help us construct configurations using at most h space. Then N lists on tape 2 all configurations in $CONFIGS_h(x)$. Since $|CONFIGS_h(x)| = n \cdot O(1)^h$, and each configuration needs $h + \log n$ space (assuming the input head position is in binary and input word is not encoded with a configuration), we use a total of $n \log n \cdot O(1)^h$ space. On a separate track below each configuration in tape 2, we will mark each configuration as having one of three possible statuses: *unseen*, *seen* and *visited*. Initially, all configurations are marked 'unseen', except for the initial configuration which is marked 'seen'. Now N performs a sequence of $\leq |CONFIGS_h(x)|$ *sweeps* of the input data, where each sweep corresponds to *visiting* a particular 'seen' configuration; after the sweep, that 'seen' configuration will have a 'visited' status. To pick a configuration for sweeping, N picks out the leftmost 'seen' configuration C on tape 2. When visiting configuration C, N first generates all configurations derivable from C, say C_1 and C_2, and puts them on two different tapes. Then N goes through tape 2 again, this time to locate the occurrences of C_1 and C_2 in tape 2. This can be done in time $n \log n \cdot O(1)^h$. If C_i is located in tape 2, we mark it 'seen' if it is currently 'unseen'; otherwise do nothing. Then we mark C itself as 'visited'. If we ever mark an accepting configuration as 'seen', we can accept at once; otherwise, when there are no more 'seen' configurations, we proceed to stage $h + 1$ (thus we will not halt if the input is not accepted). The total time in stage h is seen to be $n^2 \log n \cdot O(1)^h$. The correctness of this simulation follows from the fact that the input x is accepted by M iff N accepts at some stage h, $h \leq s(|x|)$. The total time to accept is

$$\sum_{h=1}^{s(n)} n^2 \log n \cdot O_1(1)^h = n^2 \log n \cdot O_2(1)^{s(n)}.$$

We now improve on the $O(n^2 \log n)$ factor. The idea [7] is to represent the input head positions implicitly. Let us call a **storage configuration** to be a 'standard' configuration except that information about the input tape (*i.e.*, the input word and the input head position) is omitted. Thus, for a configuration that uses h space, the corresponding storage configuration is just the current state, the contents of the work tapes and the positions of the work heads. This can be stored in $O(h)$ space. In stage h, we first compute on tape 2 a list L of all the storage configurations using space h. This list can be represented by a string of length $h \cdot O(1)^h = O(1)^h$. Then we duplicate L $n + 2$ times. So tape 2 contains the string

$$T = \underbrace{L\#L\#\cdots\#L}_{n+2}.$$

This string represents every 'standard' configuration C as follows. If the input head of C is at position i ($i = 0, 1, \ldots, n+1$) and its storage configuration is S then C is represented by the $i+1$st occurrence of S in T. Again, we assume that the all the configurations in T are initially 'unseen' except the initial configuration is 'seen'. As long as T has any 'seen' configuration, we do repeated **sweeps** as before to 'visit' the leftmost 'seen' configuration. To visit C, we must mark as 'seen' any successor C' of C that is still 'unseen'. Note that there are $O(1)$ such successors and these are within $O(1)^h$ distance from C. Hence the marking process takes only $O(1)^h$ time.

To find the next 'seen' configuration to visit, we would also like to spend $O(1)^h$ steps in finding it. It turns out that this can only be done in the amortized sense. The idea is to keep track of the leftmost and rightmost 'seen' configuration: for each 'seen' configuration, we can store two flags, **leftmost** and **rightmost**, which will be set to true or false according as the configuration is leftmost and/or rightmost. Note that if the next leftmost 'seen' configuration lies to the left of the configuration being visited, then we can find it in $O(1)^h$ steps. This is a consequence of our policy to always visit the leftmost unseen configuration first. But if it lies to the right, there is

basically no bound (*i.e.*, the maximum number $n \cdot O(1)^h$ of steps may be needed). However, we shall charge the cost to traverse an entire block L of configurations to the last visited configuration on that block. Then it is not hard to see that each visited configuration is charged at most once (assuming as we may, that each configuration has at most two successors). Thus the charge of $O(1)^h$ per visited configuration is maintained. The total charges over all visited configurations is then $n \cdot O(1)^h$.

There is one other detail to mention: since the input head position in a configuration of T is implicit, we need to know the current input symbol for any configuration that we visit in T. This is easy to take care of by using the input head of the simulator to track the position of head 2. When head 2 is inside the ith copy of L in T, we assume that the simulator is scanning the ith cell in its input tape. This concludes our description. There are $n \cdot O(1)^h$ sweeps and each sweep takes $O(1)^h$ time, the total time of stage h is $n \cdot O(1)^h$. This implies the stated bound in our theorem. □

There is another basic method of doing a time simulation of nondeterministic space. First we review the connection between Boolean matrices and digraphs. A digraph G on n vertices can be represented by its adjacency matrix A where A is an $n \times n$ matrix with its (i,j)th entry $A_{i,j} = 1$ if and only if there is an edge from the vertex i to vertex j. We assume that $A_{i,i} = 1$ for all i. Recall that the product of two Boolean matrices is defined as in ordinary matrix multiplication except that addition becomes logical-or '$\vee$' and multiplication becomes logical-and '$\wedge$'. It is easy to see that if $B = A^2$ (A squared) then $B_{i,j} = 1$ if and only if there is a vertex k such that $A_{i,k} \wedge A_{k,j} = 1$, i.e. there is a path of length at most 2 from i to j. Arguing the same way, with $C = B^2 = A^4$, we see that $C_{i,j} = 1$ if and only if there is a path of length at most 4 from i to j. Therefore we see that A^* given by

$$A^* := A^{2^k}$$

(where k is the least integer such that $2^k \geq n$) then $A^*_{i,j} = 1$ if and only if there is a path from i to j. The matrix A^* is called the *transitive closure* of A. There is a well-known method attributed to Warshall and to Floyd for computing the transitive closure A^*. For $s = 0, 1, \ldots, n$, let $A^{(s)}$ be the matrix defined as follows: $A^{(s)}_{i,j} = 1$ if and only if there is a path from i to j such that the intermediate nodes (excluding i and j) lie in the set $\{1, \ldots, s\}$. Hence $A^* = A^{(n)}$. Clearly $A^{(0)} = A$. It is also easy to see that $A^{(s+1)}$ is obtained as follows:

$$A^{(s+1)}_{i,j} = A^{(s)}_{i,j} \vee \left(A^{(s)}_{i,s+1} \wedge A^{(s)}_{s+1,j} \right).$$

It is easy to see that this algorithm has complexity $O(n^3)$ on a random-access machine. It turns out (Exercise) that we can accomplish this on a Turing machine also, assuming that the matrix A is given in row-major order. We exploit this connection between paths and transitive closure. In particular, if we regard configurations as nodes of a graph, and edges correspond to the $\vdash$ relation, then deciding if a word is accepted amounts to checking for a path from the start configuration to any final configuration. We leave the details as an exercise. This alternative method is slightly less efficient with complexity $(nO(1)^{s(n)})^3 = n^3 O(1)^{s(n)}$. We will encounter the transitive closure method again.

A consequence of this theorem are the following two inclusions among the canonical classes:

$$NLOG \subseteq P\,, \qquad PSPACE \subseteq DEXPTIME. \tag{5}$$

Although theorem 16 remains valid if we replace *NSPACE* with *NTIME*, the next theorem will show something slightly stronger. But first we note a fact similar to lemma 13.

LEMMA 17. *Let* M *have k-tapes and x an input word. For any $h \geq 1$, let $CONFIGS'_h(x)$ denote the set of configurations of M that can be reached from the initial configuration on input x using at most h steps. Then*

$$|CONFIGS'_h(x)| = O_{\mathrm{M}}(1)^h.$$

Proof. As before, if q is the number of states in M and d is the number of tape symbols of M, plus one, then $|CONFIGS'_h(x)| \leq q \cdot (h+2)^{k+1} \cdot d^h = O(1)^h$. The difference is that now, the input head position n_0 satisfies $n_0 \leq h + 1$. □

Thus $|CONFIGS'_h(x)|$ does not depend on $n = |x|$.

THEOREM 18. $NTIME(t) \subseteq DTIME(O(1)^t)$.

Proof. The proof proceeds as in theorem 16 but we now enumerate over configurations in $CONFIGS'_h(x)$ rather than $CONFIGS_h(x)$. □

Although theorems 16 and 18 uses accepting complexity, the results remain valid for running complexity (Exercise). The next result shows that reversal in the fundamental mode is bounded by a single exponential time.

THEOREM 19. *Let* M *be a deterministic Turing machine accepting in* $r(n)$ *reversals. Then* $AcceptTime_{\mathrm{M}}(n) = n \cdot O_{\mathrm{M}}(1)^{r(n)}$ *(assuming the left-hand side is defined).*

Proof. Let M be a k-tape machine and $\overline{C}$ an accepting computation path on an input of length $n > 0$. The theorem follows if we show that the length of $\overline{C}$ is $n \cdot O_{\mathrm{M}}(1)^{r(n)}$. Call a step $C_{j-1} \vdash C_j$ *active* if any work head moves during this step from a non-blank symbol; otherwise the step is *inactive*. Thus, during an inactive step, the input head can freely move and the work heads may move from a blank symbol. We obtain a bound on the number of active steps. By assumption, $\overline{C}$ has $r \leq r(n)$ phases. Let m_i be the number of active steps in phase i $(i = 1, \ldots, r)$.

For any configuration C_j, recall that each head h $(h = 0, \ldots, k)$ tends in some direction $d_h \in \{-1, 0, +1\}$. The *potential* of tape h in C_j is the number of non-blank cells that would be encountered as head h moves in the direction d_h, starting from its current cell in C_j. If $d_h = 0$ the potential is defined to be zero. For example, if head h is at the rightmost non-blank symbol, and $d_h = +1$ then the potential is 1; but if $d_h = -1$ then the potential would be the total number of non-blank cells on tape h. The *potential* of C_j is the sum of the potentials of all work tapes (so we discount the potential of the input tape). Clearly the potential of C_j is at most the space usage in C_j. Therefore, at the start of phase i the potential is at most

$$k \sum_{v=1}^{i-1} m_v$$

since the number of active steps until that moment is $\sum_{v=1}^{i-1} m_v$ and each active step can create at most k new non-blank cells, where k is the number of work heads. Suppose inductively that $m_v \leq cn(k+1)^v$ for some $c > 0$. This is true for $v = 1$. Now each active step in a phase consumes at least one unit of potential, so m_i is bounded by the potential at the beginning of the ith phase:

$$m_i \leq k \sum_{v=1}^{i-1} m_v \leq k \sum_{v=1}^{i-1} cn(k+1)^v < cn(k+1)^i.$$

Therefore in r phases, the total number of active steps is

$$\sum_{v=1}^{r} m_v < cn(k+1)^{r+1} = n \cdot O_{\mathrm{M}}(1)^r.$$

It remains to bound the number of inactive moves. There are $O_{\mathrm{M}}(1)$ consecutive steps in which no head moves, and there are at most n moves by the input head in any phase. Hence there are at most $n + O_{\mathrm{M}}(m_i)$ inactive moves in the ith phase. Summing over all phases gives at most $nr + O_{\mathrm{M}}(1)^r$ inactive moves. □

COROLLARY 20. $DREVERSAL(r(n)) \subseteq D\text{-}TIME\text{-}REVERSAL(n \cdot O(1)^{r(n)}, r(n))$

Remark: If reversals by input heads are not counted, then we get instead

$$DREVERSAL(r) \subseteq D\text{-}TIME\text{-}REVERSAL(n^2 O(1)^{r(n)}, r(n)).$$

2.7 Simulation by Space

We will present two techniques for simulating other resources so as to minimize deterministic space. These space simulation techniques are quite different than the time simulation techniques of the previous section.

THEOREM 21 ((Extended Savitch's theorem)). *If* $t(n) \geq 2s(n) \geq \log n$ *then*

$$N\text{-}TIME\text{-}SPACE(t(n), s(n)) \subseteq DSPACE(s(n) \log \left(\frac{t(n)}{s(n)} \right)).$$

Proof. Given a nondeterministic machine M accepting in simultaneous time-space (t, s), it suffices to show that $L(\mathrm{M})$ is accepted in space $s \log(t/s)$ by some deterministic N. (We require $t(n) \geq 2s(n)$ simply to prevent $\log(t/s)$ from vanishing.) An input x of length n is accepted by M iff there is an accepting computation path $\overline{C}$ whose time-space is $(t(n), s(n))$. As in lemma 13, for any $h \geq 1$, let $CONFIGS_h(x)$ be the set of configurations of M on input

x using space at most h. For configurations $C, C' \in CONFIGS_h(x)$, define the predicate $REACHABLE_h(C, C', m)$ to be true if there is a sub-computation path from C to C' of length $\leq m$.

Claim A: Assume that the input x is freely available, we can evaluate $REACHABLE_h(C, C', m)$ in deterministic space $O(m + h + \log n)$. *Proof of claim.* Consider the tree $T(C)$ rooted at C and whose paths comprise all those sub-computation paths from C of length $\leq m$ and involving configurations in $CONFIGS_h(x)$. It suffices to search for C' in $T(C)$ using a standard depth-first search. At any moment during the search, we are at some node labeled C'' in $T(C)$. We first check if C'' is the C' we are looking for. If not, we try to visit an unvisited successor of C'' in $CONFIGS_h(x)$, provided C'' is at depth less than m. If this is not possible, we backtrack up the tree. To backtrack, we must be able to generate the parent C''' of C'': this is easy if we had stored the instruction of M that transformed C''' into C''. The space to store this backtracking information is just $O(1)$ per level or $O(m)$ overall. We also need $O(h + \log n)$ to store the current node C'' and the target configuration C'. When we have searched the entire tree without finding C', we conclude that $REACHABLE_h(C, C', m)$ is false. This proves our claim.

Claim B: Assume $h \geq \log n$ and that the input x is freely available, we can evaluate $REACHABLE_h(C, C', m)$ in in deterministic space $O(h \log(2 + \frac{m}{h})$. *Proof of claim.* We use a following simple observation:

> (*) $REACHABLE_h(C, C', m)$ holds iff both $REACHABLE_h(C, C'', \lfloor m/2 \rfloor)$ and $REACHABLE_h(C'', C', \lceil m/2 \rceil)$ hold, for some $C'' \in CONFIGS_h(x)$.

Thus we can evaluate $REACHABLE_h(C, C', m)$ recursively: if $m \leq h$, we use claim A to directly evaluate $REACHABLE_h(C, C', m)$ in space $O(h)$, as desired. Otherwise, if $m > h$, we use observation (*) to make two recursive calls to the predicate. Here are the details: To evaluate $REACHABLE_h(C, C', m)$, assume tape 1 of the machine N stores the value h, tape 2 stores the current arguments (C, C', m) while tape 3 acts as the recursion stack. We then systematically enumerate all configurations $C'' \in CONFIGS_h(x)$. For each C'', we recursively evaluate $REACHABLE_h(C, C'', m/2)$ and $REACHABLE_h(C'', C', m/2)$ in this order. But before making first recursive call, we write the pair $(C', 0)$ on top of the recursion stack, where the Boolean value 0 indicates that this is the first of the two recursive calls. Then the current arguments on tape 2 are updated to $(C, C'', m/2)$ and we continue recursively from there. Similarly, before making the second recursive call, we write the pair $(C, 1)$ on the recursion stack and update the current arguments on tape 2 to $(C'', C', m/2)$. If either one of these recursive calls returns with "failure" (predicate is false), we generate the next configuration in $CONFIGS_h(x)$ and use it place of C''; in case C'' is the last configuration in our enumeration of $CONFIGS_h(x)$, we return from the current call $REACHABLE_h(C, C', m)$ with "failure". If both recursive calls returns with "success" (predicate is true), we return from the current call with "success".

The recursion depth in evaluating $REACHABLE_h(C, C', m)$ is $\lceil k \rceil = \lceil \log(m/h) \rceil \geq 1$. The space to store the arguments for each recursive call is $O(h + \log n)$ where the $\log n$ comes from having to represent the input head position (the actual input x need not be stored). So the total space used is $O((\log n + h) \log(m/h))$, which is $O(h \log(m/h))$ when $h \geq \log n$.

Finally, that x is accepted by M if and only if $REACHABLE_h(C_0, C_a, m)$ is true for some $h \leq s(n)$, $m \leq t(n)$, where $C_a \in CONFIGS_h(x)$ is an accepting configuration and C_0 the initial configuration on input x. By modifying M, we may assume that C_a is unique in $CONFIGS_h(x)$. Since $s(n)$ and $t(n)$ are unknown, N must search for h and m systematically, using a doubly-nested loop:

```
for p = 1, 2, ...,
    for h = 1, ..., p,
        let m = h2^⌊p/h⌋ and C_0, C_a ∈ CONFIGS_h(x).
        Accept if REACHABLE_h(C_0, C_a, m) is true,
```

The correctness follows from two remarks:
(i) For any given value of p, the maximum space used in the inner loop is $O(p)$. This is because the space to evaluate $REACHABLE_h(C_0, C_a, m)$ is order of $h \lg(m/h) \leq p$.
(ii) If input x is accepted by M then the double-loop accepts for some $p = O(s \log(t/s))$; otherwise, the double-loop runs forever. □

We obtain several interesting corollaries, including the well-known result of Savage [34].

COROLLARY 22.
(i) (Savitch's theorem) *If $s(n) \geq \log n$ then $NSPACE(s) \subseteq DSPACE(s^2)$.*
(ii) If $s(n) \geq n$ then $NTIME(s) \subseteq DSPACE(s)$.
(iii) If $s(n) \geq n$, N-SPACE-REVERSAL$(s, r) \subseteq DSPACE(s \log r)$.

Proof. (i) Use the fact that $NSPACE(s) = N\text{-}TIME\text{-}SPACE(O(1)^s, s)$ for $s(n) \geq \log n$.
(ii) This follows since $NTIME(t) \subseteq N\text{-}TIME\text{-}SPACE(2t, t)$ but $N\text{-}TIME\text{-}SPACE(2t, t) \subseteq DSPACE(t)$ by applying the theorem.
(iii) This follows since $N\text{-}SPACE\text{-}REVERSAL(s, r) \subseteq N\text{-}TIME\text{-}SPACE((n + s)r, s)$ and $N\text{-}TIME\text{-}SPACE((n + s)r, s) \subseteq DSPACE(s \log r)$ by applying the theorem. □

We remark that our formulation of the generalized Savitch theorem is motivated by the desire to combine (i) and (ii) of this corollary into one theorem.

If $s(n) = o(\log n)$, Savitch's theorem yields $NSPACE(s) \subseteq DSPACE(\log^2 n)$. Monien and Sudborough have improved this to

$$NSPACE(s) \subseteq DSPACE(s(n) \log n).$$

In chapter 7, we further improve upon Monien and Sudborough via a strengthening of theorem 21. Indeed, chapter 7 shows four distinct extensions of Savitch's theorem!

COROLLARY 23.

(i) $DSPACE(n^{O(1)}) = NSPACE(n^{O(1)})$.

(ii) $DSPACE(O(1)^n) = NSPACE(O(1)^n)$.

This corollary justifies our notation '*PSPACE*' since we need not distinguish between '*D-PSPACE*' and '*N-PSPACE*'. The notations *EXPS* and *EXPSPACE* are likewise justified.

We next simulate deterministic reversals by deterministic space. Istvan Simon [35] had shown how to simulate simultaneous time-reversal by space. Our next result strengthens his result to space-reversal.

THEOREM 24. *If* $s(n) \geq \log n$,

$$D\text{-}SPACE\text{-}REVERSAL(s, r) \subseteq DSPACE(r \log s).$$

Proof. Given a deterministic k-tape M accepting in space-reversal (s, r), it suffices to construct a deterministic N accepting $L(\mathrm{M})$ in space $r \log s$. For any configuration C in a computation path $\overline{C}$, the *trace* of C refers to the head tendencies in C (relative to $\overline{C}$) and to information that remains in C after we discard all the tape contents except for what is currently being scanned.[9] More precisely, the trace of C is defined as

$$\langle q, b_0, \ldots, b_k, n_0, \ldots, n_k, d_0, \ldots, d_k \rangle$$

where q is the state in C, the b_i's are the scanned symbols, the n_i's are the head positions, and the d_i's are the head tendencies. We remark that the n_i's are absolute positions, unlike our usual convention of using relative positions in configurations. To *simulate the jth step* means to compute the trace of the jth configuration of the computation path. We shall also need the concept of a *partial trace*: this is like a trace except any number of the b_i's can be replaced by a special symbol '$*$'. We say b_i is *unknown* if $b_i = *$. Thus a trace is also a partial trace although for emphasis, we may call a trace a *full* trace.

On tape 1 of N we assume that we had already computed a sequence

$$\tau_1, \tau_2, \ldots, \tau_m$$

where τ_i is the trace of the first configuration in phase i. To begin, we can always place the trace τ_1 of the initial configuration on tape 1. Call the last phase (phase m) represented in this sequence the *current phase*. This means that our simulation has reached some step in the mth phase.

To simulate a step of a phase, it turns out that we need to simulate k steps from previous phases and thus we can use recursion. To keep track of these recursive calls, we use tape 2 of N as a recursion stack. On the recursion stack, we will in general store a sequence of the form

$$\sigma_{i_1}, \sigma_{i_2}, \ldots, \sigma_{i_v} (v \geq 1)$$

where $m = i_1 > i_2 > \cdots > i_v \geq 1$ and each σ_{i_j} is a partial trace of a configuration in the i_jth phase. Furthermore, all the σ_{i_j}'s, except possibly for σ_{i_v}, are not full traces. Note that σ_{i_1} is always in the current phase (phase m). Intuitively, we are really trying to evaluate σ_{i_1} but its evaluation calls for the evaluation of σ_{i_2}, which calls for the evaluation of σ_{i_3}, etc. Thus σ_{i_v} is the top entry of the recursion stack.

[9] Warning: in this book we use the term 'trace' for several similar but non-identical concepts. The reader should check each context.

This is how the recursive calls arise. Suppose that at some point in our simulation we managed to compute the full jth trace τ where the jth configuration belongs to the current phase. Also assume that the recursion stack contains τ as its only entry. We now want to simulate the $(j+1)$st step. Note that we can obtain the partial trace of the $(j+1)$st configuration from the full jth trace: the only unknowns about the $(j+1)$st trace are the symbols under any head that has moved in the transition from the jth configuration to the $(j+1)$st configuration. At this moment, we replace τ on the recursion stack by partial $(j+1)$st trace σ_{i_1} (as the only stack entry).

Now inductively, assume that we have $\sigma_{i_1}, \ldots, \sigma_{i_v}$ $(v \geq 1)$ on the stack. Let

$$\sigma_{i_v} = \langle q, b_0, \ldots, b_k, n_0, \ldots, n_k, d_0, \ldots, d_k \rangle$$

be the partial trace on top of the stack. We may assume that b_0 is known (i.e. not equal to $*$) since this is on the input tape. Let $h \geq 1$ be the smallest index such that $b_h = *$ (if σ_{i_v} is a full trace, then let $h = k+1$). We take following action depending on two cases:

(1) If $h \leq k$ then we want to determine b_h. Note that we know the position $(= n_h)$ of b_h and hence from the contents of tape 1, we can determine the phase preceding phase i_v in which cell n_h was last visited. If this cell had never been visited before, we can also determine this fact and conclude that b_h is the blank symbol $\Box$. Suppose this cell was last visited by configuration C in phase i_{v+1} $(i_v > i_{v+1} \geq 1)$. Note that in the configuration following C, head h must no longer scan cell n_h. Our goal is to compute the trace of C. We say the trace of C is *sought after* by σ_{i_v}. To do this, we copy the trace of the starting configuration of phase i_{v+1} to the top of the stack. This trace, $\sigma_{i_{v+1}}$, is available in tape 1.

(2) If $h = k+1$ then σ_{i_v} is a full trace. There are two cases: (i) If this is the trace sought after by the preceding partial trace $\sigma_{i_{v-1}}$ then we can fill in one of the unknown symbols in $\sigma_{i_{v-1}}$ and erase σ_{i_v} from the top of the stack. (ii) If this is not the trace sought after by the preceding partial trace then we simply replace σ_{i_v} on the stack by its successor partial trace. By definition, if $v = 1$ then case (ii) is applied.

It is clear that we eventually convert the partial trace σ_{i_1} into a full trace, and hence repeat the cycle. We need to consider the action to take when we enter a new phase. Suppose $\tau = \sigma_{i_1}$ had just turned into a full trace. Just before we replace τ on the stack by its successor τ', as described in (2), we check whether τ is the beginning of a new phase, and if so we first store it on tape 1. In fact, the information indicating whether τ is the start of a new phase could already be determined as we initially form the partial trace corresponding to τ.

The total amount of space required on tapes 1 and 2 of N is $O(r \log s)$ since each partial trace uses $O(\log s)$ space. The assumption $s(n) \geq \log n$ is needed to represent the input head positions. $\Box$

Corollary 25.
(i) (Istvan Simon) *D-TIME-REVERSAL*$(t, r) \subseteq$ *DSPACE*$(r \log t)$.
(ii) *DREVERSAL*$(O(1)) \subseteq$ *DSPACE*$(\log n) =$ *DLOG*.
(iii) For $r(n) \geq \log n$, *DREVERSAL*$(r) \subseteq$ *DSPACE*(r^2).

To see (i) use the fact that *D-TIME-REVERSAL*$(t, r) \subseteq$ *D-SPACE-REVERSAL*(t, r). For (ii) and (iii), use the fact that

$$\textit{DREVERSAL}(r) \subseteq \textit{D-TIME-REVERSAL}(n \cdot O(1)^r, r).$$

As application of part (ii) of this corollary, we note that since the palindrome language can be accepted in linear time by an acceptor making three reversals, it can be accepted in log space.

Since time, space and reversals are not independent, it is worth pointing out some conditions under which the theorem *D-SPACE-REVERSAL*$(s, r) \subseteq$ *DSPACE*$(r \log s)$ is stronger than its corollary (i), *D-TIME-REVERSAL*$(t, r) \subseteq$ *DSPACE*$(r \log t)$. For instance, let $s(n) = \theta(n) = O(\log r(n))$. Then corollary (i) implies *D-SPACE-REVERSAL*(s, r) $\subseteq$ *DSPACE*$(r \log n)$ while the theorem yields the stronger *D-SPACE-REVERSAL*$(s, r) \subseteq$ *DSPACE*$(r \log \log n)$.

Remark: Rytter and Chrobak [33] have shown that the the converse of corollary (ii),

$$\textit{DLOG} \subseteq \textit{DREVERSAL}(O(1)),$$

holds, provided we do not count reversals made by the input head. This extra condition of Rytter and Chrobak is essential, as pointed out by Liśkiewicz[10]: Book and Yap [3] proved that all tally languages in *DREVERSAL*$(O(1))$ are regular. On the other hand, Mehlhorn and Alt (see section 11) has given a non-regular tally language. Thus corollary (ii) is a proper inclusion. Nevertheless, we cannot strengthen corollary (ii) to *DREVERSAL*$(O(1)) \subseteq$ *DSPACE*(s) for any $s(n) = o(\log n)$, because the language $\{a^n b^n : n > 0\}$ belongs to *DREVERSAL*$(O(1))$ but not to *DSPACE*(s).

[10] Private communication.

2.8 Simulation by Reversal

We consider simulation techniques that seek to minimize reversals. Reversal complexity seems much more 'powerful' than time and space complexity. For example, to double an arbitrarily large tape segment on a 2-tape Turing machines takes only 2 reversals – yet the time and space charged against this activity would be linear. Our intuition about reversals, relative to time or space, is much less developed. This is illustrated by a somewhat unexpected result from Baker and Book [2] about reversal complexity in the nondeterministic mode:

THEOREM 26. *RE* $\subseteq$ *NREVERSAL*(2).

Proof. Without loss of generality, let L be a recursively enumerable language accepted by a simple Turing machine M. Suppose $x \in L$ and let $C_0, C_1, \ldots, C_k$ be the accepting computation path of M on x. Assume that all the C_i's are encoded by strings of the same length. We construct a two-tape nondeterministic N that accepts L as follows: on input x, N guesses some sequence $C_0\#C_1\#\cdots\#C_k$ on tapes 1 and 2 (so the two tapes have identical contents). Now head 2 reverses until it is at the #-separator between C_{k-1} and C_k; while doing this N can verify if C_k is an accepting configuration. Next, after reversing the direction of head 1, we can easily check that C_{k-1} (on tape 2) and C_k (on tape 1) are consecutive configurations of some computation path. If not, N immediately rejects; in particular, N rejects if $|C_{k-1}| \neq |C_k|$. Without any further head reversals, we can continue in this fashion to check that C_{i-1} (on tape 2) and C_i (on tape 1) are consecutive configurations of some computation path, for $i = k-1, k-2, \ldots, 1$. We must ensure that C_0 is the initial configuration on input x; but it is simple to ensure this (without any additional head reversals) while we are guessing the C_i's. □

This result, combined with the earlier

$$DREVERSAL(r) \subseteq DTIME(nO(1)^r),$$

tells us that there is no fixed complexity function f such that any language accepted nondeterministically in reversal $r(n)$ can be accepted deterministically in reversal $f(r(n))$ (contrast with theorem 18 and Savitch's theorem for time and space). Thus reversal seems to be rather different than time or space. Later in this book, we see that when reversal is simultaneously bounded with either time or space, then the corresponding complexity classes are better behaved.

2.8.1 Basic Techniques for Reversals

In this subsection, we show that reversal complexity in the fundamental mode is also relatively well-behaved. These results are from [9]. First we give a series of technical lemmas.

LEMMA 27 ((Natural Number Generation)). *Given as input any integer $r > 0$ in unary, the string*

$$\bar{0}\#\bar{1}\#\bar{2}\#\cdots\#\overline{2^r - 1}\#$$

can be generated by a 2-tape Turing machine M *making $O(r)$ reversals. Here $\bar{m}$ denotes the binary representation of the integer m of length r, prefixed with 0's if necessary.*

Proof. M first generates the pattern $(0^r\#)^{2^r}$ on tape 1. This can be done within $O(r)$ reversals, using a simple doubling method. Call each portion of the tape 1 between two consecutive # symbols a *segment*. Similarly, M generates a pattern $P_1 = (01)^{2^r}$ on tape 2.

Now M uses r *stages*, making a constant number of reversals per stage, to 'fix' the r bits in each *segment* in tape 1. (The final string $\#\bar{0}\#\cdots\#\overline{2^r - 1}\#$ is going to be the contents of tape 1 at the end of the stages, so 'fixing' a bit means making it the symbol 0 or 1 that should appear finally.) For example, with $r = 3$, the rightmost bit of the $2^3 = 8$ segments alternates between 0 and 1, i.e. has the pattern $P_1 = 10101010$. The next bit has pattern $P_2 = 11001100$, and the final bit has pattern $P_3 = 11110000$.

Suppose that at the beginning of the $(i+1)$st ($i \geq 0$) stage, M has fixed the last i bits for each segment on tape 1, and suppose the pattern

$$P_{i+1} = (0^{2^i}1^{2^i})^{2^{r-i}}$$

is inductively available on tape 2. Here the first bit of a segment refers to its least significant bit. Note that the $(i+1)$st bit of the jth segment is exactly the jth bit of pattern P_{i+1}.

In the $(i+1)$st stage M needs to know which bit in the jth segment is the $(i+1)$st bit. This is solved by placing a special mark on another track below the ith bit of each segment. These marks are easily updated at each stage. Now with only one phase, M can fix the $(i+1)$st bit for each segment of P_r on tape 1.

Using a constant number of sweeps, M can generate the pattern $P_{i+2} = (0^{2^{i+1}}1^{2^{i+1}})^{2^{r-i-1}}$ from P_{i+1}. At the end of the rth stage, the string $\bar{0}\#\bar{1}\#\bar{2}\#\cdots\overline{2^r-1}\#$ is on tape 1 of M. This completes the proof. □

A string of the form

$$x_1\#x_2\#\cdots\#x_n\#, n \geq 1$$

($x_i \in \Sigma^*$, $\# \notin \Sigma$) is called a *list* of n *items* (x_i is the ith item). A list is said to be in *normal form* if n is a power of 2 and each x_i has the same length. The next lemma shows how to convert any list into one in normal form.

LEMMA 28 ((List Normalization)). *There is a 2-tape Turing machine* M *which, given a list* $x_1\#x_2\#\cdots\#x_n\#$ *of* n *items on its input tape, can construct another list* $y_1\#y_2\#\cdots\#y_{2^k}\#$ *in normal form using* $O(\log n)$ *reversals. Therefore,*

(a) $2^{k-1} < n \leq 2^k$ *;*

(b) each y_i $(i = 1, \ldots, 2^k)$ *has length* $\max_{i=1,\ldots,n} |x_i|$*; and*

(c) each y_i $(i = 1, \ldots, n)$ *is obtained by padding* x_i *with a prefix of zeroes and each* y_j $(j = n+1, \ldots, 2^k)$ *is a string of zeroes.*

Proof. First M computes the unary representation $z_0 \in \{0\}^*$ of $\max_{i=1,\ldots,n} |x_i|$ as follows: With $O(1)$ reversals, M initializes tape 1 to have all the odd numbered items from the original list and tape 2 to have all the even numbered items. So tape 1 and 2 contain the lists $x_1\#x_3\#x_5\#\cdots$ and $x_2\#x_4\#x_6\#\cdots$, respectively. In another pass, M compares x_{2i-1} on tape 1 with x_{2i} on tape 2 ($i = 1, 2, \ldots, \lceil n/2 \rceil$), marking the longer of the two words. In $O(1)$ passes, M can produce a new list $z_1\#z_2\#\cdots\#z_{\lceil n/2 \rceil}\#$ of these marked words. M repeats this process: splits the list into two with roughly the same number of items, compares and marks the longer item of each comparison, produces a new list consisting of only the marked items. After $O(\log n)$ reversals, M is left with a list containing only one item. This item has the longest length among the x_i's. It is now easy to construct z_0.

The next goal is to construct a string of the form

$$(z_0\#)^{2^k} \quad (\text{where } k = \lceil \log_2 n \rceil).$$

Suppose we already have $(z_0\#)^{2^i}$. Using $x_1\#x_2\#\cdots\#x_n\#$ and $(z_0\#)^{2^i}$, M can compare n with 2^i: if $n \leq 2^i$ then we are done, otherwise we will construct $(z_0\#)^{2^{i+1}}$ from $(z_0\#)^{2^i}$ using $O(1)$ reversals.

Finally, from $(z_0\#)^{2^k}$, we can easily construct the desired $y_1\#y_2\#\cdots\#y_{2^k}\#$. □

A more complicated problem is computing the transitive closure of a matrix. It turns out that the key to computing transitive closure is a fast matrix transposition algorithm:

LEMMA 29 ((Matrix Transposition)). *There is a 2-tape Turing machine* M *such that, given an* $n \times m$ *matrix* $A = (a_{ij})$*,* M *can compute the transpose* A^T *of* A *using* $O(\log \min\{m, n\})$ *reversals. Here we assume both* A *and* A^T *are stored in row major form.*

Proof. Because of the preceding list normalization lemma, we may assume that A satisfies the property $m = 2^k$ for some integer $k \geq 1$ and each entry of A has the same length. Let us first show how to compute A^T in $O(\log m)$ reversals.

For each $i = 0, 1, \ldots, k$, and for $j = 1, 2, \ldots, 2^i$, let $A_j^{(i)}$ denote the $n \times 2^{k-i}$ matrix consisting of the columns indexed by the numbers

$$((j-1)2^{k-i}+1), ((j-1)2^{k-i}+2), \ldots, (j2^{k-i}).$$

For example, $A_j^{(k)}$ is the jth column of A and for each i,

$$A = A_1^{(i)}|A_2^{(i)}|\cdots|A_{2^i}^{(i)}$$

(where | means concatenation of matrices). An *i-row representation* of A is a string consisting of the row-major forms of $A_1^{(i)}, A_2^{(i)}, \ldots, A_{2^i}^{(i)}$, listed in this order, separated by '$'.

Let $A^{(i)}$ denote the i-row representation of A. The lemma follows if we can show how to obtain $A^{(i+1)}$ from $A^{(i)}$ in $O(1)$ reversals. This is because the input is $A^{(0)}$ (the row major form of A) and the desired output is $A^{(k)}$ (the column major form of A).

In order to do the $A^{(i)} \to A^{(i+1)}$ transformation, we need some auxiliary data. Define the block pattern:

$$P^{(i)} = ((w_0^{2^{k-i-1}}w_1^{2^{k-i-1}}\#)^n\$)^{2^i} (i = 0, \ldots, k-1)$$

where $w_0 = 0^s$ and $w_1 = 1^s$, and s is the length of each entry of A. Each $w_0^{2^{k-i-1}}$ (respectively, $w_1^{2^{k-i-1}}$) 'marks' a left (respectively, right) half of the rows of $A_j^{(i)}$, $j = 1, 2, \ldots, 2^i$. $P^{(i)}$ helps us to 'separate' the rows of each $A_j^{(i)}$ into 'left half' and 'right half'.

We may inductively assume that each tape has two tracks with the following contents: $A^{(i)}$ is on track 1 and $P^{(i)}$ is positioned directly underneath $A^{(i)}$ on track 2. The case $i = 0$ is initialized in $O(1)$ reversals. A copy of $P^{(i+1)}$ can be made on track 2 of tape 2 using $P^{(i)}$. Now it is easy to obtain $A^{(i+1)}$ on track 1 of tape 2.

It is not hard to show that A^T can also be computed in $O(\log n)$ reversals. We leave the details as an exercise. Hence if we first determine the smaller of m and n in $O(1)$ reversals, we can then apply the appropriate matrix transposition algorithm to achieve the $O(\log \min\{m, n\})$ bound. □

LEMMA 30 ((Parallel Copying Lemma)). *There is a 2-tape Turing machine* M *which, given an input* $x = 0^n \# x_1 x_2 \cdots x_m$, *where* $x_i \in \Sigma^*$, $(i = 1, \ldots, m)$, *produces string* $y = x_1^{2^n} x_2^{2^n} \cdots x_m^{2^n}$ *within* $O(n)$ *tape reversals, where we assume that* M *can recognize the boundary between blocks* x_i *and* x_{i+1}, $i = 1, \ldots, m-1$.

Proof. Using $O(n)$ reversals, M can get a string $S = (x_1 x_2 \cdots x_m)^{2^n}$. Notice that $x_1^{2^n} x_2^{2^n} \cdots x_m^{2^n}$ is the transpose of S if we regard S as being a $m \times 2^n$ matrix. □

From the last two important lemmas, we get immediately,

LEMMA 31 ((Matrix Multiplication)). *There is a 2-tape Turing machine* M *such that, given two* $n \times n$ *Boolean matrices* $A = (a_{ij})$ *and* $B = (b_{ij})$ *in row major form,* M *computes their product* $AB = (c_{ij})$ *in* $O(\log n)$ *reversals.*

Proof. By lemma 29, we can obtain the transpose B^T of B within $O(\log n)$ reversals. Let

$$A = (a_{11} \cdots a_{1n} a_{21} \cdots a_{2n} \cdots a_{n1} \cdots a_{nn})$$

and

$$B^T = (b_{11} \cdots b_{n1} b_{12} \cdots b_{n2} \cdots b_{1n} \cdots b_{nn}).$$

By lemma 30, we can get

$$A_1 = (a_{11} \cdots a_{1n})^n (a_{21} \cdots a_{2n})^n \cdots (a_{n1} \cdots a_{nn})^n$$

and

$$B_1^T = (b_{11} \cdots b_{n1} b_{12} \cdots b_{n2} \cdots b_{1n} \cdots b_{nn})^n$$

within $O(\log n)$ reversals. Then $O(1)$ more reversals will give AB. So $O(\log n)$ reversals are enough for $n \times n$ Boolean matrix multiplication. □

LEMMA 32 ((Matrix Transitive Closure)). *There is a 2-tape Turing machine* M *such that given an* $n \times n$ *Boolean matrix* $A = (a_{ij})$, *stored in row major form,* M *computes the transitive closure* A^* *of* A *in* $O(\log^2 n)$ *reversals.*

Proof. Since $A^* = (E + A)^m$ for any $m \geq n$, where E is the identity matrix, we can reduce this problem to $\log n$ matrix multiplications. □

Sorting seems a very important tool for reversal complexity. Here we give an upper bound for sorting.

LEMMA 33 ((Sorting)). *There is a 2-tape Turing machine* M *that, given a list*

$$x_1 \# x_2 \# \cdots \# x_n \#$$

on its input tape, can construct the sorted list

$$x_{\pi(1)} \# x_{\pi(2)} \# \cdots x_{\pi(n)} \#$$

in $O(\log n)$ *reversals. Here each item* x_i *is assumed to have the form* $(k_i \& d_i)$ *where* k_i *is a binary number representing the sort key,* d_i *is the data associated with* k_i *and* & *some separator symbol.*

Proof. By lemma 28, we can assume that the input list is in normal form. We will be implementing a version of the well-known merge sort. The computation will be accomplished in k phases, where $n = 2^k$. To initialize the first phase, we use the parallel copying lemma to construct on tape 1 a string of the form

$$(x_1\$)^n \#(x_2\$)^n \# \cdots \#(x_n\$)^n \#$$

(for some new symbol \$) using $O(\log n)$ reversals.

For $i = 1, 2, \ldots, k$, let $f(i) = 2^k - \sum_{j=1}^{i-1} 2^j = 2^k - 2^i + 2$. So $f(1) = 2^k$ and $f(k) = 2$. At the beginning of the $(i+1)$st phase $(i = 0, 1, \ldots, k-1)$, we will inductively have on tape 1 a string of the form

$$B_1\#B_2\#\cdots\#B_{2^{k-i}}\#$$

where each B_j has the form

$$(x_{j,1}\$)^{f(i+1)}\%(x_{j,2}\$)^{f(i+1)}\%\cdots\%(x_{j,2^i}\$)^{f(i+1)}$$

where $x_{j,1}, x_{j,2}, \ldots, x_{j,2^i}$ are distinct items from the original list, already sorted in non-decreasing order. Call B_i the ith *block*. The substring in a block between two consecutive '%' is called a *subblock*. Observe phase 1 has been properly initialized above.

We also need some auxiliary data to carry out the induction. Let $w_0 = 0^s1$ where $s = |x_i|$ (for all i). For this, we assume that at the beginning of the $(i+1)$st phase, we also store on tape 1 the patterns

$$P_i = \left(w_0^{2^i}\#\right)^n$$

and

$$Q_i = \left(w_0^{f(i)}\#\right)^n$$

Note that Q_i can be easily obtained from P_{i-1} and Q_{i-1} within $O(1)$ reversals and P_i can be obtained from P_{i-1} in another $O(1)$ reversals. Again, phase 1 is easily initialized.

Let us now see how we carry out phase $i+1$. First we split the string $B_1\#B_2\#\cdots\#B_{2^{k-i}}\#$ into two lists containing the odd and even numbered blocks: tape 1 now contains $B_1\#B_3\#\cdots\#B_{2^{k-i}-1}\#$ and tape 2 contains $B_2\#B_4\#\cdots\#B_{2^{k-i}}\#$. This can be done using $O(1)$ reversals.

Our goal is to 'merge' B_{2j-1} with B_{2j}, for all $j = 1, 2, \ldots, 2^{k-i-1}$ in parallel. Let

$$B_{2j-1} = (y_1\$)^{f(i+1)}\%(y_2\$)^{f(i+1)}\%\cdots\%(y_{2^i}\$)^{f(i+1)}$$

and

$$B_{2j} = (z_1\$)^{f(i+1)}\%(z_2\$)^{f(i+1)}\%\cdots\%(z_{2^i}\$)^{f(i+1)}$$

We begin by comparing the first copy of y_1 with the first copy of z_1. Suppose $y_1 \geq z_1$ ('$\geq$' here is the ordering on the items, as defined by the sort key). Then we 'mark' the first copy of z_1 and move head 2 to the first copy of z_2 in the block B_{2j}. We can compare the second copy of y_1 with this copy of z_2, marking the smaller of y_1 and z_2. If y_1 is smaller, we move to the first copy of y_2 and next compare y_2 with z_2. Otherwise, z_2 is smaller and we next compare the 3rd copy of y_1 with the first copy of z_3. In general, suppose we compare (some copy of) y_i with (some copy of) z_j. We mark the smaller of the two. If y_i is smaller, we next move head 1 to the first copy of y_{i+1} and move head 2 to the next copy of z_j and proceed with comparing these copies of y_{i+1} and z_i; Otherwise, z_j is smaller and the roles of y and z are exchanged in the description. (For correctness, we will show below that there is a sufficient number of copies of each item.)

Eventually, one of the blocks is exhausted. At that point, we scan the rest of the other block and mark the first copy of each remaining item. Notice that each subblock now contains exactly one marked copy. We then proceed to the next pair of blocks (B_{2j+1} and B_{2j+2}).

After we have compared and marked all the pairs of blocks, we will get two strings S_1 and S_2, with one copy of an item in each subblock of each block marked, on the two tapes, respectively. Our goal is to produce the merger of B_{2j-1} and B_{2j} from S_1 and S_2. Call the merged result B'_j. We will scan S_1 and output on tape 2 a partially instantiated version of B'_j. Our preceding marking algorithm ensures that if a marked copy of item w in S_1 is preceded by h copies of w then w has been compared to and found larger than h other items in S_2. In the merge result, we want to place these h smaller items before w. Since each item should be repeated $f(i+2)$ times in its subblock in B'_j we must precede w with

$$h(1 + |z_1\$|f(i+2))$$

blank spaces to accommodate these h subblocks which will be placed there in another pass. With the help of the pattern Q_{i+2}, we can leave the required amount of blank spaces for each subblock in B'_j. More precisely, we make a copy of Q_{i+2} in a track of tape 2 and use it as a 'template' which has the block and subblock structure already marked out. As we scan string S_1, for each unmarked copy of an item preceding its marked version, we skip a subblock of Q_{i+2}. When we reach a marked version of item w in S_1, we will copy that $f(i+2)$ successive copies of w into a subblock of Q_{i+2}. To see that there are enough copies of w on S_1 following this marked w, observe that there are a total of $f(i+1)$ copies of w in its subblock in S_1 and since at most 2^i copies of w precedes its marked version, there are at least $f(i+1) - 2^i \geq f(i+2)$ copies left. When we finish this process, we scan the partially formed B'_j and the string S_2 simultaneously, and fill in the remaining subblocks of B'_j. We finally have a new list of blocks:

$$B'_1\#B'_2\#\cdots\#B'_{2_{k-i-1}}\#$$

where each B'_j has the required form

$$(x_{j,1}\$)^{f(i+2)}\%(x_{j,2}\$)^{f(i+2)}\%\cdots\%(x_{j,2^{i+1}}\$)^{f(i+2)}.$$

This completes the proof. □

For some applications, we need the sorting algorithm to be stable:

LEMMA 34 ((Stable Sorting)). *The above sorting algorithm can be made stable in this sense: if the output list is*

$$x_{\pi(1)}\#x_{\pi(2)}\#\cdots\#x_{\pi(n)}\#$$

where each $x_{\pi(i)}$ has the key-data form $k_{\pi(i)}\&d_{\pi(i)}$, then for each $i = 1,\ldots,n-1$, $k_{\pi(i)} = k_{\pi(i+1)}$ implies $\pi(i) < \pi(i+1)$.

Proof. To do this, we first number, in another track (say track K) of the tape, the items in the string to be sorted sequentially. This takes $O(\log n)$ reversals. Next use the method of lemma 33 to sort the string. After sorting, those items with a common sort key will form a contiguous block. We apply the sorting method again to each block of items, where we now sort each block by the numbers of items (as written on track K). To do this in $O(\log n)$ reversals, we must do this second sorting for all the blocks simultaneously − our above method can be modified for this. □

2.8.2 Reversal is as powerful as space, deterministically

We now have the tools to prove the main result of this subsection, which is an efficient simulation of a space-bounded computation by a reversal-bounded computation (all deterministic). The idea is roughly this: to simulate a machine using $s(n)$ space, we first use the Natural Number Generation lemma to generate all configurations using space $s(n)$. Then we use the Parallel Copying Lemma to obtain a string S containing *many* copies of these configurations. With two identical copies of S, we can march down these two copies in a staggered fashion (as in the proof of the Baker-Book theorem), to extract a computation path. This method is reversal efficient because most of the reversals are made when we generated S.

THEOREM 35. *Suppose $s(n) = \Omega(\log n)$. Then any language $L \in DSPACE(s(n))$ can be accepted by a 2-tape deterministic Turing machine M within $O(s(n))$ reversals.*

Proof. First we assume that $s(n)$ is reversal constructible and $s(n) = \Omega(n)$.

Let N be a deterministic Turing machine N that accepts L in space $s(n)$. We construct a 2-tape Turing machine M that accepts L in $O(s(n))$ reversals. Fix any input x of length n.

(a) We first construct an integer $s = \theta(s(n))$ in unary, using s reversals (recall the definition of reversal constructible).

(b) Given an input x of length n, M first generates a string of the form

$$W_1 = \bar{0}\#\bar{1}\#\bar{2}\#\cdots\#\overline{2^s-1}\#$$

We can suppose that all the possible configurations of N on input x are included in the string W_1.

(c) From string W_1, using $O(1)$ reversals, we can construct another string

$$W_2 = z_0\#z_1\#\cdots\#z_{2^s-1}\#$$

such that $|z_m| = s$ and z_m is the 'successor configuration' of configuration $\overline{m}$ in the string W_1, $m = 0,1,\ldots,2^s-1$. If $\overline{m}$ is a terminal configuration, we may indicate this by using some suitable z_m.

Combining strings W_1 and W_2 into one tape, we get a string

$$W_3 = u_1\#u_2\#\cdots\#u_{2^s}$$

with each u_m $(m = 1,2,\ldots,2^s)$ is a 2-track string of the form

C_m
C'_m

where C_m and C'_m are configurations of N on input x such that $C_m \vdash C'_m$.

(d) Next we construct two identical strings S_1 and S_2 on tapes 1 and 2 of M, respectively, using $O(s)$ reversals:

$$S_1 = S_2 = ((u_1\$)^{2^s} \#(u_2\$)^{2^s} \# \cdots \#(u_{2^s}\$)^{2^s}\%)^{2^s}$$

We will call each substring of the form $(u_1\$)^{2^s} \#(u_2\$)^{2^s} \# \cdots \#(u_{2^s}\$)^{2^s}$ (as determined by two consecutive %-symbols) a *block* of S_1 or of S_2. For each $j = 1, 2, \ldots, 2^s$ we call the substring $(u_j\$)^{2^s}$ a u_j *-segment.*

(e) With these two strings, we now construct the computation path

$$P = C_0 \vdash C_1 \vdash \cdots \vdash C_i \vdash \cdots$$

of N on input x in another $O(1)$ reversals as follows.

(e1) Consider the first u_{i_1}-segment in S_1 where the upper track of u_{i_1} contains the initial configuration C_0 of N on input x. Begin by marking the first copy of u_{i_1} in this segment and place head 1 between the first and second copies of u_{i_1} in the u_{i_1}-segment. (By 'marking' of a copy of u_{i_1}, we mean a special symbol is placed at the end of that copy on a separate track. This marking amounts to nothing in the following procedure, but helps the visualization of the proof.) Moreover, place head 2 at the beginning of the first block of string S_2;

(e2) Inductively, suppose we have already found and marked a sequence $u_{i_1}, u_{i_2}, \ldots, u_{i_q}$ on the string S_1 on tape 1 such that the lower track of u_{i_j} is identical to the upper track of $u_{i_{j+1}}$, $j = 1, 2, \ldots, q-1$. Hence the upper track of u_{i_l} contains the configuration C_l in the above path P for $l = 1, 2, \ldots, q$. Suppose also that head 1 is placed between the first and second copies of u_{i_q} in the qth block in S_1 and that head 2 is placed at beginning of the qth block in S_2. Our goal is to find a segment $(u_{i_{q+1}}\$)^{2^s}$ in the qth block of S_2 such that the lower track of u_{i_q} is identical to the upper track of $u_{i_{q+1}}$. Let C_{q+1} be the configuration in the lower track of u_{i_q}. If C_{q+1} is accepting (respectively, terminal), then we accept (respectively, reject) at once. Otherwise, we look for an occurence C_{q+1} in the first upper configuration of each segment of the qth block of S_2. Note that to do this without any head reversals, we use the $2^s - 1$ successive copies of C_{q+1} in the current segment of S_1.

We will surely find the desired segment $(u_{i_{q+1}}\$)^{2^s}$ in the qth block of the string S_2. In particular, if the $2^s - 1$ copies of C_{q+1} in the current segment are used up, the desired $u_{i_{q+1}}$ must be the very last segment (the u_{2^s}-segment) in the qth block. Head 2 is now between the first and second copies of $u_{i_{q+1}}$ in the $u_{i_{q+1}}$-segment.

Our goal now is to return to the inductive basis: with head 1 between the first and second copies of $u_{i_{q+1}}$ in the $q+1$st block of S_1, and head 2 at the beginning of the $q+1$st block of S_2. But the procedure just described can be applied with simple changes (viz., we reverse the roles of S_1 and S_2 and look for the first occurence of $u_{i_{q+1}}$ in the next block of S_1). When this is done, we are ready for next iteration looking for $u_{i_{q+2}}$.

If we exhaust the entire string S_1 without accepting or rejecting in the above procedure, then N must be in a loop. Then M will reject.

We now return to two assumptions we made at the beginning of this proof: (a) First we indicate how to modify the above proof for the case $s(n) = o(n)$. The above assumed that the entire input string x is stored with each configuration. This information is now omitted, although the input head position is retained. Then by sorting the string W_1 by the positions of the input head, followed by one sweep of the input x, we are able to record the 'current input symbol' into each of the modified configurations. A similar sort can be applied to W_2 in its construction. Then the entire procedure can be carried out as before. (b) Finally, we show how to avoid the assumption that $s(n)$ is reversal constructible. The problem arises because we cannot do the usual trick of trying successive values of s by increments of one. This is because reversal, unlike space, is not reusable and it would lead to $s(n)^2$ reversals overall. But we can easily fix this by doubling the value of s at each stage ($s = 1, 2, 4, 8, \ldots$, etc). □

COROLLARY 36. *For any* $s(n) = \Omega(\log n)$,

$$DSPACE(s) \subseteq DREVERSAL(s).$$

Thus, for deterministic Turing machines, reversal as a complexity resource is at least as powerful as space. Combined with Savitch's result, we obtain:

COROLLARY 37. *For any* $s(n) = \Omega(\log n)$,

$$NSPACE(s(n)) \subseteq DREVERSAL(s(n)^2).$$

Using the fact that $DREVERSAL(r) \subseteq DSPACE(r^2)$ for $r(n) = \Omega(\log n)$ (see previous section), we get important consequences:

COROLLARY 38.

$$PLOG := DSPACE(\log^{O(1)} n) = DREVERSAL(\log^{O(1)} n)$$
$$PSPACE := DSPACE(n^{O(1)}) = DREVERSAL(n^{O(1)})$$
$$EXPS := DSPACE(O(1)^n) = DREVERSAL(O(1)^n)$$

In short, in the fundamental mode of computation, space and reversals are polynomially related.

Finally, it turns out that the above results and techniques yields a tape-reduction theorem of the form:

THEOREM 39. (Tape reduction for reversals) *Let* $r(n) = \Omega(\log n)$. *If* L *is accepted by a determinisitic multitape machine within* $r(n)$ *reversals then it is accepted by a 2-tape deterministic machine within* $O(r(n)^2)$ *reversals.*

Further results on reversal complexity can be found in [8].

2.8.3 Reversal is more powerful than time, deterministically

We now prove a result of Liśkiewicz [26] showing that deterministic reversal is stronger than deterministic time by a square factor:

THEOREM 40. *For all complexity function* $t(n) \geq n$, $DTIME(t) \subseteq DREVERSAL(\sqrt{t})$.

A fundamental problem that we need to solve is the problem of retrieving data from a table. We have already seen this problem, but let us now give it an abstract description. A **table** T is just a $m \times 2$ matrix,

$$T = k_1\$d_1\#k_2\$d_2\#\cdots\#k_m\$d_m, \qquad (k_i, d_i \in \Sigma^*).$$

Each k_i is called a **key** and d_i is called the **data item associated to** k_i. We may assume that the keys k_i's are distinct. The **retrieval problem** is formulated as follows: given the table T on tape 1 and a key $k \in \Sigma^*$ on tape 2, to locate k in T and then write its associated data d out to some specified output tape; if k is not a key in T, we output some special symbol. Using the previous techniques, we can solve this problem in $O(\log m)$ reversals: first, we generate $(k\$)^m$ on tape 3 in $O(\log m)$ reversals. Then we can search for occurence of k in T in the obvious way: compare the ith copy of k on tape 3 to k_i in tape 2. If they are equal, we can output d_i on the output tape, otherwise, we go to the next i.

A key idea in Liśkiewicz's proof is what we shall call **turbo-charged retrieval**: the table T is now replaced by

$$T' = (k_1\$d_1\$)^m\#(k_2\$d_2\$)^m\#\cdots\#(k_m\$d_m\$)^m$$

and the search argument k is replaced by $(k\$)^m$. We want the output to be $(d\$)^m$ where d is the data associated to k.

LEMMA 41. *The turbo-charged retrieval problem can be solved in* $O(1)$ *reversals.*

Proof. We proceed as in the retrieval method above that uses $O(\log m)$ reversals, except that we skip the $O(\log m)$ reversals needed to make m copies of k. Once a copy of the key k is located in T', it is a simple matter to write out the desired output $(d\$)^m$ in $O(1)$ additional reversals. □

Next, we reduce the problem of simulating a deterministic machine M that accepts in time $t(n)$ to a sequence of retrievals. On input of length n, let $h = \sqrt{t(n)}$. If C, D are of configurations of M, we say that D is a h-**successor** of C if C derives D in exactly h steps, unless the computation path from C terminates in less than h steps in which case D is the terminal configuration. Intuitively, we want to construct a table T whose pairs (k_i, d_i) correspond to pairs (C, D) where D is a h-sucessor of C. Then, starting from the initial configuration C_0, we can do a sequence of h retrievals from T to obtain $C_1, C_2, \ldots, C_h$ where C_{i+1} is the h-successor of C_i. Of course, we should turbo-charge T and C_0 so that we only use $O(1)$ reversals per turbo-charged retrieval. The problem with this scenario is that the turbo-charging is too expensive because there are $\Omega(2^t)$ distinct configurations in T and this would require $\Omega(t)$ reversals for turbo-charging.

This brings us to the second idea of the simulation: we shall consider **fragments** of configurations. Let M have k work tapes. We first partition each tape (input tape as well as work tapes) into h-**blocks** where each block consists

of h consecutive cells. We may assume that an h-block occupies the cells numbered $(i-1)h+1, (i-1)h+2, \ldots, i \cdot h$ for some $i \in \mathbb{Z}$; we may call this the ith block. An **extended h-block** is just 3 consecutive blocks on some tape; these blocks are called the **left-**, **center-** and **right-blocks**, respectively, of the extended block. An **h-fragment** is a sequence

$$F = F_0 \% F_1 \% \cdots \% F_k$$

where each F_i is the contents of an extended h-block of tape i. (We simply say 'block', 'fragment', etc, when the h is understood or irrelevant.) The 'contents' of a cell shall include the current state q if that cell is currently scanned: if a cell contains symbol a and is currently being scanned, then its contents is the pair (q, a). We say that the fragment $F = F_0 \% \cdots \% F_k$ is **centered** if, for each $i = 0, \ldots, k$, some cell in the center-block of F_i is being scanned. Note that in a centered fragment, the state q is duplicated $k+1$ times. For any configuration C, there is a unique centered h-fragment associated to C.

The importances of h-fragments is that there are $O(1)^h$ of them. If F is a centered h-fragment, we define an h-fragment G to be the h-**successor** of F as follows: let F be associated to some configuration C, and D be the h-successor of C. Then G is the h-fragment obtained from D by considering the contents of the blocks used in F. Note that G need not be centered; it is also easy to see that G is uniquely defined (it is independent of our choice of C). Let T_h be the search table whose keys are precisely all the $O(1)^h$ centered h-fragments, and the data associated to a fragment F is just the h-successors G of F.

LEMMA 42. *Given h in unary, we can set up the search table T_h in $O(h)$ reversals.*

Proof. We first generate a sequence

$$k_1 \# k_2 \# \cdots \# k_m$$

of all centered h-fragments. This is done using the number generation technique and discarding those "numbers" that do not correspond to centered h-fragments. In $O(1)$ reversals, convert this to the trivial identity table:

$$T_h = k_1 \$ d_1 \# k_2 \$ d_2 \# \cdots \# k_m \$ d_m$$

where $d_i = k_i$ for each i. Next, we 'scan' this table h times, and each time we replace the data d_i by its 1-successor fragment. Each 'scan' actually takes $O(1)$ reversals. □

We then turbo-charge table T_h to

$$T'_h = (k_1 \$ d_1 \$)^m \# (k_2 \$ d_2 \$)^m \# \cdots \# (k_m \$ d_m \$)^m$$

This amounts to parallel copying, and can be done in $O(h)$ reversals. The final step is to carry out the simulation.

We need to describe the representation of configurations. Let C be a configuration of M. For $i = 0, \ldots, k$, let the nonblank contents W_i of tape i be contained in blocks ℓ to u:

$$W_i = B_\ell B_{\ell+1} \cdots B_{u-1} B_u.$$

Let us note that W_0 always uses $\lceil n/h \rceil$ blocks; but for $i = 1, \ldots, k$, we initially represent only one block in W_i (this block is all blank except for the indication of the position of head i). During the simulation, we will add more blocks of blanks as needed. Normally, we would represent C as $W_0 \% W_1 \% \cdots \% W_k$. But to turbo-charge the representation, we represent C as $C = W'_0 \% W'_1 \% \cdots \% W'_k$ where

$$W'_i = (B_\ell \$)^m \# (B_{\ell+1} \$)^m \# \cdots \# (B_u \$)^m.$$

We now return to the proof of the main theorem. We shall operate in **stages**. In stage $i = 1, 2, 3, \ldots$, we store the value $h = 2^i$ in unary on tape 1. Using the stored valued of h, we can set up the turbo-charged table T_h on tape 2, and the turbo-charged initial configuration C_0 on tape 3. Setting up tapes 2 and 3 each takes $O(h)$ reversals. We now do the following "block simulation step" for a total of h times:

Block Simulation Step. Inductively, let tape 3 contain some turbo-charged configuration C. We first extract onto tape 4 the extended blocks corresponding to head positions on each tape of M:

$$(B^0_L \$)^m \# (B^0_C \$)^m \# (B^0_R \$)^m \% \cdots \% (B^k_L \$)^m \# (B^k_C \$)^m \# (B^k_R \$)^m. \tag{6}$$

Here, $B^i_L B^i_C B^i_R$ represents an extended block of tape i. This can be done in $O(1)$ reversals. Using the matrix transpose technique, in $O(1)$ reversals, we convert this into

$$(B^0_L B^0_C B^0_R \$)^m \# (B^1_L B^1_C B^1_R \$)^m \# \cdots \# (B^k_L B^k_C B^k_R \$)^m.$$

Doing this one more time, we convert it to the "turbo-charged h-fragment"

$$(F\%)^m = (B_L^0 B_C^0 B_R^0 \# B_L^1 B_C^1 B_R^1 \# \cdots \# B_L^k B_C^k B_R^k \%)^m.$$

We can use this as the search key for a turbo-charged retrieval on the table T_h. We may assume that the retrieved data $(G\%)^m$ is placed in tape 4, replacing $(F\%)^m$. By reversing the above process, we convert $(G\%)^m$ into the form equation (6) in $O(1)$ reversals; this is still stored in tape 4. Then in $O(1)$ reversals, we can update the configuration C in tape 3 using the contents of tape 4. This ends a "block simulation step". The number of reversals in $O(1)$.

Clearly, if C is the configuration at time t then after a block simulation step, C is the configuration at time $t + h$ (or else it is the terminal configuration). Thus after h block simulation steps, we have reached terminal configuration if the computation path terminates in $\leq h^2$ steps. If a terminal configuration, we can accept or reject accordingly. Otherwise, we double h and go to the next stage. The total number of reversals in stage h is $O(h)$. Clearly if the input is accepted in $\leq t(n)$ steps then the value of h is bounded by $2\sqrt{t(n)}$. Hence the total number of reversals is easily seen to be $O(\sqrt{t(n)})$. This proves the theorem of Liśkiewicz.

2.9 Complementation of Space Classes

In this section, we prove that deterministic and nondeterministic space classes, with some reasonable technical restrictions, are closed under complement. This result for nondeterministic space classes solves a major open problem dating back to Kuroda in 1964, namely, $NSPACE(n)$ is closed under complementation. Kuroda's problem is closely related to the so-called LBAQuestion. We may identify **linear bounded automata** with nondeterministic Turing machines that operate in linear space. Hence the class $NSPACE(n)$ is also denoted LBA. The deterministic version, $DSPACE(n)$ is known as $DLBA$. The LBAQuestion asks if $LBA = DLBA$? See §4.1 for a discussion of this and related questions. There is an important alternative characterization of LBA: it is precisely the class of **context-sensitive languages** studied in formal language theory.

The solution was independently obtained by Immerman [22] and Szelepcsényi [37]. Their solution was a surprise in two ways: the solution was surprisingly simple, and many researchers had expected the opposite result. The technique for this result has wider applications that the interested reader may further pursue (e.g., [28, 38, 24, 6]).

This section uses running complexity rather than accepting complexity. Recall our notation for running complexity classes uses the subscript 'r', as in $DSPACE_r(s)$ or $NSPACE_r(s)$. It is also crucial that the complexity function $s = s(n)$ has the property that for all $n \in \mathrm{N}$, $s(n)$ is defined and $< \infty$. Suppose a machine M has such a running space complexity $s(n)$. Then for all computation paths, on any input, use only a finite amount of space. It is important to realize that this does not preclude infinite computation paths.

2.9.1 Complement of Deterministic Classes

To motivate the proof, observe that deterministic running time classes are closed under complementation:

$$\text{co-}DTIME_r(t) = DTIME_r(t) \tag{7}$$

In proof, suppose a deterministic acceptor M runs in time t. We can easily define a deterministic acceptor N that runs in time t such that co-$L(\mathrm{M}) = L(\mathrm{N})$: N simply rejects iff M accepts, but in all other respects, N behaves exactly as M. This proves (7). Note that the restriction $t(n) < \infty$ is essential and implies that M always halts. The main result of this subsection is to show the space analogue of (7) using a technique of Sipser [36].

THEOREM 43 ((Complement of deterministic space classes)). *If $s(n) < \infty$ for all n, then* co-$DSPACE_r(s) = DSPACE_r(s)$.

The above trick of reversing the roles of acceptance and rejection is insufficient because the fact that M runs in a finite amount of space does not guarantee that M halts. The proof of theorem 43 therefore hinges upon the ability to guarantee halting; indeed it is easily seen to be an immediate consequence of:

LEMMA 44. *For every deterministic Turing acceptor* M *that runs in space $s(\cdot)$, there is a deterministic* N *that runs in space s such that $L(\mathrm{M}) = L(\mathrm{N})$ and* N *halts on all computation paths.*

One idea for proving this result is to keep two running counts for the number of steps t taken and the amount of space s used: if $t > nO_{\mathrm{M}}(1)^s$, we know that we must be looping. However, this proof requires that $s(n) \geq \log n$ in order to store the value of t (why?). The following proof places no such restrictions on $s(\cdot)$.

Proof of lemma 44. Without loss of generality, assume that M has only one work tape. This means that if $C \vdash C'$ then $space(C') - space(C)$ is 0 or 1. Let us also fix an input x of M throughout the following proof. For any $h > 0$,

let $V = CONFIGS_h(x)$ be the set of configurations of M on input x using space at most h. Let $G = G_h = (V, E)$ be the digraph whose edges comprises $(C, C') \in V^2$ such that $C \vdash_M C'$. Since M is deterministic, each node of G has outdegree at most 1. It is also not hard to see that G has indegree bounded by some constant $O_{\rm M}(1)$. The graph G has many properties, as illustrated by figure 2.4.

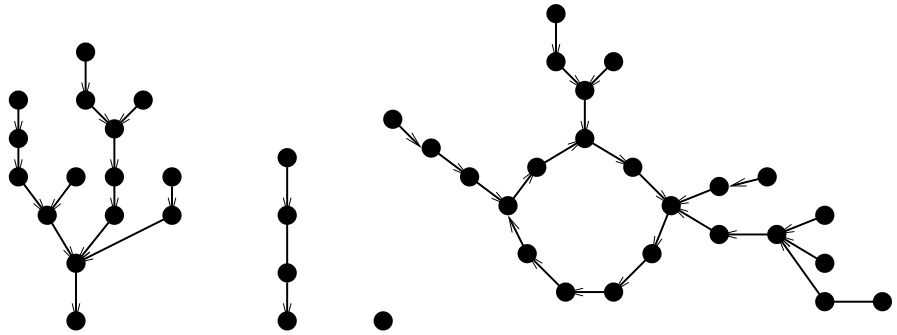

Figure 2.4: Components of a digraph G with outdegree ≤ 1.

Each acyclic component of G will just be a rooted tree. If a component is not a rooted tree, then it has a unique cycle C and has a set of trees, each rooted at some node of C. In this discussion, the edges of a tree are always directed towards the root.

For any configuration $C \in V$, let $T(C)$ denote the subgraph of G induced by the subset of configurations $C' \in V$ that can reach C. Note that by our machine conventions, if $C \vdash C'$ then $space(C) \leq space(C')$; hence the space usage of each configurations in $T(C)$ is at most $space(C)$. If C is a node in a directed cycle of G then $T(C)$ an the entire component of G; otherwise, $T(C)$ is a tree. Unfortunately, to discover if C is in a cycle may take a long time. For our application, an alternative easy-to-check condition is sufficient: we say C is *expectant* if $C \vdash C'$ and $space(C') > space(C)$. The following is immediate:

Fact. If C is expectant or terminal then $T(C)$ is a tree.

We describe a search procedure for $T(C)$ where C is expectant or accepting (hence terminal). This procedure returns "success" if the initial configuration $C_0(x)$ is in $T(C)$, and "failure" otherwise. We will use the standard depth-first search (DFS) order of visiting the nodes, but with one important difference: In the usual implementation of DFS, we need to represent the path from the root to the currently visited node, maintained on a (recursion) stack. This information is needed to return from a node to its parent. However, in a DFS on $T(C)$ we can avoid storing this information since the parent of a node C' is simply the successor of C'.

```
Search(C):
    Input: configuration C such that T(C) is a tree.
        if C = C_0(x), return("success").
        for each child C' of C
            Recursively SEARCH(C').
            If search of C' is successful, return("success").
        return("failure"). //no recursive calls were successful.
```

We now show how to implement $Search(C)$. on a Turing machine N using $O(space(C))$ amount of space. Notice that the conventional representation of C uses $O(space(C) + \log |x|)$ space. We avoid the $\log |x|$ term by assuming that x is available on the input tape, and the input head is at the correct position of the input. Hence there is no need to represent x or the input head position explicitly on the work tapes. Let N use two work tapes only. Tape 0 stores the input x; tape 1, it stores the current node C that is being visited; tape 2 is used as scratch space. At the start of the recursive call, N checks if C is equal to $C_0(x)$ and if so, returns with "success". Else, if C is a leaf, return with "failure". Otherwise, N will generate the first child C_1 of C and store it on tape 1, and recursively call *search* on C_1. We leave as an exercise to work out the details of this generation. In general, suppose we have just returned from a recursive *search* on some child C_i of C. By assumption, we a copy of C_i on tape 1. It is easy to generate C from C_i, since $C_i \vdash C$. If the *search* of C_i is a success, we return success from C; otherwise, we try to generate the next child C_{i+1} to repeat the recursive search. If C has no more children, we return "failure". This completes the description of *search*.

To conclude the proof of the lemma, we show how to use the *search* procedure. N will compute in *stages*. In *stage* h ($h = 1, 2, \ldots$) N systematically generates all configurations C that use space h. At the start of stage h, we initialize a *found* flag to false. For each expectant or accepting C, N calls the *search*(C). If the search is successful we take one of two actions: if C is accepting, we accept (terminating the entire computation). Otherwise, C is expectant, and we set the *found* flag to true. After we have examined all configurations using space h, we check the *found* flag. This flag is true iff $C_0(x)$ can reach some expectant configuration using space h. In this case N proceeds to stage $h + 1$. Otherwise, we terminate the entire computation and reject.

Why is this procedure correct? If this procedure accepts, then $C_0(x)$ reaches some accepting configuration. If this procedure rejects, $C_0(x)$ does not reach any accepting configuration using space $\leq h$, and moreover, it can never reach any configuration using more than space h. Hence we may safely reject. Thus acceptance or rejection decisions by our procedure are always correct. But we must guard against the possibility of looping. But our requirement for proceeding from stage h to the stage $h + 1$ ensures that the original machine M actually runs in at least $h + 1$ space if we reach stage $h + 1$. Since $s(n) < \infty$, we must get a decision in some finite stage. □

As an interesting consequence of the previous lemma, we have:

THEOREM 45. *If $n \leq s(n) < \infty$ is an exact space-complexity then $s(n)$ is space-constructible.*

Proof. Let M run in space exactly s. By the previous lemma, we may assume that M halts on all inputs. Now construct another acceptor N as follows: given an input of length n, N ignores the input except for noting its length. Then for each word x of length n, N simulates M on x, marking out the number of cells used. It is clear that N will mark out exactly $s(n)$ cells. The condition that $s(n) \geq n$ is needed in order to cycle through all words of length n. □

Note: Hopcroft and Ullman, [21] originally showed that $DSPACE_r(s)$ is closed under complementation under the restriction $s(n) = \Omega(\log n)$s, using the idea of a counter mentioned above. Later Hartmanis and Berman [14] proved that, in case the input alphabet is unary, we can drop the assumption $s(n) = \Omega(\log n)$. The above proof from Sipser shows that we can drop this assumption in general.

Jianer Chen has shown that deterministic reversal classes are also closed under complement:

THEOREM 46. *Let* M *accept in $f(n) < \infty$ reversals where f is time-constructible. Then there is a machine* N *that accepts L*(M) *in $O(f)$ reversals and always halts.*

2.9.2 Complement of Nondeterministic Classes

We prove the Immerman-Szelepcsényi result:

THEOREM 47 ((Complement of nondeterministic space classes)). *If $\log n \leq s(n) < \infty$ for all n then $NSPACE_r(s) =$* co-$NSPACE_r(s)$.

For the proof, let us fix M to be any acceptor running in space $s(n) \geq \log n$. We also fix the input word x. For any $h \geq 0, m \geq 0$, let $REACH_h(m)$ denote the set of all configurations of M using space at most h that can be reached from the initial configuration $C_0(x)$ in at most m steps. The *(reachability) census function* $\alpha_h(m)$ (or simply $\alpha_h(m)$) is given by

$$\alpha_h(m) = |REACH_h(m)|.$$

Note that $\alpha_h(0) = 1$ and $\alpha_h(m) \leq \alpha_h(m+1)$. Moreover, if $\alpha_h(m) = \alpha_h(m+1)$ then for all $j > m$, $\alpha_h(m) = \alpha_h(j)$. Our basic goal is to compute this census function.

But first, let us see how such a census function is used. This is seen in a nondeterministic subroutine

$$Check(C, h, m, \alpha_h(m))$$

that "accepts" iff $C \in REACH_h(m)$.

```
Check(C, h, m, α_h(m)):
    begin checking
    c_0 ← 0; // Initialize counter
    For each configuration C' using space at most h,
        begin // iteration
        Nondeterministically guess a path using ≤ h space
        and ≤ m steps, starting from C_0(x);
            If this path reaches C then accept;
            If this path reaches C' then increment c_0 by 1;
        end // iteration
    If c_0 = α_h(m) then reject;
    Else loop;
    end // checking.
```

Note that this procedure terminates in one of three ways: *accept, reject, loop.* These are indicated (respectively) by entering state q_a, state q_r or reaching any terminal configuration that is neither accepting nor rejecting. Assuming $h \geq \log |x|$, it is not hard to see that the space used by the algorithm is

$$h + \log m$$

(the $\log m$ space to count the number of steps in guessed paths). This algorithm is remarkable in three ways. First, we notice that this algorithm does not take all inputs, but only those $\langle C, h, m, p\rangle$ whose last three arguments are constrained by the equation $p = \alpha_h(m)$. Such inputs are called well-formed for this algorithm.[11] Logically speaking, the last argument is redundant because h and m determines $\alpha_h(m)$. However, from a complexity viewpoint, this argument is not redundant because it is not easy to compute $\alpha_h(m)$ from h and m. Second, we make a distinction between rejection (*i.e.*, halting in the reject state q_r) and looping. In general, looping means that the machine either does not halt or enters a terminal state that is neither accepting nor rejecting. Third, our acceptor has the property that on any well-formed input:
(a) at least one path accepts or rejects;
(b) there does not exist two paths, one accepting and the other rejecting.

In general, we call a nondeterministic acceptor M *unequivocal* if it satisfies (a) and (b) above. We then say that M *unequivocally accepts* its (well-formed) inputs.

Recall in section 2 the definition of univalent (nondeterministic) transducers that compute transformations. We want to use such transducers to compute $\alpha_h(m)$. The input k and output $\alpha_h(m)$ are represented in binary.

LEMMA 48. *Let $h \geq \log |x|$. There is a univalent transducer* N *that computes $\alpha_h(m)$ for any $h, m \geq 0$. The running space of* N *is $O(h + \log m)$.*

Proof. Suppose that we have recursively computed $\alpha_h(m-1)$; note that this value can be stored in space $O(h + \log |x|)$. This is computed univalently, meaning that at least one partial computation path that leads to the value $\alpha_h(m-1)$ and enters a special state q_1; furthermore, all paths that lead to state q_1 yield the same value. Assuming that we have reached this special state q_1, we then call the following *Count* subroutine that computes $\alpha_h(m)$ from $\alpha_h(m-1)$:

```
Count(h, m, α_h(m − 1)):
begin counting;
    c_1 ← 0; // initialize counter
    For each configuration C using space at most h,
        begin // outer iteration
        c_2 := 0; // initialize counter
        For each configuration C' using space at most h,
            begin // inner iteration
                res ← Check(C', h, m − 1, α_h(m − 1));
                If ((res="accepts") and (C' = C or C' ⊢ C)) then c_1 ← c_1 + 1 and break;
            end // inner iteration
        end // outer iteration
    Return(c_1);
end // counting.
```

[11] Cf. assumption (C) in chapter 1 (§4).

The counter c_1 is incremented if we found C that can be reached in at most m steps. Note that once counter c_1 is incremented, we break out of the inner loop. Hence, c_1 counts the number of configurations that are reachable in at most m steps. If the nondeterministic subroutine *Check* loops, then we automatically loop.

Let us verify that the procedure for computing $\alpha_h(m)$ is univalent. The key observation is that for every configuration C, the paths (restricted to the inner iteration only) have these properties:

- If $C \in REACH_h(m)$ then there is a non-looping path that increments c_1. Moreover, every path that fails to increment c_1 must loop.
- If $C \notin REACH_h(m)$ then c_1 is never incremented. Moreover, there is a non-looping path.

The space used is $O(h + \log m)$. Our lemma is proved. □

We conclude the proof of the main theorem. On input x, we show how to accept x if and only if x is not accepted by M.

MAIN PROCEDURE
0. $h \leftarrow \log |x|$. *// initialize*
1. *// Start of stage h*
2. For $m = 1, \ldots, m_0$, compute the $\alpha_h(m)$,
   where m_0 is the first value satisfying $\alpha_h(m_0) = \alpha_h(m_0 - 1)$.
3. For each accepting configuration C using space $\leq h$,
3.1. Call *Check*$(C, h, m_0, \alpha_h(m_0))$;
3.2. If *Check* accepts, we reject.
3.3. (N.B. If *Check* loops, we automatically loop.)
4. "Either *accept* or else increment h and go back to step 2."

Note that since $h \geq \log |x|$, and $m_0 = |x| O_{\mathrm{M}}(1)^h$, the previous lemma ensures that we use only $O(h)$ space. Since *Check* is an unequivocal procedure, it is clear that the main procedure is also unequivocal. We must now explain step 4. We must decide whether to accept or to repeat the procedure with a larger value of h. Perhaps *Check* did not accept within the for-loop in step 3 because h is not "large enough", *i.e.*, M could accept the input with more space. How do we know whether this is the case? Well, this is the case provided there is a path in which M attempts to use more than h space. To detect this situation, it is easy to put a 'hook' into the *Count* procedure that will set a flag *Continue* to *true* iff there exists a configuration $C \in REACH_h(m)$ where $C \vdash C'$ for some C' that uses $\geq h + 1$ space. (So C is 'expectant' in the sense of the proof of the deterministic result.) So the *Continue* flag is properly set after step 2. Now in step 4 we will accept if *Continue=false*; otherwise we increment h and go back to step 2.

This concludes our proof. An open question is whether the condition $s(n) \geq \log n$ can be removed.

2.10 *The Complexity of Palindromes

We consider the (binary) *palindrome language* over $\{0, 1\}$,

$$L_{\mathtt{pal}} = \{w : w \in \{0,1\}^* \text{ and } w = w^R\}.$$

The simultaneous time-space complexity of $L_{\mathtt{pal}}$ is remarkably well-understood (in view of what little we know about the complexity of other languages). We first prove a lower bound on the simultaneous time-space complexity of $L_{\mathtt{pal}}$ using an counting technique based on 'crossing sequences'. The notion of crossing sequences is fairly general. It was originally developed by Hennie [17] for deterministic simple Turing machines, but we shall see that the technique extends to nondeterministic computations. It has recently resurfaced in proving lower bounds on VLSI computations.

Consider a computation path $\bar{C} = (C_t \ : \ t \geq 0)$ of a Turing machine M on a fixed input w_0 of length n. Recall (§6, proof of theorem 16) that the *storage configuration* of M (at instant t) refers to the $(2k+1)$-tuple $\langle q, w_1, n_1, \ldots, w_k, n_k\rangle$ that is identical to the configuration C_t except that the input word w_0 and input head position n_0 are omitted. For integer i, the *i-crossing sequence of $\bar{C}$* is $(s_1, \ldots, s_m)$ where $m \geq 0$ is the number of times the input head crosses the boundary between the ith and the $(i+1)$st cells of tape 0, and s_j is the storage configuration of M during the jth crossing. Note that i may be restricted to $0, \ldots, n$ by our standard machine convention. Consecutive storage configurations in a crossing sequence represent crossings in opposite directions.

[11]* This optional section is independent of the rest of the book.

The sum of the lengths of crossing sequences of $\bar{C}$ at positions i, for all $i = 0, \ldots, n$, is at most the time of the computation path. Observe that crossing sequences combine elements of time and of space, so it is not surprising to see that they will be used to obtain a simultaneous lower bound on these two resources. The basic property we exploit is described in the following simple lemma:

LEMMA 49 ((fooling lemma)). *Let* M *be a nondeterministic machine that accepts the words $u = u_1u_2$ and $v = v_1v_2$. If there exist accepting computations paths π_u, π_v for u and v such that the crossing sequences of these paths at position $|u_1|$ for input u and at position $|v_1|$ for input v are identical then* M *also accepts u_1v_2 and v_1u_2.*

Proof. By induction on the length of the $|u_1|$-crossing sequence, we can splice together an accepting computation path for u_1v_2 composing of sections that come alternately from π_u and π_v. □

The following is a useful inequality:

LEMMA 50. *Let $A = (a_{i,j})$ be an $m \times n$ matrix where the $a_{i,j}$ are real numbers. Let $c_j = \frac{1}{m}\sum_{i=1}^{m} a_{i,j}$ be the average value of the numbers on the jth column. Then*

$$\max_{i=1,\ldots,m} \left(\sum_{j=1}^{n} a_{i,j} \right) \geq \sum_{j=1}^{n} c_j.$$

Proof. The left-hand side of the inequality is at least

$$\frac{1}{m}\left(\sum_{j=1}^{n} a_{1,j}\right) + \frac{1}{m}\left(\sum_{j=1}^{n} a_{2,j}\right) + \cdots + \frac{1}{m}\left(\sum_{j=1}^{n} a_{m,j}\right) = \frac{1}{m}\sum_{i=1}^{m}\sum_{j=1}^{n} a_{i,j},$$

which is seen to equal the right-hand side. □

Hennie originally showed that a simple Turing machine requires time $\Omega(n^2)$ to $L_{\texttt{pal}}$. We now generalize this result to nondeterministic multi-tape machines; in this more general setting, we only lower bound the product of time and space.

THEOREM 51. *Let t and s be non-decreasing complexity functions. If*

$$L_{\texttt{pal}} \in \textit{N-TIME-SPACE}(t, s)$$

then

$$t(n) \cdot s(n) = \Omega(n^2).$$

Proof. Suppose $L_{\texttt{pal}}$ is accepted by some nondeterministic M in time t and space s. Without loss of generality, consider palindromes of even length; the same proof can be modified for palindromes of odd length. Let S_n denote the set of palindromes of length $2n$. Thus $|S_n| = 2^n$. For each $w \in S_n$, let $\bar{C}_w$ denote some accepting computation path for w. Let R_n be the set consisting of all the $\bar{C}_w$'s, $w \in S_n$. For $i = 0, \ldots, n$, let $R_{n,i}$ denote the multiset[12] of i-crossing sequences of computation paths in R_n. For $u, v \in S_n$, let u_1 and v_1 denote the prefixes of u and v of length i. If $u_1 \neq v_1$ then the fooling lemma implies that the i-crossing sequences of $\bar{C}_u$ and $\bar{C}_v$ are distinct. Since there are 2^i distinct prefixes of length i in S_n, it follows that $|R_{n,i}| \geq 2^i$. Let $r_{n,i}$ denote the average length of crossing sequences in $R_{n,i}$ – since $R_{n,i}$ is a multiset, this is a weighted average which takes the multiplicity of a crossing sequence into account. Applying the inequality in the preceding lemma yields

$$t(2n) \geq \sum_{i=1}^{n} r_{n,i}$$

where the matrix entries $a_{i,j}$ correspond to the length of the jth crossing sequence in the ith computation path. Since at most $|R_{n,i}|/2$ of these crossing sequences have length $\geq 2r_{n,i}$, hence at least $|R_{n,i}|/2 \geq 2^{i-1}$ of these crossing sequences have length $< 2r_{n,i}$. Now there are $\sigma = O_{\mathrm{M}}(1)^{s(2n)}$ storage configurations since M accepts in space $s(2n)$. Let $\tau_{n,i} = \sigma^{2r_{n,i}}$. So there are at most $\tau_{n,i}$ crossing sequences of length $< 2r_{n,i}$. Hence $\tau_{n,i} \geq 2^{i-1}$, or $O_{\mathrm{M}}(1)^{s(2n)2r_{n,i}} \geq 2^{i-1}$. Taking logarithms, $s(2n)r_{n,i} = \Omega_{\mathrm{M}}(i)$. Hence

$$s(2n) \cdot t(2n) \geq s(2n) \sum_{i=1}^{n} r_{n,i} = \sum_{i=1}^{n} \Omega_{\mathrm{M}}(i) = \Omega_{\mathrm{M}}(n^2).$$

□

The preceding lower bound is essentially tight in the sense of the next result. Its proof is left as an exercise.

[12] We say 'multiset' instead of 'set' here to indicate that each i-crossing sequence $\chi \in R_{n,i}$ is associated with a multiplicity. In other words, each χ is tagged with the computation path $\bar{C}_w \in R_n$ from which χ is derived. The number of distinct tags for χ is its multiplicity.

THEOREM 52. *Let* $\log n \leq s(n) \leq n$ *be space-constructible. Then*

$$L_{\texttt{pal}} \in D\text{-}TIME\text{-}SPACE(\frac{n^2}{s(n)}, s(n)).$$

We might add that the above proofs go through if we had used *marked palindromes* instead; this is defined as the language over $\{0,1,\#\}$ consisting of words of the form $x\#x^R$ for all $x \in \{0,1\}^*$. The language of marked palindromes is intrinsically less complex than palindromes.[13] The above results on palindromes motivate an interesting composite-complexity class:

Definition 9. Let f be a complexity function, M a nondeterministic machine. We say M accepts in *time-space product* f if for each $w \in L(\mathrm{M})$, there is an accepting computation path of w that uses time-space (p,q) such that $pq \leq f(|w|)$. We denote the class of such languages by $NTIME \times SPACE(f)$. ■

Thus the time-space product complexity of palindromes is $\Theta(n^2)$. This definition extends to products of other pairs (or tuples) of resources. Observe that for any complexity functions t, s,

$$N\text{-}TIME\text{-}SPACE(t,s) \subseteq NTIME \times SPACE(t \cdot s).$$

In the next section we will consider space-reversal product complexity classes.

2.11 Absolute Lower Bounds

The well-known class of *regular languages* is equal to $DSPACE(0)$, i.e., those languages accepted by acceptors using no space. Recall that machines that uses no space must be 0-tape Turing acceptors: these machines are also called (nondeterministic) *finite automata*. Regular languages are very well understood. For instance, we know that

$$NSPACE(0) = NSPACE(O(1)) = DSPACE(0).$$

It is generally regarded that space bounds of $o(\log n)$ or time bounds of $o(n)$ are uninteresting for Complexity Theory. In particular, we generally consider

$$D\text{-}TIME\text{-}SPACE(n+1, \log n)$$

as the smallest interesting class for Complexity Theory. This section shows that, for a certain resource ρ, there are absolute lower limits on the complexity in the following sense: there exists a complexity function g_ρ such that if L is a non-regular language accepted in f units of resource ρ then $f = \Omega(g_\rho)$. Thus g_ρ is a lower bound on the ρ complexity of any non-regular language. We show two results of this kind. Both proofs have the same structure. But before we do that, it is instructive to dispose of the following simple observation:

LEMMA 53. *If* $t(n) \leq n$ *for any* n *then* $NTIME(t(n))$ *is regular.*

Proof. Let an acceptor M accept in time $t(n)$ and suppose $t(n_0) \leq n_0$ for some n_0. So for each word x of length n_0, M does not attempt to read the first blank symbol to the right of x. This means that for any word with prefix x, M would behave in exactly the same way. Since x is arbitrary, every input of length greater than n_0 is accepted or rejected according to the behavior of M on its prefix of length n_0. Any such language is seen to be regular. □

We now prove a result from Hong [19]. Recall the definition of product complexity classes in the last section.

THEOREM 54. *Let* f *be complexity function such that* $f(n) = o(n)$. *Then the product class* $NSPACE \times \mathrm{REVERSAL}(f)$ *is precisely the class of regular languages. In this result, reversals made by the input head are not counted.*

We remark that if reversals by input heads are counted (as in our usual definition of reversal complexity) then this theorem would have been a direct consequence of our previous lemma (using the fact that the product of space and reversal is big-Oh of time).

Proof. Suppose L is accepted by a k-tape M in space-reversal product f. To show the theorem, it suffices to prove that L is regular. Let M have as tape alphabet Γ (including the blank □), input alphabet Σ, and state set Q. Let

$$\Delta = \{first, last\} \times Q \times \Gamma^k$$

where *first* and *last* are two special symbols. Let $w \in \Sigma^*$ with $|w| \geq 2$. In the following, if $d = first$ then $w[d]$ denotes the first symbol of w and if $d = last$ then $w[d]$ denotes the last symbol. We define the binary relation R_w on Δ as follows. Let

$$t = \langle d, q, b_1, \ldots, b_k \rangle, t' = \langle d', q', b'_1, \ldots, b'_k \rangle$$

be tuples in Δ. Then $\langle t, t' \rangle \in R_w$ iff there is a sub-computation path $\bar{C} = (C_1, \ldots, C_m)$ where:

[13] The language of palindromes is context-free but not deterministic context-free. Marked palindromes are deterministic context-free.

(i) C_1 corresponds to starting M in state q with w as input, with its input head scanning $w[d]$, and the head i scanning b_i (for $i = 1, \ldots, k$).

(ii) The head on each work-tape stays stationary throughout $\bar{C}$. The input head may move. The symbols under the work heads may change.

(iii) In C_{m-1}, the input head is scanning $w[d']$ and in C_m, it is scanning a blank (so it has left the region on the tape containing w). The state and symbols scanned by the work-tape heads in C_m are $q', b'_1, \ldots, b'_k$.

For instance, R_w is empty implies that if we start the machine with w as the input word, in any initial state, with the input head at the first or last symbol of w and with any symbols under the work-tape heads, then one of two things must happen:

(1) The input-head never leaves the area of w, thus violating (iii).

(2) Eventually some work-tape head will make a move *while* the input head is still scanning w, violating (ii).

Note that since $|\Delta| = O_{\mathrm{M}}(1)$, there are $O_{\mathrm{M}}(1)$ distinct relations of the form R_w (over all w). Let α_1 denote the number of such relations R_w. We call R_w the *characteristic relation* of w.

For each $h \geq 0$, let $L^{[h]}$ denote those words w in L for which the minimal space-reversal product of any accepting computation of M for w is h. (Reversals by input head not counted.) Hence the sets $L^{[h]}$ form a partition of L. If $L^{[h]}$ is empty for all but finitely many values of h, then it is easy to show that L is regular (we leave this as an exercise). So for the sake of contradiction, let us assume $L^{[h]}$ is non-empty for infinitely many h. If $L^{[h]}$ is non-empty, then let $\mu(h)$ denote the length of the shortest word in $L^{[h]}$; else, we define $\mu(h) = 0$.

We claim that $\mu(h) < (\alpha_1 + 2)(h+1)$, where α_1 is the constant described earlier. To see this, suppose otherwise. Then for some $w \in L^{[h]}$, we have $\mu(h) = |w| \geq (\alpha_1 + 2)(h+1)$. Let $\bar{C}$ be the accepting computation path on input w that has space-reversal product of h. Let an *active* step of $\bar{C}$ refer to one in which at least one of the work heads actually moves (the input head's movement does not count). If $\bar{C}$ makes (r, s) reversal-space then the number of active steps of $\bar{C}$ is at most $rs \leq h$. Divide w into $h+1$ subwords each of length at least $\alpha_1 + 2$: a simple pigeon-hole argument shows that for some such subword u, all the work heads of M are stationary whenever the input head is scanning u. More precisely, every step of the form $C \vdash C'$ where the input head in C is scanning a symbol in u is an inactive step. Among the $\alpha_1 + 1$ of prefixes of u of length ≥ 2, there must be two prefixes v and v' with the same characteristic relation, $R_v = R'_v$. If v' is the shorter of the two prefixes, replace v with v' in w to obtain a shorter word w'. It follows from the definition of characteristic relations that w' is also accepted by M, with the same space-reversal product complexity as w. (This is essentially the fooling lemma in the previous section.) This contradicts our choice of w as the shortest word in $L^{[h]}$.

We conclude that for any shortest length word w in $L^{[h]}$, $|w| < (\alpha_1+2)(h+1)$ or $h = \Omega_{\mathrm{M}}(|w|)$. Since $f(|w|) \geq h$, and there are infinitely many such w's, we conclude $f(n) \neq o(n)$. This contradicts our assumption on f. □

We next prove a similar result of Hopcroft and Ullman [21]. To do this, we extend the notion of characteristic relations to allow the work-tape heads to move. With M, Q, Σ and Γ as above, we now define for any integer $h \geq 1$:

$$\Delta_h = \{\mathit{first}, \mathit{last}\} \times Q \times (\Gamma^h)^k \times \{1, \ldots, h\}^k.$$

(The previous definition of Δ may by identified with the case $h = 1$.) Consider a tuple

$$t = \langle d, q, w_1, \ldots, w_k, n_1, \ldots, n_k \rangle \in \Delta_h.$$

Each word w_i is of length h and, roughly speaking, denotes the contents of the ith tape of M and n_i indicates the position of head i in w_i. For any $w \in \Sigma^*$, $|w| \geq 2$, we again define a binary relation R^h_w over Δ_h as follows. Let $t \in \Delta_h$ be as given above and $t' \in \Delta_h$ be the primed version. A pair $\langle t, t' \rangle$ is in R^h_w iff there is a sub-computation path $\bar{C} = (C_1, \ldots, C_m)$ such that

(i) C_1 corresponds to starting M with w as input string, input head scanning $w[d]$, and for $i = 1, \ldots, k$: the non-blank part of tape i is w_i with head i scanning $w_i[n_i]$.

(ii) The work-tape heads stays within the w_i's throughout the computation. The input head similarly stays within w except during the last configuration C_m.

(iii) In C_{m-1}, the input head is scanning $w[d']$ and in C_m, the input head has left the region of w. The state, contents of the work-tapes and the corresponding head positions in C_m are given by the rest of the tuple t'.

Note that

$$|\Delta_h| = h^k O_{\mathrm{M}}(1)^h = O_{\mathrm{M}}(1)^h.$$

For any set X with n elements, the number of binary relations over X is 2^{n^2}. Hence if α_h denotes the number of distinct relations R_w^h, over all possible words w, then there exists some constant $C = O_{\mathrm{M}}(1)$ that does not depend on h such that

$$\alpha_h \leq 2^{C^h}$$

We can now prove the result of Hopcroft and Ullman:

THEOREM 55. *Let s be a complexity function such that $s(n) = o(\log\log n)$. Then $DSPACE(s)$ is the set of regular languages.*

Proof. Let L be accepted by some M in space $s(n)$. The proof proceeds quite similarly to the previous one. For each $h \geq 0$, let $L^{[h]}$ be the set of words in L accepted in space h but not in space $h-1$. If $L^{[h]}$ is empty for all but finitely many values of h then L is regular. So for the sake of contradiction, assume otherwise. Define $\mu(h)$ to be the length of the shortest word in $L^{[h]}$ when $L^{[h]}$ is non-empty; otherwise $\mu(h) = 0$. If w is a word of shortest length in a non-empty $L^{[h]}$ then we claim

$$|w| = \mu(h) \leq 2^{C^{s(|w|)}}. \tag{8}$$

Otherwise, there exists a pair of distinct prefixes u and v of w that have identical characteristic relations, $R_u^h = R_v^h$. Assuming that u is the longer of the two prefixes, we can decompose w into $w = vv'w' = uw'$ where $v' \neq \epsilon$. It is easy to verify that vw' is also accepted by M in the same space h (again, the fooling lemma argument). This contradicts our choice of w as the shortest word. From (8) and the fact that there are infinitely many such words w, we conclude that $s(n) \neq o(\log\log n)$. This contradicts our assumption on $s(n)$. □

We can conceive of other methods of defining languages and complexity that admit non-regular languages with space complexity bounded by functions growing slower than $\log\log n$. A simple way to do this would be to assume that the Turing machine comes equipped with a complexity function h such that on input of length n, $h(n)$ cells are automatically marked out before the computation begins, and that the Turing machine never exceeds the marked cells. Hence the machine uses space $h(n)$, which of course can be chosen arbitrarily slowly growing. Our preceding proofs would no longer be valid; indeed, it is easy to see that even non-recursively enumerable languages can be accepted this way (how?). We refer to [30] for languages with such low space complexity.

To end this section, we note that the preceding absolute lower bounds are essentially tight in the following sense:
(i) There are non-regular languages whose space-reversal product is $O(n)$. This is illustrated by the palindrome language: it is non-regular and it can be accepted in linear space and $O(1)$ reversals.
(ii) There non-regular languages in the class $DSPACE(\log\log n)$. In fact, the following language is an example.

$$L_0 = \{\bar{1}\#\bar{2}\#\cdots\#\overline{n-1}\#\bar{n} \ : \ n \text{ is a natural number}\}$$

where $\bar{m}$ denotes the binary representation of the integer m. There is another candidate language due to Mehlhorn and Alt[29]. It is rather interesting because it is over a single letter alphabet. First let $q(n)$ denote the smallest number that does not divide an integer n.

$$L_1 = \{1^n \ : \ n \text{ is a natural number and } q(n) \text{ is a power of two } \}$$

where 1^n is the unadic representation of n.

We leave it as an exercise to show that both these languages have the required properties.

2.12 Final remarks

This chapter introduces the basic model of computation that is to be the basis for comparing all future models of computation. This choice is not essential to the theory because of the computational theses in chapter 1. The results in this chapter are of two types: (I) technical ones about Turing machines in particular and (II) relationships about complexity classes. Although we are mainly interested type (II) results, some of the model-specific results are unavoidable. This would be true regardless of our choice of model of computation. Our choice of the Turing model is particularly fortunate in this regard because the literature contains a wealth of these model-specific results. We have restricted ourselves to those type (I) results that are needed later.

The three simulation sections use the very distinct techniques for each resource. They illustrate the fundamentally different properties of these resources. A common aphorism expressing this difference between time and space

says *time is irreversible but space is reusable.*[14] Both the Savitch and Istvan Simon techniques exploit the reusability of space. But what of reversals? Within a phase of a computation, all the individual actions are independent of one another. This is exactly the nature of time in parallel computation. This prompted Hong Jia-wei to identify reversals with parallel time. This identification is essentially correct except for low level complexity; we will return to the precise relationships in later chapters. Since time and space without reversal is useless (why?), we offer this complement to the above aphorism: *reversal is the ingredient to convert reusable space into time.*

The interplay of the triumvirate is best illustrated by deterministic computation: if $t_{\rm M}(n), s_{\rm M}(n)$ and $r_{\rm M}(n)$ are the exact running time, space and reversal of a halting deterministic Turing machine M, then there are constants $c_1, c_2, c_3 > 0$ that make the following fundamental inequalities true:

$$s_{\rm M} + r_{\rm M} \leq c_1 t_{\rm M} \leq c_2(n + s_{\rm M})r_{\rm M} \leq c_3 t_{\rm M}^2.$$

A basic challenge of complexity theory is to refine these inequalities and to extend them to other computational modes.

Exercises

[2.1] Construct a simple Turing machine that on input $x \in \{0,1\}^*$ converts x to $x-1$ where we regard x as a binary integer. Hence the machine decrements its input. For instance if $x = 0110100$ then the machine halts with 0110011. Write out the transition function δ of your machine explicitly.

[2.2] (i) Give a nondeterministic Turing acceptor that accepts the complement of the palindrome language L_{pal} using acceptance space of $\log n$. The exercise in last question is useful here. (Note that the acceptor in example 2 uses linear space).
(ii) Show (or ensure) that your machine accepts in linear time. (It is probably obvious that it accepts in $O(n \log n)$ time.)

[2.3] Give an algorithm for obtaining $\delta(N)$ from $\delta(\mathrm{M})$ having the nondeterministic speedup properties described in theorem 2. What is the complexity of your algorithm?

[2.4] Consider the recognition problem corresponding to multiplying binary numbers: MUL is the language comprising those triples of the form $\langle a, b, c\rangle$ where a, b, c are binary numbers and $ab = c$. Show that MUL is in *DLOG*.

[2.5] Define 'rejection complexity' for deterministic machines such that the acceptance and rejection time complexity of L are both $f(\cdot)$ iff the running time complexity of L is also f. Can this be extended to nondeterministic machines? Remark: If acceptance complexity corresponds to 'proofs' in formal proof systems then rejection complexity corresponds to 'disproofs' of (or, counter-examples to) non-theorems.

[2.6] (i) Show that if $f(n) = \omega(n)$ and f is time-constructible, then $XTIME_r(f) = XTIME(f)$, $X = D, N$.
(ii) Show a similar result for space-constructible f, but no longer assuming $f(n) = \omega(n)$.

[2.7] * Show that the following functions are time-constructible for any positive integer k:
(i) $n \log^k n$, (ii) n^k, (iii) $n!$ (factorial), (iv) k^n. (See [23] for a systematic treatment of such proofs.)

[2.8] Redo the last problem for reversal-constructible. (Recall that reversal-constructible is defined slightly differently than in the case of time or space.)

[2.9] Suppose $f(n) \geq n$ and $g(n)$ is an integer for all n, and both f and g are time-constructible.
(i) Show that the functional composition $f \circ g$ is time-constructible. (*Note:* $f \circ g(n) = f(g(n))$. It is easy to see that $f(n) + g(f(n))$ is time-constructible, but we want $f(g(n))$.)
(ii) If $f(n)$ is also an integer for all n, show that $f \cdot g$ (product) is time-constructible.

[2.10] * (i) Show that $\log n$ is space-constructible.
(ii) Define $\log^* n$ to be the largest integer m such that $\exp(m) \leq n$ where $\exp(0) = 0$ and $\exp(m+1) = 2^{\exp(m)}$. Show that $n \log^* n$ is space-constructible.
(iii) (Ladner, Freedman) Show that there exists a space-constructible f such that $f(n) = O(\log \log n)$ and, for some $C > 0$, $f(n) \geq C \log \log n$ infinitely often. **Hint:** $f(n)$ is the largest prime p such that every prime

[14] Algebraically, this means that we must sum (respectively, take the maximum of) the time (respectively, space) used by several independent computations.

$q \leq p$ divides n. Some basic knowledge of number theory (mainly the prime number theorem) is required for this problem.
(iv) (Seiferas) Show that $\log\log n$ is not space-constructible. **Hint:** prove that any space-constructible function not in $\Omega(\log n)$ must be $O(1)$ infinitely often.

[2.11] For most applications, a weaker notion of constructibility suffices: A function f is *approximately time-constructible* if there is a time-constructible function f' such that $f = \Theta(f')$. Redo the above exercises but using the "approximate" version of time-constructibility instead. It should be much easier.

[2.12] (J.C. Shepherdson, 1959) The usual definition of regular languages is that is is accepted by a finite automata. A finite automata is basically a simple Turing machine in which the sole tape head can only move from left to right. It must accept when it reaches the first blank cell. Prove that this definition gives rise to the same class *REG*. HINT: if (L, Σ) is a language, and $x, y, z \in \Sigma^*$, then $x \equiv_L y$ iff for all z, $xz \in L$ iff $yz \in L$. There are finitely many equivalence classes of the relation $\equiv_L$ when L is regular.

[2.13] Show that the two languages at the end of Section 11 are non-regular and can be accepted in $O(\log\log n)$ space. For non-regularity, the reader should use the 'pumping lemma for regular languages'.

[2.14] Show that if a function is time-constructible, then it is space-constructible.

[2.15] Prove that for any constant k, $NTIME(k)$ is a proper subclass of $NSPACE(0)$.

[2.16] There are three possible responses to the following statements: True, False or Perhaps. 'Perhaps' reflects the uncertainty from what we know (assume only the results in this chapter). If your answer is True or False, you must give brief reasons in order to obtain full credit.
(a) T/F/P: $NSPACE(n) \subseteq DTIME(O(1)^n)$.
(b) T/F/P: $NTIME(n) \subseteq DTIME(2^n)$.
(c) T/F/P: $NREVERSAL(n) \subseteq DREVERSAL(O(1)^n)$.

[2.17] (Hartmanis) For every $t(n) \geq n\log n$, if M is a simple Turing machine accepting in time t then M accepts in space $O(t(n)/\log t(n))$. **Note:** This result has been improved in two ways: First, this result has been shown for multitape Turing machine by Hopcroft, Valiant and Paul (but Adleman and Loui provided a very different proof). For simple Turing machines, Paterson shows that if a language is accepted by a simple Turing machine in time t then it can be accepted in space $t^{1/2}$. See chapter 7 for these results and references.

[2.18] (i) Construct a Turing machine that computes the Boolean matrix $\mu_{x,h}$ and then the transitive closure $\mu^*_{x,h}$ (say, using the Floyd-Warshall algorithm) in time $O(m^3)$ where the matrix μ_c is $m \times m$. (We already know that transitive closure can be computed in $O(m^3)$ time on a random access machine – the question is whether a Turing machine can achieve the same time bound. The organization of the Boolean matrix on the tapes is crucial.)
(ii) Give an alternative proof of theorem 16 using the previous result on computing transitive closure.

[2.19] Modify the proof of theorem 18 for the corresponding result for running complexity.

[2.20] * Show that $DTIME(n+1) \neq DTIME(O(n))$.

[2.21] (Hopcroft-Greibach) If $L \in NTIME(n+1)$ then L has the form $h(L_1 \cap \cdots \cap L_k)$ where the L_i's are context-free languages and h is a letter homomorphism.

[2.22] This exercise gives a characterization of $NTIME(t)$ for any complexity function t. Let $h : \Sigma \to \Gamma^*$ be a homomorphism and t a complexity function. We say h is *t-bounded on L* if for all n large enough, for all $x \in h(L)$ with $|x| \geq n$, there exists $w \in L$ such that $h(w) = x$ and $|w| \leq t(|x|)$. [15] For any complexity class K, define $H_t[K]$ to be

$$\{h(L) : L \in K \wedge h \text{ is } t\text{-bounded on } L\}.$$

Prove that for all t, $NTIME(t)$=$H_t[NTIME(n+1)]$. Conclude from this that $NTIME(n+1)$ is closed under non-erasing homomorphism. (Note: h is *non-erasing* if $h(b) \neq \epsilon$ for all $b \in \Sigma$.)

[15] Our definition here differs from the the literature because of our use of acceptance time rather than running time. The usual definition simply says: $\exists c > 0$, for all $w \in L$, $|w| \leq ct(|x|)$.

[2.23] Prove theorem 52.

[2.24] (Hennie) Show that every simple Turing acceptor for the palindrome language L_{pal} takes time $\Omega(n^2)$.

[2.25] A worktape of a Turing machine is said to be a *pushdown store* if it is initialized with one special symbol (called *bottom symbol*) which is never erased, and tape head is constrained so that (1) if it moves left, it must write a blank on the current cell (this is called a *pop move*, and (2) the blank symbol is never written under any other circumstances. When the head moves moves right, we call it *push move*. These requirements imply that the tape head is always scanning the rightmost non-blank cell on the tape (called the *top symbol*) or (temporarily for one moment only) the blank symbol on the right of the top symbol. An *auxiliary pushdown automata* (apda) is a nondeterministic 2-tape Turing acceptor in which tape 1 is a pushdown store and tape 2 is an ordinary worktape. For apda's, we only count the space usage on tape 2; that is, space on the pushdown store is not counted. Show the following to be equivalent:
(a) L is acceptd by a deterministic apda using space $s(n)$.
(b) L is acceptd by a nondeterministic apda using space $s(n)$.
(c) $L \in DTIME(O(1)^{s(n)})$.

[2.26] (Fischer, Meyer, Rosenberg) A *counter machine* (CM) is a multitape Turing machine except that instead of tapes, the CM has 'counters'. A counter contains a non-negative integer which the CM can increment, decrement or test to see if it is equal to zero. More precisely, if the CM has k counters, we may represent its transition table by a set of tuples (instructions) of the form

$$\langle q, a, z_1, \ldots, z_k, q', d_0, d_1, \ldots, d_k \rangle$$

where q, q' are states, a is an input symbol, $z_1, \ldots, z_k \in 0, 1$ and $d_0, d_1, \ldots, d_k$
$\in \{-1, 0, +1\}$. The above instruction is *applicable* if the current state is q, the input head is scanning the symbol a, and the ith counter ($i = 1, \ldots, k$) contains a zero iff $z_i = 0$. To *apply* the instruction, make the next state q', move the input head in the direction indicated by d_0, and increment or decrement the i^th counter by d_i. The *space* used by the counter on an input w is the largest integer attained by any counter during the computation. Show that for any CM accepting in space f, there is another CM accepting the same language but using space only $f^{1/2}$. (Thus counter machines exhibit a 'polynomial space speedup'.) **Hint:** show how to replace one counter by two counters where the new counters never contain a number larger than the square-root of the largest number in the original counter.

[2.27] (P. Fischer) Show a linear speedup result for simple Turing machines. **Hint:** the Linear Speedup theorem as stated fails but what stronger assumption about $t(\cdot)$ must be made?

[2.28] Show how to reduce the retrieval problem (section 2.8.3) to the sorting problem, thereby giving another $O(\log m)$ deterministic reversal solution. [Recall the retrieval problem is where we are given, say, on tape 1 an integer $h \geq 1$ in unary, on tape 2 a word $c \in \{0, 1\}^*$, and on tape 3 the table

$$c_1\$d_1\#c_2\$d_2\# \cdots c_m\$d_m\$, \qquad (c_i, d_i \in \{0, 1\}^*)$$

where $m = 2^h$, and the c_i's are assumed distinct. We want to output d_i on tape 4 where $c = c_i$ for some $i = 1, \ldots, m$. You may assume that c **does** occur among $c_1, \ldots, c_m$.

[2.29] (Chrobak and Rytter) Show that if we do not count reversals made by the input head then *DLOG* can be simulated in constant reversals.

[2.30] Give a proof that the class $DSPACE(s(n))$ is closed under complementation if $s(n) \geq \log n$, using the idea of counters described in section 8.

[2.31] Show that the census technique can also be used to show that the class $NSPACE_e(s(n))$ of languages accepted by *unequivocal* acceptors in space $s(n) \geq \log n$ is closed under complementation. Recall that on any input, the computation tree of an unequivocal acceptor has at least one accepting path or one rejecting path (rejection is by entering the special state q_r only). Moreover, the computation tree cannot have both accepting and a rejecting path.

[2.32] (Fischer, Meyer, Rosenberg) Suppose we allow each work-tape to have more than one reading head, but the number is fixed for each tape, depending on the particular machine. Furthermore, when two or more tape heads are scanning the same cell, this information is known to the machine. Formalize this *multihead multitape model* of Turing machines. Show that any such machine can be simulated by ordinary multitape machines without time loss.

[2.33] Let t, s, r be the running time, space and reversal of a halting Turing machine M. Does the fundamental inequalities $s+r = O_{\mathrm{M}}(t) = O_{\mathrm{M}}((s+n)r) = O_{\mathrm{M}}(t^2)$ (see the concluding section) hold for nondeterministic M?

[2.34] * Define the static complexity of a problem L to be the smallest sized Turing machine that will accept L. 'Machine size' here can be variously interpretated, but should surely be a function of the tape alphabet size, the number of tapes and the number of states. Does every problem L have a unique static complexity? How sensitive is this complexity to change in Turing machine conventions?

[2.35] * Is there some form of linear speedup theorem for reversal complexity?

[2.36] * Can the tape reduction theorem for reversal complexity be improved?

[2.37] * In the inclusion $\mathit{D\text{-}SPACE\text{-}REVERSAL}(s, r) \subseteq \mathit{DSPACE}(r \log s)$, can we replace the left-hand side by $\mathit{N\text{-}SPACE\text{-}REVERSAL}(s, r)$ possibly at the expense of a large space bound on the right-hand side? E.g., $\mathit{N\text{-}SPACE\text{-}REVERSAL}(s, r) \subseteq \mathit{DSPACE}(rs/\log s)$.

[2.38] * Prove or disprove: let $t(n) = \Omega(n^2)$, $r(n) = \Omega(\log t(n))$, $s(n) = \Omega(\log t(n))$, $r(n) \cdot s(n) = O(t(n))$. Then $\mathit{DTIME}(t) - \mathit{D\text{-}SPACE\text{-}REVERSAL}(s, r) \neq \emptyset$.

Bibliography

[1] S. O. Aanderaa. On k-tape versus $(k-1)$-tape real time computation. In R. M. Karp, editor, *Complexity of computation*, pages 74–96. Amer. Math. Soc., Providence, Rhode Island, 1974.

[2] B. S. Baker and R. V. Book. Reversal-bounded multipushdown machines. *Journal of Computers and Systems Science*, 8:315–322, 1974.

[3] R. Book and C. Yap. On the computational power of reversal-bounded machines. *ICALP '77*, 52:111–119, 1977. Lecture Notes in Computer Science, Springer-Verlag.

[4] R. V. Book and S. A. Greibach. Quasi-realtime languages. *Math. Systems Theory*, 4:97–111, 1970.

[5] R. V. Book, S. A. Greibach, and B. Wegbreit. Time and tape bounded Turing acceptors and AFL's. *Journal of Computers and Systems Science*, 4:606–621, 1970.

[6] S. Buss, S. Cook, P. Dymond, and L. Hay. The log space oracle hierarchy collapses. Technical Report CS103, Department of Comp. Sci. and Engin., University of California, San Diego, 1987.

[7] E.-C. Chang and C. K. Yap. Improved deterministic time simulation of nondeterministic space for small space: a note. *Info. Processing Letters*, 55:155–157, 1995.

[8] J. Chen. *Tape reversal and parallel computation time.* PhD thesis, Courant Institute, New York University, 1987.

[9] J. Chen and C.-K. Yap. Reversal complexity. *SIAM J. Computing*, to appear, 1991.

[10] M. Davis. Why Gödel didn't have Church's Thesis. *Information and Computation*, 54(1/2):3–24, 1982.

[11] P. C. Fisher. The reduction of tape reversal for off-line one-tape Turing machines. *Journal of Computers and Systems Science*, 2:136–147, 1968.

[12] J. Hartmanis. Computational complexity of one-tape Turing machine computations. *Journal of the ACM*, 15:325–339, 1968.

[13] J. Hartmanis. Tape-reversal bounded Turing machine computations. *Journal of Computers and Systems Science*, 2:117–135, 1968.

[14] J. Hartmanis and L. Berman. A note on tape bounds for sla language processing. *16th Proc. IEEE Symp. Found. Comput. Sci.*, pages 65–70, 1975.

[15] J. Hartmanis, P. M. L. II, and R. E. Stearns. Hierarchies of memory limited computations. *IEEE Conf. Record on Switching Circuit Theory and Logical Design*, pages 179–190, 1965.

[16] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Trans. Amer. Math. Soc.*, 117:285–306, 1965.

[17] F. C. Hennie. One-tape, off-line Turing machine computations. *Information and Computation*, 8(6):553–578, 1965.

[18] F. C. Hennie and R. E. Stearns. Two-tape simulation of multitape Turing machines. *Journal of the ACM*, 13(4):533–546, 1966.

[19] J. Hong. A tradeoff theorem for space and reversal. *Theor. Computer Science*, 32:221–224, 1984.

[20] J. Hong. *Computation: Computability, Similarity and Duality*. Research notices in theoretical Computer Science. Pitman Publishing Ltd., London, 1986. (available from John Wiley & Sons, New York).

[21] J. Hopcroft and J. Ullman. Some results on tape bounded Turing machines. *Journal of the ACM*, 16:168–188, 1969.

[22] N. Immerman. Nondeterministic space is closed under complement. *Structure in Complexity*, 3:112–115, 1988.

[23] K. Kobayashi. On proving time constructibility of functions. *Theor. Computer Science*, 35:215–225, 1985.

[24] K.-J. Lange, B. Jenner, and B. Kirsig. The logarithmic alternation hierarchy collapses: $A\Sigma_2^L = A\Pi_2^L$. *Proc. Automata, Languages and Programming*, 14:531–541, 1987.

[25] M. Li. On one tape versus two stacks. Technical Report Tech. Report TR84-591, Computer Science Dept., Cornell Univ., Jan. 1984.

[26] M. Liśkiewicz. On the relationship between deterministic time and deterministic reversal. *Information Processing Letters*, 45:143–146, 1993.

[27] W. Maass. Quadratic lower bounds for deterministic and nondeterministic one-tape Turing machines. *16th Proc. ACM Symp. Theory of Comp. Sci.*, pages 401–408, 1984.

[28] S. R. Mahaney. Sparse complete sets for *NP*: solution to a conjecture of Berman and Hartmanis. *Journal of Computers and Systems Science*, 25:130–143, 1982.

[29] H. A. K. Mehlhorn. A language over a one symbol alphabet requiring only $O(\log\log n)$ space. *SIGACT news*, 7(4):31–33, Nov, 1975.

[30] B. Monien and I. H. Sudborough. On eliminating nondeterminism from Turing machines which use less than logarithm worktape space. In *Lecture Notes in Computer Science*, volume 71, pages 431–445, Berlin, 1979. Springer-Verlag. Proc. Symposium on Automata, Languages and Programming.

[31] W. J. Paul, J. I. Seiferas, and J. Simon. An information-theoretic approach to time bounds for on-line computation. *Journal of Computers and Systems Science*, 23:108–126, 1981.

[32] M. O. Rabin. Real time computation. *Israel J. of Math.*, 1(4):203–211, 1963.

[33] W. Rytter and M. Chrobak. A characterization of reversal-bounded multipushdown machine languages. *Theor. Computer Science*, 36:341–344, 1985.

[34] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computers and Systems Science*, 4:177–192, 1970.

[35] I. Simon. On some subrecursive reducibilities. Technical Report Tech. Rep. STAN-CS-77-608, Computer Sci. Dept., Stanford Univ., April, 1977. (PhD Thesis).

[36] M. Sipser. Halting space-bounded computations. *Theor. Computer Science*, 10:335–338, 1980.

[37] R. Szelepcsényi. The method of forcing for nondeterministic automata. *Bull. European Association Theor. Comp. Sci.*, pages 96–100, 1987.

[38] S. Toda. $\Sigma_2 SPACE(n)$ is closed under complement. *Journal of Computers and Systems Science*, 35:145–152, 1987.

Chapter 3

Introduction to the Class NP

April 13, 2009

3.1 Introduction

According to the principles outlined in chapter 1, the class $P = DTIME(n^{O(1)})$ (resp. $NP = NTIME(n^{O(1)})$) corresponds to those problems that are time-tractable when we compute in the fundamental (resp. nondeterministic) mode. Clearly $P \subseteq NP$; the P versus NP (or '$P \neq NP$') question asks if this inclusion is proper. Since most real world computers operate in the fundamental mode, P is clearly very important. The practical significance of NP is not immediately evident. In 1971, Cook [4] proved a theorem connecting P and NP that was destined to play a decisive role in Complexity Theory. Cook showed that the well-known problem of recognizing satisfiable Boolean formulas[1] is the "hardest" problem in NP in the sense that if this problem is in P then $P = NP$. Shortly afterwards, Karp [7] showed that a large list of problems that are of practical importance in the field of operations research and combinatorial optimization are also hardest in this sense. Independently, Levin [11] proved a theorem similar to Cook's result. The list of hardest problems has since grown to include hundreds of problems arising in practically every area of the computational literature: see Garey and Johnson [5] for a comprehensive catalog up till 1979. The list highlights the fact that NP is empirically (not just theoretically) a very important class. The added significance of the Cook-Karp discoveries is that these problems in NP have defied persistent attempts by researchers to show that they are in P. Although we are probably a long way from settling the P versus NP question, most researchers believe that P is not equal to NP. Because of this belief, a proof that a certain problem is hardest for NP is taken as strong evidence of the fundamental intractability[2] of the problem. For designers of efficient algorithms, this is invaluable information: it warns that unless one has new and deep insights to these problems that escaped previous researchers, one should not expect to find polynomial time solutions.

We introduce a well-known problem that is not known to be fundamentally tractable:

Traveling Salesman Optimization Problem (TSO)	
Given:	An $n \times n$ matrix $D = (d_{i,j})$ whose entries are either non-negative integers or ∞, and $d_{i,i} = 0$ for all i.
Find:	A cyclic permutation $\pi = (\pi(1), \pi(2), \ldots, \pi(n))$ of the integers $1, 2, \ldots, n$ such that $D(\pi)$ is minimized: $D(\pi) := d_{\pi(1),\pi(2)} + d_{\pi(2),\pi(3)} + \cdots + d_{\pi(n-1),\pi(n)} + d_{\pi(n),\pi(1)}$.

The TSO represents the problem of a salesman who wants to visit each of n cities exactly once and then return to his original city. The entry $d_{i,j}$ of the matrix D represents the distance from city i to city j. Each cyclic permutation π corresponds to a *tour* of the n cities where the traveling salesman visits the cities $\pi(1), \pi(2), \cdots, \pi(n)$, and $\pi(1)$ in turn: here it does not matter which is the first city and, without loss of generality, we always let $\pi(1) = 1$. $D(\pi)$ is called the *length* of the tour.

The brute force method of solving this problem involves generating all $(n-1)!$ cyclic permutations and choosing the permutation with the minimal tour length. Thus the method takes $\Omega((n-1)!)$ time. There is a dynamic programming method due to Karp and Held that improves on this, giving $O(n2^n)$ time (see Exercises).

[1] Actually, Cook stated his theorem in terms of recognizing the set of tautologous Boolean formulas.

[2] We use the term 'fundamental intractability' for any problem (not necessarily a recognition problem) that cannot be solved in polynomial time when computing in the fundamental mode.

The reader may have observed that the TSO is not a recognition problem which is what our theory normally treats. In the next section we illustrate the empirical fact that for most optimization or functional problems Q that are not known to be fundamentally tractable, we can easily derive a corresponding recognition problem Q' with the property that Q can be solved in polynomial time iff Q' is in the class P . We say that Q and Q' are *polynomially equivalent* in this case. Note that we have not defined what it means for a functional problem like TSO to be solved in polynomial time (in the fundamental mode): this can easily be done and we leave it to the reader.

3.2 Equivalence of Functional and Recognition problems

In this section, we consider three important functional problems, and in each case give a related recognition problem that is polynomially equivalent to the functional problem. This will give evidence for our claim in chapter one that recognition problems are adequate for the study of tractability.

We have just described the traveling salesman optimization problem. The corresponding recognition problem is

Traveling Salesman Decision Problem (TSD)	
Given:	A matrix D as in TSO and an integer b.
Property:	There exists a tour π such that $D(\pi) \leq b$.

This example illustrates the typical way we describe recognition problems in this book: the well-formed inputs for the problem are described under the heading Given, where some fixed but reasonable encoding (as discussed in section 5, chapter 1) of these inputs is implicitly assumed. We then identify the recognition problem with the language whose members encode those well-formed inputs satisfying the predicate given under the heading Property. We want to show the following

Proposition A. *The TSO problem is fundamentally tractable iff the TSD is in P .*

Suppose that the TSO problem is fundamentally tractable. To solve the decision problem (on input $\langle D, b\rangle$), first find the optimal tour π_{opt}. Compute $D(\pi_{opt})$ and accept iff $D(\pi_{opt}) \leq b$. Under reasonable assumptions, and using the fact that π_{opt} can be found in polynomial time, the described procedure for the TSD clearly takes polynomial time.

Conversely, suppose that TSD is fundamentally tractable. Let R be an algorithm that solves the TSD problem in polynomial time in the fundamental mode. First, we find $b_{\min}$, the minimal tour length using a binary search of the range [0, m] where m is the sum of all the finite entries of D: each probe of the binary search involves a call to the algorithm R as subroutine. The number of probes is $O(\log m)$ which is linear in the size of the input. Next we construct the minimal tour as follows. We regard D as representing a directed graph $G(D)$ on n vertices where an edge $\{i, j\}$ is present iff $d_{i,j} < \infty$. Now we successively 'test' each edge $\{i, j\}$ (in any order) to see if the removal of that edge would result in an increase in the minimal tour length. More precisely, we replace $d_{i,j}$ with ∞ and call the subroutine R to see if the minimal tour length of the modified matrix D exceeds $b_{\min}$. If it does not increase the minimal tour length, we will permanently remove edge $\{i, j\}$; otherwise we conclude that the edge $\{i, j\}$ is necessary in any minimal tour involving the remaining edges of $G(D)$. Proceeding in this fashion, we will terminate after at most $n^2 - 2n$ tests. Note that exactly n edges will remain in the final graph, and this gives the required tour. The entire procedure is seen to require only a polynomial number of calls to R. This completes the proof of the proposition.

The transformation of the TSO problem into TSD is typical. We now illustrate a similar transformation on a graph coloring problem. Let $G = (V, E)$ be an (undirected) graph. For any positive integer k, a function $C : V \to \{1, \ldots, k\}$ is called a *k-coloring* of G if adjacent vertices in G are assigned distinct 'colors', i.e., $\{u, v\} \in E$ implies $C(u) \neq C(v)$. If a k-coloring exists for G, then G is *k-colorable*. The *chromatic number* of G denoted by $\chi(G)$ is the minimum number k such that G is k-colorable. The *graph coloring problem* asks for a coloring of an input graph G using the $\chi(G)$ colors. The *chromatic number problem* is to determine $\chi(G)$ for any input graph G. Both problems are abstractions of problems such as constructing examination time tables: the vertices represent courses such as *Computer Science I* or *Introduction to Ichthyology*. We have an edge $\{u, v\}$ if there are students taking both courses u and v. Nodes colored with the same color will hold examinations in the same time slot. So the graph is k-colorable iff there is a conflict-free examination time table with k time slots. Determining the chromatic number is clearly an optimization problem and again there is no known polynomial-time algorithm for it. We can derive a corresponding recognition problem as follows:

Graph Colorability Problem	
Given:	Graph G and integer k.
Property:	G is k-colorable.

Proposition B. *The graph coloring problem, the chromatic number problem and the graph colorability are polynomially equivalent.* (see Exercises for proof)

The last pair of problems considered in this section relate to the satisfiability problem that arises in the subjects of logic and mechanical theorem proving. The *satisfying assignment problem* is the problem where, given a CNF formula F, we are required to find an assignment that satisfies F (if F is satisfiable). The definitions for this problem appear in the appendix of this chapter. The recognition problem version of this problem is the following:

SATISFIABILITY PROBLEM (SAT)	
GIVEN :	A CNF formula F.
PROPERTY :	F is satisfiable.

Proposition C. *The satisfying assignment problem and* SAT *are polynomially equivalent.* (see Exercises for a proof)

EXERCISES

Exercise 0.1: Consider the Euclidean Traveling Salesman Problem (ETSP) where you are given a graph $G = (V, E)$ as before, and each vertex $v \in V$ is a point $v = (v_x, v_y)$ in the Cartesian plane. The distance of any edge $(u, v) \in E$ is given by its Euclidean distance $\sqrt{(u_x - v_x)^2 + (u_y - v_y)^2}$.

We will assume that the coordinates of points are integers in binary notation. The "size" of a point v is $|v| := 1 + \lg |v_x \cdot v_y|$. The size of G is $(\sum_{v \in V} |v|) + |E|$ (the contribution of E to the size is chosen to be $|E|$ because the edges $(u, v) \in E$ may be represented as a pair of numbers $(\#u, \#v)$ where $\#v \in \{1, 2, \ldots, |V|\}$ is the index of v.

(i) It is not known if ETSP is in *NP*. Prove that ETSP problem is *PSPACE*.

(ii) Prove that ETSP is *NP*-hard. ♢

END EXERCISES

3.3 Many-One Reducibility

In showing that the TSO is tractable if the TSD is tractable (or vice-versa), we have illustrated the very important idea of 'reducing one problem to another'. In this section, we formalize one version of this idea; chapter 4 embarks on a more systematic study.

We now use the concept of a transformation and transducer as defined in chapter 2 (section 2).

Definition 1. Let $(\Sigma, L), (\Gamma, L')$ be languages, and $f : \Sigma^* \to \Gamma^*$ be a transformation computed by a transducer M. We say that L is *many-one reducible* to L' *via* f (or *via* M) if for all $x \in \Sigma^*$,

$$x \in L \text{ iff } f(x) \in L'.$$

If, in addition, M runs in deterministic polynomial time, then we say L is *many-one reducible* to L' *in polynomial time* and write

$$L \leq_m^P L'.$$

■

The above reducibility is also known as *Karp reducibility*. To illustrate this reducibility, we consider a simplification of the satisfiability problem. If k is any positive integer, a kCNF formula is one whose clauses have exactly k literals. Consider the problem of recognizing satisfiable 3CNF formulas: call this problem 3SAT. Thus 3SAT $\subseteq$ SAT (as languages). We show:

LEMMA 1. SAT $\leq_m^P$ 3SAT.

In a formal proof of this lemma, we would have to construct a deterministic transducer M running in polynomial time such that for any input word x over the alphabet of SAT, if x encodes (resp. does not encode) a satisfiable CNF formula, then $f_{\mathrm{M}}(x)$ encodes (resp. does not encode) a satisfiable 3CNF formula. First of all, note that it is not hard for M to verify whether the input is well-formed (i.e. encodes a CNF formula) or not. If x is ill-formed, M can easily output a fixed unsatisfiable formula. So assume that x does encode a CNF formula F. We will show how

to systematically construct another 3CNF formula G such that G is satisfiable iff F is. The actual construction of the transducer M to do this is a straightforward but tedious exercise which we omit. (However, the reader should convince himself that it can be done.)

For each clause C in F we will introduce a set of clauses in G. Let $C = \{u_1, u_2, \ldots, u_m\}$. First assume $m > 3$. We introduce the set of clauses

$$\{u_1, u_2, z_1\}, \{\bar{z_1}, u_3, z_2\}, \{\bar{z_2}, u_4, z_3\}, \ldots, \{\bar{z}_{m-4}, u_{m-2}, z_{m-3}\}, \{\bar{z}_{m-3}, u_{m-1}, u_m\}$$

where $z_i, (i = 1, \ldots, m-3)$ are new variables. If $m = 3$, then we simply put C into G. If $m = 2$, we introduce the following two clauses:

$$\{u_1, u_2, z_1\}, \{\bar{z_1}, u_1, u_2\}.$$

If $m = 1$, we introduce the following four clauses:

$$\{u_1, z_1, z_2\}, \{u_1, \bar{z_1}, z_2\}, \{u_1, z_1, \bar{z_2}\}, \{u_1, \bar{z_1}, \bar{z_2}\}.$$

G has no other clauses except those introduced for each C in F as described above. To show that G is satisfiable iff F is satisfiable, it is easy to verify that for any satisfying assignment I to the variables occurring in F, we can extend I to the newly introduced variables (the z_i's) so that the extension also satisfies all the clauses in G. Conversely, if I is an assignment that satisfies G, then the restriction of I to the variables occurring in F will satisfy F. The reader should check these claims. This concludes the proof of lemma 1.

The following two lemmas concerning $\leq_m^P$ are basic.

LEMMA 2. *If $L \leq_m^P L'$ and $L' \in P$ then $L \in P$.*

Proof. By definition, there is a deterministic transducer M running in $p(n)$ time (for some polynomial $p(n)$) such that $L \leq_m^P L'$ via f_{M}. Also, there exists an acceptor M′ which accepts L' in $p'(n)$ time (for some polynomial $p'(n)$). We construct a machine N which accepts L by 'calling' M and M′ as 'subroutines'. On input x, N first computes $f_{\mathrm{M}}(x)$ by simulating M. Then N imitates the actions of M′ on input $f_{\mathrm{M}}(x)$, accepting iff M′ accepts. Clearly N accepts L and runs in $O(p(n) + p'(p(n)))$ time, which is still polynomial. **Q.E.D.**

LEMMA 3. *$\leq_m^P$ is transitive.*

Proof. Let $L \leq_m^P L'$ via some transducer M and $L' \leq_m^P L''$ via some M′, where both M and M′ run in polynomial time. The lemma follows if we construct another polynomial-time M″ such that L is many-one reducible to L'' via M″. This is done in a straightforward manner where M″ on input x simulates M on x, and if y is the output of M then M″ simulates M′ on y. The output of M″ is the output of M′ on y. Clearly $|y|$ is polynomial in $|x|$ and hence M″ takes time polynomial in $|x|$. **Q.E.D.**

The next definition embodies some of the most important ideas in this subject.

Definition 2. Let K be a class of languages. A language L' is *K-hard* (*under $\leq_m^P$*) if for all $L \in K$, $L \leq_m^P L'$. L' is *K-complete* (*under $\leq_m^P$*) if L' is K-hard (under $\leq_m^P$) and $L' \in K$. ■

Since we will not consider other types of reducibilities in this chapter, we will omit the qualification 'under $\leq_m^P$' when referring to hard or complete problems. We also say L is *hard* (resp. *complete*) for K if L is K-hard (resp. K-complete). From lemma 2, we have

COROLLARY 4. *Let L be NP-complete. Then $L \in P$ iff $P = NP$.*

Thus the question whether $P = NP$ is reduced to the fundamental tractability of *any* *NP*-complete problem. But it is not evident from the definitions that *NP* contains any complete problem.

3.4 Cook's Theorem

In this section we shall prove

THEOREM 5 (Cook). SAT *is NP-complete.*

This was historically the first problem shown to be *NP*-complete[3] and it remains a natural problem whereby many other problems L can be shown to be *NP*-hard by reducing SAT to L (perhaps by using transitivity). This avoids the tedious proof that would be necessary if we had to directly show that every language in *NP* can be reduced to L. We just have to go through this tedium once, as in the proof of Cook's theorem below.

To show that a problem L is *NP*-complete we have to show two facts: that L is in fact in *NP* and that every problem in *NP* can be reduced to L in polynomial time. For most *NP*-complete problems the first fact is easy to establish. In particular, the three recognition problems in section 2 are easily shown to be in *NP*: for the TSD problem, we can easily construct a Turing acceptor which on input $\langle D, b\rangle$ uses nondeterminism to guess a tour π and then deterministically computes the tour length $D(\pi)$, accepting iff $D(\pi) \leq b$. To see that this is a correct nondeterministic procedure, we note that if the input is in TSD, then there is a π which satisfies $D(\pi) \leq b$ and the acceptor will accept since it will guess π along some computation path. Conversely, if the input is not in TSD, then every choice of π will lead to a rejection and by definition, the acceptor rejects the input. Since the procedure clearly takes only a polynomial number of steps, we conclude that TSD is in *NP*. Similarly, for the graph colorability problem (resp. SAT), the acceptor guesses a coloring of the vertices (resp. an assignment to the variables) and verifies if the coloring (resp. assignment) is valid. Let us record these observations in:

LEMMA 6. *The TSD, the graph colorability problem and SAT are in NP.*

We will use simple Turing machines as defined in chapter 2. In particular, the transition table $\delta(M)$ of simple Turing acceptor M is a set of quintuples

$$\langle q, b, q', b', d\rangle$$

saying that in state q and scanning b on the tape, M can move to state q', change b to b', and move the tape-head in the direction indicated by $d \in \{-1, 0, +1\}$. In chapter 2 we showed that a multi-tape machine can be simulated by a 1-tape machine with at most a quadratic blow-up in the time usage. A very similar proof will show:

LEMMA 7. *If M is a multi-tape Turing acceptor accepting in time $t(n)$ then there exists a simple Turing acceptor N accepting the same language in time $(t(n)+n)^2$. Moreover, N is deterministic if M is deterministic.*

The import of this lemma is that for the purposes of defining the class of languages accepted in polynomial time in deterministic (resp. nondeterministic) mode we could have used the simple Turing machines instead of multitape Turing machines. This is also a simple demonstration of the polynomial simulation thesis described in chapter 1.

Proof that every language $L \in NP$ is reducible to SAT: The rest of this section is devoted to this proof. By the preceding lemma, we may assume that L is accepted by a simple Turing acceptor M in time $p(n)$ for some polynomial $p(n)$. To show that L is reducible to SAT, we must show a transducer N (computing the transformation f_N) running in polynomial time such that for each input word x over the alphabet of L, $f_N(x)$ is in SAT iff x is in L. The word $f_N(x)$ will encode a CNF formula. We shall describe this formula only, omitting the tedious construction of N. As usual, once the transformation f_N is understood, the reader should have no conceptual difficulty in carrying out the necessary construction.

For the following discussion, let x be a fixed input of M where $|x| = n$. Let

$$C_0, C_1, \ldots, C_{p(n)} \tag{1}$$

be the first $p(n)$ configurations in some computation path of M on x. (If the computation path has less than $p(n)$ configurations, then its last configuration is repeated.) The CNF formula (encoded by) $f_N(x)$ will specify conditions for the existence of an accepting computation path of the form (3.1), i.e., the formula will be satisfiable iff there exists an accepting computation path (3.1). In order to do this, $f_N(x)$ uses a large number of Boolean variables which we now describe. Each variable denotes a proposition ('elementary statement') about some hypothetical computation path given by (3.1). We will assume that the instructions (tuples) in $\delta(M)$ are numbered in some canonical way from 1 to $|\delta(M)|$.

Variable	*Proposition*
$S(q,t)$	Configuration C_t is in State q.
$H(h,t)$	The tape-Head is scanning cell h in configuration C_t.
$T(b,h,t)$	Cell h contains the Tape symbol b in configuration C_t.
$I(j,t)$	Instruction j in $\delta(M)$ is executed in the transition $C_t \vdash C_{t+1}$.

[3]Cook's theorem in Complexity Theory is comparable to Gödel's Incompleteness Theorem in Computability Theory: both play a paradigmatic role (in the sense of Kuhn [10]). In Kuhn's analysis, a *scientific paradigm* involves a system of views, methodology and normative values for doing research. A further parallel is that the relation of P to *NP* can be compared to the relation between the recursive sets and the recursively enumerable sets.

In this table, the meta-variable[4] t is the 'time indicator' that ranges from 0 to $p(n)$, the meta-variable q ranges over the states in M, and the meta-variable$b$ ranges over the tape symbols (including the blank) of M. The meta-variable h is the 'head position indication': it ranges over the $2p(n)+1$ values

$$-p(n), -p(n)+1, \ldots, -1, 0, 1, 2, \ldots, p(n)-1, p(n)$$

since in $p(n)$ steps, M cannot visit any cell outside this range. The meta-variable j is the "instruction indicator" that ranges from 1 to $|\delta(M)|$, the number of instructions in $\delta(M)$. Since the number of values for q, b and j is $O_{\mathrm{M}}(1)$, and there are $O(p(n))$ values for t and h, we see that there are $O((p(n))^2)$ variables in all. Therefore each variable can be encoded by a word of length $O(\log n)$. We emphasize that these are *Boolean* variables; so the above propositions associated with them are purely informal.

We now introduce clauses into the CNF formula $f_N(x)$ that enforce the above interpretation of the Boolean variables. By this we mean that if v is a Boolean variable that stands for the proposition P_v as described in the above table then for any satisfying assignment I to $f_N(x)$, there exists an accepting computation path of the form (3.1) such that $I(v)=1$ precisely when the proposition P_v is true of (3.1). For instance, the variable $S(q_0, 0)$ stands for the proposition

The state in configuration C_0 is the initial state q_0.

It is clear that this proposition is true of (3.1). Hence we must set-up $f_N(x)$ so that if assignment I satisfies it then necessarily $I(S(q_0, 0)) = 1$.

The clauses in $f_N(x)$ correspond to the following eight conditions. For each condition we specify a set of clauses that we will include in $f_N(x)$.

1. The configuration C_0 is the initial configuration of M on input x. Suppose that $x = a_1 a_2 \cdots a_n$ where a_i are input symbols. To enforce condition 1, we introduce the following set of clauses, each containing only one variable:

   $$\{S(q_0, 0)\}, \quad \{H(1, 0)\}$$

   $$\{T(a_1, 1, 0)\}, \{T(a_2, 2, 0)\}, \ldots, \{T(a_{n-1}, n-1, 0)\}, \{T(a_n, n, 0)\}$$

   $$\{T(\sqcup, h, 0)\} \quad \text{for} \quad h \notin \{1, \ldots, n\}$$

   For the first two clauses to be true, the variables $S(q_0, 0)$ and $H(1, 0)$ must be assigned 1. This ensures the initial state is q_0 and the head is initially scanning a_1. Similarly, the next n clauses ensure that the input x is contained in cells 1 to n. The remaining clauses ensure that the rest of the cells contain the blank symbol $\sqcup$.

2. In each configuration C_t, M is in exactly one state.
   Let S_t be the set of variables $\{S(q, t) : q \text{ is a state of M}\}$. Then condition 2 amounts to ensuring that exactly one variable in S_t is assigned '1'. For this, we have a convenient abbreviation. If X is any set of Boolean variables, let $\mathbf{U}(X)$ denote the set of clauses consisting of the clause X and the clause $\{\bar{u}, \bar{v}\}$ for each pair of distinct variables u, v in X. For example, if $X = \{x, y, z\}$ then

   $$\mathbf{U}(X) = \{\{x, y, z\}, \{\bar{x}, \bar{y}\}, \{\bar{x}, \bar{z}\}, \{\bar{y}, \bar{z}\}\}.$$

   Note that an assignment to X satisfies $\mathbf{U}(X)$ iff the assignment assigns a value of '1' to exactly one variable in X. ($\mathbf{U}$ stands for 'unique'.) Therefore, for each t, condition 2 is enforced by introducing into $f_N(x)$ all the clauses in $\mathbf{U}(S_t)$.

3. There is a unique symbol in each cell h of each C_t.
   For each h, t, condition 3 is enforced by the clauses in $\mathbf{U}(T_{h,t})$ where $T_{h,t}$ is defined to be the set $\{T(b, h, t) : b \text{ is a tape symbol of M}\}$.

4. There is a unique head position h in each C_t. This is enforced by the clauses in $\mathbf{U}(H_t)$ where $H_t = \{H(h, t) : h = -p(n), -p(n)+1, \ldots, -1, 0, 1, \ldots, p(n)\}$.

5. The last configuration is accepting.
   Introduce the single clause $\{S(q_a, p(n))\}$ where q_a is the accept state.

[4] I.e., the variable we use in describing actual Boolean variables of $f_N(x)$. These meta-variables do not actually appear in $f_N(x)$.

6. Cells that are not currently scanned in C_t must contain the same symbol as in C_{t+1}.
   This can be ensured by introducing the three-literal clause

$$\{\overline{T(b,h,t)}, \overline{T(c,h,t+1)}, H(h,t)\}.$$

   for all h, t and tape symbols b, c, where $b \neq c$.

7. For each $t < p(n)$ there is a unique instruction that causes the transition $C_t \vdash C_{t+1}$.
   This is ensured by the clauses in $\mathbf{U}(I_t)$ where $I_t = \{I(j,t) : j = 1, \ldots, |\delta(M)|\}$.

8. Changes in successive configurations follow according to the transition table of M.
   For each state q, and tape symbol b (possibly $\sqcup$), let $\delta(q,b) \subseteq \{1, \ldots, |\delta(M)|\}$ be the set instruction numbers such that $j \in \delta(q,b)$ iff the first two components of the jth instruction are q and b. In other words, $j \in \delta(q,b)$ iff the jth instruction is applicable to any configuration whose state and scanned symbol are q and b respectively.
   For each t, h, we introduce this clause:

$$\{\overline{T(b,h,t)}, \overline{S(q,t)}, \overline{H(h,t)}\} \cup \{I(j,t) \ : \ j \in \delta(b,q)\}.$$

   In addition, for each j where the jth instruction is $\langle q, b, q', b', d\rangle$, introduce three clauses:

$$\{\overline{I(j,t)}, S(q', t+1)\},$$

$$\{\overline{I(j,t)}, \overline{H(h,t)}, T(b', h, t+1)\}$$

$$\{\overline{I(j,t)}, \overline{H(h,t)}, H(h+d, t+1)\}$$

   We ought to remark that to ensure that if the machine gets stuck or halts in less than $p(n)$ steps, then we assume there are rules in $\delta(M)$ to replicate such configurations.

It is curious to observe that, with the exception of the first condition, the rest of the clauses in $f_N(x)$ are not directly related to the input x.

To show that the CNF formula $f_N(x)$ consisting of all the clauses introduced under the above eight conditions is correct, we have to show that $f_N(x)$ is satisfiable iff x is in L. Suppose $f_N(x)$ is satisfiable by some assignment I. Then we claim that an accepting computation path of the form (3.1) exists. This is seen inductively. Clearly the configuration C_1 is uniquely 'determined' by condition 1. Suppose that $C_1, \ldots, C_t$ are already defined. Then it is easy to define C_{t+1}. Hence sequence (3.1) can be defined. But condition 5 implies that $C_{p(n)}$ is accepting. Conversely, if an accepting computation path exists then it determines an assignment that satisfies $f_N(x)$.

Finally, we indicate why N can be constructed to run in polynomial time. Note that for any set X of k variables, the formula $\mathbf{U}(X)$ contains $O(k^2)$ clauses, and these clauses can be generated in $O(k^3)$ time. It is not hard to verify that all the clauses in conditions 1-8 can be generated in $O((p(n))^3)$ time. This concludes the proof of Cook's theorem.

We have the following interesting by-product of the above proof:

COROLLARY 8. *There is a polynomial-time computable transformation t of an arbitrary Boolean formula F to a* CNF *formula $t(F)$ such that F and $t(F)$ are co-satisfiable (i.e., F is satisfiable iff $t(F)$ is satisfiable).*

See the Exercises for a direct proof of this result.

3.5 Some Basic *NP*-Complete Problems in Graph Theory

We study a few more *NP*-complete problems. This will give the reader a feeling for some of the techniques used in proving *NP*-completeness. We typically prove that a language L is *NP*-hard by reducing a known *NP*-hard problem L' to L. This indirect approach is particularly effective if one chooses an L' that is rather similar to L: the list of over 300 *NP*-complete problems in [5] is a good resource when making this choice. We will assume the standard terminology of graph theory and throughout this book, the terms 'node' and 'vertex' are fully interchangeable. The following four problems will be considered in this section.

VERTEX COVER PROBLEM.	
GIVEN:	Graph $G = (V, E)$ and integer k.
PROPERTY:	G has a vertex cover of size $\leq k$.

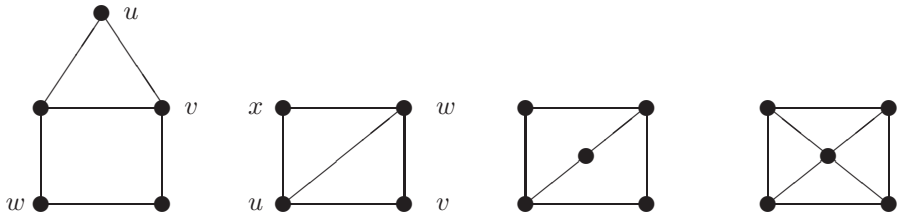

(a) (b) (c) (d)

Figure 3.1: Some graphs

A *vertex cover* of G is a subset $V' \subseteq V$ such that $\{u, v\} \in E$ implies that u or v is in V'. For example, the graph in Figure 3.1(a) has a vertex cover $\{u, v, w\}$ of size 3 but none of size 2.

CLIQUE PROBLEM.	
GIVEN:	$G = (V, E)$ and k.
PROPERTY:	There exists a clique of size $\geq k$ in G.

A *clique* of G is a subset $V' \subseteq V$ such that each pair of distinct vertices u and v in V' are adjacent in G. For example, the graph in Figure 3.1(b) has two cliques $\{u, v, w\}$ and $\{u, x, w\}$ of size 3 but none of size 4.

INDEPENDENT SET PROBLEM.	
GIVEN:	$G = (V, E)$ and k.
PROPERTY:	There exists an independent set of size $\geq k$ in G.

An *independent set* of G is a subset $V' \subseteq V$ such that no two distinct vertices u and v in V' are adjacent in G. Thus, the graph in Figure 3.1(b) has an independent set $\{v, x\}$ of size 2 but none of size ≥ 3.

HAMILTONIAN CIRCUIT PROBLEM.	
GIVEN:	$G = (V, E)$.
PROPERTY:	G has a Hamiltonian circuit.

A *Hamiltonian circuit* $C = (v_1, v_2, \ldots, v_n)$ of G is a cyclic ordering of the set of vertices of G such that $\{v_i, v_{i+1}\}$ (for $i = 1, \ldots, n$) are edges of G (here v_{n+1} is taken to be v_1). For instance, the graph in Figure 3.1(d) has a Hamiltonian circuit, but the one in Figure 3.1(c) has none.

These four problems are easily seen to be in *NP* using the usual trick of guessing and verifying. The first three problems appear very similar. The precise relationship between vertex covers, cliques and independent sets is given by the following result. The *complement* of $G = (V, E)$ is co-$G = (\bar{V}, \bar{E})$ where $\bar{V} = V$, $\bar{E} = \{\{u, v\} : u \neq v, \{u, v\} \notin E\}$.

LEMMA 9. *The following statements are equivalent:*

(a) V' is a vertex cover for G.

(b) $V - V'$ is an independent set for G.

(c) $V - V'$ is a clique in co-G.

The proof is straightforward and left to the reader. From this lemma, it is evident that the three problems are inter-reducible problems. So it is sufficient to show one of them *NP*-complete. We begin by showing that Vertex Cover is *NP*-hard; this was shown by Karp.

THEOREM 10. 3SAT $\leq_m^P$ *Vertex Cover.*

(a) pair (b) triangle

Figure 3.2: Constructions for vertex cover

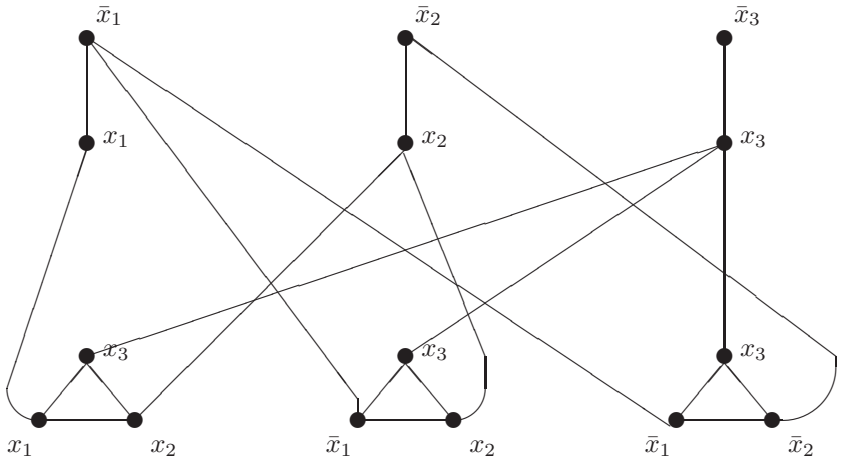

Figure 3.3: Reduction of 3SAT to Vertex Cover

Proof. Let F be a 3CNF formula with m clauses and n variables. We will describe a graph $G = (V, E)$ derived from F such that G has a vertex cover of size at most $2m + n$ iff F is satisfiable. The actual construction of a transducer to accomplish the task is a routine exercise which we omit.

The graph G contains two types of subgraphs: For each variable x in F, we introduce a pair of adjacent nodes labeled by x and $\bar{x}$, as in Figure 3.2(a).

For each clause $\{u_1, u_2, u_3\}$, we introduce a triangle with nodes labeled by the literals in the clause (Figure 3.2(b)). To complete the description of G, we introduce edges connecting a node in any pair with a node in any triangle whenever the two nodes are labeled by the same literal. For example, the graph for the formula $\{\{x_1, x_2, x_3\}, \{\bar{x}_1, x_2, x_3\}, \{\bar{x}_1, \bar{x}_2, x_3\}\}$ is shown in Figure 3.3.

We now claim that F is satisfiable iff G has a vertex cover of size $2m + n$. If F is satisfied by some assignment I, then we choose V' to consist of (a) those nodes in the pairs of G labeled by literals u where $I(u)$=1 and (b) any two nodes from each triangle of G, provided that all the nodes labeled by literals u where $I(u)$=0 (there are at most two such per triangle) are among those chosen. There are n nodes chosen in (a) and $2m$ nodes chosen in (b). Furthermore, one observes that every edge of G is incident to some node in V', and thus V' is a vertex cover. Conversely, if G has a vertex cover V' of size at most $2m + n$, then we easily see that each pair (resp. triangle) must contain at least one (resp. two) nodes in V'. This means that in fact the vertex cover has exactly one vertex from each pair and two vertices from each triangle. The assignment which assigns to each variable x a value of 1 (resp. 0) iff the node in a pair labeled by x (resp. $\bar{x}$) is in V' is seen to satisfy F. **Q.E.D.**

We next show that the Hamiltonian Circuit problem is *NP*-hard, a result of Karp.

THEOREM 11. *Vertex Cover* $\leq_m^P$ *Hamiltonian Circuit.*

Proof. Given $G = (V, E)$ and k, we construct $\bar{G} = (\bar{V}, \bar{E})$ such that G has a vertex cover with at most k nodes iff $\bar{G}$ has a Hamiltonian circuit. $\bar{G}$ has two types of vertices. The first type consists of k nodes $a_1, \ldots, a_k$ connected in a ring, i.e., a_i is adjacent to both a_{i+1} and a_{i-1} where subscript arithmetic is modulo k. See Figure 3.4(a).

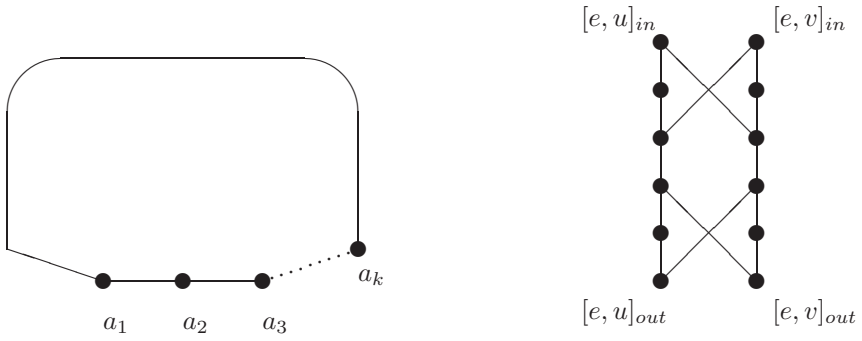

(a) ring (b) gadget

Figure 3.4: Constructions for Hamiltonian circuit

The second type of vertices is introduced by constructing a 'gadget' of 12 nodes for each edge $e = \{u, v\} \in E$ (Figure 3.4(b)). Four nodes in the gadget are specially labeled as

$$[e,u]_{in}, [e,u]_{out}, [e,v]_{in}, [e,v]_{out}.$$

So the total number of vertices in $\bar{V}$ is $k + 12|E|$. In addition to the edges in the ring and the gadgets, other edges are introduced corresponding to each node u in V thus: let $e_1, \ldots, e_m$ be the edges in E which are incident on u. We form a 'u-chain' by introducing the following edges:

$$\{[e_1,u]_{out}, [e_2,u]_{in}\}, \{[e_2,u]_{out}, [e_3,u]_{in}\}, \ldots, \{[e_{m-1},u]_{out}, [e_m,u]_{in}\}. \quad (2)$$

These edges string together the gadgets corresponding to the edges $e_1, \ldots, e_m$. This stringing of the gadgets imposes an arbitrary ordering of the e_i's. We also connect $[e_1,u]_{in}$ and $[e_m,u]_{out}$ to each of the nodes $a_i (i = 1, \ldots, k)$ in the ring. See Figure 3.5. Our description of $\bar{G}$ is now complete. It remains to show that G has a vertex cover of size at most k iff $\bar{G}$ has a Hamiltonian circuit. Suppose $U = \{u_1, \ldots, u_h\}$ is a vertex cover of G, $h \leq k$. For each node u_j in U, we define a path p_j from a_j to a_{j+1}. Suppose $e_1, \ldots, e_m$ are the edges in E which are incident on u_j. Let $e_1, \ldots, e_m$ appear in the order as determined by the u_j-chain. The path p_j will include all the edges in (3.2) and the two edges $\{a_j, [e_1,u_j]_{in}\}$ and $\{[e_m,u_j]_{out}, a_{j+1}\}$. The path p_j must also connect the nodes $[e_i,u_j]_{in}$ and $[e_i,u_j]_{out}$ for each $i = 1, \ldots, m$. We consider two ways to do this:

The first way (Figure 3.6(a)) visits every node in the gadget of e_i but the second way (Figure 3.6(b)) visits only half of the nodes. We route p_j through the gadget of e_i using the first way if the other node that e_j is incident upon is not in the vertex cover U; otherwise we use the second way. This completes the definition of the path p_j. The concatenation of $p_1, p_2, \ldots, p_h$ is seen to be a path $P(U)$ from a_1 to a_{h+1} which visits every node in each gadget. Finally, this path $P(U)$ is made into a circuit by connecting $a_{h+1}, a_{h+2}, \ldots, a_k$ and a_1 using edges in the ring structure. The result is a Hamiltonian circuit for $\bar{G}$.

Conversely, suppose there is a Hamiltonian circuit C for $\bar{G}$. The circuit is naturally broken up into paths whose end-points (but not other vertices) are in the ring. For each such path p that is non-trivial (i.e., has more than one edge), we claim that there exists a unique vertex $u(p) \in V$, such that p visits all and only the gadgets in the u-chain: this is because if p visits any gadget associated with an edge $e = \{u, v\}$, it must either enter and exit from the nodes $[e,u]_{in}, [e,u]_{out}$ or enter and exit from the nodes $[e,v]_{in}, [e,v]_{out}$. For instance, it is not possible to enter at $[e,v]_{in}$ and exit from $[e,u]_{out}$ (why?). In fact, if p enters and exits from the gadget using the nodes $[e,u]_{in}$ and $[e,u]_{out}$, then there are essentially two ways to do this, as indicated in Figure 3.6. It is easily seen that set of all such vertices $u(p)$ forms a vertex cover $U(C)$ of G. Since there are at most k such paths in C, the vertex cover $U(C)$ has size at most k. **Q.E.D.**

3.6 **Two Hard Problems

[4] ** Optional section. These problems are "hard" in the sense of being non-trivial to show in *NP*.

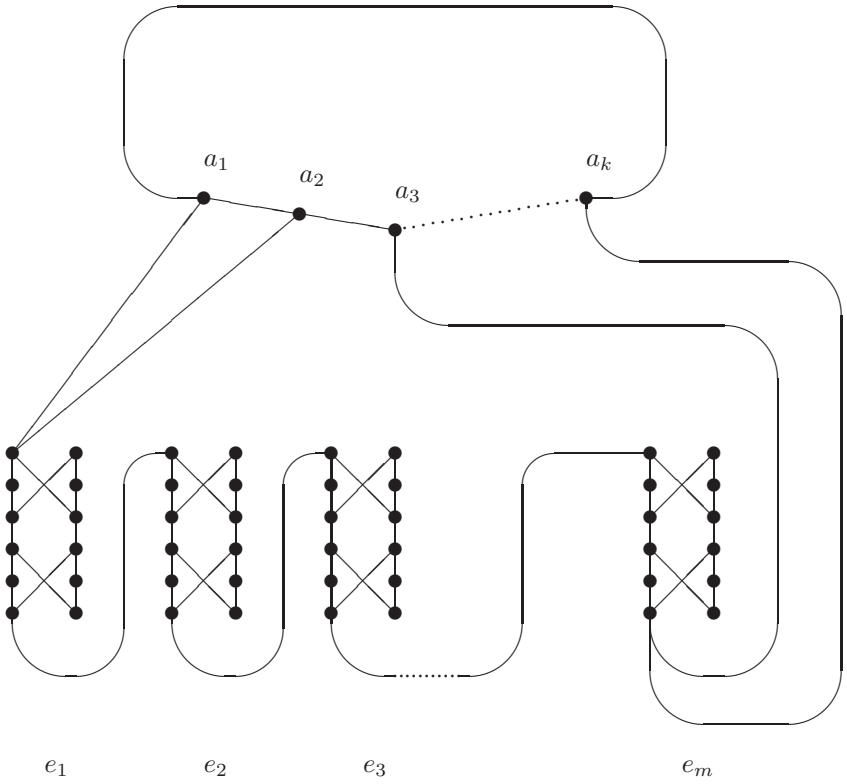

Figure 3.5: A u-chain: $e_1, \ldots, e_m$ are edges incident on node u

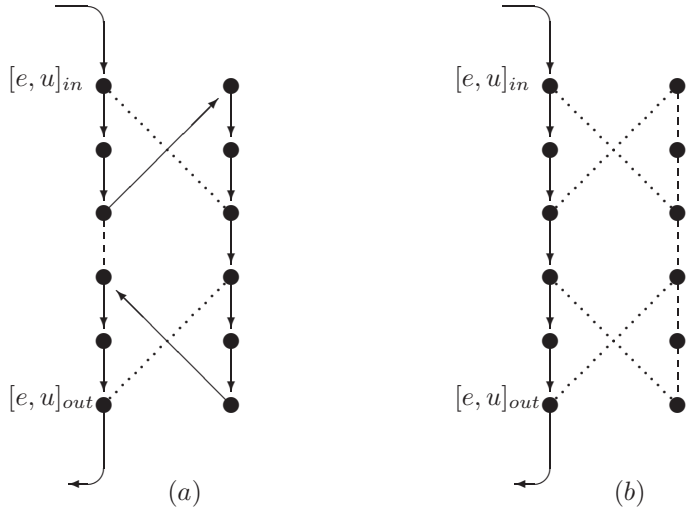

Figure 3.6: The only two ways a Hamiltonian path can route through a gadget

In all the problems seen so far, it was rather easy to show that they are in *NP*, but showing them to be *NP*-hard required a bit more work. We now consider two problems for which it is non-trivial to show that they are in *NP*.

PRIMALITY TESTING (*Primes*)	
Given:	Given a binary number n.
Property:	n is a prime.

INTEGER LINEAR PROGRAMMING (ILP)	
Given:	An integer $m \times n$ matrix A, and an integer m-vector $\mathbf{b}$.
Property:	There is an integer n-vector $\mathbf{x}$ such that the system of linear inequalities $A\mathbf{x} \geq \mathbf{b}$ hold.

These problems have practical importance. One important use of primes is in cryptographic techniques that rely on the availability of large prime numbers (several hundred bits in length). We note first that the complement of *Primes* is *Composites*, the problem of testing if a binary number is composite. It is easy to see that *Composites* is in *NP*: to test if n is composite, first guess a factor m less than n and then check if m divides n exactly. Combined with the result that *Primes* is in *NP*, we conclude

$$Primes \in NP \cap \text{co-}NP. \tag{3}$$

Now *Primes* is not known to be *NP*-complete and (3.3) gives strong evidence that it is not. This is because if it were *NP*-complete then it would be easy[5] to conclude from (3.3) that *NP*=co-*NP*. This is unlikely, for example, based on extensive experience in the area of mechanical theorem proving.

The ILP problem is also known as the *Diophantine linear programming problem* and is studied in [12]. Let *Rational LP* (or RLP) denote the variant of ILP where we allow the input numbers and output numbers to be rationals. The rational LP problem was shown to be in P in 1979 by Khachian [8]; this solved a long standing open problem. See [1] for an account. Previously, all algorithms for the rational linear programming problem had been based on the Simplex Method, and it was shown by Klee and Minty[9] that such an approach can be exponential in the worst case.[6] Khacian's algorithm and other polynomial time algorithms for RLP are called "interior point methods" because they look for solutions that may be in the interior of the defining polytope. Originally, interior point algorithms do not seem to be competitive with the simplex method(s) in practice. This situation has since changed, following extensive development of interior point algorithms especially one such algorithm discovered by Karmarkar.

3.6.1 Primality Testing

This subsection requires some elementary number theory; Hardy and Wright [6] may serve as an excellent reference. Henceforth, "number" means natural number. We follow the usual convention that the number 1 does not count as a prime: thus the smallest prime is the number 2. It is also the only even prime. An integer n **divides** another m, denoted $n|m$, if there is some integer a such that $na = m$. The greatest common divisor of m and n is denoted by $\texttt{GCD}(m, n)$, with the convention $\texttt{GCD}(0, 0) = 0$. Note that $\texttt{GCD}(0, n) = \texttt{GCD}(n, 0) = |n|$. Two integers m, n are **co-prime** if $\texttt{GCD}(m, n) = 1$. E.g., $\texttt{GCD}(18, 14) = 2$ and $16, 27$ are co-prime. For any number $n \geq 1$, let

$$\mathbb{Z}_n = \{0, 1, \ldots, n-1\} \text{ and } \mathbb{Z}_n^* = \{m : 1 \leq m \leq n, \texttt{GCD}(m, n) = 1\}.$$

Notice that we have defined $\mathbb{Z}_n^*$ in such a way that $\mathbb{Z}_n^* \subseteq \mathbb{Z}_n$ is true for all $n \geq 2$. But when $n = 1$, $\mathbb{Z}_1 = \{0\}$ and $\mathbb{Z}_1^* = \{1\}$.

Congruences. We say x and y are **congruent** modulo n, written $x \equiv y(\mod n)$, if $n|x - y$. For any m, we write $(m \bmod n)$ for the unique element in $\mathbb{Z}_n$ that is congruent to m modulo n. The following fact will be useful: given integers a, b, n, *the congruence* $ax \equiv b(\mod n)$ *has a solution in* x *iff*

$$\texttt{GCD}(a, n)|b. \tag{4}$$

[5] See chapter 4.

[6] Recently major progress have been made in understanding the average behavior of the Simplex-based algorithms. They essentially confirm the experimental evidence that, on the average, these methods take time linear in the number m of constraints, and a small polynomial in the dimension n. See [2].

In proof, if $\texttt{GCD}(a,n) = d$ and $d|b$ then $ax \equiv b(\mod n)$ iff $a'x \equiv b'(\mod n')$. $a = da', b = db', n = dn'$. Since $\texttt{GCD}(a',n') = 1$, it follows that, modulo n', a' has an inverse a'^{-1}. So a solution is $x = b'a'^{-1}$. Conversely, if d does not divide b then the equation $b = ax + ny$ does not have any solution (x,y) because d divides the right-hand side but not the left-hand side. But $b = ax + ny$ has a solution iff $ax \equiv b(\mod n)$ has a solution.

The preceding concerns linear congruences. Now suppose $A(x) = \sum_{i=0}^{d} a_i x^i$ is an polynomial ($a_i \in \mathbb{Z}$) and we want to solve the congruence $A(x) \equiv 0(\mod p)$ for some prime p. Without loss of generality, p does not divide a_d. Lagrange shows that $A(x)$ has at most d distinct solutions modulo p. In proof, if $d = 0$ or if $A(x) \equiv 0(\mod p)$ has no solutions modulo p, the result is trivial. Otherwise, let a be a solution. Then (using long division) we can write $A(x) = (x-a)B(x) + A(a)$ where $B(x)$ is an integer polynomial of degree $d-1$ and $A(a) \equiv 0(\mod p)$. For any other solution b of $A(x) \equiv 0(\mod p)$, $A(b) \equiv (b-a)B(b) \equiv 0(\mod p)$ implies b is a solution of $B(x) \equiv 0(\mod p)$. But by induction, there are at most $d-1$ distinct solutions b modulo p.

Euler's Totient Function. Euler's totient function $\phi(n)$ is defined to be $|\mathbb{Z}_n^*|$. For example,

$$\phi(0) = 0, \quad \phi(1) = \phi(2) = 1, \quad \phi(3) = \phi(4) = \phi(6) = 2, \quad \phi(5) = 4.$$

Clearly, $1 \le \phi(n) \le n-1$ with $\phi(n) = n-1$ iff n is prime. It is easy to give a formula for $\phi(n)$ in terms of the prime factorization of n. If $n = \prod_{i=1}^{k} q_i^{e_k}$ where $e_k \ge 1$ and $q_1, \ldots, q_k$ are distinct primes, then

$$\phi(n) = n \prod_{i=1}^{k} \left(1 - \frac{1}{q_i}\right) \tag{5}$$

To see this, note that the proposed product formula for $\phi(n)$ expands into

$$\phi(n) = n - \sum_{1 \le i \le k} \frac{n}{q_i} + \sum_{1 \le i < j \le k} \frac{n}{q_i q_j} - \quad \cdots \quad + (-1)^k \frac{n}{q_1 q_2 \cdots q_k}. \tag{6}$$

But the term $\frac{n}{q_i}$ is the number of numbers in $\{1, 2, \ldots, n\}$ that are divisible by q_i, and the term $\frac{n}{q_i q_j}$ is the number of numbers in $\{1, 2, \ldots, n\}$ that are divisible by both q_i and q_j, etc. By the inclusion exclusion formula, this expression gives the number of numbers in $\{1, 2, \ldots, n\}$ that is not divisible by any of the q_i's. But this number is precisely $\phi(n)$, thus validating the formula (3.5). A corollary of (3.5) is that ϕ is **multiplicative**:

$$\texttt{GCD}(m,n) = 1 \Rightarrow \phi(mn) = \phi(m)\phi(n). \tag{7}$$

In particular, when n is a prime, $\phi(n) = n-1$. On the other hand, we see that $\phi(n) < n-1$ if n is composite. For, if $k = 1$, then a composite n has the form $n = q_1^{e_1}$ for some $e_1 \ge 2$, and $\phi(n) = n - q_1^{e_1 - 1} \le n-2$. For $k > 1$, we may write n as a product $n = n_1 n_2$ of two co-prime numbers, and by induction, $\phi(n_1 n_2) = \phi(n_1)\phi(n_2) \le (n_1 - 1)(n_2 - 1) < n - 2$. We conclude: $\phi(n) \le n-1$ *with equality iff* n *is prime.*

Another identity involving $\phi(n)$ is this:

$$n = \sum_{m|n} \phi(m). \tag{8}$$

In proof, suppose $n = \prod_{i=1}^{k} q_i^{e_i}$ where q_i are distinct primes. First, we claim that the sum in (3.8) is equal to the product

$$\prod_{i=1}^{k} \left(\phi(1) + \phi(q_i) + \phi(q_i^2) + \cdots + \phi(q_i^{e_i})\right) = \sum_{d_1=0}^{e_1} \sum_{d_2=0}^{e_2} \cdots \sum_{d_k=0}^{e_k} \prod_{i=1}^{k} \phi(q_i^{d_i}). \tag{9}$$

To see this, note each summand on the right-hand side has the form $\prod_{i=1}^{k} \phi(q_i^{d_i})$ which, because of the multiplicative nature of ϕ, is equal to $\phi(\prod_{i=1}^{k} q_i^{d_i}) = \phi(m)$ where $m = \prod_{i=1}^{k} q_i^{d_i}$. Clearly $m|n$. and every m that divides n contributes a summand $\phi(m)$ to the right-hand side of (3.9). Thus (3.9) is just another way to express the sum in (3.8), as claimed. Now each factor in the left-hand side of (3.9) is equal to

$$1 + (q_i - 1) + (q_i^2 - q_i) + \cdots + (q_i^{e_i} - q_i^{e_i - 1}) = q_i^{e_i}$$

and hence the left-hand side is equal to $\prod_{i=1}^{k} q_i^{e_i} = n$. This proves (3.8).

A Characterization of Primes. The result that *Primes* is in *NP* is due to Pratt [13]. To motivate his approach, consider a little gem of number theory:

THEOREM 12 (Fermat's Little Theorem). *If p is prime and p does not divide a, then $a^{p-1} \equiv 1(\mod p)$. More generally, for any n and $a \in \mathbb{Z}_n^*$, $a^{\phi(n)} \equiv 1(\mod n)$.*

Proof. Write S for the set $\mathbb{Z}_n^*$. Then the elements in $aS = \{(a \cdot x \bmod n) : x \in S\}$ are all distinct modulo n: otherwise, for $x \neq y \in S$, we have $ax \equiv ay(\mod n)$, or, $a(x-y) \equiv 0(\mod n)$. But this is impossible since $x - y \not\equiv 0(\mod n)$ and $a \not\equiv 0(\mod n)$. Hence $S = aS$ and $\prod_{x\in S} x = \prod_{y\in aS} y = a^{\phi(n)} \prod_{x\in S} x$. Since $\prod_{x\in S} x \neq\equiv 0(\mod n)$, we conclude $1 \equiv a^{\phi(n)}(\mod n)$. **Q.E.D.**

This gives us a negative test for primality: we know n is non-prime if $a^{n-1} \not\equiv 1(\mod p)$. Unless a is given to us, this only shows that the complement of Primality is in *NP*, which we already knew. The question is whether we can strengthen this line of thinking into a complete test. That is precisely what the following key theorem provides.

We say that a **belongs to (the exponent)** d **modulo** n if d is the smallest number such that $a^d \equiv 1(\mod n)$. For example, if $n = 5$ and $a = 3$ then we have $a^2 \equiv 4, a^3 \equiv 3, a^4 \equiv 16 \equiv 1$ (all $\mod 5$). Thus a belongs to $d = 4$ modulo 5. It is easy to see that if a belongs to d modulo n, then

$$a^m \equiv 1(\mod n) \Rightarrow d|m. \tag{10}$$

This is because otherwise $m = sd + t$ for some s, t with $1 \leq t < d$ and $1 \equiv a^m \equiv a^{sd+t} \equiv a^t$, contradiction.

THEOREM 13 (Key Lemma). *Let $n > 2$. Then n is prime if and only if there exists an $a \in \mathbb{Z}_n^*$ that belongs to the exponent $n - 1$.*

We split this theorem into two halves. The easier half of this theorem is due to Lucas [6, p.72].

THEOREM 14. *Let n, a be numbers. If a belongs to the exponent $n - 1$ modulo n then n is prime.*

Proof. Since $a^{n-1} \equiv 1(\mod n)$ yields a solution to the congruence $ax \equiv 1(\mod n)$, we conclude from (3.4) that $\texttt{GCD}(a, n)$ divides 1. So a, n are co-prime. Next, $a^1, a^2, \ldots, a^{n-1}$ are distinct elements $(\mod n)$: otherwise if $a^i \equiv a^{i+m}(\mod n)$ for some $1 \leq i < i + m < n$, then $a^m \equiv 1(\mod n)$, contradiction. But $(a^i \bmod n) \in \mathbb{Z}_n^*$ for each $i = 1, \ldots, p - 1$. Hence $\phi(n) \geq n - 1$. It follows $\phi(n) = n - 1$, which characterizes n as prime. **Q.E.D.**

The harder half of theorem 13 is the statement that if p is prime then there exists an $a \in \mathbb{Z}_p^*$ that belongs to the exponent $p - 1$. Another way to say this is that $\mathbb{Z}_p^*$ is a cyclic multiplicative group with generator a. In fact, there are $\phi(p - 1)$ such generators. For instance, if $p = 17$ then $\phi(16) = 2^4 - 2^3 = 8$. So there are 8 generators of the mulplicative group $\mathbb{Z}_{17}^*$. (Exercise: determine these generators.) All these facts are a consequence of the next result.

THEOREM 15. *Let p be prime and let $d \geq 1$ divide $p - 1$. Then there are exactly $\phi(d)$ elements in $\mathbb{Z}_p^*$ that belongs to exponent d.*

Proof. Each $a \in \mathbb{Z}_p^*$ belongs to some exponent d modulo p. Let $\psi(d)$ denote the number of elements in $\mathbb{Z}_p^*$ that belongs to d modulo p. Our theorem says that if $d|p - 1$ then $\psi(d) = \phi(d)$. First note that $\psi(d) = 0$ if d does not divide $p - 1$: for, if a belongs to d modulo p, and d does not divide $p - 1$ then $p - 1 = sd + t$ for some s, t with $1 \leq t < d$. Then $1 \equiv a^{p-1}$ (Fermat's little theorem) implies $1 \equiv a^{st+t} \equiv a^t(\mod p)$. But $a^t \equiv 1$ and $t < d$ contradicts the assumption that a belongs to exponent d. This proves that

$$p - 1 = \sum_{d|p-1} \psi(d). \tag{11}$$

Observe that (3.8) implies

$$p - 1 = \sum_{d|p-1} \phi(d).$$

From these two expresions for $p - 1$, it would follow that $\psi(d) = \phi(d)$ if show that for all $d|p - 1$,

$$\psi(d) \leq \phi(d). \tag{12}$$

If $\psi(d) = 0$ then (3.12) is trivially true. Otherwise, there exists an a that belongs to d modulo p. Now, the set $S = \{1, a, a^2, \ldots, a^{d-1}\}$ contain d distinct elements modulo p (there is a similar argument in the above proof of Lucas' theorem). Moreover, each element in S is a solution to the equation $x^d \equiv 1(\mod p)$. This is because

$(a^i)^d \equiv (a^d)^i \equiv 1(\mod p)$. But Lagrange's theorem tells us that a polynomial of degree d has at most d distinct solutions modulo p. Hence S contains every element that belongs to d. It remains to show that exactly $\phi(d)$ elements in S has exponent d. This follows if we show a^i has exponent d iff i, d are co-prime. In one direction, if i, d are co-prime, then $(a^i)^j \equiv 1(\mod p)$ implies $d|ij$ (cf. (3.10)) and hence $d|j$. Thus a^i belongs to d. Conversely, if $\texttt{GCD}(i, d) = g > 1$ then $(a^i)^{d/g} \equiv (a^d)^{i/g} \equiv 1^{i/g}(\mod p)$, so that a^i belongs to exponent $\leq d/g$. **Q.E.D.**

Note that theorem 13 is not quite enough to show that primality is in *NP*: to check if n is prime, we guess an $a \in \mathbb{Z}_n$ and verify that (i) $a^{n-1} \equiv 1(\mod n)$ and (ii) for all $m = 1, \ldots, n-2$, $a^m \not\equiv 1(\mod n)$. Unfortunately, this is not polynomial time as there exponentially many choices of m. Pratt notes that we can avoid this exponential check if we have a prime factorization of $n-1$.

LEMMA 16. *Given $n > 2$ and $a \in \mathbb{Z}_n^*$, let $n - 1 = 2^r \prod_{i=1}^k q_i$ be a prime factorization of $n-1$ where the q_i's are odd primes (not necessarily distinct). Then a belongs to exponent $n-1$ modulo n if and only if*
(i) $a^{n-1} \equiv 1(\mod n)$.
(ii) $a^{(n-1)/q_i} \not\equiv 1(\mod n)$ for all $i = 1, \ldots, k$.

Proof. One direction is trivial. In the other direction, we claim that (i) and (ii) implies a belong to the exponent $n-1$. Suppose a belongs to m for some $m = 2, \ldots, n-2$. If m divides $n-1$, then $(n-1)/m$ must have the form qs for some q and s where $q \in \{q_1, \ldots, k\}$. Then $a^m = a^{(n-1)/(sq)} \equiv 1(\mod n)$ implies $a^{(n-1)/q} \equiv (a^m)^s \equiv 1(\mod n)$, contradicting (ii). But if m does not divide $n-1$, then $n-1 = sm+t$ for some numbers s, t and $1 \leq t < m$. This means $1 \equiv a^{n-1} \equiv a^{sm+t} \equiv (a^m)^s a^t \equiv a^t$, contradicting the assumption that a belongs to the exponent m. **Q.E.D.**

We finally show that *Primes* is in *NP*: on input n, we can easily say n is prime or non-prime if $n = 2$ or if n is even. Otherwise, we guess an $a \in \mathbb{Z}_n$ and verify that $\texttt{GCD}(a, n) = 1$. We then guess a prime factorization of $n-1$ as in the preceding lemma: this amounts to guessing the numbers $r, q_1, q_2, \ldots, q_k$ subject to the restriction that $r, k \leq \lg n$ and $3 \leq q_i \leq n$. We finally check that each of the following holds:

(a) $a^{n-1} \equiv 1(\mod n)$.

(b) The product $2^r \prod_{i=1}^k q_i$ is equal to $n-1$.

(c) $a^{\frac{n-1}{q_i}} \not\equiv 1(\mod n)$, for each i.

(d) Recursively, each of the q_i is a prime.

To verify (a) we must raise a to the $(n-1)$st power. Using the squaring technique, we need use only a linear number (i.e., $O(\log n)$) of multiplications of numbers whose magnitudes are each less than n (this is done by repeatedly reducing the intermediate results modulo n). Hence this takes at most cubic time. To verify (b), there are at most $\log n$ factors (counting multiplicities), and again cubic time suffices. Part (c) is similarly cubic time. Hence, if $t(n)$ is the time of this (non-deterministic) algorithm on input n, then

$$t(n) = O(\log^3 n) + \sum_{i=1}^{k} t(p_i)$$

It is easy to verify inductively that $t(n) = O(\log^4 n)$. Thus *Primes* is in *NP*. Of course, if we use faster multiplication algorithms, this bound can be improved.

3.6.2 Integer Linear Programming

We begin by showing that ILP is *NP*-hard. This is quite easy, via a reduction from 3CNF: given a 3CNF formula F consisting of the clauses

$$C_1, C_2, \ldots, C_m$$

we introduce a set of inequalities. For each variable v, if v or its negation $\bar{v}$ occurs in F, we introduce the inequalities:

$$v \geq 0, \quad \bar{v} \geq 0, \quad v + \bar{v} = 1.$$

For each clause $C = \{\alpha, \beta, \gamma\}$, introduce

$$\alpha + \beta + \gamma \geq 1.$$

It is easy to write all these inequalities in the form $A\mathbf{x} \geq \mathbf{b}$ for a suitable matrix A and vector $\mathbf{b}$. Furthermore, this system of inequalities has an integer solution iff F is satisfiable.

The original proof[7] that ILP is in *NP* is by von zur Gathen and Sieveking[14]. To show that ILP is in *NP*, we want to guess an $\mathbf{x}$ and check if the system $A\mathbf{x} \geq \mathbf{b}$ holds on input $A, \mathbf{b}$. However we cannot do this checking in polynomial time if the smallest size solution $\mathbf{x}$ is too large. The basic issue is to show that if there is a solution then there is one of small size. In this context, the *size* of a number or a matrix is the number of bits sufficient to encode it in standard ways. More precisely, we define the size of a number n as logarithm of its magnitude $|n|$; the size of a $m \times n$ matrix is defined as mn plus the sum of the sizes of its entries; the size of the input to ILP is defined as the sum of the sizes of A and $\mathbf{b}$. The proof which takes up the remainder of this subsection requires some elementary linear algebra. The basis of our estimates is the following simple inequality:

LEMMA 17. *Let B be an n square matrix whose entries each have size at most β. Then the determinant $det(B)$ is bounded by $n!\beta^n < s^{2s}$, where s is the size of B.*

This result is seen by summing up the $n!$ terms in the definition of a determinant. In the remainder of this subsection, except noted otherwise, the following notations will be fixed:

$A = (A_{i,j})$	:	an $m \times n$ matrix with integer entries with $m \geq n$.
$\mathbf{b} = (b_i)$	:	an m-vector with integer entries.
α	:	the maximum magnitude (not size) of entries in $A, \mathbf{b}$.
s	:	the size of the A plus the size of $\mathbf{b}$.
$\mathbf{a}_i$	:	the ith row of A.

LEMMA 18 (Key lemma). *Let $A, \mathbf{b}$ be given as above. If there exists an integer solution $\mathbf{x}$ to the system $A\mathbf{x} \geq \mathbf{b}$ then there is a solution $\mathbf{x}$ each of whose entries has magnitude at most $2s^{7s}$.*

It is easy to see that this implies that ILP is in *NP*: First observe that the size s of the input $\langle A, \mathbf{b}\rangle$ satisfies

$$s \geq mn + \log \alpha.$$

On input $\langle A, \mathbf{b}\rangle$, guess some vector $\mathbf{x}$ each of whose entries has magnitude at most $2s^{7s}$. Since each entry has size $O(s^2)$, the guessing takes time $O(ns^2)$. Compute $A\mathbf{x}$ and compare to $\mathbf{b}$, using time $O(mn^2s^3) = O(s^5)$. Accept if and only if $A\mathbf{x} \geq \mathbf{b}$.

We prove some lemmas leading to the Key lemma. The kth *principal submatrix* of any matrix B is the $k \times k$ submatrix of B formed by the first k columns of the first k rows of B; let $B^{(k)}$ denote this submatrix.

LEMMA 19. *Suppose the hth principal submatrix $A^{(h)}$ of A is non-singular. If $h < rank(A)$ then there exists an integer n-vector $\mathbf{z}$, $\mathbf{z} \neq \mathbf{0}$, such that*

(i) $\mathbf{a}_i\mathbf{z} = 0$ for $i = 1, \ldots, h$,

(ii) each of the first $h+1$ components of $\mathbf{z}$ has magnitude at most s^{2s}, but all the remaining components are zero.

Proof. If the $\mathbf{c}$ is the h-vector composed of the first h entries of the $(h+1)$st column of A then there is an h-vector $\mathbf{v} = (v_1, \ldots, v_h)$ such that $A^{(h)}\mathbf{v} = \mathbf{c}$. In fact, Cramer's rule tells us that the entries of $\mathbf{v}$ are given by $v_i = \pm C_i/\Delta$ where C_i, Δ are determinants of the $h \times h$ submatrices of the $h \times (h+1)$ matrix $(A^{(h)}|\mathbf{c})$. We define the required $\mathbf{z}$ by

$$z_i = \begin{cases} v_i\Delta = \pm C_i & \text{if } 1 \leq i \leq h \\ -\Delta & \text{if } i = h+1 \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that $\mathbf{z}$ is a non-zero integer vector and $\mathbf{a}_i\mathbf{z} = \mathbf{0}$ for $i = 1, \ldots, h$. Each z_i has magnitude $\leq s^{2s}$. **Q.E.D.**

LEMMA 20. *Assume $0 \leq h < rank(A)$. Suppose that the hth principal submatrix $A^{(h)}$ is non-singular and $A\mathbf{x} \geq \mathbf{b}$ has an integer solution $\mathbf{x}$ such that for $1 \leq i \leq h$, the ith row of A satisfies*

$$b_i \leq \mathbf{a}_i\mathbf{x} < b_i + s^{4s}. \tag{13}$$

Then there is a matrix A' obtained by permuting the rows and columns of A such that

[7] Our proof is based on notes by Ó'Dúnlaing. Kozen points out that the usual attribution of this result to [3] is incorrect. Borosh and Treybig showed a weaker form of the Key lemma above: namely, if $A\mathbf{x} = \mathbf{b}$ has a solution then it has a small solution. It does not appear easy to derive the Key lemma from this weaker form.

(i) the $(h+1)$st principal submatrix $A'^{(h+1)}$ is non-singular, and

(ii) Let $\mathbf{b}' = (b'_1, \ldots, b'_m)$ be the permutation of $\mathbf{b}$ corresponding to the permutation of the rows of A to derive A'. There is a solution $\mathbf{x}'$ to $A'\mathbf{x}' \geq \mathbf{b}'$ such that for all $1 \leq i \leq h+1$, the following inequality holds:

$$b'_i \leq \mathbf{a}'_i \mathbf{x}' < b'_i + s^{4s}.$$

Proof. In the statement of this lemma, the principal submatrix $A^{(0)}$ is conventionally taken to be non-singular. By permuting the columns of A if necessary, we can assume that the first $h+1$ columns of A are linearly independent. By the previous lemma, there is a non-zero vector $\mathbf{z}$ such that (i) $\mathbf{a}_i\mathbf{z} = 0$ for all $i = 1, \ldots, h$, and (ii) the first $h+1$ components of $\mathbf{z}$ are at most s^{2s} and the remaining components are 0. Since the first $h+1$ columns of A are linearly independent, there exists $i, (h < i \leq n)$, such that $\mathbf{a}_i\mathbf{z} \neq 0$. Assume that there exists some i such that $\mathbf{a}_i\mathbf{z}$ is positive (the case $\mathbf{a}_i\mathbf{z} < 0$ is similar), and let $J \subseteq \{h+1, \ldots, m-1, m\}$ consist of all those j satisfying

$$\mathbf{a}_j\mathbf{z} > 0.$$

Let δ be the smallest non-negative integer such that if $\mathbf{x}'' = \mathbf{x} - \delta\mathbf{z}$ then for some $j_0 \in J$, $b_{j_0} \leq \mathbf{a}_{j_0}\mathbf{x}'' < b_{j_0} + s^{4s}$. We claim that

(a) $\mathbf{a}_i\mathbf{x}'' = \mathbf{a}_i\mathbf{x}$ for $i = 1, \ldots, h$, and

(b) $\mathbf{a}_i\mathbf{x}'' \geq b_i$ for $i = h+1, \ldots, m$.

It is easy to see that (a) follows from our choice of $\mathbf{z}$. As for (b), the case $\delta = 0$ is trivial so assume $\delta > 0$. When $i \notin J$, (b) follows from the fact that $\mathbf{a}_i\mathbf{x}'' \geq \mathbf{a}_i\mathbf{x}$. When $i \in J$, our choice of δ implies that

$$b_i + s^{4s} \leq \mathbf{a}_i\mathbf{x} - (\delta - 1)\mathbf{a}_i\mathbf{z}$$

But $\mathbf{a}_i\mathbf{z} \leq n\alpha s^{2s} \leq s^{4s}$ implies $\mathbf{a}_i\mathbf{x}'' \geq b_i$. This proves (b). By exchanging a pair of rows of $(A|\mathbf{b}|\mathbf{x}'')$ to make the j_0th row of A the new $(h+1)$st row, we get the desired $(A'|\mathbf{b}'|\mathbf{x}')$. **Q.E.D.**

We are finally ready to prove the Key lemma. Let $\mathbf{x} = (x_1, \ldots, x_n)$ be an integer solution to $A\mathbf{x} \geq \mathbf{b}$. For each i, clearly $x_i \geq 0$ or $-x_i \geq 0$. Hence there exists an n square matrix M whose entries are all zero except for the diagonal elements which are all $+1$ or -1 such that $M\mathbf{x} \geq \mathbf{0}$. Let A' be the $(m+n) \times n$ matrix

$$\left(\frac{A}{M}\right)$$

and $\mathbf{b}'$ be the $(m+n)$-vector obtained by appending zeroes to $\mathbf{b}$. Thus $A'\mathbf{x} \geq \mathbf{b}'$ has a solution. But observe that A' has at least as many rows as columns and it has full rank (equal to the number n of columns). Hence the previous lemma is applicable to $A', \mathbf{b}'$ successively, for $h = 0, 1, \ldots, n-1$. This shows the existence of an integer solution $\mathbf{x}'$ to $A''\mathbf{x}' \geq \mathbf{b}'$, where A'' is a row and column permutation of A' and the individual components of $A''\mathbf{x}'$ are bounded by $s + s^{4s} \leq 2s^{4s}$. Since $rank(A'') = n$, there is an $n \times n$ nonsingular submatrix B of A''. If $B\mathbf{x}' = \mathbf{c}$ then $\mathbf{x}' = B^{-1}\mathbf{c}$.

We now bound the size of the entries of $\mathbf{x}'$. We claim that each entry of B^{-1} has magnitude at most s^{2s}: to see this, each entry of B^{-1} has the form $B_{i,j}/\Delta$ where Δ is the determinant of B and $B_{i,j}$ is the co-factor of the (i,j)th entry of B. Since a co-factor is, up to its sign, equal to a $(n-1) \times (n-1)$ subdeterminant of B, our bound on determinants (lemma 17) provides the claimed upper bound of s^{2s}. Now each entry of $\mathbf{c}$ has magnitude at most $2s^{4s}$. Thus each entry of $\mathbf{x}' = B^{-1}\mathbf{c}$ has magnitude at most $2ns^{6s} < 2s^{7s}$. Finally observe that some permutation of $\mathbf{x}'$ corresponds to a solution to $A\mathbf{x} \geq \mathbf{b}$. This concludes our proof of the Key lemma.

3.7 Final Remarks

This chapter not only introduces the class *NP* but opens the door to the core activities in Complexity Theory: many questions that researchers ask can be traced back to the desire to understand the relationship between *NP* (which encompasses many problems of interest) and *P* (which constitutes the feasably solvable problems in the fundamental mode). For instance, the concepts of reducibilities and complete languages will be extended and serve as subject matters for the next two chapters.

Another important motivation for *NP* comes from the theory of recursive functions: the class *P* is often compared to the recursive sets. Then the class *NP* is the analogue of the class of recursively enumerable sets.

We began this chapter with the traveling salesman problems, TSO and TSD. Now we can easily show the recognition problem TSD is *NP*-complete. It is enough to show how to transform any graph $G = (V, E)$ into an input $\langle D, b\rangle$ for the TSD problem such that G has a Hamiltonian circuit iff $\langle D, b\rangle \in$ TSD. Assume that the vertices of G are $1, 2, \ldots, n$. The $n \times n$ matrix $D = d_{i,j}$ is defined by $d_{i,j} = 1$ if $\{i, j\} \in E$; otherwise $d_{i,j} = 2$. It is easy to see that D has a tour of length $b = n$ iff G has a Hamiltonian circuit. The simplicity of this reduction illustrates an earlier remark that showing a problem L to be *NP*-hard can be facilitated if we have a closely related problem L' already known to be *NP*-hard.

Algebraic versus Bit Models of complexity In computational problems involving numbers, typically problems in computational geometry, we have two natural models of complexity. One is the Turing machine model where all numbers must ultimately be encoded as finite bit-strings; the other is where we view each number as elementary objects and we count only the number of algebraic operations. In simple situations, the only algebraic operations are the four arithmetic operations of $+, -, \times, \div$. Often $\sqrt{\cdot}$ is added. These two complexity models are usually called the *bit model* and *arithmetic model*, respectively. Complexity in the two models are often closely related. However, we point out that these models are probably independent from a complexity viewpoint. More precisely, there may be problems requiring non-polynomial time in one model but which uses only polynomial time in the other. For instance, the linear programming problem is polynomial time under the bit model, but not known to be polynomial time in the algebraic model. On the other hand, the problem of shortest paths between two points amidst polygonal obstacles is polynomial time in the algebraic model, but not known to be in polynomial time in the bit model.

Exercises

[3.1] Prove propositions B and C.

[3.2] Construct the Turing transducer N that computes the transformation in the proof of Cook's theorem.

[3.3] * Convert the satisfying assignment problem into an optimization problem by asking for an assignment which satisfies the maximum number of clauses in a given CNF formula F. Is this problem is polynomially equivalent to SAT?

[3.4] Show that 2SAT, the set of satisfiable CNF formulas with exactly two literals per clause can be recognized in polynomial time.

[3.5] Let X be a set of n Boolean variables and k be a positive integer less than n. Construct a CNF formula $T_k(X, Y)$ containing the variables in the set $X \cup Y$ where Y is an auxiliary set of variables depending on your construction satisfying the following property: an assignment $I : X \cup Y \to \{0, 1\}$ satisfies $T_k(X)$ iff I makes at least k variables in X true. Let us call $T_k(X, Y)$ the *k-threshold* formula. Your formula should have size polynomial in n. What is the smallest size you can achieve for this formula? **Hint:** it is easy to construct a formula where Y has kn variables. Can you do better?

[3.6] Show a log-space reduction of the following problems to SAT. Note that there are well-known deterministic polynomial time algorithms for these problems. However, you should give direct solutions, i.e., without going through Cook's theorem or appeal to the known deterministic polynomial time algorithms. Hopefully, your reductions to SAT here are less expensive than solving the problems directly! So your goal should be to give as 'efficient' a reduction as possible. In particular, none of your transformations should take more than quadratic time (you should do better for the sorting problem).

(i) Graph matching: input is a pair $\langle G, k\rangle$ where k is a positive integer and G is n by n Boolean matrix representing an undirected graph; the desired property is that G has a matching of cardinality k. The threshold formulas in the previous problem are useful here.

(ii) Sorting: input is a pair $\langle L, L'\rangle$ of lists where each list consists of a sequence of the form

$$\#k_1\#d_1\#k_2\#d_2\#\cdots\#k_n\#d_n\#$$

where k_i and d_i are binary strings. Regarding k_i as the 'key' and d_i as the corresponding 'data', this language corresponds to the problem of sorting the data items according to their key values. The 'input' list L is arbitrary but the 'output' list L' is the sorted version of L: L' ought to contain precisely the same (key, data) pairs as L, and the keys in L' are in non-decreasing order.

(iii) Non-planarity testing: input is a graph G and the property is that G is non-planar. (Use the two forbidden subgraphs characterization of Kuratowski.)

(iv) Network flow: input is $\langle G, s, t, C, k\rangle$ where G is a directed graph, s and t are distinct nodes of G (called the source and sink), C assigns non-negative integer values to the nodes of G ($C(u)$ is the *capacity* of node u), and $k \geq 0$ is an integer. The desired property is that the maximum value of a flow from s to t is k. Note that integers are represented in binary, and C can be encoded by a list of pairs of the form $(u, C(u))$ for each node u. (Use the max-flow-min-cut theorem and the fact that we can restrict flows to be integral.)

[3.5] Recall that a 3DNF formula has the form

$$\bigvee_{i=1}^{m} \bigwedge_{j=1}^{3} u_{i,j}$$

where $u_{i,j}$ are literals. We also let 3DNF-SAT denote the set of satisfiable (encoded) 3DNF formulas. Show that 3DNF-SAT can be recognized in polynomial time.

[3.6] (Bauer-Brand-Fisher-Meyer-Paterson) Let F be a Boolean formula. Show a systematic transformation of F to another 3CNF G such that they are co-satisfiable (i.e., both are satisfiable or both are not). Furthermore, $|G| = O(|F|)$ where $|G|$ denotes the size of the formula G. **Hint:** For each subformula H of F, introduce an additional variable α_H. We want to ensure that any assignment I to the variables of H assigns to α_H the value of the subformula H under I (i.e., if I makes H false then $I(\alpha_H) = 0$). For instance, $H = H_1 \vee H_2$. Then we introduce clauses that enforce the equivalence $\alpha_H \equiv \alpha_{H_1} \vee \alpha_{H_2}$.

[3.7] The *parity* function on n variables is $x_1 \oplus x_2 \oplus \cdots \oplus x_n$ where $\oplus$ is exclusive-or. Thus the function is 1 precisely when an odd number of its inputs are 1. Show that the smallest CNF formula equivalent to the parity function is exponential in n. (Note that this shows that we could not strengthen the previous exercise such that F and G become equivalent rather than just co-satisfiable.) **Hint:** Consider the following CNF formula $\bigwedge_{i=1}^{k} \bigvee J_i$ where each J_i is a set of literals over $x_1, \ldots, x_n$. Show that if this is the parity function then $|J_i| = n$.

[3.8] Show that the problem of graph coloring is *NP*-complete.

[3.9] Show that the problem of deciding if a graph contains a pre-specified set of vertex-disjoint paths is *NP*-complete. The input is an undirected graph together with some set $\{(u_i, v_i) \ : \ i = 1, \ldots, k\}$ of pairs of vertices. The required property is that there are k pairwise vertex-disjoint paths connecting the given pairs of vertices. **Hint:** Reduce from 3SAT.

[3.10] (Burr) A graph $G = (V, E)$ is *NMT-colorable* ('no monochromatic triangle') if it can be 2-colored such that no triangle (3-cycle) has the same color. Show that the problem of recognizing NMT-colorable graphs is *NP*-complete. *Hint*: Construct three gadgets (I), (II) and (III) with the following properties. Gadget (I) has two distinguished nodes such that any NMT-coloring of (I) must give them the same color. Gadget (II) has two distinguished nodes such that any NMT-coloring of (II) must give them distinct colors. Gadget (III) has four distinguished nodes A, B, C, D such that in any NMT-coloring of this gadget must make at least one of A, B or C the same color as D. When the various gadgets are strung together, the D node of all copies of Gadget (III) should be common.

[3.11] Prove that the 2-DISJOINT-PATH problem is *NP*-complete: given a digraph with source node s and sink node t, and a bound $B > 0$, are there two node-disjoint paths of length at most B from s to t? HINT: reduce from 3SAT or from PARTITION.

[3.12] Consider the regular expressions over some alphabet Σ involving the operators of concatenation $(\cdot)$, union $(+)$, and Kleene-star $(*)$. Each regular expression α denotes a subset of Σ^*. We now consider a class of modified regular expressions in which Kleener-star is not used but where we allow intersection $(\cap)$. Show that the problem of recognizing those modified regular expressions α where $L(\alpha) \neq \Sigma^*$ is *NP*-hard. *Hint*: imitate Cook's theorem, but use such expressions to denote those strings that *do not* represent accepting computations.

[3.13] (Cook) Show that the problem of recognizing input pairs $\langle G, G' \rangle$ of undirected graphs with the property that G is isomorphic to a subgraph of G' is *NP*-complete.

[3.14] * (Stockmeyer) Show that the problem of recognizing those graphs G that are planar and 3-colorable is *NP*-complete. (Note: we can assume that the input graph is planar because planarity testing can be done in linear time.) **Hint:** Reduce 3SAT to the present problem.

[3.15] * (Kozen) Show that the problem of recognizing valid sentences of the first-order predicate calculus with the equality symbol but *without negation* is *NP*-complete. The language contains the usual logical symbols $(\wedge, \vee, \forall, \exists)$ sans negation $(\neg)$, individual variables x_i, relation symbols R_i^m and function symbols F_i^m (where $m \geq 0$ denotes the arity of the relation or function symbol, and for $i = 0, 1, \ldots$). We further assume that R_0^m is the standard equality symbol '='. We are interested in valid sentences, i.e., closed formulas that are true in all models under all interpretations (but equality is standard). Called this the *validity problem for positive first-order logic.*

[3.16] (a) Show that the *tiling problem* is *NP*-complete. The input to this problem has the form $\langle n, k, S \rangle$ where n and k are positive integers and S is a finite set of 'tile patterns'. Each tile pattern in S is a sequence of the form $p = \langle c_1, c_2, c_3, c_4 \rangle$ where $c_i \in \{1, \ldots, k\}$. Any oriented unit square whose top, bottom, left and right edges are colored with the colors c_1, c_2, c_3 and c_4, respectively, is said to have pattern p. The problem is to decide whether it is possible to cover an n by n square area using n^2 unit tiles satisfying the following condition: (i) each unit tile has a pattern from S, and (ii) for any two tiles that abut, their two adjacent edges have the same color. **Hint:** simulate a Turing machine computation using tiles.

(b) For fixed k, let the k-tiling problem be the restriction of the tiling problem to k colors. What is the smallest k for which you can show that the k-tiling problem remains *NP*-complete? What is the largest k for which you can show that the problem is in P ?

[3.17] (Karp-Held) Give a dynamic programming solution to TSO which has a running time of $O(n2^n)$. Hint: Let $S \subseteq \{2, 3, \ldots, n\}$ be a subset of the n cities, and $k \in S$. Let $C(S, k)$ denote the minimum cost of a route which starts at city 1, visiting all the cities in S exactly once, terminating at k. Give an expression for $C(S, k)$ in terms of $C(T, i)$ for all T where $T \subseteq S$ and $|T| = |S| - 1$.

[3.18] (a) Show that if p is prime and x is an integer, $1 \leq x < p$, then p divides $\left(\begin{array}{c} p \\ x \end{array} \right)$ where $\left(\begin{array}{c} p \\ x \end{array} \right) = \frac{p(p-1)\cdots(p-x+1)}{x!}$.
(b) Conclude (using induction on k) that if p is prime then for all $x_1, \ldots, x_k$,

$$\left(\sum_{i=1}^{k} x_i \right)^p \equiv \left(\sum_{i=1}^{k} x_i^p \right) \quad (\mod p).$$

[3.19] Prove that if G is an Abelian group, and s and t are elements of G of orders m and n (respectively) then G has an element of order $\texttt{LCM}(m, n)$ $(:= mn/\texttt{GCD}(m, n))$. Hint: by replacing s with $s^{\texttt{GCD}(m,n)}$, we may assume that $\texttt{LCM}(m, n) = mn$. What is the order of st?

[3.20] In the text, we show that $\mathbb{Z}_n^*$ is a cyclic group when n is a prime.
(a) Extend this result to show that $\mathbb{Z}_n^*$ is cyclic if $n = 2, 4$ or $n = 2p^e$ $(e \geq 2)$ where p is an odd prime.
(b) Show that if n is not of the above form, then $\mathbb{Z}_n^*$ is not cyclic.

[3.21] * Suppose $a \in \mathbb{Z}_n^*$ is a primitive root of n. Is there a polynomial certificate for this fact?

[3.22] Let $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$. We call R a "certifier of primes" if it has the following property (1) $(x, y) \in R$ implies x is the binary notation for a prime number, and (2) for all prime number x in binary notation, there exists y such that $(x, y) \in R$. If $(x, y) \in R$, we call y a "certificate for x". Prove that there is a certifier of primes R which can decided in polynomial time (i.e., there is a polynomial time Turing machine with 2 input tapes which, given x, y placed on these input tapes, will halt in polynomial time and accept iff $(x, y) \in R$). Moreover, $(x, y) \in R$ implies $|y| \leq O(|x|^2)$. This last requirement says that all primes have quadratic size certificates.

[3.23] We want to demonstrate that if we allow the output tape of transducers to be a 2-way (but still read-only) tape then we have a more powerful model. Recall the non-regular language L_0 in the last section of chapter 2. Consider the problem where, given an integer n in binary, we want to output the string $\bar{1}\#\bar{2}\#\cdots\#\bar{n}$ in the language L_0. The input size is $O(\log n)$ and the output size is $O(n \log n)$. (a) Show that if the output tape is 1-way as in our standard definition of transducers, then the space used is linear, i.e. $O(\log n)$. (b) Show that if we have a 2-way read-only output tape then logarithmic space (i.e. $O(\log \log n)$) suffices.

[3.24] * (Berman) Show that if there is a tally language $L \subseteq \{1\}^*$ that is complete for NP (under $\leq_m^P$) then $P = NP$.

[3.25] * The *Euclidean Travelling Salesman's Problem* (ETS) is the following: given a set $\{p_1, \ldots, p_n\}$ of points in the plane, find the shortest tour that connects all these points. It is assumed that the points have integer coordinates and the distance between any pair of points p, q is the usual Euclidean distance:

$$d(p, q) = \sqrt{(p_x - q_x)^2 + (p_y - q_y)^2}$$

where $p = (p_x, p_y), q = (q_x, q_y)$. ETS is known to be NP-hard, but why is it not obviously NP-complete? Give a single-exponential time upper bound on the complexity of ETS.
Hint: Reduce the problem to comparing a sum of square-roots to zero.

[3.26] (Bloemer and Yap) Show that the following problem can be solved in deterministic polynomial time: given integers a_i and positive numbers n_i $(i = 1, \ldots, m)$ determine if $S = \sum_{i=1}^{m} a_i \sqrt{n_i}$ is equal to 0. NOTE: In contrast, testing if $S > 0$ is not known to be polynomial time.

Appendix A

Propositional Logic

A *Boolean variable* is a variable that assumes a value of "1" or "0" (alternatively "true" or "false"). If x is a Boolean variable, then its *negation* is denoted either $\bar{x}$ or $\neg x$. A *literal* is either a Boolean variable or the negation of one. A *(Boolean) formula* is either a variable or recursively has one of the three forms:

conjunction: $(\phi \wedge \psi)$

disjunction: $(\phi \vee \psi)$

negation: $(\neg\phi)$

where ϕ and ψ are Boolean formulas. As an example of a Boolean formula, we have

$$((x_1 \vee x_2) \vee x_3) \wedge (((\bar{x}_1 \vee x_2) \vee x_2) \wedge (\bar{x}_2 \vee \bar{x}_3)). \tag{1}$$

It is important to realize that formulas are syntactic objects, i.e., a sequence of marks laid out according to the preceding rules, where the marks comes from $\vee, \wedge, \neg, (,)$, and a unique mark for each variable x_i, $i = 1, 2, 3, \ldots$. When encoding formulas to apply our theory, we cannot assume an infinite number of marks. Hence we must encode each x_i by a sequence from a (finite) alphabet Σ. A simple choice, and the one assumed in this book, is the following: let

$$\Sigma = \{\vee, \wedge, \neg, (,), x, 0, 1\}$$

and each x_i is encoded by the symbol "x" followed by $b_0 b_1 \cdots b_m \in \{0,1\}^*$ where $b_0 \cdots b_m$ is the integer i in binary. Using this convention, and assuming that no variable x_{i+1} is used unless x_i is also used, it is easy to see that the *size* of a formula ϕ is $O(k + m \log m)$ where k is the number of Boolean operator occurrences and m is the number of variable occurrences.

A variable x *occurs* in a formula F if either x or its negation syntactically appears in F. If all the variables occurring in F are among $x_1, x_2, \ldots, x_n$, we indicate this by writing $F(x_1, x_2, \ldots, x_n)$. In this notation, some (possibly all) x_i may not occur in F. An *assignment* to a set X of Boolean variables is a function $I : X \to \{0,1\}$. If X contains all the variables occurring in F, and I is an assignment to X then I (by induction on the size of F) assigns a Boolean value $I(F)$ to F as follows: if F is a variable x, $I(F) = I(x)$; otherwise:

if F is the negation $(\neg\phi)$ then $I(F) = 1 - I(\phi)$;
if F is the conjunct $(\phi \wedge \psi)$ then $I(F) = \min\{I(\phi), I(\psi)\}$;
if F is the disjunct $(\phi \vee \psi)$ then $I(F) = \max\{I(\phi), I(\psi)\}$.

We say I *satisfies* F if $I(F) = 1$. A formula F is *satisfiable* if there exists an assignment which satisfies it. Thus the formula in equation (A.1) is satisfied by I with $I(x_1) = I(x_2) = 1, I(x_3) = 0$.

Two formulas ϕ and ψ are *equivalent*, written $\phi \equiv \psi$, if for any assignment I to the variables occurring in both formulas, $I(\phi) = I(\psi)$. The formulas are *co-satisfiable* if they are either both satisfiable or both unsatisfiable.

As seen in the example (A.1), parentheses in formulas get to be tedious. We can avoid these parenthesis by using the properties that conjunctions and disjunctions are associative, i.e.,

$$\phi_1 \wedge (\phi_2 \wedge \phi_3) \equiv (\phi_1 \wedge \phi_2) \wedge \phi_3; \quad \phi_1 \vee (\phi_2 \vee \phi_3) \equiv (\phi_1 \vee \phi_2) \vee \phi_3,$$

and commutative, i.e.,

$$\phi_1 \wedge \phi_2 \equiv \phi_2 \wedge \phi_1; \quad \phi_1 \vee \phi_2 \equiv \phi_2 \vee \phi_1,$$

and using the fact that

$$\phi \vee \phi \equiv \phi; \quad \phi \wedge \phi \equiv \phi.$$

Thus (A.1) can be simply written as

$$(x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_1 \vee x_2 \vee x_2) \wedge (\bar{x}_2 \vee \bar{x}_3). \tag{2}$$

If $S = \{u_1, \ldots, u_n\}$ is a set of literals, we also write

$$\bigvee S \text{ or } \bigvee_{u \in S} u \text{ or } \bigvee_{i=1}^{n} u_i$$

for the conjunction of the literals in S. We use $\bigwedge$ in a similar way for conjunctions. A formula that is a conjunction of disjunctions of literals is said to be in *conjunctive normal form* (CNF). The example (A.2) is such a formula. The *disjunctive normal form* (DNF) is similarly defined.

From the properties noted above, formulas in CNF can be given a convenient alternative form: A CNF formula *in clause form* is a set of clauses, where a *clause* is a set of literals. There is an obvious way to transform a CNF formula in the usual form to one in clause form; for example, the formula (A.2) in the clause form is

$$\{\{x_1, x_2, x_3\}, \{\bar{x}_1, x_2\}, \{\bar{x}_2, \bar{x}_3\}\}.$$

Note that since clauses are defined as sets, repeated literals are removed, as in the case of the second clause above. Satisfiability of CNF formulas in the clause form is particularly simple: let X be the set of variables occurring in the CNF formula F and let I be an assignment to X. Then I *satisfies* a clause C in F iff for some literal u in C we have $I(u) = 1$. I *satisfies* F iff I satisfies each clause in F. For instance, the following CNF formula

$$\{\{x_1, \bar{x}_2\}, \{\bar{x}_1, x_2\}, \{x_1\}, \{\bar{x}_2\}\}$$

is unsatisfiable.

Bibliography

[1] M. Akgül. *Topics in relaxation and ellipsoidal methods.* Pitman Advanced Publishing Program, Boston-London-Melbourne, 1981. (U. Waterloo PhD Thesis).

[2] K. H. Borgwardt. *The Simplex Method: a probabilistic analysis.* Springer-Verlag, 1987.

[3] I. Borosh and L. B. Treybig. Bounds on positive integral solutions of linear Diophantine equations. *Proc. AMS*, 55:299–304, 1976.

[4] S. A. Cook. The complexity of theorem-proving procedures. *3rd Proc. ACM Symp. Theory of Comp. Sci.*, pages 151–158, 1971.

[5] M. R. Garey and D. S. Johnson. *Computers and Intractibility: A Guide to the Theory of NP-completeness.* Freeman, New York, 1979.

[6] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers.* Oxford University Press, London, 1938.

[7] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–104. Plenum Press, New York, 1972.

[8] L. G. Khachian. A polynomial algorithm for linear programming. *Doklady Akad. Nauk USSR*, 244:5:1093–96, 1979. (tr. *Soviet Math. Doklady 20* 191-194).

[9] V. Klee and G. J. Minty. How good is the simplex algorithm? In O. Shisha, editor, *Inequalities III*, pages 159–175. Academic Press, 1972.

[10] T. S. Kuhn. *The structure of scientific revolutions.* Chicago Univ. Press, 1970.

[11] L. A. Levin. Universal sorting problems. *Problemi Peredachi Informatsii*, 9:3:265–266, 1973.

[12] C. H. Papadimitriou. On the complexity of integer programming. *Journal of the ACM*, 28:765–768, 1981.

[13] V. R. Pratt. Every prime has a succinct certificate. *SIAM J. Comput.*, 4:214–220, 1975.

[14] J. von zur Gathen and M. Sieveking. A bound on solutions of linear integer equalities and inequalities. *Proc. AMS*, 72:155–158, 1978.

Chapter 4

Reducibilities

April 13, 2009

4.1 Inclusion Questions

Many open problems of Complexity Theory are *inclusion questions* of the form:

$$\text{Is } K \text{ included in } K' \text{ ?} \tag{1}$$

Here K, K' are two classes of languages. The oldest such inclusion (or containment) question, dating back to 1960s, is the linear bounded automata (LBA) question: Is $NSPACE(n)$ included in $DSPACE(n)$? As explained in §2.9, these two classes are also known as LBA and $DLBA$, respectively. The P versus NP problem is another instance of the inclusion question. In chapter 3, we introduced Karp reducibility and showed its value for studying this question. In general, reducibilities will be the key tool for investigating inclusion questions. We begin now with a very general notion of reducibility.

Definition 1. A *reducibility* is a reflexive binary relation on languages. Let $\leq$ denote a reducibility. We use an infix notation to reducibility relations: if (L, L') belongs to the relation $\leq$, we write $L \leq L'$. A reducibility is *transitive* if it is transitive as a binary relation; otherwise it is *intransitive*. ■

In the literature, intransitive reducibilities are sometimes called "semi-reducibilities" and "reducibilities" are reserved for transitive ones. Admittedly, our use of "$\leq$" to denote semi-reducibilities can be confusing, since the inequality symbol suggests transitivity. Intransitive reducibilities typically arise when the reducibility concept is defined using nondeterministic machines (see §9).

We say L is $\leq$-*reducible* to L' if the relation $L \leq L'$ holds. Two languages L and L' are $\leq$-*comparable* if either $L \leq L'$ or $L' \leq L$; otherwise they are $\leq$-*incomparable*. We also write $L \not\leq L'$ if $L \leq L'$ does not hold; and write $L < L'$ if $L \leq L'$ and $L' \not\leq L$.

Next assume that $\leq$ is transitive. If $L \leq L'$ and $L' \leq L$ then we say L and L' are $\leq$-*equivalent*. The $\leq$-*degree* of L is the class of languages $\leq$-equivalent to L.

The following generalizes some concepts from chapter 3.

Definition 2. Let K be a class of languages and $\leq$ a reducibility. Then L is K-*hard* (*under* $\leq$) if for every L' in K, L' is $\leq$-reducible to L. L is K-*complete* (*under* $\leq$) if L is K-hard and L is in K. K is *closed under* $\leq$-reducibility if for any L and L', we have that $L \leq L'$ and $L' \in K$ implies $L \in K$. ■

In Chapter 3 we proved that the language SAT is NP-complete. This definition reminds us that, in general, we need to mention the particular reducibility concept used in such a result. The importance of complete languages as a tool for studying the question (1) comes from the following easy generalization of the corollary to lemma 2 in chapter 3.

Lemma 1 (The Basic Inclusion Lemma). *Let K and K' be language classes and $\leq$ a reducibility. Assume*
(a) K' is closed under $\leq$-reducibility, and
(b) L_0 is a K-complete language (under $\leq$).
Then

$$K \subseteq K' \iff L_0 \in K'.$$

Under hypotheses (a) and (b) of this lemma, the inclusion question (1) is equivalent to the membership of a single language L_0 in K'. Thus a K-complete language serves as a representative of the entire class K. This lemma is the basis of most applications of reducibilities to the inclusion questions.

Choice of reducibilities (I). Given that we want to study a particular inclusion question, whether $K \subseteq K'$, within the framework of this lemma, how do we choose the reducibility $\leq$? One desirable property is that $\leq$ be transitive. Then, we can show a language $L \in K$ to be K-complete under $\leq$ simply by showing a known K-complete language L_0 is $\leq$-reducible to L. This tack is well-exploited in the field of NP-completeness theory as described in Chapter 3. We will make several more remarks on the choice of reducibilities in the following section. ■

Closure under Complementation Questions. We say a class K is *closed under complementation* if $K = \text{co-}K$. The question of closure under complement is equivalent to a special kind of inclusion question since

$$K = \text{co-}K \iff K \subseteq \text{co-}K.$$

This is because $K \subseteq \text{co-}K$ iff $\text{co-}K \subseteq \text{co-}(\text{co-}K) = K$. Of the many open inclusion questions, there is often a weaker conjecture involving closure under complementation. For instance, it is not known if NP is closed under complementation. If $P = NP$ then NP must be closed under complementation. The converse is not known to be true. Hence it is conceivably easier to prove that $NP = \text{co-}NP$ than to prove that $NP = P$. Another example is the LBA-question. While the LBA-question remains open, the general result about non-deterministic space classes in §2.9 implies that LBA is closed under complementation.

We often wish to compare two reducibilities $\leq_1$ and $\leq_2$. We say $\leq_1$ is *as strong as* $\leq_2$ if $L \leq_2 L'$ implies $L \leq_1 L'$ for all L, L'. It is important to note the direction of implication in this definition since the literature sometimes gives it the opposite sense.[1] If $\leq_1$ is as strong as than $\leq_2$, and but not vice-versa, then $\leq_1$ is *stronger than* $\leq_2$. Thus, the stronger reducibility strictly contains the weaker one, if we regard a reducibility as a set of ordered pairs of languages.

Historical Notes. Reducibilities are extensively used in recursive function theory. At the subrecursive level, Cook in the proof of his theorem, is an early example of defining a reducibility that takes complexity into account. Meyer and McCreight [13] used such reducibilities in relating complexity of problems. Following Meyer, the reducibilities that take complexity into account may be generically called 'efficient reducibilities'. Systematic study of efficient reducibilities was first undertaken by Ladner, Lynch, and Selman [8, 9, 10]. See also [3, 18].

Karp reducibility is denoted $\leq_m^P$ (§3.3). The superscript P here indicates polynomial-time; similarly, the subscript m indicates Karp reducibility is a form of "many-one reducibility". Generally speaking, efficient reducibilities are classified along two lines, and this is reflected in our notational scheme

$$\leq_\tau^\chi$$

for reducibilities: the symbol τ indicates the type of the reducibility, and the symbol χ indicates the complexity characteristics (see §4.4).

Convention. The following abuse of notation is often convenient. If M (resp. T) is an acceptor (resp. transducer) we sometimes use it to also denote the language (resp. transformation) defined by the indicated machine. Thus if x, y are words, then we may say '$x \in$ M' or '$y = \text{T}(x)$' with the obvious meaning.

4.2 Many-One Reducibility

Many-one reducibility (§3.3) is now treated in a more general setting. The terminology begs the question: is there a one-one or a many-many reducibility? In recursive function theory, and occasionally in complexity theory (e.g. [19]), one-one reducibility has been studied. Nondeterministic many-one reducibility (to be introduced in section 9) may be regarded as a many-many reducibility.

Definition 3. Let Σ, Γ be alphabets, and Φ be a family of transformations. We say a language (Σ, L) is *many-one reducible to* another language (Γ, L') *via* Φ if there exists a transformation $t : \Sigma^* \to \Gamma^*$ in Φ such that for all $x \in \Sigma^*$,

$$x \in L \iff t(x) \in L'.$$

We write $L \leq_m^\Phi L'$ in this case. ■

[1] Our rationale is that the strength of the reducibility should be in direct proportion to the power of the machines used in defining it: thus the polynomial time reducibility ought to be as strong as one defined by logarithmic space bounds. The opposite sense is reasonable if we think of "logical strength of implication". Kenneth Regan suggests the term "finer" which seems avoid all ambiguity.

For any alphabet Σ, the identity transformation is $t : \Sigma^* \to \Sigma^*$ with $t(x) = x$ for all x. The following proposition is immediate.

LEMMA 2.

(i) *If Φ contains all identity transformations then $\leq_m^\Phi$ is a reducibility.*

(ii) *If Φ is, in addition, closed under functional composition then $\leq_m^\Phi$ is a transitive reducibility.*

(iii) *If Φ, Ψ are families of transformations where $\Phi \subseteq \Psi$ then $\leq_m^\Psi$ is as strong as $\leq_m^\Phi$.*

If Φ is understood or immaterial, we shall simply write '$\leq_m$' instead of '$\leq_m^\Phi$'. Recall that a language L is *trivial* if L or co-L is the empty set. It is easy to see that no non-trivial language is many-one reducible to a trivial language and conversely, a trivial language is many-one reducible to any non-trivial one via some constant transformation (i.e., $t(x) = t(y)$ for all x, y). To avoid these special cases, we tacitly restrict all discussion to only non-trivial languages when discussing many-one reducibilities. The following are some important cases of many-one reducibility.

1. If Φ is the class of all transformations computed by halting deterministic Turing transducers, we will write $\leq_m^{REC}$ for $\leq_m^\Phi$. We will say that L is **recursively many-one reducible** to L' if $L \leq_m^{REC} L'$. Recall the halting problem, HALT from Chapter 0. See the Exercise for a proof that HALT is RE-complete under $\leq_m^{REC}$.

2. In Chapter 3, we described Karp reducibility that uses the family $\Phi = \mathbf{P}$ of transformations computed by deterministic polynomial-time transducers.[2]

3. We introduce another important family, denoted **DLOG**, consisting of all transformations computed by deterministic transducers running in logarithmic space. The corresponding reducibility is denoted $\leq_m^L$ and called *log-space many-one reducibility*. The proof that $DSPACE(s) \subseteq DTIME(n \cdot O(1)^{s(n)})$ in chapter 2 (§6) can easily be adapted to show that

$$\mathbf{DLOG} \subseteq \mathbf{P} .$$

This implies that $\leq_m^P$ is as strong as $\leq_m^L$.

Choice of reducibilities (II). An important instance of (1) is whether

$$NLOG \subseteq DLOG. \tag{2}$$

Assuming that $P \neq DLOG$ (as is generally conjectured), the class $DLOG$ is not closed under $\leq_m^P$ (see Exercise). Thus hypothesis (a) of the Basic Inclusion Lemma fails and we cannot use $\leq_m^P$-reducibility to investigate question (2). We say Karp reducibility is too strong for distinguishing $DLOG$ from $NLOG$. More generally, hypothesis (a) places an upper limit on the power of the reducibility used in studying inclusion questions. On the other hand, it is easy to see that $DLOG$ is closed under $\leq_m^L$, and chapter 5 will show that $NLOG$ has complete languages under $\leq_m^L$. This is one of the motivations for introducing $\leq_m^L$. ■

We now show that $\leq_m^L$ is a transitive reducibility. This follows immediately from the following theorem of Jones [7] :

THEOREM 3. *The family* **DLOG** *of log-space transformations is closed under functional composition.*

Proof. Let $\Sigma_0, \Sigma_1, \Sigma_2$ be alphabets. For $i = 1, 2$, let $t_i : \Sigma_{i-1}^* \to \Sigma_i^*$ be computed by the transducer $\mathbb{N}_i$ that runs in space $\log n$. We shall construct a 2-tape transducer M that computes $t_0 : \Sigma_0^* \to \Sigma_2^*$ that is the functional composition of t_1 and t_2. To motivate the present proof, briefly recall the proof that $\leq_m^P$ is a transitive reducibility. A direct adaptation of that proof to the present problem would yield a machine M′ which on input x, simulates $\mathbb{N}_1$ to produce $t_1(x)$ and then simulates $\mathbb{N}_2$ on $t_1(x)$. In general such an M′ uses more than logarithmic space since $|t_1(x)|$ may have length that is polynomial in $|x|$. To ensure that only logarithmic space is used, we shall avoid storing $t_1(x)$, but rather recompute the symbols in $t_1(x)$ as often as needed.

By our convention on transducers, tape 1 of M is the output tape. Tape 2 has four tracks used as follows:

[2] We use bold letters to denote families of transformations.

Track 1:	Work-space of $\mathbb{N}_1$ on input x.
Track 2:	Work-space of $\mathbb{N}_2$ on input $t_1(x)$.
Track 3:	An integer indicating the position of the input head of $\mathbb{N}_2$ on the input $t_1(x)$.
Track 4:	Counter used when simulating $\mathbb{N}_1$.

Initially, track 3 contains the integer 1 indicating that the input head of $\mathbb{N}_2$ is scanning the first symbol of $t_1(x)$. In general, if track 3 contains an integer i, M will call a subroutine R that will determine the ith symbol of $t_1(x)$. This is done as follows: R initializes the counter on track 4 to 0 and begins to simulate $\mathbb{N}_1$ using track 1 as its work-tape. Each time $\mathbb{N}_1$ produces a new output symbol, R increments the count on track 4 and checks if the new count equals the integer i on track 3; if so, R can immediately return with the ith symbol of $t_1(x)$. The operations of M are essentially driven by a direct simulation of $\mathbb{N}_2$: to simulate a step of $\mathbb{N}_2$, M first calls subroutine R to determine the symbol in $t_1(x)$ currently scanned by the input head of $\mathbb{N}_2$. It can then update the work-tape of $\mathbb{N}_2$ represented on track 2, and increment or decrement the integer on track 3 according to the motion of the input head of $\mathbb{N}_2$. Clearly M computes $t_2(t_1(x))$. Finally, we must verify that M uses $O(\log n)$ space. Track 1 uses $\log |x|$ space. It is not hard to see that track 2, 3 and 4 uses $\log(|t_1(x)|) = O(\log |x|)$ space. **Q.E.D.**

Choice of reducibilitites (III). For most of the open inclusion questions, the consensus opinion generally favor a negative answer. Using smaller families of transformations (such as **DLOG** instead of **P**) to define reducibilities gives us sharper tools for proving negative answers to these questions. For instance, suppose L is NP-complete under some $\leq$-reducibility and P is closed under $\leq$ (thus the premises of the Basic Inclusion Lemma are satisfied). To show that NP is not contained in P , it is sufficient to prove that L is not $\leq$-reducible to some P -complete language L' (in the next chapter we shall show the P does have complete languages). Clearly, it would be easier to prove this result if $\leq$ were in fact $\leq_m^L$ rather than $\leq_m^P$. However, there are inherent limits to this tact of tool sharpening, as illustrated by the next result from [6]. ■

LEMMA 4. *Let Φ be the families of transformations computed by transducers running in space $f(n)$. If $f(n)$=$o(\log n)$ then there does not exist NP-complete languages under $\leq_m^\Phi$-reducibility.*

Proof. If t is a transformation computed by some transducer using $f(n) = o(\log n)$ space then for all x, $|t(x)| = |x|2^{O(f(|x|))} = o(|x|^2)$. Let $L_0 \in NTIME(n^k)$ for some $k \geq 1$. If $L' \leq_m^\Phi L_0$ via t then it is easy to see that $L' \in NTIME(n^{2k})$. In particular, if L_0 is NP-complete under $\leq_m^\Phi$-reducibility then $NP = NTIME(n^{2k})$. This contradicts the fact that $NP \neq NTIME(n^k)$ for any k (this result is shown in chapter 6). **Q.E.D.**

Choice of reducibilities, conclusion (IV). This lemma says if the reducibility is weakened below logarithmic space then it is useless for distinguishing P from NP. Combined with preceding discussions, we draw the lesson that the strength of the reducibility should chosen appropriately for the inclusion question of interest. It turns out that because the major questions we ask center around the canonical list, many-one log-space reducibility $\leq_m^L$ is suitable for most of our needs. ■

Despite lemma 4, there are subfamilies of the logarithmic-space transformations **DLOG** that are useful for a variety of purposes. We briefly describe three such subfamilies of **DLOG**.

(a) Consider what is essentially the smallest non-trivial family of transformations, namely, those computed by transducers that use no space (these are called *finite state transducers*). Denote this family by **FST**. This family is closed under functional composition [1]. If the input head of the transducer is constrained to be one-way, we denote the corresponding subfamily by **1FST**.[3] We will show in chapter 6 that there are complete languages for $DTIME(n^k)$ under such $\leq_m^{1FST}$-reducibilities.

(b) Let **Llin** denote the subfamily of **DLOG** where $t \in$ **Llin** implies that for all x, $|t(x)| = O(|x|)$. Meyer, Lind and Stockmeyer first considered these 'log-linear' transformations and corresponding $\leq_m^{Llin}$-reducibilities. In chapter 6, we will see applications of such transformations in lower bound proofs.

(c) Let **1DLOG** denote the family of transformations computable by transducers that use logarithmic space but where the input tapes are one-way. The corresponding reducibility $\leq_m^{1L}$ has been studied by Hartmanis and his collaborators. In particular, they point out that many of the well-known complete languages for $DLOG, NLOG, P$ and NP remain complete under $\leq_m^{1L}$ reducibilities. It is easy to see that such reducibilities are transitive.

[3]The 1-way finite state transducers are also called *generalized sequential machines* or gsm.

We next prove some useful properties about the canonical list of complexity classes of chapter 2 (§3):

DLOG, *NLOG*, *PLOG*, *P* , *NP*, *PSPACE*,
DEXPT, *NEXPT*, *DEXPTIME*, *NEXPTIME*,
EXPS, *EXPSPACE*.

THEOREM 5.

(i) *All the classes in the canonical list are closed under* $\leq_m^{Llin}$*-reducibility.*

(ii) *Each class in the canonical list, except for DEXPT and NEXPT, is closed under* $\leq_m^{L}$*-reducibility.*

Proof. (i) We prove this for the case *DEXPT*. Say $L \leq_m^{Llin} L'$ via some $t \in$ **Llin** and L' is accepted by some M in deterministic time $2^{O(n)}$. To accept L, we operate as follows: on input x compute $t(x)$ and then simulate M on $t(x)$. The time taken is clearly $O(1)^{|t(x)|} = O(1)^{|x|}$.
(ii) Proofs similar to the above can be used except that, with log-space reducibilities, the transformed output $t(x)$ could have length polynomial in $|x|$. If M accepts in time $O(1)^n$ then to simulate M on $t(x)$ takes time $O(1)^{|t(x)|}$, which is $\neq O(1)^{|x|}$. This accounts for the exclusion of the classes *DEXPT* and *NEXPT*. **Q.E.D.**

4.3 Turing Reducibility

We introduce the idea of computing relative to oracles. Such a computation is done by a multitape Turing acceptor M equipped with a special one-way **oracle tape** and with three distinguished states called QUERY, YES and NO. Such an M is called an **oracle** or **query machine** or **oracle TM**. We allow M to be nondeterministic. Let A be any language. A computation of M *relative to the oracle set* A proceeds in the usual way until the machine enters the QUERY state. Then, depending on whether the word written on the oracle tape is in A or not, M next enters the YES or NO state, respectively. This transition is called an oracle query. The oracle tape is erased (in an instant) after each oracle query, and we proceed with the computation. We require that there are no transitions into the YES and NO states except from the QUERY state. The machine M with oracle set A is denoted by $\mathrm{M}^{(A)}$. Since other oracle sets could have been used with M to produce different behavior, we denote the query machine M without reference to any specific oracle by $\mathrm{M}^{(\cdot)}$. The acceptance time, space and reversal of $\mathrm{M}^{(A)}$ are defined as for ordinary Turing machines with the provision that *space used on the oracle tape is not counted*. We do not feel justified in charging for the space on the oracle tape since the oracle tape does not allow space to be reused or even to be reexamined (so this 'space' has properties more akin to time complexity). Clearly the time, space and reversal of $\mathrm{M}^{(\cdot)}$ depend on the oracle. We say that the acceptance time of $\mathrm{M}^{(\cdot)}$ is f if for all oracles A, $\mathrm{M}^{(A)}$ accepts in time f. Similar definitions for acceptance space/reversals as well as for running complexity can be made: we leave this to the reader. Query machines define language operators (see appendix of chapter 2): the *oracle operator* $\phi = \phi_{\mathrm{M}}$ corresponding to a query machine $\mathrm{M}^{(\cdot)}$ is defined by $\phi(A)=L'$ iff $\mathrm{M}^{(A)}$ accepts L'. The alphabet of L' is taken to be the input alphabet of $\mathrm{M}^{(\cdot)}$.

Example 1. Recall the traveling salesman decision problem TSD in chapter 3. Consider the following variation, called here the *traveling salesman minimum tour problem* (abbr. TSM):

Given:	a pair $\langle D, \pi\rangle$ where D is a $n \times n$ matrix of non-negative integers and π is a permutation on $\{1, \ldots, n\}$.
Property:	π represents a minimum cost tour.

It is easy to construct a query machine $\mathrm{M}^{(\cdot)}$ that solves TSM relative to an oracle for TSD: on input $\langle D, \pi\rangle$, M begins by checking the input has the correct format and then computing the cost c of the tour π. Then it asks the oracle two questions (by writing these words on the tape): $\langle D, c\rangle$ and $\langle D, c-1\rangle$. Our machine M will accept if and only if the first answer is yes, and the second answer is no. Clearly this oracle machine runs in polynomial time and reduces the problem TSM to the problem TSD. ■

Extending our convention for acceptors and transducers, it is sometimes convenient to say '$x \in \mathrm{M}^{(L)}$' to mean that x is accepted by the query machine $\mathrm{M}^{(L)}$.

Definition 4. Let L' and L be languages and Ω a family of oracle operators. L' is *Turing-reducible* to L *via* Ω, denoted $L' \leq_T^{\Omega} L$, if there exists an oracle operator $\phi \in \Omega$ such that $\phi(L)=L'$. If M is an oracle machine computing ϕ, we also write $L' \leq_T L$ via ϕ or via M. If K any class of languages, then we let Ω^K denote the class of languages that are Turing-reducible to some language in K via some operator in Ω. If K consists of a single language A, we also write Ω^A. ■

Notation. Resource bounded complexity classes are defined by suitable resource bounded acceptors. If we replace these acceptors by oracle machines with corresponding resource bounds, where the oracles come from some class K, then the class of languages so defined is denoted in the original manner except that we add K as a superscript. For instance, if F any family of complexity functions then $DTIME^K(F)$ denotes the class of languages accepted by oracle machines $\mathrm{M}^{(A)}$ in time $t \in F$ with $A \in K$. Or again, $N\text{-}TIME\text{-}SPACE^K(n^{O(1)}, \log n)$ is the class of languages that are Turing-reducible to languages in K in simultaneous time-space $(n^{O(1)}, \log n)$. We call the classes defined in this way *relativized classes*. The classes in the canonical list can likewise be 'relativized', and we have:

$$DLOG^K, NLOG^K, P^K, NP^K, PSPACE^K, etc.$$

Any question involving complexity classes can now be asked of relativized classes. In particular, the inclusion questions can be relativized; for instance, the NP versus P problem can be relativized as follows:

Does there exist an oracle A such that $NP^A \subseteq P^A$?

Relativized classes are treated in Chapter 9.

For the time being we consider the deterministic oracle operators; in Section 9 we return to the nondeterministic case. The two major families of deterministic oracle operators are $\mathcal{P}$ and $\mathcal{DLOG}$, consisting of those oracle operators computable by some deterministic query machine that runs in polynomial time and logarithmic space (respectively).[4] If $\Omega = \mathcal{P}$, the corresponding polynomial-time reducibility is denoted $\leq_T^P$. Similarly, if $\Omega = \mathcal{DLOG}$, we denote the reducibility by $\leq_T^L$. Cook's Theorem was originally stated using $\leq_T^P$ instead of $\leq_m^P$. In the literature, $\leq_T^P$ is also known as *Cook reducibility*.

LEMMA 6.

(i) $\leq_T^P$ *is a transitive reducibility.*

(ii) $\leq_T^L$ *is a transitive reducibility.*

LEMMA 7.

(i) P *is closed under* $\leq_T^P$.

(ii) *DLOG, NLOG and PLOG are closed under* $\leq_T^L$.

The proofs of lemmas 6(i) and 7(i) are straightforward, and the proofs of lemmas 6(ii) and 7(ii) are similar to the proof of Theorem 3. The nondeterministic version of Lemma 7(i) is not known to be true: it is not clear that NP is closed under $\leq_T^P$ (see Exercises). To gain more insight into this question, it is instructive to follow the proof of our next result:

LEMMA 8. *The class* $NP \cap$ co-NP *is closed under Cook reducibility.*

Proof. Let A be accepted by an nondeterministic polynomial-time acceptor M_0, and co-A be accepted by a similar acceptor M_1. Suppose Q is a deterministic query machine that reduces B to A. Hence $B = L(Q^{(A)})$. We construct a nondeterministic polynomial-time acceptor (call it N) to accept B as follows: on any input x, we simulate Q until a query z is made. At that moment, instead of asking an oracle, we dovetail two computations: M_0 on z and M_1 on z. If M_0 accepts, we continue the simulation of Q from the YES state of the query machine. If M_1 acccepts, we continue the simulation of Q from the NO state. We accept iff Q accepts.

We claim that if $L(N) = B$, and N accepts in polynomial time. Suppose $x \in B$. There there is an accepting computation path of Q on input x, and corresponding accepting computation paths of M_0 and M_1 (depending on whether the query was a YES or NO) which can be put together to form an accepting computation path of N of polynomial length. If $x \notin B$, then it is easy to see that no computation path of N is accepting.

Since Q is deterministic, by interchanging the accept and reject states of Q, we obtain a query machine Q' such that co-$B = L(Q'^{(A)})$. Then the same argument as before proves that co-B is accepted by a nondeterministic polynomial time machine. The theorem follows. **Q.E.D.**

Turing reducibility is the strongest notion of reducibility known. For example, it is stronger than many-one reducibility if we view a transducer as a special type of oracle machine in which the oracle is asked only one question at the end of the computation (and furthermore, this oracle machine accepts if and only if the oracle says yes).

[4] We use script letters for families of oracle operators.

Alternative oracle models. The oracle machines as defined above are called *one-way oracle machines*. Although such machines will serve as the standard model of oracle computation in this book, we point out some unexpected properties. This leads to interest in alternative formulations. These properties arise in space-bounded oracle computation. Note that the oracle tape is not counted among the work-tapes. The alternative is to allow the oracle tape to behave like an ordinary work-tape until the machine enters the query state: then, as in the one-way model, the machine next enters the YES or NO state, and the oracle tape is blanked in an instant. We call this the *two-way oracle machine*. The amount of space used by the oracle tape is now included when considering space usage. We give one motivation for preferring the one-way model: we generally expect Turing reducibility to be as strong as the corresponding many-one reducibility. For instance, it seems that $\leq_T^L$ ought to be as strong as the many-one reducibility $\leq_m^L$. Note that a one-way oracle machine using log-space may make queries of polynomial length. Under the two-way oracle model the machine can only make queries of logarithmic length, and hence it is no longer evident that $\leq_T^L$ is as strong as $\leq_m^L$. In fact, $\leq_T^L$ is not even a reducibility (but $\leq_m^L$ clearly is a reducibility — Exercise) The next two results illustrate further differences between the two oracle models.

THEOREM 9. *Under the two-way oracle model, for any oracle A and any complexity function s:*

(i) *(Relativized Savitch's theorem) $NSPACE^A(s) \subseteq DSPACE^A(s^2)$.*

(ii) $NSPACE^A(s) \subseteq DTIME^A(n^2 \log n O(1)^{s(n)})$.

The relativized Savitch's theorem was noted in [18]. The proof of this theorem is an exercise but it is essentially the same as in the unrelativized cases. In contrast, with respect to the one-way machine model, there are oracles that belie the above result. We state such a result of Ladner and Lynch [9] here, but defer its proof till chapter 9.

THEOREM 10. *Let s be a complexity function that is space-constructible. Then there is an oracle A such that $NSPACE^A(s) \not\subseteq DSPACE^A(O(1)^s)$.*

This theorem contradicts the relativized Savitch's theorem in a very strong sense: there is an oracle A such $NSPACE^A(s) \not\subseteq DSPACE^A(s^k)$ for every integer k. It also contradicts the unrelativized result that $NSPACE(s)$ is included in $DTIME(O(1)^s)$, for $s(n) \geq \log n$. Such properties have suggested to researchers that our definitions of oracle space may be "pathological". To obtain a more reasonable notion, Gasarch suggests the following:

Definition 5. A *tame* oracle machine is a one-way oracle machine that has a distinguished set of states called *oracle states*. The machine only writes on the oracle tape when in these states, and transitions from oracle states are deterministic and can only go into other oracle states or enter the QUERY state. ■

Tameness implies that the oracle machine precedes an oracle query by entering the oracle states. If the machine uses s space, then it has $n2^{O(s(n))}$ configurations, where a configuration here does not include the contents on the oracle tape. Hence the length of the query word in a tame oracle machine using space $s(n)$ is $n2^{O(s(n))}$. It is now easy to prove the analog of theorem 9:

THEOREM 11. *Theorem 9 holds under the tame oracle model.*

From now on, whenever space complexity is considered, we will assume that the one-way oracles are tame.

4.4 Efficient Universal Machines

In Chapter 0, we encountered the concept of a universal Turing machine (UTM). We now extend UTM's to the multitape model (this is routine) and further more take complexity into account (slightly less obvious).

Definition 6. A **k-itape universal Turing acceptor** U is a Turing acceptor with $k \geq 0$ work-tapes and two extra read-only tapes called the **index tape** and the **input tape**. Let Σ be the input alphabet of U. Inputs for U are pairs of the form $\langle i, x \rangle$ where $i, x \in \Sigma^*$, with i on the index tape and x on the input tape. We usually inteprete i as an natural number (in $|\Sigma|$-adic notation). For each $i \in \mathbb{N}$, let L_i be the language

$$L_i = \{x \,:\, U \text{ accepts } \langle i, x \rangle\}.$$

We call U is a *universal acceptor* for a class K if:

(i) For all i, $L_i \in K$.

(ii) For each $L \in K$, there are infinitely many indices i such that $L_i = L$. Call this the **recurrence property**.

Alternatively, we say U *accepts* K or U *presents* K. Finally, if K is an arbitrary class, we say K *has universal acceptors* (or is *presentable*) if for each alphabet Σ, $K|\Sigma$ has a universal acceptor. ■

This definition updates the definition from Chapter 0 multitape Turing machines. The only twist is our requirement of the recurrence property – this will make all our subsequent argument simpler. With i fixed on the index tape, we regard U as a Turing acceptor, denoted U_i, that accepts $L(U_i)$. The number i is the **index** of the machine U_i. We can also identify U with the family $\{U_i : i = 0, 1, 2, \ldots\}$ of Turing acceptors. However not every set of Turing acceptors qualifies to be called a universal acceptor. Note that U can be deterministic or nondeterministic.

Example 2. (Standard Universal Machine Construction) Fix Σ. We construct a universal acceptor U^{RE} for the class $RE|\Sigma$. For fixed $k \geq 2$, let $\mathcal{M}_0$ be the family of all nondeterministic k-tape Turing acceptors whose input alphabet is Σ. So the machines in $\mathcal{M}_0$ accept precisely the class $RE|\Sigma$. We represent each acceptor in $\mathcal{M}_0$ by its finite transition table, listing in some arbitrary order the tuples in the table. Our universal Turing machine U^{RE} will simulate each machine from this representation. Although the input alphabet of machines in $\mathcal{M}_0$ is fixed as Σ, the tape alphabets is still arbitrary. Since U^{RE} must use a fixed alphabet, we will encode each symbol in the universal symbol set Σ_∞, say, as a word in $\{0,1\}^*$. Once this is done, each machine $M \in \mathcal{M}_0$ is naturally represented by *infinitely many binary strings* which we call **indices** of M. There are infinitely many indices because we can add any number of spurious instructions that can never be executed in a computation beginning from an initial configuration. An index i is interchangeably regarded as a natural number. However, when we write "$|i|$", we mean the length of the binary string i, *not* the absolute value of i as an integer. In case a binary string i does not encode a valid machine description, we say it encodes the **null machine** that always halt and rejects its input. The preceding conventions amount to an enumeration of the machines in $\mathcal{M}_0$:

$$\phi_0, \phi_1, \phi_2, \ldots \tag{3}$$

where each machine in $\mathcal{M}_0$ occurs infinitely often in the list. Now it is easy to construct U^{RE} that upon input $\langle i, x\rangle$ simulates the ith acceptor ϕ_i on input x in a step by step fashion; we leave the details to the reader. We can make U^{RE} nondeterministic or deterministic, as desired. To make U_i^{RE} nondeterministic, we allow the nondeterministically choice of not execute an executable instruction that is found in i but to scan i further for another executable instruction. To make U_i^{RE} deterministic, we always execute the first executable instruction we find in i. This completes our description of the (deterministic or nondeterministic) universal machine U^{RE}. Our construction has the following properties:

(P1) For each machine ϕ_i in $\mathcal{M}_0$, if ϕ_i accepts in simultaneous time-space-reversal (t, s, r) then U_i^{RE} accepts $L(\phi_i)$ in time-space-reversal $O_i(t, s, r)$.

(P2) (The **efficient recurrence property**) For each i, there are infinitely many j's such that $L(U_i^{RE}) = L(U_j^{RE})$ and, U_i^{RE} accepts in simultaneous time-space-reversal $O_i(t, s, r)$ iff U_j^{RE} accepts in simultaneous time-space-reversal $O_j(t, s, r)$.

Property (P1) follows from the observation that U_j^{RE} uses $O_i(1)$ units of its resource for each unit of the corresponding resource of the ϕ_i. In Property (P2), we allow complexity of U_i and U_j to differ up to a constant (depending on i and j). ■

Efficient Presentation. In Complexity Theory, it is not enough to say that a class K has a universal acceptor, for we also want the universal acceptor to be ""efficient". Intuitively, the machine U^{RE} in the above example is efficient because (P1) and (P2) holds. On the other hand, it is easy to construct an "inefficient" universal acceptor U for the class NP in which each U_i takes exponential time. Intuitively, U is efficient for NP if for each i, U_i accepts a language in NP in non-deterministic polynomial time. Unfortunately, NP is simply a merely class of languages, and has no inherent notion of "non-deterministic polynomial time". The later is a complexity characterization that is used in its usual definition, but it might only be an accidental fact. We therefore proceed as follows: we *explicitly* attach the complexity characterization $\chi =$"non-deterministic polynomial time" to the class NP and call the pair (NP, χ) a "characteristic class" based on NP. There are other characteristic classes based on NP. For instance, if $P = NP$, then we would have another characteristic class (NP, χ') based on NP, where $\chi' =$"deterministic polynomial time".

Definition 7. A **resource bound** is a pair $\beta = (F, \rho)$ where F is a family of complexity functions and ρ is a computational resource such as time or space. A **complexity characteristic** χ is a tuple of the form

$$(\mu, \beta) \text{ or } (\mu, F, \rho)$$

where μ is a computational mode, and $\beta = (F, \rho)$ is a resource bound. Given χ, let $\mathcal{M}(\chi)$ be the family of all multitape Turing machines with complexity characteristic χ.

A **characteristic class** is a pair (K, χ) where K is a language class and χ a complexity characteristic such that for each $L \in K$, there is some machine in $\mathcal{M}(\chi)$ that accepts L. We say (K, χ) is *based* on K. ■

Note that the definition of characteristic class (K, χ) does not say that every machine in $\mathcal{M}(\chi)$ must accept a language in K. Thus, χ is only an upper bound. In particular, we can always base K on the $\chi = (\mu, F_\infty, \rho)$ where F_∞ is the family of all complexity functions. We can also generalize complexity characteristics χ to allow several resource bounds, $\chi = (\mu, \{\beta_1, \ldots, \beta_k\})$, to discuss resource bounds $\beta_1, \ldots, \beta_k$ simultaneously. We may also vary other parameters such as: the choice between running and acceptance complexity, the number of work-tapes, etc.

We usually write F-ρ instead of (F, ρ) for resource bounds. For instance,

$$n^{O(1)}\text{-}time, \qquad \log^{O(1)} n\text{-}space$$

are resource bounds. This may be read as *poly-time* and *polylog-space*. This notation extends naturally. As further examples, the complexity characteristics

$$(nondeterministic, log(n)\text{-}reversals), (deterministic, O(1)^n\text{-}time).$$

are called *nondeterministic log-reversals* and *deterministic exponential-time*, respectively.

Example 3. The Canonical Characteristic Classes. For each of the classes in the canonical list, there is a natural complexity characteristic associated with its definition. Thus the class $DLOG$ of deterministic log-space languages is associated with the characteristic

$$(deterministic, \log n\text{-}space).$$

Indeed, the names for these classes explicitly refer to these characteristics. Henceforth, when we refer to a "characteristic class" K in the canonical list, we assume the complexity characteristic of K used in its definition. As another example, 'the characteristic class NP' refers to

$$(NP, (nondeterministic, n^{O(1)}\text{-}time)).$$

In general, for any characteristic class (K, χ), if χ is understood, we simply refer to K as the characteristic class itself. ■

In the literature, the distinction between a characteristic class and its underlying abstract abstract class of languages is often only implicit.

Example 4. The definition of characteristic classes allows considerable freedom in choosing the complexity characteristic χ to be attached to any class K. Thus, Savitch's theorem (chapter 2) tells us that the following are two characteristic classes based on $PSPACE$:

$$(PSPACE, (deterministic, n^{O(1)}\text{-}space)),$$

$$(PSPACE, (nondeterministic, n^{O(1)}\text{-}space)).$$

Below we will prove a non-trivial fact about the following interesting characteristic class:

$$(NPC, (nondeterministic, n^{O(1)}\text{-}time))$$

where $NPC \subseteq NP$ is the class of NP-complete languages under (say) Karp reducibility. If $P \neq NP$ then we see in section 7 that there exist languages in NP that are not in NPC. ■

We come to our main definition.

Definition 8. Let (K, χ) be a characteristic class, $K = K|\Sigma$, and U a universal machine. We say U is *an efficient universal acceptor for* (or is *an efficient presentation of*) the characteristic class (K, χ) if

(a) For each i, $L(U_i) \in K$.

(b) Each U_i has complexity characteristic χ. Here, reversals on the first input tape (containing i) are not counted, and space on both input tapes is not counted.

(c) U satisfies the efficient recurrence property (property (P2) ahove): for $i \in \mathbb{N}$ there exists infinitely many $j \in \mathbb{N}$ such that $L(U_i) = L(U_j)$, and for all complexity functions t, s, r, U_i accepts in simultaneous time-space-reversal (t, s, r) iff U_j accepts in simultaneous $O_j(t, s, r)$.

In general, when K is not necessarily of the form $K = K|\Sigma$, we say (K, χ) *has efficient universal acceptors* if for each alphabet Σ, there is an efficient universal machine for $(K|\Sigma, \chi)$. ■

Example 5. (Efficient universal acceptors for P) Fix an alphabet Σ. We will construct a universal acceptor U^P for the characteristic class $(P\,|\Sigma, \chi)$ where $\chi = (deterministic, n^{O(1)}\text{-}time)$. In example 2, we described a universal acceptor U^{RE} for $RE|\Sigma$. Recall that we may assume U^{RE} is deterministic. We obtain U^P from U^{RE} as follows: U^P, upon input $\langle i, x\rangle$ simulates U^{RE} on input $\langle i, x\rangle$ for at most $|x|^{|i|}$ steps (these are steps of U_i^{RE}).[5] If U_i^{RE} does not accept x within $|x|^{|i|}$ steps, then U^P rejects. U^P can keep track of the number of simulated steps since $f(n) = n^{|i|}$ is time-constructible and U^P can in parallel (i.e., on separate tapes) run for exactly $f(n)$ steps while simulating U_i^{RE}. Furthermore, we need the fact that the family of polynomials $\{n^k : k = 1, 2, \ldots\}$ are uniformly constructible, i.e., the two argument "universal-polynomial" function $p(n, k) := n^k$ is time-constructible (Exercise). We call U^P a 'clocked' version of U^{RE}.

We claim that U^P is an efficient presentation of the characteristic class $P\,|\Sigma$. Clearly U^P accepts only languages in $P\,|\Sigma$. Each language L in $P\,|\Sigma$ is accepted by some M in time $O(n^h)$ for some $h \geq 1$. By a tape reduction result in chapter 2, there is a 2-tape machine in M_1 accepting L in time $O(n^h \log n) = O(n^{h+1})$. Then property (P1) in example 2 implies that there exists an i such that U_i^P accepts L in time $O(n^{h+1})$. Finally, we must show the efficient recurrence property (P2). But this follows from our standard coding of machines. ■

We now show efficient universal machines for the characteristic classes in the canonical list. The technique is simple and generally applicable; it involves attaching a 'clock' to each Turing machine defined by a suitable universal machine. But this is precisely the technique illustrated in the last example.

Definition 9. A universal machine U is *total* if for all inputs $\langle i, x\rangle$, U eventually halts on that input. A class is *recursively presentable* if it has a total universal machine. ■

THEOREM 12. *Each characteristic class in the canonical list is efficiently presentable by a total universal machine.*

Proof. The construction for P in the above example can be generalized to all the other classes. Here we use the fact that the families of complexity functions $(n^{O(1)}, \log n$, etc) are all constructible in the corresponding complexity resource (time or space). **Q.E.D.**

Universal Transducers and Universal Oracle Machines. We can extend the preceding definitions for acceptors to transducers and oracle machines. This can be done in a rather straightforward way. Hence we are contented with a quick description. We define **universal transducers** and **universal oracle machines**, simply by giving an ordinary transducer or oracle machine an extra read-only input tape, so that input is now a pair $\langle i, x\rangle$. Let Φ and Ψ be families of transformations and oracle operators, respectively. Again we need the concept of *characteristic families* of transducers or oracles, denoted (Φ, χ) and (Ψ, χ) where χ is a complexity characteristic. For each alphabet Σ, we may define $\Phi|\Sigma$ (resp. $\Psi|\Sigma$) to be the subfamily of transformations (resp. oracle operators) whose input and output alphabets are Σ.[6] For characteristic families of transducers or oracles, there is a corresponding notion of efficient presentation for the family. Using the idea of clocks, the reader may prove as an exercise:

LEMMA 13.

(i) *The characteristic families* **DLOG** *of log-space transformations and* **P** *of polynomial-time transformations are each efficiently presentable.*

(ii) *The characteristic families $\mathcal{DLOG}$ of log-space oracle operators and $\mathcal{P}$ of polynomial-time oracle operators are each efficiently presentable.*

Definition 10. Two languages $L, L' \subseteq \Sigma^*$ are said to be *finite variants* of each other if $(L - L') \cup (L' - L)$ is a finite set. ■

We extend this definition by saying that two classes of languages $K = K|\Sigma$ and $K' = K'|\Sigma$ are finite variants of each other if for each $L \in K$, there is a finite variant of L in K' and vice-versa. An arbitrary class K is *closed under finite variations* if for each Σ,

[5] More precisely, if U_i^{RE} simulates the ith machine ϕ_i in the standard listing of $\mathcal{M}_1$, then each step of ϕ_i is simulated by $O(|i|) = O_i(1)$ steps of U^{RE}.

[6] We do not need the more general case where the input and output alphabets are different.

(a) K contains the trivial languages $\emptyset$ and Σ^* and

(b) for all $L, L' \subseteq \Sigma^*$ that are finite variants of each other, $L \in K$ if and only if $L' \in K$.

It is easy to see that all the complexity classes defined in this book are closed under finite variation.

We next prove an interesting theorem by Landweber, Lipton and Robertson [11]. It concerns the characteristic class *NPC* of example 4 and depends on the efficient presentability of the characteristic classes *NP* and **P**.

THEOREM 14. *The characteristic class NPC of NP-complete languages under Karp reducibility is efficiently presentable.*

Proof. Let $U = \{U_i\}$ be an efficient universal acceptor for the characteristic class *NP* and let $T = \{T_i\}$ be an efficient universal transducer for the polynomial time transformations. (The notation $\simeq$ next refers to the pairing function defined in the appendix.) We construct $U' = \{U'_i\}$ such that if $i \simeq \langle j, k\rangle$ then U'_i accepts $L(U_j)$ if SAT is Karp reducible to $L(U_j)$ via T_k; otherwise U'_i accepts some finite variant of SAT. We prove that U' is indeed the desired universal machine.

We show how U' operates: on input x of length n, U'_i *deterministically* verifies for each y of length $\leq \log\log n$ that

$$y \in \text{SAT} \iff T_k(y) \in L(U_j). \tag{4}$$

We say the verification for a given y *succeeds* if and only if (4) can be verified (using the procedure given below) within n steps. We say that x is a *success* if and only if the verification succeeds for each y of length $m \leq \log\log n$; otherwise x is a *failure*. U' first marks out $\lceil \log\log n \rceil$ tape-cells in $O(n)$ time. (We leave the details for this to the reader – it is similar to the proof in the appendix that $\log^* n$ can be computed in linear time.) To verify a particular y of length m, U' proceeds as follows:

(i) Determine if $y \in$ SAT. Fix any nondeterministic acceptor for SAT that takes time m^c on y, $m = |y|$, for some constant c that does not depend on the input. By results in chapter 2, we can simulate this acceptor deterministically in time $O(1)^{m^c}$.

(ii) Compute $T_k(y)$, using time m^d for some constant d that depends only on k.

(iii) Determine if $T_k(y) \in L(U_j)$. Suppose U_j on an input of length n takes nondeterministic time n^e for some e depending on j. Since $|T_k(y)| \leq m^d$, a deterministic simulation of the nondeterministic U_j takes time $O(1)^{(m^d)^e} = O(1)^{m^{de}}$.

For n sufficiently large, we see that (i–iii) can be done in n steps. Thus deciding if x is a success or failure takes at most n^2 time since there are at most $O(1)^{\log\log n} \leq n$ values of y to check. If x is a success then U_i next simulates U_j on $T_k(x)$; otherwise it simulates the fixed nondeterministic acceptor for SAT on input x. This completes our description of U'.

We now show that U' is an efficient universal acceptor for *NPC*. Suppose $i \simeq \langle j, k\rangle$. If SAT is Karp-reducible to $L(U_j)$ via T_k then we see that for *all* inputs, U'_i behaves exactly like U_j. Conversely, if SAT is not Karp-reducible to $L(U_j)$ via T_k then for *large enough* inputs, U'_i behaves like a particular acceptor for SAT. Since *NPC* is closed under finite variation, the accepted language is still in *NPC*. For the efficient recurrence property, it is not hard to show that for any $L' \in NPC$, there are infinitely many i's such that U'_i accepts L' using $O_i(n^c)$ time, for some c. **Q.E.D.**

The above proof is rather formal, but it is really quite simple in outline: On input x, $U_{\langle j,k\rangle}$ spends a small amount of time, say $f(|x|)$ steps, looking for a counter example to the assertion that SAT is reduced to $L(U_j)$ via T_k. If it fails within $f(n)$ steps, it simulates an algorithm for SAT on input x; if it succeeds, it simulates U_j on input x. For j, k fixed, $U_{\langle j,k\rangle}$ is NP provided $f(n)$ is polynomial-bounded (but clearly any unbounded non-decreasing function). It is easy to see that this machine accepts U_j or some finite variant of SAT. The point is that if the assertion is false, we will find the counterexample in finite time.

In contrast, we have the result of Chew and Machtey [4] that if $P \neq NP$ then $NP - P$ has no recursive presentation.

4.5 Truth-table Reducibilities

This section introduces a reducibility based on the idea of 'truth-tables'; it is intermediate in strength compared to Turing-reducibility and many-one reducibility. We may view many-one reducibility via a transformation t as a special case of Turing reducibility where the query machine on input x asks only one question

"Is $t(x)$ in the oracle language?"

at the end of the computation, with the decision to accept identical to the answer from the oracle. In truth-table reducibility, this is generalized to the asking more than one question with the decision to accept determined by some fixed function of the answers to these questions. The difference between this and general Turing reducibility is that the set of questions asked is solely a function of x, independent of the oracle. Thus a query machine can be decomposed into two parts: a 'query generator' which formulates a sequence of questions based on the input x, and a 'query evaluator' which evaluates a predicate on the sequence of answers. A simple example is a query generator which on any input x produces two queries (strings) $f(x)$ and $g(x)$, where f and g are transformations; the query evaluator accepts if $f(x)$ but not $g(x)$ is in the oracle.

For any language (Σ, L), define its *indicator function* as the function $\chi_L : \Sigma^* \to \{0,1\}$ such that $\chi_L(x) = 1$ iff $x \in L$.

Definition 11. Let Σ and Γ be alphabets.
(a) A *query generator* is a transformation $g : \Sigma^* \to (\Gamma \cup \{\#\})^*$, where $\#$ is a symbol not in Γ. If $g(x) = y_1\#y_2\#\cdots\#y_k$, we view the y_i's as queries to an oracle.
(b) A *query evaluator* is a language $e \subseteq \Sigma^*\#\{0,1\}^*$ (i.e., each word in e has the form $w\#x$ where $w \in \Sigma^*$ and $x \in \{0,1\}^*$).
(c) A *truth-table* is a pair (g, e) where g and e have the forms (a) and (b) respectively.
(d) Say (Σ, L) is *truth-table reducible* to (Γ, L') *via* (g, e), written

$$L \leq_{tt} L' \text{ via } (g, e),$$

if for all x in Σ^* where $g(x) = y_1\#y_2\#\cdots\#y_k$, we have

$$x \in L \iff x\#\chi_L(y_1)\chi_L(y_2)\cdots\chi_L(y_k) \in e.$$

■

Definition 12. A (deterministic) *truth-table machine* is a pair (G, E) where G is a deterministic transducer computing g and E is a deterministic acceptor for some language e, for some truth table (g, e). We say (g, e) is *polynomial-time* if it is computed by some truth-table machine (G, E) where the generator G and the evaluator E both compute in polynomial-time. In general, to say that a truth-table machine operates within some complexity bounds means that both the generator and evaluator have the same complexity bounds. ■

We defer the definition of nondeterministic truth-table machines to section 9. Let 'tt' abbreviate 'truth-table', as in 'tt-machine' or 'tt-reducibility'. We now collect the various types of reducibilities obtained by restrictions on the query generator g and evaluator e. Suppose $(\Sigma, L) \leq_{tt} (\Gamma, L')$ via (g, e).

(i) The truth-table (g, e) is *positive* if e satisfies the restriction that for all $x \in \Gamma^*$ and $b_i, b_i' \in \{0,1\}$,

$$x\#b_1b_2\cdots b_k \in e \text{ and } \bigwedge_{i=1}^{k}(b_i \leq b_i') \text{ implies } x\#b_1'b_2'\cdots b_k' \in e.$$

The corresponding truth-table reducibility is described as *positive* and denoted $\leq_{ptt}$.

(ii) The truth-table is *conjunctive* if e satisfies the restriction that for all $x \in \Gamma^*$ and $b_i \in \{0,1\}$,

$$x\#b_1\cdots b_k \in e \iff \bigwedge_{i=1}^{k}(b_i = 1).$$

Similarly, *disjunctive* truth tables are defined by using disjuncts '$\bigvee$' instead of the conjuncts '$\bigwedge$' above. The corresponding truth-table reducibilities are denoted $\leq_{ctt}$ and $\leq_{dtt}$, respectively.

(iii) Let k be a positive integer. If g makes k queries on every input x then (g, e) is called a *k-truth-table*. We then say L is k-truth-table reducible to L' via (g, e) and the corresponding notation is $L \leq_{k\text{-}tt} L'$. Similarly, any of the above restrictions on $\leq_{tt}^{P}$ can be modified by the parameter k:

$$\leq_{k-ptt}, \leq_{k-ctt}, \leq_{k-dtt}.$$

(iv) A truth-table (g, e) is *bounded* if it is a k-truth-table for some k. We then say L is *bounded truth-table reducible* to L' via (g, e), denoted $L \leq_{*\text{-}tt} L'$. Similarly, the other restrictions on $\leq^P_{tt}$ can be modified: $\leq_{*-ptt}, \leq_{*-ctt}, \leq_{*-dtt}$.

Clearly a many-one reducibility is simultaneously a positive, conjunctive and a disjunctive truth-table as well as a 1-truth-table.

Notations: If (g, e) is computed by a (deterministic) polynomial-time tt-machine and L is truth-table reducible to L' via (g, e) then we write $L \leq^P_{tt} L'$; more generally, if (g, e) satisfies any of the restrictions in (i-iv) above, we append a superscript P to the corresponding reducibility notation. Thus $\leq^P_{k\text{-}tt}$ denotes polynomial-time k-truth-table reducibility. Similarly, if we restrict the space usage of the query generator and query evaluator to logarithmic space, then we append the superscript L. Thus we have $\leq^L_{tt}, \leq^L_{ptt}, etc.$

LEMMA 15. *The following truth-table reducibilities are transitive:*

$$\leq^P_{tt}, \leq^P_{*\text{-}tt}, \leq^P_{ptt}, \leq^P_{ctt}, \leq^P_{dtt}, \leq^P_{1\text{-}tt}.$$

Among the special tt-reducibilities we have introduced, the list in this lemma omits just the k-tt-reducibilities where $k > 1$. We prove the lemma for the case of $\leq^P_{1\text{-}tt}$, leaving the rest as exercise.

Proof. Let A, B, C be languages. Suppose $A \leq^P_{1\text{-}tt} B$ via (g, e) and $B \leq^P_{1\text{-}tt} C$ via (g', e'). We will show that $A \leq^P_{1\text{-}tt} C$. Now, $x \in A$ iff $x\#\chi_B(g(x)) \in e$ and $y \in B$ iff $y\#\chi_C(g'(y)) \in e'$. Substituting $g(x)$ for y, we see that $x \in A$ holds iff either

$$x\#1 \in e \text{ and } g(x)\#\chi_C(g'(g(x))) \in e'.$$

or

$$x\#0 \in e \text{ and } g(x)\#\chi_C(g'(g(x))) \notin e'.$$

holds. That is, iff one of the following holds:

(a) $x\#1 \in e$ and $g(x)\#1 \in e'$ and $g'(g(x)) \in C$.

(b) $x\#1 \in e$ and $g(x)\#0 \in e'$ and $g'(g(x)) \notin C$.

(c) $x\#0 \in e$ and $g(x)\#1 \in e'$ and $g'(g(x)) \notin C$.

(d) $x\#0 \in e$ and $g(x)\#0 \in e'$ and $g'(g(x)) \in C$.

Now define the truth-table (g'', e'') where $g''(x) = g'(g(x))$ and e'' is given by the following two rules. A word $x\#1$ is in e'' iff either

$$x\#1 \in e \text{ and } g(x)\#1 \in e'$$

or

$$x\#0 \in e \text{ and } g(x)\#0 \in e'$$

holds. Again, a word $x\#0$ is in e'' iff either

$$x\#1 \in e \text{ and } g(x)\#0 \in e'$$

or

$$x\#0 \in e \text{ and } g(x)\#1 \in e'$$

holds.
The reader can verify that $A \leq_{1\text{-}tt} C$ via (g'', e'') and that (g'', e'') can be computed in polynomial time. **Q.E.D.**

It is also easy to show

LEMMA 16. *The following lists of reducibilities (modified by any given complexity bound β such as polynomial-time or log-space) are in order of non-decreasing strength:*

(i) $\leq^\beta_m, \leq^\beta_{k\text{-}tt}, \leq^\beta_{(k+1)\text{-}tt}, \leq^\beta_{*\text{-}tt}, \leq^\beta_{tt}, \leq^\beta_T$

(ii) $\leq^\beta_m, \leq^\beta_{ctt}, \leq^\beta_{ptt}, \leq^\beta_{tt}$

In (ii), we could have substituted $\leq^{\beta}_{dtt}$ for $\leq^{\beta}_{ctt}$.

Alternative definition of tt-reducibilities. Often the truth-table (g, e) embodies 'Boolean functions' in the sense that on any input x where $g(x) = y_1\# \cdots \#y_k$ then there is a Boolean function f_x on k variables such that $x\#b_1b_2 \cdots b_k \in e$ if and only if $f_x(b_1, \ldots, b_k) = 1$. This leads to the following alternative definition of tt-reducibilities. Fix a representation of Boolean functions as binary strings, and let f_w denote the Boolean function f_w represented by the binary string w. With query generators g defined as before, we now define an *evaluator* e as a transformation $e : \Sigma^* \to \{0,1\}^*$; intuitively, $e(x)$ represents the Boolean function $f_{e(x)}$ used to determine the final answer. We say (Σ, A) is tt-reducible to (Γ, B) via (g, e) if for all $x \in \Sigma^*$ the following is satisfied: $x \in A$ iff

$$f_{e(x)}(\chi_B(y_1), \chi_B(y_2), \ldots, \chi_B(y_k)) = 1$$

where $g(x) = y_1\# \cdots \#y_k$ for some $k \geq 1$. The particular method of representing Boolean functions could affect the strength of the reducibility. We consider three increasingly more succinct representations:

(a) A k-variable function can be represented by a Boolean 2^k-vector. In some literature, such a vector is also called a 'truth-table' which in turn lends its name to the entire family of reducibilities in this section.

(b) By Boolean formulas over the basis $\{\wedge, \vee, \neg\}$.

(c) By Boolean circuits over the same basis.

In the Exercises, we see that with respect to polynomial-time bounds, these three alternative definitions of tt-reducibilities do not lead to a different notion of tt-reducibility. We also give another definition of tt-reducibility in the Exercises that is again equivalent to our standard choice. These results lend credibility to the claim that our standard choice is robust. Unfortunately, with respect to log-space bounds, the situation seems to be different.

4.6 Strength of Reducibilities

Among the reducibilities, there are pairs $(\leq_1, \leq_2)$ where it is obvious that $\leq_1$ is as strong as $\leq_2$. For instance, $(\leq_1, \leq_2) = (\leq^P_T, \leq^P_m)$. A natural question is whether $\leq_1$ is, in fact, stronger than $\leq_2$. In this section, we illustrate some of these results. The notion of partial sets will be useful in our proofs:

Definition 13. A *partial set* (relative to a set X) is a partial function f from X to $\{0,1\}$. The *universe* of f is X, and its *domain*, $\texttt{domain}(f)$, consists of elements $x \in X$ such that f is defined at x. If f' is also a partial set with universe X, we say f' and f *agrees* if they are identical on $\texttt{domain}(f) \cap \texttt{domain}(f')$. Call f' a *partial subset* of f (or, f *extends* f') if $\texttt{domain}(f') \subseteq \texttt{domain}(f)$ and f' and f agrees. We write $f' \subseteq f$ in this case. The *complement* of f is the partial function $\overline{f}$ such that $\texttt{domain}(f) = \texttt{domain}(\overline{f})$ and $f(x) = 1 - \overline{f}(x)$ whenever $x \in \texttt{domain}(f)$. Note that $f' \subseteq f$ iff $\overline{f'} \subseteq \overline{f}$. ■

A basic diagonalization outline. The proofs in this section are based on the idea of 'diagonalization'. The reader will better understand these proofs if she keeps in mind the following outline: Suppose the result in question follows from the existence of a language L that satisfies a countably infinite list of conditions

$$C_0, C_1, C_2, \ldots$$

These conditions are finite in the sense that each condition only refers to a finite subset of L and they are 'existential' in that if L is a partial set that satisfies any condition then any extension of L also satisfies that condition. The proof involves constructing the language in stages where in the ith stage we have built a partial set L_i with finite domain such that each of the conditions $C_0, C_1, \ldots, C_i$ is satisfied by L_i. Typically, $\texttt{domain}(L_i) \subseteq \texttt{domain}(L_{i+1})$.

There are many variations to this basic outline. Perhaps there are more than one infinite list of conditions to satisfy. Perhaps we cannot guarantee satisfying a given condition C during any stage; so we may have to try to satisfy the condition C for infinitely many stages, and prove that our attempts will succeed in *some* stage. Instead using the conditions C_i to define the ith stage, we can use an "inverted diagonalization" by considering all strings in the domain: fix some total ordering of the set of input words. Usually, we use[7] the **lexicographic ordering**:

$$x_0, x_1, x_2, x_3, \ldots \tag{5}$$

[7]Despite the terminology, this is NOT the same as dictionary order which does not have strings listed in order of non-decreasing length.

That is, we list the strings in order of non-decreasing lengths, and among those of the same length we use the usual dictionary order. We decide whether x_j ought to belong to the language L in stage j. This is done by attempting to use x_j to satisfy some condition C_i in our infinite list. We may need to return a particular C_i in many stages (but hopefully C_i will be satisfied in some finite state). The advantage of this approach is that our description of the diagonalization process amounts to specifying an acceptor for the constructed language L.

We proceed with the first result that shows that 1-tt-reducibility is stronger than Karp reducibility. Note that for any language L, we have L is 1-tt-reducible to co-L. Therefore our desired result is an immediate consequence of the following from Ladner, Lynch and Selman [10]:

THEOREM 17. *There exists a non-trivial language L in DEXPT such that L is not Karp-reducible to* co-L.

Before proving this theorem, it is instructive to give a simple proof of a similar result that does not require L to be in *DEXPT*. The simpler proof will follow the basic diagonalization outline above: we construct $L \subseteq \{0,1\}^*$ in stages. Let

$$T = \{T_0, T_1, T_2, \ldots\}$$

be an efficient universal transducer for the polynomial-time transducers used in Karp reducibility. The desired result follows if L satisfies the following infinite list of conditions:

$$C_i : L \text{ is not many-one reducible to co-}L \text{ via } T_i$$

for $i = 0, 1, 2, \ldots$. Note that condition C_i is equivalent to the existence of a word x satisfying

$$x \in L \iff T_i(x) \in L.$$

Such an x is said to "cancel" T_i and x is also called a "witness" for condition C_i. Let $x_0, x_1, \ldots$, be the lexicographic ordering of all binary strings. At stage i, assume inductively that we have a partial set L_i that has a finite domain; L_i is a partial subset of the language L that we are trying to construct. Furthermore, for each $h < i$, the condition C_h is satisfied. We now wish to define L_{i+1} whose domain contains the domain of L_i. Suppose x_j is the next word (in the lexicographical ordering) not in the domain of L_i. Let $y = T_i(x_j)$ and we consider two possibilities: if y is already decided, then put x_j into the domain of L_{i+1} in such a way that x_j cancels T_i; if y is not yet decided, let us put both x_i and y into L (again cancelling T_i). It is not hard to show that condition C_i is satisfied in stage i: one only need to see that x_j exists so that the plan to cancel T_j can be carried out. One should also verify that our arguments is valid even if $x_j = y$.

This simple proof tells us little about the complexity of L (although L is evidently recursive). We need a technique due to Machtey and others to "slow down the diagonalization". To do this, we use the slow growing function $\log^* n$ and its inverse $\mathrm{Exp}(n)$ defined in the appendix.

Proof of theorem 17. We will use the "inverted diagonalization outline" in which the ith stage examines the ith word in a lexicographic listing of all words. We define a Turing acceptor M that accepts a language L over the alphabet $\Sigma = \{0,1\}$ in *DEXPT* where L is not Karp-reducible to co-L. On input $x \in \{0,1\}^*$ where $|x| = n$:

(a) If x is the empty word ϵ, accept.

(b) If x contains a '1' or if $n \neq \mathrm{Exp}(m)$ for some m then reject. Note that this step takes linear time since we can check if $n = \mathrm{Exp}(\log^* n)$ in linear time, as shown in the appendix.

(c) Hence assume $x = 0^n$ (a string of zeroes) such that $n = \mathrm{Exp}(\log^* n))$. Let $\log^* n = m \simeq \langle i, j\rangle$ for some m, i, j. The appendix shows how m, i, j can be extracted in polynomial time (actually, we will not need j). Simulate the universal transducer T on input $\langle i, x\rangle$ for 2^n steps of T. If T_i does not halt within 2^n steps, reject.

(d) Hence assume that T_i produces the output $y = T_i(x)$ within 2^n steps. If $|y| > \mathrm{Exp}(m-1)$ then reject. If $|y| \leq \mathrm{Exp}(m-1)$ then accept iff $y \in L$ (which we determine by a recursive call to M on y).

Thus step (d) uses x to explicitly[8] cancel the condition C_i; the first three steps (a-c) ensure that x is used in step (d) then x has the correct form: $x = 0^n$ where $n = \mathrm{Exp}(\log^* n)$. The construction of M is a 'slow down' of the diagonalization process since condition C_i is (explicitly) cancelled at the $\mathrm{Exp}(i)$th stage or later.

Correctness: We will show that L is non-trivial and is not Karp-reducible to co-L. To see non-triviality, L contains the empty word ϵ and does not have words containing 1's. Aiming for a contradiction, assume that L is Karp-reducible to co-L via some T_i. We may assume that T_i takes $|x|^c$ time, for some $c = c(i)$, on input $\langle i, x\rangle$. Choose j large enough so that if $m \simeq \langle i, j\rangle$ and $n = \mathrm{Exp}(m)$ then

$$2^n > n^c.$$

[8]There may be "accidental" cancellation of some conditions.

Then it is easy to see that on input $x = 0^n$, the acceptor M will proceed to step (d) after obtaining an output $y = T_i(x)$. Hence, $|y| \leq n^c < 2^n = 2^{\mathrm{Exp}(m)} = \mathrm{Exp}(m+1)$. There are two cases:

$$|y| \leq \mathrm{Exp}(m-1)$$

and

$$\mathrm{Exp}(m-1) < |y| < \mathrm{Exp}(m+1). \tag{6}$$

In the first case, by the definition of M, x is in L iff y is in L. In the second case, x is rejected so we must show that y is also rejected. In fact, (6) implies that unless y has the form $0^{\mathrm{Exp}(m)}$, y would be rejected. But if $y = 0^{\mathrm{Exp}(m)}$ then $x = y$, again showing that y is rejected. Thus the algorithm is correct.

Timing Analysis: We show that M runs in time $O(2^n)$ time. Steps (a-c) and the non-recursive part of step (d) clearly takes $O(2^n)$ steps. The recursive call to M has argument y with length $\leq \mathrm{Exp}(m-1) = \log_2 n$. If $t(n)$ is the time required by M then we have $t(n) \leq t(\log n) + O(2^n)$. The solution to this recurrence gives $t(n) = O(2^n)$. □

It should be clear that this proof works with any superpolynomial time bound f in place of 2^n, provided that both f and its 'inverse' are both easy to compute.

The fact that there is a language in *DEXPT* that distinguishes Karp- from tt-reducibility prompts the question: does there exist a language in *NP* that distinguishes them? Of course, if we can prove an affirmative answer then it is easy to see that $NP \neq P$. But even assuming that $NP \neq P$, this is an open question.

Suppose $\leq_1$ is as strong as $\leq_2$. To show that the former is in fact stronger, we have to show the existence of languages A and B such that $A \leq_1 B$ but not $A \leq_2 B$. This is usually quite simple to show, especially if we can carry out the basic diagonalization outline above. But often we would like to impose *side restrictions* on A, B. For instance, we may insist (as in the last theorem) that B does not have very high complexity.

Example 6. One reason for imposing side restrictions is that they yield more information about the relationships between two reducibilities. Ladner, Lynch and Selman define the relation 'stratifies' where $\leq_1$ is said to *stratify* $\leq_2$ if there exists L, L' such that L and L' are of the same $\leq_2$-degree but L and L' are $\leq_1$-incomparable. Clearly $\leq_1$ stratifies $\leq_2$ implies that $\leq_2$ is stronger than $\leq_1$, but intuitively, it is more than 'just' stronger: this is because there is a pair of languages that are similar from the viewpoint of $\leq_2$, while being very different from the viewpoint of $\leq_1$. ■

To illustrate another kind of side restriction, we make the following definition: let A be a language and let $\leq_1, \leq_2$ be reducibilities where $\leq_1$ is as strong as $\leq_2$. Then $\leq_1$ and $\leq_2$ are said to be *distinguishable over A* if there exists a language B such that

$$A \leq_1 B \text{ but not } A \leq_2 B.$$

The next result is from I. Simon [18]:

THEOREM 18. *For all $A \notin P$, $\leq^P_{(k+1)\text{-}tt}$ and $\leq^P_{k\text{-}tt}$ are distinguishable over A.*

Before embarking on the proof, we need the notion of characteristic families of truth-tables, and efficient universal truth-tables. By now, this can be done routinely as follows: Let Ω be any family of truth-tables, and $\chi = (\mu, F - \rho)$ any complexity characteristic. Then (Ω, χ) is a characteristic family provided that each truth-table $(g, e) \in \Omega$ is computed by some tt-machine (G, E) where G and E both have complexity characteristic χ. A *universal tt-machine* is a pair (T, U) where $T = \{T_i\}$ and $U = \{U_i\}$ are a universal transducer and a universal acceptor, respectively. For any alphabet Σ, $\# \notin \Sigma$, we again have the restriction $\Omega|\Sigma$ of Ω to Σ, consisting of $(g, e) \in \Omega$ where $g : \Sigma \to \Sigma \cup \{\#\}$ and $e \subseteq \Sigma \cup \{\#, 0, 1\}$. The characteristic family (Ω, χ) is efficiently presentable if for each Σ, there exists a universal (T, U) such that (i) for each $i \geq 0$, (T_i, U_i) computes a truth-table in $\Omega|\Sigma$, and (ii) each truth-table in $\Omega|\Sigma$ is computed by (T_i, U_i) for infinitely many $i \geq 0$ where each T_i and U_i has complexity characteristic in χ. It is an easy exercise to show that the family Ω^P_{tt} (resp. Ω^P_{k-tt}, for each k) of truth-tables computed by tt-machines with complexity characteristic

$$(deterministic, n^{O(1)}\text{-}time)$$

is efficiently presentable.

Proof of theorem 18. Let A be any language not in P; without loss of generality, assume A is over the alphabet $\{0, 1\}$. The proof follows the basic diagonalization outline above. Let (T, U) be an efficient universal truth-table with deterministic polynomial-time characteristic. We will construct a language $B \subseteq \{0, 1\}^*$ such that

$$x \in A \iff |B \cap S_k(x)| \text{ is odd} \tag{7}$$

where

$$S_k(x) := \{x0^i1^{k-i} \ : \ i = 0, 1, \ldots, k\}$$

Clearly $A \leq^P_{(k+1)\text{-}tt} B$; to show the theorem, we will show that $A \leq^P_{k\text{-}tt} B$ fails. We construct B in stages. Let B_s ($s \geq 0$) be the partial subset constructed at the end of the $(s-1)$st stage. In stage s, we extend B_s to B_{s+1} in such a way that we 'cancel' the truth-table machine (T_s, U_s), i.e., we include a word x in the domain of B_{s+1} such that $x \in A$ if and only if

$$x\#\chi_B(y_1)\cdots\chi_B(y_k) \notin U_s \tag{8}$$

where $T_s(x) = y_1\#\cdots\#y_k$ (clearly we may assume that T_s always makes exactly k queries). We will maintain the inductive hypothesis that the domain of B_s consists of all words of length $\leq h(s)$ for some $h(s) \geq 0$, and such that B_s satisfies (7) in the sense that for all x of length $\leq h(s) - k$, we have $x \in A$ iff $B_s \cap |S_k(x)|$ is odd. We describe the construction of B_{s+1} in four steps:

(a) Choose any x_0 of length $h(s)+1$; note that x_0 is not in the domain of B_s. Let $T_s(x_0) = y_1\#\cdots\#y_k$. We first ensure that each y_i is in the domain of B_{s+1}: if y_i is not in the domain of B_s then we put it in the domain of B_{s+1}, arbitrarily (say $y_i \notin B_{s+1}$).

(b) We next put x_0 in the domain of B_{s+1}, making $x_0 \in B_{s+1}$ or $x_0 \notin B_{s+1}$ so as to satisfy (8). The choices in (a) completely determine this step.

(c) Next include into the domain of B_{s+1} all the remaining words in $S_k(x_0)$ so as to ensure (7). Note that this is always achievable because at least one word in $S_k(x_0)$ has not yet been put into the domain of B_{s+1} by step (a) (note that $x_0 \notin S_k(x_0)$).

(d) To clean-up, set $h(s+1) = \max\{|x_0| + k, |y_1|, \ldots, |y_k|\}$. Let $w_1, w_2, \ldots, w_r$ order all the words not yet in the domain of B_{s+1} of length at most $h(s+1)$. We put w_i into the domain of B_{s+1} to satisfy (7). To see that this is achievable, we use two simple observations: first, for all words $w \neq w'$, the intersection $S_k(w) \cap S_k(w')$ is empty. Second, setting $Y = \{x_0, y_1, \ldots, y_k\} \cup S_k(x_0)$, we see that for each word $w \neq x_0$ of length $\leq h(s+1)$, $|S_k(w) \cap Y| \leq k$. Note that Y includes all words of length $> h(s)$ that have been put in the domain of B_{s+1} by steps (a-c).

This completes the proof. □

4.7 The Polynomial Analogue of Post's Problem

Suppose that $P \neq NP$. Then we may ask if there exists a language in $NP - P$ that is not NP-complete under Cook reducibility. This question is analogous to a famous question posed by Post in 1944: Does there exist a recursively enumerable language that is not recursive and not RE-complete under arbitrary[9] Turing reducibility? An affirmative answer was independently found in 1956 by Friedberg [5] and Muchnik [14]. In this section we show Ladner's result [8] that the polynomial analogue of Post's problem also has an affirmative solution (provided, of course, that $P \neq NP$). We remark that there are very few natural problems that are conjectured to be in $NP - P$ and not NP-complete. Perhaps the most well-known candidate[10] is graph isomorphism.

THEOREM 19 (Ladner). *If $P \neq NP$ then there exists a language in $NP - P$ that is not NP-complete under Cook reducibility.*

This theorem is a consequence of the next result.

THEOREM 20. *Let K a class of languages over the alphabet $\Sigma = \{0, 1\}$ and P denote as usual the class $DTIME(n^{O(1)})$. Suppose K has a total universal acceptor and is closed under finite variation. Let (Σ, L) be a recursive language not in $K \cup P$. Then there exists a language (Σ, L') such that*

(i) *L' is not in K,*

(ii) *L' is Karp-reducible to L,*

(iii) *L is not Cook-reducible to L'.*

[9] That is, Turing reducibility with no resource bounds on the query machines.

[10] Lesser-known candidates include Discrete Logarithm and Minimum Weight Triangulation. Famous *former* candidates include Linear Programming and Primality. In each case, the candidate has been to be in P in ground breaking work. Khacian did this for Linear Programming in 1979, and XXX for Primality in 2002.

To see that theorem 20 implies Ladner's theorem, let $K = P|\Sigma$ and $L \in (NP|\Sigma) \setminus P$. Since K has a total universal acceptor, it follows from theorem 20 that there exists an L' such that L' is not in P (by (i)), L' is in NP (by (ii)) and L' is not NP-complete under Cook reducibility (by (iii)).

Proof of theorem 20. Let

$$U = \{U_0, U_1, U_2, \ldots\}.$$

and

$$Q = \{Q_0^{(\cdot)}, Q_1^{(\cdot)}, Q_2^{(\cdot)}, \ldots\}.$$

denote (respectively) a universal acceptor for K and a universal query machine for the family of polynomial-time oracle operators whose input and output alphabet is Σ. (Note: we do not require U or Q to be efficient.)

The requirement that L' is not in K (respectively, L is not Cook-reducible to L') is equivalent to the set of even (resp. odd) numbered conditions in the following infinite list of conditions C_i $(i = 0, 1, \ldots)$ where

$$\begin{aligned} C_{2j} : &\quad L' \neq L(U_j) \\ C_{2j+1} : &\quad L \neq L(Q_j^{(L')}) \end{aligned}$$

Let

$$x_0, x_1, x_2, \ldots \tag{9}$$

be the usual lexicographic enumeration of words in Σ^*. We shall construct a transducer T that runs in linear time such that $\mathrm{T} : \Sigma^* \to \{0\}^*$. Then define

$$L' := \{x \,:\, x \in L \text{ and } |\mathrm{T}(x)| \text{ is even}\}. \tag{10}$$

This clearly shows that L' is Karp-reducible to L.

To prove the theorem it remains to describe T such that the resulting L' satisfies the list of conditions above. We first explain in an intuitive way the operation of T on input x. We say that a word x is *in stage* i if $|\mathrm{T}(x)| = i$. T has the property that $|T(x_{i+1})|$ is equal to either $|T(x_i)|$ or $|T(x_i)|+1$. Hence the words in the ordering (9) have non-decreasing stage numbers. We maintain the following invariant:

> (*) If x is in stage i, then condition C_i is not yet satisfied but each condition C_j $(j < i)$ is witnessed by some word y_j that precedes x in the lexicographical order (9).

The first thing that T does on input x is to try to determine the stage of x. It can only do this approximately in the sense that it can determine an integer i such that x is either in stage i or stage $i+1$. This uncertainty turns out not to be crucial because T will try to satisfy condition C_i and if it succeeds, then $|\mathrm{T}(x)|$ is equal to $i+1$ otherwise i. Given i, consider how T can try to satisfy C_i. There are 2 cases.

(i) Suppose that $i = 2j$. Then T systematically attempts to find a word w such that $w \in L'$ iff $w \notin U_j$. Such a w is a 'witness' for the condition C_{2j}. (Note the self-reference here: we are constructing L' and we want to check if w belongs to L'.) If T is successful in finding a witness, T outputs the word 0^{i+1}; otherwise it outputs the word 0^i. The output 0^{i+1} serves a double purpose: to inform subsequent computations that condition C_i is satisfied and to ensure that x is not in the set L' (recall the definition (10) of L'). Likewise for an output of 0^i.

(ii) Suppose $i = 2j+1$. T systematically attempts to find a witness for condition C_{2j+1}, i.e., a word w such that $w \in L$ iff $w \notin Q_j^{(L')}$. (Again note the self-reference.) If successful in finding the witness w, x is put in stage $i+1$ by the output $T(x) = 0^{i+1}$; else it stays in stage i with the output $T(x) = 0^i$.

Now we fill in the details. Let $|x| = n$. First, to determine the approximate stage of x as follows. T allocates a total of n steps to compute (by simulating itself) the values $\mathrm{T}(0^1), \mathrm{T}(0^2), \ldots, \mathrm{T}(0^m)$ in succession, where m is such that the allocated n steps are used up while in the middle of computing $\mathrm{T}(0^{m+1})$. Then x is in the approximate stage i where $i = |\mathrm{T}(0^m)|$. We next attempt to find a witness for the ith condition C_i by testing successive words in the ordering (9) until a total of n steps are used up. If a witness is found within n steps then output $T(x) = 0^{i+1}$; otherwise output $T(x) = 0^i$. Consider how we test for the witness property:

(I) Suppose $i = 2j$. To test whether any word w is a witness, T checks (a) whether $w \in U_j$ (by simulating the universal acceptor U) and (b) whether $w \in L'$ (by checking if $|\mathrm{T}(w)|$ is even and if $w \in L$ – the latter is possible since L is recursive). Note that w is a witness if the outcomes for (a) and (b) differ.

(II) Suppose $i = 2j+1$. To test whether w is a witness, we check (c) whether $w \in L$ as before and (d) whether w is accepted by $Q_j^{(\cdot)}$ using oracle L'. To do (d), we simulate the universal query machine Q using input $\langle j, w\rangle$, and whenever a query "Is z in L'?" is made, we can check this as in (b) above.

Correctness: By construction T runs in linear time. It is also easy to see that our invariant (*) holds. In particular, if T goes through stage i for all i then each condition would be satisfied and we are done. Therefore we only need to ensure that T does not remain in some stage i forever. For the sake of contradiction, we assume that for some i and some n_0, we have $|\mathrm{T}(x)| = i$ for all $|x| \geq n_0$, In other words, suppose T never leaves stage i.

(I) Suppose $i = 2j$. Since we never leave stage $2j$, by the definition of L', L' is equal to a finite variant of L. Now $L \notin K$ and K is closed under finite variation imply that L (and hence L') must differ from U_j for infinitely many inputs. Choose the first (with respect to the lexicographical ordering (9)) witness w_0 to the fact that $U_j \neq L'$. Now choose n large enough: precisely, $n > n_0$ and n exceeds the number of steps needed by T to compute $\mathrm{T}(0^1), \ldots, \mathrm{T}(0^m)$, where $|\mathrm{T}(0^m)| = 2j$, and exceeds the number of steps to test if x is a witness to condition C_{2j} for all x preceding w_0 in the ordering (9). Observe that the smallest such number n is well-defined. Then clearly $|\mathrm{T}(0^n)| = 2j+1$, a contradiction.

(II) Suppose $i = 2j+1$. Then by definition of L', L' is a finite set. Now observe that L must differ from $Q_j^{(L')}$ on infinitely many words, for otherwise $L \in P$ which would be contrary to our assumption. Let w_0 to be the first (with respect to the ordering (9)) witness word to the fact that $L \neq Q_j^{(L')}$. Again choose $n > n_0$ such that it exceeds the time to determine that 0^n is in stage $2j+1$ and exceeds the time to test if x is a witness to the condition C_{2j+1} for all x preceding w_0. Then it is easy to see that $|\mathrm{T}(0^n)| = i+1$.

This concludes our proof

An interesting observation from the proof is that L' is the intersection of L with the polynomial time computable set $G = \{x : |\mathrm{T}(x)| \text{ is even}\}$.

COROLLARY 21. *If NP $\neq$ P then for every NP-complete language L (under $\leq_T^P$) there exists a language $G \in P$ such that $G \cap L$ is in NP $-$ P but is not NP-complete.*

There is another way to view Ladner's proof: G is nothing but a 'gappy' set determined by some recursive function $r(n)$ such that $x \in G$ iff $r(2i) \leq |x| < r(2i+1)$. Here $r(n)$ is the number of steps needed to find witnesses for all the conditions up to C_n. The size of such gaps is studied in Lipton, Landweber and Robertson [11] and Chew and Machtey [4]. Schöning [17], Schmidt [16] and Regan [15] have greatly simplified and generalized these techniques. An interesting result [6] is the following: Assuming that $P \neq NP$, there is a language in $NP - P$ that is not P-hard under one-way log-space many-one reducibility $\leq_m^{1L}$.

4.8 The Structure of Polynomial Degrees

The reducibilities $\leq$ in this section are assumed to be transitive, so that the term $\leq$-degrees is meaningful. For any two languages (Σ_0, L_0) and (Σ_1, L_1), define their *join* $L_0 \oplus L_1$ to be the language

$$\{0x : x \in L_0\} \cup \{1x : x \in L_1\}.$$

over the alphabet $\Sigma_0 \cup \Sigma_1 \cup \{0,1\}$. Note that L_1 and L_2 are each $\leq_m^P$-reducible to $L_0 \oplus L_1$. Furthermore, for any L, if $L_i \leq_m^P L$ holds for $i = 0, 1$ then it is easy to see that $L_0 \oplus L_1 \leq_m^P L$. In other words, the $\leq_m^P$-degree of $L_0 \oplus L_1$ is the least upper bound of the $\leq_m^P$-degrees of L_0 and of L_1. The same holds for Cook-reducibility. In the terminology of lattice theory[11] we have just shown:

THEOREM 22. *The $\leq_m^P$-degrees form an upper semi-lattice. The same holds for $\leq_T^P$-degrees.*

Definition 14. Let K be a class of languages, and $\leq$ a transitive reducibility. We say that the $\leq$-degrees of K are *dense* if the following holds. If $L_0 < L_2$ $(L_0, L_2 \in K)$ then there exists $L_1 \in K$ such that

$$L_0 < L_1 < L_2.$$

■

[11] A partial order $(X, \leq)$ is an upper semi-lattice if for all $x, y \in X$, there is a least upper bound denoted $x \vee y$ (the least upper bound is unique if it exists).

We now address questions about the density of the above semi-lattices. The theorem of Ladner shows that if $NP \neq P$ then the Karp-degrees in $NP - P$ is dense. Similarly if $P \neq NP$ then the Cook-degrees between P and NP are dense. More generally, we have:

THEOREM 23. *If $L <_T^P L'$ and L and L' are recursive then there exist a recursive L'' such that $L <_T^P L'' <_T^P L'$.*

There have been considerable developments in such questions (e.g., see [2]).

4.9 Nondeterministic reducibilities

In this section, we consider reducibilities defined by nondeterministic machines (acceptors, transducers, query machines, tt-machines). Of these, nondeterministic tt-machines have not been defined. But first, we extend the notion of many-one reducibility to multivalued transformations.

Definition 15. Let $(\Sigma, A), (\Gamma, B)$ be languages and f be the multivalued transformation from Σ to Γ computed by some transducer T. We say that A is *many-one (nondeterministically) reducible* to B via f (or via T) if for all $x \in \Sigma^*$,

$$x \in A \iff f(x) \cap B \neq \emptyset.$$

■

A nondeterministic tt-machine (G, E) consists of a nondeterministic transducer G and a nondeterministic evaluator E. It turns out that, at least for polynomial-time bounds, we can restrict E to be deterministic (see Exercises). The corresponding pair (g, e) computed by a nondeterministic (G, E) is called a *multivalued truth-table*.

We say A is *nondeterministically tt-reducible* to B via (g, e) (or via (G, E)) if for all $x \in \Sigma^*$, there exists $y_1 \# \cdots \# y_k \in g(x)$ such that

$$x \in A \iff x \# \chi_B(y_1) \cdots \chi_B(y_k) \in e.$$

Recall that $\chi_B(x)$ is the indicator function for the set B. For nondeterministic polynomial-time bounded reducibilities, we use the superscript 'NP' in the usual way. For instance, $\leq_T^{NP}$, $\leq_{tt}^{NP}$ and $\leq_m^{NP}$ are the nondeterministic counterparts of $\leq_T^P$, $\leq_{tt}^P$ and $\leq_m^P$. These new reducibilities have rather different properties from their deterministic counterparts. The next result is by Ladner, Lynch and Selman.

THEOREM 24. *The following pairs of reducibilities form pairs of identical binary relations over the non-trivial languages:*

(i) $\leq_T^{NP}$ *and* $\leq_{tt}^{NP}$

(ii) $\leq_{ptt}^{NP}$ *and* $\leq_{ctt}^{NP}$

(iii) $\leq_{dtt}^{NP}$ *and* $\leq_m^{NP}$.

Proof. In each case, the first member of a pair is as strong as the second. Hence it is sufficient to show that the second is as strong as the first.

(i) Suppose $A \leq_T^{NP} B$ via a nondeterministic polynomial-time oracle machine M and we want to show $A \leq_{tt}^{NP} B$. Choose any $x_0 \notin B$ and $x_1 \in B$ (this is possible since we assume non-trivial languages). We will construct a nondeterministic tt-machine (G, E) such that $A \leq_{tt}^{NP} B$ via (G, E). On any input x, the query generator G simulates M on input x. G also maintains on a special work-tape T a list of values, initially empty. If the ith query ($i \geq 0$) of M to the oracle is the word y_i then G guesses the oracle's answer. If the guessed answer is yes, then G records in the tape T the pair (y_i, x_1); otherwise it records in T the pair (y_i, x_0). If M accepts, then G outputs the contents of tape T:

$$y_1 \# z_1 \# y_2 \# z_2 \# \cdots \# y_k \# z_k$$

where $z_i \in \{x_0, x_1\}$. If M rejects, G outputs $x_0 \# x_1$. We now have to describe the acceptor E. On input $x \# b_1 c_1 \cdots b_k c_k$, ($b_i, c_i \in \{0, 1\}$), E accepts iff $b_i = c_i$ for all $i = 1, \ldots, k$. It is easy to see that $A \leq_{tt} B$ via (G, E).

(ii) Let $A \leq^{NP}_{ptt} B$ via (G, E). We will construct a nondeterministic tt-machine (G′, E′). G′, on input x, simulates G on x. If G outputs $y_1\#\cdots\#y_k$, then G′ guesses a subset $I \subseteq \{1, 2, \ldots, k\}$ and checks if $x\#b_1\cdots b_k$ is accepted by E, where $b_i = 1$ iff $i \in I$. If it is accepted then G′ outputs $y_{i_1}\#y_{i_2}\#\cdots\#y_{i_m}$ where $I = \{i_1, i_2, \ldots, i_m\}$; otherwise G′ outputs x_0 where x_0 is some fixed word not in B. The evaluator E′ is essentially determined since we use conjunctive reducibility. The reader can easily verify that $A \leq^{NP}_{ctt} B$ via (G′, E′).

(iii) This follows the proof for (ii) closely.

This concludes the proof. **Q.E.D.**

We next show that some of the nondeterministic reducibilities are actually intransitive:

THEOREM 25. (i) *$\leq^{NP}_{m}$ and $\leq^{NP}_{ctt}$ are transitive.*

(ii) *$\leq^{NP}_{tt}$, $\leq^{NP}_{*\text{-}tt}$ and $\leq^{NP}_{k\text{-}tt}$ $(k \geq 1)$ are intransitive.*

Proof. The proof of (i) is straightforward, and is omitted. To show (ii), it is sufficient to construct non-trivial languages A, B, C over $\{0, 1\}$ satisfying

(a) $A \leq^{NP}_{1\text{-}tt} B \leq^{NP}_{1\text{-}tt} C$

(b) $A \not\leq^{NP}_{tt} C$.

Note that (a) follows from the following two conditions:

$$x \in A \iff (\exists y)[\ |y| = |x| \text{ and } y \notin B\ ] \tag{11}$$

$$x \in B \iff (\exists y)[\ |y| = |x| \text{ and } y \in C\ ] \tag{12}$$

Let $\mathrm{Q} = \{\mathrm{Q}_i\}$ be an efficient universal oracle machine for the characteristic family of nondeterministic polynomial-time oracle machines (it is easy to see that Q exists). By definition each Q_i runs in polynomial time; we may further assume that each Q_i uses less than 2^n steps if $n \geq i$. Since $\leq^{NP}_{T}$ and $\leq^{NP}_{tt}$ are identical relations, (b) corresponds to an infinite list of conditions where the ith condition states:

$$A \text{ is not Turing-reducible to } C \text{ via } \mathrm{Q}_i.$$

We proceed in stages where we satisfy the sth condition in stage s, $s \geq 0$. Let that the partial sets $A_{s-1}, B_{s-1}, C_{s-1}$ have been defined at the end of the $s-1$st stage. Also, inductively assume that each of the sets $A_{s-1}, B_{s-1}, C_{s-1}$ has a common domain. Moreover, for each n, all words of length n are in the common domain or all or not.

To initialize, each of the sets A_0, B_0, C_0 has domain consisting of just the empty string ϵ where $\epsilon \in A_0, \epsilon \notin B_0 \cup C_0$. In stage s pick any word x of length n where n is the smallest integer such that no word of length n is in the current common domain. We will make x a witness to condition s by considering two cases:

Case (1): x is accepted by Q_s using oracle C_s (where queries on words not in the domain of C_s are given NO answers). Let Y be the set of words queried in some arbitrary accepting computation path of $\mathrm{Q}_s^{(C_s)}$ on x. Then we extend the domains of A_s, B_s and C_s by omitting all words of length $|x|$ from A_{s+1}, including all words of length $|x|$ into B_{s+1} and including exactly one string $y \notin Y$ of length $|x|$ into C_{s+1}. The existence of y follows from the fact that $|Y| < 2^{|x|}$ since we assume that Q_s runs in time $< 2^n$ for $n \geq s$. Observe that condition s as well as (11) and (12) are satisfied.

Case (2): x is not accepted by $\mathrm{Q}_s^{(C_s)}$. Extend the respective domains by including all words of length $|x|$ into A_{s+1} but omitting them all from B_{s+1} and from C_{s+1}. Again, condition s as well as properties (11) and (12) are satisfied.

We now clean-up for stage s. For each m such that some word of length m is queried in some computation path of $\mathrm{Q}_s^{(C_s)}$ on input x, if such words are not already in the common domain of $A_{s+1}, B_{s+1}, C_{s+1}$, we put every word w of length m domains of $A_{s+1}, B_{s+1}, C_{s+1}$ so as to satisfy properties (11) and (12): more precisely, we exclude w from C_{s+1} and from B_{s+1} but include w in A_{s+1}. Note that omitting w from C_{s+1} is consistent with the replies that $\mathrm{Q}_s^{(C_s)}$ received on input x. This completes our description of stage s. **Q.E.D.**

Before concluding this section, we describe a weaker notion of nondeterministic reducibility investigated by Long [12]:[12]

[12]Long qualifies these reducibilities as 'strong' because his notion of strength of reducibilities is opposite to ours.

Definition 16. Let $(\Sigma, A), (\Gamma, B)$ be languages and f a multivalued function from Σ^* to Γ^*. We say A is *unequivocal many-one reducible to B via f* if for all $x \in \Sigma^*$, we have

$$\begin{aligned} x \in A &\Rightarrow f(x) \subseteq B, \\ x \notin A &\Rightarrow f(x) \cap B = \emptyset. \end{aligned}$$

We write $A \leq_{um} B$ *via* f in this case. ■

Recall (chapter 2, section 9) that a nondeterministic acceptor is unequivocal if for all inputs, there is at least one terminating computation path, and terminating computation paths are all accepting or all rejecting. We say A is *unequivocally Turing-reducible to B via* a query machine M if A is Turing-reducible to B via M and $M^{(B)}$ is unequivocal. We then write $A \leq_{uT} B$ *via* M.

Remark: All reducibilities seen until now in this book could equally have been made in terms of the 'abstract' concept of transformations, oracle operators or truth-tables as well as in terms of their 'concrete' counterparts of transducers, oracle machines or tt-machines. In contrast, this last reducibility apparently cannot be defined without explicit reference to the concrete concept of machines. There will be other examples.

In the usual way, we append superscripts P , NP, *etc* to the various reducibilities to indicate complexity bounds. Unequivocal many-one reducibility has proven to be useful in classifying the complexity of number-theoretic problems. The polynomial-time versions of these are called *gamma-reducibilities* (γ-reducibilities) by Adleman and Manders who introduced them. Gamma reducibilities are valuable for studying closure under complements because of the following result: there are NP-complete languages under γ-reducibilities and furthermore if L is NP-complete under γ-reducibilities then $NP =$ co-NP iff $L \in$ co-NP. In general, unequivocal reducibilities seem to have nice properties, as seen in the following result by Long.

LEMMA 26. *Unequivocal many-one reducibilities and unequivocal Turing-reducibilities with polynomial-time bounds are transitive reducibilities.*

4.10 Conclusion

In this chapter we introduced the important reducibilities and their basic properties. Many other forms of reducibility have been studied but here we are contented to give a few examples.

(1) Lynch has investigated the use of multiple oracles (by endowing machines with more than one oracle tape).

(2) *Reductions based on logical constructs.* For any complexity function f, Jones defines the *f-bounded rudimentary reducibility* based on certain simple logical predicates and quantifications bounded by f. For instance, the $\log n$-bounded rudimentary reducibility is weaker than $\leq_m^L$.

(3) An interesting concept due to J. Simon is similar to many-one reductions except for an extra 'parsimony' constraint. Roughly speaking, parsimony insists that the reducibility preserves the number of solutions. In illustration of this idea, consider the reduction of Hamiltonian circuit problem to SAT. If we have a transformation t that takes any graph G to a corresponding CNF formula F such that the number of Hamiltonian circuits in G is equal to the number of satisfying assignments to F, then t is parsimonious. To formalize this notion, we must reformulate[13] a 'problem' to be a binary relation over Σ^* for some alphabet Σ: this would allow us to count the number of solutions to a problem, and thus define parsimonious reductions among such problems. Once this is done, it can be shown that all known transformations of NP-complete problems can be modified to be parsimonious.

[13]Naturally this takes us outside the 'standard' complexity theory in the sense of departing from some basic decisions made in chapter one.

Appendix A

Useful number-theoretic functions

Pairing functions. A *pairing function* is any bijection between $\mathbb{N}\times\mathbb{N}$ and $\mathbb{N} = \{0, 1, 2, \ldots\}$. We define a particular pairing function $\langle\cdot,\cdot\rangle : \mathbb{N}\times\mathbb{N} \to \mathbb{N}$ in which the correspondence between $(i, j) \in \mathbb{N}\times\mathbb{N}$ and $n \in \mathbb{N}$ is given by the relation

$$n = i + \binom{i+j+1}{2}$$

where $\binom{x}{2} = x(x-1)/2$. Unless stated otherwise, by 'the pairing function' we refer to this definition. Using $\langle\cdot,\cdot\rangle$ to denote a pairing function is a deliberate confusion of notation since in general we use the same notation for ordered pairs. This confusion is usually without harm. Whenever a distinction is necessary, we will write $\langle i, j\rangle \simeq n$ to say that the value of the pairing function on the values i, j is n. The best way to see that this is a bijection is through figure A.1.

The bijection induces a total ordering $\preceq$ of $\mathbb{N}\times\mathbb{N}$ defined as follows: the ordered pairs $(i', j'), (i, j)$ are related as $(i', j') \preceq (i, j)$ iff $n' \leq n$ where $\langle i, j\rangle \simeq n$ and $\langle i', j'\rangle \simeq n'$. If, in fact $n < n'$ holds, then we write $\langle i', j'\rangle \prec \langle i, j\rangle$. The (first and second) *projection functions* π_1 and π_2 are defined by

$$n \simeq \langle \pi_1(n), \pi_2(n)\rangle$$

for all n. A simple property of the pairing function is that it is increasing in each of its arguments.

LEMMA 27. *The function $\langle\cdot,\cdot\rangle$, π_1 and π_2 are computable in polynomial time.*

Proof. The fact that the pairing function is polynomial-time computable follows from the fact that multiplication is polynomial-time computable. To see that π_1 is polynomial time, note that on input n, in time polynomial in the size of the input, we can determine the k such that $k^2 \leq 2n \leq (k+1)^2$. (Use a binary search for k in the range between 1 and n.) We want the m such that $m(m-1) \leq 2n \leq m(m+1)$. Clearly $k = m$ or $m-1$, and we can easily determine which is the case. Finally, $\pi_1(n) = n - \binom{m+1}{2}$. It is similarly easy to compute $\pi_2(n)$. **Q.E.D.**

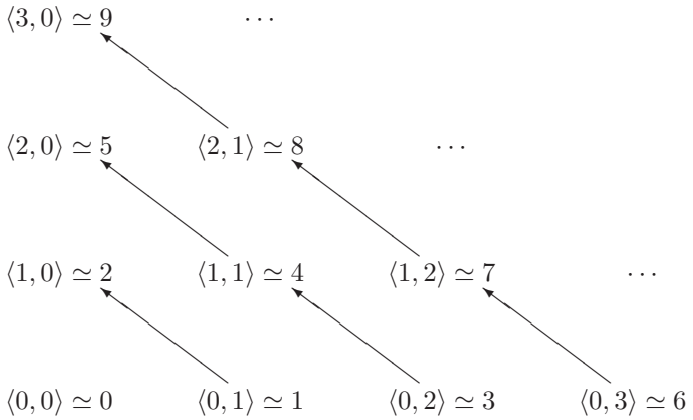

Figure A.1: The pairing function $\langle\cdot,\cdot\rangle$

Suppose $k \geq 2$ is fixed. We can use the pairing function to encode k-tuples as follows: a k-tuple $(x_1, \ldots, x_k)$, where $x_i \in \{0,1\}^*$, is encoded by

$$\langle x_1, \langle x_2, \cdots, \langle x_{k-2}, \langle x_{k-1}, x_k \rangle \rangle \cdots \rangle \rangle.$$

The projection functions $\pi_1, \ldots, \pi_k$ are easily defined. This mapping from $\mathbb{N}^k$ to $\mathbb{N}$ is a bijection. If we are willing (this is a matter of taste) to introduce extra symbols for encoding k-tuples, then it is easy to get an alternative encoding: encode $(x_1, \ldots, x_k)$ as $x_1 \# x_2 \# \cdots \# x_k$ where $\#$ is a new symbol. This poorman's encoding has the obvious simple linear-time decoding. However, it is not a bijection. But in all our applications, bijections were nice but unessential.

Finally, we could use the pairing function to encode finite sequences $\langle x_1, \ldots, x_k \rangle$ (of arbitrary length k) by explicitly encoding the length of the sequence:

$$\langle k, \langle x_1, \langle x_2, \cdots, \langle x_{k-2}, \langle x_{k-1}, x_k \rangle \rangle \cdots \rangle \rangle \rangle.$$

Note that this encoding is no longer a bijection between $\mathbb{N}^*$ and $\mathbb{N}$.

A fast growing function and its inverse.

Define the *super-exponential function* $\mathrm{Exp} : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ by $\mathrm{Exp}(0) = 0$ and $\mathrm{Exp}(n+1) = 2^{\mathrm{Exp}(n)}$. So $\mathrm{Exp}(n)$ consists of a stack of n two's.[1] In the remainder of this appendix, Exp will refer to the one argument function.

The *log-star function* $\log^* n$ is defined to be the largest integer m such that $\mathrm{Exp}(m) \leq n$. This function is one of the possible the 'inverses' of $\mathrm{Exp}(n)$.

Lemma 28.

(i) Exp *can be computed in time linear in its output size* $|\mathrm{Exp}(n)|$

(ii) $\log^* n$ *can be computed in time linear in its input size* $|n|$.

Proof. Note that $|n|$ here refers to the length of the binary representation $bin(n)$ of n, not the absolute value of n. Note that for $n > 0$, $|\mathrm{Exp}(n)|$ is $1 + \mathrm{Exp}(n-1)$ since $bin(\mathrm{Exp}(n))$ is '1' followed by $\mathrm{Exp}(n-1)$ 0's. To show (i), suppose $bin(\mathrm{Exp}(n-1))$ is available, and we want to compute $bin(\mathrm{Exp}(n))$. This amounts to writing out exactly $\mathrm{Exp}(n-1)$ zeroes. To do this, we treat $bin(\mathrm{Exp}(n-1))$ as a binary counter and we want to decrement it to zero. If $m > 0$ is the value of the binary counter and $g(m)$ denote the maximal string of zeroes that form a suffix $bin(m)$ (i.e. the suffix of $bin(m)$ has the form '$\cdots 10^{g(m)}$'), then a single decrement of the counter takes $O(g(m))$ steps. Hence the time to decrement the counter from $\mathrm{Exp}(n-1)$ to zero is order of

$$\sum_{m=1}^{\mathrm{Exp}(n-1)} g(m) = \sum_{k=1}^{\mathrm{Exp}(n-2)} k \cdot 2^{\mathrm{Exp}(n-2)-k} = O(2^{\mathrm{Exp}(n-2)}) = O(\mathrm{Exp}(n-1)).$$

Thus computing $\mathrm{Exp}(n)$ from $\mathrm{Exp}(n-1)$ is linear in $|\mathrm{Exp}(n)|$. Summing this work for $m = 1, 2, \ldots, n$ we get $\sum_{m=1}^{n} |\mathrm{Exp}(m)|$ which is $O(|\mathrm{Exp}(n)|)$, as required. To show (ii), we compute the increasing values

$$\mathrm{Exp}(0), \mathrm{Exp}(1), \mathrm{Exp}(2), .$$

until for some $m \geq 0$, $\mathrm{Exp}(m+1)$ has length greater than $|n|$. Then $\log^* n = m$, and we can convert to binary if desired. The linearity of this procedure is similarly shown as in (i). **Q.E.D.**

[1] $\mathrm{Exp}(n)$ is to be distinguished from the usual exponential function $\exp(n) := e^n$, where e is the base of the natural logarithm.

Exercises

[4.1] (i) What are the REC-complete languages under $\leq_m^{REC}$?
(ii) Prove that HALT is RE-complete under $\leq_m^{REC}$.

[4.2] Suppose there is a language L_0 that is complete for NP under Cook reducibility but not complete for NP under Karp reducibility. What consequences can you draw from the existence of L_0?

[4.3] Consider the following definition[2] of NP-completeness: a language L is NP-complete if (a) L is in NP and (b) if L is in P then $NP = P$. Clearly SAT is NP-complete under this definition. This definition appears attractive because it is simpler than any of the usual definitions. What is unusual about it? Is the language SAT NP-complete according to this definition? Is there a polynomial analogue of Post's problem here?

[4.4] (Aho-Hopcroft-Ullman) Consider the following informal notion of K-hardness: say L is *effectively K-hard* if there is an effective (i.e. recursive) procedure F such that for each $L' \in K$, $F(L')$ is a transformation such that L' is many-one reducible to L via $F(L')$. Make this into a precise definition. Verify that SAT is effectively NP-hard under your definition. Most notions of NP-hardness are based on a concept of reducibility. Is there a corresponding reducibility here?

[4.5] (i) By modifying the proof in chapter 3, show that SAT is NP-complete under many-one log-space reducibility.
(ii) In fact, show that all the NP-complete problems in chapter 3 are still NP-complete under the one-way log-space $\leq_m^{1L}$ reducibility of Hartmanis.

[4.6] Show that $LBA = NSPACE(n)$ is precisely the class of context sensitive languages.

[4.7] Prove that $DLOG$ is not closed under $\leq_m^P$-reducibility if $DLOG \neq P$.

[4.8] For any K and any family of transformations Ω, show that L is K-complete iff co-L is (co-K)-complete under $\leq_m^\Omega$-reducibility. Conclude that if L is P -complete under Karp reducibility then co-L is also P -complete.

[4.9] (Meyer) Show that NP = co-NP iff there exists an NP-complete language L in co-NP. The notion of NP-completeness depends on the choice of a reducibility $\leq$. We want you to specify the needed property of $\leq$ in order for this result to hold.

[4.10] (Jones, Lien, Laaser) Show that t is a deterministic log-space transformations iff the following language is in $DLOG$:

$$\{x\#i\#b \,:\, x, i \in \{0,1\}^* \text{ and } b \text{ is the } i\text{th symbol of } t(x)\}$$

for some transformation $|t(x)| = |x|^{O(1)}$.

[4.11] Show that the familiar functions $x+y$, $x-y$, $\max(x,y)$ and $x \cdot y$ are in **DLOG** (the family of deterministic log-space transformations). It is not known if $\lfloor x/y \rfloor$ is in **DLOG**; but obtain a small space complexity s such that division can be computed by a transducer using space s.

[4.12] What is the flaw in the following proof that NP is closed under Cook reducibility: let $L \in NP$ and L' is accepted by some deterministic polynomial time oracle machine $\mathrm{M}^{(\cdot)}$ with oracle L. To show that L' is in NP we define a machine M' which on input x simulates $\mathrm{M}^{(L)}$ until M enters the QUERY state. Then M' temporarily suspends the simulation of M and begins simulating a nondeterministic polynomial time acceptor N of L using the query word. If N accepts then M' continues the simulation of M from the YES state, otherwise it continues from the NO state. Clearly M' computes in nondeterministic polynomial time.

[4.13] Let $|\Sigma| \geq 2$. Show that $P = NP$ iff $P\,|\Sigma = NP|\Sigma$

[4.14] Say the language (Σ, L) is a *coded image* of (Γ, L') if there is a homomorphism $h : \Gamma \to \Sigma^*$ such that $h(a)$ has the same length for all a in Γ and the natural extension of h, $h^* : \Gamma^* \to \Sigma^*$, maps L' isomorphically into L. A class K is *coded in* K' if every language in K has a coded image in K'. K and K' are *codably equivalent* if they are each coded in the other. Show that P and $P\,|\Sigma$ are codably equivalent iff $|\Sigma| \geq 2$. Show the same for the other standard classes. Conclude that with respect to the $P = NP$, $DLOG = NLOG$ and $P = NLOG$ questions, our theory might as well be restricted to languages over $\{0,1\}$.

[2] we call this the 'student definition' because it often appears in oral or written exams when students are asked for a definition. Of course, teachers do it too. See discussion in [20].

[4.15] (infinitely often reducibility) Let t be a transformation $t : \Sigma^* \to \Gamma^*$. A language (Σ, L) is *infinitely often (i.o.) many-one reducible* to (Γ, L') if for infinitely many $x \in \Sigma^*$, we have $x \in L$ iff $t(x) \in L'$. We extend this definition to incorporate complexity in the obvious way: e.g. 'i.o. Karp reducibility' refers to i.o. many-one reducibility in polynomial time, denoted $\leq^P_{i.o.m}$. Investigate the properties of such efficient reducibilities.

[4.16] In this exercise we assume the two-way oracle model.
(i) Show that the relationship $\leq^L_T$ is not reflexive (and hence it is not a reducibility). **Hint:** Consider any in $DLBA - DL$.
(ii) Show that the reducibilities $\leq^{DLBA}_m$ and $\leq^{DLBA}_T$ are incomparable (here the superscript $DLBA$ denotes linear space bounds.) **Hint:** To show that $\leq^{DLBA}_T$ is not stronger than $\leq^{DLBA}_m$ pick any language $L \notin DLBA$ and show that L is not $\leq^{DLBA}_T$-reducible to where $L' = \{ww : w \in L\}$.

[4.17] * For any language L, define the language $L' = \{\text{bin}(|x|) : x \in L\} \subseteq \{0,1\}^* \subseteq \{0,1\}^*$. Give an efficient reduction of L' to L. If L is accepted in time-space-reversal (t, s, r), what are the obvious complexity bounds on L'? Find non-trivial improvements on these bounds.

[4.18] Show that if f is space-constructible then $DSPACE(f)|\Sigma$ (for any alphabet Σ) is efficiently presentable.

[4.19] * (i) Show that the two-argument *universal-polynomial* complexity function $p(n, k) := n^k$ is constructible, for integer values of the inputs. **Hint:** First show that $p'(n, k) = \Theta(n^k)$ is constructible. Actually, this is sufficient for all purposes.
(ii) Fill in the details for the construction of an efficient universal acceptor for the characteristic class P .
(iii) More generally, show that if a family F of complexity functions has a time-constructible universal-function $F^*(i, n)$, then $NTIME(F)$ is efficiently presentable. That is, for each fixed i, $F^*(i, n) \in F$ and conversely each $f(n) \in F$ is equal of $F^*(i, n)$ for some n.

[4.20] Define a *super universal machine* to be a Turing acceptor U with three input tapes. If the inputs are $\langle i, j, x\rangle$ then if i and j are fixed, we regard the machine as an ordinary Turing acceptor on input x, and denote it by $U^{(i)}_j$. If i is fixed we regard the family

$$U^{(i)} = \{U^{(i)}_0, U^{(i)}_1, \ldots\}$$

as a universal machine. We say that U efficiently presents a characteristic class K if (a) for each i, $U^{(i)}$ is an efficient presentation of K, and (b) for each efficient universal acceptor V of K, there exists an i such that $U^{(i)}$ has the same enumeration of K as V: for each j sufficiently large, $L(U^{(i)}_j) = L(V_j)$. Does the characteristic class NP have a super universal acceptor?

[4.21] Prove that for any reasonable resource bound β, say $\beta = n^{O(1)}$-*time*, $\leq^\beta_m$ and $\leq^\beta_{1\text{-}ptt}$ are identical.

[4.22] (Ladner, Lynch and Selman) Prove that (i) $\leq^P_{1\text{-}tt}$ stratifies $\leq^P_m$, and (ii) $\leq^P_{(k+1)\text{-}tt}$ stratifies $\leq^P_{k\text{-}tt}$. (see section 6 for definition of stratifies)

[4.23] Show that $\leq^P_{k\text{-}tt}$ for $k > 1$ is an intransitive reducibility.

[4.24] (Ladner, Lynch and Selman) Define the reducibility $\leq_0$ thus: $L \leq_0 L'$ if there is a query machine and a polynomial-time transformation $t : \{0,1\} \to \{\#\{0,1\}^*\}^*$ such that $M^{(L')}$ accept L and for input x, M only asks questions from the list $t(x)$ (ie. $\#x_1\#x_2\#\cdots\#x_k$ where $x_i \in \{0,1\}$ is viewed as the list $(x_1, x_2, \ldots, x_k)$). Show that $\leq_0$ is the same as polynomial-time truth-table reducibility.

[4.25] Show that if $A \leq^P_T B$ then A is tt-reducible to B in deterministic time 2^{n^k} for some $k \geq 1$.

[4.26] (Ladner and Lynch) Let L, L' be languages.
(i) Show that $L \cup L' \leq^L_{tt} L \oplus L'$. In fact, the $\leq^{FST}_{2-tt}$-reducibility is sufficient.
(ii) Show the existence of languages L, L' such that $L \cup L' \not\leq^L_m L \oplus L'$.

[4.27] (i) Prove theorem 9 (which contains the relativized Savitch's theorem) under the two-way oracle model. Why does the proof break down with the one-way oracle model?
(ii) Prove theorem 9 again, but under the tame one-way oracle model.

[4.28] Complete the proof of lemma 15.

[4.29] (Ladner, Lynch and Selman) Show that $L \leq_m^{NP} L'$ iff there exists a polynomial p and a polynomial-time transformation t such that $x \in L$ iff

$$(\exists y)[|y| \leq p(|x|) \wedge t(x,y) \in L']$$

[4.30] Show that if A is *PSPACE*-complete (under $\leq_m^P$-reducibility) then $P^A = NP^A$.

[4.31] For any set $A \subseteq \Sigma^*$, define $len(A) := \{w \in \Sigma^* : (\exists x \in A)|x| = |w|\}$. Construct a language $A \subseteq \Sigma^*$ such that $len(A) \notin P^A$.

[4.32] Prove that if $PSPACE \neq P$ then there exists L, L' in *PSPACE* such that L is Cook-reducible, but not Karp-reducible to L'. (Of course, if we can prove this result without the hypothesis that $PSPACE \neq P$ then we would have proven that $PSPACE \neq P$.)

[4.33] (Selman) A language (Σ, L) is P -selective if there is a polynomial time transformation $f : \Sigma^* \times \Sigma^* \to \Sigma^*$ such that $f(x,y) \in \{x,y\}$, and furthermore, if either x or y is in L then $f(x,y) \in L$.
(i) Show that for all non-trivial languages L, L is in P iff co-$L \leq_m^P L$ and L is P -selective.
(ii) If L is P -selective and $L' \leq_m^P L$ then L' is P -selective.
(iii) If SAT is P -selective then $NP = P$.

[4.34] Show that the following function is a pairing function:

$$\langle x, y \rangle = a^2 - 2(a-b) + \delta(x < y)$$

where $a = \max(x,y)$, $b = \min(x,y)$ and $\delta(x < y) = 1$ if $x < y$ holds and 0 otherwise. Show that this, and its associated projection functions π_1, π_2 are polynomial-time computable.

[4.35] Imitate Ladner's proof technique to obtain a language in $NEXPT - DEXPT$ (assuming this is non-empty) that is not complete.

[4.36] (I. Simon) Show that for all $A \notin P$, $\leq_T^P$ and $\leq_{tt}^P$ are distinguishable over A.

[4.37] (I. Simon) Let $\leq_1$ be stronger than $\leq_2$. We say that $\leq_1$ and $\leq_2$ are *downward distinguishable* over a language L if there exists a language L' such that $L' \leq_1 L$ but not $L' \leq_2 L$. (Note that the definition of distinguishable in the text may be qualified as 'upward'.) Let $\leq_T^{PS}$ denote Turing reducibility where the complexity bound is deterministic polynomial space. Show there exists a language $L \notin P$ such that $\leq_T^{PS}$ and $\leq_m^P$ are not downward distinguishable over L.

[4.38] (Ladner and Lynch) Show that $\leq_T^P$ is stronger than $\leq_T^L$.

[4.39] (Ladner and Lynch, open) Is $\leq_m^P$ stronger than $\leq_m^L$? Is $\leq_{tt}^P$ stronger than $\leq_{tt}^L$?

[4.40] Consider the alternative definition of tt-reducibility described at the end of section 5. Compare the relative strengths of the usual tt-reducibility with this definition, respect to each of the three choices of representing Boolean functions: (a) 'truth-tables' in the original sense of the term, (b) Boolean formulas, and (c) Boolean circuits.
(i) Show that with respect to polynomial-time bounds, there is no difference.
(ii) Show a difference (assuming $DLOG \neq P$) with respect to log-space bounds. What is the significance of the fact that the circuit value problem (CVP) is P -complete under log-space reducibility? (Note: the CVP problem, described in chapter 5, is the problem of deciding if a distinguished output of a Boolean circuit under a given input is 0 or 1.)

[4.41] Show that with respect to polynomial time, we could restrict nondeterministic tt-machines (G, E) such that E is deterministic.

[4.42] Prove the last theorem of section 8 which generalizes the Ladner's answer to the polynomial analogue of Post's problem.

[4.43] (Lind, Meyer) Show that the log-space transformations **DLOG** is closed under explicit transformation (viz., substitution by constants, renaming and identification of variables) and two-sided recursion of concatenation. The latter is defined as follows: A transformation $f : (\Sigma^*)^n + 1 \to \Delta^*$ on $n+1$ variables is defined by *two-sided recursion of concatenation* from $g : (\Sigma^*)^n \to \Delta^*$ and $h_1, h_2 : (\Sigma^*)^{n+2} \to \Delta^*$ if

$$f(\bar{w}, \epsilon) = g(\bar{w})$$
$$f(\bar{w}, xa) = h_1(\bar{w}, x, a) \cdot f(\bar{w}, x) \cdot h_2(\bar{w}, x, a)$$

for all $\bar{w} \in (\Sigma^*)^n, w \in \Sigma^*$ and $a \in \Sigma$. Note that in this context, we can regard a multi-variable transformation such as f as an ordinary one variable transformation if we encode a sequence (w_1, w_n) of words as one word $w_1\# \cdots \# w_n$ where $\#$ is a new symbol.

[4.44] * Following the outline in the concluding section, formalize the notion of parsimony by defining problems as binary relations over Σ^*. How much of standard complexity theory as used in this book carry over?

[4.45] (J. Simon) Show that SAT can be many-one reduced to Hamiltonian circuits using a parsimonious polynomial-time transformation t, i.e., for any CNF $F \in$ SAT, the number of satisfying assignments to F is equal to the number of Hamiltonian circuits in $t(F)$.

[4.46] * Investigate the basic structure of the diagonalization proofs by formalizing the concept of 'finite, existential conditions'. Refer to the basic diagonalization outline of section 6.

Bibliography

[1] A. V. Aho and J. D. Ullman. A characterization of two-way deterministic classes of languages. *Journal of Computers and Systems Science*, 4(6):523–538, 1970.

[2] K. Ambos-Spies. Sublattices of the polynomial time degrees. *Information and Computation*, 65:63–84, 1985.

[3] T. Baker, J. Gill, and R. Solovay. Relativizations of the $P =? NP$ question. *SIAM J. Computing*, 4:431–442, 1975.

[4] P. Chew and M. Machtey. A note on structure and looking back applied to the relative complexity of computable functions. *Journal of Computers and Systems Science*, 22:53–59, 1981.

[5] R. M. Friedberg. Two recursively enumerable sets of incomparable degrees of unsolvability. *Proceed. Nat. Acad. of Sciences*, 43:236–238, 1957.

[6] J. N. Hartmanis, N. Immerman, and S. Mahaney. One-way log-tape reductions. *19th Symposium FOCS*, pages 65–71, 1978.

[7] N. D. Jones. Space-bounded reducibility among combinatorial problems. *Journal of Computers and Systems Science*, 11:68–85, 1975.

[8] R. E. Ladner. On the structure of polynomial time reducibility. *JACM*, 22:1:155–171, 1975.

[9] R. E. Ladner and N. A. Lynch. Relativization of questions about log space computability. *Math. Systems Theory*, 10:19–32, 1976.

[10] R. E. Ladner, N. A. Lynch, and A. L. Selman. A comparison of polynomial time reducibilities. *Theor. Comp. Sci.*, 1:103–123, 1975.

[11] L. H. Landweber, R. J. Lipton, and E. L. Robertson. On the structure of sets in *NP* and other complexity classes. *Theor. Computer Science*, 15:181–200, 1981.

[12] T. J. Long. Strong nondeterministic polynomial-time reducibilities. *Theor. Computer Science*, 21:1–25, 1982.

[13] A. R. Meyer and E. M. McCreight. Computationally complex and pseudo-random zero-one valued functions. In Z. Kohavi and A. Paz, editors, *Theory of machines and computations*, pages 19–42. Academic Press, 1971.

[14] A. A. Muchnik. On the unsolvability of the problem of reducibility in the theory of algorithms (in russian). *Doklady Akademii Nauk SSSR*, 108:194–197, 1956.

[15] K. W. Regan. The topology of provability in complexity theory. *Journal of Computers and Systems Science*, 38:384–432, 1988.

[16] D. Schmidt. The recursion-theoretic structure of complexity classes. *Theor. Computer Science*, 38:143–156, 1985.

[17] U. Schöning. A uniform approach to obtain diagonal sets in complexity classes. *Theor. Computer Science*, 18:95–103, 1982.

[18] I. Simon. On some subrecursive reducibilities. Technical Report Tech. Rep. STAN-CS-77-608, Computer Sci. Dept., Stanford Univ., April, 1977. (PhD Thesis).

[19] O. Watanabe. On one-one polynomial time equivalence relations. *Theor. Computer Science*, 38:157–165, 1985.

[20] C. K. Yap. Logical curiosities: on definitions of *NP*-completeness. manuscript, 1985.

Chapter 5

Complete Languages

April 13, 2009

5.1 Existence of complete languages

The class *NP* has a long list of complete problems derived from almost every area of the computational literature. The theoretical unity that this fact provides for such diverse problems is one of the striking accomplishments of the field. Naturally, one is led to ask if other complexity classes such as *NLOG* and *PSPACE* have complete languages. The framework of the Basic Inclusion Lemma also motivates our search for complete languages. The answer is yes for these and many other classes; moreover, there is a systematic way to show this (e.g., [14]). The method depends on the existence of efficient universal machines; universal machines in turn depend on the existence of suitable tape-reduction theorems. In this chapter, for simplicity, the reducibility is assumed to be $\leq_m^L$ unless otherwise specified.

Convention: Let i denote a natural number. In contexts where a string is expected, we use the same symbol 'i' to denote the binary string representing the natural number i. The notation '$|i|$' always denotes the length of the binary representation of i, never the absolute value of i as a number. Also, if $\#$ is a symbol, then $\#^i$ denotes a string of i copies of $\#$.

THEOREM 1. *Each complexity class in the canonical list, with the exception of PLOG, has complete languages under log-space many-one reducibility,* $\leq_m^L$.

Proof. Here is the proof for the class P : Fix $\Sigma = \{0, 1\}$. We have shown in chapter 4 (example 5) that the characteristic class $P|\Sigma$ has an efficient universal acceptor $U^P = \{U_i^P \ : \ i = 0, 1, \ldots\}$ where each U_i^P accepts in time $n^{|i|}$. Define over the alphabet $\{0, 1, \#\}$ the language

$$L^P = \{i\#^m x \ : \ i, x \in \Sigma^*, m = |i| \cdot |x|^{|i|}, x \in L(U_i^P)\}.$$

We claim that L^P is P -complete under $\leq_m^L$. First we show that L^P is in P . This is witnessed by the acceptor M that, on input w, first checks that w has the form $i\#^m x$ for some i, x, m. Next it checks that $m = |i| \cdot |x|^{|i|}$, and if so, simulates U_i^P on x. The simulation takes time $O(m) = O(|w|)$. To show that L^P is P -hard, consider any language $(\Gamma, L) \in P$. Let h be the homomorphism that encodes each symbol of Γ as a unique binary string in Σ^* of length $\lceil \log |\Gamma| \rceil$. So $h(L)$ is a language over Σ; furthermore, it is easy to see that

$$h(L) \in P \iff L \in P \, .$$

So let $h(L)$ be accepted by U_i^P for some i. Consider the transducer T that on input x outputs $i\#^m h(x)$ where $m = |i| \cdot |h(x)|^{|i|}$. It is not hard to make T use $\log(|x|^{|i|}) = O_i(\log |x|)$ space to produce this output. Now x is in L iff T(x) is in L^P. Thus $L \leq_m^L L^P$. This shows that L^P is P -hard.

Sketch of the other cases: The proof for the class *NLOG* uses the same technique as above: we use the existence of an efficient universal acceptor U^{NLOG} for $NLOG|\Sigma$ such that for each i, U_i^{NLOG} accepts in space $|i| \log n$. The complete language is

$$L^{NLOG} = \{i\#^m x \ : \ i, x \in \Sigma^*, m = |i| \cdot |x|^{|i|}, x \in L(U_i^{NLOG})\}.$$

The proof for the class *DEXPT* uses the existence of an efficient universal acceptor U^{DXT} such that for each i, U_i^{DXT} accepts in time $|i|^n$ time. The complete language is

$$L^{DXT} = \{i\#^m x : i, x \in \Sigma^*, m = |i| \cdot |x|, x \in L(U_i^{DXT})\}.$$

The proofs for the other classes are similar. **Q.E.D.**

We make two remarks. First, the method unfortunately does not extend to $PLOG$. No complete problems for this class (under the usual reducibilities) are known. Second, the result is trivial for $DLOG$ because unrestricted log-space many-one reducibility is too powerful: it is easy to see that $DLOG$ is $\leq_m^L$-reducible to any non-trivial language. Following Hartmanis, we consider the one-way log-space many-one reducibilities, $\leq_m^{1L}$ (see section 2, chapter 4); the reader may verify that the construction for L^{NLOG} in the above proof carries over to give us a $DLOG$-complete language L^{DLOG} under $\leq_m^{1L}$-reducibility.

The next question we ask is whether classes such as

$$XTIME(n^k), XSPACE(\log^k n), XSPACE(n^k)$$

$(X = D, N)$ for each $k \geq 1$ have complete languages. The answer is yes, using a simple variation of the above proof. In particular, the characteristic class $XSPACE(\log^k n)$ has an efficient universal acceptor U' and hence it is not surprising to find complete languages for it. More precisely, we may assume that the universal acceptor U_i' accepts in space at most $|i| \log^k n$ for each i. Then a complete language for $XSPACE(\log^k n)$ is given by

$$L' = \{i\#^m x \ : \ x \in L(U_i'), m = |x|^{|i|}\}.$$

We make the observation that the language L^P in the proof of theorem 1 is actually in $DTIME(O(n))$, so it is in fact $DTIME(O(n))$-complete. To see this as a more general phenomenon, we next state a definition and lemma.

Definition. Let (Σ, L) be a language, $\#$ a symbol not in Σ, and f a number-theoretic function. Then the *f-padded version* (*with padding symbol $\#$*) of L is the language $L' = \{x\#^{f(|x|)} : x \in L\}$ over $\Sigma \cup \{\#\}$.

LEMMA 2. *Let $k \geq 1$ be any fixed integer. Under the $\leq_m^L$-reducibility, we have:*

(i) *If a language L is $DSPACE(n^k)$-complete then L is $PSPACE$-complete.*

(ii) *If L is $PSPACE$-complete then the f-padded version of L is $DSPACE(n^k)$-complete, where $f(n) = n^h$ for some $h \geq 1$.*

Proof.

(i) It is sufficient to show that every language $L' \in PSPACE$ can be reduced to L. Assume L' is accepted by some U_i in space $f(n) = n^{|i|}$, where U is an efficient universal acceptor for the characteristic class $PSPACE$. To reduce L' to L, we construct the transducer that on input x outputs $t(x) = i\#x\#^{f(|x|)}$. Clearly the language $t(L') = \{t(x) : x \in L'\}$ is in $DSPACE(n) \subseteq DSPACE(n^k)$. Since L is $DSPACE(n^k)$-complete, we have $t(L')$ is many-one reducible to L via some log-space transformation r. Hence L' is many-one reducible to L via the functional composition $r \circ t$. Since the log-space transformations are closed under functional composition (chapter 4), we conclude that $r \circ t$ is log-space computable.

(ii) Since L is $PSPACE$-complete, let L be accepted by some M in space n^h for some $h \geq 1$. Let L' be the n^k-padded version of L for some $k \geq 1$. Clearly $L' \in DSPACE(n) \subseteq DSPACE(n^k)$. To show that L' is $DSPACE(n^k)$-hard, let $L'' \in DSPACE(n^k)$ be many-one reducible to L via some log-space transformation t. The reader may verify that L'' is many-one reducible to L' via the transformation r where $r(x) = t(x)\#^{|t(x)|^h}$. Clearly r is log-space computable.

Q.E.D.

This tells us that the classes LBA, $DSPACE(n^k)$ and $PSPACE$ have essentially the same complete languages under log-space many-one reducibility.

The analog of the above lemma for $DTIME(n^k)$ can be proved in exactly the same way. Unfortunately, $DTIME(n^k)$ is not closed under log-space many-one reducibility so that we cannot apply the Basic Inclusion Lemma (chapter 4). The next result (from [4]) remedies this situation by considering the family **1FST** of transformations computed by 1-way finite state transducers (chapter 4, section 2):

THEOREM 3. *Let $k \geq 1$, and $X = D$ or N. Let K be the class $XTIME(n^k)$ or $XSPACE(n^k)$.*

(i) *K is closed under $\leq_m^{1FST}$-reducibility.*

(ii) *There exists K-complete languages under $\leq_m^{1FST}$ reducibilities.*

Proof. (i) This is a simple exercise in machine construction. (ii) We just show the result for $K = XTIME(n^k)$. Consider the language L^k consisting of all words of the form

$$w = x_1\#y\#x_2\#\cdots\#y\#x_n \tag{1}$$

where $x_i, y \in \{0,1\}^*$, each x_i has the fixed length k and y is the encoding of a Turing acceptor M_y that is clocked to run for n^k steps (with the appropriate mode X) and M_y accepts $x_1x_2\cdots x_n$. We first show that L^k is hard for $XTIME(n^k)$. Let (Σ, L) be accepted by some Turing acceptor M_y in time n^k. Let $h : \Sigma \to \{0,1\}^*$ be any encoding of the symbols of Σ using binary strings with a fixed length, and define the transformation t, depending on h and M_y, such that for any $x = a_1a_2\cdots a_n$, $t(x) = h(a_1)\#y\#h(a_2)\#y\#\cdots\#y\#h(a_n)$ where y is an encoding of M_y that is clocked to run for at most n^k steps. Clearly $t \in \mathbf{1FST}$ and L is many-one reducible to L^k via t. It remains to show that L^k is in $XTIME(n^k)$. On input w, we can verify in linear time that w has the form given by (1). Then we simulate M_y. Because there is a copy of the machine encoding y next to each 'real input symbol' $h(a_i) = x_i$, it takes only $O(|y| + |x_1|)$ steps to simulate one step of M_y. This gives $O((|y| + |x_1|) \cdot n^k) = O(|w|^k)$ steps overall. By the linear speedup theorem, we can make this exactly $|w|^k$ steps. **Q.E.D.**

Similarly, we can show that reversal classes defined by polynomial or exponential complexity functions have complete languages [7]:

THEOREM 4. *The classes $DREVERSAL(n^{O(1)})$ and $DREVERSAL(O(1)^n)$ have complete languages under log-space many-one reducibility.*

Proof. We show this for the case of $DREVERSAL(n^{O(1)})$. Using the existence of a efficient universal machine $\{U_1, U_2, \ldots\}$ for this characteristic class, we see that the language

$$\{i\#^m x \ : \ m = |x|^{|i|}, x \in L(U_i)\}$$

is complete. **Q.E.D.**

Natural Complete Languages. The complete languages generated in the above manner are artificial and it is of interest to obtain 'natural' complete problems. By natural problems we mean those arising in contexts that have independent interest, not just concocted for the present purpose. (Of course, naturalness is a matter of degree.) The advantage of natural complete problems is that they are an invaluable guide as to the inherent complexity of related natural problems.[1] We examine such languages in the remainder of this chapter. For each complexity class K studied below, we first give a direct proof that a language L_0 is complete for K under $\leq^L_m$-reducibility. Subsequently, we may show any other languages L to be K-complete by showing $L_0 \leq^L_m L$, using the fact that $\leq^L_m$-reducibility is transitive.

Notation We will conveniently use $[i..j]$ to denote the set $\{i, i+1, \ldots, j\}$ where $i \leq j$ are integers.

5.2 Complete Problems for Logarithmic Space

5.2.1 Graph Accessibility

The first problem shown to be complete for *NLOG* is the following. It was originally introduced as 'threadable mazes' by Savitch [29] but the present form is due to Jones [19].

Graph Accessibility Problem (GAP)

Given:	A pair $\langle n, G\rangle$ where G is an $n \times n$ adjacency matrix of a directed graph on the node set $[1..n]$.
Property:	There is a path from node 1 to node n.

To show that GAP is in *NLOG*, we describe a nondeterministic acceptor M that guesses a path through the graph as follows: on input x, M first verifies that x has the correct format representing the pair $\langle n, G\rangle$. Then it writes down '1' on tape 1 and on tape 2 makes a nondeterministic guess of some j, $1 \leq j \leq n$. In general, suppose tapes 1 and 2 contain the integers i and j (respectively). Then it verifies that $\langle i, j\rangle$ is an edge of G. If not, it rejects at once. It next checks if $j = n$. If so, it accepts; otherwise it copies j to tape 1 (overwriting the value i) and makes another guess k on tape 2. Now we have the pair $\langle j, k\rangle$ represented on tapes 1 and 2, and we may repeat the above

[1] Cf. comments in footnote 18, in chapter 1.

verification and guessing process. Clearly the input is in GAP iff some sequence of guesses will lead to acceptance. The space required by M is $\log n$ where the input size is $O(n^2)$.

To show that GAP is *NLOG*-hard, let L be accepted by some nondeterministic machine N in space $\log n$. We reduce L to GAP using a transducer T that on input x computes the transformation $t(x)$ as follows: let $|x| = n$ and we may assume that each configuration of N that uses at most $\log n$ space on input x. is represented by integers in the range $[2..n^c - 1]$ for some integer $c > 0$. (Some integers in this range may not encode any configuration.) Then $t(x)$ is the encoding of $\langle n^c, G\rangle$ where the entries $G_{i,j}$ of the matrix G (representing a graph) is defined as follows: $G_{i,j} = 1$ iff one of the following holds:

(1) i and j encode configurations C_i and C_j (respectively) of M and $C_i \vdash C_j$.

(2) $i = 1$ and j represents the initial configuration.

(3) i represents an accepting configuration and $j = n^c$.

It is not hard to see that T can output the successive rows of G using only space $\log n$. Furthermore, x is accepted by N iff $t(x)$ is in GAP. This completes the proof.

5.2.2 Unsatisfiability of 2CNF formulas

For the next *NLOG*-complete problem, recall from chapter 3 that the kCNF formulas ($k \in \mathbb{N}$) are those with exactly k literals per clause.

Unsatisfiability of 2CNF Formulas (2UNSAT)

Given: A 2CNF formula F.

Property: F is unsatisfiable.

With respect to a 2CNF formula F, let us write $u \to v$ if $\{\bar{u}, v\} \in F$. Thus "$\to$" here is just "logical implication". Clearly, $u \to v$ iff $\bar{v} \to \bar{u}$. The reflexive, transitive closure of $\to$ is denoted $\stackrel{*}{\to}$. Thus $u \stackrel{*}{\to} v$ iff there is a sequence $u_1, u_2, \ldots, u_k$ ($k \geq 1$) of literals such that $u_1 = u$, $u_k = v$ and $u_i \to u_{i+1}$ for $i = 1, \ldots, k-1$. The literals u_i in the sequence need not be distinct. We note that *if I is a satisfying assignment for F and $u \stackrel{*}{\to} v$ then $I(u) = 1$ implies* $I(v) = 1$. It follows that if $u \stackrel{*}{\to} \bar{u}$ and I satisfies F then $I(u) = 0$. For any set of literals X, let the closure of X be $cl(X) = \{v : (\exists u \in X) u \stackrel{*}{\to} v\}$. Hence if I satisfies F and $I(u) = 1$ for all $u \in X$ then $I(v) = 1$ for all $v \in cl(X)$. We begin with the following characterization:

THEOREM 5. *A 2CNF formula F is unsatisfiable iff there exists a literal u such that $u \stackrel{*}{\to} \bar{u}$ and $\bar{u} \stackrel{*}{\to} u$.*

Proof. If there is such a literal u. the above remarks makes it clear that F is unsatisfiable. Conversely, suppose there are no such literal. We will show that F is satisfiable. A set of literals is *consistent* if it does not contain both x and $\bar{x}$ for any variable x. We claim that for each variable x, either the set $cl(\{x\})$ is consistent or the set $cl(\{\bar{x}\})$ is consistent: otherwise let y and z be variables such that

$$\{y, \bar{y}\} \subseteq cl(\{x\}) \quad \text{and} \quad \{z, \bar{z}\} \subseteq cl(\{\bar{x}\})$$

Thus $x \stackrel{*}{\to} y$ and $x \stackrel{*}{\to} \bar{y}$. But $x \stackrel{*}{\to} \bar{y}$ implies $y \stackrel{*}{\to} \bar{x}$ so that transitivity implies $x \stackrel{*}{\to} \bar{x}$. A similar argument using $z, \bar{z}$ shows that $\bar{x} \stackrel{*}{\to} x$, contradicting our assumption that there are no such x.

We now define a sequence $U_0 \subseteq U_1 \subseteq \cdots \subseteq U_m$ (for some $m \geq 1$) of sets of literals: and $U_1 = cl(U_0)$. Let $U_0 = \emptyset$ (the empty set). Supposed U_i ($i \geq 0$) is defined. If for every variable x, either x or $\bar{x}$ is in U_i then we stop (i.e., set m to be i). Otherwise choose such a variable x and by the above observation, either $cl(\{x\})$ or $cl(\{\bar{x}\})$ is consistent. If $cl(\{x\})$ is consistent, set $U_{i+1} := U_i \cup cl(\{x\})$. Else, set $U_{i+1} := U_i \cup cl(\{\bar{x}\})$.

It is immediate that each U_i is closed, i.e., $cl(U_i) = U_i$. Suppose U_m is consistent. Then the assignment that makes each literal in U_m true is a satisfying assignment for F. To see this, suppose $\{u, v\}$ is a clause in F and u is not in U_m. This means $\bar{u}$ is in U_m. But $\bar{u} \stackrel{*}{\to} v$, so v is in U_m. This shows that every clause is satisfied, as desired.

It remains to show the consistency of U_m. We show inductively that each U_i, $i = 0, \ldots, m$, is consistent. U_0 is clearly consistent. Next assume that U_i ($i \geq 0$) is consistent but $U_{i+1} = U_i \cup cl(\{u\})$ is inconsistent. Say, $\{v, \bar{v}\} \subseteq U_{i+1}$. We may assume that $\bar{v} \in U_i$ and $v \in cl(\{u\})$. But $v \in cl(\{u\})$ implies $u \stackrel{*}{\to} v$, or equivalently, $\bar{v} \stackrel{*}{\to} \bar{u}$. Then $\bar{v} \in U_i$ implies $\bar{u} \in U_i$, contradicting our choice of u when defining U_{i+1}. **Q.E.D.**

From this lemma, we see that 2UNSAT is in *NLOG*: on input F, guess a literal u such that $u \stackrel{*}{\to} u$ witnesses the unsatisfiability of F. We then guess the sequence $u \to u_1 \to u_2 \to \cdots \to u_k \to u$. We need not store the entire

sequence, just the first literal u and a current literal u_i. If we assume that each literal is encoded by a string of length $O(\log n)$ when F has n distinct variables, the space used is clearly $\log n$. We could impose this restriction on the encoding. However, it is not hard to see that even if no such restrictions were made, logarithmic space is sufficient. The idea is that any variable in F can be represented by a counter of size $\log n$ that points to some occurrence of that variable in F. The details are left as an exercise.

We now prove that 2UNSAT is *NLOG*-hard by reducing GAP to it. Let $\langle n, G\rangle$ be the input to GAP. We can construct the following CNF formula whose variables are denoted by the integers $1, \ldots, n$:

$$1 \wedge \bar{n} \wedge \bigwedge_{\langle i,j\rangle} (\bar{i} \vee j)$$

where $\langle i, j\rangle$ range over edges in G. Note that this is not quite a 2CNF formula because the first two clauses have only one literal each.[2] To turn a one variable clause $\{u\}$ to two variables, it is easy to introduce a new variable z so that $\bar{u} \to z \to u$. If we interpret the variable i to mean "node i is reachable from node 1", then this formula says that node 1 but not node n is reachable, and if $\langle i, j\rangle$ is an edge and node i is reachable then so is node j. Clearly, if there is a path from 1 to n then the formula is unsatisfiable. Conversely, if there are no paths from 1 to n then (exercise) it is easy to see that the formula is satisfiable with this assignment: (a) assign to false all those nodes that can reach node n, (b) assign to true all the remaining nodes. Finally, the fact that this formula can be constructed from $\langle n, G\rangle$ in logarithmic space completes the proof.

Based on our intuition from the satisfiability problem SAT, it is perhaps curious that it is 2UNSAT rather than 2SAT that is *NLOG*-complete.[3]

5.2.3 Associative Generator Problem

The next problem concerns a simple algebraic system with one associative binary operation denoted $\oplus$. Without loss of generality, the elements of the system are $[1..n]$ and let the multiplication table of $\oplus$ be given by an $n \times n$ matrix $T = (T_{i,j})$. Thus $T_{i,j} = k$ means $i \oplus j = k$. For any set $S \subseteq [1..n]$, the $\oplus$-closure of S is the smallest set containing S and closed under $\oplus$. We should keep in mind that the elements in $[1..n]$ represents abstract elements, not integers.

Associative Generator Problem (AGEN)

Given:	$\langle n, T, S, w\rangle$ where T is a multiplication table for a binary operation $\oplus$ over $[1..n]$, $S \subseteq [1..n]$ and $w \in [1..n]$.
Property:	The operation $\oplus$ is associative and w is in the $\oplus$-closure of S.

We sketch the proof. To show that AGEN is in *NLOG*, we note that it is easy to check in logarithmic space that an input has the form $\langle n, T, S, w\rangle$ and that T is associative. Now w is in the $\oplus$-closure of S iff it can be expressed as $w = x_1 \oplus x_2 \oplus \cdots \oplus x_k$. It is not hard to see that the shortest such expression has length of k at most n. Hence it is easy to nondeterministically accept AGEN by guessing successive values of x_i, for $i = 1, \ldots, n$, and only keeping track of the partial sum $x_1 \oplus x_2 \oplus \cdots \oplus x_i$.

Next, we show that GAP is reducible to AGEN. Let $\langle n, G\rangle$ be the input to GAP. We describe the AGEN instance $\langle m, T, S, w\rangle$ as follows: Choose $m = 1 + n + n^2$. Instead of describing the table T in terms of $[1..m]$, it is more intuitive to interpret the elements of $[1..m]$ as the set

$$X = [1..n] \cup ([1..n] \times [1..n]) \cup \{\infty\}$$

where ∞ is a new symbol. Hence $|X| = m$. We can make a correspondence between X and $[1..m]$ as follows: Each $i \in [1..n]$ corresponds to itself in $[1..m]$; the pair $\langle i, j\rangle \in [1..n] \times [1..n]$ corresponds to $i + jn$ in $[1..m]$ and finally ∞ corresponds to m. The table T is now described by the following rules:

For all $i, j, k, u, v \in [1 \ldots n]$,

$$i \oplus \langle j, k\rangle = \begin{cases} k & \text{if } i = j \text{ and } \langle j, k\rangle \text{ is an edge of } G \\ \infty & \text{else.} \end{cases}$$

$$\langle i, j\rangle \oplus \langle u, v\rangle = \begin{cases} \langle i, v\rangle & \text{if } j = u \\ \infty & \text{else.} \end{cases}$$

[2] Since we regard clauses as sets of literals, it is no good saying that the clause with one literal can be made into a 2-literal clause by repeating the single literal.

[3] Nevertheless, this switch is typical when we change from a time class to space class.

In all other cases, the value of $x \oplus y$ is ∞. It is easy to see that $\oplus$ is associative and there is a path $(1, x_1, \ldots, x_k, n)$ from 1 to n in G iff $1 \oplus \langle 1, x_1\rangle \oplus \cdots \oplus \langle x_k, n\rangle = n$ holds. The desired AGEN instance $\langle m, T, S, w\rangle$ is now evident: m and T has been described; S is the set consisting of 1 and all pairs $\langle i, j\rangle$ representing edges of G; w is chosen to be n. Finally, one verifies that this AGEN instance can be computed in logarithmic space from $\langle n, G\rangle$.

5.2.4 Deterministic Logarithmic Space and below

To get a meaningful complete language for *DLOG*, we weaken our reducibility as in section 1. We restrict the input head on the transducers to be one-way. A natural restriction of GAP that is *DLOG*-complete is to make the outdegree of each node at most 1. However, we must also be careful about the actual encoding of the problem. In particular, we can no longer assume the adjacency matrix representation of graphs:

Deterministic Graph Accessibility Problem (1GAP)

Given:	$\langle n, G\rangle$ as in GAP except that each vertex of G has outdegree at most 1 and G is encoded by a sequence of its edges.
Property:	There is a path from node 1 to n.

It is easy to see that this problem is in *DLOG*. We leave it as an exercise to reduce an arbitrary language in *DLOG* to 1GAP: essentially we cycle through all configurations C and for each C, we output the 'edge' pair (C, C') where C' is the successor of C, $C \vdash C'$. Indeed, this proof shows that the following variant of 1GAP remains complete: where we insist that the inputs $\langle n, G\rangle$ satisfy the property that the list of edges of G be topologically sorted (i.e., all edges entering a node must be listed before edges exiting from that node).

Let us briefly consider the power of one-way log-space many-one reducibility [17]. It turns out that many natural problems which are complete for canonical classes (under $\leq_m^L$) remain complete under $\leq_m^{1L}$. This is fortuitous because it is easily shown that there are languages that are complete for P (for instance) under $\leq_m^L$ but which are not P-complete under $\leq_m^{1L}$. See Exercises.

While there is less interest sublogarithmic space complexity, it should be noted that complete problems can be defined for any nondeterministic space bound between $\log\log n$ and $\log n$ space. Monien and Sudborough define the so-called 'bandwidth limited' versions of the GAP problem [24]. To understand this, we say the *bandwidth* of a graph G is the smallest $m \geq 0$ such that each edge (i, j) satisfies $|i - j| \leq m$. We must also be careful about how to encode such graphs. Let us encode graphs on the vertices $\{1, 2, \ldots, n\}$ by a sequence of the adjacency lists of each vertex. The adjacency list for vertex i has the form $(i, m(i, 1), m(i, 2), \ldots, m(i, d_i))$ where d_i is the out-degree and the edges exiting from i go to $i + m(i, j)$ for each j. Thus $m(i, j)$ is an incremental positive or negative value, relative to i. Naturally the bandwidth of G is an upper bound on their absolute value. Now for any complexity function $f(n) \leq \log n$ define the language $GAP(f)$ to encode the family of graphs G such that if the vertices are $\{1, \ldots, n\}$ then the bandwith of G is at most $f(n)$, and such that there is a path from node 1 to node n. Then it can be shown that $GAP(f)$ is complete for $NSPACE(f)$ under $\leq_m^L$-reducibility. We may further weaken the reducibility so that the transducers use only $f(n)$ space, provided we encode the languages even more compactly: omit the first number i in the list $(i, m(i, 1), \ldots, m(i, d_i))$, and so i is implicit from the order of occurrence of its list.

5.3 Complete Problems for P

The P-complete problems are important because of an interpretation that such problems are "inherently sequential". This terminology needs explanation: just as P is viewed as the feasible time class (in the fundamental mode), the subset of P which can be solved in poly-logarithmic time by parallel computers are considered to be "parallelizable". Showing that a problem is P-hard is therefore evidence that the problem is not parallelization (or, inherently sequential). For further information about P-complete problems, we refer to [13].

5.3.1 Unit Resolution

We begin by showing a problem based on the *resolution proof system* to be P-complete. This is actually part of a theme started in the previous section where we show that 2UNSAT is *NLOG*-complete problem. That is, we will show a family of related problems arising from theorem proving, each of which is complete for some class in our canonical. The famous SAT problem (chapter 3) is part of this family.

The basic concept of resolution proofs is simple: Two clauses C and C' are *resolvable* if there is a literal u such that $u \in C$ and $\bar{u} \in C'$. Their *resolvent* is the clause $C'' = (C \cup C') - \{u, \bar{u}\}$. This rule for obtaining the resolvent is

called the *resolution rule* (also known as *annihilation rule* or *cut rule*), with the literals u and $\bar{u}$ being *annihilated*. A *unit clause* is a clause with one literal. If either C or C' is a unit clause, we say their resolvent C'' is a *unit resolvent* of C and C'. Note that if both C and C' are unit clauses then their resolvent is the empty clause which we denote by $\square$ (not to be confused with the blank symbol, but the context should make it clear). By definition, the empty clause is unsatisfiable. Let F be a CNF formula in clause form. A *deduction* (resp. *unit deduction*) of a clause C_m from F is a sequence of clauses $C_1, \ldots, C_m$ where each C_i is either in F or is the resolvent (resp. unit resolvent) of two previous clauses in the sequence. The fundamental result on resolution is the following:

Resolution Theorem. *F is unsatisfiable iff there is a deduction of $\square$ from F.*

The proof is an exercise; we do not use the theorem in the following. The following P-complete problem was introduced by Jones and Laaser [20]:

Unit Resolution for 3CNF formulas (3UNIT)

Given: A 3CNF formula F.

Property: $\square$ can be deduced from F using unit deduction.

We first show that 3UNIT is in P. Consider the following algorithm that takes as input a CNF formula F. If $\square \in F$ then we accept at once. So assume otherwise. The algorithm maintains three sets of clauses G, G' and H. Initially, G' is empty, G consists of all the unit clauses in F and H is $F - G$. In general, the sets G, G' will consist of unit clauses only; hence we may sometimes conveniently regard them as a set of literals. Also $G \cap G' = \emptyset$. The algorithm executes the following loop:

While $G \neq \emptyset$ do

1. Choose any $u \in G$. Let $G := G - \{u\}$; $G' := G' \cup \{u\}$.
2. For each clause C in H: if $\{u\}$ and C are resolvable, and their resolvent $C' = C - \{\bar{u}\}$ is not in $G \cup G' \cup H$ then add C' to G or to H according to whether C' is a unit clause or not. If $\square \in H$, accept; else reject.

It is important to realize that in line 2, the clause C is not removed from H (the Exercises explain why this is necessary for the correctness of resolution). Also, we could have accepted within the loop once the empty clause is generated, but for the sake of the proof below, we do not do this.

Correctness: We must show that the algorithm accepts iff $F \in$ 3UNIT. Note that new clauses are generated only in line 3 using the unit resolution rule; also, all clauses in $G' \cup G \cup H$ come from the original set F or are unit deducible from F. Hence if the algorithm accepts in line 2 then F is in 3UNIT. To prove the converse, suppose

$$C_0, C_1, \ldots, C_m \quad (\text{where } C_m = \square) \tag{2}$$

is a unit deduction of $\square$. We will show that the algorithm accepts. It is crucial to note that step 3 is specified in such a way that each literal u is put into G at most once. Thus the number of iterations in the while-loop is at most the number of literals appearing in the formula F. Let G_i $(i = 0, 1, \ldots)$ denote the set of clauses in G at the end of the ith iteration of the loop. G'_i and H_i are similarly defined, and let $E_i := G_i \cup G'_i \cup H_i$. Clearly, $E_i \subseteq E_{i+1}$. CLAIM: each C_j in the unit deduction (2) belongs to some E_i. In particular $C_m = \square$ appears in some E_i, so our algorithm accepts.

Complexity: As noted above, the number iterations of the while-loop is linear in the size of F. It therefore suffices to show that each iteration takes polynomial time. Note that if a clause of F has k literals, then the clause can potentially spawn 2^k subclauses. But since F is a 3CNF formula, the total number of new clauses is only linear in the number of original clauses. This in turn bounds the time to do each iteration. Thus the time of the algorithm for 3UNIT is polynomial.

We now show that 3UNIT is P-hard. The proof is similar to that for Cook's theorem. Let M accepts in deterministic time n^k for some k. Again we may assume that M is a simple Turing acceptor. We may further assume that M never moves left of its original head position (in cell 1) and that if it accepts at all then it returns to cell 1 and writes a special symbol a_0 there just before accepting. For any input x, we describe a 3CNF formula $f(x)$ such that x is accepted by M iff $f(x) \in$ 3UNIT. Let

$$C_0, C_1, \ldots, C_{n^k} \tag{3}$$

be the computation path of M on input x. As usual, if M halted before n^k steps we assume the last configuration is repeated in this path, and if M uses more than n^k steps, we will prematurely truncate the sequence. It is convenient

to encode a configuration as a word of length $m := 1 + 2n^k$ in $\Sigma^* \cdot [Q \times \Sigma] \cdot \Sigma^*$ where Σ are the tape symbols (including $\sqcup$) of M and Q are the states and $[Q \times \Sigma]$ a new set of symbols of the form $[q, a]$ where $q \in Q, a \in \Sigma$. Thus a word of the form $w_1[q, a]w_2$ (where $w_i \in \Sigma^*$) represents the configuration whose tape contents (flanked by blanks as needed) are $w_1 a w_2$ with the machine scanning symbol a in state q. We next introduce the Boolean variables $P_{i,t}^a$ that stands for the following proposition:

"symbol a is in the ith cell of configuration C_t".

A similar meaning is accorded the Boolean variable $P_{i,t}^{[q,a]}$ where $[q, a] \in [Q \times \Sigma]$. Let $x = a_1 a_2 \cdots a_n$ where each $a_j \in \Sigma$. The formula $f(x)$ is a conjunction of the form $F_0 \wedge F_1 \wedge \cdots \wedge F_{m+1}$. The last formula F_{m+1} is special and is simply given by:

$$F_{m+1} \quad : \quad \neg P_{1,m}^{[q_a, a_0]}$$

The formula F_t $(t = 0, \ldots, m)$ is a 3CNF formula asserting that conditions that the above variables must satisfy if they encode the configuration C_t in (3). Thus the first configuration C_0 is

$$F_0 \quad : \quad P_{1,0}^{[q_0,a_1]} \wedge P_{2,0}^{a_2} \wedge \cdots \wedge P_{n,0}^{a_n} \wedge \left(\bigwedge_{i=n+1}^{m} P_{i,0}^{\square} \right)$$

We indicate the form of the remaining formulas. Let

$$\partial : (\Sigma \cup [Q \times \Sigma])^3 \to \Sigma \cup [Q \times \Sigma]$$

be the function that encodes the transition table of M. Roughly, $\partial(a, b, c) = b'$ means that the symbol b in a cell whose left and right neighbors are a and c respectively will turn to b' in the next step (with understanding that the 'symbol' in a cell could mean a pair of the form $[q, a']$). For example, suppose the M has the instruction that says "on scanning symbol b in state q, write the symbol b', move the head to the right and enter the state q'". We would then have the following:

$$\partial([q, b], a, c) = [q', a], \quad \partial(a, [q, b], c) = b', \quad \partial(a, c, [q, b]) = c.$$

Also, clearly $\partial(a, b, c) = b$ for all a, b, c in Σ. Now it is easy to understand the following definition of F_t $(t = 1, \ldots, m)$:

$$F_t : \bigwedge_i \left(\bigwedge_{a,b,c} (\neg P_{i-1,t-1}^a \vee \neg P_{i,t-1}^b \vee \neg P_{i+1,t-1}^c \vee P_{i,t}^{\partial(a,b,c)}) \right).$$

We now claim:

(i) The unit clause $P_{i,t}^a$ is unit deducible from $f(x) = F_1 \wedge \cdots \wedge F_m$ iff a is the symbol or [state, symbol]-pair at cell i of C_t.

(ii) Furthermore, no other unit clause can be deduced.

The claim is clearly true for initial configuration, $t = 0$. The result can be established by a straightforward proof using induction on t; we leave this to the reader.

We now show that x is accepted by M iff $\square$ is unit deducible from $f(x)$. Claim (i) shows that x is accepted iff $P_{1,m}^{[q_a,a_0]}$ is deducible. But this clause can be unit resolved with F_{m+1} to obtain $\square$. Next claim (ii) easily implies that there are no other ways to deduce $\square$. This shows the correctness of the described formula. It is easy to construct $f(x)$ in logarithmic space.

Note that each of the clauses in the proof uses at most 4 literals. Using techniques similar to that in proving 3SAT *NP*-complete, we can replace these by clauses with exactly 3 literals per clause. $\square$

5.3.2 Path Systems

We now consider a problem that is the first one to be proved complete for P . It was invented by Cook [8] who formulated it as an abstraction of the proof method of both Savitch's theorem (chapter 2) as well as a well-known theorem that context-free languages are in $DSPACE(\log^2 n)$ (see Exercises for this connection.)

Definition. A *path system* is a quadruple $\langle n, R, S, T \rangle$ where $n > 0$ is an integer, $R \subseteq [1..n] \times [1..n] \times [1..n]$ is a relation over $[1..n]$, and S, T are subsets of $[1..n]$. Each integer $u \in [1..n]$ represents a 'node' and node u is said to be *accessible* if $u \in T$ or if there exist accessible nodes v, w such that $\langle u, v, w \rangle$ is in R. The path system is *solvable* if some node in S is accessible.

Path System Accessibility (PSA)

Given: A path system $\langle n, R, S, T\rangle$.

Property: The system is solvable.

We first claim that PSA is in P : consider the Turing acceptor that begins by computing the sets $T_0, T_1, \ldots$ of nodes in stages. In the initial stage, it computes $T_0 = T$. In stage $i \geq 1$, it computes T_i consisting of those new nodes that are accessible from T_{i-1}. It stops when no more new nodes are accessible. Clearly there are at most n stages and each stage takes time $O(n^3)$. So the whole algorithm is $O(n^4)$. Note that the size of the input can be much smaller than n (in fact as small as $O(\log n)$) if not all the integers in $[1..n]$ occur in R, S, T. Hence the algorithm could be exponential in the input size. However, with a more careful implementation (cf. the demonstration that 2UNSAT is in $NLOG$), we can easily ensure that the time is $O(m^4)$ where $m \leq n$ is the number of nodes that actually occur in the path system description. Hence our algorithm is polynomial time.

We now show that PSA is P -hard by reducing 3UNIT to it. Given a 3CNF formula F, let X be the set of all clauses that are subsets of clauses of F. We describe the corresponding path system $\langle n, R, S, T\rangle$. Let $n = |X|$ and let each integer $i \in [1..n]$ represent a clause of X. R consists of those triples $\langle i, j, k\rangle$ of clauses in X such that i is the unit resolvent of j and k. Let the set T be equal to F; let S consist of just the empty clause $\square$. It is easy to see that $\langle n, R, S, T\rangle$ is solvable if and only if F is in 3UNIT. This completes our proof that PSA is P -complete.

5.3.3 Non-associative Generator Problem

The next problem is closely related to the $NLOG$-complete problem AGEN. Basically, the AGEN turns into a P -complete problem when we remove the associative property. We now use $\otimes$ to denote the non-associative binary operation.

Generator Problem (GEN)

Given: Given $\langle n, T, S, w\rangle$ as in AGEN with T representing the multiplication table of a binary operation $\otimes$ over $[1..n]$.

Property: w is in the $\otimes$-closure of S.

The proof that GEN is in P is straightforward. To show that it is P -hard we can reduce 3UNIT to it, using a proof very similar to that for the PSA problem.

5.3.4 Other Problems

The following problem is due to Ladner [21]. Boolean circuits (which will be systematically treated in chapter 10) are directed acyclic graphs such that each node has in-degree of zero or two. The nodes of indegree zero are called *input* nodes and these are labeled by the integers $1, 2, \ldots, n$ if there are $n \geq 0$ input nodes. The other nodes, called *gates*, are each labeled by some two-variable Boolean function. If these Boolean functions are *logical-and*'s ($\wedge$) and *logical-or*'s ($\vee$) only, then the circuit is *monotone*. An *assignment* I to C is a function $I : [1..n] \to \{0, 1\}$. In a natural way, for each node v of G, we can inductively define a Boolean value $val_C(v, I)$.

Monotone Circuit Value Problem (MCVP)

Given: A monotone Boolean circuit C with a distinguished node u, and an assignment I to the input variables of C.

Property: $val_C(u, I) = 1$.

Ladner originally proved the P -completeness of the (general) circuit value problem (CVP) in which the circuit is not restricted to be monotone; the refinement to monotone circuits is due to Goldschlager [11]. Goldschlager also shows that CVP (not MCVP) remains P -complete when we restrict the underlying graph of the circuits to be planar. Closely related to MCVP is this GAP-like problem:

AND-OR Graph Accessibility Problem (AND-OR-GAP)

Given: A directed acyclic graph G on vertices $[1..n]$, where each vertex is labeled by either 'AND' or 'OR'.

Property: Node n is accessible from node 1. Here a node i is accessible from another set S of nodes if either $i \in S$ (basis case) or else i is an AND-node (resp. OR-node) and all (resp. some) of the predecessors of i are accessible from S (recursive case).

Another P -complete problem is:

Rational Linear Programming problem (RLP)

Given:	An $m \times n$ matrix A and an m-vector $\mathbf{b}$ where the entries of $A, \mathbf{b}$ are rational numbers.
Property:	There is a rational n-vector $\mathbf{x}$ such that $A\mathbf{x} \geq \mathbf{b}$.

The fact that RLP is in P is a result of Khachian as noted in chapter 3; that the problem is P -hard is due to Dobkin, Reiss and Lipton[9]. The RLP problem has a famous history because of its connection to the simplex method and because it was originally belonged to the few problems known to be in $NP \cap$ co-NP that is neither known to be in P nor NP-complete.

Goldschlager, Shaw and Staples [12] have shown that a formulation of the well-known problem Maximum Flow in graphs is P -complete.

Adachi, Iwata and Kasai [1] have defined natural problems that are complete for $DTIME(n^k)$ for each $k \geq 1$.

5.4 Complete Problems for *PSPACE*

5.4.1 Word Problems

A rich source of natural problems with high complexity arises in the study of *extended regular expressions*. Meyer and Stockmeyer [23, 33, 32] and Hunt [18] were among the first to exploit such results. Given an alphabet Σ, the set of extended regular expressions over Σ is a recursively defined set of strings over the alphabet Σ together with the following 9 additional symbols

$$\lambda, +, \cdot, \cap, ^2, ^*, \neg,), ($$

which we assume are not in Σ. Each symbol in $\Sigma \cup \{\lambda\}$ will be called an *atom*. An *extended regular expression* (over Σ) is either an atom or recursively has one of the following forms:

$$(\alpha + \beta), (\alpha \cdot \beta), (\alpha \cap \beta), (\alpha)^2, (\alpha)^*, \neg(\alpha)$$

where α and β are extended regular expressions. Each expression α over Σ denotes a language $(\Sigma, L(\alpha))$ defined as follows: if α is the atom $a \in \Sigma \cup \{\lambda\}$ then $L(\alpha)$ is the language consisting of the single word[4] a; otherwise

Union:	$L(\alpha + \beta) = L(\alpha) \cup L(\beta)$
Concatenation:	$L(\alpha \cdot \beta) = L(\alpha) \cdot L(\beta)$
Intersection:	$L(\alpha \cap \beta) = L(\alpha) \cap L(\beta)$
Squaring:	$L(\alpha^2) = L(\alpha) \cdot L(\alpha)$
Kleene-star:	$L(\alpha^*) = L(\alpha)^* = \bigcup_{i \geq 0} L(\alpha)^i$
Complement:	$L(\neg\alpha) = \Sigma^* - L(\alpha)$

Note that if α is regarded as an expression over a different alphabet Γ (provided all the atoms appearing in the expression α are in Γ) it would denote a different language; in practice, the context will make clear which alphabet is meant or else the specific alphabet is irrelevant to the discussion.

Notations: We often omit the symbol for concatenation (writing '$\alpha\beta$' for '$\alpha \cdot \beta$'). Noting that all our binary operators are associative, and assuming some precedence of operators (unary operators precede binary ones and concatenation precedes union and intersection) we can omit some of the parentheses that might otherwise be needed. If $S = \{a_1, \ldots, a_m\}$ is a set of atoms, we often use the meta-symbol S as a shorthand for the expression $a_1 + a_2 + \cdots + a_m$. If α is any expression and $k \geq 1$ (k may be an integer expression), then we write α^k as a short hand for concatenating k copies of α (note the possible confusion of this notation with applications of the squaring operator if k written as a power of 2; the context will clarify which is meant).

[4]We deliberately abuse notation here, by confusing the symbol 'λ' in extended regular expressions with the usual notation for the empty word; of course it comes to no harm since λ just stands for itself in this sense. Also, we really ought to use '$\cup$ ' for union for consistency with '$\cap$' for intersection; but it seems that the asymmetric choice has better visual aesthetics.

The *size* $|\alpha|$ of an extended regular expression α is defined recursively as follows: if α is an atom then $|\alpha| = 1$; otherwise

$$|\alpha + \beta| = |\alpha \cdot \beta| = |\alpha \cap \beta| = 1 + |\alpha| + |\beta|;$$

and

$$|\alpha^2| = |\alpha^*| = |\neg\alpha| = 1 + |\alpha|.$$

Example 1. The following expression (of size 22) denotes the set of those words over $\Sigma = \{a, b, c\}$ of length between 4 and 8 beginning with a but not terminating with b:

$$a(a + b + c)^2((a + b + c + \lambda)^2)^2(a + c)$$

■

Let Ω be a subset of the extended regular operators $\{+, \cdot, \cap, ^2, ^*, \neg\}$. Then a Ω-*expression* is a extended regular expression that uses only the operators in Ω. Then the following *word problems* are defined:

Inequality:	$\text{INEQ}(\Sigma, \Omega) = \{(\alpha, \beta) :$ α and β are Ω-expressions over Σ and $L(\alpha) \neq L(\beta)\}$
Fullness:	$\text{FULL}(\Sigma, \Omega) = \{\alpha : \alpha$ is an Ω-expression and $L(\alpha) = \Sigma^*\}$
Emptiness:	$\text{EMPTY}(\Sigma, \Omega) = \{\alpha : \alpha$ is an Ω-expression and $L(\alpha) = \emptyset\}$
Membership:	$\text{MEMBER}(\Sigma, \Omega) = \{(x, \alpha) : \alpha$ is an Ω-expression and $x \in L(\alpha)\}$

In the presence of negation, $\neg \in \Omega$, it is evident that the fullness and emptiness problems are equivalent. Furthermore, the complement of the emptiness and fullness problems are special cases of the inequality problem whenever there is an Ω-expression denoting $\emptyset$ or Σ^* (resp.). We will mainly be showing the complexity of inequality and fullness problems in this and the next section. If Σ is $\{0, 1\}$ we just denote these problems by $\text{MEMBER}(\Omega)$, $\text{INEQ}(\Omega)$, etc. If $\Omega = \{\omega_1, \ldots, \omega_k\}$ then we also write $\text{INEQ}(\omega_1, \ldots, \omega_k)$, $\text{FULL}(\omega_1, \ldots, \omega_k)$, etc., instead of $\text{INEQ}(\{\omega_1, \ldots, \omega_k\})$, $\text{FULL}(\{\omega_1, \ldots, \omega_k\})$, etc. . We begin by showing that the complexity of these problems does not really depend on Σ provided $|\Sigma| \geq 2$:

LEMMA 6. *Let Ω be any set of operators and $|\Sigma| \geq 2$.*

(i) *If Ω contains the concatenation operator then* $\text{INEQ}(\Sigma, \Omega) \equiv_m^{1FST} \text{INEQ}(\Omega)$.

(ii) *If Ω contains $\{+, \cdot, ^*\}$ then* $\text{FULL}(\Sigma, \Omega) \equiv_m^{1FST} \text{FULL}(\Omega)$.

Proof. (i) Recall the $\leq_m^{1FST}$-reducibility from chapter 4 (section 1). It is immediate that $\text{INEQ}(\Omega) \leq_m^{1FST} \text{INEQ}(\Sigma, \Omega)$. Conversely, to show $\text{INEQ}(\Sigma, \Omega) \leq_m^{1FST} \text{INEQ}(\Omega)$, we use the usual coding $h : \Sigma \to \{0, 1\}^*$ of a general alphabet by binary strings over $\{0, 1\}$ of fixed length k (for some $k \geq 2$). Then each Ω-expression α over Σ is systematically transformed to an Ω-expression $H(\alpha)$ over $\{0, 1\}$ simply by replacing each occurrence in α of the atom $a \in \Sigma$ by (the $\{\cdot\}$-expression denoting) $h(a)$. It is also easy to see that $(\alpha, \beta) \in \text{INEQ}(\Sigma, \Omega)$ iff $(H(\alpha), H(\beta)) \in \text{INEQ}(\Omega)$.

(ii) Again, one direction is trivial. To show $\text{FULL}(\Sigma, g) \leq_m^{1FST} \text{FULL}(\Omega)$, let h be the coding in (i) and let $C = \{h(a) : a \in \Sigma\} \cup \{\lambda\} \subseteq \{0, 1\}^*$. Let α be an Ω-expression over Σ, and let $H(\alpha)$ be its transformation as in (i). Clearly $L(H(\alpha)) \subseteq C^*$, with equality iff $L(\alpha) = \Sigma^*$. If we can write an Ω-expression β over $\{0, 1\}$ such that $L(\beta) = \Sigma^* - C^*$ then we see that

$$L(\alpha) = \Sigma^* \iff L(H(\alpha) + \beta) = \{0, 1\}^*.$$

It is easy to construct such an expression:

$$\beta : ((0 + 1)^k)^* \cdot ((0 + 1 + \lambda)^k - C) \cdot ((0 + 1)^k)^*$$

where k is the length of the codes $h(a)$, $a \in \Sigma$, and the subexpression '$((0 + 1 + \lambda)^k - C)$' is really a shorthand for an explicit enumeration of the words in the indicated set. These expressions involve only $\{+, \cdot, ^*\}$. **Q.E.D.**

Although it is convenient when proving upper bounds to assume a binary alphabet, it is easier to use the general alphabet Σ when proving lower bounds (i.e. hardness results). Below we shall switch between these two versions of the problem without warning.

5.4.2 Fullness Problem for Regular Expressions

Our first result is from Stockmeyer [32]:

THEOREM 7. FULL$(+,\cdot,^*)$ *is complete for the class* $LBA = NSPACE(n)$.

Note that the $\{+,\cdot,^*\}$-expressions are just the standard *regular expressions* of finite automata theory. In the following proof, we shall often exploit the fact that $NSPACE(n)$ is closed under complementation. By lemma 2, we conclude:

The fullness problem for regular languages is complete for PSPACE.

Recall that a finite automaton is a 0-tape Turing acceptor with a 1-way input tape that runs in real-time ($n+1$ steps); the *size* of the automaton is the number of tuples in its transition table. Clearly the size is at most $|Q|^2(|\Sigma|+1)$ where Q is the set of states and Σ the input alphabet (the '+1' comes from the need to consider rules for the blank symbol). To show that the fullness problem for regular expressions is in $NSPACE(n)$, we first show:

LEMMA 8. *For any regular expression* α *over* Σ*, we can construct in polynomial time and linear space a nondeterministic finite automaton that accepts the language* $(\Sigma, L(\alpha))$*; furthermore the automaton has* $\leq \max\{4, |\alpha|\}$ *states.*

Proof. We construct the transition table δ^α corresponding to α. Denote the start and accepting states for δ^α by q_0^α and q_a^α. The transition table of a nondeterministic finite automaton can be represented as a set of triples of the form $\langle$*current-state, symbol, next-state*$\rangle$. Furthermore, we assume that for any tuple $\langle q, b, q'\rangle$ in δ^α,

$$q' = q_a^\alpha \Longrightarrow b = \square.$$

In other words, all transitions into the accept state occur after the first blank symbol $\square$ is read. A state q is called *penultimate* if there is a transition from q to the accept state q_a^α. The table δ^α is defined by induction on the size of α:

(1) Suppose α is an atom b: if $b = \epsilon$ then the table δ^α consists of the single triple:

$$\langle q_0^\alpha, \square, q_a^\alpha\rangle.$$

If $b \in \Sigma$ then the table consists of two triples:

$$\langle q_0^\alpha, b, q\rangle, \langle q, \square, q_a^\alpha\rangle$$

for some state q.

(2) If α is $\beta+\gamma$ then first form the union $\delta^\beta \cup \delta^\gamma$. We then replace (in the triples) the start states of the original automata for β and γ with the start state for α, and do the same for the accept states.

(3) If α is $\beta\cdot\gamma$ then we again form the union $\delta^\beta \cup \delta^\gamma$, and do the following: if q is a penultimate state of δ^β and δ^γ contains the triple $\langle q_0^\gamma, b, q'\rangle$ then add the triple $\langle q, b, q'\rangle$ to δ^α. This ensures that after the automaton has seen a word in $L(\beta)$, it can continue to try to recognize a word in $L(\gamma)$. Also replace the start state of δ^β by the start state q_0^α and the accept state of δ^γ by the accept state q_a^α. Triples containing the accept state of δ^β and the start state of δ^γ are deleted.

(4) If α is β^* then we first take the table for β and replace its start and accept states with that for α. If q is a penultimate state of δ^β and $\langle q_0^\beta, b, q'\rangle \in \delta^\beta$ then add the triple $\langle q, b, q'\rangle$ to δ^α (this ensures that that the automaton accepts arbitrarily many copies of words in $L(\beta)$). Finally, add the new triple $\langle q_0^\alpha, \square, q_a^\alpha\rangle$ (this allows the empty string to be accepted). Thus δ^α has the same number of states as δ^β

The reader can easily show that the constructed transition table δ^α accepts $L(\alpha)$. The number of states is at most $|\alpha|$ when $|\alpha| \geq 4$. A case analysis shows that when $|\alpha| \leq 3$, δ^α has at most 4 states. The automaton δ^α can be constructed in linear space and polynomial time from α. **Q.E.D.**

We now show that FULL$(+,\cdot,^*)$ is in co-$NSPACE(n)$. It is enough to show how to accept the complement of the language FULL$(+,\cdot,^*)$ in nondeterministic linear space. On input a regular expression α, we want to accept iff $L(\alpha) \neq \{0,1\}^*$. First we construct δ^α and then try to guess a word $x \notin L(\alpha)$: nondeterministically guess successive symbols of x and simulate *all computation paths* of the automaton on the guessed symbol. More precise, we guess

successive symbols of x and keep track of the set S of all the states that can be reached by the automaton δ^α after reading the symbols guessed to this point. It is easy to see how to update S with each new guess, and since S has at most $2|\alpha|$ states, linear space suffices. When we are finished with guessing x, we make the next input symbol be the blank symbol $\sqcup$, and update S for the last time. We accept if and only if S *does not contain* the accept state. To show the correctness of this construction, we see that (i) if there exists such an x then there exists an accepting path, and (ii) if there does not exist such an x then no path is accepting. This proves $\mathrm{FULL}(+,\cdot,^*) \in$ co-$NSPACE(n)$.

To show that the fullness problem is co-$NSPACE(n)$-hard, let M be a nondeterministic simple Turing acceptor accepting in space n. We may assume that M never moves left of the cell 1 and that M enters the accepting state immediately after writing a special symbol a_0 in cell 1. For each input x of M we will construct a regular expression $E(x)$ over some alphabet Δ (see below) such that $x \notin L(\mathrm{M})$ iff $L(E(x)) = \Delta^*$. Let $t(n) = O_{\mathrm{M}}(1)^n$ be an upper bound on the running time of M on inputs of length $n = |x|$. Let

$$C_0 \vdash C_1 \vdash \cdots \vdash C_m \quad (m = t(n))$$

be a computation path on input x where as usual we assume that the last configuration is repeated as often as necessary if M halts in less that m steps. Let $I = \{1, \ldots, |\delta(\mathrm{M})|\}$ where $\delta(\mathrm{M})$ is the transition table of M. We assume that each C_i is encoded as a word w_i of length $n+2$ in

$$\Sigma^* \cdot [Q \times \Sigma \times I] \cdot \Sigma^*$$

where Σ and Q are the set of symbols and states (respectively) of M and $[Q \times \Sigma \times I]$ is the set of symbols of the form $[q, a, i]$, $q \in Q, a \in \Sigma$, $i \in I$. Note that C_i uses use $n+2$ rather than n symbols because we incorporate the adjacent blank symbols on either side of the input into the initial configuration (and thereafter assume the machine do not to exceed these squares). Intuitively, the symbol $[q, a, i]$ says that the current tape head is scanning symbol a in state q and the next instruction to be executed is the ith instruction from $\delta(\mathrm{M})$. Therefore computation path can be encoded as a word

$$\pi(x) = \#w_0\#w_1\#\cdots\#w_m\#$$

where $\#$ is some new symbol not in $\Sigma \cup [Q \times \Sigma \times I]$.

Notation. In the following, let $\Gamma = \Sigma \cup [Q \times \Sigma \times I]$ and let $\Delta = \Gamma \cup \{\#\}$. Thus $\pi(x)$ is a word over Δ.

The regular expression $E(x)$ will represent a language over Δ. In fact $E(x)$ will have the form $E_1 + E_2 + \cdots + E_6$ where the E_i's are next described.

(a) E_1 denotes the set of words over Δ that "does not begin with an $\#$, does not end with a $\#$, or has at most one $\#$". Precisely,

$$E_1 = \Gamma^* + \Gamma^* \cdot \# \cdot \Gamma^* + \Gamma \cdot \Delta^* + \Delta^* \cdot \Gamma.$$

Since $+$ takes precedence over $\cdot$, this expression is unambiguous. As mentioned, we write Γ, Δ, etc., as in the shorthand where a set of symbols X stands for the regular expression $x_1 + \cdots + x_k$ if the distinct symbols in X are $x_1, \ldots, x_k$. Thus, each occurrence of Γ^* in E_1 represents a subexpression of size $2|\Gamma|$.

The remaining regular expressions in $E(x)$ will consider strings that are not in $L(E_1)$. Such strings necessarily have the form

$$y = \#x_1\#x_2\#\cdots\#x_k\# \tag{4}$$

for some $k \geq 1$ and $x_i \in \Gamma^*$. In our informal description of the expressions below, we will be refering to (4).

(b) "Some x_i in (4) does not have a unique symbol in $[Q \times \Sigma \times I]$"

$$E_2 = \Delta^* \cdot \# \cdot \Sigma^* \cdot \# \cdot \Delta^* + \Delta^* \cdot \# \cdot \Gamma^* \cdot [Q \times \Sigma \times I] \cdot \Gamma^* \cdot [Q \times \Sigma \times I] \cdot \Gamma^* \cdot \# \cdot \Delta^*.$$

(c) "Some x_i has length different from $n+2$"

$$E_3 = \Delta^* \cdot \# \cdot \Gamma^{n+3} \cdot \Gamma^* \cdot \# \cdot \Delta^* + \Delta^* \cdot \# \cdot (\Gamma \cup \{\lambda\})^{n+1} \cdot \# \cdot \Delta^*.$$

(d) "x_1 does not represent the initial configuration on input $x = a_1 \cdots a_n$"
Note that the initial configuration is represented by one of the forms

$$\sqcup[q_0, a_1, i]a_2 \cdots a_n \sqcup$$

where $i \in I$. Let $[q, a, I]$ denotes the set $\{[q, a, i] \,:\, i \in I\}$. For any subset S of Δ, we use the shorthand '$\bar{S}$' for the regular expression denoting the set $\Delta - S$. If $S = \{a\}$, we simply write $\bar{a}$ for $\Delta - \{a\}$. We then have:

$$E_4 = \# \cdot (\overline{\sqcup} + (\overline{[q_0, a_1, I]} + [q_0, a_1, I](\bar{a}_2 + a_2(\bar{a}_3 + \cdots + a_n(\overline{\sqcup} + \Gamma^*)\cdots)))) \cdot \Delta^*.$$

(e) "There are no accepting configurations in y"
That is, none of the symbols of $[q_a, a_0, I]$ appears in y. Here q_a is the accepting state and a_0 the special symbol that is written when the machine halts. Here we assume that the transition table for M contains trivial transitions from the accept state in which 'nothing changes' (the state, scanned symbol and head position are unchanged).

$$E_5 = (\Delta - [q_a, a_0, I])^*.$$

(f) "Some transition, from x_i to x_{i+1}, is not legal"
Recall the function $\partial : \Gamma \times \Gamma \times \Gamma \to \Gamma$ defined in the proof that 3UNIT is P -hard. There, the interpretation of $\partial(a, b, c) = b'$ is that a cell containing b with neighbors a and c will change its contents to b' in the next step. For the present proof, we can easily modify the function ∂ to account for the case where a or c might be the new symbol # and where a, b, c could be an element of $[Q \times \Sigma \times I]$. Furthermore, since M is nondeterministic, $\partial(a, b, c)$ is now a subset (possibly empty) of Δ.

For instance, if b has the form $[q, b', i]$, the ith instruction is indeed executable in state q scanning b', and b' is changed to b'' and the tape head moves right. Then we put b'' in $\partial(a, b, c)$. Moreover, if the ith instruction changes state q to q', then $\partial(b, c, d)$ (for any d) contains $[q', c, j]$ for each $j \in I$. Note that we are slightly wasteful here in allowing all possible j's to appear in the next head position, but it does not matter.

If the ith instruction is not executable, we simply set $\partial(a, b, c) = \emptyset$. Thus E_6 has the form

$$E_6 = F_1 + F_2 + \cdots$$

where each F_i corresponds to a triple of symbols in Δ^3. If F_i corresponds to the triple $(a, b, c) \in \Delta^3$ then

$$F_i = \Delta^* \cdot a \cdot b \cdot c \cdot \Delta^{n+1} \cdot \overline{\partial(a, b, c)} \cdot \Delta^*.$$

We have now completely described $E(x)$. A moment's reflection will convince the reader that $L(E(x)) = \Delta^*$ iff x is rejected by M; thus $E(x)$ is in FULL$(\Delta, \{+, \cdot, ^*\})$ iff $x \notin L(\mathrm{M})$. (Exercise: why do we need the explicit introduction of the variable i in $[q, a, i]$?) The size of $E(x)$ is given by

$$|E_1| + \cdots + |E_6| + 5 = O(1) + O(1) + O(n) + O(n) + O(1) + O(n)$$

which is $O(n)$. It is easy to verify that $E(x)$ can be computed from x in logarithmic space. Thus we have shown that every language in co-$NSPACE(n)$ is reducible to FULL$(\Delta, \{+, \cdot, ^*\})$ and hence (by lemma 6(ii)) reducible to FULL$(+, \cdot, ^*)$.

5.4.3 Complexity of Games

Another rich source of problems with high complexity is the area of combinatorial games. One type of game can be defined as follows.

Definition. A *two-person game* is a quintuple

$$\langle \mathcal{I}_0, \mathcal{I}_1, p_0, \mathcal{R}_0, \mathcal{R}_1 \rangle$$

where $\mathcal{I}_0, \mathcal{I}_1$ are disjoint sets of *positions* (for *player* 0 and *player* 1, respectively), $p_0 \in \mathcal{I}_0$ is a distinguished *start* position, $\mathcal{R}_0 \subseteq \mathcal{I}_0 \times \mathcal{I}_1$ and $\mathcal{R}_1 \subseteq \mathcal{I}_1 \times \mathcal{I}_0$.

A pair $\langle p, p' \rangle$ in $\mathcal{R}_0$ is called a *move* for player 0; we say that there is a *move from p to p'*. The analogous definition holds for player 1. We call p an *end position* for player b ($b = 0, 1$) if $p \in \mathcal{I}_b$ and if there are no moves from p. A *match* is a sequence $\mu = (p_0, p_1, p_2, \ldots)$ of positions beginning with the start position such that $\langle p_i, p_{i+1} \rangle$ is a move for all $i \geq 0$ and either the sequence is finite with the last position an end position, or else the sequence is infinite. For $b = 0, 1$, we say that *player b loses a match μ* if the sequence μ is finite and the last position in μ is an end position for player b; in that case player $1 - b$ *wins the match*. In other words, the player whose turn it is to play loses if he has no move. The match is a *draw* if it is infinite. A position p is a *forced win* for Player b if (basis case) p is an end position for player $1 - b$, or else (inductive case):

(a) either $p \in \mathcal{I}_b$ and there is a move from p to a forced win for player b

(b) or $p \in \mathcal{I}_{1-b}$ and for every $p' \in \mathcal{I}_b$, if there is a move from p to p' then p' is a forced win for player b.

An example of such a game is by given Even and Tarjan [10]: Let G be a given undirected graph on the vertex set $[1..n]$ for some $n \geq 2$. One player ('short') tries to construct a path from node 1 to node n and the other player ('cut') attempts to frustrate this goal by constructing an $(1, n)$-*antipath* (i.e. a set of nodes such that every path from node 1 to n must pass through the set). The game proceeds by the players alternately picking nodes from the set $[1..n]$: the first player to achieve his goal wins.

More formally: A position is a pair $\langle S_0, S_1\rangle$ where S_0 and S_1 are disjoint sets of vertices and $S_0 \cup S_1 \subseteq \{2, 3, \ldots, n-1\}$. If $|S_0 \cup S_1|$ is even then it is a position of player 0, else of player 1. The start position is $\langle\emptyset, \emptyset\rangle$. An end position is $\langle S_0, S_1\rangle$ such that either

(a) it is a position for player 1 and there is a path in G from node 1 to node n passing only through the vertices in S_0, or

(b) it is a position for player 0 and S_1 is an $(1, n)$-antipath.

Note that (a) and (b) represents, respectively, a winning position for player 0 (short) and player 1 (cut). Also if $S_0 \cup S_1 = \{2, 3, \ldots, n-1\}$ then p must be an end position. Hence there are no draws in this game. Suppose $p = \langle S_0, S_1\rangle$ is a position for player b but p is not an end position. Then there is a move from p to $\langle S_0', S_1'\rangle$ iff for some $v \in \{2, 3, \ldots, n-1\} - (S_0 \cup S_1)$, such that $S'_{1-b} = S_{1-b}$ and $S'_b = S_b \cup \{v\}$. Thus no node is picked twice.

Generalized Hex (HEX)

Given:	An undirected graph G over $[1..n]$, $n \geq 2$.
Property:	Does player 0 have a forced win from the start position?

Remark: This game is also called *Shannon switching game on vertices*. For the analogous game where moves correspond to choosing edges the optimum strategy is considerably easier and can be determined in polynomial time.

THEOREM 9. HEX *is PSPACE-complete.*

The problem is in PSPACE: The start position p_0 is a forced win for player 0 if and only if there exists some tree T with $\leq n-1$ levels with the following properties:

(a) The root is p_0 and counts as level 0. Nodes at even levels are positions of player 0 and nodes at odd levels are positions of player 1.

(b) If there is an edge from a node p to a child node q then there is a move from p to q. If p is a position of player 0 then p has exactly one child. If p is a position of player 1 then the set of children of p represents all possible moves from position p.

The condition that positions for player 0 has exactly one child implies that such positions cannot be leaves of the tree; thus all leaves are at odd levels. Some such tree T can be searched nondeterministically using a depth-first algorithm. The search is fairly standard – basically, we need to keep a stack for the path from the root to the node currently visited. This stack has depth n, so if each node of the path requires linear storage, we have an algorithm using quadratic space. (It is not hard to see that linear space is sufficient.)

We postpone the proof that HEX is *PSPACE*-hard until chapter 9 since it is easy to reduce HEX to the problem QBF (quantified Boolean formulas) that will be introduced there.

Schaefer [30] shows several other games that are complete for *PSPACE*. Orlin [25] and Papadimitriou [26] show other methods of deriving natural *PSPACE*-complete problems.

5.5 Complete problems with exponential complexity

5.5.1 The power of squaring

The next two subsections prove two results of Meyer and Stockmeyer:

THEOREM 10. FULL$(+, \cdot, {}^*, {}^2)$ *is EXPS-complete.*

THEOREM 11. INEQ$(+, \cdot, {}^2)$ *is NEXPT-complete.*

Observe that the first result involves adding the squaring operator to the regular operators; in the second result we replace the Kleene-star operator by the squaring operator. The hardness proofs in both cases come from making simple modifications to the proof for FULL$(+,\cdot,^*)$ in the last section.[5]

In the proof for FULL$(+,\cdot,^*)$, the expression $E(x)$ contains subexpressions of the form S^k where S is one of $\Gamma, (\Gamma \cup \{\alpha\}), \Delta$, and where $k \in \{n-1, n, n+1\}$. Sets of the form S^k are called *rulers* (*of length k*) because they measure the distance between two corresponding (or neighboring) symbols in consecutive configurations. The crux of the present proofs lies in the ability to replace such expressions by squaring expressions of exponentially smaller size. The following lemma makes this precise:

LEMMA 12. *Let Σ be an alphabet and k a positive integer.*

(i) *Then there is an $\{+,\cdot,^2\}$-expression α such that $L(\alpha) = \Sigma^k$ and $|\alpha| = O_\Sigma(\log k)$. Let $[\Sigma^k]_{sq}$ denote such an expression α.*

(ii) *There is a log-space transformation from the binary representation of integer k to $[\Sigma^k]_{sq}$.*

Proof. (i) The proof uses a well-known trick: using induction on k, we have $[\Sigma^1]_{sq} = \Sigma$, $[\Sigma^{2k}]_{sq} = ([\Sigma^k]_{sq})^2$ and $[\Sigma^{2k+1}]_{sq} = \Sigma \cdot [\Sigma^{2k}]_{sq}$. (ii) We leave this as an exercise. **Q.E.D.**

5.5.2 An exponential space complete problem

We prove theorem 10. To show that FULL$(+,\cdot,^*,^2)$ is in *EXPS*, we observe that any expression α with squaring can be expanded to an equivalent expression β without squaring (just replace each expression S^2 by $S \cdot S$). Clearly $|\beta| \leq 2^{|\alpha|}$. But β is a regular expression and by the results of the previous section it can be accepted in space $O(|\beta|)$.

To show that the problem is *EXPS*-hard, suppose that M is a 2^n space-bounded deterministic acceptor (other exponential bounds for *EXPS* machines can be treated in much the same way as shown here for 2^n). Indeed, to make the following notations less cluttered, we may assume that M accepts in space $2^{n-1} - 1$ (use space compression). For each input x of length n, let

$$\pi(x) = \#w_0\#w_1\#\cdots\#w_m\# \quad (\text{where } m = 2^{2^{O(n)}})$$

represent a computation path of M where each configuration w_i has length 2^n. We show how to construct in logarithmic space an expression $\hat{E}(x)$ such that $\hat{E}(x) \in$ FULL$(+,\cdot,^*,^2)$ iff $x \notin L(\text{M})$. (This is sufficient since *EXPS* is closed under complementation.)

We need rulers of length k where $k = 2^n - \delta$, $\delta \in \{-1, 0, +1\}$. We exploit the previous lemma to express these rulers using squaring expressions of size $O(n)$. Recall the regular expression $E(x)$ in the proof for regular expressions in the previous section; the ruler subexpressions in $E(x)$ are of the form $S^{n+\delta}$ where $\delta \in \{-1, 0, +1\}$. The expression $\hat{E}(x)$ is obtained from $E(x)$ by replacing each ruler subexpression $S^{n+\delta}$ in $E(x)$ by the squaring expression $[S^{2^n+\delta}]_{sq}$ defined in the previous lemma. The reader can easily confirm that $\hat{E}(x)$ has linear size.

5.5.3 An exponential time complete problem

We prove theorem 11. It is easy to see that the $\{+,\cdot,^2\}$-expressions can only denote finite languages. We leave it as an exercise to show that the inequality problem for $\{+,\cdot,^2\}$-expressions is in *NEXPT*.

To show that the problem is *NEXPT*-hard, suppose that M is a $2^n - 2$ time-bounded nondeterministic acceptor (again, other time bounds for *NEXPT* machines can be treated in the same way). For each input x of length n, let

$$\pi(x) = \#w_0\#w_1\#\cdots\#w_m\# \quad (\text{where } m = 2^n)$$

represent a computation path of length 2^n where each configuration w_i has length 2^n. Note: we want the w_i's to represent the contents in tape cells whose absolute index (i.e., address) is $-2^n + 1$ to $2^n - 1$, so each symbol of w_i has as many tracks as the number of tapes of M, and each track encodes two tape symbols, etc. We will describe two $\{+,\cdot,^2\}$-expressions $E^1(x)$ and $E^2(x)$ corresponding to the input word x such that $L(E^1(x)) \neq L(E^2(x))$ iff x is accepted by M. The expression $E^1(x)$ is very similar to $\hat{E}(x)$ in the above proof for *EXPS*. In fact, $E^1(x)$ is obtained in two steps:

[5] An interesting remark is that the role of Kleene-star in transforming a problem complete for a time-class to a corresponding space-class is linked to the ability of forming 'transitive closure'. (Cf. [5])

(1) Since the Kleene-star operator is no longer allowed, we replace each subexpression of the form S^* (where $S = \Delta, \Gamma$, etc) with $(S \cup \{\lambda\})^{2^{2n+1}}$. The reason for the exponent '$2n+1$' is because the string $\pi(x)$ has length $(2^n+1)^2 < 2^{2n+1}$ for $n \geq 2$. We write S^{2^k} here just as a shorthand for k applications of the squaring operation. Thus the size of S^{2^k} is $O(k)$. Since the original expression $\hat{E}(x)$ has a fixed number of Kleene-star operators, the replacements maintain the linear size of the original expression.

(2) Let $\hat{E}^1(x)$ be the expression after the replacements of (a). The reader may verify that $\hat{E}^1(x)$ satisfies:

(i) Each word w in $L(\hat{E}^1(x))$ has length $\leq 5 \cdot 2^{2n+1} + 4$. (This maximum length is achieved by the modified subexpression E_2 in $E(x)$.)

(ii) A word $w \in \Delta^*$ of length $\leq 2^{2n+1}$ is in $L(\hat{E}^1(x))$ if and only if it does not represent an accepting computation of M on input x.

We obtain $E^1(x)$ as $\hat{E}^1(x) + F$ where $L(F)$ denotes all words w of length $2^{2n+1} < |w| \leq 5 \cdot 2^{2n+1} + 4$. Clearly F can be written as

$$[\Delta^{2^{2n+1}+1}]_{sq} \cdot [(\Delta + \lambda)^{2^{2n+3}+3}]_{sq}.$$

Our description of $E^1(x)$ is complete. Let $E^2(x)$ express the set of words of length at most $5 \cdot 2^{2n+1} + 4$. Therefore $L(E^1(x)) \neq L(E^2(x))$ iff x is accepted by M. Clearly $E^1(x)$ and $E^2(x)$ have linear size by our remarks. This concludes the proof that INEQ$(+, \cdot, ^2)$ is *NEXPT*-complete.

5.5.4 Other Problems

The following chart summarizes the complexity of some word problems. Most of these results are from Meyer and Stockmeyer. In the chart, we say a class K is a 'lower bound' for a problem L to mean that L is K-hard. It is intuitively clear why such a result is regarded as a 'lower bound'; in the next chapter we show that such a result can be translated into explicit lower bounds on the complexity of any Turing machine accepting L.

Problem	Lower Bound	Upper Bound
INEQ$(+, \cdot, \neg)$	*DSPACE*$(\exp(\log n))$	*DSPACE*$(\exp(n))$
FULL$(+, \cdot, ^2, ^*)$	*EXPS*	*EXPS*
INEQ$(+, \cdot, ^2)$	*NEXPT*	*NEXPT*
FULL$(+, \cdot, ^*)$	*NSPACE*(n)	*NSPACE*(n)
INEQ$(+, \cdot)$	*NP*	*NP*
INEQ$(\{0\}, \{+, \cdot, ^2, \neg\})$	*PSPACE*	*PSPACE*
INEQ$(\{0\}, \{+, \cdot, ^*\})$	*NP*	*NP*
INEQ$(\{0\}, \{+, \cdot, \neg\})$	*P*	*P*

Notes:

(i) In row 1 we refer to the super-exponential function $\exp(n)$ defined in the appendix of chapter 4. The next section considers a closely related result.

(ii) We have so far assumed an alphabet Σ of size at least two; the last three rows of this table refers to unary alphabets.

(iii) With the exception of row 1, all the upper and lower bounds agree. In other words, each language L above is K-complete for its class K.

(iv) The lower bounds for the languages FULL$(+, \cdot, ^2, ^*)$, INEQ$(+, \cdot, ^2)$ and FULL$(+, \cdot, ^*)$ have just been given in sections 4 and 5. It is easily observed that the many-one reductions in these proofs all use transformations t that are linearly bounded i.e. $|t(x)| = O(|x|)$ for all x. This fact will be used in chapter 6.

Other problems based on games that are complete for exponential time or space can be found in [6].

5.6 Elementary Problems

5.6.1 The power of negation

Recall the super-exponential function $\exp(n, m)$ given in the appendix of chapter 4.

Definition 1. A language is called *elementary* (or, elementary recursive) if it can be accepted in deterministic space $\exp(k, n)$ for some some integer $k \geq 0$. Let *ELEMENTARY* denote the class of elementary languages. ■

The class of 'elementary functions' was defined by Kalmar [27]; the characterization of the corresponding class of languages in terms of its space complexity (as in the preceding definition) is due to Ritchie [28]. Our goal in this section is to show that if we admit the negation operator, then the complexity of the word problems we study is enormously increased. More precisely, we are concerned with Ω-expressions where $\Omega = \{+, \cdot, ^*, \neg\}$; call these the *regular expressions with negation*. We show that the fullness problem for such expressions are as hard as any elementary problem. The first natural problem shown to be *ELEMENTARY*-hard is the decision problem for the so-called *weak monadic second order theory of one successor* (WS1S), due to Meyer [22].[6] Meyer showed that the emptiness problem for the so-called *gamma expressions* (γ-expressions) can be efficiently reduced to WS1S. We shall define gamma expressions[7] as those that use the operators

$$\cdot, +, ^*, \neg, \gamma$$

where only the last operator, γ, needs explanation. γ is a unary operator such that for any γ-expression α, the γ-expression $\gamma(\alpha)$ denotes the set

$$L(\gamma(\alpha)) = \{w : (\exists x)[x \in L(\alpha) \text{ and } |w| = |x|]\}.$$

Recall from the last section that the set Σ^k is called a ruler of length k; if we can give very succinct γ-expressions to denote very long rulers then the appropriate word problem for the γ-expressions is proportionally hard to decide. We now describe the ideas of Meyer for describing very long rulers by exploiting negation and γ. We introduce a special Turing machine just for the purposes of our proofs.

Definition 2. A *counting machine* is a deterministic simple Turing machine M_c with tape alphabet $\{0, 1, \&\}$, such that on input $\&x\&$, x a binary word of length $n \geq 1$, the tape head will remain within the left and right boundary markers '&' and will loop while performing the following repeatedly: treating the contents between the markers as a binary number between 0 and $2^n - 1$, the machine successively increments the binary number modulo 2^n. ■

One realization of such a machine has three states q_0, q_1 and q_2, with the following deterministic transition table:

δ	0	1	&
q_0	$-$	$-$	$(q_1, \&, +1)$
q_1	$(q_1, 0, +1)$	$(q_1, 1, +1)$	$(q_2, \&, -1)$
q_2	$(q_1, 1, +1)$	$(q_2, 0, -1)$	$(q_0, \&, 0)$

The table rows and columns are labeled by states q and symbols b, respectively. An entry (q', b', d) in the q-row and b-column says if the machine is in state q scanning symbol b, it next enters state q', changes symbol b to b' and moves its head in the direction indicated by d. The first row of this table for state q_0 (the start state) is only defined for the input symbol '&'. We assume that the input string has the form $\&x\&$ where $x \in \{0, 1\}^*$ and the machine is started in state q_0 scanning the leftmost symbol '&'. The behavior of this machine is easily described:

(i) The machine in state q_1 will move its head to the right until it encounters the right marker '&'; during this movement, it does not modify the contents of the tape. When it sees '&', it enters state q_2 and reverses direction.

(ii) The machine only enters state q_2 while at the right marker in the way indicated in (i); in state q_2, the machine will move its head leftward until it finds the first 0 or, if there are no 0's, until the left marker '&'. During the leftward motion, it changes the 0's into 1's. If it finally encounters the symbol 1, it changes it to a 0, reverses direction and enters state q_1. If it encounters the symbol & instead, it does not change the symbol but enters state q_0 while remaining stationary. Entering state q_0 indicates the start of a new cycle.

[6] This is the problem of deciding the validity of second order formulas where the second order variables vary over finite sets and the only non-logical symbol is the successor relation.

[7] Meyer's construction does not require negation or Kleene-star. We use these to simplify our illustration.

Notation and convention. From now on, when we say 'the counting machine' we refer to the particular machine just described. Let Σ_c be the alphabet for encoding computations of the counting machine, consisting of the thirteen symbols:

$$\#, \&, 0, 1, [q_j, 0], [q_j, 1], [q_j, \&]$$

where $j = 0, 1, 2$. (The symbol # will not be used until later.) If the non-blank portion of the tape of a simple Turing machine is a string of the form $w_1 b w_2$ (where $w_1, w_2 \in \{\&, 0, 1\}^*$ and b is the currently scanned symbol) and the machine is in state q, we encode this configuration as $C = w_1[q, b]w_2$. Also, whenever we write $C \vdash C'$ we assume the lengths of the encodings of the C and C' are equal: $|C| = |C'|$. We will assume that the machine is normally started with tape contents $\&0^n\&$ for some $n \geq 1$. As usual, we assume that the head of the machine is initially scanning the leftmost symbol of the input; the string encoding this initial configuration is denoted by

$$init(n) = [q_0, \&]0^n\&.$$

Using the ideas of the last two sections, we can easily construct a regular expression α_n that denotes all strings w that *fail* to satisfy at least one of the following:

(a) w has the form

$$w = C_0 C_1 \cdots C_m \quad (\text{for some } m \geq 0)$$

where the first and last symbol of each C_i is & or $[q_i, \&]$.

(b) C_0 is the configuration $init(n)$.

(c) C_i is the successor of configuration C_{i-1} $(i = 1, \ldots, m)$, and C_0 is the successor of C_m.

(d) All the C_i's are distinct.

Note that each C_i marks its own boundaries with the symbol &. The size of α_n, properly expressed, would be $O(n)$. Observe that a word w satisfying (a)-(d) has length at least $n2^n$ since the number of configurations m is at least 2^n ($O(2^n)$ is also an upper bound on m – see Exercises). Since M_c is deterministic, the complement of $L(\alpha_n)$ denotes a single word. The availability of the γ operator then gives us the expression $\gamma(\neg\alpha_n)$ which denotes a ruler of length $\geq n2^n$.

In order to recursively apply this idea to get longer rulers, we can further show:

(γ) There is a constant $c > 0$ such that, given an γ-expression ρ_n denoting a ruler of length n, we can efficiently construct a γ-expression ρ_m (for some $m \geq 2^n$) denoting a ruler of length m such that $|\rho_m| \leq c|\rho_n|$.

Hence, after k applications of this result, we can get a γ-expression of size $O(c^k n)$ denoting a ruler of length at least $\exp(k, n)$. Using such an expression, we can then describe Turing machine computations where each configuration in the path uses space $\exp(k, n)$. If the γ operator is not allowed, then it is less clear that the recursive construction can be carried out. The remainder of this section shows how this can be done, as shown by Stockmeyer [32].

For our result below, the concept of negation-depth is important. If we view an expression as a tree whose internal nodes are operands and leaves are atoms, then the negation-depth is the maximum number of negations encountered along any path of this tree. More precisely, the *negation-depth* (or ¬-depth) of an extended regular expression is recursively defined as follows: the ¬-depth of an atom is zero; the ¬-depth of

$$(\alpha + \beta), (\alpha \cdot \beta), (\alpha \cap \beta)$$

is the maximum of the ¬-depths of α and β; the ¬-depth of

$$\alpha^2, \alpha^*$$

is equal to the ¬-depth of α; finally the ¬-depth of $\neg\alpha$ is one plus the ¬-depth of α.

5.6.2 Homomorphism, rotation, smearing and padding

In the remainder of this section, 'expression' shall mean 'regular expression with negation' over a suitable alphabet Σ that varies with the context. The present subsection discusses technical tools to facilitate our rendition of Stockmeyer's construction. The reader should recall the notion of letter-homomorphisms (appendix in chapter 2) since it is the basis of many constructions.

Let Σ_1, Σ_2 be alphabets. In the following it is convenient to distinguish two ways to form composite alphabets: *horizontal composition*

$$[\Sigma_1 \times \Sigma_2] = \{[a,b] : a \in \Sigma_1, b \in \Sigma_2\}$$

and *vertical composition*

$$\left[\begin{array}{c}\Sigma_1\\ \Sigma_2\end{array}\right] = \{\left[\begin{array}{c}a\\ b\end{array}\right] : a \in \Sigma_1, b \in \Sigma_2\}.$$

A string in $\left[\begin{array}{c}\Sigma_1\\ \Sigma_2\end{array}\right]^*$ can be regarded as the contents of a two-tracked tape. Define the functions h_1 and h_2 that respectively extract the first and second component of a horizontally composite symbol: $h_i([b_1,b_2]) = b_i$, $i = 1,2$. These functions then extend in the usual manner to the letter-homomorphisms (still denoted by h_i):

$$h_i : [\Sigma_1 \times \Sigma_2]^* \to \Sigma_i^*.$$

We shall also use the inverse map

$$h_i^{-1} : \Sigma_i^* \to 2^{[\Sigma_1 \times \Sigma_2]^*}$$

where for any set X, we use 2^X to denote the set of subsets of X, and $h_i^{-1}(w)$ is defined to be the set of $x \in [\Sigma_1 \times \Sigma_2]^*$ such that $h_i(x) = w$.

For any homomorphism $h : \Sigma^* \to \Gamma^*$, and any language $L \subseteq \Sigma^*$, we call $h(L)$ the *h-image* of L. If $L' \subseteq \Gamma^*$, and $h(L') = L$ then we call L' an *h-preimage* (or *h-pre-image*) of L. Note that in general, an h-preimage of L is only a subset of $h^{-1}(L)$; we may call $h^{-1}(L)$ the *full* preimage of L.

Similarly, h_U and h_L extract the upper and lower components of a vertically composite symbol: $h_U(\left[\begin{array}{c}a\\ b\end{array}\right]) = a$, $h_L(\left[\begin{array}{c}a\\ b\end{array}\right]) = b$. Again we have the corresponding letter homomorphisms and their inverses. Since horizontal and vertical compositions have only pedagogical differences, we only need to state properties for just one of them. The following simple lemma will be one of our basic tools:

LEMMA 13. *Suppose $h : \Gamma^* \to \Sigma^*$ is a letter-homomorphism. There is a log-space computable transformation t such that for any expression α over the alphabet Σ, we can construct another expression $t(\alpha)$ over Γ such that $L(t(\alpha)) = h^{-1}(L(\alpha))$. Furthermore, the size of $t(\alpha)$ is at most $c|\alpha|$ for some c that only depends on the size of Γ.*

Proof. Replace each atom $b \in \Sigma$ occurring in α with the subexpression $b_1 + b_2 + \cdots + b_t$ where $h^{-1}(b) = \{b_1, \ldots, b_t\}$. **Q.E.D.**

We describe three operations on strings:

Definition 3. Let $x = b_1 b_2 b_3 \cdots b_m$ for some $m \geq 1$, and each $b_j \in \Sigma$ for some alphabet Σ.

(1) Then $rotate(x)$ denotes the set of strings over Σ of the form $x_1 x_2$ ($x_i \in \Sigma^*$) such that $x_2 x_1 = x$.

(2) If $m \geq 3$, then $smear(x)$ denotes the string

$$[x_m, x_1, x_2][x_1, x_2, x_3][x_2, x_3, x_4] \cdots [x_{m-2}, x_{m-1}, x_m][x_{m-1}, x_m, x_1].$$

Note that $smear(x)$ is over the alphabet $[\Sigma \times \Sigma \times \Sigma]$ and $|smear(x)| = |x|$. Also, $h_2(smear(x)) = x$ where h_2 is the homomorphism on $[\Sigma \times \Sigma \times \Sigma]$ that picks out the second component.

(3) Assume $\# \notin \Sigma$, and $k \geq 0$ is an integer. Then $pad_k(x)$ denotes the set of strings

$$b_1 x_1 b_2 x_2 \cdots x_{m-1} b_m x_m,$$

where each x_i is a string of #'s of length at least k. Note that the x_i's can all have the different lengths. The function *unpad*, which simply deletes all occurrences of '#' in its argument, acts as a partial inverse to pad_k: if w contains no '#' then $unpad(pad_k(w)) = w$.

■

The three operations extend in the natural way to sets of words: for instance, $smear(L) = \{smear(x) : x \in L \text{ and } |x| \geq 3\}$. The rotate and smear operations are seen to commute with each other, but they do not commute with the padding operation. Let

$$\begin{aligned}\Sigma_{3c} &= [\Sigma_c \times \Sigma_c \times \Sigma_c]\\ \Gamma_c &= \Sigma_{3c} \cup \{\#\}.\end{aligned}$$

Thus Σ_{3c} (resp., Γ_c) is the alphabet for representing smeared computations (resp., padded smeared computations).

5.6.3 Free computation and cyclic ruler

A key idea in Stockmeyer's construction is to modify the notion of rulers that we have used until now.

Definition 4. Recall the alphabet Σ_c for encoding computations for the counting machine M_c. A *free computation (of order n)* ($n \geq 1$) is a string over Σ_c of the form

$$\bar{C} = C_0 C_1 \cdots C_m \quad (m \geq 0)$$

such that

(1) Each C_j is (an encoding of) a configuration of the counting machine M_c, $j = 0, \ldots, m$.

(2) $C_0 = [q_0, \&]0^n\& = init(n)$.

(3) $C_{i-1} \vdash C_i$ for $i = 1, \ldots, m$, and $C_m \vdash C_0$.

■

Note that the boundary between consecutive configurations is naturally marked by '&&' (or variants containing state information). Observe that a free computation can be arbitrarily long since the configurations can repeat indefinitely. A *cycle* (of order n) is defined to be a free computation $\bar{C} = C_0 \cdots C_m$ such that all the C_i's are distinct and $C_0 = init(n)$. It follows that for each n, the cycle of order n is unique. Define the number-theoretic function g such that a cycle of order n is a string of length $g(n)$. Clearly $g(n) \geq n2^n$ and by earlier remarks, in fact $g(n) = \Theta(n2^n)$. Also the length of a free computation of order n is a multiple of $g(n)$. We use the following important notations:

$free(n)$	=	$\bigcup\{rotate(smear(x)) : x \text{ is a free computation of order } n\}$
$non\text{-}free_k(n)$	=	$\{x : x \notin pad_k(free(n))\}$
$ruler_k(n)$	=	$pad_k(rotate(smear(x)))$ where x is the cycle of order n.

We call a set of the form $ruler_k(n)$ a *cyclic ruler*. Clearly we intend to use cyclic rulers instead of the usual rulers. Note that $free(n) \subseteq (\Sigma_{3c})^*$ and $non\text{-}free_k(n) \cup ruler_k(n) \subseteq (\Gamma_c)^*$.

We now make a simple but important observation about our construction of the counting machine. Let $u_0 \in \Sigma_{3c}$ denote the special symbol

$$u_0 := [\&, [q_0, \&], 0].$$

LEMMA 14. *Let $w = C_0 C_1 \cdots C_m$ denote a free computation of order $n \geq 1$, and $C_0' C_1' \cdots C_m' = smear(w)$ where C_i' corresponds to C_i in the natural way, $h_2(C_i') = C_i$. For each i, we have: C_i' contains the symbol u_0 if and only if $C_i = init(n)$.*

Thus it is sufficient to use local information (the presence of the symbol u_0) to detect the initial configuration. The next lemma shows that we can obtain expressions for cyclic rulers from expressions for $non\text{-}free_k(n)$ at the cost of increasing the negation-depth by 1.

LEMMA 15. *Let Σ be any alphabet and h a letter homomorphism from Σ to Γ_c.*

(i) *There is a log-space transformation that, given any expression α where $h(L(\alpha)) = non\text{-}free_k(n)$, constructs an expression β where $h(L(\beta)) = ruler_k(n)$.*

(ii) *The length of β satisfies $|\beta| = |\alpha| + O(1)$, and the negation-depth of β is one greater than that of α.*

Proof. A word w in $pad_k(free(n))$ fails to be in $ruler_k(n)$ precisely when w contains more than one occurrence u_0. Hence it is easy to satisfy (i) by letting β be

$$\beta = \neg(\alpha + \Sigma^* \cdot h^{-1}(u_0) \cdot \Sigma^* \cdot h^{-1}(u_0) \cdot \Sigma^*)$$

Clearly (ii) holds. **Q.E.D.**

Thus, given an expression denoting an h-preimage of $non\text{-}free_k(n)$, this lemma shows us how to obtain another expression denoting an h-preimage of $ruler_k(n)$. To get succinct expressions for h-preimage of $non\text{-}free_k(n)$, the next lemma shows how to do this inductively.

LEMMA 16. (Key Construction) *Let Σ be any alphabet and h a letter homomorphism from Σ to Γ_c. Let $\Delta = \begin{bmatrix} \Gamma_c \\ \Sigma \end{bmatrix}$ with the usual homomorphisms $h_L, h_R : \Delta \to \Sigma$ selecting the upper and lower tracks. Let H be the letter homomorphism from Δ to Γ_c given by $H(x) = h(h_L(x))$ for all x.*

(i) *There is a log-space transformation t such that given any expression α where $h(L(\alpha)) = non\text{-}free_k(n)$, constructs the expression $t(\alpha)$ over the alphabet Δ where $H(L(t(\alpha))) = non\text{-}free_{k+1}(g(n))$. Recall the function $g(n) \geq n2_n$.*

(ii) *$|t(\alpha)| = O(|\alpha| \cdot |\Sigma|)$ and the negation-depth of $t(\alpha)$ is one greater than that of α.*

Proof. We construct the expression $t(\alpha)$ to denote all those words $w \in \Delta^*$ that are *not* of the following form:

b_1	$\#$	$\#^{i_1}$	b_2	$\#$	$\#^{i_2}$	b_3	$\cdots$
$h^{-1}(\#)$	$h^{-1}(a_1)$	$h^{-1}(\#^{i_1})$	$h^{-1}(\#)$	$h^{-1}(a_2)$	$h^{-1}(\#^{i_2})$	$h^{-1}(\#)$	$\cdots$

$\cdots$	$\#$	$\#^{i_{m-1}}$	b_m	$\#$	$\#^{i_m}$
$\cdots$	$h^{-1}(a_{m-1})$	$h^{-1}(\#^{i_{m-1}})$	$h^{-1}(\#)$	$h^{-1}(a_m)$	$h^{-1}(\#^{i_m})$

for some $m \geq 4$ where:

(a) $a_j, b_j \in \Sigma_{3c}$, and $i_j \geq k$ for all $j = 1, \ldots, m$.

(b) the bottom track contains a word u such that

$$\begin{aligned} h(u) &= \#x \\ &= \#a_1\#^{i_1+1}a_2\#^{i_2+1}\cdots\#^{i_{m-1}+1}a_m\#^{i_m} \end{aligned}$$

where $x \notin non\text{-}free_k(n)$ (i.e. $x \in pad_k(free(n))$).

(c) the top track contains a word

$$v = b_1\#^{i_1+1}b_2\#^{i_2+1}\cdots\#^{i_{m-1}+1}b_m\#^{i_m+1}$$

where $v \notin non\text{-}free_{k+1}(g(n))$ (i.e., $v \in pad_{k+1}(free(n))$).

Note that each b_i in the upper track is immediately followed by a_i in the lower track: we say b_i and a_i are 'paired'. By definition, only non-# symbols can be paired.[8] The expression $t(\alpha)$ is over the alphabet $\Delta = \begin{bmatrix} \Gamma_c \\ \Sigma \end{bmatrix}$. We shall write $t(\alpha)$ as the union $E_1 + E_2 + \cdots + E_8$. We now begin the lengthy details of this construction. A word w is in $t(\alpha)$ if it satisfies at least one of the following conditions.

(1) "The upper track of w is incoherent"
Say a string over Γ_c is *coherent* if it is of the form $pad_0(smear(x))$ for some $x \in \Sigma_c^*$. Say a pair (b, b') of symbols in Σ_{3c} is *incoherent* if they could not possibly be consecutive symbols in the smearing of some word over Σ_c. More precisely, with $h_i : \Sigma_{3c} \to \Sigma_c$ $(i = 1, 2, 3)$ the homomorphism that extracts the ith component of its argument, we say (b, b') is incoherent precisely when $h_2(b) \neq h_1(b')$ or $h_3(b) \neq h_2(b')$. Then the following expression denotes those words whose upper track is incoherent because of a consecutive pair of non-# symbols:

$$E_1' = \Delta^* \cdot h_U^{-1}\left(\sum_{(b,b')} b \cdot \#^* \cdot b'\right) \cdot \Delta^*$$

where the summation (denoting set union) is over all incoherent pairs (b, b'). Note that here and subsequently, for any expression σ over some Σ, and homomorphism $f : \Sigma \to \Sigma'$, we write $h^{-1}(\sigma)$ as a short-hand for the expression that denotes the set $h^{-1}(L(\sigma))$, as shown in lemma 13. The upper track can also be incoherent if the first and last non-# symbols do not agree:

$$E_1'' = \sum_{(b,b')} \left(h_U^{-1}(\#^* \cdot b') \cdot \Delta^* \cdot h_U^{-1}(b \cdot \#^*)\right)$$

Our first subexpression is given by $E_1 = E_1' + E_1''$.

[8] It may appear preferable to pair a_i with b_i by placing a_i directly above b_i instead of the 'staggered' fashion used here; this is possible though it causes some inconvenience elsewhere.

(2) "the h-image of the lower track of w is not of the form $\#x$ where $x \in pad_k(free(n))$"

$$E_2 = h_L^{-1}(\Sigma_c) \cdot \Delta^* + h_L^{-1}(\#) \cdot h_L^{-1}(\alpha)$$

The first summand of E_2 captures those words w whose h-image of the lower track does not begin with a #-symbol.

(3) "Some pairing of non-# symbols in the upper and lower tracks is violated"
That is, there is a consecutive pair of symbols $\cdots bb' \cdots$ occurring in w such that either $h_U(b) = \#$ and $h(h_L(b')) \neq \#$ or $h_U(b) \neq \#$ and $h(h_L(b')) = \#$.

$$E_3 = \Delta^* \cdot \left(h_U^{-1}(\#) \cdot h_L^{-1}(h^{-1}(\Sigma_{3c}))) + h_U^{-1}(\Sigma_{3c}) \cdot h_L^{-1}(h^{-1}(\#))\right) \cdot \Delta^*.$$

Note that we had to apply h^{-1} to symbols that are in the lower track by virtue of the homomorphism h.

(4) "Some right marker of configurations in the upper track is not aligned"
This condition corresponds to our intention that the upper track (before padding and rotation) represents a smeared computation path $C_0'C_1' \cdots C_h'$ where each smeared configuration $C_i' \in (\Sigma_{3c})^*$ is *aligned* in the sense that the rightmost symbol in each C_i' is paired with some symbol in $h^{-1}(u_0)$ in the lower track. Conversely, each $h^{-1}(u_0)$ in the lower track must pair with a rightmost symbol of some C_i'. Thus each C_i', if they are all aligned, has length $g(n)$. To express this, we define the set of 'right marker symbols' of the upper track:

$$RM = \{[b, \&, \&], [b, [q, \&], \&], [b, \&, [q, \&]] : b \in \Sigma_c, q \in \{q_0, q_1, q_2\}\}.$$

The expression becomes

$$E_4 = \Delta^* \cdot \left(h_U^{-1}(RM) \cdot h_L^{-1}(h^{-1}(\overline{u_0}))) + h_U^{-1}(\overline{RM}) \cdot h_L^{-1}(h^{-1}(u_0))\right) \cdot \Delta^*.$$

Here, $\overline{RM} = \Sigma_{3c} - RM$ and $\overline{u_0} = \Sigma_{3c} - \{u_0\}$. Let us introduce the analogous set of 'left marker symbols':

$$LM = \{[\&, \&, b], [\&, [q, \&], b], [[q, \&], \&, b] : b \in \Sigma_c, q \in \{q_0, q_1, q_2\}\}.$$

Let

$$NM = \Gamma_c - (LM \cup RM)$$

be the set of 'non-marker symbols'.

(5) "There is some stray &-symbol in the upper track"
This expresses our intent that between the &-symbols, the upper track contains only 0's and 1's, possibly with state information. Define the set $Z \subseteq \Sigma_{3c}$ consisting of symbols of the form:

$$[b, \&, b'], [b, [q, \&], b']$$

where $b, b' \in \Sigma_c$ and $q \in \{q_0, q_1, q_2\}$. The desired condition is given by

$$E_5 = \Delta^* \cdot h_U^{-1}\left(LM \cdot NM^* \cdot Z \cdot NM^* \cdot RM\right) \cdot \Delta^*.$$

Note that $Z \subseteq NM$, so E_5 simply asserts that at least one &-symbol occurs between the left and right markers.

Before continuing, let us observe that for any word w not in $L(E_1 + \cdots + E_5)$, the upper track of w must be of the form

$$pad_{k+1}(rotate(smear(C_0C_1 \cdots C_m))) \tag{5}$$

where each C_i (after removing the state information) encodes a binary string delimited by two &-symbols, with $|C_i| = g(n)$. We also write

$$C_0'C_1' \cdots C_m' \tag{6}$$

for $smear(C_0 \cdots C_m)$ where C_i' corresponds naturally to C_i. For the remaining subexpressions, we will refer to these 'configurations' C_i and C_i'.

(6) "Some C_i does not have a unique state symbol"
We leave this to the reader.

(7) "Some C_{i+1} is not directly derived from C_i"
A pair (b, b') in Σ_{3c} is *compatible* if there are smeared configurations C, C' such that $C \vdash C'$, b occurs in C, and b' occurs in the corresponding position in C. For example, if $b = [b_1, [q, b_2], b_3]$ and $b' = [[q', b'_1], b'_2, b'_3]$ then (b, b') is compatible iff $b_1 = b'_1$, $b_3 = b'_3$, and $\langle q, b_2, q', b'_2, -1\rangle$ is in the transition table of the counting machine. We come to the place where the recursive nature of our construction is seen. We make the following observation:

> Suppose $w \in \Delta^* - L(E_1 + \cdots + E_6)$ is a word whose upper track $h_U(w)$ has the form (5). If $\cdots bxb' \cdots$ occurs in $h_U(w)$ where $h_U(b')$ is a non-# symbol. Then $h_L(x)$ is in $ruler_k(n)$ implies that $(h_U(b), h_U(b'))$ is a compatible pair of symbols.

Thus the cyclic ruler 'measures' the distance between corresponding symbols. By the previous lemma, we can construct from α the expression β where $h(L(\beta)) = ruler_k(n)$. It is important to realize that this observation depends on the way we pair non-# symbols in the tracks in a 'staggered' fashion. We define

$$E'_7 = \Delta^* \cdot \left(\sum_{(b,b')} h_U^{-1}(b) \cdot h_L^{-1}(\beta) \cdot h_U^{-1}(b') \right) \cdot \Delta^*.$$

where the summation ranges over all pairs (b, b') that are not compatible. Because of the effect of rotation, a pair of corresponding symbols that must be compared for compatibility may appear at the opposite ends of the string w. To handle this wrap-around, we also observe:

> Suppose $w \in \Delta^* - L(E_1 + \cdots + E_6)$ is a word whose upper track has the form (5). If in addition,
>
> $$h_U(w) \in NM^* \cdot RM \cdot LM \cdot NM^* \cdot b' \cdot w' \cdot b \cdot NM^*$$
>
> where $h_U(b)$ and $h_U(b')$ are non-# symbols. Then $h_U(w') \in pad_k(free(n))$ (rather than $ruler_k(n)$!) implies that $(h_U(b), h_U(b'))$ forms a compatible pair.

To understand this observation, let b, b' be two corresponding symbols from a pair of successive configurations C, C' in the upper track of w. We need to check b and b' for compatibility. Assume that the free computation is rotated so that C splits into two parts forming a prefix and a suffix of w, respectively. If b is in the suffix part, then b' is left of b and the word w can be expressed as $p \cdot b' \cdots w' \cdot b \cdot s$ where p, w', s are words in Δ^*. Furthermore, w' must be a rotated and padded free computation, and there must be exactly one [rightmarker, leftmarker] pair in the upper track of p.

Depending on how the pair C, C' is split by the wrap-around, we may get two other possibilities:

$$h_U(w) \in NM^* \cdot b' \cdot w' \cdot b \cdot NM^* \cdot RM \cdot LM \cdot NM^*$$

and

$$h_U(w) \in LM \cdot NM^* \cdot b' \cdot w' \cdot b \cdot NM^* \cdot RM.$$

These lead to the expression

$$E''_7 = \sum_{(b,b')} (F_1(b, b') + F_2(b, b') + F_3(b, b'))$$

where

$$F_1(b, b') = h_U^{-1}(NM^* \cdot RM \cdot LM \cdot NM^* \cdot b') \cdot h_L^{-1}(\neg\alpha) \cdot h_U^{-1}(b \cdot NM^*)$$

$$F_2(b, b') = \cdot h_U^{-1}(NM^* \cdot b') \cdot h_L^{-1}(\neg\alpha) h_U^{-1}(b \cdot NM^* \cdot RM \cdot LM \cdot NM^*)$$

$$F_3(b, b') = h_U^{-1}(LM \cdot NM^* \cdot b') \cdot h_L^{-1}(\neg\alpha) \cdot h_U^{-1}(b \cdot NM^* \cdot RM)$$

We finally let $E_7 = E'_7 + E''_7$.

(8) "None of the C_i's are equal to $init(g(n))$"
It is sufficient to assert that either the upper track of w does not contain the symbol u_0 or else the symbol '1' appears in the configuration containing the u_0. This will ensure that no configuration C_i has the form $[\&, [q_0, \&], \&]0^m\&$; of course, this includes the case $m = g(n)$ which is all that we are interested in. Let $Y \subseteq \Sigma_{3c}$ be the set of symbols that contain '1' or '$[q, 1]$' as its second component. Let $E_8 = E'_8 + E''_8$ where

$$E'_8 = \Delta^* \cdot h_U^{-1}\left(LM \cdot (\overline{u_0}^*) + NM^* \cdot Y \cdot NM^*) \cdot RM\right) \cdot \Delta^*$$

and E''_8 is a similar expression to take care of wrap-around.

This concludes our construction of α. To verify (ii), note that each E_i, with the exception of E_7, has size $O(|\Delta|) = O(|\Sigma|)$. Also E_7 has size $O(|\alpha| \cdot |\Delta|)$. Similarly, the $\neg$-depth of each E_7 is one more than that of α while the other E_i has depth 0. Hence (ii) follows. This concludes our proof of lemma 16. **Q.E.D.**

Let us call our ruler construction a $g(n)$-construction, since we obtained a ruler of order $g(n)$ from one of order n. See Exercises for attempts to get a $g'(n)$-construction for faster growing functions $g'(n)$. An elegant alternative to cyclic rulers is given by Hunt [2]; however Hunt's construction requires the use of the intersection operator '$\cap$'.

5.6.4 The main result

Using the function $g(n)$ in the last section we now define the function $G(k, n)$ for $k, n \geq 0$ as follows: $G(0, n) = g(n)$ and $G(k, n) = g(G(k-1, n))$ for $k \geq 1$. Clearly $G(k, n)$ is the analog of the function $\exp(k, n)$ and $G(k, n) \geq \exp(k, n)$ for all k, n. We can apply the last lemma k times to describe the complement of free-computations (and hence, cyclic rulers) of order $G(k, n)$:

LEMMA 17. *There is a constant $c > 0$ such that for all integers $k \geq 0, n \geq 1$, there is an expression $\alpha_{k,n}$ over a suitable alphabet Δ_k and a homomorphism $h : \Delta_k \to \Gamma_c$ such that*

(i) $h(L(\alpha_{k,n})) = \textit{non-free}_k(G(k, n))$.

(ii) $\alpha_{k,n}$ *can be computed in space* $O(\log |\alpha_{k,n}|)$

(iii) $|\alpha_{k,n}| = n \cdot O(1)^{k^2}$. *The $\neg$-depth of $\alpha_{k,n}$ is k and* $|\Delta_k| = |\Gamma_c|^k$.

Proof. If $k = 0$ then we can easily construct $\alpha_{0,n}$, by modifying the proof that the fullness problem for regular languages is complete for *LBA*. If $k > 0$ then the last lemma shows how to construct $\alpha_{k,n}$ from $\alpha_{k-1,n}$. The alphabet Δ_k is $\begin{bmatrix} \Gamma_c \\ \Delta_{k-1} \end{bmatrix}$, so $|\Delta_k| = |\Delta_{k-1}| \cdot |\Gamma_c| = |\Gamma_c|^k$. The properties (ii) and (iii) easily follow from the construction. **Q.E.D.**

We next apply the cyclic rulers of length $G(k, n)$ to describe accepting computations of Turing acceptors that accept in space $G(k, n)$:

LEMMA 18. *Fix any nondeterministic acceptor* M *that accepts in space $G(k, n)$. There is a suitable alphabet Γ_{M} such that for all input words x of length n, we can construct an expression $\alpha_{\mathrm{M}}(x)$ over Γ_{M} such that*

(i) x *is accepted by* M *iff* $\alpha_{\mathrm{M}}(x) \neq \mathrm{FULL}(\Gamma_{\mathrm{M}}, \{+, \cdot, ^*, \neg\})$.

(ii) $|\alpha_{\mathrm{M}}(x)| = n \cdot O_{\mathrm{M}}(1)^{k^2}$ *and $\alpha_{\mathrm{M}}(x)$ can be constructed in space* $O_{\mathrm{M}}(\log |\alpha_{\mathrm{M}}(x)|)$.

(iii) *The $\neg$-depth of $\alpha_{\mathrm{M}}(x)$ is k.*

Proof. Let Δ_k be the alphabet of the expression $\alpha_{k,n}$ in the last lemma, and Σ be the alphabet to represent the smeared computations of M. Also, let $\Delta_{k,\#} = \Delta_k \cup \{\#\}$ and $\Sigma_\# = \Sigma \cup \{\#\}$. The expression $\alpha_{\mathrm{M}}(x)$ shall denote strings over the alphabet

$$\Gamma_{\mathrm{M}} := \begin{bmatrix} \Sigma_\# \\ \Delta_{k,\#} \end{bmatrix}.$$

A string w is in $\alpha_{\mathrm{M}}(x)$ iff it does *not* have the form:

b_1	$\#$	$\#^k$	b_2	$\#$	$\#^k$	b_3	$\#$	$\#^k$	$\cdots$	b_{m-1}	$\#$	$\#^k$	b_m	$\#$	$\#^k$
$\#$	a_1	$\#^k$	$\#$	a_2	$\#^k$	$\#$	a_3	$\#^k$	$\cdots$	$\#$	a_{m-1}	$\#^k$	$\#$	a_m	$\#^k$

where

(a) the lower track has the form $\#y$ where y is in $\textit{pad}_k(\text{non-free}_k(G(k-1, n)))$,

(b) the upper track encodes a padded, smeared (but not rotated) computation of M on input x,

(c) the left marker for each configuration on the upper track is paired with the u_0 marker on the lower track,

(d) the last configuration on the upper track is accepting.

Note that (c) ensures that the configurations in the upper track have length exactly $G(k,n)$ since the lower track is a free computation of order $G(k-1,n)$. The technique for writing down $\alpha_{\mathrm{M}}(x)$ should by now be routine. We leave this as an exercise. Hence if x is accepted by M iff there is an accepting computation iff $L(\alpha_{\mathrm{M}}(x)) \neq (\Gamma_{\mathrm{M}})^*$. **Q.E.D.**

In section 4, we show that for any regular expression α, we can construct a nondeterministic finite automaton δ^α with $\leq |\alpha|$ states, $|\alpha| \geq 2$. A well-known construction of Rabin and Scott can be applied to give a deterministic finite automaton Δ^α with $\leq 2^{|\alpha|}$ states and which accepts the same language as δ^α: the states of the deterministic automaton are sets of states of δ^α, and transitions of Δ^α are defined in a straightforward manner: if X is a set of Δ^α (so X is a set of states of δ^α) then on input symbol b, the next state of Δ^α is

$$X' = \{q' : (\exists q \in X)\ q' \text{ is the next state of } q \text{ on input } b\}.$$

We use this construction to show:

LEMMA 19. *Let α be a regular expression with negation, $|\alpha| \geq 3$.*

(i) *If the negation-depth of α is k then there is a nondeterministic finite automaton δ^α with $\exp(k, |\alpha|)$ states that accepts $L(\alpha)$.*

(ii) *Furthermore, the transition table of the automaton can be constructed in $\exp(k, |\alpha|)$ space.*

Proof. (Basis) If the negation-depth k is 0 then the result follows from our construction in lemma 8.

(Inductively) Suppose $k \geq 1$. First assume $\alpha = \beta + \gamma$. Assuming $|\beta| \geq 3$ and $|\gamma| \geq 3$, then δ^β and δ^γ can be recursively constructed with at most $\exp(k, |\beta|)$ and $\exp(k, |\gamma|)$ states respectively. Then the construction in section 4 applied to δ^β and δ^γ gives δ^α with $\leq \exp(k, |\beta|) + \exp(k, |\gamma|) \leq \exp(k, |\alpha|)$ states. If $|\beta| \leq 2$ or $|\gamma| \leq 2$, the same bound holds as well.

If $\alpha = \beta \cdot \gamma$ or $\alpha = (\alpha)^*$, the arguments are similar.

Finally, suppose $\alpha = \neg\beta$. If $|\beta| \geq 3$, then by induction, we may assume that δ^β has been formed with $\exp(k-1, |\beta|)$ states. We first carry out the Rabin-Scott construction to get a deterministic machine with $\exp(k, |\beta|)$ states. To get δ^α, we can easily modify this deterministic acceptor by interchanging the roles of acceptance and rejection. This modification does not change the number of states, and we are done. If $|\beta| \leq 2$, then a direct argument clinches the proof.

Part (ii) is routine once the constructions of part (i) is understood. **Q.E.D.**

We now state with our main result: For each $k \geq 0$ and alphabet Σ, let $\mathrm{FULL}_k(\Sigma, \{+, \cdot, ^*, \neg\})$ denote the set of regular expressions with negation α where the negation-depth of α is $\leq k$ and $L(\alpha) = \Sigma^*$. As usual, we let $\mathrm{FULL}_k(+, \cdot, ^*, \neg)$ denote the case where Σ is $\{0, 1\}$. As usual, the problem $\mathrm{FULL}_k(\Sigma, \{+, \cdot, ^*, \neg\})$ easily be reduces to the case $\mathrm{FULL}_k(+, \cdot, ^*, \neg)$.

THEOREM 20. *For each $k \geq 0$, the language $\mathrm{FULL}_k(+, \cdot, ^*, \neg)$ is complete for the class NSPACE$(\exp(k, n))$.*

Proof. In the following proof, we exploit the closure of space classes under complementation. Lemma 18 shows that $\mathrm{FULL}_k(+, \cdot, ^*, \neg)$ is co-*NSPACE*$(G(k,n))$-hard, and hence co-*NSPACE*$(\exp(k,n))$-hard. To show that $\mathrm{FULL}_k(+, \cdot, ^*, \neg)$ is in co-*NSPACE*$(\exp(k,n))$, we use the previous lemma: given a regular expression with negation α, it is easy to check whether it has negation depth at most k. We then construct the nondeterministic finite automata δ^α that accepts $L(\alpha)$ in space $\exp(k, |\alpha|)$. Using δ^α it is easy to nondeterministically decide (see the proof for regular expressions) if there is a word not in $L(\alpha)$, i.e., if $\alpha \notin \mathrm{FULL}_k(+, \cdot, ^*, \neg)$. **Q.E.D.**

COROLLARY 21. *he problem $\mathrm{FULL}(+, \cdot, ^*, \neg)$ is hard for the class ELEMENTARY.*

Remarks: Of course, in the presence of negation, we could have posed these results in terms of the problem $\mathrm{EMPTY}(+, \cdot, ^*, \neg)$. Stockmeyer also shows that the use of Kleene-star in these results is not strictly necessary (see Exercises).

5.7 Final Remarks

1. An interesting feature of many of the complete problems for the various canonical classes is that they can be grouped together into natural families. For example, the following problems based on node accessibility in graphs,

1GAP, GAP, AND-OR-GAP, HEX

are complete for $DLOG, NLOG, P, PSPACE$, respectively. Similarly, the problems

2UNSAT, UNIT, SAT, QBF

form a family of complete problems for $NLOG$, P, NP, $PSPACE$ (resp.) based on logical formulas. Alternatively, the first two problems in the above sequence can be replaced by two problems in Boolean logic:

FVP, CVP, SAT, QBF

where FVP is the *formula value problem*. In section 3, CVP is explained; FVP is a similar problem except that we replace Boolean circuits in the input by Boolean formulas. See the exercises which shows FVP to be $NLOG$-complete. Schaeffer and Lieberherr have generalized the satisfiability problem in several other directions. Galil has shown that various decision problems related to restrictions on Turing machines also form a family of problems complete for various classes in the canonical list. Such examples can be multiplied. Such families are useful as a guide to inherent complexity of problems in the 'real' world. This is because natural problems have families resemblances and, given a problem L that resembles members of a family of complete languages, we can often conjecture and confirm the complexity of L relative to the canonical ruler (see chapter 1, section 8).

2. In the family of problems involving graphs, we use directed graphs. A natural question to ask is what is the complexity of undirected graph reachability (denoted UGAP)? It is intermediate in complexity between 1GAP and GAP but is not known to be complete for $NLOG$. We shall return to this problem in a later chapter on probabilistic computation.

3. It is important to emphasize that although we have shown complete languages for many important classes, there is strong evidence that many other natural classes do not have complete languages (e.g., see [15, 31, 16]).

4. There have been interesting developments stemming from investigations about the isomorphism of complete languages. If F is a family of transformations, we say that two languages L and L' are isomorphic (modulo F) if there are transformations $t, t' \in F$ such that $t(L) \subseteq L', t'(L') \subseteq L$ and $t' \circ t$ and $t \circ t'$ are identities. An important case is the Berman-Hartmanis conjecture concerning the isomorphism of all NP-complete languages. This conjecture is confirmed for all known NP-complete languages; on the other hand, there is evidence against the truth of the conjecture. For example, a positive result about isomorphism of complete languages is [3] which shows that all complete languages for $DEXPT$ are equivalent via 1-1 length increasing polynomial time transformations (see also [35, 34]). We shall return to this topic in volume 2.

Exercises

[5.1] (i) Show complete languages for $DTIME(2^{2^{O(n)}})$ and $DSPACE(2^{n2^n})$ under suitable reducibilities.
(ii) Let f be a complexity function. Under what conditions can we say that $XSPACE(f), XTIME(f), XREVERSAL(f$ have complete languages under $\leq_m^L$?
(This exercise suggests that the special place of $\leq_m^L$ is principally derived from the special emphasis we have on the canonical list.)

[5.2] Why does the proof used for the other classes in the canonical list fail to produce a complete class for the class $PLOG$?

[5.3] Consider the structure of the proof for P and some of the other cases in theorem 1. The outline of these proofs could be described as follows: let $(K, (\mu, \rho, F))$ be any characteristic class in the canonical list where μ is a mode, ρ a resource and F a family of complexity functions. Suppose $MACHINES(\mu, \rho, F)$ has an efficient universal simulator U, where each U_i accepts in resource ρ bound $O_i(f_i)$ where $f_i \in F$. Consider the language

$$L^K = \{i\#x\#0^m : m = f_i(|x|) \wedge x \in U_i\}.$$

State the other properties required in order for L^K to be K-complete under $\leq_m^L$ (e.g., F must be 'efficiently presentable' in some sense).

[5.4] Recall the definition of the language 1GAP consisting of all directed graphs (encoded by an edge list) with outdegree 1 and where there is a path from the first node to the last node. Let 1GAP$'$ be the variant where the graphs are encoded by adjacency matrix. Show that it is impossible to reduce the 1GAP$'$ to 1GAP using $\leq_m^{1L}$-reducibility.

[5.5] Show that UGAP (undirected graph accessibility problem, defined in the concluding remarks of the chapter) is $DLOG$-hard under $\leq_m^{1L}$-reducibility.

[5.6] (Laaser) Show that the following connectivity problems on graphs are complete for *NLOG*: define CON and SCON to be the set of inputs of the form $\langle n, G\rangle$ (as in GAP) satisfying (respectively):

$$\forall m(m \in [1..n] \Rightarrow \exists \text{path from node 1to} m)$$

$$\forall m, p(m, p \in [1..n] \Rightarrow \exists \text{path from node} m \text{to} p)$$

[5.7] * Give a counter example to a 'standard' fallacy about resolution: Let C_1 and C_2 be clauses in a CNF formula F and C is their resolvent. Then F is unsatisfiable iff $(F - \{C_1, C_2\}) \cup \{C\}$ is unsatisfiable. For which direction is this true? Provide a counter example to the false direction.

[5.8] Prove the resolution theorem.

[5.9] Complete the proof that 3UNIT is in P by proving the correctness of the algorithm in the text.

[5.10] Complete the proof that 3UNIT is P-complete by replacing those clauses that do not have 3 literals by clauses that have exactly 3.

[5.11] Let M be any nondeterministic acceptor that is $s(n)$ space bounded. For each input x of length n, construct a path system $S(x)$ such that $S(x)$ is solvable iff x is accepted by M. **Hint:** Imitate the proof of Savitch's theorem.

[5.12] (Ladner) Show that the Circuit Value Problem (CVP) is P-complete under log-space many-one reducibility. To show that the problem is P-hard, give a *direct* proof (i.e., instead of reducing a known P-hard problem to CVP).

[5.13] Show the the Formula Value Problem (FVP) is *NLOG*-complete under log-space many-one reducibility.

[5.14] Show that AND-OR-GAP problem is P-complete.

[5.15] (Dobkin, Lipton, Reiss) Show that Rational Linear Programming (RLP) is P-hard under log-space many-one reducibility.

[5.16] Show direct log-space reductions between AND-OR-GAP and MCVP (there are reductions in two directions to do).

[5.17] (Reif) Show that the following decision problem related to the depth-first search algorithm is P-complete under log-space many-one reducibility. *Given*: $\langle G, u, v\rangle$ where G is a digraph over the vertices $[1..n]$, u and v are vertices. *Property*: Node u visited before v in a depth-first search of G that starts from node 1. It is important for this problem to assume that the search always choose the smallest numbered vertex to search next.

[5.18] Show that MEMBER$(+, \cdot, ^2, ^*, \neg)$ is in P.

[5.19] (Meyer, Stockmeyer) There is a family of problems that resembles the word problems of section 4: an *integer expression* involve the operators $+, \cdot, \cap, \cup$, built up recursively in the obvious way from the constants 0 and 1. For example, $((1+1+1)\cdot(1\cup(1+1))+(0\cup 1)$. Each expression now denotes a set of non-negative positive integers, with 0 and 1 denoting the integers zero and one, + denoting addition (not union, as in extended regular expressions!), $\cdot$ denoting product, $\cap$ and $\cup$ denoting intersection and union of sets. Let N-MEMBER$(+, \cup)$ be the membership problem for integer expression. More precisely, this is the problem of recognizing all pairs of the form (x, α) where x is a binary number and α is an integer $\{+, \cup\}$-expression such that x is in $L(\alpha)$. Show that this problem is *NP*-complete.

[5.20] (Meyer, Stockmeyer) Prove that INEQ$(\{0\}, \{+, \cdot, ^*\})$ is *NP*-complete.

[5.21] (Meyer, Stockmeyer) Prove that INEQ$(+, \cdot)$ is *NP*-complete.

[5.22] Give a direct procedure for deciding if $L(\alpha) = \Sigma^*$ in time $O(1)^{|\alpha|}$, avoiding the reduction to a nondeterministic finite automaton as in the text.

[5.23] (Hunt, Hopcroft) Let NONEMPTY(Ω) denote those Ω-expressions α where $L(\alpha) \neq \emptyset$. Note that if $\Omega \subseteq \{+, \cdot, ^*, ^2\}$ then $L(\alpha) \neq \emptyset$ always holds. Show that NONEMPTY$(+, \cdot, ^*, \cap)$ is *PSPACE*-complete.

[5.24] Consider the *equivalence problem* EQUIV(Σ, Ω) that consists of all pairs (α, β) of Ω-expressions such that $L(\alpha) = L(\beta)$. What is the relation between this and the inequivalence problem INEQ(Σ, Ω)?

[5.25] Construct the γ-expressions ρ_m as claimed in (γ) of section 6.

[5.26] (Stockmeyer) Show that the problem of FULL$(+, \cdot, \neg)$ is *ELEMENTARY*-hard. In other words, you have to eliminate the use of Kleene-star from the constructions of section 6.

[5.27] Show that the counting machine makes $O(2^n)$ moves. **Hint:** let c_n be the number of counting machine moves to count from 1 to $2^n - 1$. For example, we have

$$1 \to 10 \to 11 \to 100 \to 101 \to 110 \to 111$$

where the 6 transformations require 4+2+6+2+4+2=20 moves, so $c_3 = 20$. (Here we assume that the head scans the first symbol to the right of the low-order bit of the number, and returns to it after each transformation.) Also $c_1 = 0$, $c_2 = 12$. Give a recurrence for c_n.

[5.28] * The upper and lower bounds for the complexity of FULL$(+, \cdot, ^*, \neg)$ is not tight. Improve the reduction of the elementary problems to this problem. One way to do this is to obtain a '$\hat{g}(n)$-construction' corresponding to the Key lemma where $\hat{g}(n)$ grows faster than the $g(n) = n2^n$. For instance, it is easy to modify our counting machine to obtain a $(n^2 2^n)$-construction (how?). Such an improvement is inconsequential since we could have used k-ary counting and obtained a k^n-construction. What is the optimal choice of k?

[5.29] * Develop a theory of complete transformations. Let us say that a transformation $t_0 : \Sigma_0^* \to \Gamma_0^*$ is *hard for a family T of transformations under log-space reducibilities* if for all $t \in T$, $t : \Sigma_1^* \to \Gamma_1^*$, there are log-space transformations $s_0 : \Sigma_1^* \to \Sigma_0^*$ and $s_1 : \Gamma_0^* \to \Gamma_1^*$ such that for all $x \in \Sigma_1^*$, $s_1(t_0(s_0(x))) = t(x)$. Let T be the class of transformations computable by deterministic polynomial time transducer. Show complete transformations for T.

[5.30] * Little is known about natural complete problems for reversal or simultaneous complexity classes.

Bibliography

[1] A. Adachi, S. Iwata, and T. Kasai. Some combinatorial game problems require $\omega(n^k)$ time. *Journal of the ACM*, 31(2):361–377, 1984.

[2] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The design and analysis of computer algorithms*. Addison-Wesley, 1974.

[3] L. Berman. *Polynomial reducibilities and complete sets*. PhD thesis, Cornell University, 1977. PhD Thesis, Computer Science Department.

[4] R. V. Book. Translational lemmas, polynomial time, and $(\log\ n)^j$-space. *Theor. Computer Science*, 1:215–226, 1976.

[5] R. V. Book. On languages accepted by space-bounded oracle machines. *Acta Informatica*, 12:177–185, 1979.

[6] A. K. Chandra and L. J. Stockmeyer. Alternation. *17th Proc. IEEE Symp. Found. Comput. Sci.*, pages 98–108, 1976.

[7] J. Chen and C.-K. Yap. Reversal complexity. *SIAM J. Computing*, to appear, 1991.

[8] S. A. Cook. An observation of time-storage trade off. *5rd Proc. ACM Symp. Theory of Comp. Sci.*, pages 29–33, 1973.

[9] D. P. Dobkin, R. J. Lipton, and S. Reiss. Linear programming is LOG-SPACE hard for P. *Information Processing Letters*, 8:96–97, 1979.

[10] S. Even and R. E. Tarjan. A combinatorial problem which is complete in polynomial space. *Journal of the ACM*, 23:710–719, 1976.

[11] L. M. Goldschlager. The monotone and planar circuit value problems are log space complete for P. *SIGACT news*, 9(2):25–29, 1977.

[12] L. M. Goldschlager, R. A. Shaw, and J. Staples. The maximum flow problem is log space complete for P. *Theor. Computer Science*, 21:105–111, 1982.

[13] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. *Limits to Parallel Computation: P-Completeness Theory*. Oxford University Press, New York, 1995.

[14] J. Hartmanis. *Feasible Computations and Provable Complexity Properties*. S.I.A.M., Philadelphia, Pennsylvania, 1978.

[15] J. Hartmanis and L. A. Hemachandra. Complexity classes without machines: on complete languages for UP. *Theor. Computer Science*, 58:129–142, 1988.

[16] J. N. Hartmanis and N. Immerman. On complete problems for $NP \cap$ co-NP. *12th ICALP (LNCS No. 194)*, pages 250–259, 1985.

[17] J. N. Hartmanis, N. Immerman, and S. Mahaney. One-way log-tape reductions. *19th Symposium FOCS*, pages 65–71, 1978.

[18] H. B. Hunt, III. On the time and tape complexity of languages, I. *5th Proc. ACM Symp. Theory of Comp. Sci.*, pages 10–19, 1973.

[19] N. D. Jones. Space-bounded reducibility among combinatorial problems. *Journal of Computers and Systems Science*, 11:68–85, 1975.

[20] N. D. Jones and W. T. Laaser. Complete problems for deterministic polynomial time. *Theoretical Comp. Sci.*, 3:105–117, 1977.

[21] R. E. Ladner. The circuit value problem is log space complete for *P*. *SIGACT News*, 7(1):18–20, 1975.

[22] A. R. Meyer. Weak monadic second order theory of successor is not elementary-recursive. In Dold and E. (eds.), editors, *Logic Colloquium: Symposium on Logic Held at Boston University, 1972-73*, pages 132–154. Springer-Verlag, 1975.

[23] A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. *13th Proc. IEEE Symp. Found. Comput. Sci.*, pages 125–129, 1972.

[24] B. Monien and I. H. Sudborough. On eliminating nondeterminism from Turing machines which use less than logarithm worktape space. In *Lecture Notes in Computer Science*, volume 71, pages 431–445, Berlin, 1979. Springer-Verlag. Proc. Symposium on Automata, Languages and Programming.

[25] J. B. Orlin. The complexity of dynamic languages and dynamic optimization problems. *13th Proc. ACM Symp. Theory of Comp. Sci.*, pages 218–227, 1981.

[26] C. H. Papadimitriou. Games against nature. *Journal of Computers and Systems Science*, 31:288–301, 1985.

[27] R. Péter. *Recursive Functions*. Academic Press, New York, 1967.

[28] R. W. Ritchie. Classes of predictably computable functions. *Trans. AMS*, 106:139–173, 1963.

[29] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computers and Systems Science*, 4:177–192, 1970.

[30] T. J. Schaefer. On the complexity of two-person perfect-information games. *Journal of Computers and Systems Science*, 16:185–225, 1978.

[31] M. Sipser. On relativization and existence of complete sets. *9th ICALP (LNCS No. 140)*, pages 523–531, 1982.

[32] L. J. Stockmeyer. The complexity of decision problems in automata theory and logic. Technical Report Project MAC Tech. Rep. TR-133, M.I.T., 1974. PhD Thesis.

[33] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time. *5th Proc. ACM Symp. Theory of Comp. Sci.*, pages 1–9, 1973.

[34] O. Watanabe. On one-one polynomial time equivalence relations. *Theor. Computer Science*, 38:157–165, 1985.

[35] P. Young. Juris Hartmanis: fundamental contributions to isomorphism problems. *Structure in Complexity*, 3:138–154, 1988.

Chapter 6

Separation Results

April 13, 2009

6.1 Separation Results

Let K and K' be language classes. By a *separation result* for (K, K') we mean one showing the existence of a language $L \in K - K'$; L is said to *separate* K *from* K'. If the answers to the inclusion questions "Is K included in K'?" of chapter 4 are mostly negative as generally conjectured, it follows that techniques for separating classes are needed to resolve these open problems. The K-complete languages are the obvious candidates for languages L that separate K from K'. Unfortunately, given a specific L (such as a K-complete language) and a specific K', a proof that $L \notin K'$ typically appears to require powerful combinatorial methods quite beyond our present ability. Separation results based on combinatorial methods are rare (but we will see them in a later chapter). Instead, it is easier to construct a non-specific L in stages: for instance, in each stage we try to include in or exclude from L some words so as to ensure that L is not equal to each language in K'. While diagonalizing over K', we must ensure that L remains in K. We are just diagonalizing over K', of course, and chapter 4 contains such constructions. In any case, assuming that we manage to separate K from K', there are two typical consequences:

(a) We can infer that other pairs of classes must also be separated as follows. By a *translation result* we mean one with the following structure:

$$K_1 \subseteq K_2 \Rightarrow K_1' \subseteq K_2'$$

where K_i, K_i' $(i = 1, 2)$ are classes. Using this translation result, a separation result for (K_1', K_2') implies a separation result for (K_1, K_2). This method is especially useful in separating nondeterministic classes which tend to resist direct diagonalization.

(b) We can infer lower bounds on languages in K. In particular, any K-complete language (with respect to some reducibility $\leq$) is not in K', provided K' is closed under the reducibility $\leq$. For instance, we can separate *PSPACE* from *NLOG* using the diagonalization technique above. Since the fullness problem for regular expressions (chapter 5) is *PSPACE*-complete, we conclude that the problem cannot be solved in logarithmic space.

Most of the translational techniques are based on the idea of 'padding'. The idea originated with Ruby and Fischer [25] and was popularized by Ibarra [14]. We have already encountered padding in section 1 of the previous chapter: for instance, padding can 'translate' a complete problem for $NSPACE(n^{k+1})$ into a complete problem for $NSPACE(n^k)$. Padding techniques typically translate an inclusion between two classes with lower complexity into an inclusion between classes at higher complexity. For this reason, such results are sometimes called *upward translation results*. For instance, the following will be shown:

If NLOG $\subseteq$ *DLOG then LBA* $\subseteq$ *DLBA.*
If NP $\subseteq$ *P then NEXPT* $\subseteq$ *DEXPT.*

Limited forms of 'downward translation result' are known. Translational techniques will treated in sections 3 and 4.

Section 5 draws conclusions of type (b) above. The technique is based on efficient reducibilities; Meyer [20, 30] first exploited such techniques. The efficient reductions shown in the previous chapter allow us to deduce lower bounds on the various complete problems studied there.

Section 6 considers what may be described as *weak separation results*: these show that two classes K and K' must be distinct, $K \neq K'$. Hence either $K \setminus K'$ or $K' \setminus K$ (or perhaps both) must be non-empty although the proof does not indicate which is the case. For instance, we can show $DLOG \neq NTIME(n^k)$ for any k. Again the padding technique is useful. It should be noted that these weak results cannot be strengthened too easily: this will be clarified in the chapter on relativization where we show that these results can be relativized to force inclusion in either direction.

Section 7 presents *strong separation results*. The idea is that the explicit lower bounds derived in section 5 are of the 'infinitely often' type (e.g., every acceptor for such a language must, for some $c > 0$, take at least c^n steps for infinitely many n). We want to strengthen this to the 'eventually'.

Section 8 shows the limitations to the separation results. For instance, we demonstrate that there are complexity functions f, g such that f is 'arbitrarily more complex' than g and yet $DSPACE(f) = DSPACE(g)$. In other words, these two classes are inseparable.

On strong separation for characteristic classes. Let us make precise our above remarks about infinitely-often versus eventually bounds. The terminology will be useful in other contexts as well.

Definition 1. Let D be a set, θ a binary relation over D, and let $f, g : D \to D$ be partial functions with domain and range D.
(i) The relation θ is said to hold *infinitely often* between f and g if there are infinitely many $x \in D$ such that both $f(x)$ and $g(x)$ are defined, and the relation $f(x)\theta g(x)$ holds. We write '$f\theta g$ (*i.o.*)' or '$(\exists_\infty x) f(x)\theta g(x)$' in this case.
(ii) The relation θ is said to hold *eventually* between f and g if for all but finitely many $x \in D$, whenever $f(x)$ and $g(x)$ are defined then the relation $f(x)\theta g(x)$ holds. We write '$f\theta g$ (ev.)' or '$(\forall_\infty x) f(x)\theta g(x)$' in this case. ■

As an example, if f and g are total complexity functions, then f dominates g if and only if $f \geq g$ (ev.). Note that if $f\theta g$ (*i.o.*) then there are infinitely many values of $x \in D$ at which f and g are simultaneously defined. However, it is possible that $f\theta g$ (ev.) for the trivial reason that there are only finitely many of values of $x \in D$ for which both f and g are defined.

Suppose we have a language L that separates the pair of characteristic classes $(DTIME(t), DTIME(t'))$, i.e., $L \in DTIME(t) - DTIME(t')$. This implies that for any acceptor M for L, there exists a $c > 0$ such that there are infinitely many values of n such that M requires more than $t'(n)$ steps on inputs of length n:

$$AcceptTime_{\mathrm{M}}(n) > ct'(n) \text{ (}i.o.\text{)}. \tag{1}$$

Recall that $AcceptTime_{\mathrm{M}}(n)$ is undefined unless M accepts some word of length n. Note that (1) is equivalent to $AcceptTime_{\mathrm{M}}(n) \neq O(t'(n))$. We want to strengthen (1) to:

$$AcceptTime_{\mathrm{M}}(n) > ct'(n) \text{ (ev.)}. \tag{2}$$

for some $c > 0$. We could write this in the equivalent form: $AcceptTime_{\mathrm{M}}(n) = \Omega(t'(n))$.

We say that we have a *strong separation result* for the pair of characteristic classes $(DTIME(t), DTIME(t'))$ if we show the existence of an infinite[1] language L in $DTIME(t) - DTIME(t')$ such that any acceptor M for L satisfies (2). This notion of strong separation extends naturally to other characteristic classes.

6.2 The Basic Separation Results

Hartmanis and Stearns [12] began the study of time-based complexity theory and obtained the so-called *deterministic time hierarchy theorem*. Together with Lewis [11], they extended the result to space complexity. The main goal of this section is to present these two separation theorems, together with a separation theorem for reversal complexity. These theorems are considered basic because their proofs involve the simplest form of diagonalization and also because most other separation results ultimately rely on them. We also see that the diagonalization arguments here depend on the existence of tape reduction theorems (for space, time and reversal, respectively).

Our first hierarchy theorem is the following.[2]

Theorem 1 ([). *Deterministic space hierarchy)] Let s_1 and s_2 be complexity functions where $s_2(n) = \omega(s_1(n))$. If s_2 is space-constructible then*

$$DSPACE(s_2) - DSPACE(s_1) \neq \emptyset.$$

[1] To avoid the trivial interpretation of 'eventually' explained earlier.
[2] The original result requires the qualification $s_2(n) > \log n$.

Proof. Let $\Sigma = \{0, 1\}$ and fix any encoding of all 1-tape Turing machines with input alphabet Σ. Let

$$\phi_0, \phi_1, \ldots$$

be a listing of these machines in the order of their code. We may assume that the code for machine ϕ_i is the binary string (still denoted) i. Note that each machine uses states and symbols from the universal sets Q_∞ and Σ_∞, so our encoding assumes a fixed letter homomorphism h_0 from $Q_\infty \cup \Sigma_\infty$ to binary strings.

We now describe an acceptor M that diagonalizes over each machine ϕ_i. On input x of length n, M does the following: M marks off exactly $s_2(n)$ tape cells. Since $s_2(n) < \infty$ (by definition of space-constructible), the marking takes a finite number of steps. Treating x as the code of the machine ϕ_x, we would like M to "begin simulating ϕ_x on x, accepting iff either ϕ_x does not accept or if ϕ_x tries to leave the area containing the marked cells". The problem is that ϕ_x may loop without ever leaving the marked cells, and a simplistic simulation would not detect this. To overcome this difficulty, we apply Sipser's technique as shown in section 9 of chapter 2: for each accepting configuration C of ϕ_x that 'can fit' within the marked cells, M does a search of the tree $T(C)$ rooted at C. If it discovers the initial configuration $C_0(x)$ in $T(C)$, M rejects. If for all such C's, M fails to discover the initial configuration then M accepts. M can do this search in space $s_2(n)$. We should clarify the qualification that C 'can fit' within the marked cells: by this we mean that the work tapes of C can be represented using the marked cells, but the input tape of C is directly represented as the input tape of the simulating machine.

We now show that the language $L(\mathrm{M})$ separates $DSPACE(s_2)$ from $DSPACE(s_1)$. By construction, $L(\mathrm{M}) \in DSPACE(s_2)$. It remains to show that $L(\mathrm{M})$ is not in $DSPACE(s_1)$. Suppose $L(\mathrm{M})$ is accepted by some deterministic N in space s_1. By the tape reduction theorem for space, we may assume that N is 1-tape and hence, by properties of the usual coding of Turing machines, we get this property: there are infinitely many indices x such that $L(\mathrm{M}) = L(\phi_x)$ and ϕ_x accepts in space $s_1(n)$.

Using the said letter homomorphism h_0 that encodes states and symbols from the universal sets as binary strings, we may assume that each configuration of N that uses space s is encoded by a binary string of length $O_{\mathbb{N}}(s)$. Choose the length of x sufficiently large so that $O_{\mathbb{N}}(s_1(|x|)) \leq s_2(|x|)$. We now prove that M accepts x iff ϕ_x rejects x. If ϕ_x accepts x then it accepts in a configuration C with space at most $s_1(|x|)$ and C can be encoded in space $O_{\mathbb{N}}(s_1(|x|)) \leq s_2(|x|)$. Thus we will search the tree $T(C)$ and find the initial configuration, leading to rejecting. On the other hand, if ϕ_x does not accept x, then the search of the tree $T(C)$ for each accepting configuration C fails and by definition M accepts. This proves $L(\mathrm{M}) \neq L(\phi_x)$. **Q.E.D.**

If we are interested in separating the running space class $DSPACE_r(s_2)$ from $DSPACE_r(s_1)$, then we can avoid the assumption s_2 be space-constructible (Exercise).

We now show the time analogue of the previous theorem. However the hierarchy is not as tight because the tape reduction theorem for deterministic time (viz., the Hennie-Stearns theorem) incurs a logarithmic slow-down.

THEOREM 2 ((Deterministic time hierarchy)). *Let t_1 and t_2 be complexity functions with $t_1(n) \log t_1(n) = o(t_2(n))$. If t_2 is time-constructible and $t_1(n) > n$ then*

$$DTIME(t_2) - DTIME(t_1) \neq \emptyset.$$

Proof. The proof is similar to the previous one, except that we now appeal to the existence of a universal machine.[3] Let $U = \{U_0, U_1, \ldots\}$ be a universal machine simulating all 2-tape deterministic machines over some fixed input alphabet. Again we construct an acceptor M that diagonalizes over each U_i that happens to accept in time $o(t_2)$. The machine M operates as follows: on input x of length n, M first copies the input x onto two of its work-tapes. Next, M uses these two copies of x as input to simulate the universal machine U_x on input x for at most $t_2(n)$ of steps (each step of M corresponds to a step of U_x). To ensure that at most $t_2(n)$ steps are used, M will concurrently time-construct $t_2(n)$, using the original x as input. If U_x accepts within $t_2(n)$ steps then M rejects; otherwise M accepts. Clearly $L(\mathrm{M})$ is in $DTIME(2n + t_2(n)) = DTIME(t_2)$.

It remains to show that $L(\mathrm{M}) \notin DTIME(t_1)$. For the sake of a contradiction, assume that $L(\mathrm{M})$ is accepted by some acceptor in time t_1. The Hennie-Stearns theorem then implies that $L(\mathrm{M})$ is accepted by some 2-tape machine N in time $t_1 \log t_1$. From the recurrence property of universal machines, there are infinitely many indices x such that $L(\mathbb{N}) = L(U_x)$ and U_x uses time $O_{\mathrm{N}}(t_1 \log t_1)$. Choosing x such that $n = |x|$ is sufficiently large,

$$2n + O_{\mathbb{N}}(t_1(n) \log t_1(n)) \leq t_2(n).$$

If U_x accepts x, then M needs $2n + O_{\mathbb{N}}(t_1(n) \log t_1(n))$ steps to simulate U_x on x until completion. Since this number of steps is at most $t_2(n)$, M will discover that U_x accepts x and M so rejects x by our construction. If U_x

[3] We could, as in the previous proof, argue directly about the standard encoding of Turing machines, but the use of universal machines is a slightly more general (abstract) approach. Our deterministic space hierarchy theorem, however, does not seem to yield as easily to an argument by universal machines. Why?

does not accept x, then we similarly see that M will accept x. This proves $L(\mathrm{M}) \neq L(U_x) = L(\mathbb{N})$, a contradiction.
Q.E.D.

The above result can be sharpened in two ways: if t_1 is also time-constructible then Paul [22] shows that the above separation of time classes can be achieved with $t_2(n) = o(t_1(n)\log^{\epsilon} t_1(n))$ for any $\epsilon > 0$. If we consider classes defined by Turing machines with some fixed number of tapes, Fürer [10] has shown the above separation can be achieved with $t_2 = o(t_1)$ provided we restrict attention to Turing machines with a fixed number k of tapes, $k \geq 2$.

Using the space and time hierarchy theorems, we can infer separation for some of the inclusions on our canonical list:

COROLLARY 3. (a) $P \subset DEXPT \subset DEXPTIME$

(b) $NLOG \subset PSPACE \subset EXPS \subset EXPSPACE$

Finally we present a hierarchy theorem for reversal complexity[5]. Although we can give a direct proof using the tape reduction theorem for reversals, the following shorter proof follows from relationships between reversal and space complexity shown in chapter 2.

THEOREM 4 ((Deterministic reversal hierarchy)). *Let $r_1(n)$ and $r_2(n)$ be complexity functions such that $(r_1(n))^2 = o(r_2(n))$, $r_1(n) = \Omega(\log n)$ and $r_2(n)$ is reversal-constructible. Then*

$$DREVERSAL(r_2(n)) - DREVERSAL(r_1(n)) \neq \emptyset.$$

Proof. Since $(r_1(n))^2 = o(r_2(n))$, by the deterministic space hierarchy theorem, there exists a language L such that

$$L \in DSPACE(r_2(n)) - DSPACE((r_1(n))^2).$$

Using the relation $DSPACE(r_2) \subseteq DREVERSAL(r_2)$ and $DREVERSAL(r_1) \subseteq DSPACE(r_1^2)$, we conclude

$$L \in DREVERSAL(r_2(n)) - DREVERSAL(r_1(n)).$$

Q.E.D.

EXERCISES

Exercise 0.1: Let $s < \infty$ be a complexity function. Prove a $(DSPACE(s), RE)$-separation. ♢

Exercise 0.2: Prove a $(PLOG, PSPACE)$-separation. ♢

END EXERCISES

6.3 Padding Arguments and Translational Lemmas

The theme of the hierarchy theorems in the previous section is: given a complexity function $f_1(n)$ what is the smallest complexity function $f_2(n)$ such that there is a language L accepted within complexity $f_2(n)$ but not $f_1(n)$. In each case, $f_2(n)$ need not be more than quadratic in $f_1(n)$. The basic technique in these proofs requires constructing a diagonalizing machine M that can "efficiently decide" whether a simulated machine N accepts in a given amount of resource.

This approach becomes quite ineffectual when N is nondeterministic and we want to decide if N accepts within time t. The best nondeterministic method known for deciding this question amounts to a naive *deterministic* simulation of N, a procedure that takes time exponential in t. This implies that $f_2(n)$ is exponentially larger than $f_1(n)$.

To separate nondeterministic space, we could use Savitch's technique to deterministically simulate nondeterministic space-bounded computations, using quadratically more space. The next two sections show more efficient techniques to separate nondeterministic time and space.

We begin with a translational lemma under composition of complexity functions.

LEMMA 5 ((Function-composition translation for space)). *Let $s_1(n) \geq n$, $s_2 \geq n$ and $f \geq n$ be complexity functions, and assume f is space-constructible. For $X = D$ or N, $XSPACE(s_1) \subseteq XSPACE(s_2)$ implies $XSPACE(s_1 \circ f) \subseteq XSPACE(s_2 \circ f)$.*

Proof. The structure of the proof is suggested by following the arrows in this "commutative diagram":

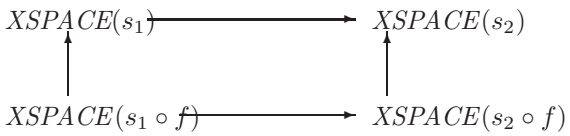

Suppose that $XSPACE(s_1) \subseteq XSPACE(s_2)$ and (Σ, L) is accepted by some M in space $s_1 \circ f$. We want to show that $L \in XSPACE(s_2 \circ f)$. First 'translate' L to $XSPACE(s_1)$ by padding. More precisely, the padded version of L is

$$L' = \{x\$^i : x \in L,\ |x\$^i| = f(|x|)\}$$

where \$ is a new symbol not in Σ. First we demonstrate that L' is in $XSPACE(s_1)$: on input $x\$^i$ of length n we check if $|x\$^i| = f(|x|)$ using space at most n, since f is space-constructible. If not, reject; else we simulate M on input x, and accept iff M accepts. Clearly the space used is at most $\max\{n, s_1(f(|x|))\} = s_1(n)$, as desired. Therefore, by assumption, L' is accepted by some M′ in space s_2. Next we demonstrate that L is in space $XSPACE(s_2 \circ f)$: on input x of length n, construct $x\$^i$ such that $|x\$^i| = f(|x|)$ using space $f(n)$. Now simulate M′ on $x\$^i$ using $s_2(|x\$^i|) = s_2(f(n))$ space. **Q.E.D.**

We illustrate the use of this lemma in the next result. In the proof, we use the fact that n^r for any rational number $r \geq 1$ is space-constructible (see Exercises in chapter 2).

LEMMA 6 ((Ibarra [14])). *For all reals $r > s \geq 1$, $NSPACE(n^r) - NSPACE(n^s) \neq \emptyset$.*

Proof. Choose positive integers a, b such that

$$r > \frac{a+1}{b} > \frac{a}{b} > s.$$

Note that $a > b \geq 1$. For the sake of contradiction, assume that $NSPACE(n^r) = NSPACE(n^s)$. Then

$$NSPACE(n^{(a+1)/b}) \subseteq NSPACE(n^{a/b}).$$

From this inclusion, an application of the previous translation lemma with $f(n) = n^{(a+1)b}$ and also with $f(n) = n^{ab}$, yields (respectively)

$$\begin{aligned} NSPACE(n^{(a+1)^2}) &\subseteq NSPACE(n^{a(a+1)}), \\ NSPACE(n^{a(a+1)}) &\subseteq NSPACE(n^{a^2}). \end{aligned}$$

Hence

$$NSPACE(n^{(a+1)^2}) \subseteq NSPACE(n^{a^2})$$

We now claim that for any $k \geq 2$ that is a power of two, the inclusion

$$NSPACE(n^{(a+1)^k}) \subseteq NSPACE(n^{a^k})$$

holds. The basis case has just been shown. For the induction, assume the inclusion holds for k. Then we can deduce

$$NSPACE(n^{(a+1)^{2k}}) \subseteq NSPACE(n^{a^{2k}})$$

by two applications of the translation lemma (the reader should verify this). If we choose $k \geq a$, then $(a+1)^k \geq 2a^k + 1$ for $a > 1$. Thus

$$\begin{aligned} DSPACE(n^{1+2a^k}) &\subseteq NSPACE(n^{1+2a^k}) \\ &\subseteq NSPACE(n^{(a+1)^k}) \\ &\subseteq NSPACE(n^{a^k}), \quad \text{(by what we had just shown)} \\ &\subseteq DSPACE(n^{2a^k}) \end{aligned}$$

where the last inclusion follows from Savitch's theorem. But $DSPACE(n^{1+2a^k}) \subseteq DSPACE(n^{2a^k})$ contradicts the deterministic space hierarchy theorem. **Q.E.D.**

This result will be improved in the next section. We now prove another translational lemma due to Savitch.

Definition 2. Let $s(n) < \infty$ be a complexity function that is defined for all sufficiently large n. We say s is a *moderately growing* if it is unbounded, $s(n) \neq O(1)$, and there is a constant c such that, eventually, $s(n) \leq c \cdot s(n-1)$. ■

Observe that the functions in $\log n$, $\log^{O(1)} n$, $n^{O(1)}$ and $O(1)^n$ are moderately growing. However functions such as 2^{2^n} and 2^{n^k} $(k > 1)$ are not.

THEOREM 7 ((Upward translation of space)). *Let s, s' be complexity functions. If s is moderately growing and space-constructible, and if $s(n) \leq s'(n) < \infty$ for all n, then*

$$NSPACE(s) \subseteq DSPACE(s) \Longrightarrow NSPACE(s') \subseteq DSPACE(s').$$

Proof. The proof is similar in structure to that for the translational lemma. Let (Σ, L) be accepted by some nondeterministic M in space s'. We shall show that L is in $DSPACE(s')$. Again we translate L to a related problem in $NSPACE(s)$ as follows: let \$ be a new symbol not in Σ and define the following padded version of L:

$$L' = \{x\$^i \ : \ \text{M accepts } x \text{ in space } \ s(|x\$^i|), i \geq 0\}.$$

Clearly, $L' \in NSPACE(s)$.

Since $NSPACE(s) \subseteq DSPACE(s)$, we infer that L' is accepted by some halting deterministic M′ in space s.

We now construct a deterministic M″ to accept L as follows: on input $x \in \Sigma^*$, simulate M′ on $x\$^i$ for $i = 0, 1, 2, \ldots$, until the first i such that M′ accepts $x\$^i$. At that point M″ accepts. Otherwise M″ runs forever.

Correctness of M″: It is easy to show that M″ accepts L. We next claim that M″ accepts in space $s'(n)$. To see this, let $x \in L$. If $s'(|x|) = s(|x|)$ then it is not hard to see that M′ accepts x in space $s'(|x|)$ space. Otherwise, $s'(|x|) > s(|x|)$ and there is a smallest $i > 0$ such that

$$s(|x\$^i|) \geq s'(|x|) > s(|x\$^{i-1}|).$$

Note that i is well-defined since s is unbounded and s' is finite. Now for any $j = 0, \ldots, i-1$, $s(|x\$^j|) < s'(|x|)$. Hence if M′ accepts $x\$^j$, then M″ accepts x in less than $s'(|x|)$ space; otherwise, surely M′ accepts $x\$^i$. This is because we see that M accepts x in space $s(|x\$^i|)$, and by definition of L', $x\$^i \in L' = L(\text{M}')$. Hence M″ accepts in space at most $s(|x\$^i|)$. But since s is moderately growing, there is some $c \geq 1$ such that $s(|x\$^i|) \leq cs(|x\$^{i-1}|)$. This proves that M″ accepts in space $cs(|x\$^{i-1}|)$. The claimed bound follows by space compression. This completes the proof. **Q.E.D.**

COROLLARY 8. *If NLOG ⊆ DLOG then LBA=DLBA.*

Similarly, we can prove translational lemmas for time complexity (Exercise) and deduce:

$$\text{If } NP \subseteq P \ \text{ then } NEXPT \subseteq DEXPT.$$

These upward translation results raises the possibility of some form of 'downward translation'. Our next result may be regarded as partial downward translation. It involves the so-called *tally* or *contentless languages*: these are languages over a single letter alphabet, say $\{1\}$. In the remainder of this section, let Σ be any alphabet with $k \geq 2$ letters. We might as well assume $\Sigma = \{1, 2, \ldots, k\}$. For any $w \in \Sigma$, let $tally(w) \in \{1\}^*$ denote the unary representation of the integer w (regarded as a k-adic number). Then for any language (Σ, L), define the language $(\{1\}, tally(L))$ where

$$tally(L) = \{tally(w) \ : \ w \in L\}.$$

Conversely, define the function *untally* that takes any unary word $w \in \{1\}^*$ to its k-adic representation $untally(w) \in \Sigma^*$. For any tally language $(\{1\}, L)$, define the language $(\Sigma, untally(L)\}$ where

$$untally(L) = \{untally(w) \ : \ w \in L\}.$$

Thus *tally* and *untally* are inverses of each other.

LEMMA 9 ((Space translation for tally languages)). *Let $X = N$ or D, let L be a language and f a complexity function with $f(n) \geq n$. Then $L \in XSPACE(f(O(n)))$ iff $tally(L) \in XSPACE(f(O(\log n)))$.*

Proof. If (Σ, L) is accepted in space $f(O(n))$ by some machine M then we can accept $tally(L)$ as follows: on input 1^n, we first compute $w \in \Sigma^*$ such that $\nu(w) = n$. This takes space $|w| = O(\log n)$. Now we simulate M on w, taking space $f(O(\log n))$. Conversely, if $tally(L)$ is accepted by some N in space $f(O(\log n))$ then we could try to accept L as follows:

> On input $w \in \Sigma^*$, compute $tally(w)$ and then simulate N on $tally(w)$ using space $f(O(\log(|tally(w)|))) = f(O(|w|))$.

The only problem is that $untally(w)$ needs space $O(1)^{|w|}$. To circumvent this, because $untally(w)$ is 'contentless', it suffices to keep track of the position of the input head of N on the virtual input $tally(w)$. The space necessary to keep track of the head position is $O(|w|)$ which is order of $f(O(|w|))$ since $f(n) \geq n$. **Q.E.D.**

We then obtain the following weak downward translation:

COROLLARY 10 ((Savitch)). *If $DLBA = LBA$ then $DLOG|\{1\} = NLOG|\{1\}$.*

Proof. Let L be a tally language in $NLOG = NSPACE(O(\log n))$. Then the above lemma implies that $untally(L) \in LBA = NSPACE(O(n))$. So by assumption $untally(L) \in DLBA$. Another application of the lemma shows that $L = tally(untally(L)) \in DLOG$. **Q.E.D.**

Combining the upward and weak downward translation results, we conclude that

$$DLBA = LBA \iff DLOG|\{1\} = NLOG|\{1\}.$$

Similar results relating to time complexity of tally languages can be shown.

6.4 Separation for Nondeterministic Classes

The reader may verify that the proof of Ibarra in the last section fails for nondeterministic time. This situation was first remedied by Cook [6] who showed the analogue of Ibarra's result [14]: for all reals $r > s \geq 1$,

$$NTIME(n^r) - NTIME(n^s) \neq \emptyset.$$

Cook's technique was generalized to a very strong form by Seiferas, Fischer and Meyer [27, 28]. Unfortunately their original proof is quite involved (using a form of recursion theorem for nondeterministic time – see Exercises). Simpler (but still delicate) proofs have been found by Žák [32] and by Li [16]; both these proofs have the added bonus of providing a tally language to separate the classes, answering an open question in [27]. Here we follow the proof of Žák.

We require some preparatory results. First, we note a simple but important consequence of the nondeterministic tape reduction theorem of Book, Greibach and Wegbreit (chapter 2):

LEMMA 11. *For any alphabet Σ, there is a 2-tape universal acceptor $U = \{U_0, U_1, \ldots\}$ for the class $RE|\Sigma$ such that for each nondeterministic acceptor* M *over the input alphabet Σ, and each complexity function $t(n) > n$, if* M *accepts in time $t(n)$ there exist infinitely many indices i such that U_i accepts L(M) in time $O_{U,\mathrm{M}}(t(n))$.*

Proof. For any machine N accepting in time $t(n) > n$, it follows from the nondeterministic tape reduction theorem that there is a 2-tape machine M accepting L(N). It is seen from our standard construction of a universal machine U for $RE|\Sigma$ that there are infinitely many indices i such that $L(U_i) = L(\mathrm{M})$ and U_i accepts in time $c \cdot t(n)$ for some c that depends on M (rather than on i) and U. **Q.E.D.**

So the point of this lemma is that we can efficiently simulate any multitape acceptor M by infinitely many 2-tape versions as given by a universal machine U.

Let us fix $\Sigma = \{0, 1\}$ and U as in the above lemma. For any complexity function t and any Turing acceptor M, define the *t-cutoff* language defined by M to be

$$L^t(\mathrm{M}) = \{x \in \{0,1\}^* \ : \ \mathrm{M} \textit{ accepts } x \textit{ in time } t(|x|)\}.$$

Note that $L^t(\mathrm{M}) \subseteq L(\mathrm{M})$ and $L^t(\mathrm{M})$ is not necessarily in $NTIME(t)$ unless t is time-constructible or if M accepts in time t. Relative to the universal machine U we define:

$$NTIME_U(t) = \{L(U_i) \ : \ U_i \textit{ accepts in time } t\}$$
$$NTIME_U^{cut}(t) = \{L^t(U_i) \ : \ i = 0, 1, \ldots\}$$

Discussion. We may call the classes $NTIME_U(t)$ and $NTIME_U^{cut}(t)$ *universal-time classes* (relative to U) since they refer to a 'universal' standard of time-keeping as defined by the steps of U. In contrast, the usual classes may be called *local-time classes* since a time step as defined by a Turing machine ϕ_i is not quite comparable to that defined by another ϕ_j when the tape alphabets and state sets of ϕ_i and ϕ_j are different. The connection between universal and local time is as follows: for each i, there is a constant $O_i(1)$ such that each step of ϕ_i is simulated by $O_i(1)$ steps of U_i. Note that the universal-time classes $NSPACE_U(t)$ do not in general enjoy the linear speedup property: it is not clear that $NTIME_U(t)$ is equal to $NTIME_U(2t)$, for instance. However, by the linear speedup theorem for local-time classes we can conclude that

$$NTIME(t) = NTIME_U(O(t)).$$

The crux of the diagonal process in our proof is captured in the following somewhat technical result. The statement of the result is somewhat long but its proof is not much longer.

LEMMA 12. *Let $K = K|\Sigma$ be a class and U be any universal acceptor for K. Suppose $1 \in \Sigma$, and there exist languages (Σ, L) and (Σ, D), and functions $\alpha : L \to \mathbf{N} = \{0, 1, 2, \ldots\}$ and $\beta : L \to \mathbf{N}$ with the following property: For all $x \in L$ and for all j $(0 \leq j \leq \beta(x))$,*

$$x1^j \in D \iff \begin{cases} x1^{j+1} \in U_{\alpha(x)} & \text{if } j < \beta(x) \\ x \notin U_{\alpha(x)} & \text{if } j = \beta(x). \end{cases} \tag{3}$$

If α is an onto function then $D \notin K$.

Let us call L the 'unpadded' set and D the 'diagonal' set. Intuitively, deciding if an unpadded word x is in D is equivalent to the question whether $U_{\alpha(x)}$ accepts the padded strings $x1^j$ (for $j = 0, \ldots, \beta(x)$). On the other hand, $U_{\alpha(x)}$ accepts $x1^{\beta(x)}$ iff x is not in the diagonal set D. These two incompatible conditions imply $D \notin K$. Observe the highly stylized nature of this translation.

Proof. Aiming towards a contradiction, assume $D = L(U_i)$ for some $i \geq 0$. Since $\alpha : L \to \mathbf{N}$ is onto, let $x \in L$ such that $\alpha(x) = i$. If $\beta(x) = 0$ then we have that $x \in D$ iff $x \notin U_{\alpha(x)}$, contradicting our assumption that $D = L(U_{\alpha(x)})$. Observe that if $\beta(x) \geq 1$ then $x \in D$ iff $x1 \in L(U_{\alpha(x)})$ iff $x1 \in D$. If $\beta(x) \geq 2$ then $x1 \in D$ iff $x1^2 \in L(U_{\alpha(x)})$ iff $x1^2 \in D$. Repeating this, we see that $x, x1, x1^2, \ldots, x1^{\beta(x)}$ are all in D or none are in D. However, $x \in D$ iff $x1^{\beta(x)} \notin L(U_{\alpha(x)}) = D$, contradiction. **Q.E.D.**

THEOREM 13. *If $t(n) > n+1$ is time-constructible then there is a tally language D in $NTIME(t(n)) - NTIME_U^{cut}(t(n-1))$.*

Proof. We use the universal machine U for $RE|\Sigma$ of lemma 11. Let Σ be the unary alphabet $\{1\}$. Let U' be an 'efficient' universal acceptor for the class $NTIME_U^{cut}(t(n-1))|\Sigma$ obtained by simulating exactly $t(n-1)$ steps of U, accepting if and only if U accepts within $t(n-1)$ steps. Note that U' can simulate U in realtime (i.e. step for step). However, U' needs an extra n steps to initially write onto another tape the word 1^{n-1} which then serves as input for the 'parallel' process to time-construct $t(n-1)$. Hence, U' on inputs of length n runs in time $n + t(n-1)$.

We inductively define an increasing sequence of integers

$$n_0 < n_1 < \cdots < n_i < \cdots$$

as follows: $n_0 = 1$ and

$$n_{i+1} = 1 + n_i + c_i$$

where c_i is the number of steps sufficient (for some fixed Turing machine) to *deterministically* simulate the behaviour of U_i' on input 1^{n_i}. Observe that c_i is well-defined because U_i' accepts in at most $t(n-1)$ steps on inputs of length n. (In general we expect c_i to be exponential in $t(n_i - 1)$, but no matter.) To apply the previous lemma, we define the 'unpadded set' L to be $\{1^{n_i} : i = 0, 1, \ldots\}$ and the functions $\alpha, \beta : L \to \mathbf{N}$ are given by:

$$\begin{aligned} \alpha(1^{n_i}) &= i, \\ \beta(1^{n_i}) &= 1 + c_i \end{aligned}$$

for $i \geq 0$. Finally, we define the diagonal language $D \subseteq \{1\}^*$ by constructing a machine M_D to accept D:

(A) On input 1^n, M_D computes in *phases* where in phase i $(i = 0, 1, \ldots)$ M_D simulates U_i' on input 1^{n_i}. M_D stops in the middle of phase $k+1$ where k is defined by the inequality

$$n_{k+1} < n \leq n_{k+2}.$$

It is easy to organize the computation of M_D so that at the end of each phase, when M_D has just finished simulating U_i' on input 1^{n_i}, M_D has a copy of $1^{n_{i+1}}$ on a separate tape, ready to be used as input for the next phase. Furthermore, the time to carry out each phase is $O(n_{i+1})$ steps (actually, $O(n_i)$ steps suffice) for $i = 0, \ldots, k$, and the partial phase $k+1$ uses only $O(n)$ time. The total time to do phases 0 to k, including the partial $(k+1)$st phase, is

$$\sum_{i=0}^{k} O(n_{i+1}) + O(n) = \sum_{i=0}^{k} O(\frac{n}{2^i}) + O(n) = O(n).$$

Here we use the fact that $c_i \geq n_i$ and so $n_{i+1} > 2n_i$.

(B) If $n = n_{k+1} - 1$ then M_D can in $O(n)$ steps discover whether $1^{n_k} \in U_k'$. This is because $n_{k+1} - 1 > c_k$ and M_D can deterministically simulate U_k' on 1^{n_k} in at most c_k steps. M_D rejects iff U_k' accepts 1^{n_k}. Note that in this case, with $x = 1^{n_k}$,

$$1^n = x1^{\beta(x)} = x1^{1+c_k} \in L(\mathrm{M}_D) \iff x \notin L(U'_{\alpha(x)}).$$

(C) If $n < n_{k+1} - 1$ then M_D simulates U_k' on 1^{n+1}, accepting if and only if U_k' accepts. Thus with $x = 1^{n_k}$ and $j = n - n_k < n_{k+1} - n_k - 1 \leq c_k$,

$$x1^j \in L(\mathrm{M}_D) \iff x1^{j+1} \in L(U'_{\alpha(x)})$$

We observe that steps (A) and (B) take $O(n)$ time; step (C) takes $1 + n + t(n)$ since U' takes $1 + n + t(n)$ steps on inputs of length $n+1$. This implies $D = L(\mathrm{M}_D) \in NTIME(t + O(n)) = NTIME(t)$, by the speedup theorem for nondeterministic time. An application of the previous lemma shows $D \notin NTIME_U^{cut}(t(n-1))$. **Q.E.D.**

Note that step (C) is where the padding takes place: it reduces the query about $x1^j$ to one above x^{j+1}. Our main result of this section follows immediately.

THEOREM 14 ((Nondeterministic time hierarchy)). *If $t(n) > n + 1$ is time-constructible then there exists a tally language in*

$$NTIME(t(n)) - NTIME(o(t(n-1))$$

Proof. For any function $t'(n)$, if $t'(n) = o(t(n-1))$ then

$$NTIME(t') \subseteq NTIME_U^{cut}(t(n-1)).$$

An application of the previous theorem yields the desired separation. **Q.E.D.**

In the statement of this theorem, we need the '-1' in defining one of the complexity classes because we need to pad at least one symbol for the induction to go through. It is not known if this '-1' can be removed or if it is essential.

Now we can infer Cook's result that $NTIME(n^r) - NTIME(n^s) \neq \emptyset$ $(r > s \geq 1)$: first choose a rational number b such that $r > b > s$. The function n^b is "approximately" time-constructible in this sense: there exists a time-constructible function t such that $t(n) = \Theta(n^b)$ (Exercise). Clearly $(n+1)^s = o(t(n))$, so the preceding theorem implies $NTIME(t) - NTIME(n^s)$ (and hence $NTIME(n^r) - NTIME(n^s)$) is non-empty.

As another application of this theorem, we infer that $NTIME(n2^n) - NTIME(2^n)$ is non-empty. On the other hand, the theorem fails to decide whether $NTIME(n2^{2^n}) - NTIME(2^{2^n})$ is empty. This remains an open problem.

We now show a corresponding separation result for nondeterministic space [27]. The proof below employs the technique used in showing that nondeterministic space is closed under complementation.[4]

THEOREM 15. *Let s_2 be space-constructible, $s_2(n) \geq \log n$. If $s_1(n) = o(s_2(n))$ then $NSPACE(s_2) - NSPACE(s_1) \neq \emptyset$.*

Proof. Let U be a universal acceptor for all nondeterministic 1-tape Turing acceptors. We describe an acceptor M to diagonalize over each U_i: on input x of length n, we mark out $s_2(n)$ cells. For each accepting configuration C that fits inside these marked cells, we call a subroutine that uses Immerman's technique to unequivocally check if C can be reached from the initial configuration. (Recall that this means that there is at least one terminating computation path and further all terminating computation paths accept or all reject.) If any subroutine call loops, then we loop; if any accepts, then we reject; if all reject, then we accept.

By now, it is a simple exercise to show that $L(M)$ separates $NSPACE(s_2)$ from $NSPACE(s_1)$. **Q.E.D.**

[4]The original proof of Seiferas-Fischer-Meyer is more involved. Immerman [15] attributes the idea of the present proof to M. Fischer.

6.5 Applications to Lower Bounds

Informally, if $L \leq L'$ where $\leq$ denotes some efficient reducibility then the complexity of L is at most the complexity of L' plus the complexity of the reducibility $\leq$. For a simple illustration of such results, we consider many-one reducibilities. First, a definition.

Definition 3. For any transformation $t : \Sigma^* \to \Gamma^*$, and f a complexity function, we say that t is *f-bounded* if for all x, $|t(x)| \leq f(|x|)$. ■

LEMMA 16. *Let L be many-one reducible to L' via a transformation g and $L' \in X$-TIME-SPACE(t, s) where $X = N$ or D, and let t, s be non-decreasing complexity functions. If g can be computed in time $u(n)$ and is $f(n)$-bounded then*

(i) $L \in XTIME(t(f(n)) + u(n))$ *and*

(ii) $L \in X$-TIME-SPACE$(u(n)t(f(n)), s(f(n)) + \log f(n))$.

Proof. Let M accept L' in time t and space s, and let T be the log-space transducer that transforms L to L'. It is straightforward to show (i) by constructing an acceptor N for L: on input x, N simulates T on x to obtain $\mathrm{T}(x)$; then it simulates M on $\mathrm{T}(x)$, accepting if and only if M accepts. Since $|\mathrm{T}(x)| \leq f(|x|)$ and t is non-decreasing, the desired time bound of $t(f(n)) + u(n)$ on N follows immediately. Note that if M is deterministic then so is N.

We show (ii). To achieve a space bound of $s(f(n))) + \log f(n)$, we modify the above construction of M by using the technique from chapter 4 (section 2): simulate the acceptor M on input $\mathrm{T}(x)$ without keeping the entire input string in storage, but use T as a subroutine to (re)compute each symbol of $\mathrm{T}(x)$ as needed by M. To do this, we need $O(\log f(n))$ space to represent the position of the input head of M on $\mathrm{T}(x)$. The space bound of (ii) follows from this '$O(\log f(n))$' plus the $s(f(n))$ space used by M on $\mathrm{T}(x)$. **Q.E.D.**

The reader can find analogous results for other types of reducibilities.

It is desirable in applying the lemma to have a small bounding function f. As seen in the transformations in chapter 5, f is typically linear, $f(n) = O(n)$; this simply means that t belongs to the class **Llin** of log-linear transformations defined in chapter 4 (section 2).

The principal application of such a lemma is to obtain lower bounds on *specific* languages. Meyer and Stockmeyer [21, 30, 19] and Hunt [13] were the first to infer such lower bounds on natural computational problems. The basic structure of such proofs is outlined next. Assume that we want to prove a lower bound of $s(n)$ on the deterministic space-complexity of a language L_0.

(a) *Show that L_0 is hard for some complexity class.* Choose a suitable class K of languages such that each $L \in K$ is efficiently reducible to L_0. For instance, suppose there is a $k \geq 1$ such that each $L \in K$ is many-one reducible to L_0 via a n^k-bounded log-space transformation.

(b) *Infer the lower bound by appeal to a separation result.* Assume that we want to show a space lower bound. Suppose that $s(n) \geq \log n$ is non-decreasing and there is a separation result for $(K, DSPACE(s))$. Then we claim that $s'(n) = s(n^{1/k})$ is an i.o. lower bound on the space-complexity of L_0. For the sake of contradiction, assume otherwise that $L_0 \in DSPACE(s')$. By our choice of K, there exists a language $L_1 \in K - DSPACE(s)$. But the above lemma implies that L_1 can be accepted in space $O(s'(n^k) + \log n) = O(s(n))$, contradiction.

Note that in this outline, we normally only have to show step (a) since step (b) is usually routine[5]. Since step (a) involves showing L_0 to be K-hard, this explains why showing a problem to be hard for a class is often called a 'lower bound proof'.

The remainder of this section illustrates such applications.

Lower bound on the fullness problem for regular languages. Recall the problem $\mathrm{FULL} = \mathrm{FULL}(+, \cdot, *)$ of checking if a given regular expression α denotes the set $\{0, 1\}^*$. In chapter 5 we showed that the problem FULL is hard for $NSPACE(n)$ under log-linear transformations. It easily follows from the outline (a) and (b), by appealing to the nondeterministic space hierarchy theorem, that

$$\mathrm{FULL} \notin NSPACE(o(n)).$$

Hence every nondeterministic acceptor for FULL must use more than linear space infinitely often. Stated in isolation, this statement should be appreciated for its depth since, as remarked in chapter 1, it is a statement

[5]Or rather, radically new separation results do not seem to be easy to derive – and so step (b) is typically an appeal to one of the separation theorems we have shown here.

about all imaginable (but mostly unimaginable) Turing machines that accept FULL. Yet, because of the long development leading up to this point, this statement may seem rather easy.

Similarly, using the log-linear reductions in section 5 of chapter 5, we conclude that any deterministic acceptor M for the problem FULL$(+,\cdot,*,^2)$ uses space more than c^n infinitely often, for some $c \geq 0$.[6]

Nondeterministic time lower bounds for complements of languages. It follows from the log-linear transformation shown in chapter 5 that any nondeterministic acceptor M for the problem INEQ = INEQ$(+,\cdot,^2)$ uses time more than c^n infinitely often, $c > 0$. Now consider what is essentially[7] the complementary problem: let EQUIV = EQUIV$(+,\cdot,^2)$ denote the set of pairs (α,β) of $\{+,\cdot,^2\}$-expressions encoded over the binary alphabet such that $L(\alpha) = L(\beta)$. We would like a lower bound on EQUIV based on a lower bound on INEQ. Towards this end, we use a nice separation result attributed to Young [29].

LEMMA 17. *If t is time-constructible then*

$$NTIME(n \cdot t(n)) - \text{co-}NTIME(t(n)) \neq \emptyset.$$

Proof. Let U be an efficient universal acceptor for the characteristic class $NTIME(t)|\Sigma$ where $\Sigma = \{0,1\}$. So for any $(\Sigma, L) \in NTIME(t)$ there are infinitely many indices i such that U_i accepts L in time $O_L(t(n))$. Using the usual encodings, we may assume that U_i accepts in time $c \cdot |i| \cdot t(n)$ for some constant $c = c(U) > 0$ that does not depend on i or n. Let L_0 consist of those words x such that U_x accepts x. Then it follows that L_0 can be accepted in $n \cdot t(n)$ using a direct simulation of U. To prove the lemma, it remains to show that $L_0 \notin \text{co-}NTIME(t(n))$. Suppose otherwise, $L_0 \in \text{co-}NTIME(t(n))$. Then let co-$L_0$ be accepted by U_x for some x; this means that $x \in$ co-L_0 iff U_x accepts x. On the other hand, by definition of L_0, we have $x \in$ co-L_0 iff U_x does not accept x. Contradiction. **Q.E.D.**

LEMMA 18. *The problem* EQUIV *requires nondeterministic time greater than c^n i.o., for some $c > 0$.*

Proof. The above lemma shows the existence of an L in $NTIME(n2^n) - \text{co-}NTIME(2^n)$. Since L is in $NEXPT$, chapter 5 shows that L is many-one reducible to INEQ via some log-linear transformation t. Furthermore, t has the property that for all x, $t(x)$ represents a well-formed pair of $\{+,\cdot,^2\}$-expressions. This implies co-L is many-one reducible to EQUIV via t. Choose $c = 2^{1/b}$ where t is bn-bounded, $b > 0$. Assume for the sake of contradiction that EQUIV $\in NTIME(c^n)$. Then lemma 16 (i), implies that co-$L \in NTIME(c^{bn}) = NTIME(2^n)$. This contradicts our assumption $L \notin \text{co-}NTIME(2^n)$. **Q.E.D.**

6.6 Weak Separation

The following result of Book [2] is useful in distinguishing between two classes:

THEOREM 19. *Let J, K be classes and $\leq$ a reducibility. Suppose K has a complete language under $\leq$-reducibility. If J is the limit of the some strictly increasing sequence*

$$J_1 \subset J_2 \subset J_3 \subset \cdots$$

where each J_i is closed under $\leq$-reducibility, then $J \neq K$.

Proof. Let L be K-complete under $\leq$-reducibility. If $J = K$ then $L \in J_i$ for some i. By the basic inclusion lemma, chapter 4, $K \subseteq J_i$. This contradicts the fact that J properly contains J_i. **Q.E.D.**

Of course, this theorem achieves only a weak separation between J and K, since it does not tell us if $J - K \neq \emptyset$ or $K - J \neq \emptyset$ (although one of these must hold true). We illustrate an application of the lemma:

THEOREM 20. *PLOG is distinct from NP and from P .*

Proof. We know that *NP* has complete languages under $\leq^L_m$. The technique (in chapter 4) for showing that *DLOG* is closed under log-space reducibility easily shows that $DSPACE(\log^k n)$ is closed under $\leq^L_m$ reducibility. *PLOG* is the limit of the increasing sequence

$$DSPACE(\log n) \subseteq DSPACE(\log^2 n) \subseteq DSPACE(\log^3 n) \subseteq \cdots .$$

[6] One is tempted to say, this problem is not in $NSPACE(o(O(1)^n))$. But notice that we have not defined the meaning of $o(E)$ where E is a big-Oh expression.

[7] EQUIV only differs from the complement of INEQ by being restricted to words that represent pairs of well-formed expressions.

By the deterministic space hierarchy theorem, we know that this sequence is strictly increasing. Hence the previous theorem applies showing that $NP \neq PLOG$. Similarly, since P also has complete languages under $\leq_m^L$, we also conclude $P \neq PLOG$. **Q.E.D.**

Consider the following attempt to show $DLOG \neq P$: we can define the strictly increasing sequence

$$DTIME(n) \subseteq DTIME(n^2) \subseteq DTIME(n^3) \subseteq \cdots$$

and using $\leq_m^L$ as our reducibility, etc., we find that one of the conditions of the lemma fails (where?).

The following chart from Book[2]: shows some of the known distinctions between the various classes.

	$PLOG$	(a)	$DLOG$	NP	(b)	(c)	P	(d)
$PLOG$								
$(a) = DSPACE(\log^k n)$	$\neq$							
$DLOG$	$\neq$	$\neq$						
NP	$\neq$	?	?					
$(b) = NTIME(n^k)$	$\neq$	$\neq$	$\neq$	$\neq$				
$(c) = NTIME(n+1)$	$\neq$	$\neq$	$\neq$	$\neq$	$\neq$			
P	$\neq$	?	?	?	$\neq$	$\neq$		
$(d) = DTIME(n^k)$	$\neq$	$\neq$	$\neq$	$\neq$	?	?	$\neq$	
$DTIME(n+1)$	$\neq$	$\neq$	$\neq$	$\neq$	$\neq$	$\neq$	$\neq$	$\neq$

Notes: k is any integer greater than 1. An entry "?" indicates that it is not known if the two classes are equal or not.

6.7 Strong Separation

In section 5 we show the 'infinitely often' (i.o.) type of lower bound on the complexity of languages. In this section we consider the 'almost every' (a.e.) version. It is important to realize that strong separation results only make sense for characteristic classes. Geske and Huynh have proved strong hierarchy theorems in this sense.

Observe that most natural problems do not seem to possess non-trivial a.e. lower bounds. For instance, the reader can easily be convinced after checking some cases that all *known* NP-complete problems have infinite subsets that are easily recognizable in polynomial time. It is unknown whether there are NP-complete problems without this property. One of the few examples of a natural problem that *may* possess a non-trivial a.e. lower bound is primality testing: it is unknown if there is an infinite subset of the prime numbers that can be recognized in deterministic polynomial time. More precisely[8], is there a language $L \in P$ such that $L \cap Primes$ is infinite?

Meyer and McCreight [20] shows that for any space-constructible complexity function $s(n) \geq n$, there exists a language whose *running* space complexity is lower bounded by $s(n)$ (a.e.). The following theorem adapts the proof for *accepting* space complexity and avoids the assumption $s(n) \geq n$.

THEOREM 21. *Let s be an non-decreasing, unbounded space-constructible function. Then there exists an infinite language L_0 in $DSPACE(s)$ such that if N is any acceptor for L_0,*

$$AcceptSpace_N(n) = \Omega(s(n)).$$

This result can be regarded as a strong version of the deterministic space hierarchy theorem for characteristic classes. Similar results have been obtained in [4, 31]. These proofs use ideas from the earlier work of Rabin [23] and Blum [1].

Proof. The language L_0 will be defined by describing an acceptor M for it. Let $U = \{U_0, U_1, \ldots\}$ be a universal machine for the class $RE|\{0,1\}$. The basic idea is for M to diagonalize over each U_i that accepts in space $cs(n)$ for some $c > 0$. More precisely, suppose the input to M is the binary string i (regarded as an integer when convenient). First we mark out exactly $s(|i|)$ tape squares on each of its work-tape. In the subsequent simulation, the computation will never exceed these marked squares. Hence $L(\mathrm{M})$ clearly is in $DSPACE(s)$.

In the following description, we see that M on input i will compute an index $\delta(i) \geq 0$ such that

[8] Recently Goldwasser and Killian show that there is an infinite subset that can be recognized in *expected* polynomial time. Expected complexity classes will be considered in a later chapter.

$$\begin{aligned}\delta(i) = \mathit{odd} &\Rightarrow \text{M accepts } i \text{ and } U_{\lfloor \delta(i)/2 \rfloor} \text{ rejects } i.\\ \delta(i) = \mathit{even} &\Rightarrow \text{M rejects } i \text{ and } U_{\lfloor \delta(i)/2 \rfloor} \text{ accepts } i.\end{aligned}$$

We say that the index $j \geq 0$ is 'cancelled' by input i if $j = \lfloor \delta(i)/2 \rfloor$; thus if j is cancelled, then $L(U_j) \neq L(\mathrm{M})$. Note that we try to cancel each j twice but of course, this is impossible if $L(U_j)$ is a trivial language. The binary string $\delta(i)$ will be written on a tape reserved for this purpose (say tape 1) at the end of the computation on i. So the machine M is doing double duty: as an acceptor as well as some kind of transducer (but only on the side).

Define $\hat{s}(n) = \max\{n, s(n)\}$ and let C_i denote the set

$$C_i = \{\delta(j) : 0 \leq j \leq \hat{s}(|i|)\}$$

Note that $C_i \subseteq C_{i+1}$ and we define C_∞ to be the union over all C_i.

Let the input to M be i. Our goal is to cancel some $k \notin C_i$. To do this, we successively submit each $k = 0, \ldots, \hat{s}(|i|)$ to the following 'test': we say k passes the test if it satisfies the following three conditions.

(i) First $k \notin C_i$.
(ii) If $k = \mathit{odd}$ then $U_{\lfloor k/2 \rfloor}$ does not accept i in space $s(|i|)$. This covers three possibilities: the machine either tries to use more than $s(|i|)$ space, or rejects within $s(|i|)$ space, or loops within $s(|i|)$ space.
(iii) If $k = \mathit{even}$ then $U_{\lfloor k/2 \rfloor}$ accepts i in space $s(|i|)$.

We claim that this test can be done in $s(|i|)$ space for each $k = 0, \ldots, \hat{s}(|i|)$. To do part (i) of the test, we check if $k = \delta(j)$ for each $j = 0, \ldots, \hat{s}(|i|)$. For each j, we determine $\delta(j)$ by recursively calling M on input j. (Note that the restriction that $j \leq \hat{s}(|i|)$ means that $j \leq |i| < i$ and hence there is no problem of self-reference; this is the reason we use $\hat{s}(n)$ instead of $s(n)$ in defining C_i.) Since $s(n)$ is non-decreasing, we do not use more than $s(|i|)$ space. For part (ii) and (iii) of the test, we see that M can decide whether $U_{\lfloor k/2 \rfloor}$ accepts i within space $s(|i|)$: in particular, we must be able to detect when $U_{\lfloor k/2 \rfloor}$ loops within space $s(|i|)$ (which we know how to do from chapter 2, section 9). Thus the test can indeed be carried out within the marked space.

If any of these k passes the test, we write this k on tape 1 (so $\delta(i) = k$) and we accept iff k is odd. Otherwise, every such value of k fails the test and we write 0 on tape 1 and reject the input i. (We may assume that the index 0 corresponds to a machine that accepts all its inputs.)

This completes the description of M. Now we must prove that our construction is correct. First we claim that M accepts infinitely many inputs because for each index k corresponding to a machine U_k that rejects all its inputs in constant space, there is some input $x = x(k)$ such that M on input x accepts and outputs $\delta(x) = 2k + 1$. This claim amounts to showing that $2k + 1 \in C_\infty$. Choose the smallest x such that C_x contains each $j < 2k + 1$ that will eventually be cancelled, i.e., C_x contains $C_\infty \cap \{0, \ldots, 2k\}$, and $\hat{s}(|x|) \geq 2k + 1$. (Such a choice can be found.) Then M on input x will eventually test $2k + 1$, and by choice of k, it will detect that U_k does not accept x within the space $s(|x|)$. Hence M accepts with output $2k + 1$ as desired.

Let N be any acceptor satisfying

$$\mathit{AcceptSpace}_{\mathrm{N}}(n) \leq c \cdot s(n) \text{ (i.o.)} \tag{4}$$

for each choice of $c > 0$. It remains show that N cannot accept the same language as M.

Since we have shown that M accepts infinitely many inputs, we may assume that $\mathit{AcceptSpace}_{\mathrm{N}}(n)$ is defined for infinitely many values of n. Now there is some $h \geq 0$ such that N accepts the language $L(U_h)$. By usual properties of universal machines,

$$\mathit{AcceptSpace}_{U_h}(n) \leq c_0 \cdot \mathit{AcceptSpace}_{\mathrm{N}}(n) \tag{5}$$

for some constant $c_0 > 0$. Choosing $c = 1/c_0$, inequalities (4) and (5) imply that $\mathit{AcceptSpace}_{U_h}(|x|) \leq c \cdot c_0 s(|x|) = s(|x|)$. So it suffices to show that h is cancelled (since $\lfloor \delta(x)/2 \rfloor = h$ implies that M on input x will accept iff U_h does not accept x in space $s(|x|)$ iff $x \notin L(U_h) = L(N)$.)

Since s is unbounded and non-decreasing, we may choose the smallest input x such that

(a) $\hat{s}(|x|) \geq 2h$,

(b) C_x contains all indices $k < 2h$ that are eventually cancelled, $C_\infty \cap \{0, \ldots, 2h - 1\} \subseteq C_x$.

(c) $x \in L(U_h)$.

Consider the action of M on such an input x: M will test each $k = 0, 1, \ldots, 2h-1$ and, by choice of x, each such k will fail the test. Thus M on input x will eventually test $2h$. If $2h \in C_x$ then h is cancelled already, and we are done. If $2h \notin C_x$ then our test calls for M to simulate U_h running on input x. But we just showed that U_h on $|x|$ uses at most $s(|x|)$ space. Since $x \in L(U_h)$, the test succeeds with M rejecting x and outputting $\delta(x) = 2h$. Thus h is cancelled after all. **Q.E.D.**

The preceding theorem shows a language L_0 that is hard (a.e.) for the characteristic class $DSPACE(s)$. In order to infer that problems reducible to L_0 are also hard (a.e.), we need some converse of lemma 16. First, we define a language (Σ, L) to be *invariant under padding* if there is a symbol $\# \in \Sigma$ such that for all $x \in \Sigma^*$, $x \in L$ iff $x\# \in L$. The following is due to Stockmeyer [29]. The present proof applies to running space only:

THEOREM 22. *Suppose L is reducible to L' via some log-linear transformation and L' is invariant under padding. If $s(n) \geq \log n$ is non-decreasing and the running space for L is at least $s(n)$ (a.e.) then for some $c > 0$, the running space for L' at least $s(cn)$ (a.e.).*

Proof. Suppose to the contrary that there is an acceptor N for L' such that for all $c > 0$, N uses running space less than $s(cn)$ (i.o.). Let $L \leq_m^{Llin} L'$ via some log-linear transformation t where $|t(x)| \leq b|x|$ for integer $b \geq 1$. We obtain an acceptor M for L from any acceptor N for L' as follows. On input x:

For $i = 0, 1, 2, \ldots$ **do**:
  For $j = 0, 1, 2, \ldots, b|x|$, **do**:
    If $|t(x)\#^j| > b|x|$
      then exit current for-loop (i.e. go to $i+1$);
    Simulate N on input $t(x) \cdot \#^j$ using only i space;
    If N attempts to use more than i space,
      then continue current for-loop (i.e. go to $j+1$);
    If N halts, **then** accept if N accepts, else reject;
  End
End

The outer for-loop is potentially infinite since i can grow arbitrarily large, but the inner for-loop is bounded by $b|x|$. It should be clear that M accepts L. Let $reduce(x)$ denote the set

$$\{t(x) \cdot \#^j : |t(x)\#^j| \leq b|x|,\ j \geq 0\}.$$

The basic idea of M is to reduce the decision on x to deciding a member of $reduce(x)$ for which N requires the least space. Because of padding, we see that $x \in L$ implies $reduce(x) \subseteq L'$ and $x \notin L$ implies $reduce(x) \cap L' = \emptyset$. To see the space usage of M, suppose $y \in reduce(x)$ requires the least space to compute. In the above simulation of N on $y = t(x) \cdot \#^j$, we assume that M does not explicitly store the word y, but uses counters to store the value j and the input head position on y. Then

$$RunSpace_{\mathrm{M}}(x) \leq RunSpace_{\mathrm{N}}(y) + O_1(\log |y|) \tag{6}$$

where the term $O_1(\log |y|)$ comes from storing the counters for j and the head position on y.

Choose $c = 1/b$. To derive our contradiction, we show that if N runs in less than $s(cn)$ (i.o.) then M will also use less than $s(n)$ space (i.o.), contradicting our assumption on L. Let E be an infinite set of 'easy lengths' for N. More precisely, for each input y with length $|y|$ in E, $RunSpace_{\mathrm{N}}(y) < s(cn)$. Let E' be the set of lengths of inputs for M that can be reduced to those with length is in E, i.e. for each input x with $|x| \in E'$ there is some $y \in reduce(x)$ with $|y| \in E$. Then (6) shows that for such x,

$$\begin{aligned} RunSpace_{\mathrm{M}}(x) &< s(c|y|) + O_1(\log |y|) \\ &= O_2(s(cb|x|)) = O_2(s(|x|)). \end{aligned}$$

Using linear compression of space, M can be modified to ensure that for all n in E',

$$RunSpace_{\mathrm{M}}(n) < s(n).$$

We also see that E' is infinite: for each $n \in E$, we have $\lceil cn \rceil \in E'$ since if $|x| = \lceil cn \rceil$ then some member of $reduce(x)$ has length n. This contradicts the assumption that the running space for L is $\geq s(n)$ (a.e.). **Q.E.D.**

The requirement that L' is invariant under padding cannot be omitted in the above theorem. To see this, suppose the language (Σ, L) requires running space $\geq s(n)$ (a.e.). Define the language

$$L' = \{xx : x \in L\} \cup \{x \in \Sigma^* : |x| = odd\}.$$

Clearly $L \leq_m^{Llin} L'$ but it is easy to design an acceptor for L' that uses no space for all inputs of odd length.

Lynch [17] shows that if a problem L is not in P then L contains a subset whose running time complexity is lower bounded by n^k (a.e.) for every $k \geq 1$; see Even, Selman and Yacobi [9] for a generalization of such results. Schnorr and Klupp [26] obtains a result similar to Lynch within a natural context: there is a subset of the language SAT that is a.e. hard. These results should remind the reader of the result of Ladner in chapter 4.

6.8 Inseparability Results

To complement the separation results, we now show some results that reveal limits to our attempts to separate complexity classes. These results are somewhat surprising because they only rely on certain rather simple properties that are seen to hold for typical complexity measures such as time and space. Blum [1] first formalized such properties of complexity measures.

Let $\Sigma = \{0, 1\}$ and let $U = \{U_i\}$ be a universal machine for some class $K = K|\Sigma$, and R be any language over the alphabet $\Sigma \cup \{\#\}$. We say that the pair (U, R) is a *Blum measure* for K if:

(B1) $x \in L(U_i)$ iff there is an m such that $i\#x\#m \in R$.

(B2) R is recursive and defines a partial function $r(i, x)$ in the sense that if $r(i, x) \downarrow$ (i.e. $r(i, x)$ is defined) then there is a unique $m \in \Sigma^*$ such that $i\#x\#m \in R$ and if $r(i, x) \uparrow$ (i.e. is undefined) then for all $m \in \Sigma^*$, $i\#x\#m \notin R$.

We also write $R_i(x)$ for $r(i, x)$. Intuitively, $R_i(x) = m$ implies that U_i on input x accepts using m units of some abstract resource. Note that if $x \notin L(U_i)$ then, consistent with our use of acceptance complexity, $R_i(x) \uparrow$. The reader may verify that we may interpret this abstract resource to be time, space, reversal or some other combinations (e.g. products) of these. One caveat is that for space resource, we must now say that the space usage is infinite or undefined whenever the machine loops (even if it loops in finite space).

The hierarchy theorems suggests a general theorem of this form: there is total recursive function $g(n)$ there exists for all total recursive functions $t(n)$, the pair

$$(DTIME(g \circ t), DTIME(t))$$

can be separated: $DTIME(g(t(n)) - DTIME(t(n)) \neq \emptyset$. Unfortunately, we now show that such a theorem is impossible without restrictions on $t(n)$. This is implied by the following theorem of Borodin [3, 18, page 148].

THEOREM 23 ((Gap lemma)). *Let (U, R) be a fixed Blum measure for all recursively enumerable languages. For any total recursive function $g(n) > n$ there exists a total recursive function $t(n)$ such for all i, there are only finitely many n such that $t(n) < R_i(n) < g(t(n))$.*

Proof. We define $t(n)$ to be the smallest value m satisfying predicate

$$P(n, m) \equiv (\forall i)[m < R_i(n) < g(m) \Rightarrow n \leq i]$$

As usual, when an integer n appears in a position (e.g. $R_i(n)$) that expects a string, we regard it as its binary representation. Note that this definition of $t(n)$, if total and recursive, would have the desired properties for our theorem. First we show that the predicate $P(n, m)$ is partial recursive. $P(n, m)$ can be rewritten as $(\forall i)[n \leq i$ or $\neg(m < R_i(n) < g(m))]$, or

$$(\forall i < n)[m \geq R_i(n) \text{ or } R_i(n) \geq g(m)].$$

It is now clear that $P(n, m)$ is decidable, and so, by searching for successive values of m, we get a partial recursive procedure for $t(n)$. To show that this procedure is total recursive, we show that for any n, there is at least one value of m satisfying $P(n, m)$. Define $m_0 = 0$ and for each $i = 1, 2, .$, define $m_i = g(m_{i-1})$. Hence $m_0 < m_1 < \cdots$. Consider the gaps between $m_0, m_1, \ldots, m_{n+1}$: there are $n + 1$ gaps so that at least one of them $[m_j, m_{j+1}]$ $(j = 0, \ldots, n)$ does not contain a value of the form $R_i(n)$, $i = 0, \ldots, n - 1$. Then $P(n, m_j)$ holds. **Q.E.D.**

For instance it easily follows that there exist complexity functions t_i $(i = 1, 2, 3)$ such that the following holds:

$$DTIME(t_1) = DTIME(2^{t_1})$$
$$DTIME(t_2) = NTIME(t_2)$$
$$DSPACE(t_3) = NSPACE(t_3)$$

Thus such pairs of classes are inseparable. Such results, like the union theorem mentioned in section 6, employ complexity functions that are highly non-constructible. Such *inseparability results* are instructive: it tells us that the constructibility conditions in our separation results cannot be removed with impunity. There are many other results of this nature (e.g., [24, 8]).

6.9 Conclusion

This chapter shows some basic separation results and uses translational methods to obtain further separations. We also illustrate techniques for stronger or weaker notions of separation. These results mostly apply to space and time classes: it would be satisfying to round up our knowledge in this 'classical' area by extending the results here to encompass reversal or simultaneous complexity classes. We also show that there are limits to such separations if we do not restrict ourselves to nice complexity functions.

Exercises

[6.1] Let t be any total time-constructible function. Show that there is time-constructible function t' such that $NTIME(t') - NTIME(t) \neq \emptyset$.

[6.2] Reprove the deterministic space hierarchy theorem for running complexity. The statement of the theorem is as before except that we do not assume s_2 to be space constructible.

[6.3] Prove a hierarchy theorem for classes of the form $D\text{-}TIME\text{-}SPACE(t, \log n)$ (for various t).

[6.4] (A distributed decrementing counter) Reconstruct Fürer's result stated in section 2. We want to simulate each $t_1(n)$-bounded machines in $t_2(n)$ steps. On input w, we want to simulate the machine M_w for $t_2(|w|)$ steps. To do this, we a construct a "decrementing-counter" initialized to the value of $t_2(|w|)$, and try to decrement this counter for each step of the simulation. The problem is that since we are restricted to $k \geq 2$ (which is fixed in Fürer's version of the theorem) tapes. So the counter must reside on one of the tracks of a work tape (which has to be used for the actual simulation of M_w. The idea is to use a "distributed representation of numbers" such that there are low order bits scattered throughout the representation (note that to decrement, it is easy to work on low order bits, and propagate the borrow to a higher order bit if necessary). So we form a balanced binary with nodes labeled by a digit between 0 and $B-1$ (B is a fixed base of the representation, and $B = 4$ suffices for us). A node at height $h \geq 0$ (the leaves are height 0) with a digit d represents a value of dB^h. The number represented is the sum of values at all nodes. Clearly there are lots of redundancy and all the leaves are low-order bits. We store the labels of the nodes in an in-order fashion so the root label is in the middle of the representation. We want to decrement at ANY leaf position. If necessary, borrows are taken from the parent of a node. When the path from any leaf to the root has only zero labels, then this representation is exhausted and we must redistribute values. Show that if we do a total of m decrements (as long as the tree is not exhausted) then the time is $O(m)$ on a Turing machine. Apply this to our stated theorem (we have to show that the tree is not too quickly exhausted and we must show how to reconstruct exhausted trees)

[6.5] (a) (Hong) Show that the transformation $w \mapsto tally(w)$ (where $w \in \{1, \ldots, k\}^*$ and $tally(w) \in \{1\}^*$) can be computed in space-reversals $O(n, \log n)$ and also in space-reversals $O(\log n, n)$.
(b) Generalize this result by giving tradeoffs between space and reversals.

[6.6] Show that for t any total time-constructible function then there exists a deterministic 1-tape acceptor M that time-constructs some t' such that

$$L(\mathrm{M}) \in DTIME(t') - \bigcup_{j=0}^{\infty} NTIME(t(n+j)).$$

[6.7] Separate the following pairs of classes:
(a) $DTIME(2^n)$ and $DTIME(3^n)$
(b) $NSPACE(2^n)$ and $NSPACE(3^n)$
(c) $NSPACE(2^{2^n})$ and $DSPACE(2^{3^n})$.

[6.8] Show the translational lemma for time resource analogous to lemma 5. Conclude that if $NP \subseteq P$ then $NEXPT \subseteq DEXPT$.

[6.9] (I. Simon) Assume the one-way oracle machines (section 3, chapter 4) for this question. Show that Savitch's upward space translation result can be relativized to any oracle. More precisely, for any oracle A, and space-constructible and moderately growing function s, $NSPACE^A(s) \subseteq DSPACE^A(s)$ implies that that for all $s'(n) \geq s(n)$, $NSPACE^A(s') \subseteq DSPACE^A(s)$.

[6.10] Where does the proof of Ibarra's theorem break down for nondeterministic time complexity?

[6.11] Complete the proof of Cook's result that $NTIME(n^r) - NTIME(n^s) \neq \emptyset$ from section 4: It is enough to show that for any $b = k2^{-m} \geq 1$ where k, m are positive integers, there exists a time-constructible function $t(n) = \Theta(n^b)$. **Hint:** Show that in linear time, you can mark out $\Theta(n^{1/2})$ cells. Extend this to $\Theta(n^{2^{-m}})$, and use the time-constructibility of the function n^k for any integer $k \geq 1$ and an exercise in chapter 2.

[6.12] (Recursion theorem for nondeterministic time) Let $\Sigma = \{0, 1, \$\}$ and let $U = \{U_i\}$ be a universal machine for the class $RE|\Sigma$. For all indices i, if U_i accepts in time t_0 then there exists an index j such that U_j accepts the language

$$\{x : j\$x \in L(U_i)\}$$

in time $O_i(1) + t_0(|j\$| + n)$. **Remark.** This lemma is used the Seiferas-Fisher-Meyer proof of the hierarchy theorem for nondeterministic time. This lemma is analogous to the second recursion theorem (also called the fixed point theorem) in recursive function theory.[9]

[6.13] (Žák) If t is time-constructible and $t(n+1) = O(t(n))$ then for each $c > 0$, $NTIME(t) - NTIME_U^{cut}(ct) \neq \emptyset$. Hint: use Žák's theorem 13.

[6.14] (Book) Show the time analogue of the space translation lemma for tally languages: Let $f(n) \geq 2^n$, $X = D$ or N. Then

$$L \in XTIME(f(O(n))) \iff tally(L) \in XTIME(f(O(\log n))).$$

[6.15] (Wilson) Show that if every tally language in $NP \cap$ co-NP belongs to P then $NEXPT \cap$ co-$NEXPT \subseteq DEXPT$. **Hint:** You may use the result of the previous exercise.

[6.16] (Book) Let $\Phi \subseteq$ **DLOG** (the logarithmic space computable transformations), $X = D$ or N. Show that if L is $XTIME(O(n))$-complete under $\leq_m^\Phi$ then L is $XTIME(n^{O(1)})$-complete under $\leq_m^{\mathbf{DLOG}}$.

[6.17] Provide the proofs for the rest of the entries in the chart at the end of section 5.

[6.18] (Book) Show that if $NP \subseteq P$ then $NEXPT \subseteq DEXPT$.

[6.19] (Book) Show that $NP|\{1\} \neq DSPACE(\log^{O(1)} n))|\{1\}$.

[6.20] Consider the family Φ of transformations computed in deterministic log-space and in linear time. So Φ is a subfamily of **Llin** in Chapter 4.
(i) Show if $L \in D\text{-}TIME\text{-}SPACE(n^i, \log n)$ and $L' \leq_m^\Phi L$ then $L' \in D\text{-}TIME\text{-}SPACE(n^{i+1}, \log n)$.
(ii) Show that $NTIME(n^2)$ has a complete language under $\leq_m^\Phi$.
(iii) Show that $NTIME(n^2) \neq DLOG$. HINT: use the hierarchy theorem in Exercise 6.3.

[6.21] (Loui) Try to imitate the proof of theorem 20 to attempt to show that $P \neq PSPACE$. Where does the proof break down?

[6.22] (open) Separate $NTIME(2^{2^n})$ from $NTIME(n2^{2^n})$.

[6.23] (open) Improve the nondeterministic time hierarchy theorem of Seiferas-Fischer-Meyer by changing the theorem statement from '$o(t(n-1))$' to '$o(t(n))$'.

[6.24] Suppose that the outline (a) and (b) in section 5 were stated for the class $DSPACE(F)$ in place of K, where F is a set of complexity functions. State the conditions on F for the same conclusion to hold.

[6.25] (Stockmeyer) Say a language (Σ, L) is *naturally padded* if there is a symbol $\# \notin \Sigma$, some integer j_0, and a logspace transformation $t : \Sigma^*\#^* \to \Sigma^*$ such that (a) $L \cdot \#^* \leq_m^L L$ via t, and (b) $|t(x\#^j)| = |x| + j$ for all $x \in \Sigma^*$ and $j \geq j_0$. Prove theorem 22 using this " natural padding" instead of the "invariance under padding" assumption.

[9] The second recursion theorem is as follows: let U be a universal machine for $RE|\{0,1\}$. Then for every total recursive function f there exists an index i such that $U_i(x) = U_{f(i)}(x)$ for all x. As pointed out in the excellent exposition in [7, chapter 11], this theorem is the quintessence of diagonal arguments; it can be used (perhaps as an overkill) to give short proofs for many diagonalization results.

[6.26] (Stockmeyer) Suppose L is Karp-reducible to L' via a polynomial time transformation t that is linearly bounded. If L has an $t(n)$ (a.e.) lower bound on its acceptance time, show a corresponding a.e. lower bound on L'.

[6.27] Define an even stronger notion of separation: the characteristic classes $DTIME(t')$ and $DTIME(t)$ are said to be *very strongly separated* if there exists an infinite language L in $DTIME(t') - DTIME(t)$ such that for any acceptor M for L, $AcceptTime_{\mathrm{M}}(x) > t(|x|)$ for almost every input $x \in L(\mathrm{M})$. (In other words, we use 'a.e. x' rather than 'a.e. n'.) Extend the proof of theorem 21 to this setting of very strong separation.

[6.28] (Even, Long, Yacobi) Say L is *easier than* L' if for all N accepting L' there is an M accepting L such that

1. $AcceptTime_{\mathrm{M}}(x) = |x|^{O(1)} + AcceptTime_{\mathrm{N}}(x)$ for all input x.
2. For some $k > 0$, $(AcceptTime_{\mathrm{M}}(x))^k \leq AcceptTime_{\mathrm{N}}(x)$ for infinitely many x.

If only (1) is satisfied then we say that L is *not harder than* L'. Show that if there exists a language $L_0 \in NP - \text{co-}NP$ then:
(a) There exists a recursive language $L_1 \notin NP \cup \text{co-}NP$ such that L_0 is not harder than L_1 and L_1 is not harder than L_0.
(b) There exists a recursive language $L_2 \notin NP \cup \text{co-}NP$ such that L_1 is easier than L_0.

[6.29] (Ming Li) Show that for all $j > 1$, $NTIME(n^j) - D\text{-}TIME\text{-}SPACE(n^j, o(n))$ is non-empty.

[6.30] (Ming Li) Show that there is a language in $NTIME(n)$ that cannot be accepted by any deterministic simple Turing machine in time $O(n^{1.366})$.

Bibliography

[1] M. Blum. A machine-independent theory of the complexity of recursive functions. *Journal of the ACM*, 14(2):322–336, 1967.

[2] R. V. Book. Translational lemmas, polynomial time, and $(\log\ n)^j$-space. *Theor. Computer Science*, 1:215–226, 1976.

[3] A. Borodin. Computational complexity and the existence of complexity gaps. *Journal of the ACM*, 19(1):158–174, 1972.

[4] A. Borodin, R. Constable, and J. Hopcroft. Dense and non-dense families of complexity classes. *10th Proc. IEEE Symp. Found. Comput. Sci.*, pages 7–19, 1969.

[5] J. Chen and C.-K. Yap. Reversal complexity. *SIAM J. Computing*, to appear, 1991.

[6] S. A. Cook. A hierarchy for nondeterministic time complexity. *Journal of Computers and Systems Science*, 7:343–353, 1973.

[7] N. J. Cutland. *Computability: an introduction to recursive function theory*. Cambridge University Press, 1980.

[8] S. Even, T. J. Long, and Y. Yacobi. A note on deterministic and nonderministic time complexity. *Information and Control*, 55:117–124, 1982.

[9] S. Even, A. L. Selman, and Y. Yacobi. Hard-core theorems for complexity classes. *Journal of the ACM*, 32(1):205–217, 1985.

[10] M. Fürer. The tight deterministic time hierarchy. *14th Proc. ACM Symp. Theory of Comp. Sci.*, pages 8–16, 1982.

[11] J. Hartmanis, P. M. L. II, and R. E. Stearns. Hierarchies of memory limited computations. *IEEE Conf. Record on Switching Circuit Theory and Logical Design*, pages 179–190, 1965.

[12] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Trans. Amer. Math. Soc.*, 117:285–306, 1965.

[13] H. B. Hunt, III. On the time and tape complexity of languages, I. *5th Proc. ACM Symp. Theory of Comp. Sci.*, pages 10–19, 1973.

[14] O. Ibarra. A note concerning nondeterministic tape complexities. *J. ACM*, 19:608–612, 1972.

[15] N. Immerman. Nondeterministic space is closed under complement. *Structure in Complexity*, 3:112–115, 1988.

[16] M. Li. Some separation results. Manuscript, 1985.

[17] N. A. Lynch. On reducibility to complex or sparse sets. *Journal of the ACM*, 22:341–345, 1975.

[18] M. Machtey and P. Young. *An Introduction to the General Theory of Algorithms*. Elsevier North Holland, New York, 1978.

[19] A. R. Meyer. Weak monadic second order theory of successor is not elementary-recursive. In Dold and E. (eds.), editors, *Logic Colloquium: Symposium on Logic Held at Boston University, 1972-73*, pages 132–154. Springer-Verlag, 1975.

[20] A. R. Meyer and E. M. McCreight. Computationally complex and pseudo-random zero-one valued functions. In Z. Kohavi and A. Paz, editors, *Theory of machines and computations*, pages 19–42. Academic Press, 1971.

[21] A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. *13th Proc. IEEE Symp. Found. Comput. Sci.*, pages 125–129, 1972.

[22] W. Paul. *Komplexitaetstheorie*. Teubner, Stuttgart, 1978.

[23] M. Rabin. Degree of difficulty of computing a function. Technical Report Tech. Report 2, Hebrew Univ., 1960.

[24] C. W. Rackoff and J. I. Seiferas. Limitations on separating nondeterministic complexity classes. *SIAM J. Computing*, 10(4):742–745, 1981.

[25] S. Ruby and P. C. Fischer. Translational methods in computational complexity. *6th IEEE Conf. Record on Switching Circuit Thoery, and Logical Design*, pages 173–178, 1965.

[26] C. P. Schnorr and H. Klupp. A universally hard set of formulae with respect to non-deterministic Turing acceptors. *IPL*, 6(2):35–37, 1977.

[27] J. I. Seiferas, M. J. Fischer, and A. R. Meyer. Refinements of the nondeterministic time and space hierarchies. *14th Annual Symposium on Switching and Automata Theory*, pages 130–136, 1973.

[28] J. I. Seiferas, M. J. Fischer, and A. R. Meyer. Separating nondeterministic time complexity classes. *Journal of the ACM*, 25(1):146–167, 1978.

[29] L. J. Stockmeyer. The complexity of decision problems in automata theory and logic. Technical Report Project MAC Tech. Rep. TR-133, M.I.T., 1974. PhD Thesis.

[30] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time. *5th Proc. ACM Symp. Theory of Comp. Sci.*, pages 1–9, 1973.

[31] B. A. Trachtenbrot. On autoreducibility. *Soviet Math. Dokl.*, 11(3):814–817, 1970.

[32] S. Žák. A Turing machine time hierarchy. *Theor. Computer Science*, 26:327–333, 1983.

Chapter 7

Alternating Choices

April 13, 2009

In this chapter, we introduce the general concept of computational choices. The general mechanism for ascribing "acceptance" to such a computation is based on the theory of valuation. Values are viewed as probability intervals $[u, v] \subseteq [0, 1]$. Computational complexity can also be based on this theory. After some general results about valuations, the rest of the chapter focuses on a simple form of choice, called alternation.

7.1 Introducing Computation with Choice

The choice-mode of computation comes in two main flavors. The first is based on probability and briefly discussed in Chapter 1 (Section 6.2). The second is a generalization of nondeterministism called *alternation*. Let us briefly see what an alternating computation looks like. Let δ be the usual Turing transition table that has choice and let $C_0(w)$ denote the initial configuration of δ on input w. For this illustration, assume that every computation path is finite; in particular, this implies that no configuration is repeated in a computation path. The computation tree $T(w) = T_\delta(w)$ is defined in the obvious way: the nodes of $T(w)$ are configurations, with $C_0(w)$ as the root; if configuration C is a node of $T(w)$ and $C \vdash C'$ then C' is a child of C in $T(w)$. Thus the leaves of $T(w)$ are terminal configurations. The description of an alternating machine M amounts to specifying a transition table δ together with an assignment γ of a Boolean function $\gamma(q) \in \{\wedge, \vee\neg\}$ to each state q in δ. This induces a Boolean value on each node of $T(w)$ as follows: the leaves are assigned 1 or 0 depending on whether the configuration is accepting or not. If C is not a leaf, and q is the state in C, then we require that the number of children of C is equal to the arity of $\gamma(q)$. For instance, if C has two children whose assigned values are x and y then C is assigend the value $\gamma(q)(x, y)$. Finally we say M accepts w if the root $C_0(w)$ is assigned value 1.

The reader will see that nondeterministic computation corresponds to the case where $\gamma(q) = \vee$ for all q. Since the introduction of alternating machines by Chandra, Kozen and Stockmeyer[3] in 1978, the concept has proven to be an extremely useful tool in Complexity Theory.

The model of probabilistic machines we study was introduced by Gill[7]. Let us rephrase the description of probabilistic computation in Chapter 1 in terms of assigning values to nodes of a computation tree. A probabilistic machine is formally a transition table δ where each configuration spawns either zero or two children. For any input w, we again have the usual computation tree $T(w)$. The leaves of $T(w)$ are given a value of 0 or 1 as in the alternating case. However, an internal node u of $T(w)$ is assigned the average $(x + y)/2$ of the values x, y of the two children of u. The input w is accepted if the root is assigned a value greater than 1/2. (The reader should be verify that this description is equivalent to the one given in Chapter 1.) The function $f(x, y) = (x + y)/2$ is called the *toss* function because in probabilistic computations, making choices is interpreted as branching according to the outcomes of tossing a fair coin.

Hence, a common feature of probabilistic and alternating modes is their systematic bottom-up method of assigning values to nodes of computation trees. One difference is that, whereas probabilistic nodes are given (rational) values between 0 and 1, the alternating nodes are assigned Boolean values. We modify this view of alternating machines by regarding the Boolean values as the real numbers 0 and 1, and generalizing the Boolean functions $\wedge$, $\vee$ and $\neg$ to the real functions min, max and $f(x) = 1 - x$ (respectively).

With this shift of perspective, we have almost accomplished the transition to a new syncretistic model that we call *probabilistic-alternating machines*. This model was first studied in [23]. A probabilistic-alternating machine M is specified by giving a transition table δ and each state is associated with one of the four real functions

$$\min(x, y), \qquad \max(x, y), \qquad 1 - x, \qquad \frac{x + y}{2}. \tag{1}$$

We require c onfiguration in state q to spawn m children where m is the arity of the function associated with q. This can be enforced by synthetic restrictions on the transition table δ. Given an input w, we construct the tree $T(w)$ and assign values to its nodes in the usual bottom-up fashion (again, assuming $T(w)$ is a finite tree). We say M accepts the input w if the value at the root of $T(w)$ is $> 1/2$.

Probabilistic and alternating machines in the literature are usually studied independently. In combining these two modes, we extend results known for only one of the modes, or unify distinct results for the separate modes. More importantly, it paves the way towards a general class of machines that we call *choice machines*. Computations by choice machines are characterized by the systematic assignment of 'values' to nodes of computation trees, relative to the functions $\gamma(q)$ which the machine associates to each state q. These functions are similar to those in (1), although an immediate question is what properties should these functions satisfy? This will be answered when the theory is developed. We call any assignment of "values" to the nodes of a computation tree a *valuation*.[1] Intuitively, these values represent probabilities and lies in the unit interval $[0, 1]$. But because of infinite computation trees, we are forced to take as "values" any subinterval $[a, b]$ of the unit interval $[0, 1]$. Such intervals represent uncertainty ranges in the probabilities. This leads to the use of a simple interval algebra. The present chapter develops the valuation mechanism needed for our theory of choice machines. We will specifically focus on alternation machines, leaving stochastic machines to the next chapter.

Other choice modes. Other authors independently proposed a variety of computational modes that turn out to be special cases of our probabilistic-alternating mode: **interactive proof systems** (Goldwasser, Micali and Rackoff [8]), **Arthur-Merlin games** (Babai [2]), **stochastic Turing machines** (Papadimitriou [16]), **probabilistic-nondeterministic machines** (Goldwasser and Sipser [9]). In general, communication protocols and game playing models can be translated as choice machines. In particular, this holds for the **probabilistic game automata** (Condon and Ladner [4]) which generalize interactive proof systems and stochastic machines[2]. Alternating machines are generalized to **logical type machines** (Hong [11]) where machine states can now be associated with any of the 16 Boolean functions on two variables. Some modes bearing little resemblance to choice machines can nevertheless be viewed as choice machines: for example, in Chapter 9 we describe a choice mode that generalizes nondeterminism in a different direction than alternation. (This gives rise to the so-called *Boolean Hierarchy*.) These examples suggests that the theory of valuation gives a proper foundation for choice modes of computation. The literature can avoid our systematic development only by restrictions such as requiring constructible time-bounds.

7.2 Interval Algebra

The above introduction to probabilistic-alternating machines explicitly avoided infinite computation trees $T(x)$. Infinite trees cannot be avoided in general; such is the case with space-bounded computations or with probabilistic choices. In particular, a computation using finite amount of space may have infinite computation paths. To see why infinite trees are problematic, recall that we want to systematically assign a value in $[0, 1]$ to each node of $T(x)$, in a bottom-up fashion. But if a node u of $T(x)$ lies on an infinite path, it is not obvious what value to assign to u.

Our solution [23] lies in assigning to u the smallest 'confidence' interval $I(u) \subseteq [0, 1]$ guaranteed to contain the 'true' value of u. This leads us to the following development of an *interval algebra*[3].

In the following, u, v, x, y, etc., denote real numbers in the unit interval [0, 1]. Let

$$INT := \{[u, v] : 0 \leq u \leq v \leq 1\}$$

denote the set of closed subintervals of $[0, 1]$. An interval $[u, v]$ is *exact* if $u = v$, and we identify the exact interval $[u, u]$ with the real number u. We call u and v (respectively) the *upper* and *lower bounds* of the interval $[u, v]$. The unit interval $[0, 1]$ is also called *bottom* and denoted $\perp$.

[1] The term 'valuation' in algebra refers to a real function on a ring that satisfies certain axioms. Despite some concern, we will expropriate this terminology, seeing little danger of a context in which both senses of the term might be gainfully employed.

[2] This game model incorporates 'partially-hidden information'. It will be clear that we could add partially-hidden information to choice machines too.

[3] *Interval arithmetic*, a subject in numerical analysis, is related to our algebra but serves a rather different purpose. We refer the reader to, for example, Moore [15].

By an *interval function* we mean a function $f : INT^n \to INT$, where $n \geq 0$ denotes the arity of the function. We are interested in six interval functions. The first is the unary function of *negation* ($\neg$), defined as follows:

$$\neg[x, y] = [1 - y, 1 - x].$$

The remaining five are binary functions:

minimum ($\wedge$), *maximum* ($\vee$),
toss ($\fatsemi$),
probabilistic-and ($\otimes$), *probabilistic-or* ($\oplus$).

It is convenient to first define them as real functions. The real functions of minimum and maximum are obvious. The toss function is defined by

$$x \fatsemi y := \frac{x + y}{2}.$$

We saw in our introduction how this function arises from probabilistic (coin-tossing) algorithms. The last two functions are defined as follows:

$$\begin{aligned} x \otimes y &:= xy \\ x \oplus y &:= x + y - xy \end{aligned}$$

Thus $\otimes$ is ordinary multiplication of numbers but we give it a new name to signify the interpretation of the numbers as probabilities. If E is the event that *both* E_1 *and* E_2 occur, then the probability $\Pr(E)$ of E occurring is given by

$$\Pr(E) = \Pr(E_1) \otimes \Pr(E_2).$$

We assume that E_1, E_2 are independent events. Similarly $\oplus$ has a probabilistic interpretation: if E is the event that *either* E_1 *or* E_2 occurs, then

$$\Pr(E) = \Pr(E_1) \oplus \Pr(E_2).$$

To see this, simply note that $x \oplus y$ can also be expressed as $1 - (1 - x)(1 - y)$. For brevity, we suggest reading $\otimes$ and $\oplus$ as 'prand' and 'pror', respectively.

We note that these 5 real functions can also be regarded as functions on $[0, 1]$ (*i.e.*, if their arguments are in $[0, 1]$ then their values remain in $[0, 1]$). We may then extend them to the subintervals INT of the unit interval as follows. If $\circ$ is any of these 5 functions, then we define

$$[x, y] \circ [u, v] := [(x \circ u), (y \circ v)].$$

For instance, $[x, y] \otimes [u, v] = [xu, yv]$ and $[x, y] \wedge [u, v] = [\min(x, u), \min(y, v)]$. Alternatively, for any continuous function $f : [0, 1] \to [0, 1]$, we extend the range and domain of f from $[0, 1]$ to INT by the definition $f(I) = \{f(x) : x \in I\}$. If f is also monotonic, this is equivalent to the above.

One easily verifies:

LEMMA 1. *All five binary functions are commutative. With the exception of* $\fatsemi$*, they are also associative.*

The set INT forms a lattice with $\wedge$ and $\vee$ as the join and meet functions, respectively[4]. It is well-known that we can define a partial order $\leq$ in any lattice by:

$$[x, y] \leq [u, v] \iff ([x, y] \wedge [u, v]) = [x, y]. \tag{2}$$

Note that (2) is equivalent to:

$$[x, y] \leq [u, v] \iff x \leq u \text{ and } y \leq v.$$

When we restrict this partial ordering to exact intervals, we get the usual ordering of real numbers. For reference, we will call $\leq$ the *lattice-theoretic ordering* on INT.

The negation function is not a complementation function (in the sense of Boolean algebra [5]) since neither $I \wedge \neg I = 0$ nor[5] $I \vee \neg I = 1$ holds for all $I \in INT$. However it is idempotent, $\neg\neg I = I$. Probabilistic-and and probabilistic-or can be recovered from each other in the presence of negation. For example,

$$I \otimes J = \neg(\neg I \oplus \neg J).$$

It easy to verify the following forms of de Morgan's law:

[4] A lattice X has two binary functions, join and meet, satisfying certain axioms (essentially all the properties we expect from max and min). Lattice-theoretic notations can be found, for instance, in [5]. The lattice-theoretic properties are not essential for the development of our results.

[5] We assume that $\neg$ has higher precedence than the binary operators so we may omit parenthesis when convenient.

LEMMA 2.

$$\begin{aligned}
\neg(I \wedge J) &= \neg I \vee \neg J \\
\neg(I \vee J) &= \neg I \wedge \neg J \\
\neg(I \oplus\!\!\!\!\ominus J) &= \neg I \oplus\!\!\!\!\ominus \neg J \\
\neg(I \otimes J) &= \neg I \oplus \neg J \\
\neg(I \oplus J) &= \neg I \otimes \neg J
\end{aligned}$$

where $I, J \in INT$.

In view of these laws, we say that the functions $\wedge$ and $\vee$ are *duals of each other* (with respect to negation); similarly for the pair $\otimes$ and $\oplus$. However, $\oplus\!\!\!\!\ominus$ is self-dual.

We verify the distributivity of $\wedge$ and $\vee$ with respect to each other:

$$\begin{aligned}
I \vee (J_1 \wedge J_2) &= (I \vee J_1) \wedge (I \vee J_2) \\
I \wedge (J_1 \vee J_2) &= (I \wedge J_1) \vee (I \wedge J_2).
\end{aligned}$$

Furthermore, $\oplus\!\!\!\!\ominus$, $\otimes$ and $\oplus$ each distributes over both $\wedge$ and $\vee$:

$$\begin{aligned}
I \oplus\!\!\!\!\ominus (J_1 \wedge J_2) &= (I \oplus\!\!\!\!\ominus J_1) \wedge (I \oplus\!\!\!\!\ominus J_2), & I \oplus\!\!\!\!\ominus (J_1 \vee J_2) &= (I \oplus\!\!\!\!\ominus J_1) \vee (I \oplus\!\!\!\!\ominus J_2) \\
I \otimes (J_1 \wedge J_2) &= (I \otimes J_1) \wedge (I \otimes J_2), & I \otimes (J_1 \vee J_2) &= (I \otimes J_1) \vee (I \otimes J_2) \\
I \oplus (J_1 \wedge J_2) &= (I \oplus J_1) \wedge (I \oplus J_2), & I \oplus (J_1 \vee J_2) &= (I \oplus J_1) \vee (I \oplus J_2)
\end{aligned}$$

However $\otimes$ and $\oplus$ do not distribute with respect to each other (we only have $x \otimes (y \oplus z) \leq (x \otimes y) \oplus (x \otimes z)$). And neither $\wedge$ nor $\vee$ distributes over $\oplus\!\!\!\!\ominus$, $\otimes$ or $\oplus$.

Another Partial Order. For our applications, it turns out that a more useful partial order on INT is $\sqsubseteq$, defined by:

$$[x, y] \sqsubseteq [u, v] \iff x \leq u \text{ and } v \leq y.$$

Clearly $\sqsubseteq$ is the reverse of the set inclusion relation between intervals: $I \sqsubseteq J \iff J \supseteq I$ as sets. With respect to the $\sqsubseteq$-ordering, all exact intervals are maximal and pairwise incomparable[6]. In view of our interpretation of intervals as 'intervals of confidence', if $I \sqsubseteq J$ then we say J has 'at least as much information' as I. For this reason, we call $\sqsubseteq$ the *information-ordering*. In contrast to the lattice-theoretic $\leq$-ordering, $\sqsubseteq$ only gives rise to a lower semi-lattice with the meet function $\sqcap$ defined by

$$[x, y] \sqcap [u, v] = [\min(x, u), \max(y, v)].$$

(The following suggestion for defining the join $\sqcup$ fails: $[x, y] \sqcup [u, v] = [\max(x, u), \min(y, v)]$.) Note that bottom $\bot$ is the least element ("no information") in the information-ordering.

Example 1. The strong 3-valued algebra described[7] by Chandra, Kozen and Stockmeyer [3] is a subalgebra of our interval algebra, obtained by restricting values to $\{0, 1, \bot\}$. See figure 7.1 for its operation tables. They only

$\wedge$	0	1	$\bot$
0	0	0	0
1	0	1	$\bot$
$\bot$	0	$\bot$	$\bot$

$\vee$	0	1	$\bot$
0	0	1	$\bot$
1	1	1	1
$\bot$	$\bot$	1	$\bot$

Figure 7.1: The strong 3-valued algebra.

[6] I and J are $\sqsubseteq$-*comparable* if $I \sqsubseteq J$ or $J \sqsubseteq I$, otherwise they are $\sqsubseteq$-*incomparable*.
[7] Attributed to Kleene.

were interested in the functions $\wedge$, $\vee$, $\neg$. Thus our interval algebra gives a model (interpretation) for this 3-valued algebra.

The contrast between $\leq$ and $\sqsubseteq$ is best exemplified by the respective partial orders restricted to this 3-valued algebra, put graphically:

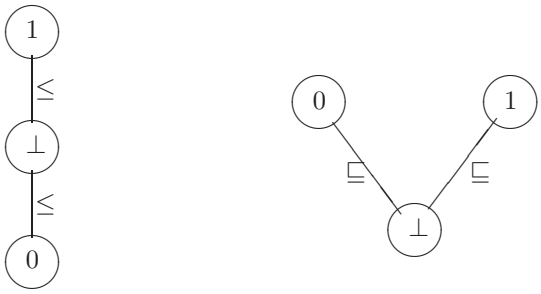

■

This information ordering gives rise to some important properties of interval functions. Given a sequence of $\sqsubseteq$ -increasing intervals

$$I_1 \sqsubseteq I_2 \sqsubseteq \cdots,$$

we define its **limit** in the natural way: $\lim_{j\geq 1} I_j$ is just $\cap_{j\geq 1} I_j$. Note that we only define this operation for increasing sequences, even though $\cap_{j\geq 1} I_j$ is defined without restrictions.

Definition 1.
(i) An n-ary function

$$f : INT^n \to INT$$

is *monotonic* if for all intervals $J_1, \ldots, J_n, J'_1, \ldots, J'_n$:

$$J_1 \sqsubseteq J'_1, \ldots, J_n \sqsubseteq J'_n \;\Rightarrow\; f(J_1, \ldots, J_n) \sqsubseteq f(J'_1, \ldots, J'_n).$$

(ii) f is *continuous* if it is monotonic and for all non-decreasing sequences

$$J_i^{(1)} \sqsubseteq J_i^{(2)} \sqsubseteq J_i^{(3)} \sqsubseteq \cdots$$

$(i = 1, \ldots, n)$, we have that

$$f(\lim_j\{J_1^{(j)}\}, \ldots, \lim_j\{J_n^{(j)}\}) = \lim_j f(J_1^{(j)}, \ldots, J_n^{(j)}). \tag{3}$$

■

Note that continuous functions are assumed monotonic. This ensures that the limit on the right-hand side of (3) is meaningful because monotonicity of f implies

$$f(J_1^{(1)}, \ldots, J_n^{(1)}) \sqsubseteq f(J_1^{(2)}, \ldots, J_n^{(2)}) \sqsubseteq f(J_1^{(3)}, \ldots, J_n^{(3)}) \sqsubseteq \cdots.$$

LEMMA 3. *The six functions* $\wedge$, $\vee$, $\overline{\wedge}$, $\otimes$, $\oplus$ *and* $\neg$ *are continuous.*

We leave the proof as an exercise. Continuity of these functions comes from continuity of their real counterpart.

Example 2. The *cut-off function* $\delta_{\frac{1}{2}}(x)$ is defined to be 1 if $x > \frac{1}{2}$ and 0 otherwise. We extend this function to intervals in the natural way: $\delta_{\frac{1}{2}}([u, v]) = [\delta_{\frac{1}{2}}(u), \delta_{\frac{1}{2}}(v)]$. This function is monotonic but not continuous. To see this, let $I_i = [0, \frac{1}{2} + \frac{1}{i}]$. Then $\lim_i I_i = [0, \frac{1}{2}]$ but

$$0 = \delta_{\frac{1}{2}}([0, \frac{1}{2}]) \neq \lim_i \delta_{\frac{1}{2}}(I_i) = \bot.$$

■

The following simple observation is useful for obtaining monotonic and continuous functions.

LEMMA 4.
(i) Any composition of monotonic functions is monotonic.
(ii) Any composition of continuous functions is continuous.

EXERCISES

Exercise 0.1: Verify the above identities of interval algebra. ♢

Exercise 0.2: Some additional properties of the lattice *INT*:
(a) Show that the two orderings $\leq$ and $\sqsubseteq$ are "complementary" in the following sense: for all I and J, either I and J are $\leq$-comparable or they are $\sqsubseteq$-comparable.
(b) Show that I and J are both $\leq$-comparable and $\sqsubseteq$-comparable iff $I \approx J$ where we write $[x, y] \approx [x, v]$ if $x = u$ or $y = v$.
(c) Extend $\wedge$ and $\vee$ to arbitrary sets $S \subseteq INT$ of intervals: denote the meet and join of S by $\wedge S$ and $\vee S$. Show that *INT* forms a complete lattice with least element 0 and greatest element 1. ♢

Exercise 0.3: Consider the 2-ary Boolean function *inequivalence* (also known as *exclusive-or*) denoted $\not\equiv$. We want to extend this function to intervals in *INT*. One way to do this is to use the equation

$$x \not\equiv y = (x \wedge \neg y) \vee (\neg x \wedge y)$$

valid for Boolean values, but now interpreting the $\wedge, \vee$ and $\neg$ as interval functions. For instance,

$$\begin{aligned} [0.2, 0.5] \not\equiv [0.6, 0.7] &= ([0.2, 0.5] \wedge [0.3, 0.4]) \vee ([0.5, 0.8] \wedge [0.6, 0.7]) \\ &= [0.2, 0.4] \vee [0.5, 0.7] = [0.5, 0.7]. \end{aligned}$$

Can you find other equations for $\not\equiv$ that are valid in the Boolean case such that the extension to intervals are not the same function? ♢

END EXERCISES

7.3 Theory of Valuations

To give a precise definition for the choice mode of computation, and its associated complexity measures, we introduce the theory of valuations. We also signal a new emphasis through a terminology change:

- **Global Decisions**. In Chapter 0, a non-accepting computation is classified as either rejecting or looping. Now, the disposition of an acceptor with respect to any input is one of the following: *acceptance*, *rejection* or *indecision*. These three dispositions are called "global decisions" because they are based on the entire computation tree, using the notion of valuations. Intuitively, looping is now replaced by indecision but the two concepts are not identical.

- **Local Answers**. Each computation path leads to a "local answer". If the path is finite, the local answer resides in the terminal configuration C of the path. Previously, the local answers are either accept or not-accept. We now have three local answers from C: **YES**, **NO** or **YO**. If the path is infinite, the local answer is **YO**. These answers are analogues of the three global decisions. In particular "YO" is neither Yes or No, i.e., indecision.

To implement the local answers, we introduce two distinguished states,

$$q_Y, q_N \in Q_\infty.$$

called the *YES-state* and *NO-state*, respectively, These are replacements for the previous accept (q_a) and reject (q_r) states. We can further arrange transition tables so configurations with these states are terminal. A terminal configuration C is called a *YES-configuration*, *NO-configuration* or a *YO-configuration*, depending on whether its state is q_Y, q_N or some other state. A computation path is called a *YES-path*, *NO-path* or *YO-path*, depending on whether it terminates in a YES, NO or otherwise. Thus a YO-path either terminates in a YO-configuration or is non-terminating.

Name	Basis B	Mode Symbol
deterministic	$\{\iota\}$	D
nondeterministic	$\{\vee\}$	N
probabilistic	$\{\oplus\!\!\!\!-\}$	Pr
alternating	$\{\wedge, \vee, \neg\}$	A
interactive proofs	$\{\oplus\!\!\!\!-, \vee\}$	Ip
probabilistic-alternating	$\{\oplus\!\!\!\!-, \wedge, \vee, \neg\}$	PrA
stochastic	$\{\oplus\!\!\!\!-, \otimes, \oplus, \neg\}$	St
stochastic-alternating	$\{\oplus\!\!\!\!-, \otimes, \oplus, \wedge, \vee, \neg\}$	StA

Figure 7.2: Some Choice Modes and their Symbols.

Naturally, the global decision is some generalized average of the local answers. This global/local terminology anticipates the quantitative study of errors in a computation (see Chapter 8). For now, it suffices to say that all error concepts are based on the discrepancies between global decisions and local answers. The seemingly innocuous introduction of YO-answers[8] is actually critical in our treatment of errors.

Let f be **interval function**, *i.e.*,

$$f \,:\, INT^n \to INT.$$

where $n \in \mathbb{N}$ is the arity of f. The 0-ary functions are *constant functions*, and the *identity function* $\iota(I) = I$ has arity 1. The constant functions are automatically continuous.

Definition 2. A set B of functions on INT is called a **basis set** if the functions in B are continuous. For any basis set B, a B-*acceptor* is a triple $M = (\delta, \gamma, \overset{M}{<})$ where δ is a Turing transition table whose state set Q_M is ordered by a total ordering $\overset{M}{<}$, and γ associates a basis function $\gamma(q)$ to each state $q \in Q_M$,

$$\gamma \,:\, Q \to B.$$

Moreover, δ has the property that if C is a configuration of δ in state q and $\gamma(q)$ has arity n, then C either is a terminal configuration or has exactly n immediate successors $C_1, \ldots, C_n$ such that the C_i's have distinct states. ■

The size of a basis set B can be finite or infinite. For example, if $B_1 = \{\iota\}$, then a B_1-machine is just deterministic Turing machine. In all the basis sets that arise naturally, the identity function ι is easily simulated, and so we might as well explicitly put them in. As another example, suppose B contains, among other things, each of the constant function $f(I) = c$, for $0 \leq c \leq 1$. Then a B-machine can assign an arbitrary value c to each leaf. However, since a B-machine has a finite number of states, only a finite number of such values c is used.

Operations of a B-machine. If the immediate successors of a configuration C are $C_1, \ldots, C_n$ such that the state of C_i is less than the state of C_{i+1} (under the ordering $\overset{M}{<}$) for each $i = 1, \ldots, n-1$, then we indicate this by writing[9]

$$C \vdash (C_1, \ldots, C_n).$$

If q is the state of C, we also write $\gamma(C)$ or γ_C instead of $\gamma(q)$. We require that $C_1, \ldots, C_n$ be distinct states because the value of the node (labeled by) C in the computation tree is given by $\gamma_C(v_1, \ldots, v_n)$ where v_i is the value of the node (labeled by) C_i. Without an ordering such as $\overset{M}{<}$ on the children of C, we have no definite way to assign the v_i's as arguments to the function γ_C. But for basis sets that we study, the functions are symmetric in their arguments and so we will not bother to mention the ordering $\overset{M}{<}$. ■

The table in Figure 7.3 collects some common classes of B-choice machines. Each basis set B gives rise to a new computational mode. The symbols for these modes are in the third column of this table. We shall say a B-machine makes B-*choices*. Thus, nondeterministic machines makes nondeterministic choices and alternating machines makes alternating choices. MIN- and MAX-choices are also called *universal choices* and *existential choices*; Coin-tossing choices are also called random choices or probabilistic choices.

[8] We owe the YO-terminology to the unknown street comedian in Washington Square Park whose response to an ongoing public campaign called "Just say NO to drugs" was: "we say YO to drugs". Needless to say, this local answer is in grave error.

[9] This notation could cause some confusion because we do not want to abandon the original meaning of "$C \vdash C'$", that C' is a successor of C. Hence "$C \vdash C'$" does not mean that C has only one successor; to indicate this, we must write "$C \vdash (C')$".

From Figure 7.3, it is evident that we differentiate between the words 'probabilistic' and 'stochastic': the adjective 'probabilistic' applies only to coin-tossing concepts – a usage that conforms to the literature. The adjective 'stochastic' is more general and includes coin-tossing concepts. We abbreviate a probabilistic-alternating machine to 'PAM', and a stochastic-alternating machine to 'SAM'.

If $\gamma(q) = \wedge$ (respectively, $\vee, \overline{\oplus}, \otimes, \oplus, \neg$) then q is called an *MIN-state* (respectively, *MAX-, TOSS-, PrAND-, PrOR-, NOT-state*). If the state of C is an MIN-state (MAX-state, etc.), then C is an *MIN-configuration* (*MAX-configuration*, etc.).

Example 3. This is an instructive exercise to gain som some facility with alternating machines. We show that an alternating machine M which accepts the palindrome language

$$L_{\texttt{pal}} = \{w : w \in \{0,1\}^*, w = w^R\}$$

in linear time, using only logarithmic space. To understand what this tells us about the power of alternation, it was shown in Chapter 2 that

$$L_{\texttt{pal}} \in \textit{D-TIME-SPACE}(O(n), O(n)),$$
$$L_{\texttt{pal}} \in \textit{D-TIME-SPACE}(O(n^2), \log n),$$

and further, if

$$L_{\texttt{pal}} \in \textit{N-TIME-SPACE}(t(n), s(n))$$

then $s(n) \cdot t(n) = \Omega(n^2)$. Hence this example shows that the alternating mode is strictly more powerful than nondeterminism. Of course, the formal definition of what it means for a choice machine to accept a word, and the notion of space and time complexity is yet to come. But if our machine halts on all paths, then these concepts are intuitively obvious: we rely on such intuitions for the reader to understand the example. The idea of the construction is that M accepts an input w iff for all $i = 1, \ldots, n = |w|$, $w[i] = w[n+1-i]$.

We describe the operations of M in two phases. Let the input be w. All the states of M are either deterministic or universal. **Phase 1:** The machine M deterministically marks out $m = \lceil \lg n \rceil$ cells where $n = |w|$. This takes $O(n)$ time, using a well-known fact that we can successively increment a binary counter from 0 to n in $O(n)$ steps of a Turing machine. **Phase 2:** First, it guesses a bit for each of the marked cells. This yields a binary number i between 0 and $2^n - 1$. Now, make a duplicate copy of the number i. Using one copy of i, M determines the bit $w[i+1]$. Using the other copy of i, it determines the bit $w[n-i]$. The reason we need two copies of i is that M "uses i" by successively decrementing i until it reaches 0; thus the original value of i is destroyed in the process. For each decrement, M moves its head position to the next bit of w (the direction of head motion depends on whether we want to look for $w[i+1]$ or $w[n-i]$). As in Phase 1, we use the fact that the successive decrementing of a binary number i to 0 takes $O(i)$ steps of M. Finally, M answers YES if "$w[i+1] = w[n-i]$", and NO otherwise. If $i \geq n-1$, then M can detect this while trying looking for $w[i+1]$, and in this case M always answer YES. This completes our description of M.

First, note that M accepts $L_{\texttt{pal}}$: only any input w, if $w \in L_{\texttt{pal}}$, then M answers YES for value of i. Thus M accepts. If $w \notin L_{\texttt{pal}}$, there is at least one NO answer. Since the states of M are universal, the overall value of the computation tree is a 0, *i.e.*, M rejects. To see that M accepts in time-space $O(n, \log n)$, note that the height of the computation tree is $O(n)$ and each configuration uses space $O(\log n)$. ■

We now want to define acceptance by choice machines. Basically we need to assign intervals $[u, v] \in \mathit{INT}$ to nodes in computation trees. The technical tool we employ is the concept of a 'valuation'.

Definition 3. Let M= (δ, γ) be a choice machine. The set of configurations of δ is denoted $\Delta(M)$. A **valuation** of M is a function

$$V \ : \ \Delta(M) \to \mathit{INT}.$$

A partial ordering on valuations is induced from the $\sqsubseteq$ -ordering on INT as follows: for valuations V_1 and V_2, define $V_1 \sqsubseteq V_2$ if

$$V_1(C) \sqsubseteq V_2(C)$$

for all $C \in \Delta(M)$. The *bottom valuation*, denoted $V_\perp$, is the valuation that always yield $\perp$. Clearly $V_\perp \ \sqsubseteq \ V$ for any valuation V. ■

Definition 4. Let $\Delta \subseteq \Delta(M)$. We define the following operator τ_Δ on valuations. If V is a valuation, then $\tau_\Delta(V)$ is the valuation V' defined by:

$$V'(C) = \begin{cases} \perp & \text{if} \quad C \notin \Delta \ \text{ or } \ C \text{ is YO-configuration,} \\ 1 & \text{else if} \quad C \text{ is a YES-configuration,} \\ 0 & \text{else if} \quad C \text{ is a NO-configuration,} \\ \gamma_C(V(C_1), \ldots, V(C_n)) & \text{else if} \quad C \vdash (C_1, \ldots, C_n). \end{cases}$$

■

For instance, we may choose Δ to be the set of all configurations of M that uses at most space h (for some h).

LEMMA 5 (Monotonicity). *$\Delta_1 \subseteq \Delta_2$ and $V_1 \sqsubseteq V_2$ implies $\tau_{\Delta_1}(V_1) \sqsubseteq \tau_{\Delta_2}(V_2)$.*

Proof. We must show $\tau_{\Delta_1}(V_1)(C) \sqsubseteq \tau_{\Delta_2}(V_2)(C)$ for all $C \in \Delta(M)$. If $C \notin \Delta_1$, then this is true since the left-hand side is equal to $\bot$. So assume $C \in \Delta_1$. If C is terminal, then $\tau_{\Delta_1}(V_1)(C) = \tau_{\Delta_2}(V_2)(C)$ $(= 0, 1$ or $\bot)$. Otherwise, $C \vdash (C_1, \ldots, C_n)$ where n is the arity of γ_C. Then

$$\begin{aligned} \tau_{\Delta_1}(V_1)(C) &= \gamma_C(V_1(C_1), \ldots, V_1(C_n)) \\ &\sqsubseteq \gamma_C(V_2(C_1), \ldots, V_2(C_n)) \\ &= \tau_{\Delta_2}(V_2)(C). \end{aligned}$$

where the $\sqsubseteq$ follows from the monotonicity of γ_C. **Q.E.D.**

For any $\Delta \subseteq \Delta(M)$ and $i \geq 0$, let τ_Δ^i denote operator obtained by the i-fold application of τ_Δ, i.e.,

$$\tau_\Delta^0(V) = V, \quad \tau_\Delta^{i+1}(V) = \tau_\Delta(\tau_\Delta^i(V)).$$

As corollary, we get

$$\tau_\Delta^i(V_\bot) \sqsubseteq \tau_\Delta^{i+1}(V_\bot)$$

for all $i \geq 0$. To see this, use induction on i and the monotonicity lemma.

Definition 5. From the compactness of the interval $[0, 1]$, we see that there exists a unique least upper bound Val_Δ defined by

$$Val_\Delta(C) = \lim\{\tau_\Delta^i(V_\bot)(C) \ : \ i \geq 0\},$$

for all $C \in \Delta$. If $\Delta = \Delta(M)$, then we denote the operator τ_Δ by τ_M, and the valuation Val_Δ by Val_M. ■

A simple consequence of the monotonicity lemma is the following:

$$\Delta_1 \subseteq \Delta_2 \quad \Rightarrow \quad Val_{\Delta_1} \sqsubseteq Val_{\Delta_2}.$$

To see this, it is enough to note that for all $i \geq 0$, $\tau_{\Delta_1}^i(V_\bot) \sqsubseteq \tau_{\Delta_2}^i(V_\bot)$.

For any operator τ and valuation V, we say V is a *fixed point* of τ if $\tau(V) = V$.

LEMMA 6. *Val_Δ is the least fixed point of τ_Δ, i.e.,*

(i) *It is a fixed point: $\tau_\Delta(Val_\Delta) = Val_\Delta$*

(ii) *It is the least such: for all valuations V, if $\tau_\Delta(V) = V$ then $Val_\Delta \sqsubseteq V$.*

Proof.

(i) If C is terminal then it is easy to see that $\tau_\Delta(Val_\Delta)(C) = Val_\Delta(C)$. For non-terminal C, if $C \vdash (C_1, \ldots, C_n)$ then

$$\begin{aligned} \tau_\Delta(Val_\Delta)(C) &= \gamma_C(Val_\Delta(C_1), \ldots, Val_\Delta(C_n)) \\ &= \gamma_C(\lim_i\{\tau_\Delta^i(V_\bot)(C_1)\}, \ldots, \lim_i\{\tau_\Delta^i(V_\bot)(C_n)\}) \\ &= \lim_i\{\gamma_C(\tau_\Delta^i(V_\bot)(C_1), \ldots, \tau_\Delta^i(V_\bot)(C_n))\} \ \text{(by continuity)} \\ &= \lim_i\{\tau_\Delta^{i+1}(V_\bot)(C)\} \\ &= Val_\Delta(C). \end{aligned}$$

(ii) $V_\bot \sqsubseteq V$, so $\tau_\Delta^i(V_\bot) \sqsubseteq \tau_\Delta^i(V) = V$ for all $i \geq 0$. Hence $Val_\Delta \sqsubseteq V$.

Q.E.D.

Example 4. To see that a fixed point of τ_Δ need not be unique, consider a binary computation tree in which all paths, with a single exception, terminate at accepting configurations. The exception is the infinite path π that always branches to the right. We could make sure that each node in this tree has a distinct configuration. Assuming that all nodes are MIN-configurations, a fixed point valuation V_1 of the computation tree is where all nodes have value 1. Another fixed point valuation V_2 assigns each nodes in π to 0 but the rest has value 1. But the least fixed point valuation V_0 assigns to the value $\perp$ to each node on the path π and the value 1 to the rest. ■

Definition 6. An interval $I \subseteq [0,1]$ is a *accepting* if $I \subseteq (\frac{1}{2}, 1]$. It is *rejecting* if $I \subseteq [0, \frac{1}{2})$; it is *undecided* if it is neither accepting nor rejecting. ■

Note that I is accepting/rejecting iff each $v \in I$ is greater/less than $\frac{1}{2}$. Similarly I is undecided iff $\frac{1}{2} \in I$.

Definition 7. (Acceptance rule for choice machines)
(i) Let w be a word in the input alphabet of a choice acceptor M, and Δ a set of configurations of M. The Δ-*value of* w, denoted $Val_\Delta(w)$, refers to $Val_\Delta(C_0(w))$ where $C_0(w)$ is the initial configuration of M on w. If $\Delta = \Delta(M)$, the set of all configurations of M, we write $Val_M(w)$ instead of $Val_{\Delta(M)}(w)$.
(ii) We say M *accepts*, *rejects* or *is undecided on* w according as $Val_M(w)$ is accepting, rejecting or undecided.
(iii) A machine is said to be *decisive* if every input word is either accepted or rejected; otherwise it is *indecisive*
(iv) The *language accepted by* M is denoted $L(M)$. The *language rejected by* M is denote $\overline{L}(M)$. Thus, M is decisive iff $\overline{L}(M) = \text{co-}L(M)$. ■

Convention. In the course of this section, we will introduce other types of fixed point valuations. It is helpful to realize that we use '*Val*' (with various subscripts) only to denote valuations that are least fixed points of the appropriate operators.

EXERCISES

Exercise 0.4: Note that the constant functions $0, 1$ and $\perp$ can be added to any basis B without changing the class of languages accepted by B-machines (under any complexity bounds). What happens if B contains all constant functions $f(I) = c$ for all $c \in [0,1]$? ♢

Exercise 0.5: Our definition of B-machines attaches a basis function $\gamma(q) \subseteq B$ to each state q. A more general definition is to attache a basis function $\gamma(q, a_0, \ldots, a_k)$ to each combination $(q, a_0, \ldots, a_k)$ where a_i are tape symbols valid for tape i. Under what conditions is it possible for our (official) choice machines to simulate the behavior of these generalized choice machines. ♢

Exercise 0.6: Generalize choice machines by allowing values from any $\sqsubseteq$ -partially ordered set F that has a $\sqsubseteq$ -mimimal element $\perp \in F$ and such that any $\sqsubseteq$ -monotonic non-decreasing sequence has a least upper bound in F. Thus, F is meant to be a replacement for *INT*. For instance, let F be the Borel sets (obtain from *INT* under the operation of intersection, difference and countable unions. ♢

END EXERCISES

7.4 Complexity and Tree Valuations

To discuss complexity in general, we need an alternative approach to valuations, called **tree valuations**. The previous notion of valuation is also called **configuration valuations**. In some sense, we are re-developing the previous section all over again. Hence some treatment will be brief.

> Configuration valuations allows us to define the notion of acceptance or rejection. They can also define space complexity: thus, we say that M accepts input w in space h if $Val_\Delta(w)$ is accepting with Δ comprising all those configurations of M that uses space $\leq h$, Unfortunately, configuration valuations are not suited for time complexity. This is because configuration valuations are unable to distinguish between different occurrences of the same configuration C in computation trees. Suppose C occurs at two nodes (say, u_1 and u_2) of a computation tree. Assume the depth of u_1 is less than the depth of u_2. In a time-limited computation, we are interested in computation trees with bounded depths. and want "valuations" V of such trees which may distinguish between u_1 and u_2. For instance, if u_2 is a descendent of u_1 then u_1 typically have "more information" than u_2, and we want $V(u_2) \sqsubseteq V(u_1)$.

The following treatment is abbreviated since it imitates the preceding development.

Definition 8. Fix a choice machine M and any input w.
(i) The *complete computation tree* $T_M(w)$ of M on w is an ordered tree whose nodes are labeled by configurations from Δ_M such that the root is labeled with the initial configuration $C_0(w)$, and whenever a node u is labeled by some C and $C \vdash (C_1, \ldots, C_n)$ then u has n children $u_1, \ldots, u_n$ which are ordered so that u_i is labeled by C_i. We write $u \vdash (u_1, \ldots, u_n)$ in this case. By abuse of terminology, we sometimes identify a node u with its label C.
(ii) A tree T' is a *prefix* of another tree T if T' is obtained from T by pruning[10] some subset of nodes of T. In particular, if T' is non-empty then the root of T' is the root of T. If T is labeled, then T' has the induced labeling.
(iii) A *computation tree* T of w is a prefix of the complete computation tree $T_M(w)$; also, $T_M(w)$ is the *completion* of T.
(iv) A *(tree) valuation* on a computation tree T is a function V that assigns a value $V(u) \in \mathit{INT}$ for each node u in the completion $T_M(w)$ of T, with the property that nodes not in T are assigned $\perp$. We also call V *a tree valuation of* w. If V, V' are valuations of w then we define

$$V \sqsubseteq V'$$

if $V(u) \sqsubseteq V'(u)$ for all $u \in T_M(w)$.
(v) The *bottom valuation*, denoted $V_\perp$, assigns each node of $T_M(w)$ to $\perp$. Clearly $V_\perp$ is the $\sqsubseteq$-minimum tree valuation, for any given w.
(vi) The operator τ_T transforms a valuation V on T to a new valuation $\tau_T(V)$ on T as follows: for each node $u \in T_M(w)$,

$$\tau_T(V)(u) = \begin{cases} \perp & \text{if } u \text{ is a YO-node, or } u \notin T, \\ 1 & \text{else if } u \text{ is a YES-node}, \\ 0 & \text{else if } u \text{ is a NO-node}, \\ \gamma_u(V(u_1), \ldots, V(u_n)) & \text{else if } u \vdash (u_1, \ldots, u_n). \end{cases}$$

Let the least fixed point of τ_T be denoted by Val_T. In fact, $\tau_T^i(V_\perp) \sqsubseteq \tau_T^{i+1}(V_\perp)$ for all $i \geq 0$ and we have

$$\mathit{Val}_T = \lim\{\tau_T^i(V_\perp) \,:\, i \geq 0\}.$$

(vii) A computation tree T is *accepting/rejecting/undecided* if $\mathit{Val}_T(u_0)$ is accepting/rejecting/undecided where u_0 is the root of T. ■

We claim that Val_T is the least fixed point without proof because it is proved exactly as for configuration valuations; furthermore, the next section gives another approach.

Let us say a word w is accepted or rejected by M in the 'new sense' if there is an accepting/rejecting tree for w. We next show the new sense is the same as the old. First we introduce a notation: if $\Delta \subseteq \Delta(M)$ then let

$$T_\Delta(w)$$

denote the largest computation tree T of w all of whose nodes are labeled by elements of Δ. It is not hard to see that this tree is uniquely defined, and is non-empty if and only if the initial configuration $C_0(w)$ is in Δ. Equivalence of the two senses of acceptance amounts to the following.

LEMMA 7. *Fix M and input w.*
(a) If T is an accepting/rejecting computation tree of w then $\mathit{Val}_\Delta(w)$ is also accepting/rejecting, where Δ is the set of labels in T.
(b) Conversely, for any $\Delta \subseteq \Delta(M)$, if $\mathit{Val}_\Delta(w)$ is accepting/rejecting then $T_\Delta(w)$ is also accepting/rejecting.

Its proof is left as an exercise. It follows that a word cannot be both accepted *and* rejected in the tree sense. For, if T is accepting/rejecting tree for w, and Δ are the labels of T, then $\mathit{Val}_\Delta(w)$ is accepting/rejecting. Hence $\mathit{Val}_M(w)$ is accepting/rejecting.

Definition 9. (Acceptance Complexity) Let r be any extended real number.
(i) We say that M *accepts x in time r* if there is an accepting tree T on input x whose nodes are at level at most r (the root is level 0).
(ii) We say M *accepts x in space r* if there is an accepting tree T on input x whose nodes each uses space at most r.
(iii) We say M *accepts x in reversal r* if there is an accepting tree T on input x such that each path in the tree

[10] To *prune* a node u from T means to remove from T the node u and all the descendents of u. Thus, if we prune the root of T, we an empty tree.

makes at most r reversals.
(iv) A computation path $C_1 \vdash C_2 \vdash \cdots \vdash C_m$ makes *(at least) r alternations* if there are $k = \lfloor r \rfloor + 1$ configurations

$$C_{i(1)}, C_{i(2)}, \ldots, C_{i(k)}$$

$1 \leq i(1) < i(2) < \cdots < i(k) \leq m$ such that each $C_{i(j)}$ $(j = 1, \ldots, k)$ is either a MIN- or a MAX-configuration, and if $k(j) < i(j)$ is the largest index such that $C_{k(j)}$ is either a MIN- or MAX-configuration then $C_{k(j)}$ and $C_{i(j)}$ makes different choices (one makes a MIN- and the other a MAX-choice) if and only if there is an even number of NOT-configurations between them along the path. We say *M accepts x in r alternations* if there is an accepting tree T on input x such that no path in the tree makes $1 + r$ alternations.
(v) Let r_1, r_2, r_3 be extended real numbers. We say *M accepts x in simultaneous time-space-reversal* (r_1, r_2, r_3) if there is an accepting tree T that satisfies the requirements associated with each of the bounds r_i $(i = 1, \ldots, 3)$ for the respective resources.
(vi) For complexity functions f_1, f_2, f_3, we say that *M accepts in simultaneous time-space-reversal* (f_1, f_2, f_3) if for each $x \in L(M)$, M accepts x in simultaneous time-space-reversal $(f_1(|x|), f_2(|x|), f_3(|x|))$. This definition extends to other simultaneous bounds. ■

We have just introduced a new resource 'alternation'. Unlike time, space and reversals, this resource is mode-dependent. For example, the machine in the palindrome example above has one alternation and nondeterministic machines has no alternations. We have a similar monotonicity property for tree valuations: if T is a prefix of T' then

$$Val_T \sqsubseteq Val_{T'}.$$

In consequence, we have:

COROLLARY 8. *If M accepts an input in time r then it accepts the same input in time r' for any $r' > r$. Similarly for the other resources.*

Our definition of accepting in time r is phrased so that the accepting tree T need not include all nodes at levels up to $\lfloor r \rfloor$. Because of monotonicity, it may be more convenient to include all nodes up to level $\lfloor r \rfloor$. But when other resource bounds are also being considered, we may no longer be free to do this.

The following result is fundamental:

THEOREM 9 (Compactness). *If a choice machine M accepts a word x then it has a finite accepting tree on input x. Similarly, if M rejects a word, then there is a finite rejecting tree.*

The proof will be deferred to the next section. Thus, if an input is accepted, then it is accepted in finite amounts of time, space, etc. This result implies that the complexity measures such as time, space, reversals or alternation are Blum measures (Chapter 6, Section 8).

On rejection and running complexity. The above definitions of complexity is concerned with accepted inputs only, and no assumptions on the computation of M are made if $w \notin L$(M). In other words, we have been discussing *acceptance complexity*. We now introduce *running complexity* whose general idea is that complexity bounds apply to rejected as well as accepted words. Should running complexity allow indecision on any input? Our definition disallows this.

Definition 10. (Running time complexity) Fix a choice machine M.
(i) We say *M rejects* an input w in k steps if there is a rejecting tree of M on w whose nodes have level at most k.
(ii) For any complexity function $t(n)$, we say *M rejects in time $t(n)$* if for all rejected inputs w, M rejects w in time $t(|w|)$.
(iii) *M runs in time* (t, t') if each input of length n is either accepted in time $t(n)$ or rejected in time $t'(n)$. If $t = t'$, we simply say M runs in time t. ■

This definition extends naturally to other resources. Note that if M has a running time that is finite, *i.e.*, $t(n) < \infty$ for all n, then it is decisive. Thus, we can alternatively say that M is *halting* if it is decisive.

Complexity classes. We are ready to define complexity classes for choice modes. Our previous convention for naming complexity classes extends in a natural way: First note that our notation for complexity classes such as *NTIME*(F) or *D-TIME-REVERSAL*(F, F') has the general format

$$Mode\text{-}Resources\ (\ Bounds\ )$$

where $Mode$ is either N or D, $Resources$ is a sublist of $time, space, reversal$ and $Bounds$ is a list of (families of) complexity functions. The complexity class defined by choice machines can be named using the same format: we only have to add symbols for the new modes and resources. The new mode symbols (see the last column of Figure 7.3) are

$$Pr, \quad A, \quad Ip, \quad PrA, \quad St, \quad StA$$

denoting (respectively) the probabilistic, alternating, interactive proof, probabilistic-alternating, stochastic, stochastic-alternating modes. We have one new resource, with symbols[11]

$$ALTERNATION \text{ or } ALT.$$

Example 5.
(i) Thus $PrTIME(n^{O(1)})$ denotes the class of languages accepted in polynomial time by probabilistic machines. This class is usually denoted $\mathbb{P}$.
(ii) The class $IpTIME(n^{O(1)})$ contains the class usually denoted IP in the literature. If we introduce (see next chapter) the notion of bounded-error decision, indicated by the subscript 'b', then we have

$$IP = IpTIME_b(n^{O(1)}).$$

(iii) If F, F' are families of complexity functions, $PrA\text{-}TIME\text{-}SPACE\,(F, F')$ denotes the class of languages that can be accepted by PAMs in simultaneous time-space (t, s) for some $t \in F, s \in F'$.
(iv) We will write $A\text{-}TIME\text{-}ALT(n^{O(1)}, O(1))$ for the class of languages accepted by alternating machines in polynomial time in some arbitrary but constant number of alternations. This class is denoted PH and contains precisely the languages in the polynomial-time hierarchy (chapter 9). ■

Example 6. Note that $\{\otimes\}$-machines (respectively, $\{\oplus\}$-machines) are equivalent to $\{\wedge\}$-machines ($\{\vee\}$-machines). A more interesting observation is that the probabilistic mode is at least as powerful as nondeterministic mode:

$$N\text{-}TIME\text{-}SPACE\text{-}REVERSAL(t, s, r) \subseteq Pr\text{-}TIME\text{-}SPACE\text{-}REVERSAL(t + 1, s, r)$$

for any complexity functions t, s, r. To see this, let N be any nondeterministic machine that accepts in time-space (t, s). Let M be the following probabilistic machine: on input w, first toss a coin. If tail, answer YES; otherwise simulate N and answer YES iff N answers YES. The jargon "toss a coin and if tail then do X, else do Y" formally means that the machine enters a TOSS-state from which there are two next configurations: in one configuration it does X and in the other choice it does Y. The reader may verify that M accepts $L(N)$ in time-space-reversal $(t + 1, s, r)$. ■

Exercises

Exercise 0.7: (Hong) Consider basis sets B that are subsets of the 16 Boolean functions on 2 variables. As usual, we assume that the identity, 0 and 1 constant functions are not explicitly mentioned when we display B. For two bases B, B', say that B *linearly simulates* B' if for every B'-choice machine M$'$, there is a B-machine M accepting the same language such that if M$'$ accepts in time $t(n)$ then M accepts in time $O_{B,B'}(t)$. Say B and B' are *linearly equivalent* if they can linearly simulate each other.
(a) Prove that every such basis set B is linearly equivalent to one of the following 5 bases:

$$B_0 := \emptyset, B_1 := \{\vee\}, B_2 := \{\wedge\}, B_3 := \{\overline{\oplus}\}, B_4 := \{\vee, \wedge\}$$

where $\overline{\oplus}$ is the exclusive-or function.
(b) By simple set inclusions, it is clear that B_0 can be linearly simulated by the others and B_4 can linearly simulate B_1 and B_2. Show that B_4 can also linearly simulate B_3. **Hint:** Use the same idea as the proof for elimination of negation.
(c)** Show that these 5 classes are distinct up to linear equivalence. (Note: it is known that B_1 is distinct from B_0.) ♢

Exercise 0.8*: A very natural function that we would like to add to basis sets is the cut-off function $\delta_{\frac{1}{2}}$ defined in Section 2. It gives us a model of oracle calls in which the complexity of the oracle machine is taken into account. Since this function is not continuous, we ask you to extend the theory of valuations to allow monotonic, piece-wise continuous functions. (A functions is piece-wise continuous if it has a finite number of discontinuities.) □

[11] Since "alternation" is the name of a mode as well as of a resource, awkward notations such as $A\text{-}ALT(f(n))$ arise.

Exercise 0.9: Suppose the basis set B comprise the rational, linear convex combinations of their arguments: more precisely, the valuation functions f have the form

$$f(x_1, \ldots, x_n) = \sum_{i=1}^{n} a_i x_i$$

where the a_i's are positive rational numbers depending on f and $\sum_i a_i = 1$. How do these machines compare to SAMs? ♢

Exercise 0.10*: Extend valuation theory for acceptors to transducers. NOTE: The question is how to assign a unique output in case different computation paths gives a different answer. Gill has addressed this problem in the case of probabilistic machines. □

End Exercises

7.5 Basic Results

In the last section, the least fixed point tree valuation Val_T for a computation tree T is obtained by repeated application of the operator τ_T to the bottom valuation $Val_\perp$. We now obtain Val_T in an alternative, top-down way.

For each integer $m \geq 0$, let T_m denote the prefix of T obtained by pruning away all nodes at level $m+1$. Thus T_0 consists of just the root of T. By monotonicity,

$$Val_{T_m} \sqsubseteq Val_{T_{m+1}}.$$

Lemma 10.
(i) For any finite computation tree T, the fixed point of τ_T is unique (and, a fortiori*, equal to the least fixed point Val_T). Moreover, this fixed point is easily computed "bottom-up" (say, by a postorder traversal of T).*
(ii) For any computation tree T (finite or not), the valuation

$$V_T^* := \lim\{\, Val_{T_m} \;:\; m \geq 0\}$$

is a fixed point of τ_T.
(iii) V_T^* is equal to the least fixed point, Val_T.
(iv) A computation tree T is accepting/rejecting if and only if it has a finite prefix that is accepting/rejecting.

Proof. (i) This is seen by a bottom-up examination of the fixed point values at each node.
(ii) If u is a leaf then clearly $\tau_T(V_T^*)(u) = V_T^*(u)$. Otherwise, let $u \vdash (u_1, \ldots, u_n)$.

$$\begin{aligned}
\tau_T(V_T^*)(u) &= \gamma_u(V_T^*(u_1), \ldots, V_T^*(u_n)) \\
&= \gamma_u(\lim_{m\geq 0}\{\, Val_{T_m}(u_1)\}, \ldots, \lim_{m\geq 0}\{\, Val_{T_m}(u_n)\}) \\
&= \lim_{m\geq 0}\{\gamma_u(\, Val_{T_m}(u_1), \ldots, \, Val_{T_m}(u_n))\} \\
&= \lim_{m\geq 0}\{\, Val_{T_m}(u)\} \\
&= V_T^*(u)
\end{aligned}$$

This proves that V_T^* is a fixed point of τ_T.
(iii) Since Val_T is the least fixed point, it suffices to show that $V_T^* \sqsubseteq Val_T$. This easily follows from the fact that $V_{T_m} \sqsubseteq Val_T$.
(iv) If a prefix of T is accepting/rejecting then by monotonicity, T is accepting/rejecting. Conversely, suppose T is accepting/rejecting. Then the lower bound of $Val_T(u_0)$ is greater/less than $\frac{1}{2}$, where u_0 is the root of T. By the characterization of Val_T in part (iii), we know that the lower/upper bound of $Val_{T_m}(u_0)$ is monotonically non-decreasing/non-increasing and is greater/less than $\frac{1}{2}$ in the limit as m goes to infinity. Then there must be a first value of m when this lower/upper bound is greater/less than $\frac{1}{2}$. This m gives us the desired finite prefix T_m of T. **Q.E.D.**

This lemma has a two-fold significance: First, part (iv) proves the compactness theorem in the last section. Second, part (iii) shows a constructive way to compute Val_T, by approximating it from below with Val_{T_m}, for increasing m. This method is constructive (in contrast to the $\tau_T^m(V_\perp)$ approximation) because T_m is finite for each m, and the proof of part (i) tells us how to compute Val_{T_m}.

The following lemma is useful for stochastic-alternating computations:

LEMMA 11. *Let T be a computation tree of a choice machine M on w, and $i \geq 0$.*
(i) If M is a probabilistic-alternating machine then $\tau_T^{i+1}(V_\perp)(u) = [x, y]$ implies $2^i x$ and $2^i y$ are integers.
(ii) If M is a stochastic-alternating machine then $2^{2^i} x$ and $2^{2^i} y$ are integers.

We leave the proof as an exercise.

We now have the machinery to show that the language accepted by a B-choice machine M is recursively enumerable provided each function in B is computable. To be precise, let X be a subset of $[0,1]$. We say X is **acceptable** if (i) $0, 1 \in X$, (ii) X is dense in $[0,1]$, and (iii) there is an encoding of X, $e : X \to \{0,1\}^*$ such that for all $u, v \in X$, the predicate $e(u) \leq e(v)$ is decidable.

For example, X can be the set of rational numbers or the set of "binary rationals" which are rationals with finite binary expansion. A suitable encoding $e : X \to \{0,1\}^*$ is easily defined in each case. Now fix an acceptable X. An interval $I = [u, v] \in INT$ is **representable** (in X) if $u, v \in X$. The encoding of I is just a pair $e(I) = [e(u), e(v)] \in (\{0,1\}^*)^2$. We say an m-ary function $f \in B$ is **X-computable** if the following holds:
(1) If $I_1, \ldots, I_m \in INT$ are representable, then $f(I_1, \ldots, I_m)$ is representable.
(2) There is a partial computable function $F : (\{0,1\}^*)^{2m} \to (\{0,1\}^*)^2$ such that for all representable $I_1, \ldots, I_m$, we have $F(e(I_1), \ldots, e(I_m)) = e(f(I_1, \ldots, I_m))$. Finally, call B a **computable basis** if there is some X such that each $f \in B$ is X-computable.

THEOREM 12. *Let B be a computable basis set.*
a) The class of languages accepted by B-machines is precisely the class RE of recursively enumerable languages.
b) The class of languages accepted by decisive B-machines is precisely the class REC.

Proof. a) Languages in RE are accepted by B-choice machines since B-choice machines are generalizations of ordinary Turing machines. Conversely, let M be a B-choice machine and w an input word. To show that $L(\mathrm{M})$ is in RE, it is sufficient to give a deterministic procedure for checking if $w \in L(\mathrm{M})$, where the procedure is required to halt only if w is in $L(\mathrm{M})$. The procedure computes, for successive values of $m \geq 0$, the value $Val_{T_m}(u_0)$ where T_m is the truncation of $T_M(w)$ below level m and u_0 the root. If T_m is accepting for any m, the procedure answers YES. If T_m is non-accepting for all m, the procedure loops. Lemma 10 not only justifies this procedure, but it also shows how to carry it out: the values $Val_{T_m}(u)$ of each node $u \in T_m$ is computed in a bottom-up fashion. The computability of the basis functions B ensures this is possible.

b) We leave this as an exercise. **Q.E.D.**

One can view the theorem as yet another confirmation of Church's thesis. Our next result shows that negation $\neg$ can be avoided in stochastic-alternating machines at the cost of an increase in the number of states. The following generalizes a result for alternation machines in [3].

THEOREM 13. *For any stochastic-alternating acceptor* M*, there is a stochastic-alternating acceptor* N *such that* N *has no NOT-states,* $L(\mathrm{M}) = L(\mathrm{N})$*, and for all w and $t, s, r, a \geq 0$:* M *accepts w in (time, space, reversal, alternation) (t, s, r, a) iff* N *accepts w in (time, space, reversal, alternation) (t, s, r, a).*

Proof. The idea is to use de Morgan's law to move negation to the leaves of the computation tree. Let M be any SAM. We construct a SAM N satisfying the requirements of the theorem. For each state q of M, there are two states q^+ and q^- for N. For any configuration C of M, let C^+ (resp., C^-) denote the corresponding configuration of N where q^+ (resp., q^-) is substituted for q. In N, we regard q_0^+, q_Y^+ and q_N^+ (respectively) as the initial, YES and NO states. However, we identify q_Y^-, q_N^- (respectively) with the NO, YES states (note the role reversal). Of course, the technical restriction does not permit two YES- or two NO-states, so we will identify them ($q_Y^+ = q_N^-, q_N^+ = q_Y^-$). The functions $\gamma(q^+), \gamma(q^-)$ assigned to the states in N are defined from $\gamma(q)$ in M as follows:

$$\gamma(q^+) = \begin{cases} \gamma(q) & \text{if } \gamma(q) \neq \neg \\ \iota & \text{if } \gamma(q) = \neg \end{cases}$$

$$\gamma(q^-) = \begin{cases} \iota & \text{if } \gamma(q) = \neg \\ \wedge & \text{if } \gamma(q) = \vee \\ \vee & \text{if } \gamma(q) = \wedge \\ \oplus & \text{if } \gamma(q) = \oplus \\ \oplus & \text{if } \gamma(q) = \otimes \\ \otimes & \text{if } \gamma(q) = \oplus \end{cases}$$

Hence N has no NOT-states. We now state the requirements on transitions of N (this easily translates into an explicit description of the transition table of N). Suppose that $C_1 \vdash C_2$. If C_1 is not a NOT-configuration then

$$C_1^+ \vdash C_2^+ \text{ and } C_1^- \vdash C_2^-.$$

If C_1 is a NOT-configuration then

$$C_1^+ \vdash C_2^- \text{ and } C_1^- \vdash C_2^+.$$

Our description of N is complete: there are no transitions besides those listed above.

Let T be an accepting computation tree of N for an input word w; it is easy to see that there is a corresponding computation tree $\widehat{T}$ for N with exactly the same time, space, reversal and alternation complexity. In fact there is a bijection between the nodes of T and $\widehat{T}$ such a node labeled C in T corresponds to one labeled C^+ or C^- in $\widehat{T}$. The fact that T and $\widehat{T}$ have identical alternating complexity comes from the carefully-crafted definition of N.

Our theorem is proved if we show that $\widehat{T}$ is accepting. Let T_m be the truncation of the tree T at levels below $m \geq 0$; $\widehat{T}_m$ is similarly defined with respect to $\widehat{T}$. For $h \geq 0$, let V_m^h denote the valuations on T_m given by h-fold applications of the operator τ_{T_m} to $\perp$:

$$V_m^h = \tau_{T_m}^h(V_\perp)$$

and similarly define $\widehat{V}_m^h = \tau_{\widehat{T}_m}^h(V_\perp)$. We now claim that for all m, h and $C \in T_m$,

$$V_m^h(C) = \begin{cases} \widehat{V}_m^h(C^+) & \text{if } C^+ \in \widehat{T} \\ \neg\widehat{V}_m^h(C^-) & \text{if } C^- \in \widehat{T} \end{cases}$$

Here, we have abused notation by identifying the configuration C with the node of T_m that it labels. But this should be harmless except for making the proof more transparent. If $h = 0$ then our claim is true

$$\perp = V_m^h(C) = \widehat{V}_m^h(C^+) = \neg\widehat{V}_m^h(C^-)$$

since $\neg\perp = \perp$. So assume $h > 0$. If C is a leaf of T, it is also easy to verify our claim. Hence assume C is not a leaf. Suppose $C^- \vdash (C_1^-, C_2^-)$ occurs in $\widehat{T}$. Then

$$\begin{aligned} V_m^h(C) &= \gamma(C)(V_m^{h-1}(C_1), V_m^{h-1}(C_2)) \\ &= \gamma(C)(\neg\widehat{V}_m^{h-1}(C_1^-), \neg\widehat{V}_m^{h-1}(C_2^-)) \quad \text{(by induction)} \\ &= \neg\gamma(C^-)(\widehat{V}_m^{h-1}(C_1^-), \widehat{V}_m^{h-1}(C_2^-)) \quad \text{(de Morgan's law for } \gamma(C)) \\ &= \neg\widehat{V}_m^h(C^-). \end{aligned}$$

Similarly, we can show $V_m^h(C) = \widehat{V}_m^h(C^+)$ if $C^+ \vdash (C_1^+, C_2^+)$ occurs in $\widehat{T}$. We omit the demonstration in case C is a NOT-configuration. Finally, noting that $V_m^{m+1} = Val_{T_m}$ and $\widehat{V}_m^{m+1} = Val_{\widehat{T}_m}$, we conclude that $Val_{T_m} = Val_{\widehat{T}_m}$. It follows $\widehat{T}$ is accepting. **Q.E.D.**

Consequence of eliminating negation. This proof also shows that negation can be eliminated in alternating machines and in PAMs. With respect to SAMs without negation, the use of intervals in valuations can be replaced by ordinary numbers in $[0, 1]$ *provided we restrict attention to acceptance complexity.* A valuation is now a mapping from $\Delta \subseteq \Delta(\mathrm{M})$ into $[0, 1]$. Likewise, a tree valuation assigns a real value in $[0, 1]$ to each node in a complete computation tree. We now let $V_\perp$ denote the valuation that assigns the value 0 to each configuration or node, as the case may be. The operator τ_Δ or τ_T on valuations is defined as before. Their least fixed point is denoted Val_Δ or Val_T as before. The connection between the old valuation V and the new valuation V' is simply that $V'(C)$ or $V'(u)$ (for any configuration C or node u) is equal to the lower bound of the interval $V(C)$ or $V(u)$. When we discuss running complexity, we need to consider the upper bounds of intervals in order to reject an input; so we are back to intervals.

Convention for this chapter. In this chapter, we only consider alternating machines, PAMs and SAMs with no NOT-states. We are mainly interested in *acceptance complexity.* In this case, we may restrict valuations take values in $[0, 1]$ instead of in *INT* (we call these real values *probabilities*). With this convention, the acceptance rule becomes:

> M accepts a word w iff the probability $Val_M(w)$ is greater than $\frac{1}{2}$.

■

It is sometimes convenient to construct SAMs *with* NOT-states, knowing that they can be removed by an application of the preceding theorem.

Suppose we generalize SAMs by allowing the k-ary versions of the alternating-stochastic functions:

$$B_k := \{\max_k, \min_k, \underset{k}{\textcircled{+}}, \otimes_k, \oplus_k\},$$

for each $k \geq 2$. For example,

$$\begin{aligned} \underset{3}{\textcircled{+}}(x,y,z) &= (x+y+z)/3, \\ \oplus_3(x,y,z) &= 1-(1-x)(1-y)(1-z). \end{aligned}$$

Consider generalized SAMs whose basis set is $\cup_{k\geq 2} B_k$. Note that even though the basis set is infinite, each generalized SAM uses only a finite subset of these functions. It is easily seen that with this generalization for the alternating choices ($\max_k$, $\min_k$ and $\otimes_k$), the time complexity is reduced by at most a constant factor. It is a little harder (Exercise) to show the same for the stochastic choices ($\textcircled{+}_k, \oplus_k, \otimes_k$).

A useful technical result is tape reduction for alternating machines. The following is from Paul, Praus and Reischuk [18].

THEOREM 14. *For any k-tape alternating machine accepting in time-alternation $(t(n), a(n))$, there is a simple alternating machine accepting the same language and time-alternation $(O(t(n)), a(n)+O(1))$. Like a simple Turing machine, a simple alternating machine has only one work-tape and no input tape.*

This leads, in the usual fashion, to a hierarchy theorem for alternating time:

THEOREM 15. *Let $t(n)$ be constructible and $t'(n) = o(t(n))$. Then*

$$ATIME(t) - ATIME(t') \neq \emptyset.$$

We leave both proofs as exercises.

THEOREM 16 (Space compression). *Let B be any basis set. Then the B-choice machines have the space compression property. More precisely, if M is a B-choice machine accepting in space $s(n)$ then there is another N which accepts the same language in space $s(n)/2$. Furthermore, N has only one work-tape.*

Proof. We only sketch the proof, emphasizing those aspects that are not present in the proof of original space compression theorem in chapter 2. As before, we compress 2 symbols from each of the k work-tapes of M into one *composite symbol* (with k tracks) of N. We show how N simulates one step of M: suppose M is in some configuration C and $C \vdash (C_1, \ldots, C_m)$. Assume that the tape head of N is positioned at the leftmost non-blank cell of its tape. By deterministically making a rightward sweep across the non-blank part of its work-tape, N can remember in its finite state control the two composite symbols adjacent to the currently scanned cell *in each track*: the cells of M corresponding to these remembered symbols constitute the *current neighborhood*. In the original proof, N makes a leftward sweep back to its starting position, updating the contents of the current neighborhood. The new twist is that there are now m ways to do the updating. N can use choice to ensure that each of these possibilities are covered. More precisely, before making the return sweep, N enters a state q such that $\gamma(q) = \gamma(C)$ and then N branches into m different states, each corresponding to a distinct way to update the current neighborhood. Then N can make a deterministic return sweep on each of the m branches. By making some adjustments in the finite state of N, we may ensure that N uses space $s(n)/2$. Further, N is also a B-machine by construction. **Q.E.D.**

For any family of functions B over INT, let B^* denote the *closure* of B (under function composition): B^* is the smallest class of functions containing B and closed under function composition.[12] For example, the function $h(x,y,z) = x \textcircled{+} (y \textcircled{+} z) = x/2 + y/4 + z/4$ is in the closure of the basis $\{\textcircled{+}\}$. The closure of a basis set is also a basis set (in fact, an infinite set).

THEOREM 17 (Linear Speedup). *Let B be an admissible family which is closed under function composition, $B = B^*$. Then the B-choice machines have the linear speedup property. More precisely, if M is a B-choice machine accepting in time $t(n) > n$ then there is another N which accepts the same language in time $n + t(n)/2$.*

[12] More precisely, if $f \in B^*$ is a function of arity k and g_i $(i = 1, \ldots, k)$ are functions of arity m_i in B^* then $f(g_1(\bar{x}_1), \ldots, g_k(\bar{x}_2))$ is a function in B^* of arity $p \leq \sum_{i=1}^k m_i$ where $\bar{x}_i$ is a sequence of m_i variables and p is the number of distinct variables in $\bar{x}_1\bar{x}_2\cdots\bar{x}_k$.

Proof. The proof is similar to that for ordinary Turing machines in chapter 2, so we only emphasize the new aspects. The new machine N has $k+1$ work-tapes if M has k. Each tape cell of N encodes up to $d > 1$ (for some d to be determined) of the original symbols. N spends the first n steps making a compressed copy of the input. Thereafter, N uses 8 steps to simulates d steps of M. In general, suppose that M is in some configuration C and d moves after C, there are m successor configurations $C_1, \ldots, C_m$ (clearly m is bounded by a function of M and d only). Suppose that the complete computation tree in question is T. The value $Val_T(C)$ is given by

$$f(\, Val_T(C_1), \ldots, \, Val_T(C_m))$$

where $f \in B$ since B is closed under function composition. We show how N can determine this f: first N takes 4 steps to determine the contents of the 'current neighborhood' (defined as in the original proof). From its finite state constrol, N now knows f and each C_i $(i = 1, \ldots, m)$. So at the end of the fourth step, N could enters a state q where $\gamma(q) = f$ and such that q has m successors, corresponding to the C_i's. In 4 more steps, N deterministically updates its current neighborhood according to each C_i. It is clear that by choosing $d = 16$, N accepts in time $n + t(n)/2$. One minor difference from the original proof: previously the updated tapes represent the configuration at the first time some tape head leaves the current neighborhood, representing at least d steps of M. Now we simply simulate *exactly* d steps and so it is possible that a tape head remain in the current neighborhood after updating. **Q.E.D.**

As corollary, if we generalize alternating machines by replacing the usual basis set $B = \{\wedge, \vee, \neg\}$ by its closure B^*, then the generalized alternating time classes enjoy the time speedup property. A similar remark holds for the other modes.

Exercises

Exercise 0.11: Prove Theorem 14. ♢

Exercise 0.12: Prove Theorem 15. ♢

Exercise 0.13: Show that if $t(n)$ is time constructible then co-$ATIME(t(n)) \subseteq ATIME(n + t(n))$. **Hint:** why do you need the "$n+$" term? ♢

Exercise 0.14: Let $B = \{\vee, \oplus, \neg\}$. Can negation be eliminated from B-machines operating in polynomial time? ♢

Exercise 0.15: (i) Consider generalized probabilistic machines in which we allow k-ary versions of the coin-tossing function, $\oplus_k$ for all $k \geq 2$. Show that these can be simulated by ordinary probabilistic machines with at most a constant factor slow-down in time.
(ii) Show the same result for stochastic machines where we now allow $\oplus_k, \oplus_k, \otimes_k$ for all k. ♢

Exercise 0.16: (Paul, Praus, Reischuk) A **simple alternating machine** is an alternating machine with one work-tape but no input tape. It is thus physically identical to simple Turing machines (STM's). As for STM, the input is presented on the sole work-tape.
(a) Construct a simple alternating machine that accepts palindromes in linear time and a constant number of alternation. **Hint:** Guess the positions (in binary) of about $\log n$ equally spaced-out input symbols, writing these directly under the corresponding symbols. The positions of all the other symbols can be determined relative to these guessed positions.
(b) Show that if a language can be accepted by an alternating Turing machine in time $t(n)$ then it can be accepted by a simple alternating Turing machine in time $O(t(n))$. **Hint:** For a computation path $C_1 \vdash C_2 \vdash, \ldots, \vdash C_m$, let its *trace* be $\tau_1, \tau_2, \ldots, \tau_m$ where τ_i contains the state, the symbols scanned on each of the tapes (including the input tape) and the direction of each tape head in the transition $C_i \vdash C_{i+1}$. The head directions are undefined for τ_m. (Note: this definition of trace does not include head positions, in contrast to a similar definition in chapter 2.) Begin by guessing the trace of the paths in an accepting computation tree – you must use universal and existential guessing corresponding to the state in τ_i. To verify correctness of the guess, the technique in part (a) is useful. ♢

Exercise 0.17: (Paterson, Paul, Praus, Reischuk) For all $t(n)$, if a language is accepted by a nondeterministic simple Turing machine in time $t(n)$ then it can be accepted by an alternating machine in time $n + t^{1/2}(n)$. ♢

End Exercises

7.6 Alternating Time versus Deterministic Space

We begin the study of alternating machines. This section points out strong similarities between alternating time and deterministic space. This motivates a variation of choice machines called the addressable-input model.

The class PH. The alternating class $ATIME(n^{O(1)})$ is an obvious class of interest. However, most interest attaches to a subclass $PH \subseteq ATIME(n^{O(1)}$ comprising the languages accepted by alternating machines in polynomial time and making only a constant number of alternations. Clearly $NP \cup \text{co-}NP \subseteq PH$. Here is an interesting problem in PH. The MIN-FORMULA problem is to recognize if a given Boolean formula F is minimal. That is, for all F', if $|F'| < |F|$ then $F' \not\equiv F$. It is not obvious that this problem is in $NP \cup \text{co-}NP$. To show that MIN-FORMULA$\in PH$: on input $F(x_1, \ldots, x_n)$, we universally guess another formula $F'(x_1, \ldots, x_n)$ such that $|F'| < |F|$ and then existentially guess an assignment $a_i \mapsto x_i$ and accepts if $F'(a_1, \ldots, a_n) \neq F(a_1, \ldots, a_n)$.

We first prove the following result of Chandra, Kozen and Stockmeyer:

THEOREM 18. *For all t, $ATIME(t) \subseteq DSPACE(t)$.*

Proof. Let M be an alternating machine accepting in time t. We describe a deterministic N that simulates M in space t. Let w be the input and N computes in successive stages. In the mth stage ($m = 1, 2, 3, \ldots$), N computes Val_{T_m} where T_m is the truncation of the complete computation tree $T_M(w)$ at levels below m. For brevity, write Val_m for Val_{T_m}. If T_m is accepting then N accepts, otherwise it proceeds to the next stage. So if w is not in $L(M)$ then N will not halt.

To complete the proof, we show that the mth stage can be carried out using $O(m)$ space. We describe a procedure to compute the values $Val_m(C)$ of the nodes C in T_m in a post-order manner. The structure of this search is standard. We inductively assume that the first four work tapes of N contain the following pieces of information when we first visit a configuration C:

(a) Tape 1 contains the configuration C. This requires $O(m)$ space. In particular, the input head position can be recorded in $O(\log m)$ space (rather than $O(\log |w|)$ space).

(b) Tape 2 contains the values m and the depth d of C in T_m. These can be represented by the string $0^{m-d}1^d$.

(c) Tape 3 contains a list $(I_1, I_2, \ldots, I_d)$ of d instructions of M. If $\pi(C) = (C_0, C_1, \ldots, C_d)$ is the computation sequence from the root of T_m to $C = C_d$, then the transition $C_i \vdash C_{i+1}$ is a result of executing the instruction I_i. The space for storing these instructions is $O(d) = O(m)$. Using instructions, we can "backup" from C to any of its ancestors C' on the path.

(d) Tape 4 contains a list $(v_0, v_1, \ldots, v_{d-1})$ where $v_i \in \{-1, 0, 1\}$. Without loss of generality, assume that each non-terminal configuration of M has exactly two children. One child of C_i is C_{i+1}; let the other child be C_i'. Then v_i is the value $Val_m(C_i')$, it is has been computed. Otherwise v_i is -1. The space used is again $O(d) = O(m)$. Note that the value $Val_m(C)$ is not known in this representation.

We maintain this information at each "step". A step (at current configuration C) either involves descending to a child of C or backing up from C to its parent. We descending to a child of C provided C is not a leaf and it has a child that is not yet been visited. Maintaining the contents of tapes 1-4 is easy in this case. So suppose we want to backup from C to its parent C'. We claim that $Val_m(C)$ can be determined at this moment. This is true if C is a leaf of T_m. Otherwise, the reason we are backing up to C' is because we had visited all the children C_1, C_2 of C. But this meant we had just backed up from (say) C_2 to C, and inductively by our claim, we have determined $Val_m(C_2)$. From (c), we also know $Val_m(C_1)$. Thus we know $Val_m(C)$. Then when we back up to the parent of C, we could either replace C by the sibling of C or record $Val_m(C)$ on tape 4, and continue the induction. Eventually we determine the value of Val_m at the root. If this value is 1, we accept. Otherwise, we go to state $m + 1$. **Q.E.D.**

Discussion. The preceding theorem motivates several research questions in alternation complexity.

(A) The theorem says that *deterministic space is at least as powerful as alternating time.* It suggests a search for the followin kinds of results: take a known simulation by deterministic space and ask if it can be improved to an alternating time simulation. This methodology has proven fruitful and has resulted in a deeper understanding of the space resource. Thus, in section 8, a known inclusion $DTIME(t) \subseteq DSPACE(t/\log t)$ was sharped to $DTIME(t) \subseteq ATIME(t/\log t)$. This strengthening lead to a simplification (!) of the original proof. This paradox is explained by the fact that the control mechanism in alternating computation is "in-built"; an alternating simulation (unlike the original space simulation) need not explicitly describe this mechanism.

(B) There is evidence to suggest that alternating time and deterministic space are very similar. For instance, we prove (§7.7) a generalization of Savitch's result, which yields the corollary

$$NSPACE(s) \subseteq ATIME(s^2).$$

This motivates another class of new results: given a known result about deterministic space, try to prove the analogue for alternating time, or vice-versa. For instance, the last section shows a tape-reduction and a hierarchy theorem for alternating-time; the motivation for these results is that we have similar results for deterministic space. We now give another illustration. In chapter 2, we show that $DSPACE_r(s)$ is closed under complementation for all s finite (*i.e.*, $s(x) < \infty$ whenever defined). We ask for a corresponding result for $ATIME_r(t)$. (Note that the subscript 'r' indicates running complexity.) As it turns out, this result[13] is rather easy for alternating time:

THEOREM 19. *For all complexity function $t(n)$,*

$$ATIME_r(t) = \text{co-}ATIME_r(t).$$

Similarly, the following time classes are closed under complementation

$$PrTIME_r(t), StTIME_r(t), PrA\text{-}TIME_r(t), StA\text{-}TIME_r(t).$$

Proof. Recall the construction in theorem 13 of a machine N without negation from another machine M that may have negation. Now let M be an alternating machine. Suppose that we make q_0^- (instead of q_0^+) the start state of N but $q_Y^+ = q_N^-$ remains the YES state. On re-examination of the proof of theorem 13, we see that this N accepts co-$L(M)$. The proof for the other time classes are similar. **Q.E.D.**

(C) By now it should be realized that the fundamental technique for space simulation is to 'reuse space'. This usually amounts to cycling through an exponential number of possibile configurations using the same space s. In alternating time, the corresponding technique is to visit exponentially many configurations using a tree of depth s, by making universal or existential choices. While a deterministic space search proceeds from what is known to what is unknown, alternating time search works in reverse: it first guesses the unknown and then try to to reduce it to the known. This remark may be clearer by the end of this chapter.

(D) We should caution that the research programs (A) and (B) have limitations: although the deterministic space and alternating time are similar, it is unlikely that they are identical. Another fundamental difficulty is that whereas sublinear deterministic space classes are important, it is easy to see that alternating machines do not allow meaningful sublinear time computations. This prompted Chandra, Kozen and Stockmeyer to suggest[14] a variation of alternating machines which we now extend to choice machines:

Definition 11. *(Addressable-input Machine Model)* An *addressable-input* choice machine M is one that is equipped with an extra *address* tape and two distinguished states called the *READ* and *ERROR* states. The address tape has a binary alphabet whose content is interpreted as an integer. Whenever M enters the READ state, the input head is instantaneously placed at the (absolute) position indicated by the address tape. If this position lies outside the input word, the ERROR state will be entered and no input head movement occurs. In addition to this special way of moving the input head, our machine machine can still move and use the input head in the standard fashion.

We assume the address tape is one-way (and hence is write-only). Hence for complexity purposes, space on the address tape is *not* counted. We also assume that after exiting from the READ state, the contents of the address tape is erased, in an instant. ■

The addressable-input model is defined so that such machines are at least as powerful as *ordinary* choice machines. However, it is not excessively more powerful; for instance the preceding theorem theorem 18 holds even with this model of alternation (Exercise). This addressable-input model now admits interesting alternating time classes with complexity as small as $\log n$ (not $\log \log n$, unfortunately). We will be explicit whenever we use this version of choice machines instead of the ordinary ones.

EXERCISES

Exercise 0.18: (i) Consider the language comprising (D, k) where D is the distance matrix as in the travelling salesman problem (TSP) in Chapter 3, and k is the length of the shortest tour. Show that this language is in PH.

(ii) Construct a "natural" language that apparently needs 4 alternations. ♢

[13] Paul and Reischuk show that if t is time-constructible then $ATIME(t) = \text{co-}ATIME(t)$.

[14] In this suggestion, they are in good company: historically, the read-only input tape of Turing machines was invented for a similar purpose.

Exercise 0.19: Give a sufficient condition on a family F of complexity functions such that $ATIME(F) = PrA\text{-}TIME(F) = DSPACE(F)$. **Hint:** Consider the case $F = n^{O(1)}$. ♢

Exercise 0.20: Show the tight complexity relationships between the ordinary SAM's and the addressable-input version of SAM's. More precisely, give efficient time/space/reversal simulations of the addressable-input machines by ordinary machines. ♢

Exercise 0.21: Rederive the various simulation of SAM's in this chapter in the case where the SAMs is the addressable-input model. In particular, show that $ATIME(t) \subseteq DSPACE(t)$. ♢

End Exercises

7.7 Simulations by Alternating Time

We present efficient simulations of other complexity resources by alternating time. We begin with an alternating time simulation of deterministic space and reversals. The results here are new, motivated by I. Simon's simulation in Chapter 2.

Theorem 20. *Let t, r be complexity functions such that $t(n) \geq 1 + n$. Under the addressable-input model,*

$$D\text{-}TIME\text{-}REVERSAL(t, r) \subseteq ATIME(O(r \log^2 t)).$$

If $r(n) \log^2 t(n) \geq n$, then this result holds under the ordinary model.

Proof. Given a deterministic M accepting in time-reversal (t, r), we show an alternating machine N accepting the same language $L(M)$ in time $O(r \log^2 t)$.

Recall the concept of a (full) trace[15] in the proof of chapter 2. On any input w, N existentially chooses some integer $r \geq 1$ and writes r full traces

$$\tau_1, \tau_2, \ldots, \tau_r,$$

on tape 1. These are intended to be the traces at the beginning of each of the r phases. On tape 2, N existentially chooses the time t_i (in binary) of each τ_i,

$$t_1 < t_2 < \cdots < t_r.$$

We may assume τ_r is the trace when the machine accepts. Note that τ_1 (which we may assume correct) is simply the trace of the initial configuration and so $t_1 = 0$. Relative to this sequence, we say that an integer $t \geq 0$ *belongs to phase j* if $t_j \leq t < t_{j+1}$.

Then N proceeds to verify τ_r. To do this, it writes on tape 3 the pairs (τ_{r-1}, t_{r-1}) and (τ_r, t_r) and invokes a procedure *TRACE*. N accepts if and only if this invocation is *successful* (i.e., the procedure accepts). In general, the arguments to *TRACE* are placed on tape 3, and they have the form

$$(\sigma, s_0), (\tau, t_0)$$

where $s_0 < t_0$ are binary integers lying in the range

$$t_i \leq s_0 < t_0 \leq t_{i+1}$$

for some $i = 1, \ldots, r-1$, and σ, τ are traces such that the head tendencies in σ and in τ_i agree, and similarly the head tendencies in τ and in τ_i agree (with the possible exception of $t_0 = t_{i+1}$ and $\tau = \tau_{i+1}$). Intuitively, $TRACE(\sigma, s_0, \tau, t_0)$ accepts if σ is the trace at time s_0, τ is the trace at time t_0, and there is a (trace of a) path from σ to τ. *TRACE* does one of two things:

(i) Suppose $s_0 + 1 < t_0$. Let $t' = \lfloor (s_0 + t_0)/2 \rfloor$. Now *TRACE* existentially chooses a trace τ' where the head tendencies in τ' agree with those of σ. Then it universally chooses to recursively call $TRACE(\sigma, s_0, \tau', t')$ and $TRACE(\tau', t', \tau, t_0)$.

[15] Briefly, the trace of a configuration in a computation path records its state and for each tape, the scanned symbol, the absolute head position and the head tendencies.

(ii) Suppose $s_0 + 1 = t_0$. *TRACE* verifies τ can be derived from σ in one step of M, and any head motion is consistent with the head tendencies in σ. Of course, we allow the head tendencies to be different but only when σ is the last trace in a phase (it is easy to determine if this is the case). Any head motion in the σ to τ transition causes the corresponding tape cell in τ to be 'marked'. Note that the marked cells were written in some previous phase (unless they are blanks), and our goal is to verify their contents. Suppose that the tape symbols and head positions in the $k+1$ tapes of τ are given by

$$b_0, \ldots, b_k, \qquad n_0, \ldots, n_k.$$

Then *TRACE* universally chooses to call another procedure *SYMBOL* with arguments (i, b_i, n_i, t_0) for each cell n_i in tape i that is marked. Intuitively, $SYMBOL(i, b_i, n_i, t_0)$ verifies that just before time t_0, the tape symbol in cell n_i of tape i is b_i.

We now implement $SYMBOL(i', b', n', t_0)$. If $i' = 0$ then we want to check that the input symbol at position n' is b'. This can be done in $O(\log n)$ steps, using the input addressing ability of our alternating machines (note that $r \log^2 t > \log n$). Otherwise, suppose that t_0 belongs to phase j_0. We then existentially choose some j,

$$j = 0, \ldots, j_0 - 1,$$

some t' and a trace σ'. Intuitively, this means that cell n' in tape i' was last visited by σ' which occurs at time t' in phase j. We want the following to hold:
(a) $t_j \le t' < t_{j+1} \le t_0$.
(b) The head tendencies σ' and in τ_j agree.
(c) The head i' is in position n' scanning symbol b' in σ'.
(d) On each tape, the head position in σ' lies in the range of possible cell positions for that phase j.
(e) On tape i', cell n' is not visited in any of phases $j+1, j+2, \ldots, j_0$.
Conditions (a)-(e) can be directly verified using the information on tapes 1 and 2. Armed with σ' and t', we then universally choose one of two actions: either call $TRACE(\tau_j, t_j, \sigma', t')$ or $TRACE(\sigma', t', \tau_{j+1}, t_{j+1})$. If $t_j = t'$ then the first call is omitted.

Correctness. Let us show that *TRACE* and *SYMBOL* are correct. Suppose input w is accepted by M. Then it is not hard to see that N accepts. To show the converse, suppose N accepts w relative to some choice of traces $\tau_1, \ldots, \tau_r$ in tape 1 and times $t_1, \ldots, t_r$ on tape 2. Suppose the call $TRACE(\sigma, s_0, \tau, t_0)$ is successful and t_0 belongs to phase j. Then this call generates a trace-path

$$(\sigma_0, \ldots, \sigma_m) \tag{4}$$

from $\sigma_0 = \sigma$ to $\sigma_m = \tau$ (where $m = t_0 - s_0$) with the following properties:

1. σ_{i-1} derives σ_i for $i = 1, \ldots, m$ according to the rules of M.

2. Each pair (σ_{i-1}, σ_i) in turn generates at most $k+1$ calls to *SYMBOL*, one call for each "marked" cell in σ_i. Each of these calls to *SYMBOL* leads to acceptance. For this reason we call (4) a 'successful' trace-path for this call to *TRACE*.

3. Some of these calls to *SYMBOL* in turn calls *TRACE* with arguments belonging some phase ℓ $(1 \le \ell < j)$. We call any such phase ℓ a *supporting phase* of $TRACE(\sigma, s_0, \tau, t_0)$.

Notice that if phase ℓ is a supporting phase for some accepting call to *TRACE* then there must be two successful calls of the form

$$TRACE(\tau_\ell, t_\ell, \sigma', t') \quad \text{and} \quad TRACE(\sigma', t', \tau_{\ell+1}, t_{\ell+1}) \tag{5}$$

for some σ', t'. We claim:

(a) If $TRACE(\sigma, s_0, \tau, t_0)$ accepts and phase ℓ is a supporting phase of $TRACE(\sigma, s_0, \tau, t_0)$, then the traces $\tau_1, \ldots, \tau_{\ell+1}$ on tape 1 and times $t_1, \ldots, t_{\ell+1}$ on tape 2 are correct (i.e., τ_i is the trace at the beginning of the ith phase at time t_i for $i = 1, \ldots, \ell+1$.)
(b) If, in addition, we have that $\sigma = \tau_j$ and $s_0 = t_j$ for some $j = 1, \ldots, r$, then τ is indeed the trace of the t_0th configuration in the computation path of M on input w. (Note that (b) implies the correctness of *SYMBOL*.)

We use induction on ℓ. Case $\ell = 1$: τ_1 is always correct and $t_1 = 0$. By (5), we see directly that there must be two successful calls of the form $TRACE(\tau_1, t_1, \sigma', t')$ and $TRACE(\sigma', t', \tau_2, t_2)$. One immediately checks that this implies τ_2, t_2 are correct. This proves (a). Part (b) is immediate.

Case $\ell > 1$: for part (a), again we know that there are two successful calls of the form (5). But notice that phase $\ell - 1$ is a supporting phase for the first of these two calls: this is because in τ_ℓ, some tape head made a reversal that this means that this head scans some symbol last visited in phase $\ell - 1$. Hence by induction hypothesis, the traces $\tau_1, \ldots, \tau_\ell$ and times $t_1, \ldots, t_\ell$ are correct. Furthermore, as in (4), we have a successful trace-path from τ_ℓ to $\tau_{\ell+1}$. Each trace (except for τ_ℓ) in the trace-path in turn generates a successful call to $SYMBOL$ with arguments belonging to some phase less than ℓ, and by induction (b), these are correct. Thus $\tau_{\ell+1}$ and $t_{\ell+1}$ are correct. For part (b), we simply note that $j - 1$ is a support phase for such a call to $TRACE$ by the preceding arguments. So by part (a), τ_j and t_j are correct. Then we see that there is a trace-path starting from τ_j as in (4) that is correct. Part (b) simply asserts that the last trace in (4) is correct. This completes our correctness proof.

Complexity. The guessing of the traces τ_i and times t_i on tapes 1 and 2 takes alternating time $O(r \log t)$. If the arguments of $TRACE$ belongs to phase j, then $TRACE$ may recursively call itself with arguments belonging to phase j for $O(\log t)$ times along on a computation path of N. Then $TRACE$ calls $SYMBOL$ which in turn calls $TRACE$ but with arguments belonging to phase $< j$. Now each call to $TRACE$ takes $O(\log t)$ alternating steps just to set up its arguments (just to write down the head positions). Thus it takes $O(\log^2 t)$ alternating steps between successive calls to $SYMBOL$. In the complete computation tree, we make at most r calls to $SYMBOL$ along any path. This gives an alternating time bound of $O(r \log^2 t)$. **Q.E.D.**

COROLLARY 21. $DREVERSAL(r) \subseteq ATIME(r^3)$

We next show that alternating time is at least as powerful as probabilistic time. The proof is based on the following idea: suppose a probabilistic machine accepts an input x in $m \geq 0$ steps and T is the computation tree. If T is finite and all its leaves happen to lie a fixed level $m \geq 0$ ($m = 0$ is the level of the root) then it is easily seen that T is accepting iff the number of accepting leaves is more than half of the total (i.e. more than 2^{m-1} out of 2^m). In general, T is neither finite nor will all the leaves lie in one level. But we see that if T_m is the truncation of T to level m, then T_m if accepting iff the sum of the "weights" of accepting leaves in T_m is more than 2^{m-1}. Here we define a leaf at level i ($0 \leq i \leq m$) to have a weight of 2^{m-i}. It is now easy to simulate a probabilistic machine M that uses time $t(n)$ by a deterministic machine N using space $t(n)$, by a post-order traversal of the tree $T_{t(n)}$. But we now show that instead of deterministic space $t(n)$, alternating time $t(n)$ suffices.

THEOREM 22. *For all complexity functions t, $PrTIME(t) \subseteq ATIME(t)$.*

Proof. Let M be a probabilistic machine that accepts in time t. We describe an alternating machine N that accepts in time t. Let x be an input and T_m be the computation tree of M on x restricted to configurations at level at most m. For any configuration C in T_m, define

$$VAL_m(C) = 2^{m-level(C)}\, Val_{T_m}(C)$$

where Val_{T_m} is, as usual, the least fixed point valuation of T_m. We abuse notation with the usual identification of the nodes of T_m with the configurations labeling them. Thus if C is the root then $VAL_m(C) = 2^m\, Val_\Delta(C)$. If C is not a leaf, let C_L and C_R denote the two children of C. Observe that

$$VAL_m(C) = VAL_m(C_L) + VAL_m(C_R).$$

Regarding $VAL_m(C)$ as a binary string of length $m + 1$, we define for $i = 0, \ldots, m$,

$$\begin{aligned} BIT_m(C, i) &:= i^{th} \text{ bit of } VAL_m(C) \\ CAR_m(C, i) &:= i^{th} \text{ carry bit of the summation } VAL_m(C_L) + VAL_m(C_R) \end{aligned}$$

where we assume that $i = 0$ corresponds to the lowest order bit. It is easy to see that the following pair of mutually recursive formulas hold:

$$\begin{aligned} BIT_m(C, i) &= BIT_m(C_L, i) \overline{\oplus} BIT_m(C_R, i) \overline{\oplus} CAR_m(C, i-1) \\ CAR_m(C, i) &= \lfloor \frac{BIT_m(C_L, i)) + BIT_m(C_R, i) + CAR_m(C, i-1)}{2} \rfloor \end{aligned}$$

Here, $\overline{\oplus}$ denotes[16] the exclusive-or Boolean operation: $b \overline{\oplus} b' = 1$ iff $b \neq b'$. If $i = 0$, $CAR_m(C, i-1)$ is taken to be zero.

[16]Normally, $\oplus$ denotes exclusive-or. We put a bar over $\oplus$ to distinguish it from the probabilistic-or operator.

If C is a leaf, we need not define $CAR_m(C,i)$ but

$$BIT_m(C,i) = \begin{cases} 1 & \text{if } i = m - level(C) \text{ and } C \text{ answers YES;} \\ 0 & \text{otherwise.} \end{cases}$$

To simulate M on input x, N first guesses the value $m = t(|x|)$ in unary in tapes $1, 2$ and 3. Note that M accepts iff $VAL_m(x) > 2^{m-1}$, iff there exists an i, $0 \leq i < m-1$, such that

$$\text{Either} \quad BIT_m(C_0(x), m) = 1 \tag{6}$$
$$\text{or} \quad BIT_m(C_0(x), m-1) = BIT_m(C_0(x), i) = 1. \tag{7}$$

N checks for either condition (6) or (7) by an existential choice. In the latter case, N makes a universal branch to check that $BIT_m(C_0(x), m-1) = 1$ and, for some existentially guessed unary integer $0 < i < m-1$, $BIT_m(C_0(x), i) = 1$.

It remains to describe the subroutine to verify $BIT_m(C,i) = b$ for any arguments C, i, b. It is assumed that just before calling this subroutine the following setup holds. N has the first m cells on tapes 1,2 and 3 marked out. Head 1,2 and 3 are respectively keeping track of the integers i, $level(C)$ and $i + level(C)$, in unary. Moreover, because of the marked cells, it is possible to detect when these values equals 0 or m. The configuration C is represented by the contents and head positions of tapes 4 to $k+3$ (k is the number of work-tapes of M) and the input tape. This setup is also assumed when calling the subroutine to verify $CAR_m(C,i) = b$.

With this setup, in constant time, N can decide if C is a leaf of T_m (i.e. either C is terminal or $level(C) = m$) and whether $i = m - level(C)$. Hence, in case C is a leaf, the subroutine can determine the value of $BIT_m(C,i)$ in constant time. If C is not a leaf, say $C \vdash (C_L, C_R)$, then N guesses three bits, b_1, b_2 and c such that

$$b = b_1 \overline{\oplus} b_2 \overline{\oplus} c.$$

It then universally branches to verify

$$BIT_m(C_L, i) = b_1, \ BIT_m(C_R, i) = b_2, \ CAR_m(C, i-1) = c.$$

It is important to see that N can set up the arguments for these recursive calls in constant time. A similar subroutine for CAR_m can be obtained.

It remains to analyze the time complexity of N. Define the function t_m to capture the complexity of BIT_m and CAR_m:

$$t_m(d,i) = \begin{cases} 1 & \text{if } d = m \\ 1 + t_m(d+1, 0) & \text{if } (d < m) \wedge (i = 0) \\ 1 + \max\{t_m(d+1, i), t_m(d, i-1)\} & \text{else.} \end{cases}$$

The last case corresponds to the recursive calls for BIT_m and CAR_m when $i > 0$. An examination of the recursive equations for BIT_m and CAR_m reveals that the times to compute $BIT_m(C,i)$ and $CAR_m(C,i)$ are each bounded by $O(t_m(d,i))$ where $d = level(C)$. On the other hand, it is easily checked that from the recursive equations that t_m satisfy

$$t_m(d,i) = m - d + i + 1 = O(m)$$

since i and d lie in the range $[0..m]$. This proves that the time taken by N is $O(m) = O(t(|x|))$. **Q.E.D.**

This theorem, together with that in section 5, imply that deterministic space is at least as powerful as probabilistic or alternating time separately:

$$PrTIME(t) \cup ATIME(t) \subseteq DSPACE(t)$$

It is not clear if this can be improved to showing that deterministic space is at least as powerful as probabilistic-alternating time. If this proves impossible, then the combination of probabilism and alternation is more powerful than either one separately. This would not be surprising since probabilism and alternation seems to be rather different computing concepts. Our current best bound on simulating probabilistic-alternating time is given next.

THEOREM 23. *For all complexity functions t, PrA-TIME$(t) \subseteq ATIME(t \log t)$.*

Proof. We use the same notations as the proof of the previous theorem. Let M be a PAM that accepts in time t. Fix any input x and let T_m be the complete computation tree of M on x restricted to levels at most m, and define VAL_m, BIT_m and CAR_m as before. There is one interesting difference: previously, the values m and i in

calls to the subroutines to verify $BIT_m(C,i) = b$ and $CAR_m(C,i) = b$ were encoded in unary. We now store these values in binary (the reader is asked to see why we no longer use unary).

Consider the verification of $BIT_m(C,i) = b$ for any inputs C, i, b, using alternating time. If C is a leaf this is easy. Suppose C is a TOSS-configuration. Then we must guess three bits b_1, b_2, c and verify that $BIT_m(C_L, i) = b_1$, $BIT_m(C_R, i) = b_1$ and $CAR_m(C, i-1) = c$. Here we use an idea from the design of logic circuits: in circuits for adding two binary numbers, the carry-bits can be rapidly generated using what is known as the 'carry-look-ahead' computation. In our context, this amounts to the following condition:

$$CAR_m(C,i) = 1 \iff \text{if there is a } j \ (j = 0, \ldots, i) \text{ such that } BIT_m(C_L, j) = BIT_m(C_R, j) = 1 \text{ and for all } k = j+1, \ldots, i, \text{ either } BIT_m(C_L, k) = 1 \text{ or } BIT_m(C_L, k) = 1.$$

In $O(\log m)$ alternating steps, we can easily reduce these conditions to checking $BIT_m(C', j) = b'$ for some j, b' and C' a child of C. (We leave this detail as exercise.)

Finally, suppose C is a MIN-configuration (a MAX-configuration is handled similarly). By definition $BIT_m(C,i) = BIT_m(C_L, i)$ iff $VAL_m(C_L) < VAL_m(C_R)$, otherwise $BIT_m(C,i) = BIT_m(C_R, i)$. Now $VAL_m(C_L) < VAL_m(C_R)$ iff there exists a j $(0 \le j \le m)$ such that

$$\begin{aligned} BIT_m(C_L, h) &= BIT_m(C_R, h), \text{ (for } h = j+1, j+2, \ldots, m), \\ BIT_m(C_L, j) &= 0, \\ BIT_m(C_R, j) &= 1. \end{aligned}$$

Again, in $O(\log m)$ time, we reduce this predicate to checking bits of $VAL_m(C')$, C' a child of C.

To complete the argument, since each call to check a bit of $VAL_m(C)$ is reduced in $O(\log m)$ steps to determining the bits of $VAL_m(C')$ where $level(C') = level(C) + 1$, there are at most m such calls on any computation path. To generate a call to C', we use $O(\log m)$ time, so that the length of each path is $O(m \log m)$. **Q.E.D.**

EXERCISES

Exercise 0.22: Obtain the tight bound of $\Theta(\log n)$ on the alternating time complexity for the multiplication problem define as follows:

$$L_{mult} = \{x\#y\#z\# : x, y, z \in \{0,1\}^*, x \cdot y = z\}.$$

Use the addressable-input model of machine. ♢

Exercise 0.23*: Can the above alternating simulation of deterministic reversals be improved? In particular, is it true that $DREVERSAL(r) \subseteq ATIME(r^3)$? □

Exercise 0.24: What is the number of alternations in the proof of $DTIME(t) \subseteq ATIME(t/\log t)$? Can this be improved? ♢

END EXERCISES

7.8 Further Generalization of Savitch's Theorem

Savitch's theorem says that for all $s(n) \ge \log n$, $NSPACE(s) \subseteq DSPACE(s^2)$. Chapter 2 gives a generalization of Savitch's theorem. In this section, we further improve this along three directions: (i) by using alternating time, (ii) by allowing small space bounds $s(n)$, i.e., $s(n) < \log n$, and (iii) by extending the class of simulated machines from nondeterministic to alternating machines.

Consider what happens when $s(n) < \log n$. Savitch's proof method gives only the uninteresting result $NSPACE(s) \subseteq DSPACE(\log^2 n)$. Monien and Sudborough [14] improved this so that for $s(n) < \log n$,

$$NSPACE(s) \subseteq DSPACE(s(n) \log n).$$

Using addressable-input alternating machines, Tompa [21] improved the Monien-Sudborough construction to obtain:

$$NSPACE(s) \subseteq ATIME(s(n)[s(n) + \log n])$$

for all $s(n)$. Incorporating both ideas into the generalized Savitch's theorem of chapter 2, we prove the following new result:

THEOREM 24. *For all complexity functions $t(n) > n$,*

$$N\text{-}TIME\text{-}SPACE(t,s) \subseteq ATIME(s(n)\log\frac{n\cdot t(n)}{s(n)})$$

where the alternating machine here is the addressable-input variety.

Proof. Let M be a nondeterministic machine accepting in time-space (t,s). We describe a addressable-input alternating machine N to accept $L(\mathrm{M})$. Let x be any input, $|x| = n$. M begins by existentially guessing $t = t(n)$ and $s = s(n)$ and marking out s cells on each work tape.

We number the n cells of the input tape containing x as $0, 1, \ldots, n-1$ (rather than the conventional $1, \ldots, n$). We will divide the cells $0, 1, \ldots, n-1$ of the input tape into intervals I_w (subscripted by words $w \in \{L,R\}^*$) defined as follows:

$i \in I_w \iff$ the most significant $|w|$ bits in the binary representation of i corresponds to w

where the correspondence between words in $\{L,R\}^*$ and binary strings is given by $L \leftrightarrow 0$ and $R \leftrightarrow 1$. It is also assumed here that the binary representation of i is expanded to exactly $\lceil \log n \rceil$ bits. Clearly I_w is an interval of consecutive integers. For example: with $n = 6$,

$$I_\epsilon = [0..5],\ I_L = [0..3],\ I_R = [4..5],\ I_{RR} = \emptyset.$$

Observe that

$$I_w = I_{wL} \cup I_{wR} \text{ and } |I_{wL}| \geq |I_{wR}| \geq 0.$$

The *L-end* (resp., *R-end*) *cell* of a non-empty interval is the leftmost cell (resp., rightmost) cell in that interval.

A storage configuration is a configuration in which the contents as well as head position of the input tape are omitted. Let S, S' be storage configurations of M, $d, d' \in \{L,R\}$, $w \in \{L,R\}^*$. Let $conf(S,d,w)$ denote the configuration in which the contents and head positions in the work-tapes are specified by S, with input tape containing the fixed x and the input head scanning the d-end cell of interval I_w. In the course of computation, N will evaluate the two predicates *REACH* and *CROSS* defined next. The predicate

$$REACH(S, S', d, d', w, m)$$

holds if there is a computation path π of length at most m from $conf(S,d,w)$ to $conf(S',d',w)$ where the input head is restricted to the interval I_w throughout the computation, and the space used is at most s. Recall that s is the guessed value of the maximum space usage $s(|x|)$. Let $\overline{L}$ denote R and $\overline{R}$ denote L. Then the predicate

$$CROSS(S, S', d, d', w, m)$$

holds if there is a computation path of length at most m from $conf(S,\overline{d},wd)$ to $conf(S',\overline{d'},wd')$ where the input head is restricted to the interval I_w throughout the computation, and the space used is at most s. Observe that the intervals I_{wd} and $I_{wd'}$ used in this definition are adjacent and $I_{wd} \cup I_{wd'} \subseteq I_w$ (if $d = d'$ then this inclusion is proper). We assume in this definition I_{wR} is non-empty; this automatically implies I_{wL} is non-empty. For instance: $CROSS(S, S', L, R, RLR, m)$ holds means there is a path from $conf(S,R,RLRL)$ to $conf(S',L,RLRR)$ of length at most m, as illustrated in the following figure.

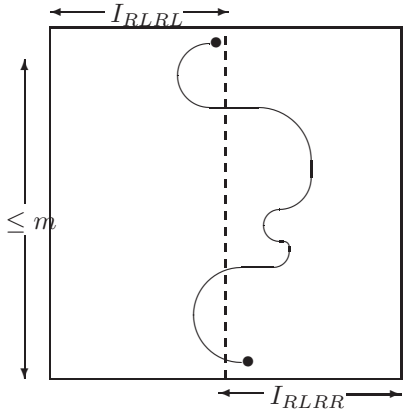

Figure 7.1 The input head positions implied by $CROSS(S, S', L, R, RLR, m)$

We may assume that when M halts, its input head returns to cell 0. The simulation of M amounts to the evaluation of $REACH(S_0, S_f, L, L, \epsilon, t)$ where S_0 is the (trivial) initial storage configuration (independent of x) and S_f is some existentially guessed accepting storage configuration that uses at most s space, λ being the empty string.

We describe the recursive evaluation of $REACH(S, S', d, d', w, m)$. It is assumed that the following setup holds at the time of calling the procedure: S and S' are represented by the contents and head positions on designated tapes of N (if M has k work-tapes, N uses $2k$ tapes for S and S'). The argument w is on the address tape of the N. The value m is written in binary on another tape. One of the following three cases hold:

(i) $|w| = \lceil \log n \rceil$: this condition can be checked in constant time (for instance, by keeping a tally of the length of w) provided that before the simulation begins, the unary representation of $\lceil \log n \rceil$ is guessed and verified once and for all. In this case $|I_w| = 1$ and N enters the read state to read the input symbol $x[w]$ indexed by w. From the definition of *REACH*, from now on N can ignore its input tape and call the predicate

$$REACHABLE(conf(S, d, w), conf(S', d', w), m)$$

where *REACHABLE* is the predicate in the original proof of the generalized Savitch's theorem (chapter 2, section 7). Recall that the predicate $REACHABLE(C, C', m)$ holds if there is a computation path from configuration C to configuration C' of at most m steps using at most s space. The original simulation in chapter 2 uses $O(s \log \frac{m}{s}) = O(s \log \frac{t}{s})$ space, but this is easily modified to give us alternating time $O(s \log \frac{t}{s})$ (Exercise). Another slight modification is to make *REACHABLE*reject at once if input head ever moves at any time during the simulation.

(ii) $|I_{wR}| = 0$: then call $REACH(S, S', d, d', wL, m)$. Note that $|I_{wR}| = 0$ iff the binary number corresponding to wR is greater than $n - 1$. This is easily checked, for instance, by entering the READ state and seeing if we next enter the ERROR state.

(iii) $|I_{wR}| \geq 1$ (so $|w| < \lceil \log n \rceil$): N existentially guesses whether there is a computation path π from $conf(S, d, w)$ to $conf(S', d', w)$ with the input head restricted to I_{wd}. If it guesses 'no' (and it will not make this guess unless $d = d'$) then it next calls

$$REACH(S, S', d, d, wd, m)$$

If it guesses 'yes' (it could make this guess even if $d = d'$) then it chooses existentially two storage configurations S'', S''' in the computation path π and then chooses universally to check one of the following:

$$REACH(S, S'', d, \overline{d}, wd, m),$$
$$CROSS(S'', S''', d, d', w, m),$$
$$REACH(S''', S', \overline{d'}, d', wd', m).$$

N existentially guesses one of the cases (i)-(iii), and then universally checks that its guess is correct as well as performs the respective actions described under (i)-(iii). This completes the description of *REACH*.

The subroutine for $CROSS(S, S', d, d', w, m)$ has two cases:

(i)' $m \leq s$: in this case, N can check the truth of the predicate in time s (since a nondeterministic machine is just a special case of alternation).

(ii)' $m > s$: N guesses two storage configurations S'', S''' and a value $d'' \in \{L, R\}$ and universally branches to check

$$CROSS(S, S'', d, d'', w, m/2),$$
$$REACH(S'', S''', \overline{d''}, \overline{d''}, wd'', m) and$$
$$CROSS(S''', S', d'', d', w, m/2).$$

We should explain this choice of S'', S''', d'': if there is a computation path from $conf(S, \overline{d}, wd)$ to $conf(S', \overline{d'}, wd')$ that makes $CROSS(S, S', d, d', w, m)$ true then this path can be broken up into several disjoint portions where the input head is confined to I_{wL} or to I_{wR} in each portion. Consider the portion π of the path that contains the configuration at time $m/2$: let π be confined to the interval $I_{wd''}$ for some d'', and let the storage configurations at the beginning and end of π be S'' and S''', respectively. With this choice of d'', S'', S''', it is clear that the recursion above is correct.

Note that it is unnecessary to check for $m \leq s$ in *REACH* (since *REACH* does not reduce m in making recursive calls); likewise it is unnecessary to check for $|I_w| = 1$ in *CROSS* (since it is called by *REACH* only with $|I_w| \geq 2$). Observe that every two successive calls to *REACH* or *CROSS* result in either $|w|$ increasing by at least one or m decreasing to at most $m/2$. Hence in $2(\lceil \log n \rceil + \log(t/s)) = O(\log(nt/s))$ recursive calls, we reduce the input to the 'basis' cases where either $m \leq s$ or $|I_w| = 1$. Each recursive call of *REACH* or *CROSS* requires us to guess the intermediate storage configurations S'', S''' in time $O(s)$. Hence in $O(s \log(nt/s)])$ steps we reach the basis cases. In these basis cases, the time used is either $O(s)$ time or that to compute $REACHABLE(C, C', m)$. The latter is $O(s \log(t/s))$ as noted before. The total time is the sum of the time to reache the basis cases plus the time for basis cases. This is $O(s \log(nt/s)])$. **Q.E.D.**

The structure of the preceding proof involves dividing at least one of two quantities in half until the basis case. An immediate consequence of the above results is this:

Corollary 25.
(i) $NSPACE(s) \subseteq ATIME(s^2)$.
(ii) $PrTIME(n^{O(1)}) \subseteq ATIME(n^{O(1)}) = PrA\text{-}TIME(n^{O(1)}) = PSPACE$.

Borodin [3] observed that Savitch's theorem is capable of generalization in another direction. Incorporating Borodin's idea to the previous theorem yields the following "super" Savitch's theorem. Recall the definition of alternating complexity in section 3.

Theorem 26. *Let $t(n) > n, s(n)$ and $a(n)$ be any complexity functions. Then*

$$A\text{-}TIME\text{-}SPACE\text{-}ALTERNATION(t, s, a) \subseteq ATIME(s(n)[a(n) + \log \frac{n \cdot t(n)}{s(n)}])$$

where alternating machines are the addressable-input variety.

Proof. Suppose an alternating machine M accepts in time, space and alternating complexity of $t(n), s(n)$ and $a(n)$. On input x, the machine N begins by guessing the values $t_0 = t(|x|)$ and $s_0 = s(|x|)$. Let $T(x) = T_{t_0,s_0}(x)$ be the computation tree on x restricted to nodes at level $\leq t_0$ and using space $\leq s_0$. There are two procedures involved: The main procedure evaluates a predicate $ACCEPT(C)$ that (for any configuration C as argument) evaluates to true if $C \in T(x)$ and the least fixed point value $Val_T(C)$ of C is equal to 1. The other procedure we need is a variation of the predicate $REACH(S, S', d, d', w, m)$ in the proof of theorem 24. Define the new predicate

$$REACH'(S, S', v, v', w, m)$$

where S, S' are storage configurations, $v, v', w \in \{L, R\}^*$ and $m \geq 1$ such that $|wv| = |wv'| = \lceil \log n \rceil$. Let $conf(S, w)$ where $|w| = \lceil \log n \rceil$ denote the configuration whose storage configuration is S and input tape contains x and the input head is at position indicated by w. The predicate *REACH'* evaluates to true provided there is a computation path π of length $\leq m$ from $conf(S, wv)$ to $conf(S', wv')$ such that all the intermediate configurations C use space $\leq s$ and has input head restricted to the interval I_w. We further require that

(a) C and $conf(S, wv)$ have opposite types, i.e., C is a MIN-configuration if and only if $conf(S, wv)$ is a MAX-configuration. Note that we assume M has no NOT-configurations.

(b) The computation path π we seek must have only configurations of the same type as C with the sole exception of its last configuration (which is equal to $conf(S, wv)$, naturally).

It is clear that we can compute *REACH'* in alternating time $O(s \log t/s)$ as in the case of *REACH*.

The procedure $ACCEPT(C)$ proceeds as follows: suppose C is an MAX-configuration (resp., MIN-configuration). Then the algorithm existentially (resp., universally) chooses in time $O(s)$ a configuration C' with opposite type than C. Let $C = conf(S, v)$ and $C' = conf(S', v')$. Regardless of the type of C, the algorithm existentially chooses to call the following subroutines:

(1) $ACCEPT(C')$

(2) $\neg REACH'(S, S', v, v', \epsilon, t_0)$ where the values S, S', v, v' are related to C, C' as above. Of course, by $\neg REACH'$ we mean that the procedure first enters a NOT-state and then calls *REACH'*. (Here is an occasion where it is convenient to re-introduce NOT-states.) The reader should easily see that the procedure is correct.

We now analyze the complexity of the procedure $ACCEPT$. For any configuration C let T_C be the subtree of configurations reachable from C. Define $depth(C)$ to be the minimum k such that there is prefix T' of T_C such that T' is accepting and each path in T' has at most k alternation. In particular, observe that if x is accepted by M then $depth(C_0(x))$ is at most $a(|x|)$, with $C_0(x)$ the initial configuration. Let $W(k)$ be the (alternating) time required by the procedure for $ACCEPT(C)$ on input C with depth k. Then we have

$$W(k) = O(s) + \max\{s \log \frac{t}{s}, W(k-1)\}.$$

To see this, suppose C is an MAX- (resp., MIN-) configuration. Then $O(s)$ is the time to existentially (resp., universally) choose the configurations C' reachable from C; $s \log t/s$ is the time to decide the predicate $REACH'$; and $W(k-1)$ is the time to recursively call $ACCEPT(C')$. It is easy to deduce that $W(k)$ has solution given by:

$$W(k) = O(s \cdot [k + \log t/s]).$$

The theorem follows immediately. **Q.E.D.**

This is still not the last word on extensions of Savitch's theorem! We return to this in the next chapter.

7.9 Alternating Time is More Powerful than Deterministic Time

This section proves the following important result:

THEOREM 27. *For all t, $DTIME(t) \subseteq ATIME(\frac{t}{\log t})$.*

Tompa and Dymond [6] obtained this result by adapting the result of Hopcroft, Paul and Valiant [12] showing $DTIME(t) \subseteq DSPACE(t/\log t)$. Adleman and Loui [1] gave an interesting alternative proof of the Hopcroft-Paul-Valiant result. The Hopcroft, Paul and Valiant achievement showed for the first time that space is a more powerful resource than time in a general model of computation (namely, multitape Turing machines). For restricted models of computation, Paterson [17] already established that space is more powerful than time for simple Turing machines. Theorem 27 is also an improvement of the result of Paul and Reischuk who simulated deterministic time t using alternating time $O(t \log\log t/\log t)$. In the following, we shall assume the addressable-input model of alternation in case $t/\log t = o(n)$ in the theorem, but otherwise, the regular model suffices.

As an interesting corollary, in conjunction with the alternating time hierarchy theorem at the end of section 4, is that there are languages in $DLBA$ that cannot be accepted in deterministic linear time.

7.9.1 Reduction of Simulation to a Game on Graphs.

First consider the simpler problem of simulating a deterministic Turing machine using as little deterministic space as possible. A key step is the reduction of this problem to a combinatorial question on graphs. Suppose a deterministic k-tape machine M accepts an input in $t > 0$ steps. Our goal is to describe a deterministic machine N that simulates M using as little space as possible.

Let $B = B(t) > 0$ be the *blocking factor*, left unspecified for now. For $i = 0, 1, \ldots, \lceil t/B \rceil$, let

$$\tau_i := iB$$

be *time samples*, and let the cells of each work-tape be grouped into *blocks* consisting of B consecutive cells. For each block b, let *neighborhood*(b) denote the set of 3 blocks consisting of b and the two adjacent blocks on either side of b. We construct a directed acyclic graph $G = (V, E)$ with node set

$$V = \{0, 1, \ldots, \lceil t/B \rceil\}$$

and *labels* for each node. The label for a node i consists of the following two pieces of information:

(i) positions $h_0, \ldots, h_k$ of the $k+1$ tape heads and

(ii) a state q.

We say that this label of node i is *correct* if at time sample τ_i, the machine is in state q and the heads are at positions given by $h_0, \ldots, h_k$. We may say that block b is *visited* in time sample τ_j if the label of node j says that there is a tape head somewhere in b. Note that this definition is relative to the labeling, regardless of its correctness. Once the labels are given, we can define an edge set E as follows. The edges in E are those (i, j) satisfying one of two requirements:

(a) There is a tape block b visited at time sample τ_j, another tape block b' visited time sample τ_i such that $neighborhood(b) \cap neighborhood(b')$ is non-empty and, previous to sample time τ_j, b' is last visited in time sample τ_i.
(b) There is a tape block b visited in time sample τ_j such that $neighborhood(b)$ contains a block that has never been visited in time samples before τ_j, and $i = 0$.

If $(i, j) \in E$ then necessarily $i < j$, and if the labeling is correct then $(i, i + 1)$ must be in E. Let $neighborhood(j)$ denote the union of the blocks in $neighborhood(b)$ where b range over all blocks visited in time sample τ_j. Clearly $neighborhood(j)$ has exactly $3k$ blocks. Let b be visited in time sample τ_j. Then there ≤ 5 blocks b' such that $neighborhood(b') \cap neighborhood(b)$ is non-empty. Each such b' contributes an edge of the form $(i, j) \in E$ for some i. This implies that the indegree of each node in G is $\leq 5k$. The outdegree of G (with the exception of node 0) is similarly bounded by $5k$.

A description of G together with its labels can be written down using at most

$$\frac{t \log t}{B}$$

space. This space is less than t if B is larger than $\log t$. We attempt to find such a graph G by testing successively larger values of t, and for each t, cycling through all ways of assigning labels. It remains to show how to verify a proposed labelling. The idea is that each node in G can be 'expanded' in the following sense: the *expansion* of a node $i \in G$ consists of the contents of the blocks in $neighborhood(i)$ in time sample τ_i. Note that the expansion of i can be encoded using $O(B)$ space. The edges of G define a predecessor-successor relationship: (i, j) is an edge mean that i is a *predecessor* of j, and j the *successor* of i. Next we make an important but elementary observation:

(*) If we already have the expansions of all the predecessors of node $i \geq 1$ then we may expand node i simply by simulating the machine starting from time $\tau_{i-1} = (i - 1)B$.

To do this, we first reconstruct the contents of blocks in $neighborhood(i - 1) \cup neighborhood(i)$, using the expansions of the predecessors of i. (There could be overlap among the predecessor expansions, but it is easy to only use the contents of the most recent version of a block.) Now simulate M starting from time sample τ_{i-1} to time sample τ_i. At the end of the simulation, we may assume that the expansion of node i is now available, in addition to the previous expansions. Details can be filled in by the reader. Let us say that a node i is *verified* if we confirm that its label (i.e., head positions and state) is correct.

(**) If the predecessors of node i are expanded and verified then we can also expand and verify node i.

This is because we can compare the state and head positions in the expansion of i with the labels of node i.

Now we can give a nondeterministic procedure to verify G: nondeterministically expand nodes, one at a time. At any moment, the tapes of the simulator contain some number of expanded nodes. Those nodes whose only predecessor is node 0 can be expanded at any moment; for any other node i, we can only expanded i if all its predecessors are expanded. At the end of expanding node i, we verify the label of i. We may nondeterministically *contract* any previous expansion if we wish; contraction is just the inverse of expansion. Of course we may contract a node only to re-expand it later. The space used by this procedure is $O(B)$ times the maximum number of expanded nodes at any moment. So to minimize space usage, we should contract nodes "at suitable moments". The graph G is said to be verified if its final node $\lceil t/B \rceil$ is verified in this process; we might as well assume that the label of $\lceil t/B \rceil$ always contains the accept state.

It is not hard to see that M accepts its input x iff there is a graph G that is verified by this procedure. We can make this procedure deterministic by cycling through all nondeterministic choices used in the expansion/contraction above. For a space-efficient method of verifying G, Hopcroft, Paul and Valiant showed a general strategy that never store more than $\frac{t}{B \log t}$ expanded nodes at any moment during the verification process. This means that the strategy never use more than $\frac{t}{\log t}$ space since each expanded node uses $O(B)$ space. This proves that $DTIME(t) \subseteq DSPACE(t/\log t)$. The details of this will not be explicitly described since it is essentially subsumed in the Tompa-Dymond alternating time implementation of the strategy, shown next.

7.9.2 A Pebble Game.

Now we transcribe the previous expansion and contraction process for verifying G into a combinatorial game on graphs. We are given a directed acyclic graph $G = (V, E)$ together with a *goal* node $i_0 \in V$. There is only one player in this game. There is an infinite supply of indistinguishable pebbles and each node of G can hold a single pebble. A node is said to be *pebbled* if there is a pebble in it; it is *empty* otherwise. Initially, all the nodes are empty. A

pebbling step consists of placing a pebble on an empty node u, provided all predecessors of u are already pebbled. In particular, we can always pebble an empty *source node* (i.e., a node with no predecessors). An *unpebbling step* consists of removing a pebble from a pebbled node. A *play* is simply a sequence of pebbling or unpebbling steps, with the last step being the pebbling of the goal node i_0. At any moment during the play, there is some number of pebbles on the graph, and our aim (as the player) is to choose the steps in a play in order to minimize the maximum number k of pebbled nodes at any time during the play. This number k is called *the pebble cost* of the play.

The reader will have no difficulty making the connection between this pebble game and the simulation described earlier: pebbling (unpebbling) a node corresponds to expansion (contraction) of nodes.

A *game strategy* is a rule to play the game for any graph. A trivial game strategy is to pebble the nodes in topological order and never to unpebble any nodes. On an n node graph, the pebble cost is n with this strategy. Can we do better in general? The key result here says: *for any directed acyclic graph G on n nodes with in- and out-degree at most d, the strategy yields a play with pebble cost $O_d(n/\log n)$.*

We want an 'alternating version' of playing this pebbling game. As usual, alternation turns the problem inside-out (or rather, bottom-up): instead of proceeding from the source nodes to the goal node i_0, we first ask what is required to pebble the goal node. This is viewed as a "challenge" at node i_0. The challenge at a node u in turn spawns challenges at other nodes (which must include all predecessors of u). This is roughly the idea for our key definition:

Definition 12. A *pebbling tree* for a directed acyclic graph G with goal node i_0 is a finite tree T satisfying the following.[17] Each vertex u of T is associated with a triple $[i, X, Y]$ where X and Y are subsets of nodes of G, and i is a node of G. We called i the *challenged node*, X the *pebbling set*, Y the *unpebbling set* (at vertex u). At each vertex u, define the **current set** $C(u)$ of (currently) pebbled nodes at u by induction on the level of u: if u is the root then $C(u)$ is simply the pebbling set at u; otherwise if u is a child of u' then $C(u) = (C(u') - Y) \cup X$ where X (resp., Y) is the pebbling (resp., unpebbling) set at u. We require these properties:

(i) The challenged node at the root is the goal node i_0.

(ii) At each vertex u associated with $[i, X, Y]$, either $i \in X$ or else all the predecessors of i are contained in the current set $C(u)$.

(iii) If the pebbling set at vertex u is X, then u has $|X|$ children, and the set comprising the challenged nodes at these children is precisely X.

■

Remarks Note that (iii) implies that the pebbling set at a leaf must be empty; (ii) further implies that the predecessors of a challenged node at a leaf u is in $C(u)$. The concept of "pebble" in pebbling trees is distinct from the pebbles in the original pebbling game: in the literature, this distinction is made by giving colors (black and white) to the different concepts of pebbles. We may call the "pebbles" in pebbling trees "challenge pebbles", because they are targets to be achieved in the original pebbling game.

Interpretation: This tree is an abstract description of an alternating computation tree that verifies the labels of a graph G in the sense of the Hopcroft-Paul-Valiant simulation of a deterministic time t machine M. To make this precise, we first describe an alternating machine N that on input a labeled graph G with goal node i_0 behaves as follows: initially, N existentially guesses some expansion of node i_0 and writes this onto tape 1; tape 2 is empty. In general, N is in the following 'inductive stage':

Tape 1 contains the expansion $e(i')$ some node i',
Tape 2 holds some number of expansions of nodes in G.

Then N existentially deletes some expansions in tape 2 and existentially writes some (possibly zero) new expansions in tape 3. If no expansion of node i' is found in tapes 2 and 3 then N tries to produce one: first N checks that all the predecessors of i' are in tapes 2 and 3 (otherwise it rejects) and then simulate M from sample time $\tau_{i'-1}$ to sample time $\tau_{i'}$ and, as usual, assume that we now have an expansion $d(i')$ of i'. N can now verify if the expansion $e(i')$ agrees with the found or newly constructed $d(i')$ (if not, N rejects). To continue, either tape 3 is empty (in which case N accepts) or else N universally chooses to copy one of the expanded nodes from tape 3 to tape 1 and the rest onto tape 2. This completes the 'inductive stage'.

[17] To avoid confusing the nodes of T with those of G, we will refer to the nodes of T as 'vertices'.

We claim that N accepts iff there is a pebbling tree T. Suppose T exists. To show that N accepts, we describe an accepting computation tree T' of N that is modeled after T: each inductive stage of N corresponds to a vertex u of T. If u is associated with the triple $[i, X, Y]$ then tape 1 contains the expansion of node i and tape 2 contains the expansions of the set $C(u)$ of pebbled nodes at u. Then the sets Y, X corresponds (respectively) to the expansions that are deleted from tape 2 or existentially guessed in tape 3. If we choose T' in such a way that the guessed expansions in tape 3 are always correct, then T' would be an accepting computation tree. Conversely, if N accepts, then we can construct the pebbling tree by reversing the above arguments. ■

With this interpretation, it is easy to understand the following definition of complexity. The *pebbling time* at any vertex with label $[i, X, Y]$ is given by $1 + |X| + |Y|$. The pebbling time of a path of T is the sum of the pebbling times of nodes along the path. The *pebbling time* of T is the maximum pebbling time over all paths in T. The *pebbling space* of T is the maximum of $|C(u)|$ over all vertices u. These corresponds to alternating time and alternative space, respectively. Since we do not care about minimizing alternating space in the following proof, we may assume each unpebbling set Y is empty. We leave the following as an exercise:

LEMMA 28. *Let G be any directed acyclic graph and i_0 be a node in G. If there is a pebbling tree for (G, i_0) with pebbling time t, then we can pebble (G, i_0) using t pebbles.*

We come to the main lemma:

LEMMA 29. *Let G be any directed acyclic graph with m edges, and whose indegree and outdegree is at most d. For any goal node i_0 in G, there exists a pebbling tree for (G, i_0) with pebbling time of $O_d(m/\log m)$*

Proof. Let G be any graph described by the lemma and i_0 is any node in G, We describe a pebbling tree for (G, i_0) whose pebbling time is at most $P(m)$, where $P(m) = O_d(m/\log m)$. We may suppose m is sufficiently large. First partition the nodes V of G into two disjoint sets, $V = V_a \cup V_b$ such that

(i) There are no edges from V_b ('nodes below') to V_a ('nodes above'). So edges of G that cross between V_a and V_b must descend from above to below. Let G_a, G_b be the induced subgraphs with nodes V_a, V_b, respectively.

(ii) The number of edges m_b in V_b satisfies

$$\frac{m}{2} - \frac{m}{\log m} \leq m_b < \frac{m}{2} - \frac{m}{\log m} + d.$$

To see that such a partition exists, we offer a construction. Starting with V_b as the empty set, keep adding nodes into V_b in topological order until the number of edges of G_b satisfies the above inequalities (this is possible because additional node in V_b increases the number of edges by at most d).

Let $B \subseteq V_b$ comprise those nodes with at least one predecessor in V_a. Consider three cases.

CASE 1 Suppose the goal node i_0 is in V_a. Then a pebbling tree for (G_a, i_0) is also a pebbling tree for (G, i_0). This tree has a pebbling time at most

$$P(m - m_b) \leq P(\frac{m}{2} + \frac{m}{\log m}).$$

Assume $i_0 \in V_b$ in the remaining cases.

CASE 2 Suppose $|B| < 2m/\log m$. Then we construct the following pebbling tree T for (G, i_0). The pebbling set at the root of T is $B \cup \{i_0\}$. At each child u of the root, we consider two possibilities. If the challenged node at u is $i \in B$, then inductively construct a pebbling tree for (G_a, i) with pebbling time $P(m/2 + m/\log m)$. Otherwise the challenged node is i_0 and we inductively construct a pebbling tree for (G_b, i_0) with pebbling time $P(m/2 - m/\log m + d)$. The result is a pebbling tree for (G, i_0) with pebbling time at most

$$\frac{2m}{\log m} + P(\frac{m}{2} + \frac{m}{\log m}).$$

CASE 3 Suppose $|B| \geq 2m/\log m$. Consider a pebbling tree T_b for (G_b, i_0) with pebbling time $\leq P(m/2 - m/\log m + d)$. We convert T_b into a pebbling tree for (G, i_0): let u be any leaf of T_b with challenged node i. The predecessors of i in G_b are contained in $C(u)$, by definition of T_b. But i may bave predecessors in G but not in G_b: let $X(i)$ be this set of predecessors. We make $X(i)$ the pebbling set at u and create $|X(i)| \leq d$ children for u. Each child v of u has a challenged node $j \in V_a$. We can attach to v a pebbling tree T_j for (G_a, j). Notice G_a has at most $m - m_b - |B| \leq (m/2) - (m/\log m)$ edges since are at least $|B| \geq 2m/\log m$

edges from V_a to V_b are not counted in G_a or G_b. Hence the pebbling time for T_j is at most $P(m/2-m/\log m)$. This completes our description of the pebbling tree for (G, i_0). The pebbling time of this tree is equal to the pebbling time of T_b plus the pebbling time of any T_j's plus at most d. This is at most

$$2P(\frac{m}{2} - \frac{m}{\log m}) + d.$$

Taking the worst of these three cases, we obtain

$$P(m) \leq \max\{P(\frac{m}{2} + \frac{m}{\log m}) + \frac{2m}{\log m}, 2P(\frac{m}{2} - \frac{m}{\log m}) + d\}$$

We want to show that there is a constant $c = c(d) \geq 5$ such that for m', $P(m') \leq cm'/\log m'$. By making c sufficiently large, we may assume that the truth has been established for m large enough. Inductively, we have the the following derivation:

$$\begin{aligned} P(\frac{m}{2} + \frac{m}{\log m}) + \frac{2m}{\log m} &= P(\alpha m) + \frac{2m}{\log m} \quad (\text{where } \alpha = \tfrac{1}{2} + \tfrac{1}{\log m}) \\ &\leq \frac{c\alpha m}{\log(\alpha m)} + \frac{2m}{\log m} \\ &\leq \frac{cm}{\log m}\left(\frac{\alpha \log m}{\log(\alpha m)} + \frac{2}{c}\right) \\ &\leq \frac{cm}{\log m}. \end{aligned}$$

We also have

$$\begin{aligned} 2P(\tfrac{m}{2} - \tfrac{m}{\log m}) + d &\leq \frac{2c\left(\frac{m}{2} - \frac{m}{\log m}\right)}{\log\left(\frac{m}{2} - \frac{m}{\log m}\right)} + d \\ &\leq \frac{cm\left(1 - \frac{2}{\log m}\right)}{\log m + \log\left(\frac{1}{2} - \frac{1}{\log m}\right)} + d \\ &\leq \frac{cm\left(1 - \frac{2}{\log m}\right)}{\log m - 1.1} + d \\ &\leq \frac{cm}{\log m}\left(\frac{\log m - 2}{\log m - 1.1}\right) + d \\ &\leq \frac{cm}{\log m}. \end{aligned}$$

Q.E.D.

We may now complete the proof of the main result showing a deterministic M that accepts in time t can be simulated in alternating time $O(t/\log t)$. In applying the above lemma, we may recall that the graph G obtained from a computation of M by using some blocking factor B has bounded in- and out-degrees except for node 0. To fix this, we can simply place a pebble at node 0 and the rest of the graph is now effectively bounded degree.

(1) Reserve tapes 1, 2 and 3 for later use. First we existentially choose the time t (tape 4) and blocking factor B (tape 5). Then we existentially choose a labeling of the graph G with nodes $V = \{0, \ldots, t/B\}$ (tape 6), and an edge set E (tape 7). Since each label uses $O(\log t)$ space, all this (when choice is correct) takes time $O(\frac{t \log t}{B})$.

(2) Universally choose to verify that E is correct relative to node labels, and to verify that the label of node t/B is correct. It takes time $O(\frac{t \log t}{B})$ to verify E. Verification of the label at node t/B is recursive as shown next.

(3) The verification of node $\lceil t/B \rceil$ amounts to simulating a pebbling tree T for $(G, \lceil t/B \rceil)$ (i.e., with $\lceil t/B \rceil$ as the goal node of G). We do this along the lines given by the "Interpretation" above. As before, each 'inductive

stage' of our simulation of T corresponds to a vertex u of T: if $[i, X, Y]$ is associated with u then an expansion of the challenged node i is available on tape 1. The set of nodes previously expanded are available on tape 2. Since the pebbling time of T can be assumed to be $t/(B\log(t/B))$, we may assume that tape 2 has at most $t/(B\log(t/B))$ nodes. Since each expansion uses $O(B)$ space, tape 2 uses $O(t/\log(t/B))$ space. We existentially choose the pebbling set X at u and also their expansions, writing down these guesses on tape 3. (As noted before, we may assume Y is empty.) We then verify the challenged node i (it must either appear in tapes 2 or 3 or has all its predecessors expanded so that it can be simulated directly). This non-recursive verification of node i takes time $O(t/\log(t/B))$. To get to the next inductive stage, we universally choose to transfer one of the expanded nodes on tape 3 to tape 2, which is now the challenged node.

We have seen that the non-recursive parts of step (3) takes alternating time $O(t/\log(t/B))$. This time must be added to the total alternating time in the recursive parts of the computation. The recursive part of the computation, we claim is $O(B)$ times the pebbling time P. This is because each unit of pebbling time P can be associated with the guessing of an expanded node. But each expansion, when correct, takes space $O(B)$. It follows that the recursive part of the computation takes time $O(t/\log(t/B))$ since the pebbling time for the optimal tree is at most $O(t/[B\log(t/B)])$ (by preceding lemma, with $n = t/B$). Finally, if B is chosen to be $\log^2 t$, we get a time bound of $O(t/\log t)$ for steps (1),(2) and (3). (We could choose B as large as t^ϵ for any constant $\epsilon > 0$.) This concludes the proof of our main theorem.

Finally, we note that the space bound just obtained is the best possible in the sense that $\Omega(t/\log t)$ is a lower bound on the worst case pebbling time for the class bounded in-degree graphs [19].

EXERCISES

Exercise 0.25: Recall the graph $G = (V, E)$ that attempts to represent the time-sampled computation of a deterministic TM on a fixed input. The vertex set is $V = \{0, 1, \ldots, \lceil t/B\rceil\}$. An alternative definition of the edge set E is this: "$(j, i) \in E$ iff $j < i$ and there is some block b visited in the ith time interval and last visited in the jth time interval." What modifications are needed in the rest of the proof which uses G to show that $DTIME(t) \subseteq DSPACE(t/\log t)$? ♢

Exercise 0.26: In the graph G, we say node v is an ancestor of node w if there is a path from v to w (so "ancestor" is the reflexive transitive closure of the "predecessor" relation).
(i) Show that in any pebbling tree T for (G, i_0), it is possible to restrict the pebbling sets at each vertex to the ancestors of i_0.
(ii) Prove lemma 28. **Hint:** Suppose $u_1, \ldots, u_k$ are the children of the root of T and x_i is the challenged node at u_i. Let T_j be the subtree of T rooted at u_j. Renumber the indices so that if x_i is an ancestor of x_j then $i < j$. You strategy need to take into account this topologically sorted sequence on $x_1, \ldots, x_k$. Is it true that T_j is a pebbling tree of (G, x_j)? ♢

Exercise 0.27: Show that if a graph with goal node (G, i_0) has a pebbling tree with pebbling time t then it can be pebbled with t pebbles. Is the converse true? ♢

Exercise 0.28: (Paul, Tarjan, Celoni, Cook)
(a) A *level graph* is a directed acyclic graph with bounded in-degree such that the vertices can be partitioned into 'levels' and edges only go from level i to level $i + 1$. Show that every level graph on n vertices can be pebbled using $O(\sqrt{n})$ pebbles.
(b) Show an infinite family of graphs with indegree 2 that requires $\Omega(\sqrt{n})$ pebbles to pebble certain vertices. ♢

END EXERCISES

7.10 Alternating Space

We show two results from Chandra, Kozen and Stockmeyer that relate alternating space and deterministic time.

Note that for a nondeterministic machine M, if there is an accepting path then there is one in which no configuration is repeated. The next lemma shows an analogous result for alternating machines.

LEMMA 30.
(a) Let M be any choice machine. If T is an accepting computation tree for an input w then there is an accepting computation tree T' for w with the following properties:

- *each computation path in T' is a (prefix of a) computation path in T*
- *if $u, v \in T'$ are vertices such that u is a proper ancestor of v and both u and v are labeled by the same configuration, then $Val_{T'}(v) \sqsubset Val_{T'}(u)$ (strict ordering).*

(b) If, in addition, M is an alternating machine then we can assume that v is a leaf of T'.

Proof. (a) The result follows if we show another accepting computation tree T' with fewer vertices. Let T_v denote the subtree of T rooted at v consisting of all descendents of v. There are two cases: if $Val_T(u) \sqsubseteq Val_T(v)$ then we can form T' from T by replacing T_u with T_v. By monotonicity, T' is still accepting.
(b) This simply follows from part (a) since for alternating machines, $Val_T(v) \sqsubset Val_T(u)$ implies that $Val_T(v) = \perp$. In that case, we might as well prune away all proper descendents of v from T.

Q.E.D.

THEOREM 31. *For all complexity functions s,*

$$ASPACE(s) \subseteq DTIME(n^2 \log n O(1)^s).$$

Proof. Let M be an ordinary alternating machine accepting in space s. Later we indicate how to modify the proof if M has the addressable-input capability. We will construct a deterministic N that accepts $L(M)$ in the indicated time. On an arbitrary input w, N proceeds in stages: in the mth stage, N enumerates (in tape 1) the set Δ_m defined to be all the configurations C of M on input w where C uses at most m space. Note that each configuration can be stored in $m + \log n$ space, and there are $nO(1)^m$ configurations, so we use

$$(m + \log n) nO(1)^m = n \log n O(1)^m$$

space on tape 1. Then N enumerates (in tape 2) the edges of the computation tree T_m whose nodes have labels from Δ_m and where no node $u \in T_m$ repeats a configuration that lie on the path from the root to u, except when u is a leaf. Clearly this latter property comes from the previous lemma. Using the information in tape 1, it is not hard to do this enumeration of edges in a 'top-down' fashion (we leave the details to the reader). Furthermore each edge can be produced in some constant number of scans of tape 1, using time $n \log n O(1)^m$. Thus the overall time to produce all $nO(1)^m$ edges is $n^2 \log n O(1)^m$. Now we can compute the least fixed point $Val_{T_m}(u)$ value at each node $u \in T_m$ in a bottom-up fashion, again $O(n \log n O(1)^m)$ per node for a total time of $n^2 \log n O(1)^m$. This completes our description of the mth stage.

The previous lemma shows that if M accepts w then at some mth stage, $m \leq s(|x|)$, T_m is accepting. Since the time for the mth stage is $n^2 \log n O_1(1)^m$, the overall time over all stages is

$$\sum_{m=1}^{s(n)} n^2 \log n O_1(1)^m = n^2 \log n O_2(1)^{s(n)}.$$

It remains to consider the case where M has the addressable-input capability. We first note that we never have to use more than $O(\log n)$ space to model the address tape (if M writes more than $\log n$ bits, we ignore the tape from that point onwards since it will lead to error when a READ is attempted). Hence the above space bounds for storing a configuration holds. Furthermore, the time to generate the contents of tapes 1 and 2 remains asymptotically unchanged. Similarly for computing the least fixed point Val_{T_m}. This concludes the proof. **Q.E.D.**

THEOREM 32. *For all $t(n) > n$, $DTIME(t) \subseteq ASPACE(\log t)$.*

Proof. Let M accept in deterministic time t. We describe an alternating machine N to accept $L(M)$ in space $\log t$. For this simulation, N can be the ordinary variety of alternating machine. We may assume that M is a simple Turing machine and M never moves its tape head to the left of its initial position throughout the computation. (Otherwise, we first convert M into a simple Turing machine accepting in time $t(n)^2$ with these properties. How?) Let x be any word accepted by M. Let $C_0, C_1, \ldots, C_m$, $m = t(|x|)$, be the unique accepting computation path of M on x. We assume that the final accepting configuration is repeated as many times as needed in this path. Let each C_i be encoded as a string of length $m + 2$ over the alphabet

$$\Gamma = \Sigma \cup [Q \times \Sigma] \cup \{\sqcup\}$$

where Σ is the tape alphabet of M, Q the state set of M, and $[Q \times \Sigma]$ is the usual composite alphabet. Furthermore, we may assume the the first and last symbol of the string is the blank symbol $\sqcup$. Let $\alpha_{i,j}$ denote the jth symbol in configuration C_i $(i = 0, \ldots, m; j = 1, \ldots, m + 2)$.

N will be calling a subroutine $CHECK(i, j, b)$ that verifies whether $\alpha_{i,j} = b$ where $b \in \Gamma$. N begins its overall computation by existentially choosing the integer m, the head position h ($1 \leq h \leq m+2$) and a symbol $b' = [q_a, c]$ and then it calls $CHECK(m, h, b')$. The subroutine is thus verifying that that M is scanning the symbol c at position h when M enters the accept state q_a. All integer arguments are in binary notation.

In general, the subroutine to verify if $\alpha_{i,j} = b$ (for any i, j, b) operates as follows: if $i = 0$ or $j = 1$ or $j = m$, N can directly do the checking and accept or reject accordingly. Otherwise, it existentially chooses the symbols b_{-1}, b_0, b_{+1} such that whenever b_{-1}, b_0, b_{+1} are consecutive symbols in some configuration of M then in the next instant, b_0 becomes b. Now N universally chooses to call

$$CHECK(i-1, j-1, b_{-1}), CHECK(i-1, j, b_0), CHECK(i-1, j+1, b_{+1}).$$

It is important to realize that even if b does not contain the tape head (i.e., $b \notin [Q \times \Sigma \times I]$), it is possible for b_{-1} or b_{+1} to contain the tape head. The space used by N is $O(\log m)$.

Correctness. If M accepts then it is clear that N accepts. The converse is not obvious. To illustrate the subtlety, suppose $CHECK(i, j, b)$ accepts because $CHECK(i-1, j-\epsilon, b_\epsilon)$ accepts for $\epsilon = -1, 0, 1$. In turn, $CHECK(i-1, j-1, b_{-1})$ accepts because $CHECK(i-2, j-1, b')$ (among other calls) accepts for some b'; similarly $CHECK(i-1, j, b_0)$ accepts because $CHECK(i-2, j-1, b'')$ (among other calls) accepts for some b''. But there is no guarantee that $b' = b''$ since these two calls occur on different branches of the computation tree. Another source of inconsistency is that the existential guessing of the b_j's may cause more than one tape head to appear during one configuration. Nevertheless, we have a correctness proof that goes as follows. First observe that if $CHECK(i, j, b)$ is correct if $i = 0$. Moreover, given j there is a unique b that makes $CHECK(0, j, b)$ accept. Inductively, assume $CHECK(i, j, b)$ is correct for all j, b and that $CHECK(i, j, b)$ and $CHECK(i, j, b')$ accept implies $b = b'$. Then it is easy to see that $CHECK(i+1, j, b)$ must be correct for all j, b; moreover, because of determinism, $CHECK(i+1, j, b)$ and $CHECK(i+1, j, b')$ accept implies $b = b'$. [This is why *CHECK* must universally call itself three times: for instance, if *CHECK* only makes two of the three calls, then the correctness of these two calls does not imply the correctness of the parent.] So we have shown that the symbols $\alpha_{i,j}$ are uniquely determined. In particular $\alpha_{m,h} = b'$ in the initial guess is correct when the machine accepts. **Q.E.D.**

The reader should see that this proof breaks down if we attempt to simulate nondeterministic machines instead of deterministic ones.

Combining the two theorems yields the somewhat surprising result:

COROLLARY 33. *For $s(n) \geq \log n$, $ASPACE(s) = DTIME(O(1)^s)$.*

EXERCISES

Exercise 0.29: Obtain a lower bound on the space-time product of alternating Turing machines which accepts the palindrome language L_{pal}. ♢

END EXERCISES

7.11 Final Remarks

This chapter introduced valuations as a mathematical framework to study most of the known choice models of computation. We can extend the basic framework to other value sets S (analogous to the role of INT) provided S is equipped with a partial order $\sqsubseteq$ such that limits of $\sqsubseteq$-monotonic chains are in S and S has a $\sqsubseteq$-least element (such an S is called a *complete partial order*). The reader familiar with Scott's theory of semantics will see many similarities. For relevant material, see [22].

We have a precise relationship between alternating space and determistic time: for $s(n) \geq \log n$,

$$ASPACE(s) = DTIME(O(1)^s). \tag{8}$$

What is the precise relationship between alternating time and deterministic space? Although we have tried to emphasize that alternating time and deterministic space are intimately related, they are not identical. We know that

$$ATIME(s) \subseteq DSPACE(s) \subseteq NSPACE(s) \subseteq ATIME(s^2). \tag{9}$$

for $s(n) \geq \log n$. How 'tight' is this sequence? It is unlikely that that the first two inclusions could be improved in the near future.

From (8) and (9), we get find new characterizations of some classes in the canonical list:

$$\begin{aligned} P &= ASPACE(\log n) \\ PSPACE &= ATIME(n^{O(1)}) \\ DEXPT &= ASPACE(n) \\ EXPS &= ATIME(O(1)^n). \end{aligned}$$

What is the relationship of alternating reversal with deterministic complexity? Of course, for alternating machines, we must take care to simultaneously bound reversal with either time or space in order to get meaningful results. Other complexity measures for alternating machines have been studied. Ruzzo [20] studied the *size* (i.e., the number of nodes) of computation trees. In particular, he shows

$$A\text{-}SPACE\text{-}SIZE(s(n), z(n)) \subseteq ATIME(s(n) \log z(n)).$$

King [13] introduced other measures on computation trees: *branching* (i.e., the number of leaves), *width* (below), *visit* (the maximum number of nodes at any level). Width is not so easy to motivate but in the special case of binary trees (which is all we need for alternating machines), it is the minimum number of pebbles necessary to pebble the root of the tree. Among the results, he shows (for $s(n) \geq \log n$),

$$\begin{aligned} A\text{-}SPACE\text{-}WIDTH(s(n), w(n)) &= NSPACE(s(n)w(n)), \\ A\text{-}SPACE\text{-}VISIT(s(n), v(n)) &\subseteq ATIME(s^2(n)v(n)). \end{aligned}$$

Exercises

Exercise 0.30: A **probabilistic-alternating finite automaton** (pafa) is a PAM with no work-tape and whose input tape is restricted so that in each step, the input head must move to the right. Moreover, the machine must halt upon reading the first blank symbol after the input word. The special cases of alternating finite automaton (afa) or probabilistic finite automaton (pfa) is defined analogously.
(i) (Chandra-Kozen-Stockmeyer) Prove that an afa accepts only regular languages.
(ii) (Starke) Show the following language (well-known to be non-regular) $L = \{0^m 1^n : m \geq n\}$ can be accepted by a pfa. ♢

Exercise 0.31: Recall an alternating finite automaton (afa) defined in the previous question. Let us define a **generalized afa** (gafa) to be an afa in which each state q has an arity $k(q) \geq 0$ and is assigned a **generalized Boolean function**

$$\gamma(q) : \{0, 1, \perp\}^{k(q)} \to \{0, 1, \perp\}.$$

(i) Let M be an alternating Turing acceptor with no work-tapes and whose input head cannot move right. For the purposes of this question, assume that M has only $\wedge$- and $\vee$-states. So M looks very similar to an afa, except that it can make "ϵ-moves": this is a terminology from automata theory: an ϵ-move is one where the automaton can change state without moving its input head, and this is performed non-deterministically. Show that such moves can be eliminated in the sense that there is a generalized afa that accepts the same language as M.
(ii) (Open) Characterize the kinds of generalized Boolean functions that can arise in the generalized afa in part (i).

♢

Exercise 0.32: (Freivalds) Show a probabilistic finite automata to recognize the language $\{0^n 1^n : n \geq 1\}$ with bounded error. **Hint:** We are basically trying to check that the number of 1's and number of 0's are equal. Show that the following procedure works:
a. Choose a coin with probability $p << \frac{1}{2}$ of getting a *head*.
b. Toss coin for each 0 in input and each 1 in input. If we get all *heads* for 0's and at least one *tail* for the 1's then we say we have a 0-*win*. If we get at least one *tail* for the 0's and all *heads* for the 1's, we have a 1-*win*. All other outcomes result in a *tie*.
c. Repeat this experiment until we have at least k 0-wins or k 1-wins. We accept if and only if there is at least one 1-win and at least one 0-win.
(For more information, see [10].) ♢

Exercise 0.33: (Ruzzo, King) Show the following simulations among measures for alternating computation, as stated in the concluding section: for $s(n) \geq \log n$,

$$\begin{aligned} \mathit{A\text{-}SPACE\text{-}SIZE}(s(n), z(n)) &\subseteq \mathit{ATIME}(s(n) \log z(n)), \\ \mathit{A\text{-}SPACE\text{-}WIDTH}(s(n), w(n)) &= \mathit{NSPACE}(s(n)w(n)), \\ \mathit{A\text{-}SPACE\text{-}VISIT}(s(n), v(n)) &\subseteq \mathit{ATIME}(s^2(n)v(n)). \end{aligned}$$

♢

Exercise 0.34: (King) Recall the definitions of branching and width resource for alternating machines in the concluding section. Show that the branching resource (simultaneously bounded with space) has the linear complexity reduction property: for $s(n) \geq \log n$,

$$\mathit{A\text{-}SPACE\text{-}BRANCH}(s(n), b(n)) = \mathit{A\text{-}SPACE\text{-}BRANCH}(s(n), b(n)/2).$$

Show the same result for width resource: for $s(n) \geq \log n$,

$$\mathit{A\text{-}SPACE\text{-}WIDTH}(s(n), w(n)) = \mathit{A\text{-}SPACE\text{-}WIDTH}(s(n), w(n)/2).$$

♢

End Exercises

Bibliography

[1] L. Adleman and M. Loui. Space-bounded simulation of multitape Turing machines. 14:215–222, 1981.

[2] L. Babai and S. Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computers and Systems Science*, 36:254–276, 1988.

[3] A. Chandra, D. Kozen, and L. Stockmeyer. Alternation. *J. ACM*, 28:1:114–133, 1981.

[4] A. Condon and R. Ladner. Probabilistic game automata. *Journal of Computers and Systems Science*, 36:452–489, 1988.

[5] P. Crawley and R. Dilworth. *Algebraic theory of lattices*. Prentice-Hall, 1973.

[6] P. Dymond and M. Tompa. Speedups of deterministic machines by synchronous parallel machines. *Journal of Computers and Systems Science*, 30:149–161, 1985.

[7] J. T. Gill. Computational complexity of probabilistic Turing machines. *SIAM J. Comp.*, 6(4):675–695, 1977.

[8] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proofs. *17th ACM Symposium STOC*, pages 291–304, 1985.

[9] S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. *18th ACM Symposium STOC*, pages 59–68, 1986.

[10] A. G. Greenberg and A. Weiss. A lower bound for probabilistic algorithms for finite state machines. *Journal of Computer and System Sciences*, 33:88–105, 1986.

[11] J. Hong. *Computation: Computability, Similarity and Duality*. Research notices in theoretical Computer Science. Pitman Publishing Ltd., London, 1986. (available from John Wiley & Sons, New York).

[12] J. E. Hopcroft, W. J. Paul, and L. G. Valiant. On time versus space. *Journal of the ACM*, 24:332–337, 1977.

[13] K. N. King. Measures of parallelism in alternating computation trees. 13:189–201, 1981.

[14] B. Monien and I. H. Sudborough. On eliminating nondeterminism from Turing machines which use less than logarithm worktape space. In *Lecture Notes in Computer Science*, volume 71, pages 431–445, Berlin, 1979. Springer-Verlag. Proc. Symposium on Automata, Languages and Programming.

[15] R. E. Moore. *Interval Analysis*. Prentice-Hall, Englewood Cliffs, N.J., 1966.

[16] C. H. Papadimitriou. Games against nature. *Journal of Computers and Systems Science*, 31:288–301, 1985.

[17] M. S. Paterson. Tape bounds for time-bounded Turing machines. *Journal of Computers and Systems Science*, 6:116–124, 1972.

[18] W. J. Paul, E. J. Praus, and R. Reischuk. On alternation. *Acta Informatica*, 14:243–255, 1980.

[19] W. J. Paul, R. E. Tarjan, and J. R. Celoni. Space bounds for a game on graphs. 11:239–251, 1977. (See corrections in Math. Systems Theory, 11(1977)85.).

[20] W. L. Ruzzo. Tree-size bounded alternation. 11:352–359, 1979.

[21] M. Tompa. An improvement on the extension of Savitch's theorem to small space bounds. Technical Report Technical Report No. 79-12-01, Department of Computer Sci., Univ. of Washington, 1979.

[22] K. Weihrauch. *Computability*. Springer-Verlag, Berlin, 1987.

[23] C. K. Yap. On combining probabilistic and alternating machines. Technical report, Univ. of Southern California, Comp. Sci. Dept., January 1980. Technical Report.

Contents

Chapter 8

Stochastic Choices

April 13, 2009

We continue investigating the choice-mode of computation. This chapter focuses on the stochastic choices, viz., coin-tossing $\oplus$, probabilistic-and $\otimes$ and probabilistic-or $\oplus$. In contrast to alternation, we see that a rich theory arises when we restrict computational errors in stochastic computation.

In this chapter, it is essential to revert to the use of intervals when discussing valuations. For convenience, an appendix on basic probabilistic vocabulary is included.

8.1 Errors in Stochastic Computation

Two new phenomena arise with stochastic choices:

- Infinite loops in stochastic computation is an essential feature, rather than one we seek to eliminate (cf. alternation). This will be evident when we study space-bounded computations in section 3.
- An extremely rich theory arises from quantifying the forms of computational error. We think of error as a new computational resource.

Forms of computational error. Let M be a choice machine and suppose $Val_M(w) = [b, c]$ where w is an input word. If M accepts w (*i.e.*, $b > \frac{1}{2}$) then both $1 - b$ and $1 - c$ are useful measures of error. Since $\frac{1}{2} > 1-b \geq 1-c \geq 0$, we may call $1-b$ and $1-c$ (respectively) the **pessimistic acceptance error** and **optimistic acceptance error**. Similarly, if M rejects w, then b and c are (respectively) the **optimistic rejection error** and **pessimistic rejection error**. We have two basic paradigms for classifing errors: for optimistic errors, we say M has "zero error" if for all inputs, its optimistic error is 0. For pessimistic errors, we say M has "bounded-error" if its pessimistic errors are bounded away from $\frac{1}{2}$ by a positive constant.

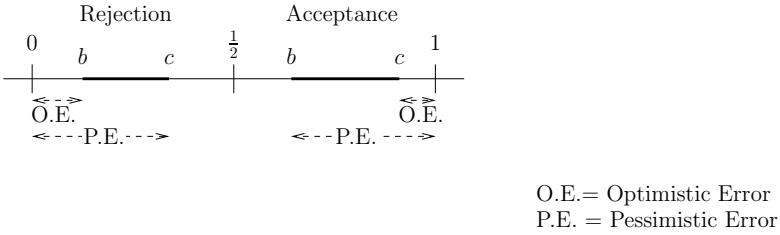

Figure 8.1: Forms of Error.

Stochastic computation has been studied since the early days of automata theory [9]. One original motivation[1] is the fundamental question of synthesizing reliable components from unreliable ones [22]. Stochastic computation in complexity theory started with the work of Gill [12].

Example 1. Gill was in turn motivated by some surprising probabilistic algorithms for primality testing algorithms due to Rabin [23] and Solovay-Strassen [33]. These primality testing algorithms share a common property: every computation path terminates and at each terminal configuration C,

[1] Modern computer hardware is remarkably reliable (compared to software). In the early days of computing, this reliability could not be taken for granted.

1) if C answers YES (i.e., claims that the input is prime) then there is a "small probability" that it made an error;

2) if C answers NO (i.e., claims that the input is composite) then it is surely correct.

The algorithm does not have YO-answers. It is not so easy to quantify the "small probability" of the YES answers. But we may rephrase this property, taking the viewpoint of the inputs:

1′) if the input is prime then all local answers are YES;

2′) if the input is composite then the local answer is NO with "high probability".

The reader should verify that 1′) and 2′) (the global perspective) are just reformulations of 1) and 2) (the local perspective). We elaborate on this because it is easy to become confused by a mixture of these two perspectives. The global formulation is more useful because we can explicitly quantify the "high probability" in it: it means with probability more than 3/4. We see that these primality algorithms make no errors whenever they (globally) accept.

> To remember which way the primality algorithms may have errors, it is helpful to know this: the algorithms answers NO (*i.e.*, claims that an input is composite) only if it finds a "witness" of compositeness. Hence NO-answers are never wrong by virtue of such witnesses. In contrast, the algorithm answers YES (*i.e.*, claims an input is prime) based on its failure to find a witness – such a YES answer could be wrong because an individual path does not exhaustively search for witnesses. In fact the algorithm looks for witnesses by randomly picking candidate witnesses and then checking each for the witness property. An example of a witness of the compositeness of a number n is a factor m $(1 < m < n)$ of n. It is easy to check if any proposed witness m is a factor. Unfortunately, there may be as few as two witnesses for some composite numbers; in this case, the method of randomly selecting candidate witnesses for testing is unreliable. The algorithms of Rabin and of Solovay-Strassen solve this by using more sophisticated concepts of "witnesses" where it can be shown that each composite number has a positive fraction of witnesses among the candidates, and a prime number has no witness.

This example will be used again to illustrate error concepts. ■

We now classify Turing machines according to their error properties.

Definition 1.
(i) A non-empty interval g containing the value $\frac{1}{2}$ is called an *error gap*. Thus g has one of the forms

$$[a,b], \quad (a,b], \quad [a,b), \quad (a,b)$$

where $0 \leq a \leq \frac{1}{2} \leq b \leq 1$. We call a and b (respectively) the **lower** and **upper bounds** of g.
(ii) A choice machine M **accepts with error gap** g if for all accepted inputs w,

$$Val_M(w) \cap g = \emptyset.$$

Nothing is assumed if w is rejected or undecided. Similarly, we define **rejection with error gap** g. Finally, M **has error gap** g if for all input w, it accepts w or rejects with error gap g. Thus M is decisive.
(iii) We say M **accepts with bounded-error** if there exists e $(0 < e \leq \frac{1}{2}$ such that for all accepted inputs w, $Val_M(w) \cap [\frac{1}{2} - e, \frac{1}{2} + e] = \emptyset$. Similarly, M **rejects with bounded-error** if for all rejected inputs w, $Val_M(w) \cap [\frac{1}{2} - e, \frac{1}{2} + e] = \emptyset$. Also, M has **bounded-error** if for all inputs, $Val_M(w) \cap [\frac{1}{2} - e, \frac{1}{2} + e] = \emptyset$. Say M has **unbounded-error** if it does not have have bounded-error. ■

While bounded error focuses on pessimistic errors, the next set of definitions focus on optimistic errors.

Definition 2. (Continued)
(iv) We say M **accepts with zero-error** if for all accepted w, the upper bound of $Val_M(w)$ is 1. Similarly, M **rejects with zero-error** if for all rejected w, the lower bound of $Val_M(w)$ is 0. We say [2] M has **zero-error** if it is decisive, and it has zero-error acceptance and zero-error rejection.
(v) We combine bounded-error and zero-error: M has **bounded zero-error rejection** if it has bounded-error and zero-error rejection. Bounded zero-error rejection is also called **one-sided error**. By symmetry, we say M has **bounded zero-error acceptance** it has bounded-error and zero-error acceptance. Finally, M has **bounded zero-error** if it has bounded-error and zero-error. ■

[2] We emphasize that "zero-error" does not imply the absence of all errors: it only refers to the optimistic errors. It also seems that "errorless" would be preferable to "zero-error" in this set of terminology.

Let us briefly discuss the significance of these forms of error and their motivations. Although our error concepts treat acceptance and rejection with an even hand, we still favor acceptance when it comes to defining languages and complexity classes: for any machine M, the notation

$$L(M)$$

continues to refer to the language **accepted** by M. So a word $w \notin L(M)$ is either rejected or undecided by M.

Bounded Errors. Note that undecided intervals in INT are error gaps. Clearly a machine has error gap if and only it has the minimal error gap $[\frac{1}{2}, \frac{1}{2}]$, if and only if it is decisive. In this sense, bounded-error is a strengthening of halting deterministic computations (which are decisive). Students sometimes confuse the definition of "decisiveness" (a global condition) with the local condition that every path gives a YES or NO answer (and hence each path is finite). For instance, a decisive Turing machine could have infinite computation trees on each of its inputs.

To understand the significance of bounded-error, note that in general, acceptance and rejection errors can get arbitrarily close to $\frac{1}{2}$. This behavior is forbidden by bounded-error. The next section shows that with bounded-error we can modify a machine to yield error gaps of the form $[\epsilon, 1-\epsilon]$ for any desired constant $0 < \epsilon < \frac{1}{2}$, at a cost of increasing the computational time by a constant factor depending on ϵ. This yields a very important conclusion: *assuming we can tolerate constant factor slowdowns, bounded-error algorithms are practically as good as deterministic algorithms.* This is because any physical realization of a deterministic algorithm will still be subject to a small $\epsilon^* > 0$ probability of error, from quantum effects, manufacturing imperfections, etc. Thus, all real computer programs are bounded error algorithms.

Nondeterministic vs. probabilistic machines. Let N be a nondeterministic machine. The valuation function Val_N has values in $\{0, 1, \perp\}$. It follows that N accepts with no pessimistic error. Next, let us see what happens when N is converted into a probabilistic machine M, simply by declaring each state a toss-state. If N does not accept an input, M also does not accept. Hence we have

$$L(M) \subseteq L(N).$$

A further modification to M yields a probabilistic machine M' that accepts the same language as N: let M' begin by tossing a coin in the start state, and on one outcome it immediately answers YES, and with the other outcome it simulates N. So M' is undecided on input $w \notin L(N)$, since the lower bound of $Val_{M'}(w)$ is exactly $\frac{1}{2}$. This shows

$$NTIME(t) \subseteq PrTIME(t+1). \tag{1}$$

A simple modification to M' will ensure that it has zero-error acceptance: simply make all terminal which answer NO to answer YO instead. Notice that N, M and M' are not decisive in general.

Zero-error computation. The primality testing algorithms above accept with zero-error. The concept of zero-error is best understood in terms of stochastic machines with no negation states: in this case acceptance with zero-error means that if an input is accepted then no computation path leads to a NO-configuration. Similarly, "rejecting with zero-error" means that if an input is rejected, then no computation path leads to a YES-configuration. In either case, YO-configuration (or looping) is not precluded. Because of monotonicity properties, if a complete computation tree $T_M(w)$ accepts with zero-error then any prefix T' of $T_M(w)$ also accepts with zero-error, if T' accepts at all.

One-sided error. One-sided error is also motivated by the primality algorithms. These algorithms have no pessimistic errors on prime inputs, by virtue of property $1'$), and have bounded error on composite inputs, by property $2'$). Thus such an algorithm has bounded zero-error acceptance. Now with a trivial change, we can regard the same algorithm as a recognizer of composite numbers: it answers YES iff it finds a witness of compositeness. This new algorithm has bounded zero-error rejection, i.e., it has one-sided error.

This begs the question as to why we define one-sided error in favor of "zero-error rejection" over "zero-error acceptance". We suggest that the (non-mathematical) reason has to do with our bias towards nondeterministic machines: *a probabilistic machine M with one-sided error can be regarded as a nondeterministic machine N.* Let us clarify this remark. For, if M accepts w then some terminal configuration gives a YES-answer, and so N accepts w. Conversely, if M does not accept w, there is no YES-configuration because M has zero-error rejection. Hence, N does not accept w. We conclude that $L(M) = L(N)$. This proves

$$PrTIME_1(t) \subseteq NTIME(t). \tag{2}$$

The subscript "1" in "$PrTIME_1(t)$" refers to one-sided error (the general convention is described below).

The literature often discuss errors in the context of probabilistic machines that run in time $t(n)$ for some time-constructible t (e.g., t is a polynomial). In this case, we can simplify. For instance, we need not deal with intervals: if a stochastic machine does not terminate within t steps, we simply answer NO. This avoids the value $\perp$ and the pessimistic and optimistic errors coincide. This approach does not work in space-bounded computations, for instance.

Error-Restricted Complexity. Since error is viewed as another computational resource, we combine error bounds with other resource bounds. The number of logical possibilities is large, but happily, only a few forms are important. For instance, we exclusively use constant gap functions in complexity characteristics.

The following illustrates the combination of acceptance time with bounded-error or with zero-error:

Definition 3. Let t be a complexity function. A choice machine M **accepts in time** t **with bounded-error** if there is an error gap $g = [a, b]$, $a < \frac{1}{2} < b$ such that for all accepted inputs w, there is an accepting computation tree T such that $Val_T(w) \cap g = \emptyset$ and each path in T has length at most $t(|w|)$. ■

Note that the time bound and error gap are simultaneously achieved in a single computation tree T. To appreciate this, suppose another computation tree T' of height $t' < t$ is a prefix of T. Because of monotonicity, we have $Val_{T'}(w) \sqsubseteq Val_T(w)$. So it is possible that $Val_{T'}(w) \cap g$ is non-empty and yet T' is accepting. So M accepts in time t' but not necessary with the same gap g. In general, all complexity characteristics we specify for a machine are assume to be simultaneous unless otherwise stated.

We can add any of the error restrictions (bounded-error, zero-error, one-sided error, etc) to any of the usual complexity resource (time, space, reversal, simultaneous time-space, etc) bounds. Thus we can speak of M accepting in polynomial space with one-sided error.

We can also extend acceptance complexity to "rejection complexity" and to "running complexity". For rejection complexity, just replace acceptance by rejection. For running complexity, we make decisiveness a pre-requisite. Then a running complexity bound is simply a common bound on both accepting complexity and rejecting complexity.

Notation for error-restricted complexity classes. For simplity, we assume that *the machines used in defining error-restricted classes must be decisive and that running complexity is used.* However, we do not assume our machines to be halting (*i.e.*, halts on all computation paths). A machine can be decisive without being halting; conversely, it can be halting without being decisive. Indeed, in space bounded computation, halting is not necessarily desirable. To refer to such classes, it is sufficient to augment our previous convention for complexity classes, simply by introducing new subscripts. Until now, we have only one kind of subscript, 'r', denoting running complexity. We now introduce three new subscript z,

$$z \in \{b, 0, 1\},$$

to indicate the following restrictions: *bounded-error* ($z = b$) or *one-sided error* ($z = 1$) or *bounded zero-error* ($z = 0$). Note a linear hierarchy in these subscripts: for instance, for most complexity functions t, we have

$$PrTIME_0(t) \subseteq PrTIME_1(t) \subseteq_1 PrTIME_b(t+2) \subseteq PrTIME_r(t+2) \subseteq PrTIME(t+2).$$

The only case in the above inclusions where t must be restricted is $PrTIME_1(t) \subseteq_1 PrTIME_b(t+2)$ where we require t be to time-constructible. This inclusion follows from the following general result:

We could replace *PrTIME* here by other suitable mode-resource pairs. These notations are illustrated in the last column of the following table.

The table below lists some important time-feasible (i.e., polynomial time) complexity classes, under the various choice and error modes:

Some Polynomial Time Stochastic Classes

Error Mode	*Choice Mode*	*Common Symbol*	*Generic Notation*
Unbounded error	$\{\oplus\}$	$\mathbb{P}$	$PrTIME(n^{O(1)})$
Bounded error	$\{\oplus\}$	*BPP*	$PrTIME_b(n^{O(1)})$
One-sided error	$\{\oplus\}$	*RP* (also denoted *VPP* or *R*)	$PrTIME_1(n^{O(1)})$
Bounded zero-error	$\{\oplus\}$	*ZPP*	$PrTIME_0(n^{O(1)})$
Bounded error	$\{\oplus, \vee\}$	*IP*,	$IpTIME_b(n^{O(1)})$
Bounded error	$\{\oplus, \vee\}$	*AM*	$IpTIME\text{-}_(n^{O(1)}, O(1))$

Remark: $\mathbb{P}$ is the only class in this table that is not error-restricted. Hence $\mathbb{P}$-machines are not *á priori* required to be decisive, and such machines have polynomial acceptance (not running) time.

Relationships Among Feasible Time Stochastic Classes. Let us show some known inclusions among these classes and the connection to canonical classes such P and NP. The following are trivial relationships.

$$P \subseteq ZPP \subseteq RP \subseteq BPP \subseteq \mathbb{P} \subseteq PSPACE,$$

$$BPP \subseteq IP \cap \mathbb{P}.$$

We also have

$$RP \subsetneq_1 NP \subsetneq_2 \mathbb{P}.$$

The inclusion ($\subsetneq_2$) follows from equation (1) while inclusion ($\subsetneq_1$) follows from equation (2).

From the previous chapter (corollary 25), we have:

$$\mathbb{P} \subseteq IP \subseteq PSPACE.$$

However, we shall see that $IP = PSPACE$. Recall that $Primes \in$ co-RP.

We now show the following, attributed to Rabin:

$$ZPP = RP \cap \text{co-}RP. \tag{3}$$

Some preliminary remarks are in order: recall that $\oplus$-operator (see Section 7.1) satisfies De Morgan's law: $(1 - x) \oplus (1 - y) = 1 - (x \oplus y)$. This means that if M is a halting probabilistic machine, and we negate its local answers (interchanging YES and NO), then we obtain a machine $\overline{M}$ such that $L(M) \cap L(\overline{M}) = \emptyset$. If M is decisive, then $L(\overline{M}) = \text{co-}L(M)$. Moreover, if M has 1-sided error, then $\overline{M}$ has bounded zero-error acceptance.

One direction of (3) is clear: $ZPP \subseteq RP \cap \text{co-}RP$. Conversely, suppose $L \in RP \cap \text{co-}RP$. By the preceding remarks, L is accepted by bounded error polynomial-time probabilistic machines M and N where M has 1-sided error and N has bounded zero-error acceptance. We may assume that M and N are halting. We then construct a probabilistic machine that dovetails the computation of M and N in a step-for-step fashion until one of the following two events: (a) if N answers YES, we answer YES; (b) if M answers NO, we answer NO. If N answers NO or M answers YES, we simply continue simulation of the other machine. If both machines halt without events (a) or (b) occurring, we loop. This dovetail process essentially gives us a computation tree whose paths can be made to correspond in a bijective fashion with the set of all pairs of paths (p, p') where p comes from M and p' from N. Then it is not hard to see that the simulation runs in polynomial time with zero-error (Exercise).

Figure 8.2 summarizes the relationship of the main feasible-time classes based on stochastic choices with the canonical classes:

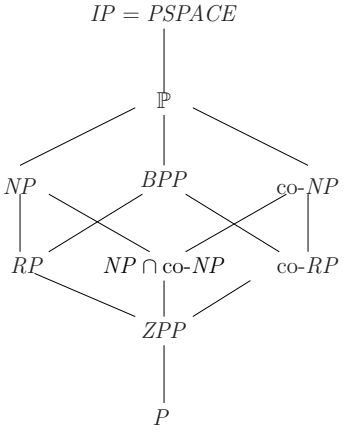

Figure 8.2: Feasible classes in the stochastic mode

Remark: The "error" terminology can be confusing but it is unfortunately well-established in the literature. We have attempted to systematize these concepts in the light of interval values. Our classification of errors into pessimistic versus optimistic may be helpful: thus, "error gap" and "bounded-error" *always* refers to pessimistic errors, and "zero-error" *always* refers to optimistic errors. Since pessimistic and optimistic errors are independent constraints on a computation, the two concepts can be distinguished in our terminology: the word "error" is always accompanied by key words such as "bounded", "gap" or "zero" which unambiguously indicate the form of error.

Most of the literature do not discuss these concepts in the generality given here, mostly focusing on polynomial-time classes. Indeed, while *BPP* is standard notation, there is no standard notation for classes such as $PrTIME_b(t(n))$.

EXERCISES

Exercise 0.1: Consider an alternative approach that distinguishes YO-answers from looping: assign the value $\frac{1}{2}$ to YO-configurations. (Thus, basis sets B for choice machines are required to the new constant function, $\frac{1}{2}$ in addition to the others. Also, the function $\frac{1}{2}$ adds nothing new if the toss-function $\textcircled{T}$ is already in B.) In the standard treatment of YO-configurations, it is not hard to see that a $\{\otimes, \oplus\}$-machine amounts to an alternating machine. In what sense does the alternative treatment of YO-configurations apparently increase the power of such machines? Is this apparent increase real? ♢

Exercise 0.2: Suppose a stochastic machine is not decisive. Prove that if it accepts in time bound $t(n)$ where t is time-constructible, then we can convert it into one accepting in time $O(t)$ with the minimal error gap $[\frac{1}{2}, \frac{1}{2}]$. Prove the same for space bounded acceptance.

♢

Exercise 0.3: Let $t(n)$ be any complexity function and $K_1 = PrTIME_1(t(n))$. Also let K_2 be the class of languages accepted by probabilistic choice machines with bounded error and zero-error acceptance, running in time $t(n)$. Under what conditions on t would we obtain $K_1 = \text{co-}K_2$? ♢

Exercise 0.4: Complete the arguments showing $ZPP = RP \cap \text{co-}RP$. ♢

Exercise 0.5: (Gill)
(i) Show that $\mathbb{P}$ and *BPP* are closed under complementation.
(ii) Show that *BPP* and *RP* are closed under union and intersection.

♢

Exercise 0.6: (Gill)

We consider probabilistic transducers. For any probabilistic transducer M and input w, we may talk of the probability that M on w produces x as output. Let this (least fixed point) probability be denoted $\Pr\{M(w) = x\}$. Let t_M be the partial transformation such that for all w, $t_M(w) = x$ if $\Pr\{M(w) = x\} > 1/2$; and $t_M(w) \uparrow$ if there is no such x. Clearly t_M is uniquely determined by M and is called the transformation computed by M. Show that if M is $s(n)$-space bounded then $x = t_M(w)$ implies $|x| \leq f(s(|w|))$ where $f(n)$ is the number of configurations of M using space at most n.

♢

END EXERCISES

8.2 How to Amplify Error Gaps

In this section, we seek techniques to convert a machine with error gap G into one with a strictly larger error gap G', $G \subset G'$. We now describe three such techniques, depending on the form of error and type of machine. The importance of such techniques is clear: if we can make such error gaps approach $(0, 1)$ (while keeping the running time feasible) then probabilistic algorithms may become indistinguishable from deterministic ones. For instance, a "theoretical" algorithm with a error probability of less than 2^{-1000} is much more reliable than any implemented algorithm will ever be, since implemented algorithms will contain many other (non-mathematical) sources of error such as the unreliability of hardware, not to speak of software.

(I) Halting Zero-error Rejection Machines. Let M be a halting probabilistic machine with zero-error rejection. Suppose M accepts[3] if with error gap $[0, b]$. We can boost the error gap very easily as follows. Fix some $k \geq 1$. Let N be the following machine:

> On any input w, simulate M on w for k times. This yields k answers, since M is halting. If any of the simulation answers YES, N answers YES. If all the k answers are NO or YO, N answers NO.

[3]By definition, $b \geq 1/2$. But in some sense, this technique works as long as $b > 0$.

If w is rejected by M, then N always answer NO (this is because N has zero-error rejection). If w is accepted by M, the probability of a YES-computation path of M is at least b. Hence the probability that N accepts is at least $1-(1-b)^k$. So N accepts with error gap

$$[0, 1-(1-b)^k]$$

which strictly includes $[0, b]$. Thus, *halting zero-error rejection implies bounded-error acceptance.* The technique will not work if M can have pessimistic errors in rejection.

(II) Bounded-error Probabilistic Machines. For bounded-error machines, we can use the "majority voting" scheme: repeat for an odd number of times an experiment with binary outcome; we take the majority outcome (*i.e.*, the outcome occurring more than half the time) as output. We justify this procedure with a lemma [29]:

LEMMA 1. *(a) Consider an experiment in which an event E occurs with probability*

$$p \geq \frac{1}{2} + e$$

for some $0 < e < \frac{1}{2}$. Then in $2t+1$ independent trials of the experiment, the probability that E is the majority outcome is greater than

$$1 - \frac{1}{2}(1-4e^2)^t.$$

(b) Similarly, if E occurs with probability

$$p \leq \frac{1}{2} - e$$

then the probability that E is the majority outcome is less than

$$\frac{1}{2}(1-4e^2)^t.$$

Proof. (a) Let $q = 1-p$ and $i = 0, \ldots, t$. Then the probability p_i that E occurs exactly i times out of $2t+1$ is given by the binomial distribution,

$$\begin{aligned}
p_i &= \binom{2t+1}{i} p^i q^{2t+1-i} \\
&= \binom{2t+1}{i} (\tfrac{1}{2}+e)^i (\tfrac{1}{2}-e)^{2t+1-i} \left[\frac{p}{\frac{1}{2}+e}\right]^i \left[\frac{q}{\frac{1}{2}-e}\right]^{2t+1-i} \\
&= \binom{2t+1}{i} (\tfrac{1}{2}+e)^i (\tfrac{1}{2}-e)^{2t+1-i} \left[\frac{pq}{(\frac{1}{2}+e)(\frac{1}{2}-e)}\right]^i \left[\frac{q}{\frac{1}{2}-e}\right]^{2t+1-2i} \\
&\overset{*}{\leq} \binom{2t+1}{i} (\tfrac{1}{2}+e)^i (\tfrac{1}{2}-e)^{2t+1-i} \\
&\leq \binom{2t+1}{i} (\tfrac{1}{2}+e)^i (\tfrac{1}{2}-e)^{2t+1-i} \left[\frac{\frac{1}{2}+e}{\frac{1}{2}-e}\right]^{t-i} \\
&= \binom{2t+1}{i} (\tfrac{1}{2}+e)^t (\tfrac{1}{2}-e)^{t+1} \\
&< \binom{2t+1}{i} (\tfrac{1}{4}-e^2)^t \frac{1}{2}.
\end{aligned}$$

Note that our derivation is careful to justify the transition ($\overset{*}{\leq}$) from p to $\frac{1}{2}+e$. Therefore the probability that E occurs in more than t trials is at least

$$\begin{aligned}
1 - \sum_{i=0}^{t} p_i &> 1 - \sum_{i=0}^{t} \binom{2t+1}{i} (\tfrac{1}{4}-e^2)^t \frac{1}{2} \\
&= 1 - 2^{2t} (\tfrac{1}{4}-e^2)^t \frac{1}{2} \\
&= 1 - \frac{1}{2}(1-4e^2)^t.
\end{aligned}$$

(b) Similarly, if $p \leq \frac{1}{2} - e$ then for $i \geq t+1$ we have

$$p_i \leq \binom{2t+1}{i}(\frac{1}{4} - e^2)^t \frac{1}{2}$$

and hence the probability that E occurs in more than t trials will be at most

$$\begin{aligned} \sum_{i=t+1}^{2t+1} p_i &\leq 2^{2t}(\frac{1}{4} - e^2)^t \frac{1}{2} \quad (4) \\ &= \frac{1}{2}(1-4e^2)^t. \quad (5) \end{aligned}$$

Q.E.D.

Using this, we can boost an error gap $G_e = [\frac{1}{2} - e, \frac{1}{2} + e]$ $(0 < \frac{1}{2} < e)$ to $[\frac{1}{2} - e', \frac{1}{2} + e']$ where

$$e' = \frac{1}{2}(1-4e^2)^t$$

if we do the majority vote for $2t+1$ trials. For instance, with $e = 1/4$ and $t = 8$, we have $e' = \frac{1}{2}(3/4)^t < 0.051$.

An **error gap function** G assigns an error gap $G(n)$ to each $n \in \mathbb{N}$. A machine M has error gap G if for all inputs w, $Val_M(w) \cap G(|w|) = \emptyset$. Let G_0 be the error gap function given by

$$G_0(n) = [2^{-n}, 1 - 2^{-n}].$$

We have the following useful lemma:

LEMMA 2. *Each language in BPP is accepted by a probabilistic acceptor that runs in polynomial time with error gap G_0.*

Proof. We may assume that the language is accepted by some M that runs in time n^d with error gap $G = [\frac{1}{2} - e, \frac{1}{2} + e]$ for some $d \geq 1$ and $0 < e < \frac{1}{2}$. Applying the lemma, we want to choose t satisfying

$$\begin{aligned} 2^{-n} &\geq \frac{(1-4e^2)^t}{2} \\ 2^{n-1} &\leq \frac{1}{(1-4e^2)^t} \\ n-1 &\leq t \log\left(\frac{1}{1-4e^2}\right) \\ t &\geq \frac{n-1}{\log(1/(1-4e^2))}. \end{aligned}$$

The desired machine N, on each computation path, simulates M for at most $2t+1 = O(n)$ times and outputs the majority outcome. Clearly N runs in time $O_e(n^{d+1})$ with error gap G_0. **Q.E.D.**

Let us give an application of this lemma:

THEOREM 3. (Ko, 1982) *If NP $\subseteq$ BPP then NP = RP.*

Proof. Since *RP* $\subseteq$ *NP*, it suffices to show inclusion in the other direction. It is easy to see that *RP* is closed under polynomial-time many-one reducibility, and hence we only have to show that the *NP*-complete language SAT belongs to *RP*. Suppose we want to check if a given CNF formula $F = F(x_1, \ldots, x_n)$ on n variables is satisfiable. For any sequence of Boolean values $b_1, \ldots, b_k$ $(k \leq n)$, let $F_{b_1 b_2 \cdots b_k}$ denote the formula F with x_i replaced by b_i, for $i = 1, \ldots, k$. We show how to construct a sequence $b_1, \ldots, b_n$ such that if F is satisfiable then $F_{b_1 \cdots b_n}$ is true with very high probability. By our assumption that *NP* $\subseteq$ *BPP*, there is a bounded-error probabilistic machine M accepting SAT in polynomial time. Moreover, by the preceding lemma, we may assume that M has error gap function G_0 and that M halts on every path in polynomial time.

We shall operate in n stages. At the start of stage k $(k = 1, \ldots, n)$, inductively assume that we have computed a sequence of Boolean values $b_1, \ldots, b_{k-1}$. It will be shown that $F_{b_1, \ldots, b_{k-1}}$ is probably satisfiable. In stage k, we compute b_k:

1.	Call M on input $F_{b_1 \cdots b_{k-1} 0}$.
2.	If M answers YES, then set $b_k = 0$ and go to DONE.
3.	Else call M on input $F_{b_1 \cdots b_{k-1} 1}$.
4.	If M answers NO again, we answer NO and return.
5.	Else set $b_k = 1$.
6.	DONE: If $k < n$ we go to stage $k+1$.
7.	Else answer YES if $F_{b_1,\ldots,b_n} = 1$, otherwise answer NO.

Let us analyze this procedure. It is clearly polynomial time.

If $k < n$, we either terminate in stage k with a NO answer, or we proceed to stage $k+1$. If $k = n$, we will surely terminate in stage k with answer YES or NO, and this answer is never in error. Thus our YES answers are never wrong. So if F is unsatisfiable, we answer NO on every path. Thus we have zero-error rejection.

Finally, let us prove that if F is satisfiable, then our procedure answer YES with probability $> 1/2$. Write F_k for $F_{b_1,\ldots,b_k}$, assuming that $b_1, \ldots, b_k$ are defined. Let the event A_k correspond to " no mistakes up to stage k", *i.e.*, F_k is defined and satisfiable. Similarly, let event E_k correspond to "first mistake at stage k", *i.e.*, $E_k = A_{k-1} \cap \overline{A_k}$.

CLAIM: $\Pr(E_k) \leq 2 \cdot 2^{-|F|+1})$.

Proof: Note that $\Pr(E_k) \leq \Pr(E_k | A_{k-1})$. We will bound $\Pr(E_k | A_{k-1})$. Assuming A_{k-1}, we consider 2 cases: (A) CASE $F_{b_1 \cdots b_{k-1} 0}$ is not satisfiable. Then $F_{b_1 \cdots b_{k-1} 1}$ is satisfiable. With probability $\geq (1 - 2^{-|F|})$, M will (correctly) answer NO the first time we invoke M. Then with probability $\geq (1 - 2^{-|F|})$, M will (correctly) answer YES the second time. So $\Pr(A_k | A_{k-1}) \geq (1 - 2^{-|F|})^2$ and

$$\Pr(E_k | A_{k-1}) \leq 1 - (1 - 2^{-|F|})^2 \leq 2^{-|F|+1}).$$

(B) CASE $F_{b_1 \cdots b_{k-1} 0}$ is satisfiable. This case is even easier, and yields $\Pr(E_k | A_{k-1}) \leq 2^{-|F|}$. This proves the claim.

To conclude the theorem, the probability of making mistake at any stage is at most

$$\sum_{k=1}^{n} \Pr(E_k) \leq 2n \cdot 2^{-|F|} \leq 2n \cdot 2^{-2n}.$$

This is less than $1/2$ for n large enough. Hence F will be accepted. **Q.E.D.**

See Exercise for another proof.

(III) Stochastic machines. We now introduce a third technique that is applicable to stochastic machines. Motivated by a paper of Valiant, we introduce the following probability functions:

- $P(x) := x \otimes x = x^2$
- $Q(x) := x \oplus x = x(2 - x)$.
- $A(x) := Q(P(x)) \oplus P(Q(x))$.

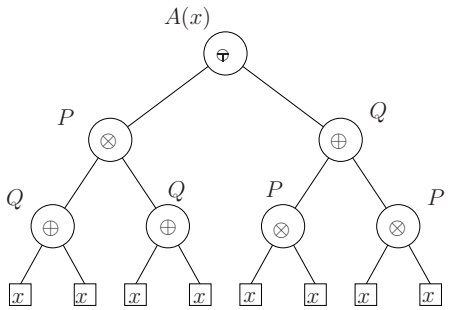

Figure 8.3: The operator $A(x)$

Thus,

$$A(x) = \frac{x^2(2 - x^2) + x^2(2 - x)^2}{2} = x^2(3 - 2x).$$

These operators are extended to *INT* in the usual way: thus, $A([u, v]) = [A(u), A(v)]$. The exercises show other properties of these functions. We show that $A(x)$ has the following "amplification property":

LEMMA 4. *If* $0 \leq e \leq \frac{1}{2}$, *then*

$$x \notin [\frac{1}{2} - e, \frac{1}{2} + e] \quad \Rightarrow \quad A(x) \notin [\frac{1}{2} - e', \frac{1}{2} + e']$$

where

$$e' = e(\frac{3}{2} - 2e^2) \geq e.$$

Proof. Since $\frac{dA}{dx} = 6x(1-x)$, $A(x)$ is monotone increasing for $0 \leq x \leq 1$. The lemma then follows from a calculation,

$$\begin{aligned} A(\frac{1}{2} + e) &= \frac{1}{2} + e[\frac{3}{2} - 2e^2] \\ A(\frac{1}{2} - e) &= \frac{1}{2} - e[\frac{3}{2} - 2e^2]. \end{aligned}$$

Q.E.D.

Note that the error gap has an "amplification factor" $\alpha(e) = \frac{3}{2} - 2e^2$ which decreases monotonically from 3/2 to 1 as e increases from 0 to 1/2. Note that

$$\alpha(1/4) = \frac{11}{8}.$$

So if the error bound e for a value x is less than $1/4$, then the error bound for $A(x)$ is at least $11e/8$. We conclude:

THEOREM 5.
Let $t(n)$ *be a time-constructible function. Then*

$$PrTIME(t) \subseteq StTIME_b(O(t)).$$

In particular,

$$\mathbb{P} \subseteq StTIME_b(n^{O(1)}).$$

Proof. Suppose M is a probabilistic machine that decides its inputs in running time t. Since t is time-constructible, we can modify it always halt in time $O(t)$ and have error gap

$$G(n) = [\frac{1}{2} - 2^{-t}, \frac{1}{2} + 2^{-t}].$$

Let $s = \frac{t}{\log(11/8)}$. We construct a stochastic machine N that, on input w, operates by first "computing the iterated function $A^s(x)$" (the meaning should be clear) and then simulating M on w. One checks that N has gap at least $[\frac{1}{4}, \frac{3}{4}]$. Note that N still runs in time $O(t)$. **Q.E.D.**

EXERCISES

Exercise 0.7: This is a slight variation on boosting error gaps for machines with zero-error rejection. Let M be a machine with error gap $G = (0, b]$. We construct the following machine N to boost the error gap G:

> 1. Simulate M on input from beginning until it halts.
> (If M loops then we loop)
> 2. If M answers YES then answer YES; else toss a coin.
> 3. If coin-toss is heads then answer NO; else go to 1.

(i) Show that N has error gap $G' = (0, b']$ where $b' = \frac{2b}{1+b}$.
(ii) For any $\epsilon > 0$, show how to modify step 3 so that $1 - b' \leq \epsilon$.

♢

Exercise 0.8: Let M be a bounded-error halting probabilistic machine. We construct a probabilistic machine N that keeps simulating M on its input until the "running score" reaches $+2$ or -2. The running score (at any moment) is defined to be the number of YES answers minus the number of NO answers up to that moment. If we reach $+2$, we answer YES, and if we reach -2, we answer NO.
(i) What is the language accepted by N?
(ii) Analyze the complexity and errors of N.

♢

Exercise 0.9: Give another proof of Ko's theorem that $NP \subseteq BPP$ implies $NP = RP$ in §2. Recall the event A_k in the first proof.
(i) Show that $\Pr(A_n) \geq (1 - 2^{-2n})^{2n}$.
(ii) Show that $(1 - 2^{-n})^n \geq \frac{1}{2}$ for large enough n. HINT: you may use the following facts:

- $e^t \geq 1 + t$ for all $t \in \mathbb{R}$ with equality iff $t = 0$.
- $\left(1 + \frac{t}{n}\right)^n \geq e^t \left(1 - \frac{t^2}{n}\right)$ for all $t, n \in \mathbb{R}$, $n \geq 1$ and $|t| \leq n$.

♢

Exercise 0.10:
(i) Redraw the class inclusion diagram in Figure 8.2 assuming that $NP \subseteq BPP$.
(ii) What consequences for the diagram of Figure 8.2 can you draw if we assume that $BPP \subseteq NP$?

♢

Exercise 0.11: Let $g = [1/3 - e, 1/3 + e]$ for some $0 < e < 1/3$. Give an analog of majority voting to amplify this gap.

♢

Exercise 0.12: Let $P(x)$ and $Q(x)$ be as defined for boosting error gaps in stochastic machines. For all $n = 0, 1, 2, \ldots$, let

$$P_n(x) := \begin{cases} x & \text{if} \quad n = 0 \\ P(P_{n-1}(x)) & \text{if} \quad n = odd \\ Q(P_{n-1}(x)) & \text{if} \quad n = even \end{cases}$$

$$Q_n(x) := \begin{cases} x & \text{if} \quad n = 0 \\ Q(Q_{n-1}(x)) & \text{if} \quad n = odd \\ P(Q_{n-1}(x)) & \text{if} \quad n = even \end{cases}$$

For example, $P_2(x) = Q(P(x)) = x^2(2 - x^2)$, and $Q_2(x) = P(Q(x)) = x^2(2 - x)^2$. The function $Q_2(x)$ was used by Valiant to give a non-constructive proof that the majority function for Boolean functions has a monotone formula of size $O(n \log n)$. The amplification function in the text is just $A(x) = P_2(x) \oplus Q_2(x)$. Now write $p_n^+(e)$ for $P_n(\frac{1}{2} + e)$, $p_n^-(e)$ for $P_n(\frac{1}{2} - e)$, and similarly for $q_n^+(e), q_n^-(e)$ relative to Q_n. For example,

$$\begin{aligned}
p_1^+(e) &= \frac{1}{4} + e + e^2 \\
q_1^+(e) &= \frac{3}{4} + e - e^2 \\
p_1^-(e) &= \frac{1}{4} - e + e^2 \\
q_1^-(e) &= \frac{3}{4} - e - e^2 \\
p_2^+(e) &= \frac{7}{16} + \frac{3e}{2} + \frac{e^2}{2} - 2e^3 - e^4 \\
q_2^+(e) &= \frac{9}{16} + \frac{3e}{2} - \frac{e^2}{2} - 2e^3 + e^4 \\
p_2^-(e) &= \frac{7}{16} - \frac{3e}{2} + \frac{e^2}{2} + 2e^3 - e^4 \\
q_2^-(e) &= \frac{9}{16} - \frac{3e}{2} - \frac{e^2}{2} + 2e^3 + e^4 \\
p_3^+(e) &= \frac{49}{256} + \frac{21e}{16} + \frac{43e^2}{16} - \frac{e^3}{4} - \frac{53e^4}{8} - 5e^5 + 3e^6 + 4e^7 + e^8 \\
q_3^+(e) &= \frac{207}{256} + \frac{21e}{16} - \frac{43e^2}{16} - \frac{e^3}{4} + \frac{53e^4}{8} - 5e^5 - 3e^6 + 4e^7 - e^8
\end{aligned}$$

Show
(i) $p_n^+(e) + q_n^-(e) = 1$

(ii) $p_n^-(e) + q_n^+(e) = 1$
(iii) $p_n^+(e) - p_n^-(e) = q_n^+(e) - q_n^-(e)$
(iv) $x = \frac{1}{2}$ is a fixed point of $A_n(x) = P_n(x) \oplus Q_n(x)$, *i.e.*, $A_n(\frac{1}{2}) = \frac{1}{2}$. Are there other fixed points?
(v) The exact relationships between the coefficients of $p_n^+(e), q_n^+(e), p_n^-(e)$ and $q_n^-(e)$.

♢

Exercise 0.13: Show that $RP = NP$ iff some NP-complete language is in RP. ♢

Exercise 0.14: Another proof that $NP \subseteq BPP$ implies $NP = RP$. Recall the event A_k in the first proof.
(i) Show that $\Pr(A_n) \geq (1 - 2^{-2n})^{2n}$.
(ii) Show that $(1 - 2^{-n})^n \geq \frac{1}{2}$ for large enough n. HINT: you may use the following facts:

- $e^t \geq 1 + t$ for all $t \in \mathbb{R}$ with equality iff $t = 0$.
- $\left(1 + \frac{t}{n}\right)^n \geq e^t \left(1 - \frac{t^2}{n}\right)$ for all $t, n \in \mathbb{R}$, $n \geq 1$ and $|t| \leq n$.

♢

End Exercises

8.3 Probabilistic Feasible Time

Feasible time simply means polynomial time, of course. Here we consider two versions of feasible time, depending on the error concept: unbounded error (the class $\mathbb{P}$) and bounded error (the class BPP).

8.3.1 Unbounded Error Polynomial Time

We derive some basic properties of the class $\mathbb{P}$.

Lemma 6. *For any $\mathbb{P}$-machine M, there exists another $\mathbb{P}$-machine N such that $L(N) = L(M)$ where N is halting, decisive, and always answers YES or NO.*

Proof. (Sketch) The machine N simulates M. To make N halting, we clock the machine M for a polynomial number of steps and answer NO if M does not answer YES by the time the clock runs out. This means an YO answer is also turned into a NO answer. To ensure that N is decisive, we "shift" the gap $(\frac{1}{2}, \frac{1}{2} + e)$ to some error gap $(\frac{1}{2} - e', \frac{1}{2} + e')$. Here $e = e(n)$ depends on the length n of the input, and can be assumed to be a binary rational which is polynomial time computable. **Q.E.D.**

The class $\mathbb{P}$ satisfies the following remarkable closure property:

Lemma 7. *If $A, B \in \mathbb{P}$ then the symmetric difference $A \oplus B \in \mathbb{P}$.*

Proof. Let M_A, M_B be $\mathbb{P}$-machines that accept A and B (resp.). By the previous lemma, assume M_A and M_B are halting and decisive and always answer YES or NO. Consider the machine N that, on any input x, simulates M_A on x and then M_B on x. N will answer YES if exactly one of M_A and M_B answers YES, otherwise N answers NO. We note that

- If both M_A and M_B answered correctly (answers YES iff the global answer is ACCEPT), then N's answer is correct.
- If both M_A and M_B answered incorrectly then N's answer is (still) correct.
- If exactly one of M_A or M_B answered incorrectly, then N's answer is incorrect.

Assume that the probability of M_A (resp., M_B) being correct on any input x is *exactly* $(1/2) + \varepsilon_A$ (resp., $(1/2) + \varepsilon_B$). Notice that $\varepsilon_A > 0$ and $\varepsilon_B > 0$ exists since the machines are decisive. Moreover ε_A and ε_B does not depend on x. Then the probability of N being correct on input x is exactly

$$(\frac{1}{2} + \varepsilon_A)(\frac{1}{2} + \varepsilon_B) + (\frac{1}{2} - \varepsilon_A)(\frac{1}{2} - \varepsilon_B) = \frac{1}{2} + 2\varepsilon_A\varepsilon_B.$$

Q.E.D.

An alternative argument is as follows: if probability of M_A (M_B) answering YES is a (b), then the probability of N answering YES is $c = a(1-b)+(1-a)b$. Note that we can add $a(1-b)$ and $(1-a)b$ because these correspond to disjoint events. We see that

$$\begin{aligned}\frac{1}{2} - c &= \frac{1}{2}\left(\frac{1}{4} - \frac{a+b}{2} + ab\right) \\ &= \frac{1}{2}(\frac{1}{2} - a)(\frac{1}{2} - b).\end{aligned}$$

Since the machines are decisive, $a \neq 1/2$ and $b \neq 1/2$, and so $1/2 - c \neq 0$. Moreover, $1/2 - c < 0$ iff exactly one of the inequalities, $a < 1/2$ and $b < 1/2$, hold. This is exactly what we want to show.

The following is left as an exercise.

LEMMA 8. *$\mathbb{P}$ is closed under complementation.*

The next result is from Beigel, Reingold and Spielman [3].

THEOREM 9. *$\mathbb{P}$ is closed under intersection.*

The proof uses an interesting property of polynomials, omitted here.

We now investigate complete languages for $\mathbb{P}$. It should be noted that among the stochastic feasible classes $(ZPP, RP, BPP, \mathbb{P})$, only $\mathbb{P}$ has known complete languages. For any Boolean formula F, let $\#(F)$ denote the number of satisfying assignments to variables that occur in F. Consider three related languages (we include the standard SAT in this table for comparision):

$$\begin{aligned}\text{MAJ} &:= \{\langle F\rangle : \#(F) > 2^{m-1} \text{ where } m \text{ is the number of variables in } F\} \\ \#\text{SAT} &:= \{\langle F, k\rangle : F \text{ is 3CNF and } \#(F) \geq k\} \\ \#\text{SAT}_0 &:= \{\langle F, k\rangle : F \text{ is 3CNF and } \#(F) = k\} \\ \text{SAT} &:= \{\langle F\rangle : F \text{ is 3CNF and } \#(F) \geq 1\}\end{aligned}$$

Here $\langle F\rangle$ and $\langle F, k\rangle$ denote any reasonable encoding of F and (F, k) as binary strings.

LEMMA 10. MAJ *and* #SAT *are both $\mathbb{P}$-hard under Karp reducibility.*

Proof. (a) It is easy to see that MAJ $\in \mathbb{P}$: construct a $\mathbb{P}$ machine which accepts a Boolean formula F with probability equal to $\#(F)/2^m$ (if F has m variables. Thus the machine accepts precisely the formulas of MAJ.
(b) We show that MAJ is $\mathbb{P}$-hard. If M is a $\mathbb{P}$-machine and x is an input, we want a 3CNF formula F such that $\#(F)$ is equal to the number of YES-paths in the computation tree for x. But we note that the proof of the *NP*-hardness of SAT (Cook's Theorem in Chapter 3) produces such a formula. See Exercise.
(c) Clearly #SAT is $\mathbb{P}$-hard since we can reduce MAJ to #SAT.
(d) Finally, to show #SAT $\in \mathbb{P}$ it is sufficient to construct a transformation $t : \langle F, k\rangle \mapsto t(F, k)$ such that $\langle F, k\rangle \in$ #SAT iff $t(F, k) \in$ MAJ. Suppose F has $m+1$ free variables. If $k = 0$, then clearly $\langle F, k\rangle \in \#\text{SAT}_1$. Hence we may assume $1 \leq k \leq 2^{m+1}$. We construct a CNF forumula C_k with these $m+1$ variables of F such that exactly $2^{m+1} + 1 - k$ assignments satisfy C_k. Let z be a new variable. Then

$$t(\langle F, k\rangle) = (z \wedge F) \vee (\overline{z} \wedge C_k).$$

It is easy to see that $\langle F, k\rangle \in$ #SAT iff $t(\langle F, k\rangle) \in SAT_0$. The result is not quite CNF yet. Assuming F and C_k are CNF, then $z \wedge F$ and $\overline{z} \wedge C_k$ are each CNF. In general, to create a CNF formula G from the disjunction $G_1 \vee G_1$ of two CNF formulas G_1, G_2, we let each clause of G be written as the disjunct of a clause of G_1 with a clause of G_2. Of course, G need not be 3-CNF (it is a bit more work to achieve this, while ensuring that $G \in SAT_0$).

If $k = 1$, then C_k can be taken to be 1 (always true). Otherwise $k \geq 2$ and we consider the binary expansion of $2^{m+1} + 1 - k = \sum_{j=0}^{m} b_j 2^j$ ($b_j = 0$ or 1). We leave the details to the reader. **Q.E.D.**

The above techniques do not work with $\#\text{SAT}_0$, and we also do not know any complete languages for any error-limited complexity classes such as *RP* or *BPP*.

EXERCISES

Exercise 0.15: Prove lemma 8. ♢

Exercise 0.16: The proof of Cook's theorem in Chapter 3 reduces any language $A \in NP$ to SAT. If M is an NP-machine for A, for any input w, the proof constructs a 3CNF formula F_w such that F_w is satisfiable iff $w \in A$. We claim that something stronger is true: $\#F_w$ is equal to the number of accepting computations of M on input w. Actually, this claim needs a mild condition on M. What is it? Prove this claim under this mild condition. **Hint:** if you do not see what this condition might be, we suggest the strategy of ignoring it at first, and trying to prove the claim unconditionally. ♢

Exercise 0.17*: Does ZPP, RP or BPP have complete languages under any reasonable reducibilities? □

End Exercises

8.3.2 Bounded Error Polynomial Time

See Balcazar' book.

This should go to Constructive Hierarchies Chapter:
(1) BPP is in $\Sigma_2 \cap \Pi_2$.
(I think the current results here are not the most general... see the proof)
(2) Approximate Counting
(3) $NTIME(n) \neq DTIME(n)$.
(4) BPP has polynomial size circuits
We use the following fact:

Lemma 11. *Let M be a choice machine with only binary choices, $L(M) \subseteq \mathbb{B}^*$, and M runs in time $t(n)$. Then there is a polynomial $p(n)$ such that for all n, there is a Boolean circuit C_n on $n + t(n)$ inputs such that*
(1) C_n has size $p(t(n))$, and
(2) for all $x \in \mathbb{B}^n$ and all $y \in \mathbb{B}^{t(n)}$, $C_n(x, y) = 1$ iff M on input x and making the choices specified by y leads to a YES answer.

Proof. Sketch proof: There is a fixed size circuit that computes the "local transition" function, as given in Papadimitriou's book (theorem 8.1, p.168). Note: this function does not depend on M having only one worktape. **Q.E.D.**

Theorem 12. *Every binary language in BPP has a polynomial size circuit.*

Proof. [As in Papadimitriou] Let $L \subseteq \mathbb{B}^*$ be accepted by a BPP-machine M. We will construct for each n, a Boolean circuit C_n^* of polynomial size such that $x \in L$ iff $C_n^*(x) = 1$.

We may assume that for any input x, $\Pr\{M(x) = error\} \leq 1/4$. As in the previous lemma, assume M runs in time $t(n)$ and there are circuits C_n of size $p(t(n))$ which simulates M on any input $x \in \mathbb{B}^n$ and any path choice $y \in \mathbb{B}^{t(n)}$. Let $T_n = \mathbb{B}^{12(n+1)t(n)}$. We consider a string $Y \in T_n$ as a $(12(n+1)$-tuple $(y_1, \ldots, y_{12(n+1)}$ of choices for a computation on an input of length n. We say that Y is "good" for $x \in \mathbb{B}^n$ if more than half of the $\{C_n(x, y_i) : i = 1, \ldots, 12(n+1)\}$ is correct. Here, $C_n(x, y)$ is correct if $C_n(x, y) = 1$ iff $x \in L$. For any x, and for a random choice of Y, the probability that $y \in \mathbb{B}^{t(n)}$ is incorrect for x is $\leq 1/4$. Hence the probability that $Y \in T_n$ is not good for x is, by Chernoff's bound at most $e^{-(n+1)}$.

[See Papadimitriou, or my lecture notes]

Now the probability that Y is bad for more any $x \in \mathbb{B}^n$ is therefore at more $2^n e^{-(n+1)} < 1/2$. Hence, there exists $Y^* \in T_n$ that is good for all $x \in B_n$. We now construct a circuit $C_n^*(x)$ to be simply $12(n+1)$ copies of C_n in which each C_n has the form $C_n(x, y_i^*)$ where $Y^* = (y_1^*, \ldots, y_{12(n+1)}^*)$. Finally, the output of C_n^* is just the majority output of all the C_n's. **Q.E.D.**

Note: a random construction of C_n^* will succeed with probability more than $1/2$. So there is a BPP algorithm for this construction? Not yet: you need to check if a given choice of Y is good for all x.

8.4 Average Time for Probabilistic Machines

So far, we have not discussed the very natural concept of "average time" for probabilistic machines. We systematically develop the concept as follows (the appendix contains the probabilistic terms used here). Let M be a probabilistic machine and $T_M(w)$ denotes the usual complete computation tree on input w. Without loss of generality, assume $T_M(w)$ is a full binary tree, *i.e.*, a binary tree in which every internal node has two children. Let T any prefix of $T_M(w)$.

We construct associate with T a probability space

$$(\Omega_T, \Sigma_T, \Pr_T)$$

in which the sample space Ω_T comprises all complete paths of T. A subset of Ω_T is a **basic set** if it is the collection of complete paths of T all sharing some common initial prefix π; denote this set by $B_T(\pi)$. In particular, Ω_T is the basic set $B_T(\epsilon)$ where $\pi = \epsilon$ is the empty path. Any singleton set consisting of a complete path is also a basic set. Let Σ_T^0 comprise all finite union and complement of basic sets: clearly Σ_T^0 forms a field. Let Σ_T be the Borel field generated by Σ_T^0. The probability measure $\Pr_T$ assigns to each basic set $B_T(\pi)$ the probability $\Pr_T(B_T(\pi)) = 2^{-|\pi|}$ where $|\pi|$ is the length of the path π. E.g., $\Pr_T(\Omega_T) = 2^0 = 1$, as expected. Notice that every element of Σ_T^0 is a finite union of disjoint basic sets. Extend the definition of $\Pr_T$ to sets in Σ_T^0 so as to preserve finite additivity of pairwise disjoint unions. One checks that the extension does not depend on how we partition sets in Σ_T^0 into a countable union of basic sets. Finally, a theorem of Carathéodory (appendix) gives us a unique extension of $\Pr_T$ to all of Σ_T.

Definition 4. (i) We introduce three random variables for Ω_T. For any complete paths of T,

$$\begin{aligned} \mathrm{Time}_T(\pi) &= |\pi| \text{ possibly infinite,} \\ \mathrm{Accept}_T(\pi) &= 1 \text{ iff } \pi \text{ ends in a YES-node in } T, \\ \mathrm{Halt}_T(\pi) &= 1 \text{ iff } \pi \text{ is a finite path.} \end{aligned}$$

When T is the complete computation tree for input w, we write

$$\mathrm{Time}_w, \mathrm{Accept}_w, \mathrm{Halt}_w, \Omega_w, etc,$$

for Time_T, Accept_T, etc.

(ii) The **average time** of T is the expected value of Time_T. A machine M **accepts/rejects in average time** $t(n)$ if for all accepted/rejects inputs w of length n,

$$\Pr\{\mathrm{Time}_w \leq t(n), \mathrm{Accept}_w = 1\} > \frac{1}{2}.$$

It **runs in average time** $t(n)$ if it is decisive and accepts and also rejects in average time $t(n)$.

■

An equivalent definition of average time is this: the average time of a computation tree T is equal to the sum over the weights of each edge in T where an edge from level $\ell - 1$ to level ℓ (the root is level 0) has weight $2^{-\ell}$. Naturally, **the probability that** M **halts on** w is the expected value of the random variable Halt_w. If M is halting (*i.e.*, there are no infinite computation paths) then it halts with probability 1; the converse is false. When we use running complexity for probabilistic machines, they are sometimes designated 'Monte Carlo algorithms'; when average complexity is used, they are designated 'Las Vegas algorithms'.

We first note an observation of Gill: *Every recursively enumerable language can be accepted by a probabilistic machine with constant average time.* In proof, suppose M is a deterministic Turing machine. Without loss of generality assume that M accepts whenever it halts. We construct N to simulate M as follows:

```
repeat
    Simulate one step of M;
    if the simulated step halts, we answer YES or NO following M;
until head = cointoss();
if head = cointoss() then answer YES else answer NO.
```

Here cointoss() is a random function that returns *head* or *tail* and there are two separate invocations of this function above. We note that if the input word is not accepted by M then N can only reject (since the probability of N saying YES is equal to the probability of saying NO). Hence N has zero-error rejection. If an input word is accepted by M, then we see that the probability of its acceptance by N is more than $\frac{1}{2}$ since each NO path can be uniquely paired with an YES path of the same length, but there is one YES path that is not paired with any NO path. These remarks are easy to see once we unroll the computation tree of N as shown in Figure 8.4.

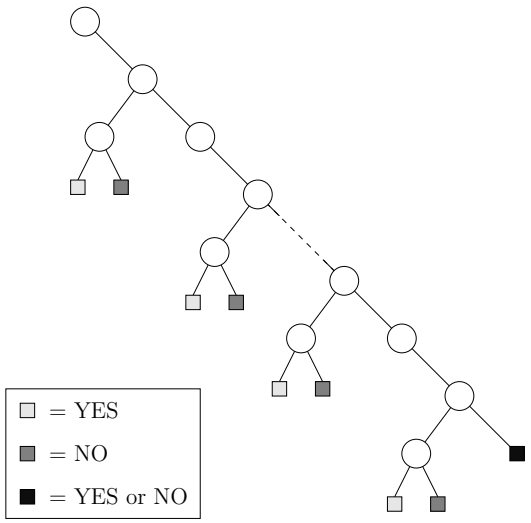

Figure 8.4: Computation Tree of N, simulating a deterministic Turing acceptor M.

The average time $\bar{t}$ spent in the repeat-loop satisfy the inequality

$$\bar{t} \leq 2 + \frac{\bar{t}}{2}$$

where the term '2+' comes from simulating a step of M and tossing a coin to decide on continuing inside the loop (it takes no time to decide to say YES if M says YES). Thus $\bar{t} \leq 4$. The average time of N is at most $1 + \bar{t} \leq 5$ (where the '1' for the final cointoss).

Gill notes that such pathological behaviour does not happen with bounded-error machines:

LEMMA 13. *Let M be a probabilistic machine accepting/rejecting with bounded-error. There is a constant $c > 0$ such that if M accepts/rejects in average time $\bar{t}(n)$ and accepts/rejects in time $t(n)$ then*

$$\bar{t}(n) \geq \frac{t(n)}{c}.$$

Proof. Suppose M accepts with probability at least $\frac{1}{2} + e$ for some $0 < e < \frac{1}{2}$. (The proof if M rejects with probability at most $\frac{1}{2} - e$ is similar.) Fix any input of length n and let $\bar{t} = \bar{t}(n)$. If $\bar{t} = \infty$ there is nothing to prove. Otherwise, let T be the complete computation tree. Since Time_T is non-negative, Markov's inequality yields

$$\Pr\{\text{Time}_T \geq c\bar{t}\} \leq \frac{1}{c}$$

for any $c > 0$. Choosing $c = \frac{2}{e}$,

$$\begin{aligned} \Pr\{\text{Time}_T < c\bar{t}, \text{Accept}_T = 1\} &\geq \Pr\{\text{Time}_T < c\bar{t}\} - \Pr\{\text{Accept}_T = 0\} \\ &\geq (1 - \frac{1}{c}) - (\frac{1}{2} - e) \\ &\geq \frac{1}{2} + \frac{e}{2}. \end{aligned}$$

This proves that T, truncated below $c\bar{t}$, accepts with bounded error. **Q.E.D.**

In view of this lemma, let

$$AvgTIME(t(n))$$

denote the class of languages accepted by probabilistic machines M where M has bounded-error and M runs in average time $t(n)$. Note that both properties here are independently defined for the entire computation tree. We have thus proved:

COROLLARY 14. *For any $t(n)$,*

$$AvgTIME(t) \subseteq PrTIME_b(O(t)). \tag{6}$$

As further example, we provide deterministic time upper bounds for languages accepted in average time $t(n)$ with bounded-error.

COROLLARY 15. *If a bounded-error probabilistic machine M accepts with average time $\bar{t}(n)$ then $L(M) \in DTIME(O(1)^{\bar{t}(n)})$.*

Proof. We can simulate M by computing the least fixed point of a computation tree of depth $O(\bar{t}(n))$. **Q.E.D.**

LEMMA 16. *Let $s(n) \geq \log n$ be space constructible. Let M be any nondeterministic machine that accepts in space s. Then there is a probabilistic machine N with zero-error that accepts $L(M)$ in space s.*

Proof. Choose $c > 0$ such that there are at most $c^{s(n)}$ configurations using space at most $s(n)$. Fix any input of length n and let $s = s(n)$. First we mark out exactly s cells. The probabilistic machine N proceeds as follows:

repeat forever
1. Initialize M to its initial configuration.
2. Simulate M for c^s steps. Nondeterministic choices of M become coin-tossing choices of N.
3. If M answers YES in this simulation, we answer YES. (If M answers NO or YO or does not halt in c^s steps, then we go to back to 1.)

end

Clearly N loops if M does not accept. If M accepts then the probability of N answering YES is easily seen to be 1. **Q.E.D.**

This lemma implies that probabilistic space-bounds with zero-error is as powerful as nondeterministic space:

THEOREM 17. *For any space-constructible $s(n) \geq \log n$,*

$$NSPACE(s) = PrSPACE_0(s) = PrSPACE_1(s).$$

Proof. The above lemma shows that $NSPACE(s) \subseteq PrSPACE_0(s)$. The converse is easy since for any probabilistic machine M with zero error, when viewed as a nondeterministic machine N, accepts the same language with the same space bound. We check that the same construction applied to probabilistic one-sided error machines in place of nondeterministic machines show $PrSPACE_1(s) \subseteq PrSPACE_0(s)$, and hence they are equal. **Q.E.D.**

This result can be generalized to log-space alternating machines, but we now have two-sided error [28].

The simulation in the above proofs can be modified so that the simulating machine N halts with probability 1. However, N is no longer zero-error. The technique will be introduced in section 6.

Probabilistic Simulating of Alternation. Consider how a probabilistic machine can simulate an alternating machine M. We want our probabilistic machine to have bounded error. Suppose C is a configuration of M and $C \vdash (A, B)$. Let T_C, T_A, T_C denote the subtree at these nodes.

Inductively, assume that our recursive construction gives us probabilistic computation trees $T_{A'}$ and $T_{B'}$ (rooted at A' and B') which emulates T_A and T_B (respectively) with error gap

$$g_0 = [1/4, 3/4].$$

This means that if A accepts, then the value of A' is at least $3/4$ and if A rejects, the value of A' is at most $1/4$. Similarly for B and B'. Let us see how to carry the induction through.

CASE: C is a $\wedge$-configuration. Let C_0 be a $\oplus$-configuration such that $C_0 \vdash (A', B')$ (see figure 8.5). Then the value at C_0 has an error gap of

$$[5/8, 3/4].$$

This is because, if at least one of A' or B' rejects, then value of C_0 is at most $(1/4) \oplus 1 = 5/8$. And if both A' and B' accepts, the value is at least $3/4$. Then, we 'shift' the gap so that it is centered, by averaging it with the value $5/16$. This gives us a new gap (of half the size!)

$$g_0 = [15/32, 17/32].$$

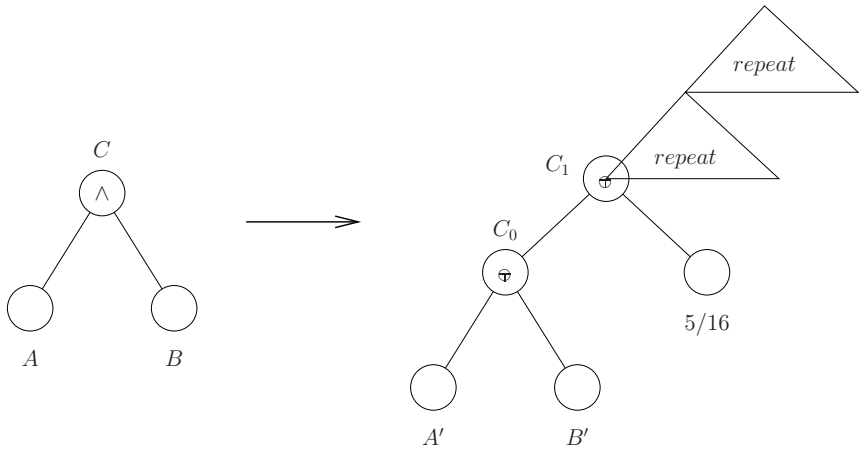

Figure 8.5: Simulating a $\wedge$-configuration C.

We now use majority voting to boost this gap back to at least $[1/4, 3/4]$. We need to take the majority of $2k+1$ votes, where k is a constant that can be computed.

CASE: C is a $\vee$-configuration. In this case, C_0 has error gap of

$$[1/4, 3/8].$$

Now we shift this gap by averaging it with $11/16$, yielding the gap g_0 above. We again boost it back to $[1/4, 3/4]$.

Note that if the original tree has height h and this procedure produces a tree of height $\tau(h)$, then

$$\tau(h) = k(\tau(h-1) + 2) \leq (2k)^h.$$

Of course, we need to make this recursive transformation something that can be carried out by a suitable probabilistic machine. This is left as an exercise. We have thus shown:

THEOREM 18.

$$ATIME(t) \subseteq PrTIME_b(2^{O(t)}).$$

EXERCISES

Exercise 0.18: Let $t(n)$ be time-constructible. Determine the smallest function $t'(n)$ as a function of $t(n)$ such that

$$\text{co-}NTIME(t(n)) \subseteq PrTIME_b(t'(n)).$$

Hint: Simulate each step of a universal machine, but at each step, ensure bounded-error by using majority votes. ♢

Exercise 0.19: Let M be a probabilistic machine that runs in time $t(n)$ and which uses $\leq \log t(n)$ coin tosses along each computation path. Give an upper bound for the function $t'(n)$ such that $L(M) \in PrTIME_b(t'(n))$. ♢

END EXERCISES

8.5 Interactive Proofs

This section introduces interactive proofs [13] and Arthur-Merlin games [2]. The class of languages recognized in polynomial time by such machines is denoted IP. We show that IP contains two languages NONISO and #SAT. This gives some indication of the power of IP because NONISO is not known to be in NP, and #SAT is complete for $\mathbb{P}$. The main result is $IP = PSPACE$, which provides yet another characterizaiton of $PSPACE$.

Graph Non-Isomorphism. A motivating example is the graph isomorphism problem: given a pair $\langle G_0, G_1\rangle$ of graphs, decide if they are isomorphic,

$$G_0 \sim G_1. \tag{7}$$

We could use digraphs or bigraphs in the following. To be specific assume bigraphs (undirected graphs). Let $\mathcal{G}_n$ denote the set of bigraphs on the vertex set $\{1, 2, \ldots, n\}$. Formally define the languages

$$\begin{aligned} \text{ISO} &= \{\langle G_0, G_1\rangle \in \mathcal{G}_n^2 : n \geq 1, G_0 \sim G_1\}, \\ \text{NONISO} &= \{\langle G_0, G_1\rangle \in \mathcal{G}^2 : n \geq 1, G_0 \not\sim G_1\}. \end{aligned}$$

These are basically complementary languages. Let S_n denote the set of n-permutation (*i.e.*, permutations of $\{1, \ldots, n\}$). For $\pi \in S_n$, let $\pi(G_0)$ denote the graph G_0 when its vertices are renamed according to π: (i, j) is an edge of G_0 iff $(\pi(i), \pi(j))$ is an edge of $\pi(G_0)$. Then (7) holds iff there exists π such that $\pi(G_0) = G_1$. Call π a **certificate** for $\langle G_0, G_1\rangle$ in this case. Thus $\langle G_0, G_1\rangle \in$ ISO iff $\langle G_0, G_1\rangle$ has a certificate. This concept of certificate has two important properties:

- (Succinctness) The certicate π has size polynomial in n.
- (Verifiability) There is a deterministic polynomial-time algorithm V to decide if a given π is a certificate for $\langle G_0, G_1\rangle$.

These two properties characterize languages in *NP*. Hence, ISO $\in$ *NP*. We can next try to improve our upper bound on ISO (is it in *P* ?) or prove a lower bound (is it *NP*-hard?). Both of these questions are open. Another related upper bound question (is it in co-*NP*) is also open.

It easily follows NONISO $\in$ co-*NP*. Unfortunately, co-*NP* does not have a characterization by certificates. Although certificates are not necessarily easy to find, they are easy to verify. In this sense, they have practical utility. We next introduce a generalization of certificates verifiability, and eventually show that NONISO is verifiable in this more general sense.

Concept of Interactive Proofs. We generalize certificate verifiability in two ways: first, we allows the verifying algorithm V to be probabilistic, and second, we allow interaction between the verifying algorithm with another algorithm called the "prover", denoted P. Thus there are two communicating processes (sometimes called *protocols*), an *interactive prover* P and an *interactive verifier* V which are Turing machines that send each other messages,

$$m_0, m_1, m_2, \ldots.$$

Message m_i is written by V if i is even, and by P if i is odd, and these are written on a common worktape. We assume some convention for each process to indicate that it is done writing its message (say, by entering a special state) and for some external agent to prompt the other process to continue. The computation ends when V answers YES or NO. The input is originally on V's input tape. We place complexity bounds on V alone. Thus the time and space in a (V, P)-computation refer solely to the time and space incurred by V alone. We place no restriction on P (which need not even have to be computable), except that P must respond to each message of V in finite time. For instance, suppose if we say that V accepts in polynomial time, and P turns out to write exponentially long messages, then V will not be able to read the long messages of P. Intuitively, V is sceptical about what the process P is communicating to it, and needs to be "convinced" (with high probability). For any input w, let $\Pr(V, P, w)$ be the probability that V accept. We will assume that V halt on every computation path, thus avoiding any discussion of probability intervals. Languages will be defined with respect to V alone: writing

$$\Pr(V, w) := \sup_P \Pr(V, P, w),$$

then the language accepted by V is

$$L(V) := \{w : \Pr(V, w) > 1/2\}.$$

Say V has **bounded-error** if, in addition to the preceding requirements, we have $\Pr(V, w) \geq 2/3$ or $\Pr(V, w) \leq 1/3$ for all input w, The class *IP* comprises those languages that are accepted by bounded-error polynomial-time verifiers.

Interactive Verifier for Graph Non-Isomorphism. We want to describe an interactive verifer V such that $L(V) =$ NONISO. Here is a well-known V_0 from the literature:

INPUT: STRING w
1. Reject unless $w = \langle G_0, G_1\rangle \in \mathcal{G}_n^2$, $n \geq 1$.
2. Randomly generate an n-permutation π and a binary bit b.
3. Let $H \leftarrow \pi(G_b)$.
4. Send message $m_0 = \langle H, G_0, G_1\rangle$. This message asks P whether $H \sim G_0$ or $H \sim G_1$.
5. (Pause for P to reply with message m_1)
6. If $b = m_1$ answer YES, else answer NO.

Note that V_0 concludes in two message rounds (sends and receives a message). Assume $w = \langle G_0, G_1\rangle$. There are two cases to consider.

- $w \in$ NONISO: We claim $\Pr(V_0, w) = 1$. To see this, suppose P_0 is the prover who sends the message $m_1 = c$ such that $H \sim G_c$. Since c is unique, V_0 always answer YES, so $\Pr(V_0, P_0, w) = 1$.
- $w \notin$ NONISO: We wish to claim $\Pr(V_0, w) = 1/2$. Intuitively, an "honest prover" P_0 cannot distinguish whether the answer should be $H \sim G_0$ or $H \sim G_1$. It is reasonable for P_0 to flip a coin and answer $m_1 = 0$ and $m_1 = 1$ with equal probability. This will establish our claim. But suppose we have a "dishonest prover" P_1 whose goal is to mislead V_0 into accepting w. P_1 knows something about about π and b it may be able to mislead V_0. For instance, if P_1 knows the value of b, then it will always fool V_0. How can we be sure that such information has not leaked in our definition of message m_0? This justification is non-trivial (see [20, p. 175]) and may be based on the so-called Principle of Deferred Decisions.

This example points out that informal descriptions of interactive proofs (with suggestive language such as "V is convinced", "P knows", etc) can be tricky to formalize. For this reason, we prefer to view interactive proofs as choice computations. The idea is this: we can combine V and P into one choice machine denoted, loosely,

$$\mathrm{M} = \text{“}V + P\text{”},$$

where the states of M is the disjoint union of the states of V and P (so each state of M may be classified as a P-state or a V-state). We will let the choice function at each V-state q be $\gamma(q) = \oplus$. But what about P? We need to simulate all possible behavior for P (recall that we define $\Pr(V, w)$ as the maximization of $\Pr(V, P, w)$). This is not hard (we can essentially make all possible choices for the message). Furthermore, we let the choice function at each P-state q be $\gamma(q) = \vee$. Thus M is a $\{\oplus, \vee\}$-machine. Unfortunately, there are two issues.

One issue is that, although P is powerful, it seems that we do not want it to know about the coin tosses of V. Such a verifier is said to use "private coins". The formulation "$\mathrm{M} = V + P$" apparently use "public coins". As noted, the use of public coins in V_0 above would be disastrous for the NONISO protocol above. Verifiers with private coins seems more powerful. It turns out, for polynomial-time computations, a verifier with private coins can be simulated by one with public coins, at the cost of of two extra rounds [14]:

$$IP[k] \subseteq AM[k+2]. \tag{8}$$

The parameters k and $k+2$ bound the number of message rounds; the full explanation for this notation is given below,

The second issue is this: the 0/1-message m_1 returned by the prover is not easily modeled by $\oplus$- and $\vee$-choices alone. The ability to pass a 0/1-message from P to V seems more powerful than simply allowing V to ask P question and receiving a YES/NO answer. For instance, V upon receipt of the Boolean message m_1, can trivially compute the negation of m_1. But a $\{\oplus, \vee\}$-computation cannot trivially negate the value at a node. Thus it seems we need a $\{\oplus, \vee, \neg\}$-machine (equivalently, a $\{\oplus, \vee, \wedge\}$-machine) to efficiently simulate an interactive proof. But this would make the interactive prover for NONISO uninteresting (NONISO is trivially accepted by a $\wedge$-machine in polynomial-time.) It turns out $\neg$ can be avoided, but this is a non-trivial result. We can avoid all these complications of interactive proofs by using the Arthur-Merlin formulation of Babai and Moran. The advantage, besides its simplicity, is the direct connection to choice computation.

Arthur-Merlin Games. Let M be an $\{\oplus, \vee\}$-machine, and $\pi = (C_0, C_1, \ldots, C_m)$ be a computation path of M. A computation sequence

$$\pi' = (C_i, C_{i+1}, \ldots, C_j), \quad (1 \leq i \leq j \leq m)$$

is called a **$\oplus$-round** (or **Arthur round**) if it contains at least one $\oplus$-configuration but no $\vee$-configurations. Notice that π' could contain deterministic configurations. Similarly, π' is called a **$\vee$-round** (or **Merlin round**) if we interchange $\oplus$ and $\vee$. We say π has k **rounds** if π can be divided into k subpaths,

$$\pi = \pi_1; \pi_2; \cdots; \pi_k \quad (k \geq 1)$$

where ";' denotes concatenation of subpaths such that π_i is an Arthur round iff π_{i+1} is a Merlin round. Note that k is uniquely determined by π. The definition generalizes in a natural way to $k = 0$ and $k = \infty$. We say M is a k**-round** *Arthur-Merlin game* if

- M has bounded error and runs in polynomial time
- every computation path of M has at most k rounds in which the first round (if any) is an Arthur round

A k**-round** *Merlin-Arthur game* is similarly defined, with the roles of Arthur and Merlin interchanged. Let

$$AM[k] = \{L(\mathrm{M}) : \mathrm{M} \text{ is an Arthur-Merlin game with at most } k \text{ rounds}\}$$

The class $MA[k]$ is similarly defined using Merlin-Arthur games instead. Of course, we can generalize this to $AM[t(n)]$ and $MA[t(n)]$ where $t(n)$ is a complexity function.

We can identify[4] Arthur with the verifier V, and Merlin with the prover P, of interactive proofs. We can similar define the classes $IP[k]$ and $IP[t(n)]$ accepted by interactive proofs in k or $t(n)$ rounds. This is the notation used in (8).

So it turns out that the apparent gap between interactive provers (V, P) and choice machines is non-existent. But this result is non-trivial since the pair (V, P) can communicate as in true parallelism. As discussed in Chapter 1, the choice mechanism is, in general, weaker than true parallelism.

8.5.1 Arthur-Merlin Game for Graph Non-Isomorphism.

The new idea for checking non-isomorphism is as follows. Let $G_0, G_1 \in \mathcal{G}_n$. Consider the set

$$\texttt{ALIKE}_1(G_0, G_1) = \{H : H \sim G_0 \text{ or } H \sim G_1\}.$$

Intuitively, $|\texttt{ALIKE}_1(G_0, G_1)|$ is either $n!$ or $2(n!)$, depending on whether $G_0 \sim G_1$ or not. Unfortunately, this is not always true and relates to the notion of an automorphism: we call $\pi \in S_n$ an **automorphism** of G if $\pi(G) = G$. Of course, the identity $\mathbf{1}$ permutation is always an automorphism, but this is the trivial case. We initially assume that G_i $(i = 0, 1)$ has only the trivial automorphism, so that $|\texttt{ALIKE}_1(G_0, G_1)| = n!$ or $2(n!)$,

Example: Let $n = 3$ and G_0 has the single edge $(1, 2)$ and G_1 has the single edge $(2, 3)$. Thus the transposition $(1, 2)$ is an automorphism of G_0. We see that the set $\texttt{ALIKE}_1(G_0, G_1) = \{G_0, G_1, G_2\}$ where G_2 has the single edge $(1, 3)$. So $|\texttt{ALIKE}_1(G_0, G_1)| = 3 < 3!$. Thus, our assumption that the G_i's have only the trivial automorphism is essential.

Continuing, suppose we randomly pick $H \in \mathcal{G}_n$, then there is some constant c such that $\Pr\{H \in \texttt{ALIKE}_1(G_0, G_1)\}$ is c or $2c$, depending on whether $G_0 \sim G_1$ or not. This probabilistic gap is the basis for recognizing NONISO. However, the constant c is exponentially small in n since $|\mathcal{G}_n| = 2^{\binom{n}{2}}\, n!$. We need to modify this gap by a hashing trick: the idea is to map $\mathcal{G}_n$ into a smaller set of size $\Theta(n!)$.

Let us understand the problem more generally: we have a set C whose size is either $n!$ or $2(n!)$. You want to computationally determine whether it is $n!$ or $2(n!)$. We cannot count explicitly since $n!$ is too large. Without loss of generality, let $C \subseteq \mathbb{B}^m$ (bit strings of length m). Here $\mathbb{B} = \{0, 1\}$ is the Boolean field of two elements. If $n!$ is comparable to $|\mathbb{B}^m| = 2^m$, say $3(n!) < 2^m < 4(n!)$, then we can randomly sample elements $x \in \mathbb{B}^m$ and test if $x \in C$ (we assume testing membership in C is easy). If $|C| = 2(n!)$, probability that $x \in C$ is greater than $1/2$. If $|C| = n!$, then the probability is less than $1/3$. This gap implies that our problem is in BPP. In our application, $n!$ is not comparable to 2^m but exponentially smaller. So we define a hash function $h : \mathbb{B}^m \to \mathbb{B}^k$ where $n!$ is comparable to $|\mathbb{B}^k| = 2^k$. If h is $1-1$ when restricted to C, then the previous BPP-algorithm will work. Unfortunately, we do not know how to find such an h efficiently. The next lemma say that if h is a random hash function, we can achieve much the same result.

Lemma 19 (Boppana). *Let B be a $k \times m$ Boolean matrix, and let*

$$h_B : \mathbb{B}^m \to \mathbb{B}^k$$

be defined by $h_B(x) = B \cdot x$ where all arithmetic is in $\mathbb{B}$. Fix $C \subseteq \mathbb{B}^m$ and write $h_B(C) = \{h_B(x) : x \in C\}$. Assume $c = |C|/2^k$ and $k \geq 2$. If $z \in \mathbb{B}^k$ and $B \in \mathbb{B}^{m \times k}$ are both random then

$$\Pr\{z \in h_B(C)\} > c - \frac{c^2}{2} = c\left(1 - \frac{c}{2}\right).$$

[4] As in the legend of King Arthur, the magician Merlin is more powerful than Arthur. Merlin, as the $\vee$-player, can use existential guesses (which is magically correct). Arthur, all too human, can only roll his dice and take his chances.

Proof. We will show that for all $x \neq y$,
(a) $\Pr\{z = h_B(x)\} = 1/2^k$, and
(b) $\Pr\{z = h_B(x) = h_B(y)\} \leq 1/4^k$.
The lemma then follows by the inclusion-exclusion principle,

$$\Pr\{z \in h_B(C)\} \geq \sum_{x \in C} \Pr\{z = h_B(x)\} - \sum_{\{x,y\} \in \binom{C}{2}} \Pr\{z = h_B(x) = h_B(y)\}.$$

(a) There are two cases for showing $\Pr\{z = h_B(x)\} = 1/2^k$: fix $x = (x_1, \ldots, x_m)$. If $x = \mathbf{0}_m$ (an m-vector of 0's) then $\Pr\{z = h_B(x)\} = \Pr\{z = \mathbf{0}_k\} = 1/2^k$. Otherwise, we may assume $x_1 \neq 0$. Let B_i denote the ith column of B. First fix $z' \in \mathbb{B}^k$; then for any choices of $B_2, \ldots, B_m$ there is a unique choice of B_1 such that $Bx = z'$. Thus

$$\Pr\{Bx = z'\} = \sum_{B_2, \ldots, B_m} 2^{-mk} = 2^{(m-1)k} 2^{-mk} = 2^{-k}. \tag{9}$$

If z is also random, then $\Pr\{Bx = z\} = \sum_{z'} 2^{-k} \Pr\{Bx = z'\} = 2^{-k}$. This proves (a).

(b) There are also two cases in showing $\Pr\{z = Bx = By\} \leq 1/4^k$. First, suppose $x = \mathbf{0}_m$. Then

$$\begin{aligned} \Pr\{z = Bx = By\} &= \Pr\{z = \mathbf{0}_k, By = \mathbf{0}_k\} \\ &= \Pr\{z = \mathbf{0}_k\} \Pr\{By = \mathbf{0}_k\} \\ &= 2^{-k} 2^{-k}, \end{aligned}$$

where the last equality uses part (a). Next suppose $x \neq \mathbf{0}_m$ and, by symmetry, $y \neq \mathbf{0}_m$. In this case, since $x \neq y$, there is some $1 \leq i < j \leq m$ such that

$$M = \begin{bmatrix} x_i & y_i \\ x_j & y_j \end{bmatrix}$$

is non-singular. To see this, without loss of generality, assume $i = 1, j = 2$ and $x_i = 1, y_i = 0$. Then we can choose $y_2 = 1$, and it does not matter what x_2 is. For any fixed $z' \in \mathbb{B}^k$ and any choice of columns $B_3, \ldots, B_m$, there is a constant $k \times 2$ matrix C such that the equation $z' = Bx = By$ can be rewritten as

$$[B_1 | B_2] M = C.$$

Since M is invertible, this uniquely determines B_1, B_2. There are $2^{(m-2)k}$ choices for $B_3, \ldots, B_m$ and hence $\Pr\{z' = Bx = By\} = 4^{-k}$. Again, when z is randomly chosen, this leads to $\Pr\{z = Bx = By\} = 4^{-k}$. This proves (b). **Q.E.D.**

The next idea is to get rid of the assumption that G_i has only the trivial automorphism. We use elementary facts of group theory. For any digraph G, let $\mathtt{aut}(G)$ denote the automorphism group of G. It is standard to look at the cosets of $\mathtt{aut}(G)$: these are sets of the form

$$\pi \circ \mathtt{aut}(G) = \{\pi \circ \sigma : \sigma \in \mathtt{aut}(G)\}$$

for each $\pi \in S_n$. These cosets form a partition of S_n. Here is the standard argument: (1) It is clear that S_n is the union of these cosets since $\mathbf{1} \in \mathtt{aut}(G)$. (2) $\pi \circ \mathtt{aut}(G)$ and $\pi' \circ \mathtt{aut}(G)$ are either disjoint or equal: for, if $\pi \circ \sigma = \pi' \circ \sigma'$ for some $\sigma, \sigma' \in \mathtt{aut}(G)$, then $\pi = \pi' \circ \sigma' \circ \sigma^{-1}$ so that $\pi \in \pi' \circ \mathtt{aut}(G)$. Then $\pi \circ \mathtt{aut}(G) \subseteq \pi' \circ \mathtt{aut}(G)$; reversing the argument implies $\pi \circ \mathtt{aut}(G) = \pi' \circ \mathtt{aut}(G)$.

We also note that each coset C has size $|\mathtt{aut}(G)|$ since $\pi^{-1} C = \mathtt{aut}(G)$ for some π. It follows that the number d of distinct cosets must divide $|S_n| = n!$. In fact, $d = n!/|\mathtt{aut}(G)|$. If the distinct cosets are C_i $(i = 1, \ldots, d)$, choose a representative $\pi_i \in S_n$ for each i such that $\pi_i \mathtt{aut}(G) = C_i$. Now consider the

$$\mathtt{iso}(G) = \{\pi_i(G) : i = 1, \ldots, d\}$$

It is clear that $\mathtt{iso}(G)$ is a complete list of all the graphs isomorphic to G. Moreover, it is not hard to check that $\mathtt{aut}(\pi_i(G)) = \pi_i \mathtt{aut}(G) \pi_i^{-1}$ [pf: $\sigma(\pi_i(G)) = \pi_i(G)$ iff $\pi_i^{-1} \sigma \pi_i(G) = G$].

It follows that the set

$$|\mathtt{iso}(G) \times \mathtt{aut}(G)| = n! \tag{10}$$

So if we count each graph $H \in \mathtt{iso}(G)$ with multiplicity equal to $|\mathtt{aut}(G)|$, the size of the multiset $\mathtt{iso}(G)$ would be $n!$. Equivalently, we "tag" each H with one of its automorphisms π, and count the corresponding set $\mathtt{iso}(G) \times \mathtt{aut}(G)$ set instead of $\mathtt{iso}(G)$. This tagging idea will provide a suitable replacement for $\mathtt{ALIKE}_1(G_0, G_1)$: define $\mathtt{ALIKE}(G_0, G_1)$ to be

$$\mathtt{ALIKE}(G_0, G_1) := \{\langle H, \pi \rangle : \pi(H) = H \text{ and } H \sim G_0 \text{ or } H \sim G_1\}.$$

LEMMA 20.

$$|\mathtt{ALIKE}(G_0, G_1)| = \begin{cases} n! & \text{if } G_0 \sim G_1, \\ 2 \cdot n! & \text{if } G_0 \not\sim G_1. \end{cases} \tag{11}$$

Proof. If $G_0 \sim G_1$, then this is just a restatement of (10). If $G_0 \not\sim G_1$, then the $\mathtt{ALIKE}(G_0, G_1)$ is the disjoint union of $\mathtt{iso}(G_0) \times \mathtt{aut}(G_0)$ and $\mathtt{iso}(G_1) \times \mathtt{aut}(G_1)$. **Q.E.D.**

We are ready to prove that NONISO belongs to AM.

THEOREM 21. NONISO $\in AM[2]$.

Proof. Assume each element of $\mathcal{G}_n$ is given by a string in $\mathbb{B}^{\binom{n}{2}}$. Hence each element of $\mathtt{ALIKE}(G_0, G_1)$ can be represented by a string in $\mathbb{B}^m$ where $m = \binom{n}{2} + \lceil n \lg n \rceil$. Choose $k = \lceil \lg(n!) \rceil + 2$. As in Boppana's lemma, let $z \in \mathbb{B}^k$ and $B \in \mathbb{B}^{m \times k}$ be randomly chosen. Let us define $c := 2(n!)/2^k$. Note that if $G_0 \sim G_1$, then $|\mathtt{ALIKE}(G_0, G_1)| \leq n!$ and

$$\Pr\{z \in h_B(\mathtt{ALIKE}(G_0, G_1))\} \leq \frac{n!}{2^k} = c/2.$$

On the other hand, if $G_0 \not\sim G_1$ and applying lemma 19 with $C = \mathtt{ALIKE}(G_0, G_1)$ yields

$$\Pr\{z \in h_B(\mathtt{ALIKE}(G_0, G_1))\} \geq c\left(1 - \frac{c}{2}\right)$$

where $|C|/2^k = c$. Since $c \leq \frac{1}{2}$, we have

$$\Pr\{z \in h_B(\mathtt{ALIKE}(G_0, G_1))\} \geq 3c/4.$$

This gives rise to the following $\{\oplus, \vee\}$-machine M_0:

INPUT: $\langle G_0, G_1 \rangle \in \mathcal{G}_n^2$.
1. Randomly choose B and z, as above.
2. Existentially choose a bit b, $H \in \mathcal{G}_n$ and two n-permutations π, σ.
3. Answer YES if $\pi(H) = H$ and $\sigma(H) = G_b$ and $z = h_B(H)$; else Answer NO.

This machine has probability gap $(c/2, 3c/4)$. It does not accept NONISO yet because the gap is not an error gap (which would contain $\frac{1}{2}$). We need to "shift" this gap by modifying M_0 as follows: begin by tossing a coin to spawn 2 branches. On one branch, perform the computations of M_0. On the other, perform a probabilistic computation whose valuation is exactly $1-(5c/8)$. Let us see how the second branch can be accomplished. First, we can compute the binary representation of c in polynomial time in a straightforward manner (in fact, $O(n^3 \log^2 n)$ time suffices). Next, compute the binary representation of $1 - (5c/8)$. Once this is available, we can make a series of coin tosses to achieve a valuation of exactly $p = 1 - (5c/8)$. In general, let the binary representation of a fraction $0 \leq p \leq 1$ is $0.p_1 p_2 \ldots p_k$ for some k. Note that $1 - (5c/8)$ is of this form. Then following randomized algorithm accepts with probability exactly equal to p:

L_0 : for i=1 to k
        Randomly choose labels L_0 or L_1
L_1 : if $p_i = 1$ then Answer YES else Answer NO.

Note that the modified machine M_1 has the error gap

$$\left(\frac{c}{2}, \frac{3c}{4}\right) \oplus \left[1 - \frac{5c}{8}, 1 - \frac{5c}{8}\right] = \left(\frac{1}{2} - \frac{c}{16}, \frac{1}{2} + \frac{c}{16}\right).$$

In fact, $c/16 \geq 2^{-6}$. Thus M_1 is an $AM[2]$-game. **Q.E.D.**

Note that using (8) and the original interactive verifier V_0, we only infer NONISO $\in AM[4]$. Hence this direct proof yields a sharper result.

8.5.2 Arithmetization of Quantified Boolean Formulas

The next two subsections uses a technique called "arithmetization" of quantified Boolean formulas. We recall the notion of a **quantified Boolen formula** F: In the base case, F is a Boolean variable x or the constants 0 or 1. Inductively, if F_1, F_2 are quantified Boolean formulas, then F has the form in the left column of the following table,

F	$\|F\|$
$x, 0, 1$	1
$\neg F_1$	$1 + \|F_1\|$
$(F_1 \vee F_2)$	$\|F_1 + \|F_2\|\|$
$(F_1 \wedge F_2)$	$\|F_1\| + \|F_2\|$
$(\exists x)[F_1]$	$1 + \|F_1\|$
$(\forall x)[F_1]$	$1 + \|F_1\|$

F_1, F_2 are quantified Boolean formulas. The **length** $|F|$ of F is given in the right column of the preceding table.

In the following, we simply say "formulas" for quantified Boolean formulas. Parenthesis or brackets may be dropped from formulas if this does not lead to ambiguity. Any occurence of the variable x in F_1 is said to be **bound** in $(\exists x)F_1$ and $(\forall x)F_1$. An occurence of the variable x is **free** if it is not bound. When we write $F = F(x_1, \ldots, x_n)$, this means that any variable that occurs free in F is among $x_1, \ldots, x_n$ (but it does not mean that each x_i actually occurs free in F). In this case, when we write $F(a_1, \ldots, a_n)$, it means we replace each free occurence of x_i in F by a_i. When F has no free variables, it is called a **quantified Boolean sentence** (or simply, "sentence"). We assume that the reader knows what it means for sentence to be true. If $F(x_1, \ldots, x_n)$, then F is **valid** if for every choice of Boolean values $a_i \in \mathbb{B}$, $F(a_1, \ldots, a_n)$ (which is a sentence) is true. A formula is **quantifier-free** when it does not have any quantifiers ($\forall$ or $\exists$). We say F is in **prenex form** if it has the form

$$(Q_1 x_1)(Q_2 x_2) \cdots (Q_m x_m)[\phi]$$

where $Q_i \in \{\forall, \exists\}$ and ϕ is quantifier-free.

Arithmetization. Suppose $F = F(x_1, \ldots, x_n)$ is a formula. Its "arithmetization" is an integer polynomial $\widetilde{F}(x_1, \ldots, x_n) \in \mathbb{Z}[x_1, \ldots, x_n]$ defined as follows:

- (Base Case) If $F \in \{0, 1, x_1, \ldots, x_n\}$, then $\widetilde{F} = F$.
- (Induction) If F_1, F_2 be quantified Boolean formulas, and $F(x_1, \ldots, x_n)$ has the form in the left column, then $\widetilde{F}$ has the form on the right column:

F	$\widetilde{F}$
$\neg F_1$	$1 - \widetilde{F_1}$
$(F_1 \wedge F_2)$	$\widetilde{F_1} \otimes \widetilde{F_2} = \widetilde{F_1}\widetilde{F_2}$
$(F_1 \vee F_2)$	$\widetilde{F_1} \oplus \widetilde{F_2} = 1 - (1 - \widetilde{F_1})(1 - \widetilde{F_2})$
$(\forall x)[F_1]$	$\widetilde{F_1}(x_1, \ldots, x_{n-1}, 0) \otimes \widetilde{F_1}(x_1, \ldots, x_{n-1}, 1)$
$(\exists x)[F_1]$	$\widetilde{F_1}(x_1, \ldots, x_{n-1}, 0) \oplus \widetilde{F_1}(x_1, \ldots, x_{n-1}, 1)$

For example, if $F = (\neg x)(x \vee y)$ then

$$\widetilde{F} = (1 - x)(x \oplus y) = x + y - x^2 - 2xy + x^2 y. \tag{12}$$

Here we use a common convention where x_1, x_2, x_3 are synonymous with x, y, z. Recall that $\mathbb{B} = \{0, 1\}$. Our polynomials have two simple properties:
(i) If each $a_i \subseteq \mathbb{B}$ value then $\widetilde{F}(a_1, \ldots, a_n) \in \mathbb{B}$.
(ii) $\widetilde{F}(a_1, \ldots, a_n) = 1$ iff the sentence $F(a_1, \ldots, a_n)$ is true.

Define $\mathcal{P}_n$ to be the set of polynomials $\widetilde{P}$ that arise from formulas F over the variables $x_1, \ldots, x_n$. Alternatively, $\mathcal{P}_n$ is the smallest set of integer polynomials such that (1) $\{0, 1, x_1, \ldots, x_n\} \subseteq \mathcal{P}_n$, and (2) if $p, q \in \mathcal{P}_n$ then $1 - p \in \mathcal{P}_n$ and $pq \in \mathcal{P}_n$. Viewing a polynomial p as a sum of **terms** (or monomials) where each term has the form $t = c \prod_{i=1}^n x_i^{e_i}$ $(c \neq 0)$. Let $e = (e_1, \ldots, e_n)$ and $|e| := \sum_{i=1}^n e_i$. Then we call $e = (e_1, \ldots, e_n)$ the **exponent**, $|e|$ the **degree**, $c \in \mathbb{Z}$ is the **coefficient** and $\prod_{i=1}^n x_i^{e_i}$ $(e_i \geq 0)$ is a **power product** of the term t. We also write x^e for the power product $x_1^{e_1} x_2^{e_2} \cdots x_n^{e_n}$. The **degree** of p, denoted $\deg(p)$, is defined as the maximum of the degrees of its terms.

Let $\deg_i(p)$ be the largest exponent $d \geq 0$ such that x_i^d divides a term in p Then the **maximum degree** of p is the vector

$$\texttt{MaxDeg}(p) = (e_1, \ldots, e_n)$$

where $e_i = \deg_i(p)$ for each i. We say p is **principal** when $x^{\texttt{MaxDeg}(p)}$ occurs in p; then the coefficient of $x^{\deg(p)}$ in p is called the **leading coefficient** of p. Define[5]

$$M(e) := \frac{|e|!}{e!}$$

where $e! := e_1!e_2! \cdots e_n!$ and $|e| = \sum_{i=1}^n e_i$ as above. For instance if $e = (0, 2, 1, 2)$ then $M(e) = \frac{5!}{0!2!1!2!} = 30$. We note a basic identity of multinomials:

LEMMA 22. *Let $g \in \mathbb{N}^n$ and $|g| = a + b$ $(a, b \in \mathbb{N})$. Let $I = I(g, a, b)$ be the set of all pairs of the form (e, f) such that $e, f \in \mathbb{N}^k$, $|e| = a, |f| = b$ and $e + f = g$.*

$$M(g) = \sum_{(e,f)\in I} M(e)M(f). \tag{13}$$

Proof. This proof exploits a counting interpretation of $M(g)$: $M(g)$ is the number of ways to color a set of $|g|$ elements with k colors. Let A, B be disjoint sets of a and b elements, respectively. The number of ways to color $A \cup B$ with k colors is therefore $M(g)$. But for each $(e, f) \in I$, there are $M(e)$ ways to k-color A and $M(f)$ ways to k-color B. Combining them, this gives rise to $M(e)M(f)$ ways to k-color $A \cup B$. Let $C(e, f)$ be the k-colorings of $A \cup B$ that is associated in this way with (e, f). Note that $C(e, f) \cap C(e', f') = \emptyset$ for $(e, f) \neq (e', f')$, and hence

$$M(g) \geq \sum_{(e,f)\in I} M(e)M(f).$$

It is also easy to see that every k-coloring of $A \cup B$ is a member of $C(e, f)$ for some (e, f). Hence the inequality is in fact an equality. **Q.E.D.**

A **homogeneous polynomial** is one in which every term has the same degree. For any polynomial p which does not involve x_0, id $d = \deg(p)$, we define its **homogeneous version** $\widehat{p}$ to be the result of replacing each term $c_e x^e$ in p by the term $c_e x^e x_0^{d-|e|}$. For instance, if $p(x, y) = p(x_1, x_2)$ is the polynomial in (12), then

$$\widetilde{p} = (x_0 - x_1)(x_0x_1 + x_0x_2 - x_1x_2) = x_0^2x_1 + x_0^2x_2 - x_0x_1^2 - 2x_0x_1x_2 + x_1^2x_2.$$

If $e = (e_1, \ldots, e_n)$ and $|e| \leq d$, define

$$M_d(e) = \frac{d!}{(d - |e|)!e_1! \cdots e_n!}$$

For each term with exponent e in a polynomial p of degree d, the corresponding term in $\widetilde{p}$ has exponent $e' = (d - |e|, e_1, \ldots, e_n)$. Then $M_d(e)$ is just $M(e')$. It is easy to see that

$$M(e) \leq M_d(e). \tag{14}$$

LEMMA 23. *Let $p \in \mathcal{P}_n$.*
(i) p is principal.
(ii) The leading coefficient of p is ± 1, and its constant term is either 0 or 1.
(iii) If $c_e x^e$ is a term in p, then $|c_e| \leq M_d(e)$ where $d = \deg(p)$.
(iv) For any formula F, the coefficients of $\widetilde{F}$ has bit size $O(|F| \log |F|)$.

Proof. (i) and (ii) are immediate.
(iii) (Basis) If $|e| = 0$, then $M_d(e) = 1$, and so the result holds when $p = 0$ or $p = 1$. If $|e| = 1$, then $M_d(e) \geq 1$, and thus it holds when $p = x_i$ or $p = 1 - x_i$. (Induction) Suppose $p, q \in \mathcal{P}_n$. By induction, the result is true for $1 - p$, and we want to verify the result for pq. Consider their homogeneous versions $\widehat{p}, \widehat{q}$. For $e = (e_0, \ldots, e_n) \in \mathbb{N}^{n+1}$ let $X^e = x_0^{e_0}x_1^{e_1}, \ldots, x_n^{e_n}$ and α_e be the coefficient of X^e in $\widehat{pq}$. Similarly, let β_e and γ_e be the coefficients of X^e in $\widehat{p}$

[5] $M(e)$ is also known as a **multinomial**, but this is terminology should be distinguished from the "monomial" terminology.

and in q, respectively. Then for any $e \in \mathbb{N}^{n+1}$,

$$\begin{aligned} \alpha_e &= \sum_{f+g=e} \beta_f \gamma_g, \\ |\alpha_e| &\leq \sum_{(f,g)\in I} |\beta_f| \cdot |\gamma_g| \\ &\leq \sum_{(f,g)\in I} M(f)M(g) \\ &= M(e), \end{aligned}$$

by the previous lemma. But if $e = (e_0, e_1, \ldots, e_n)$ and $e' = (e_1, \ldots, e_n)$ then $M(e)$ is just $M_d(e')$, as desired.
(iv) Let F be a formula and $d \leq \deg(\widetilde{F})$. If $e \in \mathbb{N}^n$ then from (iii), the coefficient of x^e in $\widetilde{F}$ has bit size at most $\log M_d(e) \leq \log(d!)$. The lemma follows since $d = O(|F|)$. **Q.E.D.**

8.5.3 *IP* contains $\mathbb{P}$.

In order to prove $\mathbb{P} \subseteq IP$, it is enough to show that a $\mathbb{P}$-complete problem (under many-one polynomial-time reducibility, say) is in *IP*. The same approach is used in the next section to show that $PSPACE \subseteq IP$, so this proof serves as warm-up. Our main result here is.

THEOREM 24. #SAT*is in IP.*

Let F be a quantifier-free formua over $x_1, \ldots, x_n$. From the definition of $\#F$ as the number of satisfying assignments to the formula F, we see that

$$\#F = \sum_{a_1=0}^{1} \sum_{a_2=0}^{1} \cdots \sum_{a_n=0}^{1} \widetilde{F}(a_1, \ldots, a_n).$$

We define the following polynomials for $i = 0, \ldots, n$:

$$F_i(x_1, \ldots, x_i) := \sum_{a_{i+1}\in\mathbb{B}} \sum_{a_{i+2}\in\mathbb{B}} \cdots \sum_{a_n\in\mathbb{B}} \widetilde{F}(x_1, \ldots, x_i, a_{i+1}, a_{i+2}, \ldots, a_n). \tag{15}$$

Notice that

- $F_n(x_1, \ldots, x_n) = \widetilde{F}(x_1, \ldots, x_n)$.
- For $i = 1, 2, \ldots, n$, $F_{i-1}(x_1, \ldots, x_{i-1}) = F_i(x_1, \ldots, x_{i-1}, 0) + F_i(x_1, \ldots, x_{i-1}, 1)$.
- $F_0 = \#F$.

Clearly, each $F_i \in \mathbb{P}_n$ and $\deg(F_i) \leq \deg F_n$. In particular, the coefficients of F_i has bit size $O(|F| \log |F|)$. The idea of the algorithm is to reduce the search for F_i to F_{i+1}, and hence reducing the search for $\#F = F_0$ to F_n (which we know). Moreover, since F_i is a i-variate polynomial, by random substitution for $i-1$ of these variables, we can reduce the checking to a univariate problem.

An *IP*-Algorithm for #SAT**.** Let the input for #SAT be $\langle F, k\rangle$ where F is a Boolean formula. The following algorithm is to accept iff $\#F \geq k$. Assume the polynomial $\widetilde{F}$ has n variables and degree $\leq d = |F|$.

1. Construct the polynomial
$$F_n = \widetilde{F}(x_1, \ldots, x_n) \in \mathbb{Z}[x_1, \ldots, x_n].$$
We represent F_n as an expression using the inductive rules in the definition of $\widetilde{F}$. Note that we do not expand F_n into an explicit sum of terms, which would be prohibitive. In the following, we refer to $F_i(x_1, \ldots, x_i)$ $(i = 0, \ldots, n-1)$, defined as in (15). Note that these F_i's will are never computed.
Let $S \subseteq \mathbb{Z}$ be any finite set of size at least $8dn$. For instance, we may let $S = \{0, 1, \ldots, 2^{\lceil lg(8dn)\rceil} - 1\}$. In the following, whenever we choose a random number r, it will be taken from S.

2. Stage 0: Existentially guess some k_0 where $k \leq k_0 \leq 2^n$. Our main task is to verify $F_0 = k_0$ (recall $F_0 = \#F$). How can we do this, seeing that we do not have F_0? We this indirectly, by guessing the coefficients of a polynomial $G_1(x) \in \mathbb{Z}[x]$ of degree $\leq d = |F|$. We intend[6] $G_1(x)$ to be equal to $F_1(x)$. Unlike the multivariate F_n, we can afford to represent the univariate G_1 as a sum of its terms. We then check if $G_1(0) + G_1(1) = k_0$. If this fails, we answer NO (our guess is wrong). If our guess were correct, then $G_1(0) + G_1(1) = F_0$ and our main task is next reduced to to verifying $G_1(x) = F_1(x)$, which is addressed in stage 1.

3. Stage $i = 1, \ldots, n-1$: Inductively assume that, at the beginning of stage i, we have a univariate polynomial $G_i(x)$ and a sequence of elements $r_1, \ldots, r_{i-1} \in \mathbb{Z}$. Each r_j is randomly chosen a previous stage (in fact, stage j). The task for this stage is to verify that

$$G_i(x) = F_i(r_1, \ldots, r_{i-1}, x). \tag{16}$$

   Here, $G_i(x)$ is explicitly represented by its sequence of coefficients, but $F_i(r_1, \ldots, r_{i-1}, x)$ is implicitly represented by $F_n(x_1, \ldots, x_n)$ (from step 2) and the values $r_1, \ldots, r_{i-1}$.

   To verify (16), we *existentially guess* (the coefficients of) a polynomial $G_{i+1}(x) \in \mathbb{Z}[x]$ of degree $\leq d$. Then we *randomly choose* $r_i \in S$. Notice that the modes for choosing $G_{i+1}(x)$ and for choosing r_i are different. Again, $G_{i+1}(x)$ is intended to be $F_{i+1}(r_1, \ldots, r_i, x)$. Since $F_i(r_1, \ldots, r_i) = F_{i+1}(r_1, \ldots, r_i, 0) + F_{i+1}(r_1, \ldots, r_i, 1)$, instead of (16), we verify that

$$G_i(r_i) = G_{i+1}(0) + G_{i+1}(1). \tag{17}$$

   If this fails, we answer NO. Otherwise, we proceed to the next stage.

4. Stage n: from stage $n-1$, we have inherited $G_n(x)$ and $r_1, \ldots, r_{n-1}$. We now need to verify $G_n(x) = F_n(r_1, \ldots, r_{n-1}, x)$ (cf. equation (16)). We randomly[7] guess $r_n \in S$, and answer YES if $G_n(r_n) = F_n(r_1, \ldots, r_{n-1}, r_n)$, otherwise answer NO. This final check is possible because we can easily evaluate $F_n(r_1, \ldots, r_{n-1}, r_n)$ from our representation of $F_n(x_1, \ldots, x_n)$ and $r_1, \ldots, r_n$.

Correctness. We prove that this procedure accepts #SAT. One direction is easy: suppose $\langle F, k\rangle \in \#\text{SAT}$. We may assume that all our existential guesses are correct. In particular, k_0 is correctly chosen to be $\#F$, and at each stage, the guessed polynomial $G_i(x)$ is indeed equal to $F_i(r_1, \ldots, r_{i-1}, x)$. Under these assumptions, regardless of the choice of the r_i's, we always answer YES in stage n. In fact, this shows that we accept $\langle F, k\rangle$ with no pessimistic acceptance error. Moreover, there is a computation tree of polynomial height, since each G_i has degree at most $d = |F|$ and its coefficients are $O(|F| \log |F|)$ bits.

In the harder direction, where $\langle F, k\rangle \notin \#\text{SAT}$, we have $\#F \neq k_0$ for all choices of $k_0 \geq k$. Fix any k_0. We will show that $\Pr(E_n) \geq 3/4$ where E_n be the event that we answer NO in stage n. For $i = 0, \ldots, n-1$, define E_i to be the event that we enter stage $i+1$ with

$$G_{i+1}(x) \neq F_{i+1}(r_1, \ldots, r_i, x). \tag{18}$$

For events A and B, we have $\Pr(A) \geq \Pr(A|B)\Pr(B)$. Iterating this,

$$\begin{aligned}
\Pr(E_n) &\geq \Pr(E_n|E_{n-1})\Pr(E_{n-1}) \\
&\geq \Pr(E_n|E_{n-1})\Pr(E_{n-1}|E_{n-2})\Pr(E_{n-2}) \\
&\geq \cdots \\
&\geq \Pr(E_n|E_{n-1})\Pr(E_{n-1}|E_{n-2}) \cdots \Pr(E_1|E_0)\Pr(E_0) \\
&= \Pr(E_0)\prod_{i=1}^{n} \Pr(E_i|E_{i-1}).
\end{aligned} \tag{19}$$

In stage 0, the fact that $G_1(0) + G_1(1) = k_0$ and $\#F = F_1(0) + F_1(1) \neq k_0$ implies that $G_1(x) \neq F_1(x)$. Hence $\Pr(E_0) = 1$.

In stage 1, we $\exists$-guessed $G_2(x)$ of degree $\leq d$, and randomly choose a number r_1 such that $G_1(r_1) = G_2(0) + G_2(1)$. Since the total degrees of G_1, F_1 are $\leq d$, the fact that $G_1(x) \neq F_1(x)$ implies follows that there are at most d choices of $r \in \mathbb{Z}$ such that $G_1(r) = F_1(r)$. This follows from a simple fact of algebra that a non-zero polynomial $p(x) \in \mathbb{Z}[x]$ of degree d has at most d zeros. Here, $p(x) = G_1(x) - F_1(x)$. Thus $\Pr(E_1|E_0) \geq 1 - (d/p)$.

[6] Mnemonic: think of the G's as "guesses" for the F's.

[7] We could actually perform this step in deterministic polynomial time.

The same argument works for stage $i+1$ ($i = 1, \ldots, n-2$): assuming (18), $\Pr\{G_{i+1}(r_{i+1}) \neq F_{i+1}(r_1, \ldots, r_i, r_{i+1})\} \leq d/p$. Thus $\Pr(E_{i+1}|E_i) \geq 1 - (d/p)$. It is also true that $\Pr(E_n|E_{n-1}) \geq 1 - (d/p)$. From (19), we conclude that

$$\begin{aligned} \Pr(E_n) &\geq \left(1 - \frac{d}{p}\right)^n \\ &>_1 (e^{-2d/p})^n \\ &>_2 1 - \frac{2nd}{p} \\ &> 3/4 \qquad (\text{since } \; p > 8dn). \end{aligned}$$

See the Appendix for the inqualities $>_1$ and $>_2$. This proves the correctness of our *IP*-algorithm. Since #SAT is $\mathbb{P}$-complete (say, under Karp-reducibility) and *IP* is clearly closed under Karp-reducibility, we have shown:

THEOREM 25.

$$\mathbb{P} \subseteq IP.$$

Neither of the inclusions $NP \subseteq IP$, and co-$NP \subseteq IP$ are obvious because of the requirement of bounded error in *IP*. But as $3SAT$ is trivially Karp-reducible to #SAT, it follows that $NP \subseteq IP$. Since *IP* is closed under complementation, this also means that co-$NP \subseteq IP$.

COROLLARY 26.

$$NP \cup \text{co-}NP \subseteq IP.$$

Remark: In the above proof, the equality of polynomials is always taken in the abstract mathematical sense, not in terms of their representations. Thus $F(a_1, \ldots, a_n)$ is equal to 0 or 1 when the $a_i \in \mathbb{B}$. On the other hand, we need to address the issue of representation when we construct explicit polynomials (e.g., F_n or the guessed G_i's).

8.5.4 *IP = PSPACE*

In Chapter 7, we proved $PrA\text{-}TIME(t) \subseteq ATIME(t \log t)$. Thus,

$$IP \subseteq PrA\text{-}TIME(n^{0(1)}) \subseteq ATIME(n^{O(1)}) = PSPACE.$$

We now prove the converse, $PSPACE \subseteq IP$, a result of Shamir [30]. The algebraic technique for showing $\mathbb{P} \subseteq IP$. in the previous section will be extended. In particular, we now show that a particular *PSPACE*-complete language belongs to *IP*. Define the set of **valid quantified Boolean formulas** to be

$$\text{QBF} = \{F : F \text{ is a valid formula}\}.$$

This can be viewed as a language, after choosing some standard encoding of formulas. We may assume $F \in$ QBF are in prenex form.

LEMMA 27. QBF *is PSPACE-complete.*

The proof is left as an Exercise. Hence our desired result amounts to showing an *IP*-algorithm for QBF. We initially try to imitate the previous proof, as this helps to locate the new difficulties. Let F be a formula. Unlike the previous proof, F may now have quantifiers. We may assume that

$$F = (Q_1 x_1 Q_2 x_2 \cdots Q_n x_n)[\phi] \tag{20}$$

where $\phi = \phi(x_1, \ldots, x_n)$ is quantifier-free. Define formulas F_i for $i = 0, \ldots, n$ as follows:

$$\begin{aligned} F_n(x_1, \ldots, x_n) &= \phi \\ F_{n-1}(x_1, \ldots, x_{n-1}) &= (Q_n x_n)[F_n] \\ &\vdots \\ F_{i-1}(x_1, \ldots, x_{i-1}) &= (Q_i x_i \cdots Q_n x_n)[\phi] = (Q_i x_i)[F_i] \\ &\vdots \\ F_1(x_1, x_2) &= (Q_3 x_3 Q_4 x_4 \cdots Q_n x_n)[\phi] \\ F_1(x_1) &= (Q_2 x_2 Q_3 x_3 \cdots Q_n x_n)[\phi] \\ F_0() &= F \end{aligned}$$

If $\widetilde{F}_i$ is the arithmetization of F_i, then we have

$$\widetilde{F}_i(x_1,\ldots,x_i) = \begin{cases} \widetilde{F}_{i+1}(x_1,\ldots,x_i,0) \otimes \widetilde{F}_{i+1}(x_1,\ldots,x_i,1) & \text{if } Q_i = \forall \\ \widetilde{F}_{i+1}(x_1,\ldots,x_i,0) \oplus \widetilde{F}_{i+1}(x_1,\ldots,x_i,1) & \text{if } Q_i = \exists \end{cases}$$

The problem is that the degree of $\widetilde{F}_i$ is double that of $\widetilde{F}_{i+1}$. Hence the degree of $\widetilde{F}$ may be exponential in $|F|$.

A Linearization Quantifier. To solve this problem, we introduce a new kind of quantifier denote L. The class of quantified Boolean formulas are now enlarged to include this new quantifier. In particular, for any formula $\phi(x_1,\ldots,x_n,x)$ and variable x, the following is also a formula,

$$F = Lx[\phi].$$

Furthermore, its arithmetization is defined via

$$\widetilde{F} := (1-x)\widetilde{\phi}(x_1,\ldots,x_n,0) + x\widetilde{\phi}(x_1,\ldots,x_n,1).$$

Note the following properties:

- If $a \in \mathbb{B}$, then $\widetilde{F}(x_1,\ldots,x_n,a) = \widetilde{\phi}(x_1,\ldots,x_n,a)$.
- $\widetilde{F}$ has the same set of free variables as $\widetilde{\phi}$ (namely, $x_1,\ldots,x_n,x$).
- $\widetilde{F}$ is linear in x.

The last two properties are quite unlike the arithmetization of the other two quantifiers. Indeed, the name of the L-quantifier is taken from the last property. An Exercise below shows that the linearlization transformation

$$\phi(x_1,\ldots,x_n,x) \mapsto (1-x)\phi(x_1,\ldots,x_n,0) + x\phi(x_1,\ldots,x_n,1)$$

amounts to replacing any power of x by a plain x in every term of the polynomial $\phi(x_1,\ldots,x_n,x)$. We now transform the formula F in (20) to a new formula H,

$$\begin{aligned} H &:= (Q_1x_1)(Lx_1Q_2x_2)(Lx_1Lx_2Q_3x_3)\cdots(Lx_1Lx_2\cdots Lx_{n-1}Q_nx_n)[\phi] \\ &= (\overline{Q}_1y_1\overline{Q}_2y_2\cdots\overline{Q_m}y_m)[\phi] \end{aligned}$$

where $\overline{Q}_i \in \{L,\forall,\exists\}$ and $y_i \in \{x_1,\ldots,x_n\}$. Here $m = \binom{n+1}{2}$ is the number of quantifiers in H. These m quantifiers are placed into n groups, as indicated by the matched pairs of parentheses. In each group (reading from right to left), a normal quantifier Q_i is followed by $i-1$ linearization operators to ensure that the degrees of the remaining variables are linear. As usual, let

$$\begin{aligned} H_m(x_1,\ldots,x_m) &= \phi \\ H_{m-1}(x_1,\ldots,x_m) &= (\overline{Q}_m y_m)[H_m] \\ H_{m-2}(x_1,\ldots,x_m) &= (\overline{Q}_{m-1}y_m)[H_{m-1}] = (\overline{Q}_{m-1}y_{m-1}\overline{Q}_m y_m)[H_m] \\ &\cdots \\ H_1(x_1) &= (\overline{Q}_2 y_2)[H_2] \\ H_0() &= (\overline{Q}_1 y_1)[H_1] = H \end{aligned}$$

Note that $H_i = H_i(x_1,\ldots,x_{n(i)})$ for some $n(i)$. The H_i's are formulas, but in our algorithm, we will simulate their arithmetized[8] versions $\widetilde{H}_i$. In particular, we will make $\exists$-guesses of $G_i = G_i(x_1,\ldots,x_{n(i)})$ which are intended to be equal to $\widetilde{H}_i$. The critical fact is that

$$\deg(\widetilde{H}_i) \leq |F|^2.$$

To see this, let $d_i = \deg(\widetilde{H}_i)$. Then we have $d_m \leq |F|$. However, d_{m-1} could be as large as $(d_m)^2$. However, the next $n-1$ linearization quantifiers reduces this degree back to at most $|F|$. This pattern is repeated for each group of linearization quantifiers that follows the next $\forall$ or $\exists$ quantifier (Q_{m-1}). We now see why the linearization quantifier is an antidote to the exponential degree of F_i's.

Here now is the *IP*-algorithm for QBF. The description will be abbreviated where similarities to the previous *IP*-algorithm for #SAT is clear.

[8] We really ought to write "$\widetilde{H_i}$". But this is uglier than the "$\widetilde{H}_i$" which we will use.

1. The input is a quantified Boolean formula F. Let H be the corresponding linearized version, and the formulas $H_0, \ldots, H_m$ is defined as above. Again, we assume a finite set S of size $|S| \geq$???. Whenever we choose a random number, it will be taken from S.

2. Stage 0: We guess $G_1(x)$, intended to be equal to $\widetilde{H}_1(x)$. Now $H = H_0 = (\overline{Q}_1 x_1)[H_1]$ and $\overline{Q}_1 = \forall$ or $\exists$. If $\overline{Q}_1 = \forall$, then we check if $G_1(0) = G_1(1) = 1$. If $\overline{Q}_1 = \exists$, then we check if $G_1(0) = 1$ or $G_1(1) = 1$. If the check fails, we answer NO; else we go to Stage 1.

3. Stage i ($i = 1, \ldots, m-1$). Inductively, we have random values $r_1, \ldots, r_{n(i)-1}$ and $G_i(x)$. Intuitively, $G_i(x)$ is meant to be $H_i(r_1, \ldots, r_{n(i)-1}, x)$. Now $H_i = (\overline{Q}_{i+1} y_{i+1})[H_{i+1}(x_1, \ldots, x_{n(i+1)})]$.

   Note that if $\overline{Q}_{i+1} = L$ then $n(i+1) = n(i)$; otherwise $\overline{Q}_{i+1} = \forall$ or $\exists$ and $n(i+1) = n(i) + 1$. Existentially guess $G_{i+1}(x)$. If $\overline{Q}_{i+1} \neq L$ then we randomly choose $r_{n(i+1)-1} = r_{n(i)}$. Regardless of $\overline{Q}_{i+1}$, the intention is that $G_{i+1}(x)$ is $H_i(r_1, \ldots, r_{n(i+1)-1}, x)$.

   We consider three cases:
   (i) If $\overline{Q}_i = \forall$, we check if $G_i(r_{n(i)}) = G_{i+1}(0) \otimes G_{i+1}(1)$
   (ii) If $\overline{Q}_i = \exists$, we check if $G_i(r_{n(i)}) = G_{i+1}(0) \oplus G_{i+1}(1)$
   (iii) If $\overline{Q}_i = \exists$, we check if $G_i(r_{n(i)}) = G_{i+1}(r_{n(i)})$

4.

Notes. Our approach to interactive proofs via choice computation is unconventional, but it provides a more satisfactory and general foundation than the customary treatment. This allows interactive proofs to seen to be a natural part of a wide spectrum of computational choices. A survey on interactive proofs may be found in [24]. The place of ISO in the complexity landscape is explored in the monograph of Köbler, Schöning and Torán [17]. Another closely related question is GRAPH AUTOMORPHISM: Given G, does it have a non-trivial automorphism?

Exercises

Exercise 0.20: Show that *IP* is closed under Karp- and Cook-reducibility. ♢

Exercise 0.21: This exercise helps you gain some facility with the group theoretic ideas in the NONISO $\in$ *IP* proof. Let $V = V_n = \{1, \ldots, n\}$ and S_n be the set of permutations on V_n. The trivial permutation is denoted $\mathbf{1}_n$ (or simply $\mathbf{1}$). Write the composition of $\sigma, \sigma' \in S_n$ in the form of a product $\sigma\sigma'$, instead of $\sigma \circ \sigma'$.
(i) Let $2 \leq k \leq n$. If $\{a_1, \ldots, a_k\} \subseteq \binom{V}{k}$, then $(a_1, \ldots, a_k) \in S_n$ denotes the permutation which takes each a_i to $a_{(i+1 \mod k)}$, called a **cyclic permutation**. Two special cases are $k = 2$ or $k = n$. Then $(a_1, \ldots, a_k)$ is **transpose** or a **Hamiltonian permutation**, respectively. Two cyclic permutations $(a_1, \ldots, a_k)$, $(b_1, \ldots, b_\ell)$ are disjoint if $a_i \neq b_j$ for all i, j. For instance, $(132)(45) = (45)(132)$ is a product of two disjoint cycles. The order of writing disjoint products does not matter. Show that every non-trivial permutation is a product of disjoint permutations.
(ii) Let G_0 be the digraph shown in Figure 8.6.

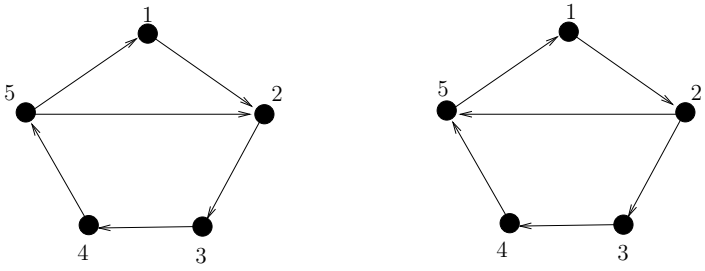

(a) G_0 (a) G_1

Figure 8.6: Two labeled digraphs G_0, G_1.

Determine $\texttt{iso}(G_0)$ and $\texttt{iso}(G_1)$. What the sizes of these two sets?
(iii) Determine $\texttt{aut}(G_0)$ and $\texttt{aut}(G_1)$. What the sizes of these two sets? ♢

Exercise 0.22: If $p \in \mathcal{P}_n$ has degree d, then there is a constant $C > 0$ such that the magnitude of each coefficient of p is at most

$$C\left(\frac{d}{\exp(1)}\right)^{n+(1/2)}\left(\frac{n+1}{d}\right)^{n+1}.$$

This sharpens the $d!$ bound used in the text. ♢

Exercise 0.23: Prove that QBF is *PSPACE*-complete. Since this proof can be found in the literature, we will enforce some originality in your solution by asking you to use the *same* framework as the proof of Cook's theorem in Chapter 3. **Hint:** The additional idea you need is found in the proof of Savage's theorem: if $C \vdash^{2m} C'$ (i.e., there is an $2m$-step path from C to C' then $C \vdash^{m} C''$ and $C'' \vdash^{m} C'$ for some C''. ♢

Exercise 0.24: Show that for any polynomial $\phi(x_1, \ldots, x_n, x)$, the linearization transformation

$$\phi(x_1, \ldots, x_n, x) \mapsto (1-x)\phi(x_1, \ldots, x_n, 0) + x\phi(x_1, \ldots, x_n, 1)$$

amounts to replacing any power of x by a plain x in every term of the polynomial $\phi(x_1, \ldots, x_n, x)$. ♢

Exercise 0.25: Are there reasons to believe that co-$NP \subseteq AM$, based on NONISO $\subseteq AM[2]$? Derive nontrivial consequences under the assumption co-$NP \subseteq AM$. ♢

Exercise 0.26: If $L \in AM[k]$ then L is accepted by an Arthur-Merlin game in $k+1$ rounds with zero-error acceptance. ♢

Exercise 0.27: How does one amplify error gaps for languages in $AM[k]$? ♢

End Exercises

8.6 Markov Chains and Space-bounded Computation

We want to study computations by space-bounded probabilistic machines. The behavior of such computations can be analyzed in terms of finite[9] *Markov chains*. We develop the needed results on Markov chains (see also the appendix in this chapter). For further reference on Markov chains, see [16, 10].

The main result of this section is

Theorem 28. *For all $s(n) \geq \log n$,*

$$PrSPACE(s) \subseteq DSPACE(s^2)$$

Notice that this result is yet another strengthening of Savitch's theorem! We follow the proof of Borodin, Cook and Pippenger [6]; Jung [15] independently obtained the same result using different techniques[10]. This result improves earlier simulations by Gill [12] (which uses exponential space) and by J. Simon [31] (which uses $s(n)^6$ space).

A sequence of non-negative real numbers $(p_1, p_2, \ldots, p_i, \ldots)$ is *stochastic* if the sum $\sum_{i\geq 1} p_i = 1$; it is *substochastic* if $\sum_{i\geq 1} p_i \leq 1$. A matrix is stochastic (resp., substochastic) if each row is stochastic (substochatic). In general, stochastic sequences and stochastic matrices may be denumerably infinite although we will only consider finite matrices. An n-state *Markov process* (or *Markov chain*) is characterized by an $n \times n$ stochastic matrix $A = (p_{i,j})_{i,j=1}^n$. Call A the *transition matrix*. The states of A will be called *Markov states*, as distinguished from machine states. We interpret this chain as an n-state finite automaton where $p_{i,j}$ is the probability of going from state i to state j. For any integer $k \geq 0$, the kth power $A^k = (p_{i,j}^{(k)})_{i,j}^n$ of A is defined inductively: A^0 is the identity matrix and $A^{k+1} = A \cdot A^k$. It is easy to check the product of stochastic matrices is stochastic; hence each A^k is stochastic. Clearly $p_{i,j}^{(k)}$ denotes the probability of a transition from state i to state j in exactly k steps.

Markov states admit a straight forward combinatorial classification. From the transition matrix $A = (p_{i,j})_{i,j=1}^n$ of the Markov chain, construct the Boolean matrix $B = (b_{i,j})_{i,j=1}^n$ where $b_{i,j} = 1$ iff $p_{i,j} > 0$. We view B as the

[9] *i.e.*, discrete time, homogeneous Markov processes, with finitely many states.

[10] Borodin, Cook and Pippenger uses redundant arithmetic techniques while Jung uses modular arithmetic techniques. Borodin, Cook and Pippenger states the result in a stronger form (in terms of circuit depth, see chapter 10), but it has essentially the proof to be presented here.

adjacency matrix of a directed graph G_A, called the *underlying graph* of the matrix A. We may form the transitive closure $B^* = (b^*_{i,j})^n_{i,j=1}$ of B (see chapter 2, section 6). As usual, define states i and j to be *strongly connected* if

$$b^*_{i,j} = b^*_{j,i} = 1.$$

This is easily seen to be an equivalence relationship and the equivalence classes form the *(strongly connected) components* of G_A. These strongly connected components in turn are related by the *reachability relation*: if C and C' are components, we say C *can reach* C' if there are states $i \in C$ and $j \in C'$ such that $b^*_{i,j} = 1$. It is easy to see that this definition does not depend on the choice of i and j. Furthermore, if C can reach C' and C' can reach C then $C = C'$. Thus the reachability relation induces an acyclic graph F on the components where F has an edge from C to C' iff C can reach C'. Those components C that cannot reach any other components are called *essential components* and the states in them known as *essential states*. The other components are called *inessential components* and their members known as *inessential states*.[11] We say state i is *absorbing* if $p_{i,i} = 1$. Such a state is clearly essential and forms a component by itself. A Markov chain is *absorbing* if all essential states are absorbing.

The above classification depends only on the underlying graph G_A. Let us now classify states by their stochastic properties. These notions properly belong to a subarea of probability theory called renewal theory. We introduce an important concept in renewal theory: let $f^{(n)}_{i,j}$ denote the probability that, starting from state i, we enter state j for the first time after n steps. We call these $f^{(n)}_{i,j}$ the *first entrance probabilities*. Write $f^{(n)}_i$ for $f^{(n)}_{i,i}$. It is not hard to see that for $n = 1, 2, \ldots$,

$$f^{(n)}_{i,j} = p^{(n)}_{i,j} - \sum_{k=1}^{n-1} f^{(k)}_{i,j} p^{(n-k)}_{j,j}$$

or,

$$p^{(n)}_{i,j} = \sum_{k=0}^{n} f^{(k)}_{i,j} p^{(n-k)}_{j,j} \tag{21}$$

where we conventionally take

$$f^{(0)}_{i,j} = 0, \quad p^{(0)}_{i,j} = \delta_{i,j}.$$

Here $\delta_{i,j}$ is Kronecker's delta function that assumes a value of 1 or 0 depending on whether $i = j$ or not. The sum

$$f^*_{i,j} = \sum_{n=1}^{\infty} f^{(n)}_{i,j}$$

clearly denotes the probability of ever reaching state j from i. Let f^*_i abbreviate $f^*_{i,i}$. We now define a state to be *recurrent* if $f^*_i = 1$ and *nonrecurrent* if $f^*_i < 1$.

LEMMA 29. *An inessential state is nonrecurrent.*

Proof. By definition, if state i is inessential, there is a finite path from i to some state outside the component of i. Then $f^*_i \leq 1 - c$ where $c > 0$ is the probability of taking this path.

Q.E.D.

The converse does not hold in general (Appendix and Exercise). But in the case of finite Markov chains, essential states are recurrent. To show this result, we proceed as follows: let $g^{(n)}_{i,j}$ denote the probability of the event $G^{(n)}_{i,j}$ that starting from state i we will visit state j at least n times. Note that $G^{(n+1)}_{i,j} \subseteq G^{(n)}_{i,j}$ and so we may define the limit

$$g_{i,j} := \lim_{n\to\infty} g^{(n)}_{i,j} = \Pr(\bigcap_{n=0}^{\infty} G^{(n)}_{i,j}).$$

It is not hard to see that $g_{i,j}$ is the probability that starting from state i we visit state j infinitely often. Again, let $g_{i,i}$ be abbreviated to g_i.

LEMMA 30.
(i) $g_i = 1$ or 0 according as i is recurrent or not.
(ii) In a finite Markov chain, essential states are recurrent.

[11] The reader should be aware that the classification of Markov states are not all consistent in the literature. The essential/inessential distinction is due to Chung [8]. His terminology is justified in the sense that every chain has at least one essential component; but it also seems to reflect an attitude in probabilistic theory that the most interesting phenomena occur in the essential components. This is unfortunate because we will see that the inessential components are more interesting for us!

Proof. (i) Note that

$$g_i^{(n+1)} = f_i^* g_i^{(n)}.$$

Since $g_i^{(1)} = f_i^*$, we get inductively

$$g_i^{(n+1)} = (f_i^*)^n.$$

Taking limits as $n \to \infty$, we see that $g_i = 1$ if $f_i^* = 1$ and $g_i = 0$ if $f_i^* < 1$.
(ii) Let $E_i^{(n)}$ be the event that starting from state i, there are no returns to state i after n steps. Clearly

$$E_i^{(1)} \subseteq E_i^{(2)} \subseteq \cdots$$

and $E_i := \bigcup_{n\geq 0} E_i^{(n)}$ is the event that there are only finitely many returns. But

$$\Pr(E_i^{(n)}) \leq 1 - e$$

where $e > 0$ is the minimum probability that any state in the component of i can get to state i. (To see this,

$$\Pr(E_i^{(n)}) = \sum_j p_{i,j}^{(n)}(1 - f_{j,i}^*) \leq (1-e)\sum_j p_{i,j}^{(n)}$$

which is at most $1 - e$.) Hence $\Pr(E_i) \leq 1 - e < 1$. But $g_i = 1 - \Pr(E_i)$. Hence $g_i > 0$ and by part (i), $g_i = 1$. This means state i is recurrent. **Q.E.D.**

We now see that for finite Markov chains, the combinatorial classification of essential/inessential states coincides with the stochastic classification of recurrent/nonrecurrent states. The appendix describe some refined classifications.

The *stochastic completion* of $A = (p_{i,j})_{i,j=1}^n$ is the matrix $A^* = (p_{i,j}^*)_{i,j=1}^n$ where

$$p_{i,j}^* = \sum_{k=0}^{\infty} p_{i,j}^{(k)}$$

with the understanding that the sum is ∞ when it diverges. The completion operation is defined even if A is not a stochastic matrix.[12]

The entries of A^* has this natural interpretation:

LEMMA 31.
(i) $p_{i,j}^*$ is the expected number of steps that the automaton spends in state j if it started out in state i.
(ii) Furthermore, if j cannot return to itself in one or more steps then $p_{i,j}^*$ is the probability that the automaton ever enters state j.

Proof. Interpretation (i) follows when we note $p_{i,j}^{(n)}$ is the expected fraction of time that the automaton spends in state j during nth unit time period, assuming that it started out in state i. For (ii), under the stated assumptions on state j, we see that $p_{i,j}^{(n)} = f_{i,j}^{(n)}$ and hence $p_{i,j}^* = f_{i,j}^*$. **Q.E.D.**

Let us introduce the following generating functions (see appendix) for state i:

$$\begin{aligned} F_i(s) &:= \sum_{n=0}^{\infty} f_i^{(n)} s^n \\ G_i(s) &:= \sum_{n=0}^{\infty} p_i^{(n)} s^n \end{aligned}$$

Using the relation (21), we see that

$$G_i(s) - 1 = F_i(s)G_i(s)$$

or,

$$G_i(s) = \frac{1}{1 - F_i(s)}$$

Now if we take the limit as $s \to 1^-$, the left hand side approaches $p_{j,j}^*$ and the right hand side approaches $\frac{1}{1-f_j^*}$. This proves

[12]The terminology is from in [6]. The notation $p_{i,j}^*$ is not to be confused with the limiting value of $p_{i,j}^{(k)}$ as $k \to \infty$. Unfortunately, a stochastic completion is no longer a stochastic matrix. This is obvious from the interpretation of $p_{i,j}^*$ as the expected number of steps in state j.

LEMMA 32.

$$p^*_{j,j} < \infty \iff f^*_j < 1.$$

To relate this to other values of $p^*_{i,j}$, we have

LEMMA 33.

$$p^*_{i,j} = \delta_{i,j} + f^*_{i,j} p^*_{j,j}.$$

Proof.

$$\begin{aligned} p^*_{i,j} &= \sum_{n=0}^{\infty} p^{(n)}_{i,j} \\ &= \delta_{i,j} + \sum_{n=1}^{\infty} \sum_{k=1}^{n} f^{(k)}_{i,j} p^{(n-k)}_{j,j} \\ &= \delta_{i,j} + \sum_{k=1}^{\infty} f^{(k)}_{i,j} \sum_{n=k}^{\infty} p^{(n-k)}_{j,j} \\ &= \delta_{i,j} + \sum_{k=1}^{\infty} f^{(k)}_{i,j} \sum_{n=0}^{\infty} p^{(n)}_{j,j} \\ &= \delta_{i,j} + f^*_{i,j} p^*_{j,j} \end{aligned}$$

Q.E.D.

COROLLARY 34. *For all i, j, if $f^*_{i,j} > 0$ then*

$$p^*_{i,j} < \infty \iff p^*_{j,j} < \infty.$$

We need one more basic fact [16].

LEMMA 35. *Let A be a square matrix such that $A^n \to 0$ as $n \to 0$, i.e., each entry of the nth power of A approaches zero as n approaches infinity. Then the matrix $I - A$ is nonsingular where I is the square matrix with the same dimensions as A. Moreover the infinite sum*

$$\sum_{n=0}^{\infty} A^n$$

converges, and this sum is given by

$$(I - A)^{-1} = \sum_{n=0}^{\infty} A^n.$$

Proof. We begin with the identity

$$(I - A)(I + A + A^2 + \cdots + A^n) = I - A^{n+1}.$$

Now the right hand side approaches I for large n. So for sufficiently large n, $\det(I - A^{n+1}) \neq 0$. This means $\det(I - A) \det(I + A + \cdots + A^n) \neq 0$. Thus $I - A$ is nonsingular, as asserted. So we may multiply both sides of the identity on the left with $(I - A)^{-1}$, giving a new identity

$$(I + A + A^2 + \cdots + A^n) = (I - A)^{-1}(I - A^{n+1}).$$

Now as $n \to \infty$, the left hand side approaches the infinite sum of the lemma and the right hand side approaches $(I - A)^{-1}$. Since the right hand side approaches a definite limit, so the left hand side approaches the same limit.
Q.E.D.

For any transition matrix A, let B be obtained by deleting the rows and columns of A corresponding to essential states. Following Kemeny and Snell, we call B the *fundamental part* of A. Note that B is a substochastic matrix. Then after permuting the rows and columns of A, we have

$$A = \begin{pmatrix} B & T \\ 0 & C \end{pmatrix}$$

where '0' denotes a matrix of zeroes of the appropriate dimensions. Moreover, the nth power of A is

$$A^n = \begin{pmatrix} B^n & T_n \\ 0 & C^n \end{pmatrix}$$

where B^n, C^n are the nth powers of B, C (respectively) and T_n is some matrix whose form need not concern us. Hence, the stochastic completion

$$A^* = \begin{pmatrix} B^* & T_* \\ 0 & C^* \end{pmatrix}$$

where B^*, C^* are the stochastic completions of B, C (respectively). In our applications, we only need B^*. Note that the entries in C^*, T_* are 0 or ∞, by what we have proved.

From the above development, the entries $p_{i,j}^{(n)}$ in B^n satisfy the property $\sum_{n=0}^{\infty} p_{i,j}^{(n)} < \infty$. This means $p_{i,j}^{(n)} \to 0$ as $n \to \infty$. Hence $B^n \to 0$ as $n \to \infty$ and the preceding lemma shows that B^* converges to $(B - I)^{-1}$. Hence computing B^* is reduced to the following:

THEOREM 36. *Let A be the transition matrix of an absorbing chain. The stochastic completion of the fundamental part B of A can be computed in deterministic space $\log^2 n$ where B is n by n and each entry of B are rational numbers represented by a pair of n-bit binary number.*

The proof of this theorem requires several preparatory results that are interesting in their own right. Therefore we defer it to the next section. We are now ready to prove the main result (theorem 28) of this section. Although we only compute the fundamental part B^*, with a bit more work, one can compute all the remaining entries of the stochastic closure in the same complexity bounds (see [6]).

Proof of main result (theorem 28). Basically the proof amounts to reducing a space-bounded probabilistic computation to computing the stochastic closure of the fundamental part of an absorbing Markov chain.

Let M be a probabilistic machine accepting in space $s(n)$. We analyze the probability of M accepting an input w by considering the Markov chain whose (Markov) states correspond to those configurations of M on w using space at most $s(|w|)$. We introduce an extra Markov state. Number these Markov states from 1 to r, where we may assume that Markov state 1 is the initial configuration on input w, and r is the extra Markov state (viewed as a NO-configuration of M). The corresponding transition matrix is $A = (p_{i,j})_{i,j=1}^{r}$ where

$$p_{i,j} = \begin{cases} \frac{1}{2} & \text{if configuration } i \text{ non-uniquely derives configuration } j \\ & \quad \text{(i.e., } i \vdash (j,k) \text{ for some } k \neq j) \\ 1 & \text{if either configuration } i \text{ uniquely derives } j, \text{ i.e., } i \vdash (j,j) \\ & \quad \text{or if } i = j \text{ and } i \text{ is terminal} \\ 0 & \text{else.} \end{cases}$$

This is not quite all: how shall we treat state i if $i \vdash (j,k)$ where j and/or k uses more than $s(|w|)$ space? Now assign for such an i, $p_{i,r} = \frac{1}{2}$ or 1, depending on whether one or both of j, k use more than $s(|w|)$ space.

We derive from A an absorbing Markov chain with transition matrix B as follows: say a Markov state in A is **useful** if it is has a path to a YES-state. Clearly useful states are inessential, but some inessential states may not be useful. In B, we retain all the useful states of A and also their transition probabilities among themselves. We renumber the useful Markov states from 1 to some $m - 1$ ($m < r$). In addition to the $m - 1$ useful states inherited from A, B has two essential states, m and $m + 1$. Basically, we collapse all essential YES-states into m and the remaining states in A (essential or not) are collapsed into $m + 1$. States m and $m + 1$ are both absorbing. More precisely, for each useful state $i = 1, \ldots, m - 1$, if the sum of the transition probabilities into the YES-states is p then we set the $(i, m+1)$th entry $[B]_{i,m} := p$. If the sum of the transition probabilities from i to the non-YES and non-useful states is q then we make $[B]_{i,m+1} = q$. Also, we have

$$[B]_{m,m} = [B]_{m+1,m+1} = 1$$

We do one more small transformation: let now C be the matrix that is identical to B except that

$$[C]_{m,m} = 0, \quad [C]_{m,m+1} = 1.$$

So state m is now a transient state. For future reference, call C the *reduced transition matrix* (for input w). If D is the fundamental part of C (obtained by deleting the last row and last column of C) then by theorem 36, we

can compute the stochastic completion D^* in $O(\log^2 m)$ space. Now m is $O(1)^{s(|w|)}$ and hence $O(\log^2 m) = O(s^2)$ space suffices.

Our 'interpretation' (lemma 31) of the entries of a stochastic completion suggests that the entry $[D^*]_{1,m}$ is the probability that starting out in state 1 we reach m (since state m cannot return to itself in 1 or more steps, by construction). It is instructive to carry out the proof that $[D^*]_{1,m}$ is indeed the least fixed point value $Val_\Delta(w)$ where Δ is the set $\{1, \ldots, r\}$ of configurations that uses space at most $s(|w|)$. A valuation V on Δ amounts to an r-vector $V = (v_1, \ldots, v_r)$ where v_i is the value of configuration i. (We may assume here that the values v_i are real numbers in $[0, 1]$ rather than intervals.) The valuation operator τ_Δ is the linear transformation given by the transition matrix A, and $\tau_\Delta(V)$ is equal to $A \cdot V^T$ (V^T is the column vector obtained by transposing V). Let $V_0 = \tau_\Delta(V_\perp)$ be the row vector that assigns a 1 to the YES state and a zero to the NO state. (Note that V_0 is not stochastic in general.) The valuation $\tau^n_\Delta(V_0)$ is given by $V_n = A^n \cdot V_0^T$. We conclude: *$Val_\Delta(w)$ is the limiting value of the first component of $A^n \cdot V_0^T$, as $n \to \infty$.* Alternatively, if the set of YES-configurations are $S \subseteq \Delta$, then $Val_\Delta(w)$ is the limiting value of $\sum_{i \in S}[A^n]_{1,i}$.

It is not hard to see that our transformation of A to B does no harm and we have a slightly simpler picture: $Val_\Delta(w)$ is given by the limiting value of $[B^n]_{1,m}$ (Exercise).

The transformation from B to C is less obvious. Let us compare their nth powers, B^n and C^n, for each $n \geq 0$. The first $m-1$ columns of both powers are seen to be identical. We claim that the first $m-1$ entries in the mth column of B^n is equal to the corresponding entry in the sum $\sum_{i=1}^n C^i$: *for each $j = 1, \ldots, m-1$, and for all $n \geq 0$,* $[B^n]_{j,m} = \sum_{\ell=1}^n [C^\ell]_{j,m}$. In proof, this is clearly true for $n = 1$. For $n \geq 1$, we have

$$\begin{aligned}
[B^{n+1}]_{j,m} &= \sum_{k=1}^{m} [B^n]_{j,k}[B]_{k,m} \\
&= [B^n]_{j,m}[B]_{m,m} + \sum_{k=1}^{m-1} [B^n]_{j,k}[B]_{k,m} \\
&= [B^n]_{j,m} + \sum_{k=1}^{m-1} [C^n]_{j,k}[C]_{k,m} \\
&= [B^n]_{j,m} + [C^{n+1}]_{j,m} \\
&= \sum_{\ell=1}^{n+1} [C^\ell]_{j,m}
\end{aligned}$$

This essentially gives us the theorem.

There are several other details that we must defer to the next section: in particular we cannot afford to explicitly store the matrices A or C. Instead, they are represented 'procedurally' in the sense that each entry of such matrices can be obtained by invoking suitable subroutines. For instance, this means that our 'application' of theorem 36 is really not in the form as stated. We need to show that the techniques implicit in that theorem can be modified to accomodate the implicit representation of C. Another clarification is needed: to form matrix C, we need to determine the useful states in A (one efficient way to detect such states uses the the original Savitch's theorem technique). Modulo these details, we are done with the proof.

Exercises

Exercise 0.28: (1-dimensional random walk)

Analyze the 1-dimensional random walk with parameter $0 < p < 1$.
(i) Show that the generating functions $G(s)$ for the probability $p_{0,0}^{(n)}$ $(n = 0, 1, \ldots)$ of returning to the origin in n steps is given by $G(s) = (1 - 4pqs^2)^{1/2}$ where $q = 1 - p$.
(ii) Conclude that each Markov state is recurrent if and only if $p = \frac{1}{2}$.
(iii) In case $p = \frac{1}{2}$, show that the mean recurrence time is infinite. **Hint:** Use the relation that the generating function $F(s)$ for first re-entrance probability $f_{0,0}^{(n)}$ is related to $G(s)$ by $G(s) - 1 = F(s)G(s)$ and the mean recurrence time is given by $\lim_{s \to 1^{-1}} \frac{dF(s)}{ds}$. ♢

Exercise 0.29: (Erdös, Feller, Pollard)

Let $(f_0, f_1, \ldots)$ be a stochastic vector with $f_0 = 0$ and period is equal to 1. (The period is the largest d such that $f_n > 0$ implies d divides n.) Now define $p_0 = 1$ and $p_n = \sum_{k=0}^n f_k p_{n-k}$. Prove that $\lim_{n \to \infty} p_n = \frac{1}{\mu}$ where $\mu = \sum_{n=0}^\infty n f_n$. **Hint:** Note that the relation between pair of sequences $\{f_n\}$ and $\{p_n\}$ is identical with

the relation between the first entrance probabilities $f_i^{(n)}$ and n-step transition probabilities $p_{i,i}^{(n)}$ for a fixed state i in section 3. ♢

Exercise 0.30: Define a transition matrix $A = (p_{i,j})_{i,j\geq 0}$ to be *doubly stochastic* if the row sums as well as column sums are equal to 1. Show that each space-bounded computation of a reversible probabilistic Turing machine gives rise to such a matrix. Study the properties of such machines.

♢

End Exercises

8.7 Efficient Circuits for Ring Operations

By an *algebraic structure* we means a set A together with a finite set of partial functions f_i $(i = 1, \ldots, k)$ of the form

$$f_i : A^{\alpha(i)} \to A$$

where $\alpha(i) \geq 0$ are integers called the *arity* of f_i. We call f_i a *constant* if $\alpha(i) = 0$ and in this case, f_i is identified with an element of A. We write $(A; f_1, \ldots, f_k)$ for the algebraic structure. This is abbreviated to 'A' when the functions f_i are understood. In general, by an *operation f over an algebraic structure A* we mean a partial function $f : A^m \to A$ for some $m \geq 0$, where m is called the arity of f.

Example 2. a) Of course, for any set A, there is the trivial algebraic structure on A which no functions at all.
b) The integers $\mathbb{Z} = \{0, \pm 1, \pm 2, \ldots\}$ with the usual operations ($+$, $-$, $\times$ and $0, 1$) is an algebraic structure. It is, in fact, a unitary commutative ring (see below).
c) The rational numbers $\mathbb{Q}$, with the operations of $\mathbb{Z}$ but also including the division $\div$ operation, is an algebraic structure called a commutative field. Here division is a partial function since division by zero is not defined.
d) The set of all n-square matrices with entries from $\mathbb{Q}$ forms a matrix ring with the usual matrix operations of $+$, $-$ and $\times$.
e) Consider the Boolean algebra on two elements $(\{0, 1\}; \vee, \wedge, \neg, 0, 1)$ where $\vee, \wedge, \neg$ are interpreted as the usual Boolean operations.
f) A class of finite rings is $\mathbb{Z}_n = \{0, 1, \ldots, n-1\}$ with usual arithmetic operations modulo n. In case n is a prime, $\mathbb{Z}_n$ is a field, also called $GF(p)$. The case $GF(2)$ has special interest. ■

We are mainly interested in computing over various unitary commutative rings R, henceforth simply called 'rings'[13]. Our goal is to show how operations in common rings can be implemented efficiently. A computational model that is appropriate for algebraic structures is circuits.

The following definitions gather most of our notations and terminology related to circuits in one place. It will serve as reference beyond just the immediate concern of this section.

Definition 5. Let $(A; f_1, \ldots, f_k)$ be an algebraic structure.

(i) Any set Ω of operations over A obtained by functional composition from $f_1, \ldots, f_k$ is called a *basis* of A. In general Ω may have infinite cardinality.

(ii) A *circuit C for A over the basis Ω* a finite directed acyclic graph (called the *underlying graph* of C) with the following properties: A node with indegree 0 is called an *input* node, and if there are n input nodes, we label each with a distinct integer between 1 and n. The remaining nodes are called *gates* and are each labeled by a basis function $f \in \Omega$. If a gate is labeled by f, we say its *type* is f or, equivalently, it is called an *f-gate*. Each f-gate u has indegree exactly equal to the arity $\alpha(f)$ of f. Furthermore the incoming edges to u are labeled with distinct integers between 1 and $\alpha(f)$. So we may speak of the jth *incoming edge* of u for $j = 1, \ldots, \alpha(f)$.

[13]Commutative rings are simply algebraic structures satisfying certain axioms. The student unfamiliar with rings simply need remember two main examples of such structures given above: the integers $\mathbb{Z}$ and the set of n-square matrices with rational number entries. So a ring comes with the five total operation $+, -, \times, 0, 1$ with the usual properties (inverse relation between plus and minus, associativity, commutativity, distributivity, and properties of 0 and 1) are satisfied. If one writes down these properties, they would constitute an axiomatization of unitary commutative rings (it is a good exercise to try this and compare your results with a standard algebra book). Here 'unitary' serves to to warn that, in general, rings are defined without assuming the existence of element 1.

(iii) Each node u of a circuit C is said to *compute* the function

$$\texttt{res}_C(u) \ : \ A^n \to A$$

defined as follows: an input node u labeled by i computes the projection function $\texttt{res}_C(u)(x_1, \ldots, x_n) = x_i$. An f-gate u computes the function

$$\texttt{res}_C(u)(\overline{x}) = f(\texttt{res}_C(u_1)(\overline{x}), \ldots, \texttt{res}_C(u_m)(\overline{x}))$$

where $\overline{x} = (x_1, \ldots, x_n)$ and the ith incoming edge of u leads from node u_j $(j = 1, \ldots, m)$, m is the arity of f.

(iv) A *circuit family* (over the basis Ω) is an infinite sequence of circuits $\overline{C} = (C_n)_{n=0}^{\infty}$ such that each C_n is an n-input circuit over Ω.

(v) A *problem instance of size n* (over A) is a set P_n of functions $g \ : \ A^n \to A$ (so each $g \in P_n$ has arity n). An *aggregate problem* $P = (P_n)_{n=0}^{\infty}$ over A is an infinite sequence of problem instances P_n, each P_n of size n. When no confusion arises, we may omit the qualification 'aggregate'. Often, $P_n = \{f_n\}$ is a singleton set, in which case we simply write $P = (f_n)_{n \geq 0}$.

(vi) Let P_n be a problem instance of size n over A. A circuit C over A is said to *solve* or *realize* P_n if if C has n inputs and for each $g \in P_n$, there is a node $u \in C$ such that $\texttt{res}_C(u) = g$. A circuit family $\overline{C} = (C_n)_{n=0}^{\infty}$ is said to *solve* a problem $P = (P_n)_{n=0}^{\infty}$ if C_n solves P_n for each n.

(vii) The *size* of a circuit C is the number of gates in the circuit[14]. The *size* of a circuit family $\overline{C} = (C_n)_{n=0}^{\infty}$ is the function $SIZE_{\overline{C}}$ where $SIZE_{\overline{C}}(n)$ is the size of C_n.

(viii) Two other complexity measures for circuits are as follows: the *depth* of C is the length of the longest path in C. The *width* of C is the maximum cardinality of an edge anti-chain[15] in C. As for the size-measure, we let $DEPTH_{\overline{C}}(n)$ and $WIDTH_{\overline{C}}(n)$ denote the depth and width of C_n, where $\overline{C} = (C_n)_{n \geq 0}$.

(ix) For any problem instance P_n, let $SIZE(P_n)$ denote the smallest sized circuit that realizes P_n. If $P = (P_n)_{n \geq 0}$ is an aggregate problem, the size function $SIZE_P$ is the function given by $SIZE_P(n) = SIZE(P_n)$. Similarly for

$$DEPTH(P_n), WIDTH(P_n), DEPTH_P, WIDTH_P.$$

(x) For any complexity function $f(n)$, let $SIZE(f)$ denote the family of aggregate problems $\{P \ : \ \forall n, SIZE_P(n) \leq f(n)\}$. Similarly for $DEPTH(f)$, $WIDTH(f)$. We can extend this to simultaneous measures, for instance $SIZE - DEPTH - width(f_1, f_2, f_3)$.

(xi) For any non-negative integer k, a circuit family $\overline{C}$ is said to be NC^k if $SIZE_{\overline{C}}(n) = n^{O(1)}$ and $DEPTH_{\overline{C}}(n) = O(\log^k n)$. An aggregate problem is said to be NC^k if it can be realized by an NC^k circuit family.

■

Remark: We are often interested in problems for which there is really no problem instances of size n for certain values of n (e.g., multiplying square Boolean matrices only has interesting input sizes of the form $2n^2$). In these cases, we artificially create the trivial problem instance of size n, such as the identically zero function of arity n. Also, the above definition of circuits do not allow constant values as inputs. The definition can trivially be changed to accommodate this.

We are mainly interested in circuits for two types of algebraic structures: (a) where A is a ring and (b) where $A = \{0, 1\}$ is the Boolean algebra in the above example. We call a circuit for A an *arithmetic circuit* or a *Boolean circuit* in cases (a) or (b), respectively.

[14] It is unfortunate that we have to use the word 'size' for problem instances as well as for circuits, both of which appear in the same context. Since the usage is well accepted and there seems to be no better alternative, we will continue this usage. But for emphasis, we could say 'circuit size' or 'problem size'.

[15] An *edge anti-chain* in a directed acyclic graph is a set of edges such that no two edge in the set belongs to a common path. Of course, one can define *node anti-chain* as well.

8.7.1 The Parallel Prefix Problem.

We begin with a basic but important technique from Ladner and Fischer [19] for the so-called *parallel prefix problem*. In this problem, we assume that we are given an algebraic structure $(A; \circ)$ where the only operation $\circ$ is a binary associative operation (which we will call 'multiplication' or 'product'). As is usual with multiplicative notations, when convenient, we may write xy and $\prod_{i=1}^{n} x_i$ (respectively) instead of $x \circ y$ and $x_1 \circ x_2 \circ \cdots \circ x_n$. We call a circuit over such an A a *product circuit*. The parallel prefix problem instance (of size $n \geq 0$) amounts to computing the set of n functions

$$f_i(x_1, \ldots, x_n) := x_1 \circ x_2 \circ \cdots \circ x_i, \quad \text{for each } i = 1, \ldots, n.$$

We may call these f_i's the set 'iterative-$\circ$ functions' (on n variables). We shall apply this in two cases: where $\circ$ is addition and where $\circ$ is multiplication in a ring. Then, we call parallel prefix the *iterative addition* and *iterative multiplication* problems, respectively.

LEMMA 37. *There is a recursive construction of a family of product circuits $(C_n)_{n=0}^{\infty}$ of linear size and logarithmic depth that solves the parallel prefix problem.*

Proof. We may assume that n is a power of 2. C_1 is trivial, consisting of just an input node for x_1. So let $n > 1$. The following figures shows the recursive construction of C_n from $C_{n/2}$:

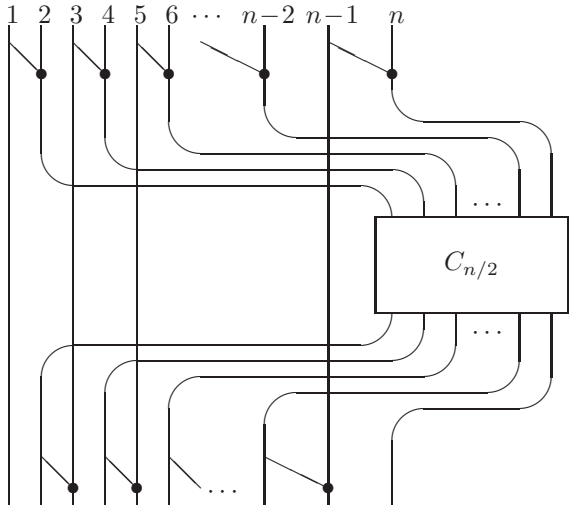

Figure 8.1 Construction of C_n ('•' indicates a gate)

In this figure, the edges of the circuit are implicitly directed downwards. It is easy to see that the size $s(n)$ of a circuit C_n satisfies the relation $s(1) = 1$ and $s(n) = s(n/2) + n - 1$. The solution is $s(n) = 2n - \log n - 2$. The depth of C_n is also easily seen to be $2 \log n$. **Q.E.D.**

8.7.2 Detecting Useless Markov States.

Recall in the previous section, in the context of a Markov chain whose nodes are machine configurations, we define a Markov state to be useless if it cannot reach a YES-configuration. To detect such useless states, we can formulate the following general problem: given the adjacency matrix of a digraph G, we want to determine its "useless nodes", defined to mean that those nodes that cannot reach a distinguished node in G.

LEMMA 38. *There is an NC^2 Boolean circuit family which computes, on any input n by n matrix A_n which represents the adjacency matrix of directed graph, the set of 0/1-functions $\{f_i : i = 1, \ldots, n\}$ where $f_i(A_n) = 1$ iff i cannot reach node n.*

Proof. This is rather straight forward: we can compute the products of Boolean matrices in NC^1. The transitive closure of A_n is given by $(A_n)^m$ for any $m \geq n$. By the usual doubling method, the transitive closure is obtained by $\log n$ matrix multiplications, hence by NC^2 circuits. Finally a node i can reach node n if and only if $A^*(i, n) = 1$. **Q.E.D.**

8.7.3 Computing the characteristic polynomial.

The determinant of an $n \times n$ matrix $A = (a_{i,j})$ can be expanded by its ith row (for any i) in the standard fashion:

$$\det(A) = a_{i,1}D_{i,1} + a_{i,2}D_{i,2} + \cdots + a_{i,n}D_{i,n}$$

where $(-1)^{i+j}D_{i,j}$ is the determinant[16] of the $(n-1)$-square matrix obtained by deleting the ith row and the jth column of A. $D_{i,j}$ is called the (i,j)-*cofactor* (or (i,j)-*complement*) of A. Let the *adjoint* $\mathtt{adj}(A)$ of A be the $n \times n$ matrix whose (i,j)th element $[\mathtt{adj}(A)]_{i,j}$ is the (j,i)-cofactor of A (notice the transposition of subscripts i and j). It is not hard to see that[17] that the following is valid for all A:

$$A \cdot \mathtt{adj}(A) = \mathtt{adj}(A) \cdot A = \det(A) \cdot I$$

(We only have to see that the off-diagonal elements of $A \cdot \mathtt{adj}(A)$ are of the form

$$a_{i,1}D_{j,1} + a_{i,2}D_{j,2} + \cdots + a_{i,n}D_{j,n}$$

where $i \neq j$. But this sum is also seen to be zero since it is the determinant of the singular matrix obtained by replacing the jth row of A with the ith row. On the other hand, the diagonal entries are equal to $\det(A)$, which may be zero if A is singular.) The *characteristic polynomial* $P_A(x)$ of A is the determinant

$$\begin{aligned} P_A(x) &:= \det(xI_n - A) \\ &= x^n + \lambda_1 x^{n-1} + \cdots + \lambda_{n-1}x + \lambda_n \end{aligned}$$

where I_n is the $n \times n$ identity matrix. (We will omit the subscript in I_n when convenient.) A fundamental identity is the Cayley-Hamilton theorem (Exercise) that states that A is a root of the polynomial $P_A(x)$

$$P_A(A) = A^n + \lambda_1 A^{n-1} + \cdots + \lambda_{n-1}A + \lambda_n I = 0.$$

Our goal is to develop a space-efficient algorithm for computing the characteristic polynomial. To compute the characteristic polynomial of A means to determine the above coefficients $\lambda_1, \ldots, \lambda_n$. Note that this computation is a generalization of the problem of computing determinant since the constant term in P_A is (up to sign) equal to $\det(A)$. We present an efficient parallel implementation of Samuelson's method[18] by Berkowitz [4]. For this, we use the notation $red(A)$ ('reduction' of A) which is the $(n-1)$-square matrix obtained from A by deleting the first row and first column. Define the column $(n-1)$-vector $\mathtt{col}(A)$ and row $(n-1)$-vector $\mathtt{row}(A)$ so that

$$A = \begin{pmatrix} a_{1,1} & \mathtt{row}(A) \\ \mathtt{col}(A) & \mathtt{red}(A) \end{pmatrix}$$

LEMMA 39. *The characteristic polynomial of A is related to the matrix $\mathtt{red}(A)$ as follows:*

$$P_A(x) = (x - a_{1,1}) \det(xI_{n-1} - \mathtt{red}(A)) - \mathtt{row}(A) \cdot \mathtt{adj}(xI_{n-1} - \mathtt{red}(A)) \cdot \mathtt{col}(A). \tag{22}$$

Proof.

$$\begin{aligned} P_A(x) &= \det(xI_n - A) \\ &= (x - a_{1,1}) \det(xI_{n-1} - A) + \sum_{j=2}^{n} a_{1,j}D_{1,j} \end{aligned}$$

where $D_{1,j}$ is the $(1,j)$-cofactor of $xI_n - A$. Write

$$D[i_1, i_2, \ldots; j_1, j_2, \ldots]$$

[16] This determinant is called the (i,j)-*minor* of A

[17] See [11] for most basic results on matrices.

[18] The importance of this method (as opposed to an earlier method of Csanky that has comparable complexity) is that it uses no divisions, so it is applicable to any unitary commutative ring.

for the determinant of the square matrix obtained by deleting rows $i_1, i_2, \ldots$ and columns $j_1, j_2, \ldots$. The lemma follows from the following identity

$$\begin{aligned}
\sum_{j=2}^{n} a_{1,j} D_{1,j} &= \sum_{j=2}^{n} a_{1,j}(-1)^{1+j} D[1;j] \\
&= \sum_{j=2}^{n} a_{1,j}(-1)^{1+j} \sum_{i=2}^{n} a_{i,1}(-1)^{i} D[1,i;1,j] \\
&= -\sum_{j=2}^{n}\sum_{i=2}^{n} a_{1,j} a_{i,1}(-1)^{i+j} D[1,i;1,j] \\
&= -\sum_{j=2}^{n}\sum_{i=2}^{n} a_{1,j} a_{i,1} D'_{i,j}
\end{aligned}$$

where $D'_{i,j}$ is the $(i-1, j-1)$-cofactor of $\mathtt{red}(xI_n - A) = xI_{n-1} - \mathtt{red}(A)$. Since $D'_{i,j}$ is the $(j-1, i-1)$th entry of $\mathtt{adj}(xI_{n-1} - \mathtt{red}(A))$, the lemma follows. **Q.E.D.**

The adjoint of $xI - A$ can be expressed with the help of the the next lemma.

LEMMA 40.

$$\mathtt{adj}(xI - A) = \sum_{i=0}^{n-1} B_i x^{n-1-i} \tag{23}$$

where

$$B_i = A^i + \lambda_1 A^{i-1} + \cdots + \lambda_{i-1} A + \lambda_i I$$

and λ_i are the coefficients of the characteristic polynomial

$$P_A(x) = x^n + \lambda_1 x^{n-1} + \cdots + \lambda_{n-1} x + \lambda_n.$$

Proof. We observe that $B_0 = I_n$ and for $i = 1, \ldots, n$,

$$B_i = AB_{i-1} + \lambda_i I.$$

Hence

$$\begin{aligned}
&\det(xI - A) \cdot I \\
&= [x^n + \lambda_1 x^{n-1} + \cdots + \lambda_{n-1} x + \lambda_n] \cdot I \\
&= [x^{n-1} B_0](xI - A) + x^{n-1} AB_0 + [\lambda_1 x^{n-1} + \lambda_2 x^{n-2} + \cdots + \lambda_{n-1} x + \lambda_n] \cdot I \\
&= [x^{n-1} B_0 + x^{n-2} B_1](xI - A) + x^{n-2} AB_1 + [\lambda_2 x^{n-2} + \lambda_3 x^{n-3} + \cdots + \lambda_n] \cdot I \\
&= \cdots \\
&= \left[\sum_{i=0}^{n-1} x^{n-1-i} B_i\right](xI - A) + x^0 AB_{n-1} + \lambda_n I \\
&= \left[\sum_{i=0}^{n-1} x^{n-1-i} B_i\right](xI - A)
\end{aligned}$$

where the last equality follows from

$$x^0 AB_{n-1} + \lambda_n I = B_n = P_A(A) = 0$$

by the Cayley-Hamilton theorem. On the other hand $\det(xI - A) \cdot I$ can also be expressed as

$$\det(xI - A) \cdot I = \mathtt{adj}(xI - A) \cdot (xI - A).$$

Since $xI - A$ is nonsingular (x is a indeterminate), we can cancel $xI - A$ as a factor from the two expressions for $\det(xI - A) \cdot I$, giving the desired equality for the lemma. **Q.E.D.**

The last two lemmas show that the characteristic polynomial $P_A(x)$ of A can be computed from the characteristic polynomial $P_{\mathtt{red}(A)}(x)$ of $\mathtt{red}(A)$ and the matrix products $\mathtt{red}(A)^i$ (for $i = 1, \ldots, n-1$). Let us derive this precisely.

Let the coefficients of $P_{\mathtt{red}(A)}(x)$ be given by μ_i $(i = 0, \ldots, n-1)$ where $P_{\mathtt{red}(A)}(x) = \sum_{i=0}^{n-1} \mu_{n-1-i} x^i$. Then we see that

$$\begin{aligned}
P_A(x) &= (x - a_{1,1}) P_{\mathtt{red}(A)}(x) - \mathtt{row}(A) \cdot \left(\sum_{i=0}^{n-1} B_i x^{n-1-i} \right) \cdot \mathtt{col}(A) \\
&= (x - a_{1,1}) \sum_{i=0}^{n-1} \mu_i x^{n-1-i} + \sum_{i=0}^{n-1} x^{n-1-i} \mathtt{row}(A) \cdot \left(\sum_{j=0}^{i} (\mathtt{red}(A))^j \mu_{i-j} \right) \cdot \mathtt{col}(A) \\
&= \sum_{i=0}^{n-1} \mu_i x^{n-i} + \sum_{i=0}^{n-1} x^{n-1-i} \left(-a_{1,1}\mu_i + \sum_{j=0}^{i} \beta_j \mu_{i-j} \right) \\
&\qquad (\text{where } \beta_j = \mathtt{row} \cdot (A)(\mathtt{red}(A))^j \cdot \mathtt{col}(A)) \\
&= \sum_{j=0}^{n} \lambda_{n-j} x^j
\end{aligned}$$

where

$$\lambda_j = \begin{cases} \mu_0 & \text{if } j = 0 \\ \mu_j - a_{1,1}\mu_{j-1} + \sum_{k=0}^{j-1} \beta_k \mu_{j-1-k} & \text{if } j = 1, \ldots, n-1 \\ -a_{1,1}\mu_{n-1} + \sum_{k=0}^{n-1} \beta_k \mu_{n-1-k} & \text{if } j = n \end{cases}$$

We can write this in matrix form. Let C_0 be the following $(n+1) \times n$ matrix:

$$C_0 := \begin{pmatrix} 1 & & & & & & \\ \beta_0 - a_{1,1} & 1 & & & & & \\ \beta_1 & \beta_0 - a_{1,1} & 1 & & & & \\ \beta_2 & \beta_1 & \beta_0 - a_{1,1} & & & & \\ \vdots & & & \ddots & & & \\ \beta_{n-2} & \cdots & & & \beta_0 - a_{1,1} & 1 & \\ \beta_{n-1} & \beta_{n-2} & \cdots & & & \beta_0 - a_{1,1} & 1 \\ 0 & \beta_{n-1} & \beta_{n-2} & \cdots & & \beta_1 & \beta_0 - a_{1,1} \end{pmatrix}$$

Then we have

$$(\lambda_0, \ldots, \lambda_{n-1}, \lambda_n)^T = C_0 \cdot (\mu_0, \ldots, \mu_{n-1})^T$$

where $(\cdots)^T$ indicates matrix transpose. We repeat this procedure to express the μ_i's in terms of the characteristic polynomial of $\mathtt{red}(\mathtt{red}(A)) = \mathtt{red}^2(A)$, etc. In general, let C_i be the $(n+1-i) \times (n-i)$ matrix that reduces the characteristic polynomial of $\mathtt{red}^i(A)$ to that of $\mathtt{red}^{i+1}(A)$. The entries[19] of C_i consists of 0, 1, the diagonal elements $a_{i+1,i+1}$ of the matrix A, and elements that can be constructed from

$$\mathtt{row}(\mathtt{red}^i(A)) \cdot \mathtt{red}^{i+1}(A) \cdot \mathtt{col}(\mathtt{red}^i(A)).$$

Putting these together, we get

$$(\lambda_0, \ldots, \lambda_n)^T = C_0 C_1 \cdots C_{n-1}$$

LEMMA 41. *Given an $n \times n$ matrix A we can construct in deterministic log-space an arithmetic circuit C of depth $O(\log^2 n)$ and size $n^{O(1)}$ such that C computes (the coefficients of) $P_A(x)$.*

Proof. We proceed as follows:

1. We first construct a circuit to compute the set of polynomials

$$\{(\mathtt{red}^i(A))^j \ : \ i = 1, \ldots, n-1 \text{ and } j = 1, \ldots, i\}.$$

Note that to compute a matrix means to compute each of its entries and to compute a polynomial means to compute each of its coefficients. It suffices to show that for each i, we can compute $\{(\mathtt{red}^i(A))^j \ : \ j = 1, \ldots, i\}$

[19] In fact C_i has the property that each (j,k)th entry is equal to the $(j+1, k+1)$th entry provided, of course, that the $(j+1, k+1)$th entry is defined. This property is the defining characteristic of *Toeplitz matrices* and it is known that the multiplication of Toeplitz matrices can be done more efficiently than we care to exploit here (Exercise).

with a circuit of polynomial size and depth $O(\log^2 n)$. But this amounts to the parallel prefix computation on i copies of the matrix $\texttt{red}^i(A)$. Parallel prefix, we saw, uses an $O(\log n)$ depth product circuit. Each gate of the product circuit is replaced by an arithmetic circuit of depth $O(\log n)$ since we can multiply two $n \times n$ matrices in this depth (straightforward). Hence the overall depth is $O(\log^2 n)$.

2. Next, we compute the elements

$$\{\texttt{row}(\texttt{red}^{i-1}(A) \cdot (\texttt{red}^i(A))^j \cdot \texttt{col}(\texttt{red}^{i-1}(A) \ : \ i = 1, \ldots, n-1 \text{ and } j = 1, \ldots, i\}.$$

This takes $O(\log n)$ depth. We have now essentially computed the entries of the matrices $C_0, \ldots, C_{n-1}$.

3. We can compute the product $\prod_{i=0}^{n-1} C_i$ using a balanced binary tree of depth $O(\log n)$ to organize the computation. Each level of this binary tree corresponds to the multiplication (in parallel) of pairs of matrices. Since each matrix can be multiplied in $O(\log n)$ depth, the overall circuit depth is $O(\log^2 n)$.

One can easily verify that the circuit is polynomial in size. **Q.E.D.**

8.7.4 Computing the Matrix Inverse

We want to compute the inverse of a matrix A. As before, let

$$P_A(x) = x^n + \lambda_1 x^{n-1} + \cdots + \lambda_{n-1} x + \lambda_n$$

be the characteristic polynomial. Since $P_A(x) = \det(xI - A)$ we see that

$$\lambda_n = P_A(0) = \det(-A) = (-1)^n \det(A).$$

Next, by lemma 40, we have that

$$\begin{aligned} \texttt{adj}(-A) &= B_{n-1} \\ &= A^{n-1} + \lambda_1 A^{n-2} + \cdots + \lambda_{n-2} A + \lambda_{n-1}. \end{aligned}$$

Note that $\texttt{adj}(-A) = (-1)^{n-1}\texttt{adj}(A)$. Putting these together, we get

$$\begin{aligned} A^{-1} &= \frac{1}{\det(A)}\texttt{adj}(A) \\ &= -\frac{1}{\lambda_n}(A^{n-1} + \lambda_1 A^{n-2} + \cdots + \lambda_{n-2} A + \lambda_{n-1}). \end{aligned}$$

It should now be easy to deduce:

LEMMA 42. *We can detect if an n-square matrix A is nonsingular using an NC^2 arithmetic circuit. Moreover, in case A is nonsingular, we can compute its inverse A^{-1} with another NC^2 arithmetic circuit.*

8.7.5 Balanced p-ary Notations

In applying the preceding results, we will use matrices A whose entries are rational numbers. Assuming the usual binary representation of integers, if the integers involved in the computation are k-bits long, then an $O(\log^2 n)$ depth arithmetic circuits translate into Boolean circuits of depth $\Omega(\log^2 n \log k)$, at least (and in our applications $k = \Omega(n)$). To obtain a depth of $O(\log^2 n)$ on Boolean circuits for computing characteristic polynomials, we need one new idea: by a suitable choice of representing integers, we can implement the operations of addition with constant depth (i.e., NC^0) circuits. This in turn yields an improved depth for several other problems. Circuits in this subsection shall mean Boolean circuits unless otherwise stated.

Definition 6.
(i) A *representation* r of a algebraic structure A is an onto function

$$r \ : \ \{0,1\}^* \to A.$$

We say $u \in \{0,1\}^*$ is an *r-representative* of $r(u)$.
(ii) We say an operation $f \ : \ A^n \to A$ is in NC^k with respect to r if there is an NC^k *Boolean* circuit family $\overline{C} = (C_m)_{m \geq 0}$ such that for each m, C_m computes f in the sense that for all $u_1, \ldots, u_n \in \{0,1\}^m$,

$$r(C_m(u_1, \ldots, u_n)) = f(r(u_1), \ldots, r(u_n))$$

■

We usually like to ensure that each element can be represented by arbitrarily large binary strings. For instance, as in the binary representation of numbers, we may have the property $r(0u) = r(u)$ for all $u \in \{0,1\}^*$. In other words, we can left pad an representation by 0's. In practice, ring elements have some natural notion of "size" and we may insist that the representation r 'respect' this size function within some bounds (e.g., $|r(x)| \geq size(x)$). We shall not concern ourselves with such considerations but the interested reader may consult [6].

Our goal is to find a representation of integers so that addition and negation is in NC^0, multiplication and iterated addition is in NC^1 and iterated multiplication is in NC^2. We resort to the *balanced p-ary representation* of integers of Avizienis [1]. Here $p \geq 2$ is an integer and a *balanced p-ary number* is a finite string

$$u_1 u_2 \cdots u_n$$

where u_i is an integer with $|u_i| \leq p-1$. This string represents the integer $(u_1, \ldots, u_n)_p$ given by

$$(u_1, \ldots, u_n)_p := \sum_{i=0}^{n-1} u_{i+1} p^i.$$

So u_1 is the least significant digit. Clearly, the usual p-ary representation of a number is a balanced p-ary representation of the same number. This representation is redundant in a rather strong sense. We will implicitly assume that strings such as $u_1, \ldots, u_n$ are ultimately encoded as binary strings, so that they fit our formal definition of representations.

LEMMA 43. *With respect to the balanced p-ary representation of integers, for any $p \geq 3$:*
(i) Addition and negation of integers are in NC^0.
(ii) Iterated addition and multiplication of integers are in NC^1.
(iii) Iterated multiplication is in NC^2.

Proof. (i) Suppose we want to add $(u_1, \ldots, u_n)_p$ to $(v_1, \ldots, v_n)_p$. Note that for each $i = 1, \ldots, n$, we can express the sum $u_i + v_i$ as

$$u_i + v_i = p x_i + y_i$$

with $|x_i| \leq 1$ and $|y_i| \leq p-2$. To see this, note that $|u_i + v_i| \leq 2p-2$ and if $|u_i + v_i| \leq p-2$ or $|u_i + v_i| \geq p$ then it is clear that the desired x_i, y_i can be found. The remaining possibility is $|u_i + v_i| = p-1$. In that case we could let $x_i = 1$ and $y_i = \pm 1$, but note that this is possible only because $p \geq 3$ (we would violate the constraint $|y_i| \leq p-2$ if $p = 2$). Now the sum of the two numbers is given by $(w_1, \ldots, w_n w_{n+1})_p$ where

$$w_i = x_{i-1} + y_i$$

for $i = 1, \ldots, n+1$ (taking $x_0 = y_{n+1} = 0$). Clearly this can be implemented by an NC^0 circuit.
(ii) To show iterated addition is in NC^1 we simply use part (i) and the technique for parallel prefix.

Now consider multiplication of integers. Suppose we want to form the product of the numbers $(u_1, \ldots, u_n)_p$ and $(v_1, \ldots, v_n)_p$. For each $i, j = 1, \ldots, n$, we form the product $u_i v_j$ and we can express this in the form

$$u_i v_j = p x_{i,j} + y_{i,j}$$

where $|x_{i,j}| \leq p-1$ and $|y_{i,j}| \leq p-1$ (since $|u_i v_j| \leq (p-1)^2 \leq p(p-1) + (p-1)$). For each $i = 1, \ldots, n$, form the number

$$X_i = (00 \cdots 00 x_{i,1} x_{i,2} \cdots x_{i,n-1} x_{i,n})_p$$

where X_i has a prefix of i zeroes. Similarly, form the number

$$Y_i = (00 \cdots 00 y_{i,1} y_{i,2} \cdots y_{i,n-1} y_{i,n})_p$$

where Y_i has a prefix of $i-1$ zeroes. It is then easy to see that the product is given by the sum

$$\sum_{i=1}^{n} (X_i + Y_i).$$

But each summand has at most $2n$ digits and there are $2n$ summands. We can form a balanced binary tree T on $2n$ leaves to organize this summation process: each leaf is labeled with one of these summands and each interior node is labeled with the sum of the labels at the leaves below. Clearly the root of T has the desired product. This tree converts into a Boolean circuit of depth $O(\log n)$. This shows multiplication is in NC^1.
(iii) We leave this as exercise. **Q.E.D.**

Next we extend the lemma to matrix rings. By the *balanced p-ary representation of matrices with integer entries* we mean that each entry is encoded by balanced p-ary notation, and matrices are stored (to be specific) in row-major order.

LEMMA 44. *With respect to the balanced p-ary representation ($p \geq 3$) of integer matrices:*
(i) Addition and negation are in NC^0.
(ii) Multiplication is in NC^1.
(iii) Interated multiplication is in NC^2.

Proof. It is clear addition and negation of integer matrices can be implemented efficiently by the previous lemma. For multiplication, suppose we want to compute the product of two $n \times n$ matrices A and B, and each entry has at most n bits. We note that each entry of AB is the sum of at most n products of pairs of entries. These individual products can be viewed at the sum of at most $2n$ numbers of n bits, as revealed in the proof of the previous lemma. So for each entry, we need to sum $O(n^2)$ numbers, each of n bits. Again, we can arrange these as a balanced binary tree of depth $O(\log n)$. This gives us the efficient NC^1 circuit we seek.

To get an NC^2 circuit family for iterated multiplication of integer matrices we simply apply parallel prefix to the previous part. **Q.E.D.**

Finally we consider matrices whose entries are rational numbers. A rational number is represented by a pair of balanced p-ary representation, extended to matrices as before. Unfortunately, we no longer know how to do addition of rational numbers in NC^0. Nevertheless, we have the following:

LEMMA 45. *With respect to the balanced p-ary ($p \geq 3$) representation of matrices with rational number entries:*
(i) Iterated multiplication is in NC^2.
(ii) Characteristic polynomial computation is in NC^2.

Proof. (i) Suppose we want to compute the iterated product

$$A_1, A_1A_2, \ldots, A_1A_2 \cdots A_n$$

where each A_i is a $n \times n$ matrix with rational number entries, and each entry is represented by pairs of n-bit integers. We first convert each A_i to integer matrices B_i and compute an integer D_i such that $A_i = \frac{1}{D_i}B_i$. To do this, first form the product $D_{i,j}$ of the denominators in the jth row of A_i; then multiply each entry in the jth row of A_i by $D_{i,j}$. Doing this for all rows, we get B_i; of course D_i is just the product of all the $D_{i,j}$'s. It is clear that we can obtain each of the $D_{i,j}$ and D_i by iterated multiplication in NC^2. Notice that $D_{i,j}$ and D_i are $O(n^3)$-bit integers and so we can compute B_i from A_i and $D_{i,j}$'s in NC^1.

Next, we compute the iterated integer products $\{D_1, D_1D_2, \ldots, D_1D_2 \cdots D_n\}$ in NC^2. Similarly, we compute the iterated matrix product $\{B_1, B_1B_2, \ldots, B_1B_2 \cdots B_n\}$ in NC^2. It is clear that

$$A_1A_2 \cdots A_i = \frac{1}{D_1D_2 \cdots D_i}B_1B_2 \cdots B_i$$

for each $i = 1, \ldots, n$. This can be computed in NC^1 since each of the integer involved in polynomial in size.
(ii) We imitate the proof of lemma 41. Details are left as exercise. **Q.E.D.**

One more computational problem: we need to be able to check the sign of a balanced p-ary number. (In our application, we want to compare such a number with $\frac{1}{2}$, which is easily reduced to checking if a number is positive.) But after the preceding development, the reader should have no trouble devising an NC^1 solution (Exercise).

Putting it all together. We must tidy up the loose bits in the proof of the main theorem in the last section. In particular, we must address the issue of how to implicitly construct and represent the reduced transition matrices C described in the last section. Let A be the transition matrix from which we derive C. Since we want to do all this using $O(s^2)$ space, we cannot afford to write A or C explicitly. Then we must see how the techniques given in this section can be adapted to some implicit representation. All this is tedious but it is crucial that the reader understands how this can be done. So let us assume a given probabilistic machine M accepting in space $s(n)$, and w is an input. Let us begin by constructing matrix C: the proof of lemma 38 shows a transitive closure circuit of depth $O(\log^2 n)$ applied to the underlying graph G_A of the transition matrix A to determine the inessential states. But note that the transitive closure circuit is relatively systematic that we can assume some numbering of its gates such that given any gate number g, we can determine the gates at the other end of incoming as well as outgoing edges at g, and given g, we also know the Boolean function labeling g. The gate numbers can be stored in $O(s)$ space and we can determine these information also in $O(s)$ space. It is now not hard to see that we can determine the output of any gate in $O(s)$ space, given that we know the input graph G_A. This is not hard to do (we basically store one Boolean value at each gate along the path from the output gate to the current position – a similar proof using this technique is shown in chapter 10.) In this way, in $O(s)$ space, we can determine if any given i is inessential.

The basis of the preceding argument is the observation that the transitive closure circuit is quite systematic and hence in space $O(s)$ we can answer basic questions such as connections between gates, etc. Similarly for all the other circuits in this section. The student should carefully work out some other cases. With this, we conclude.

Remark: In Chapter 10, we will study the property stated above, that all the circuits in this section are 'systematically constructed' so that we can essentially determine gates and their interconnections in circuits efficiently. (These are called *uniformity* conditions.) For this reason we are contented with a somewhat sketchy outline here.

Exercises

Exercise 0.31:

(i) (Generalized Bezout's theorem) Let

$$F(x) = F_0 x^m + F_1 x^{m-1} + \cdots + F_m \quad (F_0 \neq 0)$$

be a matrix polynomial where each F_i is an $n \times n$ matrix. The *right value* of $F(x)$ at an $n \times n$ matrix A is given by

$$F(A) = F_0 A^m + F_1 A^{m-1} + \cdots + F_m.$$

(The *left value* $\hat{F}(A)$ is similarly obtained except that A is multiplied from the left.) Show that if $F(x)$ is divided by $xI - A$ from the left, the remainder is the right value $F(A)$. Hint: the proof is a direct long division of $F(x)$.

(ii) Let $B(x)$ be the adjoint of $xI - A$. Conclude that $(xI - A)B(x) = P_A(x)I$.

(iii) Infer the Cayley-Hamilton theorem from the Generalized Bezout's theorem. ♢

Exercise 0.32: (Fisher-Ladner) Improve the $2 \log n$ depth for the parallel prefix circuit C_n in the text. **Hint:** consider a recursive construction of a circuit C'_n as illustrated in the following figure where, with the proper wires added, we have

$$DEPTH(C'_n) = 1 + \max\{DEPTH(C_{n/2}), DEPTH(C'_{n/2})\}.$$

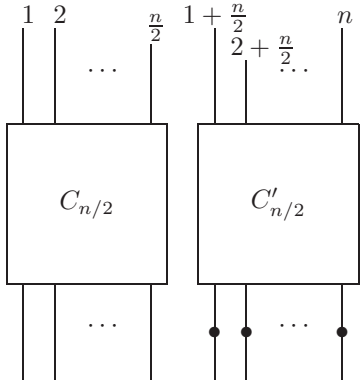

Figure 8.2 Construction of C'_n (the connections between $C'_{n/2}$ and $C_{n/2}$ not shown)

Note that the original C_n is used in this construction.

(a) Give exact expressions for the size and depth of C'_n. (Hint: the exact size of C'_n involves the Fibonacci numbers.)

(b) Compare the fan-out degree of C_n and C'_n. ♢

Exercise 0.33: (Balanced p-ary representation)

(a) Suppose we have a fixed finite state automaton M that reads an input sequence of symbols $a_1 \cdots a_n$ in real time. Let q_i be the state of M after reading a_i, starting from the start state q_0. Show an NC^1 Boolean circuit for the problem of computing $\{q_1, \ldots, q_n\}$ for any input sequence $a_1 \cdots a_n$. **Hint:** Apply parallel prefix.

(b) Apply the above to determine the sign of a balanced p-ary number in NC^1 ($p \geq 3$).

(c) Show how to convert a balanced p-ary number to a p-ary number in NC^1 (assume part (b)). ♢

Exercise 0.34: (General representation)
Let us fix $b \geq 1, p \geq 0, q \geq 1$. We say a word $u \in \{-p, -p+1, \ldots, q-1, q\}\star$ represents the integer

$$\sum_{i=0}^{n} u_i b^i \tag{24}$$

in the (b, p, q)-notation. Note that b-ary notations are $(b, 0, b-1)$-notations; b-adic notations are simply $(b, 1, b)$-notations; balanced p-ary numbers are $(p, -p, p)$-notations. Show that for $(3, 0, 3)$-notations, addition and multiplication of natural numbers (since we cannot represent negative numbers) is in NC^0 and NC^1, respectively. For what values of (b, p, q)-notations are these results true? ♢

End Exercises

8.8 Complement of Probabilistic Space

This section proves that probabilistic space is closed under complementation. We begin with a useful result:

Lemma 46. *Let B be the fundamental part of the transition matrix of a Markov chain. If B is $m \times m$ and the entries of B are taken from $\{0, \frac{1}{2}, 1\}$, then the stochastic closure B^* has the property that dB^* is an integer matrix for some integer $d < m!2^m$.*

Proof. We know that $B^* = (I - B)^{-1}$. Thus $B^* = \frac{1}{\det(I-B)}\mathtt{adj}(I - B)$. Clearly each entry of $2^{m-1}\mathtt{adj}(I - B)$ is an integer. Also, $\frac{c}{\det(I-B)}$ is an integer for some $c \leq m!$. The result follows. **Q.E.D.**

Theorem 47. *Let* M *be any probabilistic machine that accepts in space $s(n) \geq \log n$. Then there is a $c > 0$ such that every accepted input of length n is accepted with probability more than*

$$\frac{1}{2} + 2^{-c^{s(n)}}.$$

Proof. At the end of section 3, we showed that the probability of accepting any input is given by a suitable entry of D^*, where D is the fundamental part of a reduced transition matrix C. D is $m \times m$ with $m = nO(1)^{s(n)} = O(1)^{s(n)}$. We then apply the previous lemma. **Q.E.D.**

Lemma 48. *Let $s(n) \geq \log n$. For any probabilistic machine that runs in space $s(n)$, there is a probabilistic machine* $\mathbb{N}$ *accepting* L(M) *and runs in space s with error gap $(0, \frac{1}{2}]$. Morever,* $\mathbb{N}$ *halts with probability* 1 *and has average time $2^{2^{O(s)}}$.*

Proof. Let $c = \max\{c_1, c_2\}$ where c_1 is chosen so that there are at most c_1^s configurations using space at most s, and c_2 is chosen so that (by the previous theorem) if M accepts an input w then M accepts with probability greater than $\frac{1}{2} + 2^{-c_2^{s(|w|)}}$. Now $\mathbb{N}$ is obtained by modifying the proof of lemma 16 in the last section:

repeat forever
1. Simulate M for another c^s steps. Nondeterministic choices of M become coin-tossing choices of $\mathbb{N}$.
2. If M answers YES then we answer YES.
   If M answers NO, we rewind our tapes to prepare for a restarted simulation.
3. Toss $2c^s$ coins and answer NO if all tosses turn up *heads*.

end

(Notice that unlike in lemma 16, each iteration of the loop continues the simulation from where the last iteration left off, provided the last iteration did not halt.) Now we see that $\mathbb{N}$ halts with probability 1. The average time $\bar{t}$ is seen to satisfy the bound

$$\bar{t} \leq 3c^s + (1 - 2^{-2c^s})\bar{t}$$

which gives us $\bar{t} = 2^{2^{O(s)}}$ as desired.

If M rejects then clearly $\mathbb{N}$ rejects. So assume that M acccepts. Let us call a configuration C of M *live* if there is a computation path starting from C into a YES configuration. Similarly, a configuration C of N is *live* if there is a computation path from C into one in which the simulated M answers YES. If a configuration is not alive, we say it is *dead*.

For $k \geq 1$, define the following events for the probabilistic spaces $\Omega_w^{\mathbb{N}}$ and Ω_w^{M}:

$$\begin{array}{lcl} R_k^{\mathbb{N}} & = & \{\text{N answers NO in the } k\text{th iteration}\} \\ D_k^{\mathbb{N}} & = & \{\text{N became dead during the } k\text{th iteration}\} \\ D_k^{\mathrm{M}} & = & \{\text{M became dead between the } (k-1)c^s\text{th and the } kc^s\text{th step}\} \\ A_k^{\mathbb{N}} & = & \{\text{N is alive at the end of the } k\text{th iteration}\} \\ A_k^{\mathrm{M}} & = & \{\text{M is alive at the end of the } kc^s \text{ step}\} \end{array}$$

Then the rejection event of N corresponds to

$$\begin{array}{rcl} \bigcup_{k\geq 1}\{R_k^{\mathbb{N}}\} & = & \bigcup_{k\geq 1}\left(\{R_k^{\mathbb{N}}, A_k^{\mathbb{N}}\} \cup \bigcup_{j=1}^{k}\{R_k^{\mathbb{N}}, D_j^{\mathbb{N}}\}\right) \\ & = & (\bigcup_{k\geq 1}\{R_k^{\mathbb{N}}, A_k^{\mathbb{N}}\}) \cup (\bigcup_{j\geq 1}\bigcup_{k\geq j}\{R_k^{\mathbb{N}}, D_j^{\mathbb{N}}\}) \end{array}$$

Clearly the probability that M remains alive after c^s steps is at most $1-e$ where

$$e := 2^{-c^s}.$$

Since the probability of getting $2c^s$ heads in a row is e^2, the probability that N remains alive through one iteration is at most $(1-e)(1-e^2)$. We claim that

$$\Pr\{R_k^{\mathbb{N}}, A_k^{\mathbb{N}}\} = (1-e)^k(1-e^2)^{k-1}e^2.$$

(In this proof, it is instructive to set up the connection between the probabilistic spaces Ω_w^N and Ω_w^M for input w.) Hence

$$\begin{array}{rcl} \Pr(\bigcup_{k\geq 1}\{R_k^{\mathbb{N}}, A_k^{\mathbb{N}}\}) & = & \sum_{k\geq 1}\Pr\{R_k^{\mathbb{N}}, A_k^{\mathbb{N}}\} \\ & \leq & e^2\sum_{k\geq 1}(1-e)^k(1-e^2)^{k-1} \\ & < & e^2\sum_{k\geq 1}(1-e)^{k-1} \\ & = & e. \end{array}$$

Next,

$$\begin{array}{rcl} \Pr(\bigcup_{j\geq 1}\bigcup_{k\geq j}\{R_k^{\mathbb{N}}, D_j^{\mathbb{N}}\}) & \leq & \Pr(\bigcup_{j\geq 1}\{D_j^{\mathbb{N}}\}) \\ & = & \sum_{j\geq 1}\Pr\{D_j^{\mathbb{N}}\} \\ & \leq & \sum_{j\geq 1}\Pr\{D_j^{\mathrm{M}}\} \\ & = & \Pr\{\text{M rejects}\} \\ & \leq & \frac{1}{2} - e \end{array}$$

by our choice of the constant $c \geq c_2$. Above we have used the inequality $\Pr\{D_j^{\mathbb{N}}\} \leq \Pr\{D_j^{\mathrm{M}}\}$ relating across two different probability spaces. Hence the probability that N rejects is less than the sum of $e + (\frac{1}{2} - e)$. We conclude that N accepts with probability greater than $\frac{1}{2}$. **Q.E.D.**

The technique in this lemma can be viewed as using $m = c^s$ coin tosses to control a loop so that the expected number of iterations is 2^m. Since we need only $\log m$ space to control this loop, we are able to probabilistically

achieve a number of iterations that is double exponential in the space used. This technique, due to Gill, demonstrates one of the fundamental capabilities of coin-tossing that distinguishes space-bounded probabilism from, say, space-bounded alternating computations. The expected number of iterations is achieved in the worst case: if M is a machine that does not halt, then $\mathbb{N}$ has expected time $2^{2^{\Omega(s(n))}}$.

We are ready to show that probabilistic space is closed under complementation.

THEOREM 49. *If $s(n) \geq \log n$ is space-constructible then*

$$PrSPACE(s) = \text{co-}PrSPACE(s).$$

This result was shown by Simon [32]; the proof here is essentially from [28]. This result is almost an immediate consequence of lemma 48.

Proof. Given any probabilistic machine accepting in space $s(n)$, lemma 48 gives us another probabilistic machine accepting the same language in the same space bound with error gap $(0, \frac{1}{2}]$; moreover, M halts with probability 1. Then let N be the complement of M: i.e., N answers YES iff M answers NO. For any input w, the probability that N accepts w plus the probability that M accepts w is equal to the probability of halting, i.e., 1. Hence, N has the error gap $[\frac{1}{2}, 1)$. It follows that N accepts if and only if M rejects. **Q.E.D.**

8.9 Stochastic Time and Space

In this section, we give upper bounds on the complexity of time and space-bounded stochastic computations. Stochastic space is especially interesting in view of the tremendous computational power that seems inherent in it. Also, instead of Markov chains we now turn to the study of discrete time dynamical systems.

THEOREM 50. *For all t, StA-$TIME(t) \subseteq ATIME(t^3)$.*

The basic idea for this result is the bit-counting technique that was used quite effectively for simulating probabilistic alternating machines (chapter 7, section 6). It turns out that several new technical problems arises.

Let M be a stochastic alternating machine that accepts in time $t(n)$. We construct an alternating machine N to simulate M. For the rest of this proof, we fix an input w that is accepted by M and let $t = t(|w|)$. We may assume that N has guessed t correctly and let T_0 be the accepting computation tree of M on w obtained by truncating the complete computation tree $T_{\mathrm{M}}(w)$ at levels below t (as usual, root is level 0). To demonstrate the essential ideas, we assume that M has $\oplus$- and $\otimes$-states only. As in chapter 7 we 'normalize' the least fixed point values $Val_{T_0}(C)$ for each configuration in T_0:

$$VAL_0(C) := 2^{2^{t-\ell}}\, Val_{T_0}(C)$$

where $\ell = level(C)$. Thus T_0 is accepting if and only if $VAL_0(C_0) > 2^{2^t - 1}$ where C_0 is the root of T_0. It is easy to see that $VAL_0(C)$ is an integer between 0 and 2^{2^t}. We shall think of $VAL_0(C)$ as a 2^t digit number in the balanced 4-ary notation. Although the balanced 4-ary notation is highly redundant, we will want to refer to the 'ith digit of $VAL_0(C)$' in an unambiguous manner. We will show how to uniquely choose a balanced 4-ary representation for each $VAL_0(C)$.

Let us note that in fact $VAL_0(C)$ can be explicitly defined as follows: if $\ell = level(C)$ and $C \vdash (C_L, C_R)$ (provided C is not a leaf) then

$$VAL_0(C) \quad = \quad \begin{cases} 0 & \text{if } C \text{ is a non-YES leaf} \\ 2^{2^{t-\ell}} & \text{if } C \text{ is a YES leaf} \\ 2^{2^{t-\ell-1}-1}(VAL_0(C_L) + VAL_0(C_R)) & \text{if } C \text{ is an } \oplus\text{-configuration} \\ VAL_0(C_L) \cdot VAL_0(C_R) & \text{if } C \text{ is an } \otimes\text{-configuration} \end{cases}$$

It follows that each $VAL_0(C)$ has at most $2^{t-\ell}$ digits of significance.

The alternating simulation of N effectively constructs a tree T_1 that is an 'expansion' of T_0. To describe T_1, we need to define a certain product of trees:

Definition 7. Let T, T' be any two trees. For nodes $i, j \in T$, write $i \to j$ to mean that i is the parent of j. Their *product* $T \times T'$ consists of nodes (i, i') where $i \in T$ and $i' \in T'$ such that $(i, i') \to (j, j')$ if and only if either
(a) $i = j$ and $i' \to j'$, or
(b) $i \to j$, i' is the root of T' and j' is a leaf of T'. ■

Clearly,

$$T \times T' = T' \times T \iff T = T',$$
$$SIZE(T \times T') = SIZE(T' \times T) = SIZE(T) \cdot SIZE(T'),$$

and

$$DEPTH(T \times T') = DEPTH(T' \times T) = DEPTH(T) + DEPTH(T').$$

Define tree T_1 to be $T_0 \times T^t$ where T^t is defined as the binary tree in which every internal node has 2 children and every path from the root to a leaf has length exactly $t+1$. Hence T^t has exactly 2^{t+1} leaves. Let us assume the nodes in T^t are strings $s \in \{L, R\}^*$ such that the root of T^t is the empty string ϵ, and each internal node $s \in T^t$ has two children, sL and sR.

We want an assignment function VAL_1 on T_1 in analogy to VAL_0. Write $VAL_1(C, s)$ (instead of $VAL_1((C, s))$, which is correct but pedantic) for the value assigned to the node $(C, s) \in T_1$, $C \in T_0, s \in T^t$. Instead of integers, $VAL_1(C, s)$ is a balanced 4-ary number.

First, we need one more definition: recall that in section 5, we compute the product of two balanced p-ary numbers $u = u_1 \cdots u_n$, $v = v_1 \cdots v_n$ as the sum of $2n$ balanced p-ary numbers X_i, Y_i $(i = 1, \ldots, n)$. Let us call X_i and Y_i the ith and $(n+i)$th *summand* of the product of u and v. Clearly the summands depend on the particular representation of the numbers $(u)_p$ and $(v)_p$. These $2n$ summands can be regarded as $2n$ digits numbers although they each have at most $2n-1$ digits of significance.

Suppose C is a leaf in T_0, $level(C) = \ell$. Then for any $s \in T^t$, $VAL_1(C, s)$ is defined to be the (ordinary) 4-ary representation of

$$2^{2^{t-\ell}} \text{ or } 0$$

depending on whether C is YES or not. Next assume $C \vdash (C_L, C_R)$. There are two cases:

(1) C is a $\otimes$-configuration.

(1.1) If s is a leaf of T^t. Then s is a string in $\{L, R\}^*$ of length $t+1$. Then $VAL_1(C, s)$ is the sth summand of the product of $VAL_1(C_L, \epsilon)$ and $VAL_1(C_R, \epsilon)$. Here s is interpreted as the 2-adic number (see chapter 1, section 4.2) where the symbols L, R in s are (respectively) interpreted as 1,2. Note that s ranges from 0 to 2^{t+1} and by definition the 0th summand is 0.

(1.2) If s is not a leaf of T^t then $VAL_1(C, s) = VAL_1(C, sL) + VAL_1(C, sR)$.

(2) C is a $\oplus$-configuration.

(2.1) If s is a leaf of T^t then let

$$VAL_1(C, s) = 2^{2^{t-\ell-1}-1}[VAL_1(C_L, \epsilon) + VAL_1(C_R, \epsilon)].$$

Note that we are multiplying by a power of 2 and this is relatively trivial in balanced 4-ary notation. Basically we must reexpress each balanced 4-ary digit as a pair of balanced 2-ary digit, shift these to the right by $2^{t-\ell-1} - 1$ positions. Then we recombine into balanced 4-ary digits.

(2.2) If s is not a leaf then $VAL_1(C, s) = VAL_1(C, sL) (= VAL_1(C, sR))$.

It is not hard to show that for all $C \in T_0$,

$$VAL_0(C) = VAL_1(C, \epsilon).$$

We note that VAL_1 is uniquely defined since the balanced p-ary representation at the leaves are uniquely specified, and this propagates to all other nodes using the above rules.

Now it should be easy for the reader to use the technique of chapter 7 to provide an alternating machine that guesses the tree T_1 in order to determine the predicate

$$DIGIT(C, s, i, b)$$

that is true if the ith digit of $VAL_1(C, s)$ is equal to b, for $|b| \leq 3$ and $i = 1, \ldots, 2^t$. To invoke this procedure, we may assume the work tapes of N contains the following information:

1. i in binary

2. C
3. s
4. *level*(C) in unary
5. $t - \mathit{level}(C) + |s|$ in unary

Note that each of these uses $O(t)$ space. Moreover, from the above information, we can generate the arguments for the recursive calls in $O(t)$ steps. So the total work spend along any path in T_1 is $O(t^3)$ since T_1 has $O(t^2)$ levels. It is now clear that *DIGIT* can be solved in $O(t^2)$ alternating time.

The final work to be done is to compare $VAL_1(C_0, \epsilon)$ to $\frac{1}{2}$ where C_0 is the root of T_0. Basically, our goal is to convert the balanced 4-ary number

$$VAL_1(C_0, \epsilon) = u_1 u_2 \cdots u_m \quad (m = 2^t)$$

to an ordinary 4-ary number

$$v_1 v_2 \cdots v_m.$$

We begin with a simple remark: since each of the digits v_i must be non-negative, if u_i is negative, we must borrow one unit from the next digit u_{i+1}. Let $b_i = 1$ or 0 depending on whether we need to borrow from u_{i+1} or not in order to make v_i non-negative. Of course, we must also take into account the borrow b_{i-1} from u_i. This gives us the equation

$$b_i = \begin{cases} 1 & \text{if } u_i - b_{i-1} < 0 \\ 0 & \text{if } u_i - b_{i-1} \geq 0 \end{cases}.$$

We set $b_0 = 0$. Note that this method is correct because we know that *apriori* that the number $(u_1 \cdots u_m)_4$ is non-negative: so the most significant non-zero digit is positive. It is not hard to reduce the above to: $b_i = 1$ *iff for some* j $(1 \leq j \leq i)$, $u_j < 0$ *and for all* $k = j+1, \ldots, i$, $u_k = 0$. Hence we can in $O(t^2)$ alternating time check the value of any b_i for $i = 1, \ldots, 2^t$: we simply guess j and then universally check that $u_j < 0$ and $u_k = 0$ for $k = j+1, \ldots, i$. Of course, checking if $u_k = b$ is nothing but the subroutine $\mathit{DIGIT}(C_0, \epsilon, k, b)$ which can be determined in $O(t^2)$ time.

Since we can check for the value of borrow bits b_i, we can check the digits v_i in the 4-ary representation of $VAL_1(C_0, \epsilon)$ via $v_i = u_i - b_{i-1} + 4b_i$. Now it is easy to determine if $VAL_1(C_0, \epsilon)$ is greater than $2^{2^t - 1}$.

We must address one more detail. The above construction did not consider the other basis functions of a stochastic alternating machine: $\oplus, \wedge, \vee$. However, it should be evident that since we know how to add and to multiply, we can also compute the value

$$VAL_0(C) = 2^{2^{t-\ell-1}-1}(\,VAL_0(C_L) + \,VAL_0(C_R)) - \,VAL_0(C_L)\,VAL_0(C_R)$$

where $\ell = \mathit{level}(C)$, C is a $\oplus$-configuration and $C \vdash (C_L, C_R)$. The remaining MIN- and MAX-configurations are also easy. This completes our proof of theorem 50.

Space-bounded Stochastic computation. We consider an $s(n)$ space-bounded stochastic machine M. In analogy with Markov chains, we set up a finite number of *dynamical states*, each corresponding to a configuration of the machine M using space at most s. If the set of states is taken to be $\{1, \ldots, n\}$, a valuation V can be viewed as an n-vector

$$V = (v_1, \ldots, v_n) \in [0, 1]^n.$$

Let V_0 denote the vector of zero elements. We have the usual operator $\tau = \tau_\Delta$ where $\Delta = \{1, \ldots, n\}$. Thus

$$\tau(V) = (\tau_1(V), \tau_2(V), \ldots, \tau_n(V))$$

where $\tau_i(V) = 0, 1$ or $v_j \circ v_k$ for some j, k depending on i, $\circ \in \{\oplus, \otimes, \oplus\}$. In chapter 7, we showed that the limit of the sequence

$$\tau(V_0), \tau^2(V_0), \tau^3(V_0), \ldots$$

is the least fixed point of our system. Let

$$V_\tau^* = (v_1^*, \ldots, v_n^*)$$

denote this limit. Clearly any fixed point $V = (v_1, \ldots, v_n)$ of τ satisfies the following set of equations: each $i = 1, \ldots, n$,

$$v_i = f_i(v_{j(i)}, v_{k(i)}) \tag{25}$$

where $f_i(x,y)$ is one of the stochastic functions $0, 1, x \textcircled{+} y, x \otimes y, x \oplus y$. Let $\Sigma(V)$ denote the set of equations (25). We can then characterize the least fixed point property V_τ^* as follows:

$$LFP(V_\tau^*) \equiv \Sigma(V_\tau^*) \wedge (\forall V)[\Sigma(V). \Rightarrow .V_\tau^* \leq V].$$

We are now ready to prove the following, by appeal to some results on the complexity of real closed fields. (cf. Renegar [25, 26, 27]).

THEOREM 51. *For all $s(n)$,*

$$StSPACE(s) \subseteq DTIME(2^{2^{O(s)}}).$$

Proof. Suppose M is a stochastic machine that accepts in space $s(n)$. We show how to decide if M accepts any input w in space $s(|w|)$. As usual, we can assume $s = s(|w|)$ is known, and let there be n configurations of M that uses space at most s. Without loss of generality, let these configurations be identified with the integers $1, 2, \ldots, n$ and 1 denotes the initial configuration on input w. Let τ be the operator corresponding of these configurations. Hence we want to accept iff the least fixed point V_τ^* of τ has the property that its first component $[V_\tau^*]_1$ greater than $\frac{1}{2}$. This amounts to checking the validity of the following sentence:

$$(\exists V_\tau^*)(\forall V)[\Sigma(V_\tau^*) \wedge [V_\tau^*]_1 > \frac{1}{2} \wedge (\Sigma(V). \Rightarrow .V_\tau^* \leq V)].$$

This sentence can be decided in time $2^{O(n^4)} = 2^{2^{O(s)}}$, using the above results of Renegar. **Q.E.D.**

Appendix A

Probabilistic Background

The original axiomatic treatment of probability theory due to Kolmogorov [18] is still an excellent rapid introduction to the subject. We refer to [5, 34] for more advance techniques useful for complexity applications. This appendix is a miscellany of quick reviews and useful facts.

Basic Vocabulary. A *Borel field* (or sigma-field) is a set system (Ω, Σ) where Ω is a set and Σ is a collection of subsets of Ω with three properties (i) $\Omega \in \Sigma$, (ii) $E \in \Sigma$ implies $\Omega - E \in \Sigma$ and (iii) $\{E_i\}_{i\geq 0}$ is a countable collection of sets in Σ then the countable union $\cup_{i\geq 0} E_i$ is in Σ. If (iii) is replaced by the weaker condition that $E_1, E_2 \in \Sigma$ imples $E_1 \cup E_2 \in \Sigma$ then we get a *field*. For any collection S of subsets of Ω, there is a unique smallest Borel field that contains S, called the Borel field *generated by* S. The most important example is the *Euclidean Borel field* (R^1, B^1) where $R^1 = \mathbb{R}$ is the real line and B^1 is the Borel field generated by the collection of intervals $(-\infty, c]$ for each real c, $-\infty < c < +\infty$. Members in B^1 are called *Euclidean Borel sets*.

A *probability measure on a Borel field* (Ω, Σ) is a function $\Pr : \Sigma \to [0, 1]$ such that (a) $\Pr(\Omega) = 1$, (b) if $\{E_i\}$ is a countable collection of pairwise disjoint sets in Σ then $\Pr(\cup_{i\geq 0} E_i) = \sum_{i\geq 0} \Pr(E_i)$. A *probability space* is a triple $(\Omega, \Sigma, \Pr)$ where (Ω, Σ) is a Borel field and Pr is a probability measure on (Ω, Σ).

The elements in Ω are often called *elementary events* or *sample points*. Elements of Σ are called *events* or *measurable sets*. Thus Ω and Σ are called (respectively) the *sample space* and *event space*. $\Pr(E)$ is the *probability* or *measure* of the event E. A simple example of probabilistic space is the case $\Omega = \{H, T\}$ with two elements and Σ consists of all subsets of Ω (there are only 4 subsets), and Pr is defined by $\Pr(\{H\}) = p$ for some $0 \leq p \leq 1$.

A *random variable* (abbreviation: r.v.) X of a probability space $(\Omega, \Sigma, \Pr)$ is a real (possibly taking on the values $\pm\infty$) function with domain Ω such that for each real number c, the set

$$X^{-1}((-\infty, c]) = \{\omega \in \Omega \,:\, X(\omega) \leq c\}$$

belongs to Σ. We may simply write

$$\{X \leq c\}$$

for this event. In general, we write[1]

$$\{\ldots X \ldots\}$$

for the event $\{\omega : \ldots X(\omega) \ldots\}$. It is also convenient to write $\Pr\{X \in S\}$ instead of $\Pr(\{X \in S\})$. The intersection of several events is denoted by writing the defining conditions in any order, separated by commas: $\{X_1 \in S_1, X_2 \in S_2, \ldots\}$. If $f(x, y)$ is a real function and X, Y are random variables, then $f(X, Y)$ is a new function on Ω given by $f(X, Y)(\omega) := f(X(\omega), Y(\omega))$. If $f(x, y)$ is "nice", then $f(X, Y)$ will be a new random variable. In particular, the following are random variables:

$$\max(X, Y), \quad \min(X, Y), \quad X + Y, \quad X - Y, \quad XY, \quad X^Y, \quad X/Y.$$

The last case assumes Y is non-vanishing. Similarly, if X_i's are random variables, then so are

$$\inf_i X_i, \quad \sup_i X_i, \quad \liminf_i X_i, \quad \limsup_i X_i.$$

Each X induces a probability measure $\Pr_X$ on the Euclidean Borel field determined uniquely by the condition $\Pr_X((-\infty, c]) = \Pr\{X \leq c\}$. We call $\Pr_X$ the *probability measure* of X. The *distribution function* of X is the real

[1]This '$\{\cdots\}$' notation for events reflects the habit of probabilists to keep the event space implicit. Notice that while probability measures are defined on Σ, random variables are defined on Ω.

function given by $F_X(c) := \Pr_X((-\infty, c])$. Note that $F_X(-\infty) = 0, F_X(\infty) = 1$, F_X is non-decreasing and right continuous. In general, any F with these properties is called a distribution function, and determines a random variable. A set of random variables is *identically distributed* if all members shares a common distribution function F. A finite set of random variables $\{X_i : i = 1, \ldots, n\}$ is *independent* if for any Euclidean Borel sets B_i $(i = 1, \ldots, n)$,

$$\Pr(\cap_{i=1}^n \{X_i \in B_i\}) = \prod_{i=1}^n \Pr\{X_i \in B_i\}.$$

An infinite set of random variables is independent if every finite subset is independent. An important setting for probabilistic studies is a set of independent and identically distributed random variables, abbreviated as *i.i.d.*.

Let (Ω, Σ) be a field, and we are given $m : \Sigma \to [0, 1]$ such that for any countable collection of pairwise disjoint sets $\{E_i \in \Sigma : i \in I\}$,

$$E = \cup_{i \in I} E_i \in \Sigma \text{ implies } m(E) = \sum_{i \in I} m(E_i).$$

Then a standard theorem of Carathéodory says that m can be uniquely extended to a probability measure on (Ω, Σ^*), the Borel field generated Σ.

A standard construction shows that for any countable set of probability measures $\{m_i : i \geq 0\}$ on the Euclidean Borel field, we can construct a probability space $(\Omega, \Sigma, \Pr)$ and a collection of random variables $\{X_i : i \geq 0\}$ such that for each i, m_i is the probability measure of X_i. *Sketch*: We let Ω be the product of countably many copies of the real line $R^1 = \mathbb{R}$, so a sample point is $(w_0, w_1, \ldots)$ where $w_i \in R^1$. A *basic set* of Ω is the product of countably many Euclidean Borel sets $\prod_{i \geq 0} E_i$ where all but a finite number of these E_i are equal to R^1. Let Σ_0 consists of finite unions of basic sets and then our desired Borel field Σ is the smallest Borel field containing Σ_0. It remains to define Pr. For each basic set, define $\Pr(\prod_{i \geq 0} E_i) := \prod_{i \geq 0} \Pr(E_i)$ where only a finite number of the factors $\Pr(E_i)$ are not equal to 1. We then extend this measure to Σ_0 since each member of Σ_0 is a finite union of disjoint basic sets. This measure can be shown to be a probability measure on Σ_0. The said theorem of Carathéodory tells us that it can be uniquely extended to Σ. This concludes our sketch.

A random variable is *discrete* if it takes on a countable set of distinct values. In this case, we may define its *expectation* of X to be $E[X] := \sum_i a_i \Pr\{X = a_i\}$ where i range over all the distinct values a_i assumed by X. Note that $E[X]$ may be infinite. The *variance* of X is defined to be $Var[X] := E[(X - E[X])^2]$. This is seen to give $Var[X] = E[X^2] - (E[X])^2$.

A fundamental fact is that $E[X + Y] = E[X] + E[Y]$ where X, Y are *arbitrary* random variables. Using this simple fact, one often derive surprisingly consequences. In contrast, $Var[X + Y] = Var[X] + Var[Y]$ and $E[XY] = E[X]E[Y]$ are valid provided X, Y are independent random variables.

A random variable X that is 0/1-valued is called a *Bernoulli random variable*. The distribution function of such an X is denoted $B(1, p)$ if $\Pr\{X = 1\}$ is p. If $X_1, \ldots, X_n$ is a set of i.i.d. random variables with common distribution $B(1, p)$ then the random variable $X = X_1 + \cdots + X_n$ has the *binomial distribution* denoted by $B(n, p)$. It is straightforward to calculate that $E[X] = np$ and $Var[X] = np(1 - p)$ if X has distribution $B(n, p)$. Note that Bernoulli random variables is just another way of specifying events, and when used in this manner, we call the random variable the *indicator function* of the event in question. Furthermore, the probability of an event is just the expectation of its indicator function.

Estimations. Estimating probabilities is a fine art. There are some tools and inequalities that the student must become familiar with.

(a) One of these, Stirling's formula in the form due to Robbins (1955), should be committed to memory:

$$n! = \left(\frac{n}{e}\right)^n e^{\alpha_n} \sqrt{2\pi n}$$

where

$$\frac{1}{12n + 1} < \alpha_n < \frac{1}{12n}.$$

Sometimes, the alternative bound $\alpha_n > (12n)^{-1} - (360n^3)^{-1}$ is useful [10]. Using these bounds, it is not hard to show [21] that for $0 < p < 1$ and $q = 1 - p$,

$$G(p, n) e^{-\frac{1}{12pn} - \frac{1}{12qn}} < \binom{n}{pn} < G(p, n) \tag{1}$$

where

$$G(p, n) = \frac{1}{\sqrt{2\pi pqn}} p^{-pn} q^{-qn}.$$

(b) The 'tail of the binomial distribution' is the following sum

$$\sum_{i=\lambda n}^{n} \binom{n}{i} p^i q^{n-i}.$$

We have the following upper bound [10]:

$$\sum_{i=\lambda n}^{n} \binom{n}{i} p^i q^{n-i} < \frac{\lambda q}{\lambda - p} \binom{n}{\lambda n} p^{\lambda n} q^{(1-\lambda)n}$$

where $\lambda > p$ and $q = 1 - p$. This specializes to

$$\sum_{i=\lambda n}^{n} \binom{n}{i} < \frac{\lambda}{2\lambda - 1} \binom{n}{\lambda n} 2^{-n}$$

where $\lambda > p = q = 1/2$.
(c) A useful fact is this: for all real x,

$$e^{-x} \geq 1 - x$$

with equality only if $x = 0$. If $x \geq 1$ then this is trivial. Otherwise, by the usual series for the exponential function, we have that for all real x

$$e^{-x} = \sum_{i=0}^{\infty} \frac{(-x)^i}{i!} = (1 - x) + \frac{x^2}{2!}(1 - \frac{x}{3}) + \frac{x^4}{4!}(1 - \frac{x}{5}) + \cdots .$$

The desired bound follows since $x < 1$. Similarly, we have

$$e^{-x} = (1 - x + x^2/2) - \frac{x^3}{3!}(1 - \frac{x}{4}) - \frac{x^5}{5!}(1 - \frac{x}{6}) - \cdots .$$

Then

$$e^{-x} < 1 - x + x^2/2 = 1 - x(1 - x/2)$$

provided $0 \leq x \geq 4$. If $0 \leq x \leq 1$ then we conclude $e^{-x} < 1 - x/2$. (d) Jensen's inequality. Let $f(x)$ be a convex real function. Convexity of $f(x)$ means $f(\sum_i p_i x_i) \leq \sum_i p f(x_i)$ where $\sum_i p_i = 1, p_i \geq 0$ for all i, and i ranges over a finite set. If X and $f(X)$ are random variables then $f(E[X]) \leq E[f(X)]$. Let us prove this when X has takes on finitely many values x_i with probability p_i: so $E[X] = \sum_i p_i x_i$ and

$$f(E[X]) = f(\sum_i p_i x_i) \leq \sum_i p_i f(x_i) = E[f(X)].$$

For instance, if $r > 1$ then $E[|X|^r] \geq (E[|X|])^r$.

(e) Markov's inequality. Let X be a non-negative random variable, $e > 0$. Then we have the trivial inequality $H(X - e) \leq \frac{X}{e}$ where $H(x)$ (the Heaviside function) is the 0-1 function given by $H(x) = 1$ if and only if $x > 0$. Taking expections on both sides, we get

$$\Pr\{X > e\} \leq \frac{E[X]}{e}.$$

(f) Chebyshev's inequality. Let X be a discrete random variable, $\Pr\{X = a_i\} = p_i$ for all $i \geq 1$, with finite second moment and $e > 0$. Then

$$\Pr\{|X| \geq e\} \leq \frac{E[X^2]}{e^2}.$$

We say this gives an upper bound on tail probability of X. In proof,

$$\begin{aligned} E[X^2] &= \sum_{i \geq 1} p_i a_i^2 \\ &\geq e^2 \sum_{|a_i| \geq e} p_i \\ &= e^2 \Pr\{|X| \geq e\} \end{aligned}$$

Another form of this inequality is

$$\Pr\{|X - E[X]| > e\} \leq \frac{Var(X)}{e^2}$$

where $|X - E[X]|$ measures the deviation from the mean. We could prove this as for Markov's inequality, by taking expectations on both sides of the inequality

$$H(|X - E[X]| - e) \leq \left(\frac{X - E[X]}{e}\right)^2.$$

(g) Chernoff's bound [7] is concerned a set of i.i.d. random variables $X_1, \ldots, X_n$. Let $X = X_1 + \cdots + X_n$ and assume $E[X]$ is finite. Define

$$M(t) := E[e^{tX_1}]$$

and

$$m(a) = \inf_t E[e^{t(X_1 - a)}] = \inf_t e^{-at} M(t).$$

Then Chernoff showed that

$$E[X] \geq a \Rightarrow \Pr\{X \leq na\} \leq [m(a)]^n$$

and

$$E[X] \leq a \Rightarrow \Pr\{X \geq na\} \leq [m(a)]^n.$$

In particular, if X has distribution $B(n, p)$ then it is not hard to compute that

$$m(a) = \left(\frac{p}{a}\right)^a \left(\frac{1-p}{1-a}\right)^{1-a}.$$

Since it is well-known that $E[X] = np$, we obtain for $0 < e < 1$:

$$\Pr\{X \leq (1-e)np\} \leq \left(\frac{1}{1-e}\right)^{(1-e)np} \left(\frac{1-p}{1-(1-e)p}\right)^{n-(1+e)np}.$$

Markov Chains. We continue the discussion of Markov chains from section 3. The *period* of a state i in a chain A is defined in a combinatorial way: it is the largest positive integer d such that every cycle in the underlying graph G_A that contains i has length divisible by d. A state is *periodic* or *aperiodic* depending on whether its period is greater than 1 or not. It is left as an exercise to show that the period is a property of a component.

Recall that state i is recurrent if $f_i^* = 1$. In this case, there is certainty in returning to state i, and under this condition, we may speak of the *mean recurrence time for state* i μ_i, defined as follows:

$$\mu_i = \sum_{n=0}^{\infty} n f_i^{(n)}$$

Using the mean recurrence time, we may introduce new classification of states: state i is *null* if $\mu_i = \infty$, and *non-null* otherwise.

To illustrate the classification of states, we consider the (1-dimensional) random walk with parameter p_0 ($0 < p_0 < 1$): this is the Markov chain whose states are the integers, and the transition probability is given by $p_{i,i+1} = p_0$ and $p_{i,i-1} = 1 - p_0$, for all i. It is clear that every state is essential. It can be shown that each Markov state if recurrent or transient depending on whether $p_0 = \frac{1}{2}$ or not (Exercise). So state 0 is recurrent iff $p = \frac{1}{2}$. Thus $p_0 \neq \frac{1}{2}$ provide examples of essential but transient states. In the recurrent situation, the mean recurrence time is infinite (Exercise). So this illustrates recurrent but null states.

Generating functions. A (real) formal power series is an infinite expression $G(s) = \sum_{n=0}^{\infty} a_n s^n$ in some indeterminate s, where $a_0, a_1, \ldots$ are given real numbers. We say that $G(s)$ is the (ordinary) generating function for the seqence $a_0, a_1, \ldots$. We can manipulate $G(s)$ algebraically: we may add, multiply or (formally) differentiate power series in the obvious way. One should think of $G(s)$ as a convenient way to simultanueously manipulate all the elements of a sequence; hence the terms s^n are just 'place-holders'. These operations reflect various combinatorial operations on the original series. Using well-known identities we can deduce many properties of such series rather transparently. Although we have emphasized that the manipulations are purely formal, we occasionally try to sum the series for actual values of s; then one must be more careful with the analytic properties of these series. Most elementary identities involving infinite series reduces (via the above manipulations) to the following most

fundamental identity $(1-x)^{-1} = \sum_{i\geq 1} x^i$. For example, the student observes that all the results from sections 3 and 4 involving limits is basically an exploitation of this identity.

Rate of Convergence of Substochastic matrices. In section 3, we showed that the entries of the nth power of the fundamental matrix of an absorbing chain approaches 0. We now give a more precise bound on the rate of convergence. For any matrix B, let $\delta_*(B)$ and $\delta^*(B)$ denote the smallest and largest entries in B. Let $\delta(B) = \delta^*(B) - \delta_*(B)$. We have the following simple lemma (cf. [16]):

LEMMA 52. *Let $A = (a_{i,j})$ be a stochastic matrix each of whose entries are at least e for some $0 < e < 1$. For any $n \times m$ non-negative matrix $B = (b_{i,j})$, we have*

$$\delta(AB) \leq (1-2e)\delta(B).$$

Proof. Consider the (i,j)th entry $\sum_{k=1}^n a_{i,k}b_{k,j}$ of AB. Without loss of generality, assume that $a_{i,1} \leq a_{i,2}$. To obtain a lower bound on the (i,j)th entry, assume wlog that $\delta^*(B) = \max\{b_{1,j}, b_{2,j}\}$. Then

$$\begin{aligned}\sum_{k=1}^n a_{i,k}b_{k,j} &\geq a_{i,1}b_{1,j} + a_{i,2}b_{2,j} + (\sum_{k=3}^n a_{i,k})\delta_*(B)\\ &\geq a_{i,1}\delta^*(B) + a_{i,2}\delta_*(B) + (\sum_{k=3}^n a_{i,k})\delta_*(B)\end{aligned}$$

where the last inequality must be justified in two separate cases (in one case, we use the simple fact that $a \geq b$ and $a' \geq b'$ implies $aa' + bb' \geq b' + a'b$). Thus

$$\begin{aligned}\sum_{k=1}^n a_{i,k}b_{k,j} &\geq a_{i,1}\delta^*(B) + (\sum_{k=2}^n a_{i,k})\delta_*(B)\\ &= e\delta^*(B) + (\sum_{k=2}^n a_{i,k})\delta_*(B) + (a_{i,1}-e)\delta^*(B)\\ &\geq e\delta^*(B) + (1-e)\delta_*(B)\end{aligned}$$

To obtain an upper bound on the (i,j)th entry, Assuming wlog that $\delta_*(B) = \min\{b_{1,j}, b_{2,j}\}$, we have

$$\begin{aligned}\sum_{k=1}^n a_{i,k}b_{k,j} &\leq a_{i,1}b_{1,j} + a_{i,2}b_{2,j} + (\sum_{k=3}^n a_{i,k})\delta^*(B)\\ &\leq a_{i,1}\delta_*(B) + a_{i,2}\delta^*(B) + (\sum_{k=3}^n a_{i,k})\delta^*(B)\\ &\leq a_{i,1}\delta_*(B) + (\sum_{k=2}^n a_{i,k})\delta^*(B)\\ &= e\delta_*(B) + (\sum_{k=2}^n a_{i,k})\delta^*(B) + (a_{i,1}-e)\delta_*(B)\\ &\leq e\delta_*(B) + (1-e)\delta^*(B).\end{aligned}$$

The lemma follows since the difference between the largest and smallest entry of AB is at most

$$(e\delta_*(B) + (1-e)\delta^*(B)) - (e\delta^*(B) + (1-e)\delta_*(B)) \leq (1-2e)\delta(B).$$

Q.E.D.

In the exercises, we show how to extend this to substochastic matrix B, *i.e.*, each row sum in B is at most 1.

EXERCISES

Exercise 0.35: Show the following inequalities ((i)-(iv) from [Kazarinoff]):

(i) $(1+\frac{1}{n})^n < (1+\frac{1}{n+1})^{n+1}$.

(ii) $(1+\frac{1}{n})^n < \sum_{k=0}^n \frac{1}{k!} < (1+\frac{1}{n})^{n+1}$.

(iii) $n! < (\frac{n+1}{2})^n$ for $n = 2, 3, \ldots$.

(iv) $(\sum_{i=1}^n x_i)(\sum_{i=1}^n \frac{1}{x_i}) \geq n^2$ where x_i's are positive. Moreover equality holds only if the x_i's are all equal.

(v) $n! < \left(\frac{12n}{12n-1}\right)(2\pi n)^{1/2}e^{-n}n^n$. (Use Robbin's form of Stirling's formula.) ♢

Exercise 0.36:
(i) (Hölder's Inequality) If X and Y are random variables, and $1 < p < \infty$, $\frac{1}{p} + \frac{1}{q} = 1$ then

$$E[XY] \leq E[|XY|] \leq E[|X|^p]^{1/p} E[|Y|^q]^{1/q}.$$

When $p = 2$, this is the Cauchy-Schwartz inequality. (In case $Y \equiv 1$ we have $E[|X|] \leq E[|X|^p]^{1/p}$, which implies the Liapounov inequality: $E[|X|^r]^{1/r} \leq E[|X|^s]^{1/s}$ for $1 < r < s < \infty$.)
(ii) (Minkowski's Inequality)

$$E[|X+Y|^p]^{1/p} \leq E[|X|^p]^{1/p} + E[|Y|^p]^{1/p}.$$

(iii) (Jensen's inequality) If f is a convex real function, and suppose X and $f(X)$ are integrable random variables. Then $f(E[X]) \leq E[f(X)]$. (Note that convexity means that if $\sum_{i=1}^n c_i = 1$, $c_i > 0$, then $f(\sum i = 1^n c_i x_i) \leq \sum_{i=1}^n c_i f(x_i)$.) ♢

Exercise 0.37: Construct the probability space implicitly associated with a Markov chain. ♢

Exercise 0.38: For any positive integer k, construct a finite Markov chain with states $0, 1, \ldots, n$ such that states 0 has the value $k \leq p^*_{0,0} < k+1$. Try to minimize $n = n(k)$. ♢

Exercise 0.39: In this exercise, we do not assume the Markov chain is finite. Show that the following are properties, though defined for individual states, are characteristics of components:
(i) Period of a Markov state.
(ii) Nullity of a Markov state. ♢

Exercise 0.40: Show that $g_{i,j} = f^*_{i,j} g_{j,j}$. (From this, conclude that $g_{i,j} > 0$ if and only if $g_{i,j} = f^*_{i,j}$.) **Hint:** Write $g_{i,j} = \sum_{n=0}^\infty \Pr(A_n B_n C_n | D)$ where D is the event that the state at time 0 is i, A_n is the event that the states at times $1, \ldots, n-1$ are not equal to j, B_n is the event that the state at time n is equal to j, C_n is the event that the state at time s is equal to j for infinitely many $s > n$. Then $\Pr(A_n B_n C_n | D) = \Pr(C_n | A_n B_n D) \Pr(A_n B_n | D)$ But the Markov property implies $\Pr(C_n | A_n B_n D) = \Pr(C_n | D)$. ♢

Exercise 0.41: Above, we proved a bound on the rate of convergence of stochastic matrices. Extend it to substochastic matrices. ♢

Exercise 0.42:
(a) Prove equation (1) in the appendix.
(b) Say $f(n) \sim g(n)$ if $f(n)/g(n)$ approaches 1 as $n \to \infty$. Conclude that for $k = 1, \ldots, n/2$, $p = k/n$ and $q = 1 - p$, then as $k \to \infty$ and $n - k \to \infty$:

$$\binom{n}{k} \sim \frac{1}{\sqrt{2\pi pqn}(p^p q^q)^n}$$

(c) Let $0 < p < 1$ and $q = 1 - p$. Show that the probability that a Bernoulli random variable with mean p attains k successes in n trials is

$$\binom{n}{k} p^k q^{n-k} \sim \frac{1}{\sqrt{2\pi pqn}}$$

♢

Exercise 0.43: Show
(a)

$$\left(\frac{1-p}{1-\delta p}\right)^{1-\delta p} \leq e^{\delta - 1}$$

for $0 < \delta < 2$.
(b)

$$\left(\frac{1}{1+e}\right)^{1+e} \leq e^{-e-(e^2/3)}$$

for $0 < e \leq 1$.
(c)

$$\left(\frac{1}{1-e}\right)^{1-e} \leq e^{e-(e^2/2)}$$

for $0 < e \leq 1$.
(d) Conclude that in the binomial case of Chernoff's inequality,

$$\Pr\{X \geq (1+e)np\} \leq e^{-(e^2/3)np}$$

and

$$\Pr\{X \leq (1-e)np\} \leq e^{-(e^2/2)np}.$$

♢

Exercise 0.44: Deduce from Chernoff's bound the following estimate on the tail of binomial distribution:

$$\sum_{i=\lfloor t/2 \rfloor}^{t} \binom{t}{i} p^i q^{t-i} \leq (4pq)^{t/2}.$$

♢

End Exercises

Bibliography

[1] A. Avizienis. Signed-digit number representation for fast parallel arithmetic. *Inst. Radio Engr. Trans. Electron. Comput.*, 10:389–400, 1961.

[2] L. Babai and S. Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computers and Systems Science*, 36:254–276, 1988.

[3] R. Beigel, N. Reingold, and D. Spielman. PP is closed under intersection. *J. of Computer and System Sciences*, 50(2):191–202, 1995.

[4] S. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18:147–150, 1984.

[5] B. Bollobás. *Random Graphs*. Academic Press, 1985.

[6] A. Borodin, S. Cook, and N. Pippenger. Parallel computation for well-endowed rings and space-bounded probabilistic machines. *Information and Computation*, 58:113–136, 1983.

[7] H. Chernoff. A measure of asymptotic efficiency for tests of hypothesis based on sum of observations. *Ann. of Math. Stat.*, 23:493–507, 1952.

[8] K. L. Chung. *Markov Chains with stationary transition probabilities*. Springer-Verlag, Berlin, 1960.

[9] K. de Leeuw, E. F. Moore, C. E. Shannon, and N. Shapiro. *Computability by probabilistic machines*. Automata Studies. Princeton, New Jersey, 1956.

[10] W. Feller. *An introduction to Probability Theory and its Applications*. Wiley, New York, 2nd edition edition, 1957. (Volumes 1 and 2).

[11] F. R. Gantmacher. *The Theory of Matrices*. Chelsea Pub. Co., New York, 1959. Volumes 1 and 2.

[12] J. T. Gill. Computational complexity of probabilistic Turing machines. *SIAM J. Comp.*, 6(4):675–695, 1977.

[13] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proofs. *17th ACM Symposium STOC*, pages 291–304, 1985.

[14] S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. *18th ACM Symposium STOC*, pages 59–68, 1986.

[15] H. Jung. Relationships between probabilistic and deterministic tape complexity. *10th Sympos. om Mathematical Foundations of Comp. Sci.*, pages 339–346, 1981.

[16] J. G. Kemeny and J. L. Snell. *Finite Markov chains*. D. Van Nostrand, Princeton, N.J., 1960.

[17] J. Köbler, U. Schöning, and J. Torán. *The Graph Isomorphism Problem: Its Structural Complexity*. Birkhauser, 1993.

[18] A. N. Kolmogorov. *Foundations of the theory of probability*. Chelsea Publishing Co., New York, 1956. Second English Edition.

[19] R. E. Ladner and M. J. Fischer. Parallel prefix computation. *Journal of the ACM*, 27(4):831–838, 1980.

[20] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[21] W. W. Peterson and J. E. J. Weldon. *Error-Correcting Codes*. MIT Press, 1975. 2nd Edition.

[22] N. Pippenger. Developments in "The synthesis of reliable organisms from unreliable components". manuscript, University of British Columbia, 1988.

[23] M. O. Rabin. Probabilistic algorithms. In J. F. Traub, editor, *Algorithms and Complexity: New Directions and Recent Results*, pages 21–39. Academic Press, New York, 1976.

[24] J. Radhakrishnan and S. Saluja. Lecture notes: Interactive proof systems. Research Report MPI-I-95-1-007, Max-Planck-Institut für Informatik, Im Stadtwald, D-66123 Saarbrücken, Germany, March 1995. This is a full TECHREPORT entry.

[25] J. Renegar. On the Computational Complexity and Geometry of the First-Order Theory of the Reals, Part I: Introduction. Preliminaries. The Geometry of Semi-Algebraic Sets. The Decision Problem for the Existential Theory of the Reals. *J. of Symbolic Computation*, 13(3):255–300, March 1992.

[26] J. Renegar. On the Computational Complexity and Geometry of the First-Order Theory of the Reals, Part II: The General Decision Problem. Preliminaries for Quantifier Elimination. *J. of Symbolic Computation*, 13(3):301–328, March 1992.

[27] J. Renegar. On the Computational Complexity and Geometry of the First-Order Theory of the Reals, Part III: Quantifier Eliminiation. *J. of Symbolic Computation*, 13(3):329–352, March 1992.

[28] W. L. Ruzzo, J. Simon, and M. Tompa. Space-bounded hierarchies and probabilistic computations. *Journal of Computers and Systems Science*, 28:216–230, 1984.

[29] U. Schöning. *Complexity and Structure*, volume 211 of *Lecture Notes in Computer Science*. Springer-Verlag, 1986.

[30] A. Shamir. IP=PSPACE. *J. of the ACM*, 39(4):869–877, 1992.

[31] J. Simon. On tape bounded probabilistic Turing machine acceptors. *Theor. Computer Science*, 16:75–91, 1981.

[32] J. Simon. Space-bounded probabilistic turing machine complexity are closed under complement. *13th Proc. ACM Symp. Theory of Comp. Sci.*, pages 158–167, 1981.

[33] R. Solovay and V. Strassen. A fast monte-carlo test for primality. *SIAM J. Computing*, 6:84–85, 1977.

[34] J. Spencer. *Ten Lectures on the Probalistic Method*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, 1987.

Lecture XXII
QUANTUM COMPLEXITY

§1. Introduction

Quantum computation is an exciting area of current research. It is fundamentally different from classical computation because different laws of physics are used. The idea of quantum computing was first suggested in the early 1980's by Paul Benioff [3, 4, 5], a physicist from Argonne National Laboratory. Feynman [18] showed how to build a classical computer based on quantum principles. Deutsch [15, 16] introduced the universal quantum computer and quantum circuits. A comprehensive and accessible treatment of quantum computation including quantum information theory is Nielsen and Chuang [22].

Related questions about the physics of information and computing date back earlier: since information is physically represented, what does the laws of physics say about fundamental limits of computation? Rolf Landauer (1927-1999) was interested in minimum energy computation. He noted (1961) that erasure of information is necessarily a dissipative process – heat is lost from the computer into the environment. Thus, if we erase one bit of information, the overall entropy increases by $k \ln 2$. At temperature T, the work expended is $kT \ln 2$. The converse to this so-called "Landauer Principle" is that if we compute reversibly (with no erasure of information) then no dissipation or power consumption is needed. In 1973, Bennett [7] showed that such computation is always possible in theory (but it would be a very slow computation!). Reversible computation was further investigated by Tofolli and Fredkin [25, 19]. Although reversible computing is based on classical laws of physics, it can be viewed as a precursor to quantum computation. For a history of reversible computing, see [8].

Two developments in the early 1990s help to push quantum computing out of the curiosity stage. One was the discovery of quantum algorithms that have major implications for cryptography. In 1994, Peter Shor at AT&T Labs showed that the problems of integer factorization and discrete logarithm can be solved by quantum computers with high probability in polynomial time [24]. As it is widely believed that both these problems are non-polynomial time on classical computers, and the security of many cryptographic protocols depend on these assumptions, this suggests a "killer app" for quantum computing. As a result, the subject holds real interest for agencies such as DARPA and the National Security Agency. The other development is the experimental demonstration of techniques that could be used to build quantum computers. A major challenge here is to isolate the quantum bits (qubits) from environment (to keep the system "coherent"). Several competing technologies are being investigated. Seth Lloyd (1993) showed that a quantum computer could be built from an array of coupled two-state quantum systems. An implementation proposed by Chuan and Gershenfeld, and independently by Cory, Fahmy and Havel, is based on spins in the nucleus of atoms. Such nuclear qubits are naturally isolated from the external world by its clouds of electrons, and they may be assembled naturally as molecules. The technology for manipulating nuclear spins is NMR (nuclear magnetic resonance), a well-developed technology routinely used in medicine. In August of 2000, a 5-qubit computer was announced by IBM corporation. The ion trap approach of Ignacio Cirac and Peter Zoller [13] is based on confining cold ions along a line ("linear Paul trap"). The quantum state of each ion is a superposition of its ground state $|0\rangle$ and some relatively long-lived excited state $|1\rangle$, representing the qubits (see below). Laser beams directed at the individual ions can achieve the transitions within each ion but how can the qubits interact in the quantum mechanical sense? Cirac and Zoller showed with proper tuning of the lasers, the controlled XOR gate (see below) can be implemented with 5 laser pulses. Such devices have been constructed [23, 21]. The speed of such a device depends on the frequency of the fundamental vibrational modes of the ions; current technology can perhaps achieve 10^4 steps/second (see [2]). Even if current approaches do not lead realistic quantum computers, they are nevertheless useful for demonstrating the principles of quantum computation [2].

Quantum computing also has implications for the fields of information theory, coding theory and cryptography. See the survey [11]. It calls for new foundations for each of these areas. Among other things, one goal in this chapter is to present Shor's algorithm for factoring.

§2. Quantum Bits

¶1. The basic unit of information in a classical computers is the **bit**, an entity that can assume exactly one of two distinct values, denoted 0 and 1. The analogous **quantum bit** (or "qubit") also has two values which we identify with the classical values 0 and 1. These are called **eigenvectors** or **eigenstates**. Following a standard notation[1] in physics these eigenstates are denoted $|0\rangle$ and $|1\rangle$. In general, if Ψ is the "name" of a quantum state, we

[1] This is the **ket** notation; there is a corresponding **bra** notation which has the form $\langle y|$. The pair $|x\rangle$ and $\langle y|$ can be composed as $\langle y||x\rangle$, viewed as the inner product of two vectors (in this inner product, y must first be conjugated before forming the scalar product with x). The "bra-ket notation" is from the physicist Dirac.

write $|\Psi\rangle$ to denote the state. What we choose for the name is not important. For instance, we may use suggestive symbols such as $\uparrow, \downarrow, +, -$, or more readable names such as *up, down, left, right*. However, the value of a qubit is a **quantum state** of the form

$$c_0|0\rangle + c_1|1\rangle \tag{1}$$

where $c_0, c_1 \in \mathbb{C}$ (complex numbers) satisfying $|c_0|^2 + |c_1|^2 = 1$. For instance, if $c_0 = 0, c_1 = 1$, then the quantum state is just $|1\rangle$, a eigenstate. We say that the quantum state is a **superposition** of the eigenstates.

Mathematically, an eigenstate is just a basis vector in some chosen basis. Thus the states of a qubit lives in the complex two-dimensional vector space in which all the vectors have unit length. We could discuss all our results using only this mathematical model, but in the following we will suggest some physical intuitions and possible interpretations of the mathematics.

The qubit can be realized by a variety of physical quantum systems. Any of these systems, in principle, can be the basis for constructing quantum computers. For instance, an electron or other spin-$\frac{1}{2}$ particle can have one of two distinct spins, called spin-up and spin-down. Thus the state of a qubit can be encoded by the state of an electron spin. Alternatively, the photon (light particle) is a massless, spin-1 particle that can have one of two independent polarizations. Using a photon as a qubit, we can manipulate its state by rotating the polarization of the photon.

¶2. Quantum states as probability distributions. So far, we identified an eigenstate with a classical value (0 or 1). But what is the classical analogue of the quantum state (1)? To answer this, we introduce the concept of a **measurement**. This is an operation that is applicable to quantum states. When the quantum state (1) is measured, it "collapses" to the eigenstate $|i\rangle$ with probability $|c_i|^2$ where $i = 0, 1$. Note that, unlike classical bits which can be measured without affecting its value, any measurement of qubits is destructive. In any case, we can now view quantum states as the generalization of another classical computational concept: the notion of **random bit**. A random bit is just a (discrete) probability distribution (p_0, p_1) where $p_0 + p_1 = 1$ and $0 \leq p_i \leq 1$. This random bit assumes the value i with probability p_i where $i = 0, 1$. Thus the quantum state (1) may be associated with a probability distribution $(p_0, p_1) = (|c_0|^2, |c_1|^2)$. But infinitely many different quantum states are associated to a given probability distribution. For instance, if $c_0 = c_1 = \frac{1}{\sqrt{2}}$ then the quantum states $c_0|0\rangle + c_1|1\rangle$ and $c_0|0\rangle - c_1|1\rangle$ are both associated to the probability distribution $(\frac{1}{2}, \frac{1}{2})$.

¶3. Stern-Gerlach Experiment (1922) We describe some physical experiments which elucidate the nature of quantum states. In particular, we want to see why quantum states are not simply probability distributions.

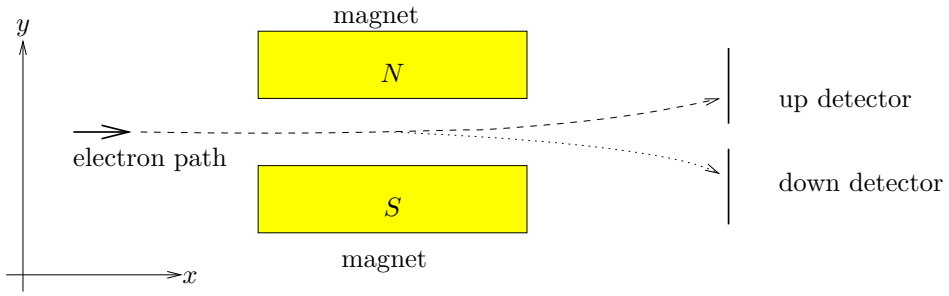

Figure 1: Stern-Gerlach measuring device

The physical interpretation of measurement in spin models is illustrated in an apparatus known as a Stern-Gerlach device. See Figure 1 for a 2-dimensional rendition. The apparatus comprises a pair of magnets (N, S) that surrounds the path of an electron, together with two detector plates at the exit paths of the electron. When an electron passes through the Stern-Gerlach device, the path of the electron will deflect up or down with certain probability (depending on the source of the electrons). For a fixed electron source, this probability is empirically well-defined. (The original experimental particle was silver atoms, but later hydrogen atoms.) We assume that the electrons deflect up or down with equal probability. To explain this, we postulate that an electron has two **spin states**, called up-spin or down-spin. To explain the deflection, we could say the electron state is in some superposition of these spin states, or we could say that the electron state is a probability distribution over the spin states. The experiment so far has no preference for either explanation.

The Stern-Gerlach device illustrates the idea that measurements are relative to a choice of basis: the magnetic flux is conventionally said to flow perpendicularly to the plane defined by the magnets (i.e., perpendicularly out of

the page). In Figure 1, the electron deflects up (positive y direction) or down (negative y direction) with suitable probability. If we rotate the device by 90° about the axis of the electron path, so that the magnet N (S) lies above (below) the page, resulting measurement would be different: the same electron would now deflect left or right with suitable probability. In Figure 1, left/right means into/out of the page. This amounts to a new measurement basis. We normally have the freedom to choose a basis that is most convenient for the application.

We now extend the previous experiment in two ways. First, suppose we pass the up-spin electrons from the first Stern-Gerlach device into another identical Stern-Gerlach device. What do we expect to see? Well, the electrons will only deflect up. This is fully expected.

Let us do a different experiment, by passing the up-spin electrons through a second Stern-Gerlach device whose magnets have been rotated by 90°, to cause a left-right deflection. We will see that the electrons will deflect to the left and right with equal probability. This might be surprising, if we had expected the up-spin electrons to have no left- or right-spin components. So perhaps electrons have independent spin components for up/down as well as left/right spins. If we continue by passing the left-spin electrons through a third Stern-Gerlach device with the up-down orientation as the first device, we again see an equal probability of the electrons deflecting up or down. This would be a surprise if we had expected to see up-spins only, as only up-spin electrons were sent through the second device.

To explain all this, we postulate that the eigenstates in the basis of the up/down measurement device are $|up\rangle$ and $|down\rangle$, respectively. But relative to the left/right measurement, the eigenstates are $|left\rangle = (|up\rangle + |down\rangle)/\sqrt{2}$ and $|right\rangle = (|up\rangle - |down\rangle)/\sqrt{2}$. Thus the electrons sent into the second device are in state $|up\rangle$, but in the measurement basis, this appears as $(|left\rangle + |right\rangle)/\sqrt{2}$, and thus they have equal probability of going left or right. The electrons sent into the third device are in the state $|left\rangle$, but relative to the measurement basis of up/down, they have equal probabilities of going up or down.

¶4. Mach-Zehnder Interferometer. The preceding discussion of measure basis already suggests that superposition of eigenstates is different from a probability distribution over eigenstates. Another series of experiments[2] based on the **Mach-Zehnder interferometer** confirms this. Refer to Figure 2. In experiment A, we reflect a photon off a half-silvered mirror M_1, we will detect the photon at detector 1 or 2 with 50% probability each (we choose a photon source with this property). The classical explanation is that the photon has equal probability of taking either path. The quantum mechanical explanation is that both paths are taken at once (superposition), but the detectors cause a collapse of the state. But so far, we have no basis to prefer one explanation over the other (in fact, the classical one should be preferred for its simplicity). Incidentally, note that this experiment uses photons (instead of the electrons of Stern-Gerlach) to demonstrate quantum effects. The two eigenstates of photons are called polarizations instead of spins.

In experiment B, we place two fully silvered mirror at the positions of detectors 1 and 2, but arranged so that the two photon paths (P_1, P_2) recombine. We place a second half-silvered mirror M_2 at the point of recombination. We also place the detectors 3 and 4 to measure the reflection or non-reflection from M_2. It turns out, the photon reaches detector 3 with 100% probability and never reach detector 4. This is impossible to explain classically. The quantum mechanical view accounts for this: the photon must have travelled both paths P_1 and P_2, and when recombined, it is able to distinguish the two choices at M_2 and only chose (by way of interference) the "correct path". In experiment C, we confirm this explanation by placing a barrier (K) in path P_2 of the previous experiment. Now, we have 50% probability of detecting the photon at detectors 3 and 4.

§3. Quantum Words

¶5. Real world computers operate on fixed length sequence of bits, called a **word**. A word in modern computers is typically 32 bits or 64 bits long (circa 2000). Similarly, a finite sequence of qubits will be called a **quantum word** ("quword" for short). In the literature, a quword is also known as a "quantum register". An array of n qubits is called a n-quword.

Classically, transition from bits to words is trivial. But quantum words introduce new situations with no classical analogue. This is the phenomenon of quantum interference and phase information. This is already evident for $n = 2$. Let A and B be the qubits in a 2-quword. If A is in the eigenstate $|0\rangle$ and B in $|1\rangle$, then the state of the quword is written $|0\rangle \otimes |1\rangle$ (tensor product), or more compactly, $|01\rangle$ and sometimes $|0\rangle|1\rangle$. This looks like Cartesian product, but tensor product is more than this. The hint that something else is going on is the following basic property of tensors which we will use often: for any scalar α,

$$\alpha(u_1 \otimes u_2) = (\alpha u_1) \otimes u_2 = u_1 \otimes (\alpha u_2). \tag{2}$$

[2]See "Un saut d'echelle pour les calculateurs", by A. Barenco, A.Ekert, A. Sanpera and C.Machiavello, in La Recherche, Nov 1996. Adapted article by Barenco may be found in http://www.qubit.org/intros/comp/comp.html. See also [14].

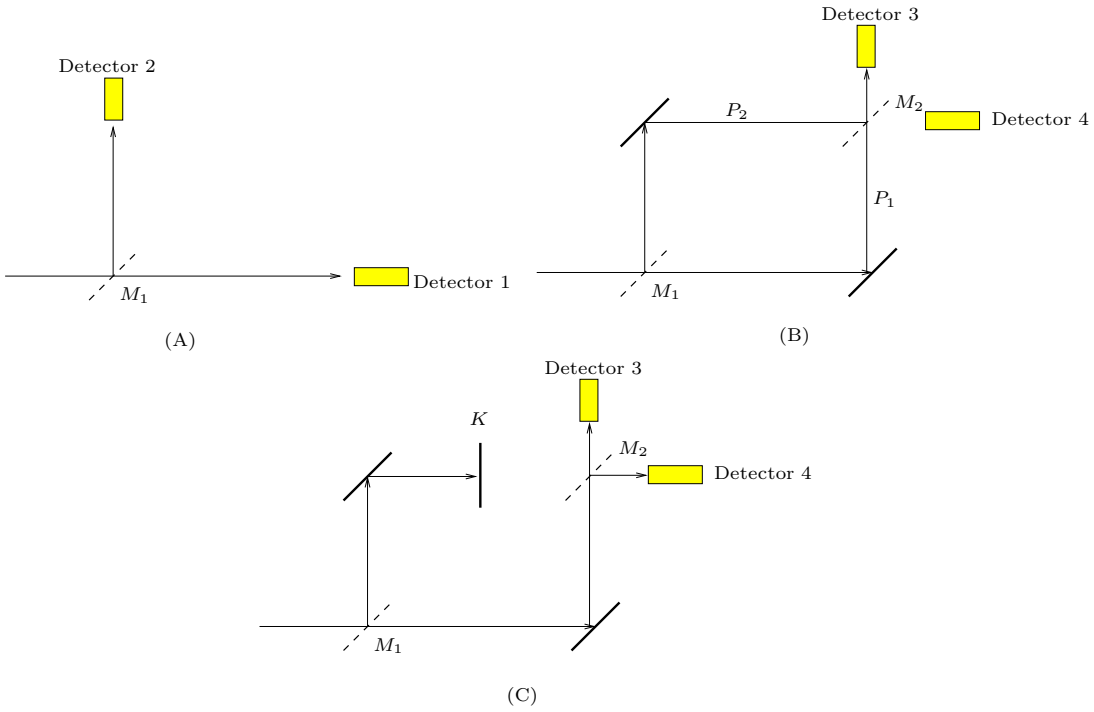

Figure 2: Experiments A, B, C.

If A, B are in the superposition of eigenstates $x = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $y = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, respectively, then their joint state is given by

$$|x\rangle \otimes |y\rangle = \frac{1}{2}(|00\rangle - |01\rangle + |10\rangle - |11\rangle).$$

Next supposed that the quword is in the state $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ ("Bell State"). These two qubits are "entangled": when you measure one of the qubits, then the other qubit would also collapse to the same value. Intuitively, they are maximally entangled (or correlated); quantum information theory is the subject that shed light on this phenomenon.

¶6. Eigenstates. In general, an n-quword has 2^n eigenstates of the form $|b_1 \cdots b_n\rangle$ where $b_i \in \{0, 1\}$. For instance, if $n = 2$, then the possible eigenstates are $|00\rangle, |01\rangle, |10\rangle, |11\rangle$. Interpreting $b_1 \cdots b_n$ as a binary representation of a natural number, the eigenstates can be denoted $\{|i\rangle : i = 0, 1, \ldots, 2^n - 1\}$. A quantum state (or[3] **pure state**) is again a superposition of these eigenstates, and can be represented by a unit length vector in $c \in \mathbb{C}^{2^n}$. Unit length means that $\sum_{i=0}^{2^n-1} |c_i|^2 = 1$ where $c = (c_0, \ldots, c_{2^n-1})$. The quantum state corresponding to c is $\sum_{i=0}^{2^n-1} c_i|i\rangle$.

¶7. Bit Notation and Indexing Notation. We have just established 2 conventions for writing eigenstates: $|x\rangle$ where x is either a binary string or a natural number. The first convention of using binary strings is called the **bit notation** and the second convention of using natural numbers is called the **indexing notation**.

In practice, these two conventions are used interchangeably, whichever is more convenient. The indexing notation is convenient when we treat each eigenstate wholistically, and is used for writing superposition as a summation: $\sum_{x=0}^{2^n-1} c_x|x\rangle$. The bit notation is useful for describing composite eigenstates: if x, y representing two disjoint quwords, the ket-notation admits the operation $|x\rangle \otimes |y\rangle = |xy\rangle$. When x, y are superposition of states for two disjoint quwords, $|x\rangle \otimes |y\rangle$ is obtained as a Cartesian product of the separate eigenstates.

¶8. Measurement. Quwords offer a twist to the concept of measurement. Naturally, if $|x\rangle = \sum_{i=0}^{2^n-1} c_i|i\rangle$ and it is measured, then the state collapses to $|i\rangle$ with probability $|c_i|^2$ for each i. But we can also measure any individual qubit, or more generally, any subset Y of qubits. If the set of n qubits is X, let $Z = X \setminus Y$ be the complementary

[3]This curious terminology will become clearer when we discuss ensembles of quantum states.

set of qubits. Each eigenstate $|i\rangle$ of X can be written as $|y\rangle \otimes |z\rangle$ where $|y\rangle, |z\rangle$ are eigenstates of Y and Z. Denote the Y-projection of $|i\rangle$ by $\pi_Y(|i\rangle)$, and thus

$$|i\rangle = \pi_Y(|i\rangle) \otimes \pi_Z(|i\rangle).$$

Let $P_Y(y) = \{i = 0, \ldots, 2^n - 1 : \pi_Y(|i\rangle) = |y\rangle\}$. When we measure the Y-bits of a state $|x\rangle$, we will see each eigenstate $|y\rangle$ of Y with probability $p_x(y) := \sum_{i \in P_Y(y)} |c_i|^2$. After a measurement of $|x\rangle$ which revealed a particular state $|y\rangle$ in Y-qubits, the state of the quword becomes

$$\frac{\sum_{i \in P_Y(y)} c_i |i\rangle}{\sum_{i \in P_Y(y)} |c_i|^2}.$$

Let us illustrate this for a 3-quword whose eigenstates are labeled $|0\rangle, \ldots, |7\rangle$ as usual. If Y refer to the middle bit, then $P_Y(0) = \{0, 1, 4, 5\}$ and $P_Y(1) = \{2, 3, 6, 7\}$. Upon measuring the y-bit of state $|x\rangle = \sum_{i=0}^{7} c_i |i\rangle$, we see the value of $|0\rangle$ with probability $p_x(0) = |c_0|^2 + |c_1|^2 + |c_4|^2 + |c_5|^2$. The new state of the quword is $(c_0|0\rangle + c_1|1\rangle + c_4|4\rangle + c_5|5\rangle)/p_x(0)$. There is an analogous case where we see the value $|1\rangle$ for the Y-bit.

¶9. Observables. The general treatment of measurements is in terms of an **observable**, $M = \{M_i : i \in I\}$ where I is the set of possible outcomes of the observable M, and M_i is an **measurement operator** corresponding to outcome i. The operators (i.e., matrices in concrete representation) satisfy the equation

$$\sum_{i \in I} M_i^* M_i = I \tag{3}$$

where I is the identity operator. We can **measure** any state $|x\rangle$ using an observable M. This measurement changes $|x\rangle$ to one of $|I|$ possible new states called **collapsed states**, as follows. The probability of outcome $i \in I$ is[4] given by

$$p_i(x) := \langle x | M_i^* M_i | x \rangle$$

($p_i(x) = \|M_i|x\rangle\|^2$, see below). The corresponding collapsed state is

$$\frac{M_i|x\rangle}{\sqrt{\langle x|M_i^* M_i|x\rangle}} = \frac{M_i|x\rangle}{\sqrt{p_i(x)}}.$$

The Equation (3) ensures that

$$\sum_{i \in I} p_i(x) = \sum_{i \in I} \langle x | M_i^* M_i | x \rangle = 1.$$

If we are not interested in describing collapsed states, only the probability of various outcomes, then it is enough to specify the matrices $E_i = M_i^* M_i$ ($i \in I$). The E_i's are called positive operator-valued measure (POVM) operators.

¶10. Example. Let us reformulate our discussion about measuring a subset of qubits of a quword in the framework of observable. Say we want to "observe" the last qubit (least significant bit) in a 2-quword. So there are only two outcomes, $I = \{0, 1\}$ and $M = \{M_0, M_1\}$.

$$M_0 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \qquad M_1 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Clearly, $M_0 + M_1 = I$. For instance, if $|y\rangle = (|0\rangle + |1\rangle) \otimes |0\rangle/\sqrt{2}$, then $\langle y|M_0^* M_0|y\rangle = \langle y|M_0|y\rangle = 1$. Thus the probability of observing a zero in the last qubit in $|y\rangle$ is 1, as we would expect. Moreover, the outcome of the observation M_0 is $|y\rangle$. As expected, $\langle y|M_1^* M_1|y\rangle = 0$.

Each measurement operator M_i is represented by a $n \times n$ Hermitian matrix. Using the spectral representation theorem, $M_i = \sum_{j=1}^{n} \lambda_j P_j$ where P_j are projection operators.

Exercises

[4]Note that we could parse the expression $\langle x|M_i^* M_i|x\rangle$ in several equivalent ways: $(\langle x|M_i^*)(M_i|x\rangle) = \langle x|(M_i^*(M_i|x\rangle)) = \langle x|((M_i^* M_i)|x\rangle)$. Also, $(M_i|x\rangle)^* = \langle x|M_i^*$.

Exercise 3.1: Suppose $|y\rangle = e^{\mathbf{i}\theta}|x\rangle$. Show that there is no observable difference between $|x\rangle$ and $|y\rangle$. ♢

Exercise 3.2: Suppose $X = (x_2x_1x_0)$ be three qubits whose eigenstates represents the binary numbers between 0 and 7. Describe an observable with two outcomes $M = \{M_0, M_1\}$ such that for eigenstate $|x\rangle$, we have $p_i(x) = 1$ iff the $|x\rangle$ corresponds to a prime number (i.e, $2, 3, 5$ or 7). ♢

End Exercises

§4. Axiomatic Quantum Mechanics

¶11. The simplest approach to quantum mechanics is an axiomatic one. We postulate the basic entities, operators and their properties, and investigate them mathematically. It is the physical interpretation that validates the mathematical model, but after the model has been accepted, it can be studied fully on its own merits. This often lead to consequences that are so surprising that we need physical experiments to confirm the model prediction.

As seen above, quantum states are some kind of complex vectors. Following von Neumann, we postulate the space of quantum states as a suitable complex vector space S. Elements $\varphi, \psi \in S$ of S are called **states**. In the Dirac notation, we would use "$|\varphi\rangle$" instead of a plain φ. Each φ defines a dual vector (see below) which in the Dirac notation is written "$\langle\varphi|$" (the bra-notation). We emphasize that these are stylistic conventions. We saw that those complex numbers $\alpha \in \mathbb{C}$ with absolute value $|\alpha| \leq 1$ are very important: such α's are called **probability amplitudes** (or, complex probabilities).

¶12. Hilbert Space. We outline some main properties we will need; refer to the appendix for additional mathematical background. Let S be a complex vector space, endowed with an inner product $\langle\cdot,\cdot\rangle : S \times S \to \mathbb{C}$. Consistent with the Dirac bra-ket notation, the inner product of $\psi, \varphi \in S$ will be written as "$\langle\psi|\varphi\rangle$" instead of the conventional notation for scalar product, $\langle\psi, \varphi\rangle$. This suggests that we view the inner product as the application of a dual vector $\langle\psi|$ to the state $|\varphi\rangle$. The inner product satisfies three axioms:

1. (Positivity) $\langle\psi|\psi\rangle$ is real and non-negative for all ψ. Furthermore, it is strictly positive iff $\psi \neq 0$.
2. (Linearity) $\langle\psi|a\varphi_1 + b\varphi_2\rangle = a\langle\psi|\varphi_1\rangle + b\langle\psi|\varphi_2\rangle$.
3. (Skew Symmetry) $\langle\psi|\varphi\rangle = \langle\varphi|\psi\rangle^*$.

Define the **norm** of $\varphi \in S$ by $\|\varphi\| := \sqrt{\langle\varphi|\varphi\rangle}$. Two vectors φ, ψ are **orthogonal** if $\langle\varphi|\psi\rangle = 0$. A set of vectors $\{|e_i\rangle : i \in B\}$ is called a **basis** if every $|\varphi\rangle \in S$ can be *uniquely* written as a linear combination $\sum_{i\in B} c_i|e_i\rangle$. Uniqueness implies that the $|e_i\rangle$'s are linearly independent. If each $|e_i\rangle$ has unit norm, and they are mutually orthogonal, the basis is **orthonormal**. A basic fact about vector spaces is that the cardinality of B is a property of S, and is called the **dimension** of S. Dimension can be finite or infinite. Relative to the norm $\|\cdot\|$, a sequence $\{x_0, x_1, \ldots\}$ in S is **Cauchy** if for every $\varepsilon > 0$ there exists n such that $\|x_i - x_j\| < \varepsilon$ whenever $i, j > n$. We say S is **complete** relative to this norm if every Cauchy sequence $\{x_0, x_1, \ldots\}$ in S with respect to the norm converges to an element of S.

A **Hilbert space** is a complex vector space S with an inner product $(x, y) \in S^2 \mapsto \langle x|y\rangle \in S$ that is complete. Hilbert spaces has a natural topology induced by the metric $d(x, y) = \|x - y\|$. So the notion of continuity of a function $f : S \to \mathbb{C}$, etc, is meaningful.

¶13. Examples. Let $S = \ell_2(n)$ denote the vector space $\mathbb{C}^n$ where $x, y \in \mathbb{C}^n$ has inner product $\langle x|y\rangle = \sum_{i=1}^n x_i^* y_i$, x_i, y_i being the components of x, y, respectively. We can extend this to the infinite dimensional space $S = \ell_2(\infty)$ comprising vectors $(x_i)_{i=0}^\infty$ with $\sum_{i=1}^\infty |x_i|^2 < \infty$. Another infinite dimensional space is $L_2(a, b)$ for reals a, b with $-\infty \leq a < b \leq \infty$. This space comprises all $f : (a, b) \to \mathbb{C}$ such that $\int_a^b |f(t)|^2 dt$ is defined and $< \infty$. The inner product is defined by $\langle f|g\rangle = \int_a^b f^*(t)g(t)dt$. The elements in $L_2(a, b)$ are equivalence classes of functions where $f \equiv g$ if $\|f - g\| = 0$.

¶14. Linear operators and functionals. A **linear functional** over S is a continuous linear function $L : S \to \mathbb{C}$. Linearity means $L(ax + by) = aL(x) + bL(y)$ for all $x, y \in S$ and $a, b \in \mathbb{C}$. For any $x \in S$, we can obtain a linear functional $L_x : S \mapsto \mathbb{C}$ where $L_x(y) = \langle x|y\rangle$. A basic result about Hilbert space is that *all* linear functionals are of this form. Let S^* be the space of all linear functionals over S. We make S^* a vector space by defining $aL_x + bL_y = L_{ax+by}$ for all $x, y \in S$ and $a, b \in \mathbb{C}$. This S^* is also called the **dual space** of S, and elements of S^* are **dual vectors**. If a state y is written $|y\rangle$, its dual vector is denoted $\langle y|$. Applying a linear functional L_x to y yields $\langle x|y\rangle$.

A **linear operator** of S is a linear function $A : S \to S$. Linearity means $A(ax + by) = aA(x) + bA(y)$. When S is finite dimensional, A can be represented by an $n \times n$ matrix with complex entries. The identity operator I is represented by the identity matrix. The operator A is invertible iff its matrix has non-zero determinant. When invertible, the inverse operator A^{-1} is represented by its matrix inverse. An important class of operators are the unitary ones: an operator U is **unitary** if it has an inverse which given by its conjugate transpose $U^{-1} = U^*$. Equivalently, unitary means $U^*U = I$.

See the appendix for more information. All our functionals and operators are linear, and hence we normally drop the "linear" adjective.

Let $|x\rangle \in S$ and $|y\rangle \in T$ are vectors in two Hilbert spaces. Consider the map $L : S \to T$ that maps any $|v\rangle \in S$ to $\langle x|v\rangle|y\rangle$. We verify that L is a linear operator. A convenient notation for L is $|x\rangle\langle y|$, called the **outer product** of $|x\rangle$ and $\langle y|$. Applying L to $|v\rangle$ is then written as

$$(|x\rangle\langle y|)(|v\rangle) = |x\rangle\langle y|v\rangle = \langle y|v\rangle|y\rangle.$$

¶15. Quantum Mechanical System Postulates. Each quantum mechanical system (QMS) lives in some Hilbert space S. At each moment in time, the QMS has a definite **state**. These states can change in one of two ways: by **transformation** or **measurement**. The mathematical properties of these concepts are given in three postulates:

I. State Postulate: The (quantum) **states** of the quantum mechanical system (QMS) are elements $|x\rangle$ of a Hilbert space S with unit norm.

II. Transformation Postulate: **Transformations** of states correspond to applying a unitary operator U, taking $|x\rangle$ to $U|x\rangle$.

III. Measurement Postulate: **Measurement** of states correspond applying an observable $M = \{M_i : i \in I\}$ (see ¶8). When we measure a state $|x\rangle$ using M, this **collapses** the state to a new state $M_i|x\rangle/\sqrt{p_i(x)}$ with probability $p_i(x)$ (for each $i \in I$). Here, $p_i(x) = \langle x|M_i^* M_i|x\rangle$. (Check: $M_i|x\rangle/\sqrt{p_i(x)}$ is a state.)

IV. Composition of Quantum States: If S and T are the state space of two QMS's. Then their combined state space is the tensor product $S \otimes T$, another Hilbert space.

Remarks:
1. Where is time in these postulates? It is implicit in our postulates that U and M are fixed operators and they act "instantaneosuly". Time is just a totally ordered sequence of states that results from a transformation or a measurement. Thus time is **discrete** rather than continuous, and is implicitly recorded by the sequence of states. This is the computer science view. In physics it is more common to assume that a unitary operator U is a function of a continuous time variable t, written $U = U_t$. See [22, p. 83] to see how continuous time can be reduced to discrete time.
2. Postulates II and III concern change of states. While the change by unitary transformations is reversible (since U^{-1} is also unitary), the change by measurement is not. One way to describe the former is that U corresponds to transformations in a **closed system**.
3. For our purposes, it is sufficient to focus on the finite dimensional Hilbert spaces S. Henceforth, we identify S with $\mathbb{C}^m$ for some m. Relative to some basis set $e_1, \ldots, e_m$ for S, states of S have the form $\psi = \sum_{i=1}^m c_i e_i$ where $\sum_{i=1}^m |c_i|^2 = 1$.
4. The observable M corresponds to physical concepts such as linear momentum, energy or mass.
5. The fourth postulate introduces the tensor structure into quantum states – this leads to the wonderful strange world of quantum phenomenon, such as the Bell states. In particular, it rejects the classical method of combining states based on Cartesian product.

¶16. Structure of quantum states: tensor products. The states of complex QMS are obtained by combining simpler QMS's. The simplest example is the state of a quword, made up of the states of its qubits. The mathematical description of such combinations is the tensor product.

Given two finite dimensional Hilbert spaces S, T, we define a new Hilbert space $S \otimes T$ called their **tensor product**. The basic intuition of tensors is understood in the concrete setting of n-vectors: if $x \in \mathbb{C}^m$ and $y \in \mathbb{C}^n$, then their tensor product is $x \otimes y = (x_i y_j : i = 1, \ldots, m, j = 1, \ldots, n) \in \mathbb{C}^{mn}$. Note that when we write out $x \otimes y$ as a mn-vector, $z = (z_1, \ldots, z_{mn})$, we must have some fixed convention for identifying each component z_k with some $x_i y_j$. That is, we need a bijection $b : \{1, \ldots, m\} \times \{1, \ldots, n\} \to \{1, \ldots, mn\}$ so that $z_k = x_i y_j$ if $b(i, j) = k$. A standard bijection is $b(i, j) = (i - 1)n + j$.

Thus, the property (2) above is easily verified. Also, we have

$$(x + y) \otimes z = xz \otimes yz, \qquad z \otimes (x + y) = zx \otimes zy$$

Associativity of tensor products is clear. In terms of the state notation, we tend to write $|xy\rangle$ instead of $|x\rangle \otimes |y\rangle$. If H_i $(i = 1, \ldots, k)$ are Hilbert spaces, so is $H = H_1 \otimes H_2 \otimes \cdots \otimes H_k$. Moreover, if B_i is an orthonormal basis for H_i, then $B_1 \otimes B_2 \otimes \cdots \otimes B_k$ is an orthonormal basis for H.

In the Exercise, we see that tensor products $S \otimes T$ can be viewed as the space of bilinear forms on $S \times T$.

Let A, B be linear transformations on S and T respectively. We can also define their tensor produce $A \otimes B$ which would be a linear transformation on $S \otimes T$. In the concrete setting of n-vectors, A is an $m \times m$ matrix, and B B is an $n \times n$ matrix. Their **tensor product** $A \otimes B$ is then given a $mn \times mn$ matrix which is called the **Kronecker product** of A and B. If $A = (a_{ij})_{i,j=1}^m$ and $B = (b_{k\ell})_{k,\ell=1}^n$ then $A \otimes B$ is the block matrix given by $(a_{ij}B)_{i,j=1}^m$ (or $(b_{k\ell}A)_{k,\ell=1}^n$). It is easily verified that $A \otimes B$ is a linear transformation of $S \otimes T$. We also write $A^{\otimes n}$ for $A \otimes A \otimes \cdots \otimes A$ (n-fold tensor product of A). See below for an example, where $H_n = H^{\otimes n}$ where H is the Hadamard matrix.

Exercises

Exercise 4.1: (i) Determine the eigenvalues λ_1, λ_2 of H. Then determine the corresponding unit eigenvectors u_i for λ_i $(i = 1, 2)$.
(ii) Give the spectral representation of H.
(iii) Suppose you are given the qubit $|0\rangle$. Using the observable $H = \{H_0, H_1\}$, what are the possible outcomes and their probabilities? ♢

Exercise 4.2: (i) Prove that if A, B are linear transformations on $S = \mathbb{C}^m$ and $T = \mathbb{C}^n$ then $A \otimes B$ is a linear transformation on $S \otimes T$.
(ii) Prove the same result when the vectors in S, T are restricted to unit length. ♢

Exercise 4.3: An (S, T)-**bilinear form** is a function $f(x, y)$ taking $(x, y) \in S \times T$ to a complex number and satisfying the equation

$$f(ax + a'x', by + b'y') = abf(x, y) + ab'f(x, y') + a'bf(x', y) + a'b'f(x', y').$$

(i) Show the set $B(S, T)$ of all bilinear forms is a Hilbert space.
(ii) Show that $S \otimes T$ is isomorphic to $B(S, T)$. ♢

End Exercises

§5. Reversible Circuits

¶17. We know that quantum computation amounts to "unitary transformation of states in Hilbert spaces". But to organize such transformations into useful algorithms, we need computational structures. Computer science provides many of these structures; perhaps the lowest level structure here is at the circuit level. A circuit is basically a directed acyclic graph whose nodes represent the unitary transformations (called unitary gates). But before we study unitary gates, we propose to treat the special case of reversible gates.

As noted in the introduction, the study of reversible computation actually predates quantum computing. Another rational for introducing reversible logic in quantum computing is based on the fact that *any reversible computation can be simulated on a quantum computer with essentially the same complexity.* So an upper bound on the reversible complexity of a problem yields an upper bound of its quantum complexity. This is useful because reversible computation, being classical, is easier to understand.

¶18. Functions realizable by circuits. Let $\mathbb{B} = \{0, 1\}$. A **Boolean function** has the form $f : \mathbb{B}^m \to \mathbb{B}^n$. A (classical) **circuit** is a directed, ordered[5] acyclic graph whose nodes are partitioned into three sets: the **input nodes** (those with indegree 0), the **output nodes** (those with outdegree 0), and the rest, called **internal nodes**. At each internal node with indegree k and outdegree ℓ, we associate it with Boolean function $g : \mathbb{B}^k \to \mathbb{B}^\ell$. We call g a **gate** of the circuit. We further require the input nodes to have outdegree 1, and the output nodes to have indegree 1.

The graph edges are called **wires**. Each input or output node in incident to a unique wire; thus we may refer to these as the input and output wires. Given an assignment $\mathbf{a}$ of Boolean values to the input wires, we can inductively associate a Boolean value $w(\mathbf{a})$ to each wire w of the graph: call $w(\mathbf{a})$ the value computed at w.

A circuit C is said to **realize** a function $f : \mathbb{B}^m \to \mathbb{B}^n$ if the input nodes can be partitioned into two sets, $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{x}' = (x'_1, \ldots, x'_k)$, and the output nodes can be partitioned into two sets, $\mathbf{z} = (z_1, \ldots, z_n)$ and $\mathbf{z}' = (z'_1, \ldots, z'_\ell)$, and there exists an assignment of values $\mathbf{c} \in \mathbb{B}^k$ to $\mathbf{x}'$ such that, for all assignment of values $\mathbf{a} \in \mathbb{B}^m$ to $\mathbf{x}$, the values computed at $\mathbf{y}$ is given by $f(\mathbf{a})$.

In this definition, the input nodes $\mathbf{x}'$ are called **control nodes** and the output nodes $\mathbf{y}'$ are called **junk nodes**. The value $\mathbf{c}$ that we assign to the control nodes is the **preparation**. Intuitively, C realizes f means that, with suitable preparation of the control nodes, we can compute f. We say that C **strongly realizes** f if we can use an arbitrary preparation $\mathbf{c}$.

A **basis** B is just any set of Boolean functions. A circuit **over** a basis B is one whose gates are taken from B. The basis is **universal** if every Boolean function can be realized by a circuit over B.

A well-known universal basis is the set of Boolean functions $\{AND_{m,n}, OR_{m,n}, NOT : m, n \in \mathbb{N}\}$. Here $AND_{m,n}$ is the m input, n output function where each output line is 1 iff all the input lines are 1. The $OR_{m,n}$ function is similarly defined. These functions may also be denoted by the symbols

$$\wedge_{m,n}, \quad \vee_{m,n}, \quad \neg.$$

Remark: Our definition of circuits is slightly unconventional because of the need to handle control inputs and junk outputs, and because of its adaptation to reversible circuits.

¶19. Reversible functions. A Boolean function $f : \mathbb{B}^m \to \mathbb{B}^n$ is **reversible** if it is a $1-1$ function. Thus $m \leq n$. In particular, if $m = n$, then the reversible function is actually a permutation.

If f, g are Boolean functions, we shall say that f is **reducible** to g (denoted $f \leq g$) if for all circuits C, if C realizes g then C realizes f. Intuitively, f is reducible to g means that, by setting some inputs of g as control, and some outputs of g as junk, we obtain a realization of f.

Lemma 1. *Every Boolean function $f : \mathbb{B}^m \to \mathbb{B}^n$ is reducible to some permutation.*

Proof. If f is not reversible, we can reduce to a reversible function f'. The idea is to just reproduce the input in the output: define the function f'

$$f' : \mathbb{B}^m \to \mathbb{B}^{m+n}$$

where $f'(x) = (x, f(x))$. Clearly, $f \leq f'$.

Next, we reduce a reversible f' into some permutation f'' where

$$f'' : \mathbb{B}^{m+n} \to \mathbb{B}^{m+n}$$

where $f''(x, 0^n) = f'(x) = (x, f(x))$. Note that f'' is under-specified at this point, since we have left unspecified the values $f''(x, y)$ for $y \neq 0^n$. But it is clear that such f'' exists. Again, we see that $f' \leq f''$. **Q.E.D.**

¶20. Reversible XOR gate. The classical XOR gate is defined as follows: $XOR(x, y) = 0$ iff $x = y$. In infix notation, we write $x \oplus y$ for $XOR(x, y)$. It may be easier to understand XOR as addition modulo 2 (the symbol $\oplus$ is highly suggestive of this). This gate is clearly non-reversible. A simple variation gives us the **reversible exclusive-or** (denoted $T_2(x, y)$) with 2-inputs and 2-outputs where

$$T_2(x, y) = (x, x \oplus y).$$

This is also called the **controlled-NOT** gate as we can think of the first bit as a "control bit", and the second bit as "data bit", which is negated or not, depending on the control bit. The diagramatic representation[6] of this gate is seen in Figure 3(a).

[5] The graph is **ordered** if, at each internal node u, the set of incoming edges at u has a total ordering, and likewise the set outgoing edges at u has a total ordering. This ordering is needed for a correct definition of computation at the gates.

[6] Feynman [18] introduced this notation, but with the $\oplus$ node is written as an "X". The X-gate is a common alternative notation for the controlled-not gate.

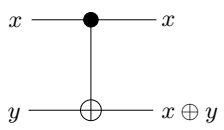

(a) Reversible $XOR(x, y)$

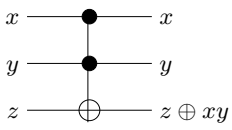

(b) Toffoli gate $T(x, y, z)$

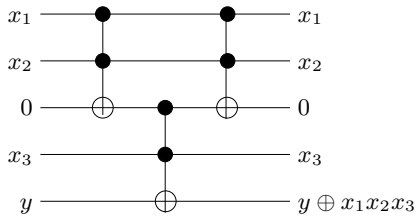

(c) $T_4(x_1, x_2, x_3, y; 0)$

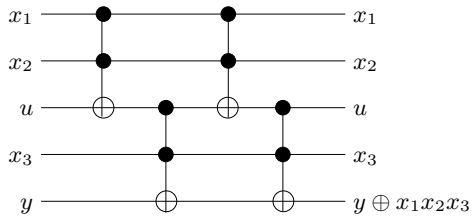

(d) $T'_4(x_1, x_2, x_3, y; u)$

Figure 3: (a) Controlled XOR, $T_2(x, y)$, (b) Toffoli Gate, $T(x, y, z)$, (c) Circuit for T_4, (d) Alternative circuit for T_4

There are three simple uses of T_2-gate. First, it can perform negation: $T_2(1, x) = (1, \neg x)$. Second, it can function as a "copier" of values. If the second wire of T_2 is prepared as 0, then the both output wires will contain the value of the first wire:

$$(x, 0) \mapsto T_2(x, 0) = (x, x). \tag{4}$$

A third application is for exchanging any two values. This is done[7] by arranging three of these gates in series, as

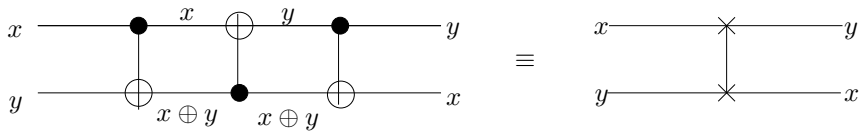

Figure 4: Exchange of x and y-bits, and corresponding gate symbol

in Figure 4 with the middle gate exchanging the roles of the control bit and data bit:

$$(x, y) \mapsto (x, x \oplus y) \mapsto (y, x \oplus y) \mapsto (y, x). \tag{5}$$

It follows that any permutation of the input bits can be achieved using this gate, as any permutation can be obtained as a product of transpositions (pairwise exchanges). Note that being able to permute bits does not mean the T_2-gate is universal: there are $n!$ bit functions that permute bits of $\mathbb{B}^n$ but there are $(2^n)!$ permutations of the set $\mathbb{B}^n$. Indeed, it can be shown that the set of all 1-bit and 2-bit reversible gates is not universal (see Exercise).

¶21. Scheme for drawing reversible circuits. We shall work mainly with bases B in which each gate $g \in B$ has input arity and output arity that are equal, $g : \mathbb{B}^k \to \mathbb{B}^\ell$ where $k = \ell$. A circuit over such bases will have the same number of input wires as output wires. Such circuits can be systematically diagramed as in Figure 3:

- If there are m input wires, we draw m horizontal lines. These **lines** can be given a useful interpretation – they represent a unit of storage space.
- Each gate is represented by a vertical bar: when the bar crosses a horizontal wire that is used by the gate, we mark this intersection with a suitable "intersection symbol". For the gates in Figure 3, note that we need

[7] Matt Johnson points out a programming analogy: in the `C` programming language, the instruction $x \mathrel{\wedge}= y \mathrel{\wedge}= x \mathrel{\wedge}= y$ will exchange the values of the Boolean variables x, y.

only two kinds of intersection symbols – a small black circle and somewhat larger $\oplus$ symbol. Each of the small black circle could be replaced by a white circle, indicating that the control line is activated iff the line contains the state $|1\rangle$. Also Figure 4 shows that an exchange gate can be indicated with two "x" symbols.

- Thus, the wires in such a diagram are represented by the horizontal segments between two consecutive intersection symbols.
- We conventionally think of values in the wires as flowing in the left-to-right direction. Call this the "forward" direction. So the input (output) wires are on the left (right) end of the lines. When the gates are reversible, we could equally have the values flow in the "backward" direction.

¶22. Toffoli gate. We now consider a 3-bit input gate called the **Toffoli gate** defined by

$$T(x, y, z) = (x, y, z \oplus xy).$$

It is clearly a generalization of the controlled-NOT gate; in turn we can generalize this to the n-input Toffoli gate,

$$T_n(x_1, x_2, \ldots, x_{n-1}, y) = (x_1, x_2, \ldots, x_{n-1}, y')$$

where $y' = y \oplus (x_1 x_2 \cdots x_{n-1})$. In other words, the last bit is complemented provided the first $n-1$ bits are 1. Thus the Toffoli gate may also be denoted T_3.

It is easy to see that $T_n \leq T_{n+1}$ for $n \geq 2$. This reduction is achieved by setting one of inputs of T_n to 1: $T_n(1, x_2, \ldots, x_n) = (1, T_{n_1}(x_2, \ldots, x_n))$. In particular, it means that T_n can negate bits, copy bits and exchange any two bits, just as T_2. Also, $T_n(T_n(x_1, \ldots, x_n)) = (x_1, \ldots, x_n)$, *i.e.*, T_n is its own inverse.

Figure 3(c) shows a circuit over T_3 that realizes T_4. Note that there is a control line and a junk line: the control line is prepared with 0, and this value is preserved in the junk line. The role of the third T-gate is to preserve this 0 bit. In Figure 3(d), we use another realization of T_4, using an extra Toffoli gate. In this case, the control line can be arbitrary (i.e., T_4 is strongly realized by this circuit). It is easy to generalize these constructions to obtain T_n.

Theorem 2. *The basis $\{T\}$ comprising only the Toffoli gate is universal.*

Proof. Since every function can be realized by a reversible function, it is sufficient to prove that every reversible $f : \mathbb{B}^m \to \mathbb{B}^n$ $(m \leq n)$ can be realized by toffoli gates. For $m = 0$, this is trivial: we do not need any gates, just n control wires and n output wires. The control wires are set to exactly the desired outputs.

Let $i = 0, 1$. For $m > 0$, let $f_i : \mathbb{B}^{m-1} \to \mathbb{B}^n$ be the functions given by $f_i(\mathbf{x}) = f(0, \mathbf{x})$ for all $\mathbf{x} \in \mathbb{B}^{m-1}$. By induction, let $F_i(x_2, \ldots, x_m)$ be a circuit over T realizing f_i.

We construct an "controlled exchange circuit" $E(c, a, b)$ that is analogous to the circuit in Figure 4, except that we replace the T_2 gates by Toffoli gates. Thus, an exchange occurs iff the first bit is 1, i.e.,

$$E(c, a, b) = \begin{cases} (1, b, a) & \text{if } \ c = 1, \\ (0, a, b) \text{ if } \ c = 0. \end{cases}$$

We use m such exchange circuits to select among the outputs of C_0 or C_1. More precisely, we construct a circuit $F(x_1, \ldots, x_m) = (y_1, \ldots, y_n)$ for $f : \mathbb{B}^m \to \mathbb{B}^n$ as follows. Assume we have two circuits $F_0(x_2, \ldots, x_m)$ and $F_1(x_2, \ldots, x_m)$ realizing f_0 and f_1, as described above. For each $j = 1, \ldots, n$, let $F_i[j] = F_i(x_2, \ldots, x_m)[j]$ denote the jth output wire of $F_i(x_2, \ldots, x_m)$. We put these wires into the exchange circuit as follows:

$$E(x_1, F_0[j], F_1[j]) \tag{6}$$

The jth output of $F(x_1, \ldots, x_m)$ is defined to be the second output of (6), i.e.,

$$y_j = E(x_1, F_0[j], F_1[j])[2].$$

Thus, if $x_1 = 0$, then $y_j = F_0[j]$, and if $x_1 = 1$, then $y_j = F_1[j]$. **Q.E.D.**

Suppose f is invertible. Then we would like the realizing circuit to have control wires which can be prepared with 0's, and whose junk wires only carry 0's. This is left as an exercise.

Exercises

Exercise 5.1: Prove the following extension of the universality theorem: suppose f is an invertible Boolean function. Then f is realized by a circuit over the basis $B = \{\neg, T_3\}$ in which all control wires are prepared with 0's and all junk bits are 0's. ♢

Exercise 5.2: Consider the function C_n with two n-quwords $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_n)$ as inputs. The output is given by $C_n(\mathbf{x}, \mathbf{y}) = (\min\{\mathbf{x}, \mathbf{y}\}, \max\{\mathbf{x}, \mathbf{y}\})$. The comparison is lexicographic (or you may interpret it as comparison of binary numbers). This function is not reversible.
(i) Construct a reversible circuit that realizes C_1. HINT: Be sure that you use only reversible gates in your construction (it is not enough to just ensure that the overall circuit computes a reversible function).
(ii) Generalize your solution in (i) to C_n. ♢

End Exercises

§6. Bennett's Scheme for Reversible Computation

¶23. Bennett [7] shows that *any algorithm for computing a function $F(x) = y$ can be converted into a reversible algorithm for computing the related $\widetilde{F}(x) = (x, y)$. Moreover, if F is computing in time T then $\widetilde{F}$ can be computed in time $O(T)$.* Bennett's proof uses the same ideas as the above proof on the universality of Toffoli gates, except that it is now expressed as a result on general computation. But we shall see that Bennett's computation scheme achieves an extra property.

As before, the basic idea is to keep a history of the computational steps in performing F. Begin with any machine M that computes some function

$$M(input_m) = (output_n).$$

The subscript m, n tells us how many bits are in the input and output arguments. We may think of M as a Turing machine. As is well-known in complexity theory, the computation of M can be encoded by a Boolean circuit

We first replace the circuit by a reversible circuit M_1. For instance, if C uses only AND and OR gates only, we can make them reversible by adding extra (control or junk) wires. This reversible circuit can be further realized using Toffoli gates. Also, each original input wire must be duplicated, as many times as it is used in the Boolean circuit. Although we intend to ignore the junk wires, we can think of the junk as a record of the computation. Hence, denote these junk wires by $record_k$ (there are k wires). We may denote the transformation of M_1 by

$$M_1 : (input_m, zero_{n+k+\ell}) \mapsto (output_n, record_k, zero_{m+\ell}). \tag{7}$$

Here we assume that $\ell \geq 0$ extra lines are needed to realize the original computation reversibly. This representation is completely general because in the circuits realized by Toffoli gates, we may assume that extra control wires is prepared with 0's, and the junk wires eventually contain zero.

In some sense, we are done. But let us impose an additional requirement: we want to reset all junk wires to zero. This property will be useful later: if we want to implement our scheme using quantum computers, then resetting $record_k$ to 0 is needed to to avoid interference when we read the desired $output_n$.

In (7), the values in $record_k$ are considered junk, and they are not necessarily zero. Hence, to satisfy our extra requirement, we add another n control wires, prepared to 0 – call these extra inputs $zeros_n$. They will hold our eventual output (recall the interpretation of "lines" in ¶21). We describe the overall scheme in three stages:

- First, we call M_1, but we re-interprete this transformation as

$$(input_m, zeros_k, zeros_n) \mapsto (input_m, zeros_k, output_n).$$

$$M_1 : (input_m, zero_n, zero_{n+k+\ell}) \mapsto (output_n, zero_n, record_k, zero_{m+\ell}), \tag{8}$$

  i.e., adding $zero_n$ as new control wires.

- Next, we make an extra copy of $output_n$ in the control wires $zeros_n$. This copying of values can be done using[8] T_2 gates, as in (4). So Stage B produces the transformation

$$M_2 : (output_n, zero_n, record_k, zero_{m+\ell}) \mapsto (output_n, output_n, record_k, zero_{m+\ell}). \tag{9}$$

- Finally, we run the reversible circuit M_1 backwards, with $(output_n, recork_k, zero_{m+\ell})$ as input! This produces $(input_m, zero_n, zero_{k+\ell})$. With suitable reordering, we write this transformation as

$$M_3 : (output_n, output_n, record_k, zero_{m+\ell}) \mapsto (input_m, output_n, zero_{n+k+\ell}). \tag{10}$$

[8]Note that T_2 requires no control lines. But if we use T_3 instead of T_2, we need an extra control line.

- Composing these three stages, we have the transformation

$$(input_m, zero_n, zero_{n+k+\ell}) \mapsto (input_m, output_n, zero_{n+k+\ell}).$$

By viewing the scratch space $zeros_{n+k+\ell}$ as internal to the machine, we have shown the reversible computation of the desired function

$$\widetilde{F} : (input_m, zeros_n) \mapsto (input_m, output_n).$$

Shor noted that Lecerf [20] obtained a result similar to Bennett's. Note that although time is not asymptotically increased by this scheme, space (i.e., extra wires) may be as bad as $\Omega(T)$. It is possible to improve the space utilization by trading it for increased time [9].

Exercises

Exercise 6.1: Consider reversible circuits to compute the following transformation: $f : \mathbb{B}^{2n} \to \mathbb{B}^{2n}$ where $f(x, y) = (x, y^2)$ where $0 \leq x, y < 2^n$ are n-bit binary numbers and y^2 is taken modulo 2^n. Construct the circuit explicitly for the case $n = 3$ using the family of T_n gates. ♢

Exercise 6.2: Design a full adder using T_2 and T_3 gates. This circuit has 4 wires a, b, c, d where a, b, c are the two input bits plus a carry-in bit, and d is set to 0. The output wires are a', b', c', d' where c' is the sum and d' the carry-out. Consider a', b' to be junk wires. Note: 4 gates suffices. ♢

Exercise 6.3: Show by a direct argument that T_3 cannot be constructed out of circuits involving T_2 and other two-input gates. ♢

Exercise 6.4: We will be computing in $\mathbb{Z}_2$ (modulo 2) in the following.
(i) Enumerate the 6 invertible 2×2 matrices.
(ii) Show that all invertible 2-input functions are linear: $(x, y) \mapsto (x, y)M + (a, b)$ where M is an invertible matrix from (i) and a, b are constants.
(iii) Show that any circuit composed of linear gates is linear.
(iv) Conclude that the set of reversible 2-inputs and 1-input gates is not universal.
(v) Give an nonlinear invertible function $f : \mathbb{B}^3 \to \mathbb{B}^3$. Show that it is nonlinear as well as give an implementation of f by Toffoli gates. ♢

Exercise 6.5:
(i) Show the recursive construction of T_n from T in which we use only one extra bit of scratch space. What is the number of T-gates used? (ii) Show that with 3 scratch bits, initiallized to 0, we can construct T_n with $2n - 5$ T-gates. ♢

Exercise 6.6: Show that any 4-input circuit composed from $T(x, y, z)$ must compute an even permutation of $\mathbb{B}^4$. Conclude that such a circuit cannot compute T_4. ♢

End Exercises

§7. Quantum Circuits

¶24. Unlike classical circuits, a quantum circuit has qubits instead of bits on each wire. It also uses quantum gates, which are a generalization of reversible gates. Each quantum gate is a unitary transformation on a fixed number of qubits. These unitary transformations are chosen from some finite basis. At the end of the computation, a measurement is made of some of the qubits. This "measurement at the end" is a canonical choice; in practice, one may wish to make measurements throughout the computation. Quantum circuits were introduced by David Deutsch in 1981.

A gate or function that manipulates n qubits is represented by a $2^n \times 2^n$ unitary matrix. The following **Hadamard matrix** corresponds to the Hadamard gate which operates on 1 qubit:

$$H := \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \tag{11}$$

It is easy to check that the matrix[9] H is unitary: $H^*H = I$. This gate transforms the eigenstates $|0\rangle$ and $|1\rangle$ as follows:

$$\begin{aligned} H : |0\rangle &\mapsto \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \\ |1\rangle &\mapsto \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle, \end{aligned}$$

The transformed states $H|0\rangle$ and $H|1\rangle$ are sometimes denoted $|+\rangle$ and $|-\rangle$, respectively. By viewing $|0\rangle$ and $|1\rangle$ as the column vectors $e_0 = (1,0)^T$ and $e_1 = (0,1)^T$, respectively, then transformations by H is just matrix-vector multiplication.

Recall that in the Stern-Gerlach device figure 1, we could change the basis of the measurement by just rotating the apparatus along the axis of the electron path. Rotating the device by 90° amounts to using the basis $H|0\rangle$ and $H|1\rangle$.

We can generalize H to

$$H_n := \underbrace{H \otimes H \otimes \cdots \otimes H}_{n}$$

which is a $2^n \times 2^n$ matrix. This is now a unitary operator on n-quwords. For instance,

$$H_2 = \frac{1}{2}\left[\begin{array}{cc|cc} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ \hline 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{array}\right]. \tag{12}$$

Transformation H_n will be used in various quantum algorithms to create a superposition of all possible 2^n states for subsequent computation. The gate representation of H_n amounts to placing a H-gate on each of the n wires, in parallel, in an n-quword circuit.

¶25. Convention: The operations of a $2^n \times 2^n$ unitary matrix M on states of n-quwords uses a natural convention which may be worth spelling out. The eigenstates are ordered lexicographically, in the order $|0^n\rangle, |0^{n-1}1\rangle, \ldots, |1^n\rangle$ or equivalently,

$$|0\rangle, |1\rangle, |2\rangle, \ldots, |2^n - 1\rangle.$$

This ordering is used to label the rows and columns of M. For $i = 0, \ldots, 2^n - 1$, the elementary 2^n-vector e_i which has 1 in the ith position and 0 elsewhere represents the eigenstate $|i\rangle$. The state x of the quword is just a 2^n-vector of length 1. The transformation M operates on state x by matrix-vector multiplication, Mx. For instance,

$$H_n|\underbrace{00\cdots0}_{n}\rangle = \frac{1}{2^{n/2}} \sum_{i=0}^{2^n-1} |i\rangle.$$

More generally:

Lemma 3.

$$H_n(|j_1 j_2 \cdots j_n\rangle) = \frac{1}{2^{n/2}} \sum_{i=0}^{2^n-1} (-1)^{\langle i,j\rangle} |i\rangle.$$

where $\langle i,j\rangle = \sum_{k=1}^n i_k j_k$ *denotes scalar product of two* n*-vectors.*

We leave this demonstration to an Exercise.

¶26. Unitary Matrices of Reversible Circuits. We said that reversible circuits are special cases of quantum circuits. It is instructive to see[10] what this means.

Lemma 4. *Reversible gates are quantum gates.*

[9] Hadamard matrices are often defined as non-singular matrices whose entries are ±1. But to get unitary matrices here, we scale the ±1 entries by the factor $1/\sqrt{2}$. Another name for Hadamard transformations is "Hadamard-Walsh transforms".

[10] There is also the converse statement "quantum gates are reversible", which is also true. Taken together, this might yield the wrong conclusion that quantum gates and reversible gates are identical concepts! The reason for this is, of course, is the quirk in our terminology: "g is a reversible gate" has a special meaning and means *more* than just saying "gate g is reversible".

Proof. Formally, a reversible gate computes a classical function $f : \mathbb{B}^n \to \mathbb{B}^n$. This gives rise to a permutation matrix U_f of size $2^n \times 2^n$. Permutation matrices, by definition, have exactly one 1 on each row and each column, and has 0's everywhere else. This lemma amounts to the claim that $U = U_f$ is unitary, *i.e.*, $UU^* = I$. This means $(UU^*)_{ij} = 1$ if $i = j$ and 0 otherwise. If $i = j$ then the ith row of U and the jth column of U^* are the same and hence their scalar product is 1. If $i \neq j$, then the unique 1 on the ith row will not be matched up with the unique 1 on the jth column, and so their scalar product is 0. **Q.E.D.**

Another way to interprete the lemma is to say that a classical reversible gate specifies transformation on eigenstates, and this becomes a quantum gate if we extend the transformation to superposition of eigenstates, using linearity.

In illustration, consider the Toffoli gate $T(x, y, z) = T_3(x, y, z)$. The corresponding unitary 8×8 matrix U_3 operating on linear combinations of the eigenstates $|xyz\rangle = |x\rangle \otimes |y\rangle \otimes |z\rangle$. We have

$$U_3(|x\rangle \otimes |y\rangle \otimes |z\rangle) = U_3(|xyz\rangle) = \begin{cases} |xyz\rangle & \text{if } x = 0 \text{ or } y = 0, \\ |111\rangle & \text{if } x = y = 1, z = 0, \\ |110\rangle & \text{if } x = y = z = 1. \end{cases}$$

With the usual labeling conventions, U_3 becomes the following matrix

$$U_3 = \left[\begin{array}{cccccc|cc} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{array}\right] \tag{13}$$

When we compose gates into circuits, we are building unitary matrices from the unitary matrices represented by the individual gates in a natural way. For instance, if $|x\rangle = |x_1x_2x_3\rangle$ is the state of 3 qubits that are being operated on by T_3, and $|y\rangle = |x_4 \cdots x_n\rangle$ is the state of the rest of the qubits, then the overall unitary matrix corresponding to this transformation is the tensor product

$$V_3 = U_3 \otimes I$$

where I is a $2^{n-3} \times 2^{n-3}$ identity matrix. We note: permutation matrices are closed under tensor products. The universality of T_3 implies that every permutation matrix can written as a product of matrices such as V_3 (which are extensions of U_3 to perform transformations on n qubits).

Deutsch shows the existence of a universal quantum gate $D(x, y, z)$. It is a generalization of Toffoli's gate: if $x, y, z \in \mathbb{B}$ then $D(|x, y, z\rangle) = |x, y, z\rangle$ if both $x \wedge y \neq 1$; otherwise $D(x, y, z) = |x, y, U(z)\rangle$ where U is a 2×2 unitary transformation.

Exercises

Exercise 7.1: A $2^n \times 2^n$ unitary matrix U represents a classical gate iff it is a permutation matrix. ♢

End Exercises

§8. Structure of 2×2 Unitary Matrices.

¶27. Bloch Sphere. This is a scheme for visualizing the states of a qubit as a point on the sphere unit 2-sphere S^2. Let

$$|x\rangle = c_0|0\rangle + c_1|1\rangle$$

where $|c_0|^2 + |c_1|^2 = 1$. Then we can write $c_0 = e^{\mathbf{i}\gamma} \cos(\theta/2)$ and $c_1 = e^{\mathbf{i}(\gamma+\phi)} \sin(\theta/2)$ for some γ, ϕ, θ. Hence

$$|x\rangle = e^{\mathbf{i}\gamma} \left(\cos(\theta/2)|0\rangle + e^{\mathbf{i}\phi} \sin(\theta/2)|1\rangle \right). \tag{14}$$

Note that we use "$\theta/2$" instead of "θ" in this formula. This expression is unique unless $c_0 = 0$. Call γ the **phase angle** of $|x\rangle$. The state $|x\rangle$ is normalized by removing this phase angle; wwo states as **equivalent** if they are

equal upon normalization. Define the map $\mu : |x\rangle \mapsto (\phi, \theta)$ which acts on equivalence classes of states. The map is uniquely well-defined except when $|x\rangle \equiv |0\rangle$ or $|x\rangle \equiv |1\rangle$: in these exceptional cases, μ is multivalued, with $\mu(|1\rangle) = (\phi, \pi)$, and $\mu(|0\rangle) = (\phi, 0)$ for any ϕ.

Figure 5 shows the unit 2-sphere, S^2. Let us first establist an (x, y, z)-coordinate system. Pick two distinguished points on the sphere: N (North Pole) and E ("East Pole") where the distance between N and E along a great arc is $\pi/2$. We may introduce points S (South Pole) and W, diametrically opposite N and E (respectively). The North Pole N points in the positive z direction, and E with the positive y direction. The positive x direction must be orthogonal to the y and z directions; this is uniquely determined by requiring (x, y, z) to be right-hand system (see Figure 5).

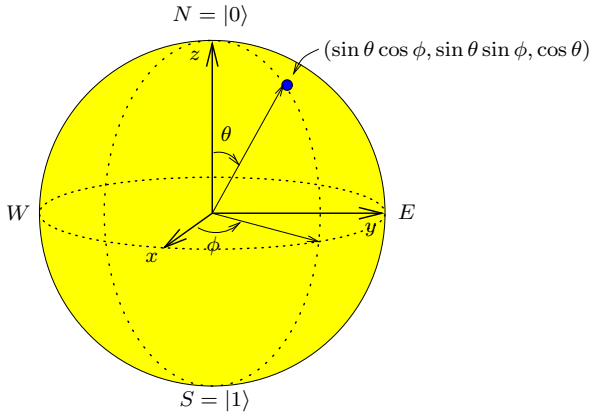

Figure 5: Bloch sphere. [MUST FIX: replace ϕ by its 90 deg-complement, label $(\cos\phi, \sin\phi, 0)$ on the Equator, and $(\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$ corresponds to (ϕ, θ)]

Like the surface of the earth, S^2 can be given a coordinate system based on two angles. This is the (ϕ, θ) coordinate system. We may identify ϕ with **longitude lines** and θ with **latitude lines** as in conventional earth maps. The longitude $\phi = 0$ corresponds to the half-circle from N to S passing through the x-axis. This is the Greenwich meridian. The longitude $\phi = \pi/2$ passes through the point E. However, we measure latitude θ somewhat unconventionally, with the Equator (the great circle passing through E and W) corresponding to $\theta = \pi/2$; also $\theta = 0, \pi$ corresponding to N and S. Without loss of generality, assume $\phi \in [0, 2\pi)$ and $\theta \in [0, \pi]$.

So each qubit state $|x\rangle$ corresponds to a point $\mu(|x\rangle)$ on S^2; call S^2 the **Bloch sphere** under this correspondence. In particular, observe that the states $|0\rangle$ and $|1\rangle$ correspond to the North and South Poles, respectively.

In the (x, y, z)-coordinate system, the point (ϕ, θ) becomes $(\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$.

¶28. Decomposition of qubit gates. Recall that a quantum gate for a single qubit corresponds to a 2×2 unitary matrix. An example is the Hadamard gate H. We now derive a decomposition of such matrices Barenco et al [1]. In particular, each unitary 2×2 matrices will be written as a product of three simpler unitary matrices

Let $U = [x|y]$ where $x, y \in \mathbb{C}^2$ are column vectors. Suppose $x = (re^{\mathbf{i}\theta}, r'e^{\mathbf{i}\theta'})^T$ where r, θ, r', θ' are real and $\mathbf{i} = \sqrt{-1}$. Now U is unitary implies $|x|^2 = |y|^2 = 1$ and $x^*y = 0$. From $|x|^2 = 1$, we obtain $r^2 + r'^2 = 1$ and hence $x = (\cos\alpha e^{\mathbf{i}\theta}, \sin\alpha e^{\mathbf{i}\theta'})^T$ for some real α. Similarly, let $y = (\sin\beta e^{\mathbf{i}\psi}, \cos\beta e^{\mathbf{i}\psi'})^T$ for some real β, ψ, ψ'. Next, $x^*y = 0$ implies $\cos\alpha e^{-\mathbf{i}\theta}\sin\beta e^{\mathbf{i}\psi} + \sin\alpha e^{-\mathbf{i}\theta'}\cos\beta e^{\mathbf{i}\psi'} = 0$, or in matrix notation,

$$\begin{bmatrix} \cos(\psi-\theta) & \cos(\psi'-\theta') \\ \sin(\psi-\theta) & \sin(\psi'-\theta') \end{bmatrix} \begin{pmatrix} \cos\alpha\sin\beta \\ \sin\alpha\cos\beta \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Thus the determinant of the 2×2 matrix vanishes,

$$\cos(\psi-\theta)\sin(\psi'-\theta') - \sin(\psi-\theta)\cos(\psi'-\theta') = 0.$$

Writing $2\delta = \theta' - \theta$ and $2\delta' = \psi' - \psi$, this gives $\sin(2\delta' - 2\delta) = 0$. Thus $\delta = \delta'$ and we have

$$x = e^{\mathbf{i}(\theta+\delta)}(\cos\alpha e^{-\mathbf{i}\delta}, \sin\alpha e^{\mathbf{i}\delta})^T, \qquad y = e^{\mathbf{i}(\psi+\delta)}(\sin\beta e^{-\mathbf{i}\delta}, \cos\beta e^{\mathbf{i}\delta})^T.$$

The matrix $U = [x|y]$ may be written as a product $U = ABC$ where A, B, C are the matrices

$$A = \begin{bmatrix} e^{-\mathbf{i}\delta} & 0 \\ 0 & e^{\mathbf{i}\delta} \end{bmatrix}, \qquad B = \begin{bmatrix} \cos\alpha & \sin\beta \\ \sin\alpha & \cos\beta \end{bmatrix}, \qquad C = \begin{bmatrix} e^{\mathbf{i}(\theta+\delta)} & 0 \\ 0 & e^{\mathbf{i}(\psi+\delta)} \end{bmatrix}.$$

Since A, C are unitary, it follows from $I = U^*U = C^*B^*A^*ABC = C^*B^*BC$ that B must be also be unitary. Thus $\cos\alpha \sin\alpha + \cos\beta \sin\beta = 0$. This implies $\alpha = -\beta$.

$$B = \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix}.$$

Note that $\det(U) = \pm 1$. But if we further assume that U is special, *i.e.*, $\det(U) = 1$. Since $\det(A) = \det(B) = 1$, we have $\det(C) = 1$. This means $e^{\mathbf{i}(\theta+\psi+2\delta)} = 1$ or $\theta + \delta = -(\psi + \delta)$. By renaming $\theta + \delta$ to θ and $\psi + \delta$ to ψ, we obtain the result of Bloch:

Theorem 5. *Every unitary matrix $U \in \mathbb{C}^{2\times 2}$ has the form*

$$U = \begin{bmatrix} e^{-\mathbf{i}\delta} & 0 \\ 0 & e^{\mathbf{i}\delta} \end{bmatrix} \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \begin{bmatrix} e^{\mathbf{i}\theta} & 0 \\ 0 & e^{\mathbf{i}\psi} \end{bmatrix}.$$

If U is special, then $\psi = -\theta$.

¶29. Transformations of the Bloch sphere. We can now visualize each qubit gate or 2×2 unitary matrix: *it corresponds to a transformation of the Bloch sphere.* For instance, the Hadamard gate H corresponds to rotating S^2 through $90°$ about the y-axis, and then reflecting across the xy-plane. But without some tools, it is non-trivial to verify such a remark, not to speak of discovering such an intepretation of H or any given unitary matrix.

The tool to help us is the exponential map: we define the **exponentiation** of a square matrix A as

$$e^A = \exp(A) := \sum_{n=0}^{\infty} \frac{A^n}{n!} = I + A + \frac{A^2}{2} + \frac{A^3}{3!} + \cdots,$$

where I is the identity matrix. This is a direct analogy with the exponential function $e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!}$ for any complez number z. In case $A^2 = I$ and x is any real number, we may verify that

$$\exp(\mathbf{i}Ax) = \cos(x)I + \mathbf{i}\sin(x)A. \tag{15}$$

Note that if A is Hermitian ($A^* = A$) and unitary ($A^*A = I$), then $A^2 = I$.

The following three unitary matrices are called the **Pauli matrices**:

$$\sigma_x = X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad \sigma_y = Y = \begin{bmatrix} 0 & -\mathbf{i} \\ \mathbf{i} & 0 \end{bmatrix}, \qquad \sigma_z = Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \tag{16}$$

Sometimes they are also denoted $\sigma_1, \sigma_2, \sigma_3$, respectively. Also, σ_0 can be taken as the identity matrix.

Observe that the Pauli matrices are Hermitian as well as unitary. Hence, if we define

$$R_j(\theta) = e^{-\mathbf{i}\sigma_j\theta/2}$$

where $j = x, y, z$, then we have

$$\left.\begin{array}{rcll} R_x(\theta) & = & \cos(\theta/2)I - \mathbf{i}X\sin(\theta/2) & = \begin{bmatrix} \cos(\theta/2) & -\mathbf{i}\sin(\theta/2) \\ -\mathbf{i}\sin(\theta/2) & \cos(\theta/2) \end{bmatrix} \\ R_y(\theta) & = & \cos(\theta/2)I - \mathbf{i}Y\sin(\theta/2) & = \begin{bmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{bmatrix} \\ R_z(\theta) & = & \cos(\theta/2)I - \mathbf{i}Z\sin(\theta/2) & = \begin{bmatrix} e^{-\mathbf{i}\theta/2} & 0 \\ 0 & e^{\mathbf{i}\theta/2} \end{bmatrix} \end{array}\right\} \tag{17}$$

The matrix $R_j(\theta)$ corresponds to rotation about the j-axis by angle θ.

Example: let $v = (x, y)^T \in \mathbb{C}^2$ be a unit vector representing the quantum state $x|0\rangle + y|1\rangle$. and let $\mu(x, y) = (\phi, \theta)$. Consider the matrix $R_x(\pi) = \begin{bmatrix} x & x \\ y & \\ y & \end{bmatrix}$. Then $R_x v =$??.

We can now re-interpret Theorem 5 using such rotation matrices:

Theorem 6 (YZ-Decomposition). *Every unitary matrix $U \in \mathbb{C}^{2\times 2}$ can be written as the product*

$$U = e^{\mathbf{i}\beta} R_z(2\delta) R_y(2\alpha) R_z(2\gamma) \tag{18}$$

where $\alpha, \beta, \gamma, \delta$ are real numbers. In other words, U is the composition (in suitable order) of a phase shift, two rotations about the Z-axis and a rotation about the Y-axis.

Proof. By comparing $R_y(\theta)$ and $R_z(\theta)$ in (17) with the decomposition $U = ABC$ of Theorem 5, we see that $A = R_z(2\delta)$ and $B = R_y(2\alpha)$. What about C? First, we choose $\beta := \frac{\theta+\psi}{2}$ and $\gamma := \frac{\psi-\theta}{2}$. Hence $\beta - \gamma = \theta$ and $\beta + \gamma = \psi$ and we may rewrite C as

$$C = \begin{bmatrix} e^{\mathbf{i}\theta} & 0 \\ 0 & e^{\mathbf{i}\psi} \end{bmatrix} = e^{\mathbf{i}\beta} \begin{bmatrix} e^{-\mathbf{i}\gamma} & 0 \\ 0 & e^{\mathbf{i}\gamma} \end{bmatrix} = e^{\mathbf{i}\beta} R_z(2\gamma).$$

Q.E.D.

Remark: multiplying by $e^{\mathbf{i}\beta}$ is called a **phase shift** by β.

¶30. Quantum Toffoli gates. Let $U = \begin{bmatrix} a_0 & b_0 \\ a_1 & b_1 \end{bmatrix}$ be any unitary matrix. Following [1], for any $n \geq 1$, we define the $2^n \times 2^n$ matrix

$$T_{n-1}(U) = \begin{bmatrix} 1 & & & & & & \\ & 1 & & & & & \\ & & \ddots & & & & \\ & & & 1 & & & \\ & & & & a_0 & b_0 & \\ & & & & a_1 & b_1 & \end{bmatrix}.$$

This is clearly unitary. This generalizes the Tofolli gate $T_n = T_n(x_1, \ldots, x_{n-1}, y)$ above. Indeed, T_n is just $T_{n-1}(X)$ where $X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ is a Pauli matrix. We can similarly describe the action of $T_{n-1}(U)$ on an eigenstate $|x_1 \cdots x_n\rangle$ as follows: $T_n(U)(|x_1 \cdots x_{n-1}, y\rangle) = |x_1 \cdots x_{n-1}, y'\rangle$ where

$$y' = \begin{cases} x_n & \text{if } x_i = 0 \text{ for some } i = 1, \ldots, n-1 \\ a_{y_n}|0\rangle + b_{y_n}|1\rangle & \text{if } x_i = 1 \text{ for all } i = 1, \ldots, n-1 \end{cases}$$

Thus the last qubit is transformed according to U iff the first $n-1$ qubits are 1. The diagramatic representation of such gates is shown in figure 6(a).

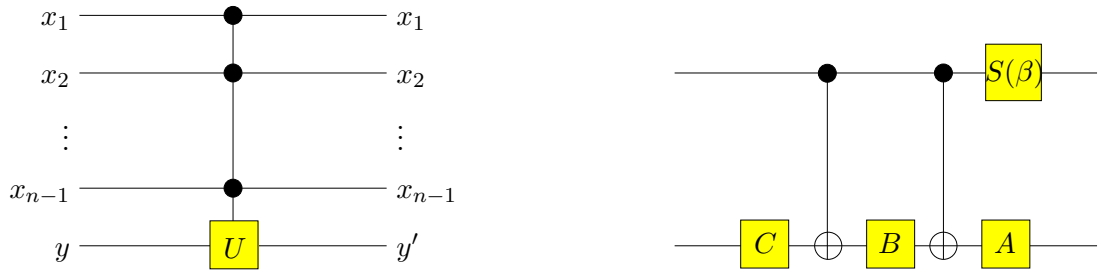

(a) $T_{n-1}(U)$ (b) $T_1(U)$

Figure 6: Quantum Toffoli gate: (a) $T_n(U)$, (b) realization of $T_1(U)$.

¶31. Implementation of the Control-U gate. Given a unitary U matrix, how can we realize the gate $T_n(U)$? It is an easy exercise to realize $T_n(U)$ using a $T_1(U)$ gate and some control-NOT's. The following theorem shows how to realize $T_1(U)$ assuming we have gates corresponding Y- and Z-rotations associated with U:

Theorem 7. *If $U \in \mathbb{C}^{2\times 2}$ is unitary, then there exists unitary matrices A, B, C and a real $\beta > 0$ such that $ABC = I$ and*

$$U = e^{\mathbf{i}\beta} AXBXC$$

where $X = \sigma_x$. Moreover, there are real numbers α, γ, δ such that

$$\begin{aligned} A &= R_z(2\delta)R_y(\alpha), && (19)\\ B &= R_y(-\alpha)R_z(-\gamma-\delta), && (20)\\ C &= R_z(\gamma-\delta). && (21) \end{aligned}$$

Proof. We use the decomposition (18) of Theorem 6. Let $\alpha, \beta, \gamma, \delta$ be chosen as in that theorem. It is easily checked that $ABC = I$. It remains to show that $U = e^{\mathbf{i}\beta} AXBXC$.

We exploit the ability of the X matrix to reverse the angle of rotation. Note that YX exchanges the columns of Y, and $X(YX)$ exchanges the rows of YX. Hence we see that $XYYX = -Y$. Similarly, $XZZX = -Z$. Then

$$\begin{aligned} XR_y(\theta)X &= X\left(\cos(\theta/2)I - \mathbf{i}\sin(\theta/2)Y\right)X \\ &= \cos(\theta/2)I + \mathbf{i}\sin(\theta/2)Y \\ &= R_y(-\theta). \end{aligned}$$

Similarly, $XR_z(\theta)X = R_z(-\theta)$.

From these elementary properties, we see that $XBX = R_y(\alpha)R_z(\gamma+\delta)$. This yields our desired conclusion

$$AXBXC = A(XBX)C = [R_z(2\delta)R_y(\alpha)] \cdot [R_y(\alpha)R_z(\gamma+\delta)] \cdot [R_z(\gamma-\delta)] = R_z(2\delta)R_y(2\alpha)R_z(2\gamma).$$

Q.E.D.

We can now realize the $T_1(U)$ gate using gates for A, B, C and some control-NOT (i.e., control-X) gates. But we also need a phase shift gate,

$$S(\beta) = \begin{bmatrix} 1 & 0 \\ 0 & e^{\mathbf{i}\beta} \end{bmatrix}$$

This is seen in Figure 6(b).

In the Exercise, we show that $T_2(U)$ is universal for almost all U. In contrast, there are no 2-input reversible gates that are universal for reversible circuits. Also, no work-bits are needed in quantum circuits, in contrast to some reversible circuits.

Exercises

Exercise 8.1: Give the Bloch decomposition (theorem 5) of the Hadamard matrix H and also the Pauli matrices $\sigma_x, \sigma_y, \sigma_z$. ♢

Exercise 8.2: Verify that the matrix $R_j(\theta)$ $(j = x, y, z)$ corresponds to rotation about the j-axis by angle θ. ♢

Exercise 8.3: Consider the Pauli matrices σ_i $(i = x, y, z$ or $i = 1, 2, 3)$. Show
(i) The eigenvalues of σ_i are ± 1.
(ii) $\sigma_{i-1}\sigma_i = \mathbf{i}\sigma_{i+1}$, where subscript addition is modulo 3. ♢

Exercise 8.4: Give the following variation on Theorem 5, where U is written as a product of R_x and R_y matrices. ♢

Exercise 8.5: (i) What points on the Bloch sphere correspond to the following states:

$$\begin{aligned} |+\rangle = H|0\rangle &= \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \\ |-\rangle = H|1\rangle &= \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \\ |a\rangle &= \sqrt{\frac{3}{4}}|0\rangle + \sqrt{\frac{1}{4}}|1\rangle \\ |b\rangle &= \sqrt{\frac{3}{4}}|0\rangle \end{aligned}$$

(ii) Prove that the action of the Hadamard gate H on the Bloch sphere has the description (rotation followed by a reflection) given in the text.
(iii) Similarly, describe the actions of the Dirac matrices $\sigma_x, \sigma_y, \sigma_z$ on the Bloch sphere. ♢

Exercise 8.6: Prove that $T_2(U)$ is universal for "most" unitary 2×2 U. ♢

End Exercises

§9. No Cloning Theorem

Before we leave this topic of quantum circuits, we prove a simple but fundamental result about unitary transformations from Wootters and Zurek (1982). This result says that we cannot copy an unknown quantum state perfectly. But first, recall the controlled XOR gate, $T_2(x, y)$. We have

$$T_2(x, 0) = (x, x).$$

Thus, the bit x has been copied or cloned. The next result shows that this is impossible with a quantum gate.

Theorem 8 (No Cloning). *Let $S = S_1 \otimes S_1$ be a Hilbert space where S_1 is of dimension at least 2. There does not exist a unitary transformation U such that for all $x \in S_1$,*

$$U(|x, 0\rangle) = |x, x\rangle$$

Proof. By way of contradiction, suppose U exists. Let $|x\rangle, |y\rangle$ be two linearly independent vectors of S_1. Then $U(|x, 0\rangle) = |x, x\rangle$ and $U(|y, 0\rangle) = |y, y\rangle$. If $z = \frac{1}{\sqrt{2}}(|x\rangle + |y\rangle)$, then

$$U(|z, 0\rangle) = \frac{1}{\sqrt{2}}(|x, x\rangle + |y, y\rangle)$$

by linearity of U. However, we note that

$$|z, z\rangle = |z\rangle \otimes |z\rangle = \frac{1}{2}(|x\rangle + |y\rangle) \times (|x\rangle + |y\rangle) = \frac{1}{2}(|x, x\rangle + |y, x\rangle + |y, x\rangle + |y, y\rangle).$$

But the vectors $|x, x\rangle, |y, x\rangle, |y, x\rangle, |y, y\rangle$ are linearly independent in S means that any two linear combinations of these vectors are equal if and only if their coefficients are identical. Thus, $U(|z, 0\rangle) \neq |z, z\rangle$. **Q.E.D.**

This result seems to conflict with our previous result that says reversible gates such as T_2 can be viewed as quantum gates. Yet, we said T_2 can copy the first bit. The resolution of this apparent paradox is again linguistic: when viewed as a classical gate T_2 can copy the first bit. This breaks down when T_2 is viewed as a quantum gate. Thus $T_2(|x, 0\rangle) \neq |x, x\rangle$ when x is no longer a classical bit.

One consequence of this theorem is in quantum cryptography: it implies that secure quantum key generation is possible. However, the no-cloning theorem says nothing about the possibility of good but imperfect copying. Indeed such "weak copying" techniques have been proposed.

Exercises

Exercise 9.1: Fill in a detail in the No Cloning theorem: if $|x\rangle, |y\rangle$ are linearly independent in S_1 then $|xx\rangle, |xy\rangle, |yx\rangle, |yy\rangle$ are linearly independent in S. ♢

End Exercises

§10. Superdense Coding and Teleportation

We give two surprising applications of Bell states or EPR pairs.[11] These two applications are converses of each other, so it is instructive to see them together. First, we must distinguish two kinds of channels for transmitting information: a channel is **classical** if it transmits classical bits, and **quantum** if it transmits quantum states. We have many examples of classical channels, from radio, emails to telephones. Quantum channels have been

[11] Named after John S. Bell (1964) and A. Einstein, N. Rosen and B. Podolsky (1935).

demonstrated in experimental setups. The two applications concern this question: *if we have one kind of channel, how can we use it to effectively transmit data of the other kind?*

In both applications, there is a preliminary setup. There is a 2-quword that is in the Bell state

$$|\beta_{00}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle).$$

Alice holds the first qubit (named a), and Bob holds the second qubit (named b). It does not matter how Alice and Bob come into possession of these two bits. Perhaps they met before to get their respective bits entangled in this way. Or, we could imagine a third party preparing this ab pair, and sending a to Alice and b to Bob. It also does not matter how far apart are Alice and Bob; perhaps they are in different galaxies.

One way to produce this Bell state is to begin with the ab-quword in state $|00\rangle$, subject a to the H-gate and then subject b to the control-NOT (*i.e.*, T_2) gate (using a as control wire). See figure 7 (a).

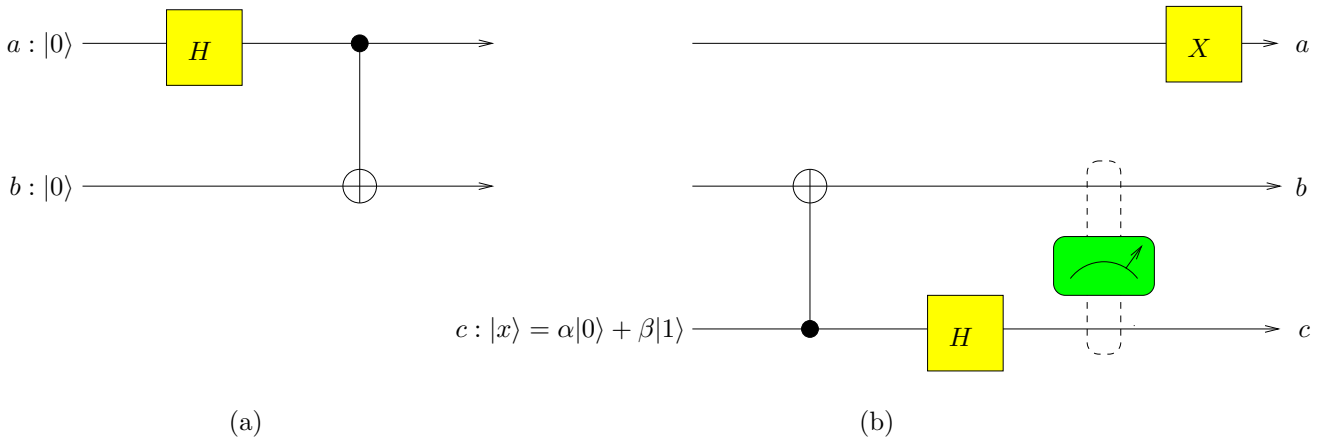

Figure 7: (a) Producing a Bell State, (b) Teleporting state $|x\rangle$.

The important thing to remember is that as long as the two qubits remain isolated from the outside world, they remain "entangled" regardless of physical distance. For instance if Bob measures his b-qubit and sees a 0, then Alice's a-qubit will instantaneously turn into a 0. In some sense, this information has been transmitted faster than the speed of light. This, by itself, does not seem to do anything useful in real life.

¶32. Transporting One Quantum Bit. Let us begin with the easier problem: suppose Alice and Bob shares a quantum channel, and Alice wants to transmit two classical bits to Bob. We show Alice only need to transmit one qubit. This scheme is from Bennett and Wiesner (1992), is called **superdense coding**. Experiments by Mattle, Weinfurter, Kwiat, Zeilinger (1996) have partially verified a form of this concept.

Suppose Alice wants to transmit to Bob the two classical bits $x, y \in \{0, 1\}$. She subjects her a qubit to the following transformation: If $x = 1$, then phase flip her bit. Next, if $y = 1$, then negate her bit.

[Draw Circuit here]

x	y	Results of Alice's Operations
0	0	$\frac{1}{\sqrt{2}}(\|00\rangle + \|11\rangle)$
0	1	$\frac{1}{\sqrt{2}}(\|10\rangle + \|01\rangle)$
1	0	$\frac{1}{\sqrt{2}}(\|00\rangle - \|11\rangle)$
1	1	$\frac{1}{\sqrt{2}}(\|10\rangle - \|01\rangle)$

Now she sends her qubit a to Bob. Upon receiving a, Bob subjects the ab quword in his possession to the following transformation

$$B = \begin{bmatrix} \frac{1}{\sqrt{2}} & & & \frac{1}{\sqrt{2}} \\ & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \\ \frac{1}{\sqrt{2}} & & & -\frac{1}{\sqrt{2}} \\ & -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \end{bmatrix}.$$

We omitted the 0 entries of the matrix. Check that B is unitary and Bob's quword ends up in the state $|xy\rangle$. He can now observe each qubit in turn, to recover x and y.

¶33. Transporting Two Classical Bits. Conversely, suppose Alice and Bob has a classical channel, and Bob wishes to send a qubit to Alice. This application is called **teleportation**, and comes from Bennett, Brassard, Crépeau, Jozsa, Peres and Wootters (1993), and various experiments have verified this.

As before, Alice and Bob holds one half of an ab quword in the Bell state $|\beta_{00}\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. Bob wants to send Alice a quantum state $|x\rangle$ of a brand new qubit named c. We could assume that Bob knows nothing about this state. Is there any way for Bob to convey this state to Alice by using a classical channel? We show that it is enough for Bob to send two classical bits to Alice.

Here is the outline of the protocol. First, Bob entangles the b and c qubits, then performs a measurement of b and c, and finally sends the results of this measurement to Alice. Based on this information, Alice performs a suitable transformation of her a-qubit and behold, the state of the a qubit would be the unknown $|x\rangle$.

Here are the details. Refer to the illustration in Figure 7(b). Let $|x\rangle = \alpha|0\rangle + \beta|1\rangle$ for unknown $\alpha, \beta \in \mathbb{C}$. Then the state of the abc-quword is initially

$$\begin{aligned} |x_0\rangle &:= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \otimes (\alpha|0\rangle + \beta|1\rangle) \\ &= \frac{1}{\sqrt{2}}\left\{\alpha(|000\rangle + |110\rangle) + \beta(|001\rangle + |111\rangle)\right\}. \end{aligned}$$

- Bob first performs a control-NOT to the b-qubit (using the c-qubit as control). This produces the state

$$|x_1\rangle := \frac{1}{\sqrt{2}}\left\{\alpha(|000\rangle + |110\rangle) + \beta(|011\rangle + |101\rangle)\right\}.$$

- Then Bob subjects c to an H-gate, to produce the state

$$\begin{aligned} |x_2\rangle &:= \frac{1}{2}\left\{\alpha(|00\rangle + |11\rangle) \otimes (|0\rangle + |1\rangle) + \beta(|01\rangle + |10\rangle) \otimes (|0\rangle - |1\rangle)\right\} \\ &= \frac{1}{2}\left\{(\alpha|0\rangle + \beta|1\rangle) \otimes |00\rangle + (\alpha|0\rangle - \beta|1\rangle) \otimes |01\rangle + (\alpha|1\rangle + \beta|0\rangle) \otimes |10\rangle + (\alpha|1\rangle - \beta|0\rangle) \otimes |11\rangle\right\} \end{aligned}$$

  where we have grouped the a-states by each possibe outcome for the bc-qubits.

- Now Bob measures the bc-qubits (in Figure 7(b), this is indicated by the meter). This yields the bits $b_1c_1 \in \{00, 01, 10, 11\}$, each choice occurring with probability 1/4. The a-qubit simultaneously collapses to the corresponding state, which is (respectively)

$$(\alpha|0\rangle + \beta|1\rangle)/\sqrt{2}, \qquad (\alpha|0\rangle - \beta|1\rangle)/\sqrt{2}, \qquad (\alpha|1\rangle + \beta|0\rangle)/\sqrt{2}, \qquad (\alpha|1\rangle - \beta|0\rangle)/\sqrt{2}$$

- Bob sends the two classical bits b_1c_1 to Alice using radio waves.

- Based on the received b_1c_1, Alice can apply a simple unitary transformation $X = X_{b_1c_1}$ to her a-qubit to recover the state $|x\rangle$. For instance, if $b_1c_1 = 00$, then the a-qubit is already in the required state (so X is the identity matrix in this case).

Note that the requirement of a classical channel guarantees that this protocol has not violated any fundamental laws such as transmitting information faster than the speed of light.

Summary: the deep strangeness of both phenomenon comes from the fact that entanglement acts across space – the two entangled bits in the Bell state are not required to be in the same locality. But it seems that a physical realization of entanglement requires the two bits to be in the same locality, at least initially.

Exercises

Exercise 10.1: Describe the unitary matrices $X_{b_1c_1}$ that Alice must use to recover the state $|x\rangle$ in the quantum teleportation protocol. ♢

Exercise 10.2: Show how to produce the following Bell states: $\frac{1}{2\sqrt{2}}(|000\rangle + |111\rangle)$ starting from $|000\rangle$. What if you start from $|111\rangle$? ♢

End Exercises

§11. Quantum Algorithms

¶34. There is a basic intuition that "quantum computers are inherently more powerful than classical computers". This can be formalized as questions about inclusions among complexity classes. Unfortunately, we are unlikely to resolve such questions directly because it would imply the resolution of some deep conjectures in classical complexity theory. A more modest goal is to demonstrate individual problems that could be solved more efficiently using quantum computers than with classical computers. Even here, most positive results must be qualified in the sense that a quantum computer can solve a problem more efficiently than any *known* classical algorithm. We should also remember that quantum algorithms are inherently probabilistic, and hence we should only compare them to classical randomized algorithms.

But what kinds of problems can exploit the special capabilities of quantum computers? There is an obvious candidate task where quantum computers can do more efficiently than classical computers: the simulation of quantum systems! This was noted by Feynman. But what else? Let us note that the apparent advantage of quantum computers is the ability to simultaneously maintain (exponentially) many possible states. Hence the obvious way to exploit this is to evolve these states simultaneously. On the other hand, if we measure the quantum system, we only get one of these states (possibly with exponentially small probability). So in order for useful computation to be carried out, we need to have these states interfere in some controlled manner, to bias the probabilities towards the desirable states. One view of quantum computation (see Cleve et al [14]) is that it is basically an application of interferometry to multiparticle systems. This is so different from classical computers that completely new computational techniques must be developed before we can exploit its power.

¶35. Deutsch's Problem. Deutsch [15] gave a simple demonstration of the power of quantum computers over classical computers. Assume we have a quantum black box B that computes the 2-qubit transformation

$$B : (x, y) \mapsto B(x, y) = (x, y \oplus f(x))$$

for some unknown function $f : \mathbb{B}^1 \to \mathbb{B}^1$. There are only 4 possibilities for f: constant 0 function ($f(0) = f(1) = 0$), constant 1 function ($f(0) = f(1) = 1$), identity function ($f(x) = x$), or negation function ($f(x) = 1 - x$). Deutsch's problem is to determine whether $f(0) = f(1)$ (f is constant) or $f(0) \neq f(1)$ (f is balanced). If this were a classical tranformation, we would need to access the black box twice, with $x = 0$ and $x = 1$. But being quantum, we now show a one-access solution.

Consider the Hadamard matrix H in (11). If we first prepare our input to be $|0\rangle \otimes |1\rangle$, then subject it to the transformation $H_2 = H \otimes H$ (see (12)), we obtain

$$|x\rangle \otimes |y\rangle := \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle).$$

If we pass this through our black box, the overall state becomes

$$\begin{aligned} B(|x\rangle \otimes |y\rangle) &= |x\rangle \otimes |y \oplus f(x)\rangle \\ &= |x\rangle \otimes \frac{1}{\sqrt{2}}(|0 \oplus f(x)\rangle - |1 \oplus f(x)\rangle) \\ &= |x\rangle \otimes \frac{1}{\sqrt{2}}(-1)^{f(x)}(|0\rangle - |1\rangle) \\ &= (-1)^{f(x)}|x\rangle \otimes \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \\ &= (-1)^{f(x)}|x\rangle \otimes |y\rangle \end{aligned}$$

The first qubit is quite interesting:

$$\begin{aligned} (-1)^{f(x)}|x\rangle &= \frac{1}{\sqrt{2}}(-1)^{f(x)}(|0\rangle + |1\rangle) \\ &= \frac{1}{\sqrt{2}}((-1)^{f(0)}|0\rangle + (-1)^{f(1)}|1\rangle) \\ &= (-1)^{f(0)}\frac{1}{\sqrt{2}}(|0\rangle + (-1)^{f(1)+f(0)}|1\rangle) \\ &= \begin{cases} (-1)^{f(0)}\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) & \text{if } f(x) \text{ is constant} \\ (-1)^{f(0)}\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) & \text{if } f(x) \text{ is balanced} \end{cases} \\ &= \begin{cases} (-1)^{f(0)}|x\rangle & \text{if } f(x) \text{ is constant} \\ (-1)^{f(0)}|y\rangle & \text{if } f(x) \text{ is balanced} \end{cases} \end{aligned}$$

Thus, if we measure this qubit using the basis $(|x\rangle, |y\rangle)$, we obtain $|x\rangle$ with probability 1 if $f(x)$ is constant, and $|y\rangle$ with probability 1 if $f(x)$ is balanced. Note that the $(-1)^{f(0)}$ is just phase information which does not affect the probability. Alternatively, we can apply H to the first qubit again (recall that $H^{-1} = H$), the possible states are $|0\rangle$ or $|1\rangle$ depending on the nature of f. Now we measure in the original basis. The quantum circuit to perform this computation is shown in figure 8(a). We have just described the improved solution of Cleve et al [14].

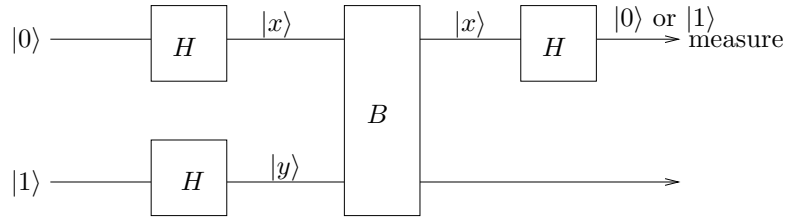

(a) Simple

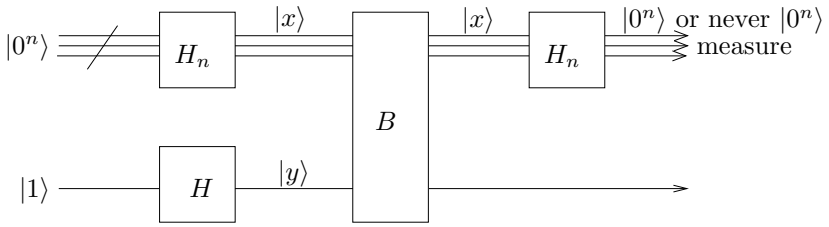

(b) General case

Figure 8: Deutsch's Problem: (a) simple case, (b) general case

¶36. Generalization. Deutsch and Jozsa [17] extended the above example to allow arbitrarily large input sizes. This Deutsch-Jozsa problem is as follows. Suppose $f : \mathbb{B}^n \to \mathbb{B}$ is an unknown function, but it is either a constant function ($f(x)$ is always 0 or always 1) or it is balanced (the number of times $f(x) = 0$ is to 2^{n-1}). We need to decide whether f is constant or balanced. We are given a quantum blackbox B_n which takes a $(n+1)$-quword as input and which applies the function f as follows:

$$B_n(|x_1, \ldots, x_n y\rangle) = |x_1, \ldots, x_n\rangle \otimes |y \oplus f(x_1, \ldots, x_n)\rangle.$$

Using a deterministic algorithm, it seems that we need to make at least $2^{n-1} + 1$ evaluations of B_n, in the worst case. A randomized procedure can do better (Exercise). But we now show that a single call to the black-box is sufficient using a quantum transformation.

The method is a straightforward generalization of the original solution. We prepare the input to be $|0^n 1\rangle$ and apply $H_{n+1} = H \otimes \cdots \otimes H$ to this input. This produces the state

$$\left(\frac{1}{2^{n/2}} \sum_{x=0}^{2^n-1} |x\rangle \right) \otimes \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle). \tag{22}$$

Then we pass this state through B_n. Suppose $y = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ and $x = x_1, \ldots, x_n$ is a eigenstate of the first n bits. Then

$$B_n(|x_1, \ldots, x_n y\rangle) = |x\rangle \otimes \frac{1}{\sqrt{2}}(|f(x) \oplus 0\rangle - |f(x) \oplus 1\rangle) = |x\rangle \otimes \frac{1}{\sqrt{2}}(-1)^{f(x)}(|0\rangle - |1\rangle).$$

Hence the overall state of the quword, after applying B_n to the state (22), is

$$\left(\frac{1}{2^{n/2}} \sum_{x=0}^{2^n-1} (-1)^{f(x)} |x\rangle \right) \otimes |y\rangle.$$

Finally, apply H_n to the first n qubits to get

$$\left(\frac{1}{2^{n/2}}\sum_{x=0}^{2^n-1}\left(\frac{1}{2^{n/2}}\sum_{j=0}^{2^n-1}((-1)^{f(x)}(-1)^{x\cdot j}|j\rangle)\right)\right)\otimes|y\rangle,$$

using Lemma 3 (where $x\cdot j$ denotes scalar product of two n-vectors). Reordering the two summations, we obtain

$$\left(\frac{1}{2^{n}}\sum_{j=0}^{2^n-1}|j\rangle\left(\sum_{x=0}^{2^n-1}((-1)^{f(x)+x\cdot j})\right)\right)\otimes|y\rangle,$$

Note that the state $|j\rangle=|0^n\rangle$ of the first n qubits has amplitude $\frac{1}{2^n}\sum_{x=0}^{2^n-1}(-1)^{f(x)}$. If $f(x)$ is constant, then this amplitude is 1 or -1 (depending on whether $f(x)=0$ or 1). If $f(x)$ is balanced, this amplitude is 0. Therefore, if we measure the first n qubits, we either observe the state $|0^n\rangle$ with probability 1 (if f is constant) or we never observe the state $|0^n\rangle$ (if f is balanced). This concludes our proof.

Exercises

Exercise 11.1: Let $0<\varepsilon<1$. Give an randomized classical algorithm to solve the Deutsch-Jozsa problem with probability of success of $1-\varepsilon$ using $O(\log(1/\varepsilon))$ calls to the blackbox. HINT: Let E_i $(i=0,1)$ be the event that $f(x_1,\ldots,x_n,1)=i$ where $x_1,\ldots,x_n$ are random. What is the probability that E_0 occurs in k trials? ♢

End Exercises

§12. Quantum Discrete Fourier Transform

We introduce the quantum analogue of the well-known Discrete Fourier Transform (DFT). The quantum version of this (QFT) is obtained as a natural adaptation. We derive a quantum circuit for computing QFT. This algorithm will be used later for integer factorization.

¶37. Discrete Fourier Transform. Let $N\in\mathbb{N}$, and $x=(x_0,\ldots,x_{N-1})^T\in\mathbb{C}^N$. The classic **discrete Fourier transform** of x is $DFT(x)=Fx$ (a matrix vector product), where F is the following $N\times N$ matrix

$$F=F_N=\frac{1}{\sqrt{N}}\begin{bmatrix}1 & 1 & 1 & \cdots & 1\\ 1 & \omega & \omega^2 & \cdots & \omega^{N-1}\\ 1 & \omega^2 & \omega^4 & \cdots & \omega^{2(N-1)}\\ & \vdots & & & \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \cdots & \omega^{(N-1)^2}\end{bmatrix} \tag{23}$$

and $\omega=\omega_N:=e^{\mathbf{i}2\pi/N}$. Note that F is symmetric.

Let e_j denote the jth elementary vector e_j (with the jth component equal to 1 and everywhere else 0). Then

$$DFT(e_j)=\frac{1}{\sqrt{N}}\sum_{k=0}^{N-1}\omega^{jk}e_k. \tag{24}$$

We now prove that $F=F_N$ is unitary. Two preliminary remarks are useful: In general, the conjugate of $e^{\mathbf{i}\theta}$ is $\overline{e^{\mathbf{i}\theta}}=e^{-\mathbf{i}\theta}$, and thus $\overline{\omega}=e^{-\mathbf{i}2\pi/N}$. Second, note that ω is an Nth **root of unity** meaning that $\omega^N=1$. But more is true: ω is actually a **primitive root of unity**, meaning that if $\omega^m=1$ then m is a multiple of N.

Lemma 9. *F is unitary: $F^*F=I$.*

Proof. We must show that the (i,j)th entry of F^*F is 1 iff $i = j$:

$$\begin{aligned}(F^*F)_{ij} &= \frac{1}{N}\sum_{k=0}^{N-1} \overline{\omega^{ki}}\omega^{kj} \\ &= \frac{1}{N}\sum_{k=0}^{N-1} \omega^{-ki}\omega^{kj} \\ &= \frac{1}{N}\sum_{k=0}^{N-1} \omega^{k(j-i)} \\ &= \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases}\end{aligned}$$

The last equation is clearly true when $i = j$, since the sum is N in this case. When $i \neq j$, the sum $\sum_{k=0}^{N-1} \omega^{k(j-i)}$ vanishes because of the simple identity

$$(1-x)\sum_{k=0}^{N-1} x^k = 1 - x^N. \tag{25}$$

Plugging $x = \omega^{j-i}$, we conclude that the righthand side is 0 (since ω is a N-th root of unity). But $1-x = 1-\omega^{j-i} \neq 0$ (since ω is a primitive N-th root of unity). This implies that the sum $\sum_{k=0}^{N-1} x^k$ must vanish. **Q.E.D.**

This lemma also shows that we may define **inverse discrete Fourier transform** $IDFT(x)$ by the equation $IDFT(x) = F^*x$.

¶38. Omega notation. We now choose $N = 2^n$. Previously, we wrote $\omega = e^{\mathbf{i}2\pi/N}$. Now we shift to some more permanent notations, fixed for the rest of the chapter.

Let $\omega_\ell := e^{\mathbf{i}2\pi/2^\ell}$ for all $\ell \in \mathbb{N}$. For instance, $\omega_0 = e^{\mathbf{i}2\pi} = 1$ and $\omega_1 = e^{\mathbf{i}\pi} = -1$. Note that we often write "ω_ℓ^{jk}" for $j, k, \ell \in \mathbb{N}$. This should be parsed as $(\omega_\ell)^{jk} = e^{\mathbf{i}2\pi jk/2^\ell}$. We use another convention involving "ω" without any subscripts: for any complex number ϕ, we write "ω^ϕ" as[12] a short hand for $e^{\mathbf{i}2\pi\phi}$. When ϕ is real, then note that $\omega^\phi = \omega^{(\phi \,\mathbf{mod}\, 1)}$ where $\phi \,\mathbf{mod}\, 1$ is the fractional part of ϕ, $0 \leq (\phi \,\mathbf{mod}\, 1) < 1$. When E is a complex expression, we may write $\exp(E)$ instead of e^E.

¶39. Quantum Fourier Transform. The quantum analogue of DFT is just the above DFT, couched in the special setting of quantum circuits. Indeed, since $DFT(x)$ is a unitary transformation, we should be able to compute it with a quantum circuit. The circuit turns out to be fairly simple, but to verify its correctness, we need some preparation.

To define the **quantum Fourier transform (QFT)**, it is enough to define its action on an eigenstate $|x\rangle = |j\rangle$ where $j = 0, \ldots, N-1$:

$$QFT(|j\rangle) = \frac{1}{\sqrt{N}}\sum_{k=0}^{N-1} \omega_n^{jk}|k\rangle. \tag{26}$$

[12] This notation must be properly understood as a textual substitution: "ω^ϕ" should become expanded into "$e^{\mathbf{i}2\pi\phi}$" before the usual interpretation of these mathematical symbols take place. For instance, it is not the same as "$(\omega)^\phi$", since $\omega = e^{\mathbf{i}2\pi}$ evaluates to 1 and so $(\omega)^\phi = 1^\phi$, which is not officially defined.

In general, suppose $|x\rangle = \sum_{j=0}^{N-1} x_j |j\rangle$. Then, by linearity of QFT, we obtain

$$\begin{aligned} |y\rangle = QFT(|x\rangle) &= \sum_{j=0}^{N-1} x_j QFT(|j\rangle) \\ &= \sum_{j=0}^{N-1} x_j \left(\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \omega_n^{jk} |k\rangle \right) \\ &= \sum_{k=0}^{N-1} \left(\frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j \omega_n^{jk} \right) |k\rangle \\ &= \sum_{k=0}^{N-1} y_k |k\rangle \end{aligned}$$

where

$$y_k = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} x_j \omega_n^{jk}. \tag{27}$$

The following important lemma expands the right-hand side of (26) as a tensor product.

Lemma 10.

$$QFT(|j_1 \cdots j_n\rangle) = \frac{1}{\sqrt{N}} \bigotimes_{\ell=1}^{n} (|0\rangle + \omega_\ell^j |1\rangle) \tag{28}$$

$$= \frac{1}{\sqrt{N}} (|0\rangle + \omega_1^j |1\rangle) \otimes (|0\rangle + \omega_2^j |1\rangle) \otimes \cdots \otimes (|0\rangle + \omega_n^j |1\rangle) \tag{29}$$

Proof. Write $|j\rangle = |j_1, \ldots, j_n\rangle = |j_1\rangle \otimes \cdots \otimes |j_n\rangle$.

$$\begin{aligned} QFT(|j\rangle) &= \tfrac{1}{\sqrt{N}} \textstyle\sum_{k=0}^{N-1} \omega_n^{jk} |k\rangle, && \text{(from (26))} \\ &= \tfrac{1}{\sqrt{N}} \textstyle\sum_{k_1=0}^{1} \sum_{k_2=0}^{1} \cdots \sum_{k_n=0}^{1} \omega_n^{j(\sum_{\ell=1}^{n} k_\ell 2^{n-\ell})} |k_1 \cdots k_n\rangle, && \text{(rewriting the sum over } k\text{)} \\ &= \tfrac{1}{\sqrt{N}} \textstyle\sum_{k_1=0}^{1} \sum_{k_2=0}^{1} \cdots \sum_{k_n=0}^{1} \bigotimes_{\ell=1}^{n} \omega_n^{j k_\ell 2^{n-\ell}} |k_\ell\rangle, && \text{(rewriting a tensor)} \\ &= \tfrac{1}{\sqrt{N}} \textstyle\sum_{k_1=0}^{1} \sum_{k_2=0}^{1} \cdots \sum_{k_n=0}^{1} \bigotimes_{\ell=1}^{n} \omega_\ell^{j k_\ell} |k_\ell\rangle, && \text{(by definition of } \omega_\ell\text{)} \\ &= \tfrac{1}{\sqrt{N}} (|0\rangle + \omega_1^j |1\rangle) \otimes (|0\rangle + \omega_2^j |1\rangle) \otimes \cdots \otimes (|0\rangle + \omega_n^j |1\rangle), && \text{as desired.} \end{aligned}$$

Q.E.D.

¶40. The notation in (28), although essentially equivalent to (29), can be ambiguous unless we understand the convention that a tensor product of the form $\otimes_{\ell=1}^{n}$ has to be taken in order of increasing ℓ. This remark may become clearer when we describe the quantum circuit to compute QFT.

¶41. The coefficient ω_ℓ^j in this lemma needs to be decoded:

$$\begin{aligned} \omega_\ell^j &= e^{\mathbf{i} 2\pi 2^{-\ell} (\sum_{i=1}^{n} j_i 2^{n-i})} \\ &= e^{\mathbf{i} 2\pi (\sum_{i=1}^{\ell} j_{n+i-\ell} 2^{-i})}, \qquad \text{(keeping only the fractional part of the exponent of } e^{\mathbf{i}2\pi}\text{.} \\ &= e^{\mathbf{i} 2\pi (0.j_{n+1-\ell} j_{n+2-\ell} \cdots j_n)} \end{aligned}$$

where the $(0.j_{n+1-\ell} j_{n+2-\ell} \cdots j_n)$ is a binary rational. In other words, the notation "ω_ℓ^j" is meant to suggest that this coefficient only depends on the last ℓ bits of $j_1, j_2, \ldots, j_n$. In the following, we shall write

$$\omega^{(0.j_1 j_2 \cdots j_\ell)} \tag{30}$$

for $e^{\mathbf{i} 2\pi (0.j_1 j_2 \cdots j_\ell)}$. For example, we have

$$\omega_1^j = e^{\mathbf{i} 2\pi (0.j_n)} = \begin{cases} 1 & \text{if } j_n = 0 \\ e^{\mathbf{i}\pi} & \text{if } j_n = 1 \end{cases} = \omega_1^{j_n}.$$

¶42. Inverse QFT. We also need to compute the inverse QFT (or IQFT). This simply amounts to using the complex conjugate of ω_n in place of ω_n: Note that the complex conjugate is $\overline{\omega_n} = \overline{e^{\mathbf{i}2\pi 2^{-n}}} = e^{-\mathbf{i}2\pi 2^{-n}}$. We may write $\overline{\omega}_n$ for this complex conjugate.

Lemma 11. *For all states $|x\rangle$ $IQFT(QFT(|x\rangle)) = QFT(IQFT(|x\rangle))$.*

Note that we only need to verify this lemma for the case where $|x\rangle$ is an eigenstate. We leave this verification as an exercise.

¶43. One Stage of QFT Circuit. The QFT circuit will consist of n stages. It is sufficient to understand a single stage of this process. In fact the first stage is representative of all the other stages, and is illustrated in Figure 9.

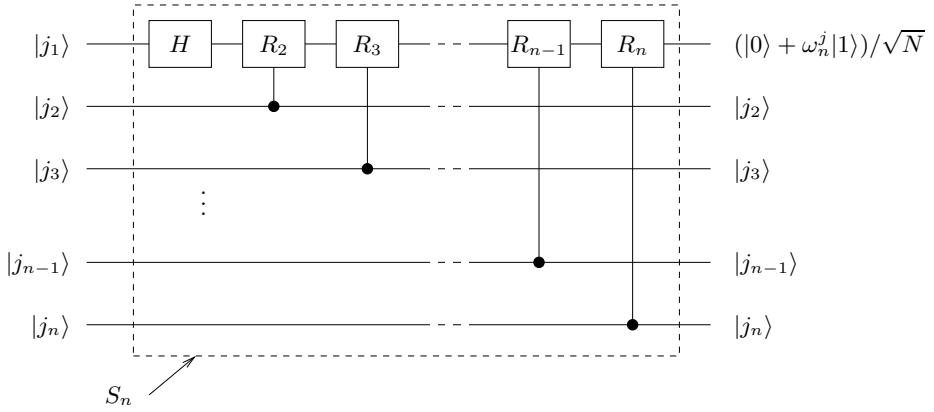

Figure 9: S_n: Stage One of the QFT circuit

The first gate is represented by the standard H-matrix. This H-gate transforms a qubit eigenstate $|j_1\rangle$ as follows: $H(|0\rangle) = (|0\rangle + |1\rangle)/\sqrt{2}$ if $j_1 = 0$ and $H(|1\rangle) = (|0\rangle - |1\rangle)/\sqrt{2}$ if $j_1 = 1$. We summarize this by writing

$$H(|j_1\rangle) = \frac{1}{\sqrt{2}}(|0\rangle + \omega_1^{j_1}|1\rangle). \tag{31}$$

Viewing the first H-gate as a transformation of all the qubits of $|j\rangle$, we obtain

$$H(|j\rangle) = \frac{1}{\sqrt{2}}(|0\rangle + \omega_1^{j_1}|1\rangle) \otimes |j_2 \cdots j_n\rangle.$$

If $n = 1$, we are done. If $n \geq 2$, we introduce $n-1$ transformations of the following type. For $\ell \geq 2$, the transformation R_ℓ is a 2-qubit gate that achieves a "controlled rotation", and is given by the matrix

$$R_\ell := \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & \omega_\ell \end{bmatrix} \tag{32}$$

In Figure 9, R_2 is applied to the first and second qubits of $|j_1 \cdots j_n\rangle$. Note that in the figure, we write R_2 as a box is applied to line 1 (prepared in state $|j_1\rangle$), using line 2 (prepared as $|j_2\rangle$) as the control line. But from the matrix (32), we see that the roles of the two input lines are completely symmetric (so we could view line 1 as the control line if we wish). Thus

$$\begin{aligned} R_2(|00\rangle) &= |00\rangle \\ R_2(|01\rangle) &= |01\rangle \\ R_2(|10\rangle) &= |10\rangle \\ R_2(|11\rangle) &= \omega_2|11\rangle \end{aligned}$$

The output of H on lines 1 and 2 is $\frac{1}{\sqrt{2}}(|0\rangle + \omega_1^{j_1}|1\rangle) \otimes |j_2\rangle$. Considering the case $j_2 = 0$ and $j_2 = 1$ separately, R_2 yields:

$$\begin{aligned} R_2(\frac{1}{\sqrt{2}}(|0\rangle + \omega_1|1\rangle) \otimes |0\rangle) &= \frac{1}{\sqrt{2}}(|0\rangle + \omega_1|1\rangle) \otimes |0\rangle, \\ &= \frac{1}{\sqrt{2}}(|0\rangle + \omega^{(0.10)}|1\rangle) \otimes |0\rangle, \\ R_2(\frac{1}{\sqrt{2}}(|0\rangle + \omega_1|1\rangle) \otimes |1\rangle) &= \frac{1}{\sqrt{2}}(|0\rangle + \omega_2\omega_1|1\rangle) \otimes |1\rangle, \\ &= \frac{1}{\sqrt{2}}(|0\rangle + \omega^{(0.11)}|1\rangle) \otimes |1\rangle, \end{aligned}$$

This proves that

$$\begin{aligned} R_2(H(|j_1 j_2\rangle)) &= \frac{1}{\sqrt{2}}(|0\rangle + \omega^{(0.j_1 j_2)}|1\rangle) \otimes |j_2\rangle) \\ &= \frac{1}{\sqrt{2}}(|0\rangle + \omega_2^{j_1 j_2}|1\rangle) \otimes |j_2\rangle). \end{aligned}$$

Continuing the same way, each R_ℓ-gate simply transforms the factor $(|0\rangle + \omega^{0.j_1 j_2 \cdots j_{\ell-1}}|1\rangle)$ into $(|0\rangle + \omega^{0.j_1 j_2 \cdots j_{\ell-1} j_\ell}|1\rangle)$ (by appending j_ℓ to the the binary fraction in the exponent). The final result is that line 1 has the value,

$$\frac{1}{\sqrt{2}}(|0\rangle + \omega^{(0.j_1 j_2 \cdots j_n)}|1\rangle).$$

By comparing this result to Lemma 10, we see that this result should really appear in line n in the QFT circuit. But we may postpone this transposition until the end. Lines 2 to n have their original values unchanged.

¶44. Putting together the stages. Let S_n be the circuit represented by stage 1. It is now clear that we can continue this process on lines 2 to line n, but applying the circuit S_{n-1} instead. as a result, line 2 will have the value $(|0\rangle + \omega_{n-1}^j|1\rangle)/\sqrt{2}$, Finally, line n is transformed by S_1 which is just a single H-gate, yielding $(|0\rangle + \omega_1^j|1\rangle)/\sqrt{2}$ on this line. The resulting circuit is seen in Figure 10.

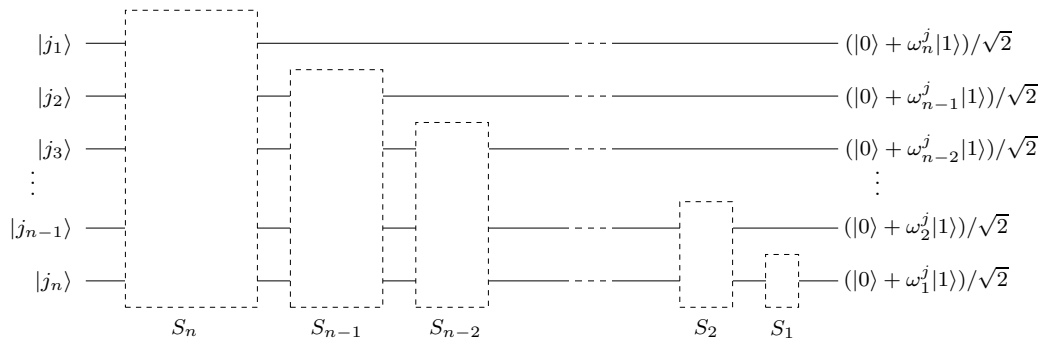

Figure 10: Quantum circuit for Reverse Fourier transform

We are almost done, except the values in the output lines of Figure 10 are in reverse order (compare Lemma 10). But we know that transposing any two bits (quantum or classical) can be achieved by three T_2 gates. We can therefore introduce $\lfloor n/2 \rfloor$ such gates to exchange the outputs of line i and line $n - i + 1$. This completes the description of the QFT circuit.

¶45. Complexity Analysis. Since each stage uses $\Theta(n)$ gates, the overall number of gates is $\Theta(n^2)$.

At this point, we should pause and consider what has been accomplished: we can compute QFT using $O(n^2)$ quantum gates. This is exponentially smaller than any classical construction which surely need $\Omega(N) = \Omega(2^n)$ gates. But exploiting this result is hardly obvious. First, we need to physically prepare an arbitrary quantum state $|x\rangle$. Further, we need to extract the resulting values stored in $|y\rangle = QFT(|x\rangle)$. These issues will come up as we seek to exploit QFT.

Exercises

Exercise 12.1: Explain why QFT is really the quantum analogue of DFT. ♢

Exercise 12.2: Prove Lemma 11. ♢

Exercise 12.3: Give a simple upper bound on the number of classical gates to compute DFT (assuming each gate can perform a single complexity arithmetic operation in constant time). ♢

End Exercises

§13. Phase Estimation Problem

¶46. The phase estimation problem is another basic task in quantum algorithms. Suppose U is a unitary operator U, given as blackbox. This simply means that we can use it as a primitive gate in our quantum circuits. Furthermore, we are given an eigenvector $|v\rangle$ of U. If the associated eigenvalue is $\omega^\phi = e^{\mathbf{i}2\pi\phi}$, then

$$U|v\rangle = \omega^\phi|v\rangle.$$

Call ϕ the **phase angle**, and we may assume $0 \leq \phi < 1$. The problem of **Phase Estimation** is, given U and $|v\rangle$ and also $n \in \mathbb{N}$ and $\varepsilon > 0$, to compute an approximation of ϕ to n bits of precision, with probability of failure $\leq \varepsilon$. The approximation $\widetilde{\phi}$ has n bits of precision if $|\widetilde{\phi} - \phi| \leq 2^{-n}$. However note that this does not mean that $\widetilde{\phi}$ is represented by a binary rational number with n-bits after the mantissa; below, we will use about $\log(1 + \varepsilon^{-1})$ more bits.

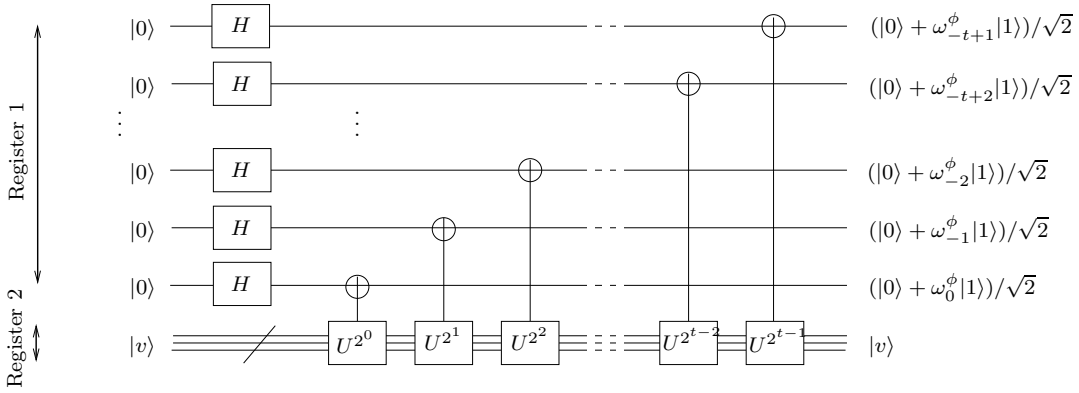

Figure 11: Stage 1 of Phase Estimation

¶47. Stage One of phase estimation. There are two stages in phase estimation. The first stage is illustrated in Figure 11. Note that there are two registers (quwords) in the circuit:
(1) Register 1 comprises t qubits that are used to readout the final estimate of ϕ. At the end, we will measure this register 1, and this value, interpreted as a binary rational $0.j_1 j_2 \cdots j_t$ is regarded as the estimate of ϕ. The value of t will be determined later, as it depends on the input parameters n and ε.
(2) Register 2 is used to carry the eigenvector $|v\rangle$.

Note that U^k is just k-fold application of U. Thus $U^k(|v\rangle) = U(U(\cdots U(|v\rangle)\cdots)) = \omega^{k\phi}|v\rangle$. In the following, we actually use the control-U gate, which has an extra control line to turn turn on or off the actual actions of the U-gate. However, we continue to use "U" to denote the control-U gate. Consider the result of applying control-U^{2^ℓ} ($\ell \geq 0$) to the output of H:

$$\begin{aligned}
U^{2^\ell}(|0\rangle \otimes |v\rangle) &= |0\rangle \otimes |v\rangle \\
U^{2^\ell}(|1\rangle \otimes |v\rangle) &= \omega^{2^\ell\phi}|1\rangle \otimes |v\rangle \\
U^{2^\ell}((|0\rangle + |1\rangle) \otimes |v\rangle) &= (|0\rangle + \omega^{2^\ell\phi}|1\rangle) \otimes |v\rangle \\
&= (|0\rangle + \omega^\phi_{-\ell}|1\rangle) \otimes |v\rangle
\end{aligned}$$

This justifies the output on Register 1 as specified in the lines of Figure 11.

¶48. Phase Two. We motivate the construction of the second phase. Suppose

$$\phi = 0.\phi_1\phi_2\cdots\phi_t, \qquad \phi_i \in \{0,1\}.$$

Then we observe the output lines of Register 1 in Figure 11 is simply the outputs specified by Lemma 10. For instance, line one's output in Figure 11 is $(|0\rangle + \omega^{\phi}_{-t+1}|1\rangle)/\sqrt{2}$, which is equal to $(|0\rangle + \omega^{0.\phi_t}|1\rangle)/\sqrt{2}$. Line one of Lemma 10 has output $(|0\rangle + \omega^{j}_{1}|1\rangle)/\sqrt{2}$, which is equal to $(|0\rangle + \omega^{0.j_n}|1\rangle)/\sqrt{2}$. So the two outputs are the same once we identify n with t and $j_1,\ldots,j_n$ with $\phi_1,\ldots,\phi_t$.

This means, if we apply the inverse QFT circuit of the previous section as our second phase, we would obtain as output the eigenstate $|\phi_1\cdots\phi_t\rangle$. This inverse QFT is indeed our second phase, even when ϕ is not a binary rational. But the proof that it yields a good estimate of ϕ in general is more subtle and will be taken up next.

Finally, the estimate of ϕ is obtained by measuring Register 1. This yields t bits which is interpreted as a binary rational $0.b_1b_2\cdots b_t$ ($b_i \in \{0,1\}$) and taken as the estimate of ϕ.

¶49. Error Analysis. Let $\phi < 1$ be an arbitrary real. Suppose $\widetilde{\phi} = 0.\phi_1\cdots\phi_t$ is a binary rational such that

$$\delta := |\phi - \widetilde{\phi}| \leq 2^{-t}. \tag{33}$$

Remark that if ϕ is not a binary rational, then $\widetilde{\phi}$ is uniquely determined by t. Ideally, we would like our final measurement to yield $|\phi_1\cdots\phi_t\rangle = |2^t\widetilde{\phi}\rangle$ in Register 1. But short of this, our current goal is to obtain eigenstate $|j\rangle = |j_1\cdots j_t\rangle$ such that

$$\Pr\{|\phi - 0.j_1\cdots j_t| < 2^{-n}\} \geq 1-\varepsilon. \tag{34}$$

where n and ε are user specified parameters. We show how to choose $t > n$ so that we obtain the guarantee (34).

The output of Phase 1 is

$$\frac{1}{2^{t/2}}(|0\rangle + \omega^{\phi}_{-t+1}|1\rangle) \otimes (|0\rangle + \omega^{\phi}_{-t+2}|1\rangle) \otimes \cdots \otimes (|0\rangle + \omega^{\phi}_{0}|1\rangle) = \frac{1}{2^{t/2}}\sum_{\ell=0}^{2^t-1}\omega^{\phi\ell}|\ell\rangle$$

The inverse QFT is basically the same as QFT, except that we replace ω_t (taking $N = 2^t$) by its conjugate $\overline{\omega_t}$. When we apply the inverse QFT, we get

$$\begin{aligned} QFT^{-1}\left(\tfrac{1}{2^{t/2}}\textstyle\sum_{\ell=0}^{2^t-1}\omega^{\phi\ell}|\ell\rangle\right) &= \tfrac{1}{2^{t/2}}\textstyle\sum_{\ell=0}^{2^t-1}\omega^{\phi\ell}\tfrac{1}{2^{t/2}}\left(\sum_{k=0}^{2^t-1}\omega_t^{-\ell k}|k\rangle\right) \quad \text{(by (26))} \\ &= \tfrac{1}{2^t}\textstyle\sum_{\ell=0}^{2^t-1}\sum_{k=0}^{2^t-1}(\omega^{\phi-k2^{-t}})^{\ell}|k\rangle. \end{aligned}$$

The amplitude of $|k\rangle$ is therefore

$$\begin{aligned} \alpha_k &:= \frac{1}{2^t}\sum_{\ell=0}^{2^t-1}(\omega^{\phi-k2^{-t}})^{\ell} \\ &= \frac{1}{2^t}\frac{1-(\omega^{\phi-k2^{-t}})^{2^t}}{1-\omega^{\phi-k2^{-t}}} \qquad \text{(sum of geometric series)}. \end{aligned}$$

We use the following simple upper and lower bound on $|1-e^{\mathbf{i}\theta}|$, where $|\theta| \leq \pi$:

$$2 \geq \left|1-e^{\mathbf{i}\theta}\right| \geq 2\sin|\theta/2| \geq |\theta|. \tag{35}$$

This bound is illustrated in Figure 12.

Thus $|1-\omega^x| = |1-e^{\mathbf{i}2\pi x}| \geq 2\pi|x|$ when $|x| \leq 1/2$. Hence

$$|\alpha_k| \leq \frac{1}{2^t}\frac{2}{2\pi\,|\phi-k2^{-t}|} = \frac{1}{2^t\pi\,|\phi-k2^{-t}|}.$$

Writing β_k for $\alpha_{k+2^t\widetilde{\phi}}$, we have

$$|\beta_k| \leq \frac{1}{2^t\pi\left|\phi-\widetilde{\phi}-k2^{-t}\right|} \leq \frac{1}{2^t\pi\,|\delta-k2^{-t}|}. \tag{36}$$

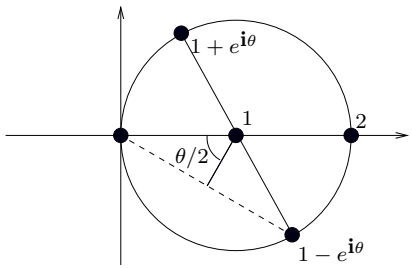

Figure 12: Bounds on $\left|1 - e^{\mathbf{i}\theta}\right|$

At the end of Phase 2, we measure Register 1 and suppose we obtain the state $|j\rangle = |j_1 \cdots j_t\rangle$. We declare $j2^{-t}$ as our estimate for ϕ. Since j is a random variable, we now want to upper bound the probability that $|\phi - j2^{-t}| > 2^{-n}$ (this is the error probability because we want an n-bit approximation to ϕ). We will let $\Delta = 2^{t-n} - 1$ in the following derivation. We now bound the probability $\Pr\{|\phi - j2^{-t}| > 2^{-n}\}$, which is equal to:

$$\begin{array}{rcll}
\Pr\{|2^t\phi - j| > 2^{t-n}\} & \leq & \Pr\{|2^t\widetilde{\phi} - j| > 2^{t-n} - 1\} & (\text{since } |2^t\phi - 2^t\widetilde{\phi}| \leq 1) \\
& = & \displaystyle\sum_{k:|k-2^t\widetilde{\phi}|>\Delta} |\alpha_k|^2 & (\Delta = 2^{t-n} - 1) \\
& = & \displaystyle\sum_{k':|k'|>\Delta} |\beta_{k'}|^2 & (\text{put } k' = k - 2^t\widetilde{\phi}) \\
& \leq & \displaystyle\sum_{k:|k|>\Delta} \frac{1}{2^{2t}\pi^2|\delta - k2^{-t}|^2} & (\text{by (36)}) \\
& = & \frac{1}{\pi^2} \displaystyle\sum_{k:|k|\geq\Delta} \frac{1}{|2^t\delta - k|^2} & \\
& \leq & \frac{1}{\pi^2} \displaystyle\sum_{k:k>\Delta} \frac{1}{(k-1)^2}, & (2^t\delta \leq 1) \\
& < & \frac{1}{\pi^2} \displaystyle\sum_{k:k>\Delta} \frac{1}{(k-1)(k-2)} & \\
& = & \frac{1}{\pi^2} \displaystyle\sum_{k:k>\Delta} \left(\frac{1}{(k-1)} - \frac{1}{(k-2)}\right) & \\
& < & \frac{1}{\pi^2\Delta} & (\text{by telescoping}).
\end{array}$$

We are almost there: recall that we want the probability of error to be at most ε. Hence it is enough that $\varepsilon \geq 1/(\pi^2\Delta)$. It suffices if $\varepsilon \geq 1/(2^{t-n})$. Hence we may choose t to be

$$t = n - \lfloor \lg(\varepsilon) \rfloor . \tag{37}$$

We are done with phase estimation.

¶50. In summary: given a blackbox U and an eigenvector $|v\rangle$, and a precision bound of n and error bound of $\varepsilon > 0$, if we construct the above 2-phase quantum circuit using the parameter t in Equation (37), then the measured value $0.j_1 \cdots j_n$ in Register 1, is an n-bit approximation of ϕ with probability $1 - \varepsilon$. Here, ϕ is given by $U|v\rangle = \omega^\phi|v\rangle$.

REMARKS:

1. There is an important caveat in actual applications: notice that our circuit requires the control-U^{2^i} circuts. If we only assume the availability of the control-U blackboxes, then the only way to construct the control-U^{2^i} circuits is to cascade a sequence of 2^i control-U gates. This is not efficient enough in our applications below, because it requires exponentially (as a function of i) many gates. It turns out that, in fact, we will be able to construct the control-U^{2^i} circuits using a polynomial number of gates.
2. We can extend this technique to the case where an eigenvector $|v\rangle$ is not directly available, but we have a fixed state $|x\rangle = \sum_i c_i|v_i\rangle$ that is a superposition of eigenvectors $|v_i\rangle$, each with a phase ϕ_i. When we measure Register 1, for each i, we obtain an estimate of ϕ_i with probability $|c_i|^2$ (and on measuring Register 2, we find out the $|v_i\rangle$ whose phase we measured).
3. The Exercise develops another approach to phase estimation, based on estimating the bias of a coin tossing.

Exercises

Exercise 13.1: We want to estimate the phase ϕ_1 of an eigenvector $|v_1\rangle$ of a blackbox unitary operator U. We outlined a method to estimate ϕ_1 when we can only prepare a state $|x\rangle$ that contains $|v_1\rangle$ as one of its components. Compute the probability of correct measurements of ϕ_1, and discuss its impact on complexity. ♢

Exercise 13.2: Let U be a unitary transformation such that $U|v\rangle = \omega^\phi|v\rangle$, and we want to estimate $0 \leq \phi < 1$, in the usual phase estimation sense. Consider the circuit in Figure 13.

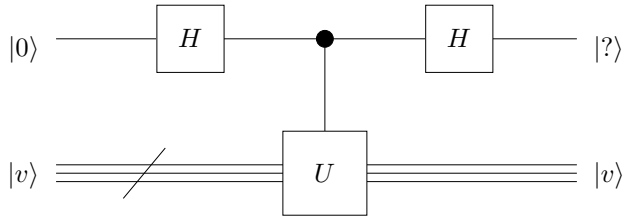

Figure 13: Alternative primitive for phase estimation

(i) Show that the probability of measuring a 0 on the control line is $p_0 = (1+\cos 2\pi\phi)/2$; and, the probability of measuring a 1 is $p_1 = (1-\cos 2\pi\phi)$.
(ii) Suppose you repeat this measurement k times, and get the outcome $|i\rangle$ in k_i $(i = 0, 1)$ of the time. Let $\widetilde{p}_i = k_i/(k_0+k_1) = k_i/k$. What is the probability that $|p_0 - \widetilde{p}_0| > 1/4$?
(iii) Show how to generalize this to estimate p_0 to higher order accuracy. HINT: think of (ii) as measuring the first bit of p_0. ♢

End Exercises

§14. Integer Factoring Problem

Our main goal is integer factorization. So far, we have developed some tools from quantum algorithms: QFT and Phase Estimation. We now address the number theoretic background for integer factorization.

The factorization problem is easy to state: given an integer $N > 2$, to find distinct primes p_i and exponents $e_i \geq 1$ $(i = 1, \ldots, m)$ such that

$$N = \prod_{i=1}^{m} p_i e^i, \tag{38}$$

as guaranteed by the Fundamental Theorem of Arithmetic (see below). This problem can easily be solved in time polynomial in N (see Exercise), but this is not considered a polynomial-time algorithm because the size of the input is only $n = \lfloor \lg(1+N) \rfloor$ (tbe number of bits in the binary representation of N). The current record for integer factorization takes time $\Theta((N \log^2 N)^{1/3})$ using sophisticated number field sieve methods. Our goal is to describe a quantum factoring algorithm that takes time $\Theta(\log^4 N) = \Theta(n^4)$.

¶51. Some Number Theory. When we say "numbers" in Number Theory, it means a natural number $n \in \mathbb{N}$. The starting point of Number Theory is the **divisibility relation** on integers: we say m **divides** n, and write $m|n$, if $ma = n$ for some $a \in \mathbb{Z}$. Also, let $m \not| n$ denote the negation of divisibility. We say $n \in \mathbb{Z}$ is **prime** if it is divisible by exactly two numbers. This definition excludes the numbers 0 (being divisible by all numbers) and 1 (being divisible by one number). Hence 2 is the smallest prime number, and it also has the distinction of being the only even prime. The Fundamental Theorem of Arithmetic says that every number $n \geq 2$ has a unique representation of the form

$$n = \prod_{i=1}^{m} p_i^{e_i}, \qquad m \geq 1 \tag{39}$$

where $p_1 < p_2 < \cdots p_m$ are distinct primes, and $e_i \geq 1$. We write "$x \equiv y \pmod n$" if $n|(x-y)$. We assume that students are familiar with the concept of **modulo** n **arithmetic**: we may add and multiply modulo n. E.g., $5 + 6 = 2 \pmod 9$. Thus, the set $\mathbb{Z}_n = \{0, 1, \ldots, n-1\}$ is a ring under modulo n arithmetic. We shall write "$x \,\mathbf{mod}\, n$" for the unique value in $\mathbb{Z}_n$ that is equivalent to x modulo n. Thus **mod** is a notation for an infix operator. E.g., $11 \,\mathbf{mod}\, 9 = 2$.

The greatest common divisor (GCD) of m, n is denoted $\texttt{GCD}(m, n)$, and is the largest number that divides both m and n. When $m = n = 0$, we define $\texttt{GCD}(m, n) = 0$; otherwise, it is clear that $\texttt{GCD}(m, n) \geq 1$. If $\texttt{GCD}(m, n) = 1$, we say m, n are **coprime** (or, m is coprime to n). We can compute $\texttt{GCD}(m, n)$ in polynomial time, for instance, using Euclid's algorithm. Euclid's algorithm is simple to describe: if $n_0 > n_1 \geq 1$ are numbers, then we compute $\texttt{GCD}(n_0, n_1)$ be generating the following "Euclidean sequence"

$$(n_0, n_1, n_2, \ldots, n_{k-1}, n_k), \qquad k =\geq 1, \tag{40}$$

where

$$n_{i+1} = n_{i-1} \bmod n_i, \qquad i = 1, \ldots, k-1. \tag{41}$$

The termination condition for the sequence is given by

$$n_{k-1} \bmod n_k = 0. \tag{42}$$

It is easily seen from (41) that the following holds:

$$\texttt{GCD}(n_i, n_{i+1}) = \texttt{GCD}(n_{i-1}, n_i).$$

But $n_k = \texttt{GCD}(n_{k-1}, n_k)$, by the termination condition (42) Hence n_k is equal to $\texttt{GCD}(n_0, n_1)$. This proves the correctness of Euclid's algorithm. E.g., $\texttt{GCD}(22, 15) = 1$ follows from the Euclidean sequence

$$22, 15, 7, 1.$$

Again, $\texttt{GCD}(22, 18) = 2$ because $(22, 14, 8, 6, 2)$ is an Euclidean sequence. We can extract a valuable piece of information from the Euclidean algorithm: it is easy to verify by induction that in the Euclidean sequence (40), there exists integers s_i, t_i such that

$$n_{i+1} = s_i n_0 + t_i n_1, \qquad i = 1, \ldots, k-1.$$

In particular, there exists $s, t \in \mathbb{Z}$ such that $\texttt{GCD}(m, n) = sm + tn$. Thus $sm \equiv \texttt{GCD}(m, n) \pmod n$. When m, n are coprime, we conclude that $sm \equiv 1 \pmod n$ and $tn \equiv 1 \pmod m$. Thus s is the (multiplicative) **inverse** of m modulo n, and similarly t is the inverse of n modulo n. We sometimes write $s = m^{-1} \pmod n$. Summarizing: *every element m that is coprime to n has an inverse modulo n.* It is easy to modify the Euclidean algorithm to compute s_k, t_k as well (Exercise). This is usually called the Extended Euclidean algorithm. Thus we can compute multiplicative inverses.

For $n \in \mathbb{N}$, let

$$\mathbb{Z}_n^* = \{i \in \mathbb{Z}_n : \texttt{GCD}(i, n) = 1\}.$$

Since every element in $\mathbb{Z}_n^*$ has an inverse modulo n, and 1 is clearly the identity of multiplication modulo n, we conclude: *$\mathbb{Z}_n^*$ is a group under multiplication modulo n.* The size of this group is $|\mathbb{Z}_n^*| = \phi(n)$ where $\phi(n)$ is Euler's totient function. Thus $\phi(n)$ is the number of distinct values in $\mathbb{Z}_n$ that are coprime to n. For instance, if p is prime, $\phi(p) = p - 1$ since $\mathbb{Z}_p^* = \{1, 2, \ldots, p-1\}$. We can easily generalize this to $\phi(p^e) = p^{e-1}(p-1)$ because the elements in $\mathbb{Z}_{p^e} \setminus \mathbb{Z}_{p^e}^*$ are precisely the multiples of p and there are p^{e-1} of these. [In proof: each $i \in \mathbb{Z}_{p^{e-1}}$ gives rise to a unique multiple $pi \in \mathbb{Z}_{p^e}$.] Thus $\phi(p^e) = p^e - p^{e-1} = p^{e-1}(p-1)$. Further, if m, n are coprime, we can show (Exercise) that $\{xx' : x \in \mathbb{Z}_n^*, x' \in \mathbb{Z}_m^*\} = \mathbb{Z}_{mn}^*$. Hence $\phi(mn) = \phi(m)\phi(n)$. This gives us a formula for $\phi(n)$ when we know the complete factorization of n, similar to (39). For instance, $\phi(24) = \phi(2^3)\phi(3) = 8$. The Fermat-Euler theorem says that for all $x \in \mathbb{Z}_n^*$, $x^{\phi(n)} = 1 \pmod n$. A **power** is a number of the form m^e for some $m, e \in \mathbb{N}$. When m is prime, then m^e is called a **prime power**.

¶52. Reductions of the Factoring Problem. Henceforth, we simply say "factorization" for "integer factorization". We reduce the factorization problem to its computational "core". Here, we say a problem P is reducible to another problem Q if we can construct a polynomial time algorithm A_P for P from any polynomial time algorithm A_Q for Q. Moreover, this reduction is randomized if A_P is randomized (regardless of whether A_Q is randomized or not). In general we need randomized algorithms.

1. **Reduction to Simple Factoring:** First we may reduce the problem to finding any non-trivial factor M of N. That is, either declare N a prime or find an M such that $M|N$ and $1 < M < N$. Call this the **simple factoring problem**, as opposed to the original version which we call the **complete factoring problem**, represented by (39). The complete factoring problem is reduced to at most $\lg N$ simple factoring problem. This is because in (39), we have $N = \prod_{i=1}^m p_i^{e_i} \geq \prod_{i=1}^m 2^{e_i} \geq 2^e$ where $e = \sum_{i=1}^m e_i$. Hence $e \leq \lg N$.

2. **Reduction to an Odd Non-power:** Given an N, we want to detect if it is a power, and if so, completely factor it, $N = M^e$. This can easily be done in polynomial time (Exercise).

   We can further assume N is odd. This is trivial, but the reader will rightfully wonder if this is a just an adhoc decision: why only exclude multiples of 2? We could likewise exclude any multiples of 3 or 5, or any number we like. The reason this decision is not adhoc is related to the next reduction.

3. **Reduction to Finding a Squareroot of Unity modulo N.** Call x a **squareroot** of unity (*i.e.*, 1) modulo N if $x^2 \equiv 1 \pmod N$. Clearly, $x = 1$ and $x = N - 1$ (usually denoted -1) are squareroots of 1 modulo N. But these are the **trivial squareroots** of 1; we want nontrivial squareroots where $1 < x < N - 1$. Armed with such an x, we see that $x^2 - 1 = (x - 1)(x + 1)$ must be divisible by N. Hence, a prime factor p of N must divide either $x - 1$ or $x + 1$. This means p divides either $\texttt{GCD}(N, x - 1)$ or $\texttt{GCD}(N, x + 1)$, *i.e.*, either $\texttt{GCD}(N, x - 1)$ or $\texttt{GCD}(N, x + 1)$ is a nontrivial factor of N.

   In other words, if we could find nontrivial squareroots of unity modulo N, then we can factor N. Part of what we need to establish is that when N is composite, there will be plenty of such squareroots of unity!

4. **Reduction to Order Finding.** We can reduce finding non-trivial squareroots of unity to order finding. For $x \in \mathbb{Z}_n^*$, the **order of x modulo n** (or, the **n-order of x**, or $\texttt{ord}_n(x)$) is the smallest $r \geq 0$ such that $x^r = 1 \pmod n$. The **order finding problem** is[13] to find r, given x, n. By the Euler Fermat theorem, $x^{\phi(n)} \equiv 1 \pmod n$ and hence $r \leq \phi(n)$. Indeed, $r|\phi(n)$ because if $\phi(n) = ar + b$ where $0 \leq b < r$, then $x^b \equiv x^{b+ar} = x^{\phi(n)} \equiv 1 \pmod n$, which is a contradiction unless $b = 0$. The order finding problem is not known to be solvable in polynomial time. The reduction to order finding is nontrivial and will be taken up next.

¶53. Squareroots of unity: cyclic case The above reduction prompts a closer examination of the squareroots of 1. Let x be a squareroot of 1 modulo n. An obvious question is: when is x nontrivial? To answer this question, we must look into the group structure of $\mathbb{Z}_n^*$.

1. The simplest groups are the cyclic ones. By definition, a group G is cyclic if there exists $g \in G$ (called a **generator**) that generate the entire group by repeated multiplication: $G = \{g^i : i \in \mathbb{N}\}$. It is known that $\mathbb{Z}_n^*$ is cyclic if and only if $n = 2, 4$ or $n = p^m$ or $2p^m$ where p is an odd prime and $m \geq 1$. E.g., $\mathbb{Z}_8^* = \{1, 3, 5, 7\}$ is not cyclic since $3^2 \equiv 5^2 \equiv 7^2 \equiv 1 \pmod 8$.

2. *If $\mathbb{Z}_n^*$ is cyclic then the only squareroots of 1 modulo n are the trivial ones.* In proof, let $g \in \mathbb{Z}_n^*$ be a generator of the cyclic group and let $x = g^e$ for some e $(0 < e < \phi(n))$. Since $g^{2e} \equiv x^2 \equiv 1 \pmod n$, we have $\phi(n)|2e$ and so $\phi(n)/2 \leq e$. But $\phi(n)/2 < e$ is not possible because it yields the contradiction that $g^{2e-\phi(n)} \equiv 1$ and $1 \leq 2e - \phi(n) < \phi(n)$. This proves our claim.

   As corollary,if n is an odd prime power, there are no nontrivial squareroots of unity.

3. *The number of generators in a cyclic group G of size n is $\phi(n)$.* This fact is not used below, and may be skipped. In proof, let g be any generator of G. We claim that g^e is also a generator iff $\texttt{GCD}(e, n) = 1$. If $\texttt{GCD}(e, n) = m > 1$ then $(g^e)^{n/m} = 1$ and so g^e cannot generate G. If $\texttt{GCD}(e, n) = 1$ and if r is the order of g^e, then $g^{re} = 1$. Since G is cyclic, this implies $n|re$. As e and n are coprime, we conclude that $n|r$. Since $n \geq r$, we conclude that $r = n$. As there are $\phi(n)$ choices for e, we have shown that G has exactly $\phi(n)$ generators.

   As corollary, the number of generators of $\mathbb{Z}_{p^e}^*$ is $\phi(\phi(p^e)) = \phi(p^{e-1}(p-1))$. If $e \geq 2$, this is $p^{e-2}(p-1)\phi(p-1)$.

Example. Suppose $n = 45 = 9 \times 5$. Then $\phi(n) = \phi(9)\phi(5) = 6 \times 4 = 24$. Now, $\phi^2(9) = \phi^2(5) = 2$ and so $\mathbb{Z}_9^*$ and $\mathbb{Z}_9^*$ each has 2 generators each. We may check that 2 and 5 are the generators of $\mathbb{Z}_9^* = \{1, 2, 4, 5, 7, 8\}$, and 2 and 3 are generators of $\mathbb{Z}_5^* = \{1, 2, 3, 4\}$.

¶54. Group structure of $\mathbb{Z}_n^*$ for odd nonpowers. We now assume that n is an odd nonpower. Then $n = \prod_{i=1}^m q_i$ $(m \geq 2)$ where each q_i is a prime power and the q_i's are coprime.

1. The group $\mathbb{Z}_n^*$ is thus isomorphic to
$$G^* := \mathbb{Z}_{q_1}^* \times \cdots \times \mathbb{Z}_{q_m}^* \tag{43}$$
because of the Chinese Remainder Theorem (Exercise). Indeed the isomorphism is the natural one, $h^* : \mathbb{Z}_n^* \to G^*$ where $h^*(x) = (x_1, \ldots, x_n)$ and $x_i \equiv x \pmod{q_i}$.

[13] Since orders are also known as "periods", it is also known as the "period finding problem".

2. The group operation $\circ$ in G^* is componentwise multiplication, modulo q_i in the ith component. Thus, $h^*(ab) = h^*(a) \circ h^*(b)$. Let us note that

$$h^*(1) = (1, 1, \ldots, 1), \quad h^*(-1) = (-1, -1, \ldots, -1).$$

Note that there are 2^m squareroots of 1 corresponding to elements of the form $(\pm 1, \pm 1, \ldots, \pm 1) \in G^*$.

3. In the following, write $h^*(x) = (x_1, \ldots, x_m)$, and fix g_i to be any generator of $\mathbb{Z}_{q_i}^*$. Then $x_i = g_i^{e_i}$ for some $e_i \in \mathbb{Z}_{\phi(q_i)}$. Hence there is an isomorphism $\overline{h}$ from the multiplicative group G^* to the additive group

$$\overline{G} := \mathbb{Z}_{\phi(q_1)} \times \mathbb{Z}_{\phi(q_2)} \times \cdots \times \mathbb{Z}_{\phi(q_m)}, \tag{44}$$

where $\overline{h}(x_1, \ldots, x_m) = (e_1, \ldots, e_m)$. The group operation in $\overline{G}$ is componentwise addition, modulo $\phi(q_i)$ in the ith component. Remark that this is essentially the discrete logarithm map, since we may write $e_i = \log_{g_i}(x_i)$.

4. *If* $\mathtt{ord}_n(x) = r$ *and* $\mathtt{ord}_n(x_i) = r_i$, *then* r *is equal to* $\ell := \mathtt{LCM}(r_1, \ldots, r_m)$. In proof, note that $x_i^r \equiv 1 \pmod n$ implies $x_i^r \equiv 1 \pmod{q_i}$ and so $r_i | r$. Thus $r \geq \ell$ (by definition of LCM). But for each i, $x^\ell \equiv (x^{r_i})^{\ell/r_i} \equiv 1 \pmod{q_i}$ and so $h^*(x^\ell) = (1, 1, \ldots, 1)$. Thus $x^\ell \equiv 1 \pmod n$. This implies $r \leq \ell$ (by definition of order). This proves that $r = \ell$.

5. *The fraction of elements in* $\mathbb{Z}_n^*$ *of odd order modulo* n *is at most* 2^{-m}. From $x^r \equiv 1 \pmod n$, we conclude that $g_i^{e_i r} \equiv 1 \pmod{q_i}$ and so $\phi(q_i) | e_i r$. But $\phi(q_i)$ is even. If e_i is odd, then r must be even. Therefore, a necessary condition for r to be odd is that every e_i be even. The fraction of elements x in $\mathbb{Z}_n^*$ such that $x \mapsto (e_1, \ldots, e_m)$ with all e_i even is exactly $1/2^m$.

¶55. Probability of finding nontrivial squareroots of unity. We introduce a useful notation: for $n \in \mathbb{N}$, let $\mathrm{v}(n)$ be the largest $d \geq 0$ such that $2^d | n$. This function is not defined for $n = 0$. We can generalize this to any nonzero rational number n/m: $\mathrm{v}(n/m) := \mathrm{v}(n) - \mathrm{v}(m)$, which can be negative or positive.

¶56. Now let $\mathbb{Z}_n^k := \{i \in \mathbb{Z}_n : i > 0, \mathrm{v}(i) = k\}$ and $f_n(k) = |\mathbb{Z}_n^k|/n$. For instance, $f_{16}(k) = \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, 0$ for $k = 0, 1, 2, 3, 4$. We may verify (Exercise) that for all $k \geq 0$:

$$\begin{aligned}
f_n(0) &\leq 1/2, \text{ with equality iff } n \text{ even}, && (45)\\
\sum_{\ell \geq 1} f_n(\ell) &\leq 1/2, \text{ with equality iff } n \text{ odd}. && (46)\\
f_n(k+1) &\leq f_n(k) && (47)\\
f_n(k) &\leq \frac{1}{2^k + 1} && (48)\\
f_n(k) &\geq \frac{1}{3 \cdot 2^k}, \text{ provided } f_n(k) > 0. && (49)
\end{aligned}$$

It should be remarked that in our application, n will be the $\phi(q_i)$'s in (44).

¶57. The elements of $\mathbb{Z}_n^*$ with even order is particularly interesting to us: for if $\mathtt{ord}(x) = r$ is even, then $x^{r/2}$ is a squareroot of unity. Let us define $E = E_n := \{x \in \mathbb{Z}_n^* : \mathtt{ord}_n(x) = \text{even}\}$ and for $x \in E$, let $s(x) := x^{\mathtt{ord}(x)/2} \bmod n$. From the preceding, we know that $|E| \geq \phi(n)(1 - 2^{-m}) \geq \frac{3}{4}\phi(n)$.

¶58. CLAIM. *At most* $\phi(n)/2$ *of the elements in* E *give rise to trivial squareroots of unity.* Let us call $x \in E$ a "bad element" if $s(x) = -1$. Hence, if $s(x)$ is a trivial squareroot of unity, then x must be a bad element since we already know that $s(x) \neq 1$. Let $\mathtt{ord}(x) = r$ and $\overline{h}(x) = (g_1^{e_1}, \ldots, g_m^{e_m})$. Then x is bad implies that

$$g_i^{e_i r/2} \equiv -1 \pmod{q_i}$$

for all $i = 1, \ldots, m$. Clearly, $e_i r > 0$ so that $\mathrm{v}(e_i r)$ is defined. Now $g_i^{e_i r} \equiv 1 \pmod{q_i}$ implies $\phi(q_i) | e_i r$, and so

$$\mathrm{v}(\phi(q_i)) \leq \mathrm{v}(e_i r). \tag{50}$$

If $g_i^{e_i r/2} \equiv -1 \pmod{q_i}$ then $\phi(q_i) \nmid \frac{e_i r}{2}$, and so from (50), we conclude that

$$\mathrm{v}(\phi(q_i)) = \mathrm{v}(e_i r). \tag{51}$$

Comparing this equation for $i = 1$ and $i = 2$, we obtain

$$\mathrm{v}(e_1/e_2) = \mathrm{v}(\phi(q_1)/\phi(q_2)). \tag{52}$$

The righthand side is a constant k. Without loss of generality, assume $k \geq 0$ (otherwise, consider $\mathrm{v}(e_2/e_1)$).

But for each choice of $\ell = 0, 1, \ldots$, the number of $(e_1, e_2) \in \mathbb{Z}_{\phi(q_1)} \times \mathbb{Z}_{\phi(q_2)}$ such that $\mathrm{v}(e_2) = \ell$ and $\mathrm{v}(e_1) = \mathrm{v}(e_2) + k = \ell + k$ is at most $f_{\phi(q_2)}(\ell)\phi(q_1)/2$, using the bounds (45) and (46). Summing over all ℓ, the number of $(e_1, e_2) \in \mathbb{Z}_{\phi(q_1)} \times \mathbb{Z}_{\phi(q_2)}$ such that (52) holds is at most $\sum_{\ell \geq 0} f_{\phi(q_1)}(\ell)\phi(q_2)/2 = \phi(q_1)\phi(q_2)/2$.

Under the isomorphism that maps $x \in \mathbb{Z}_n^*$ to $h^*(\overline{h}(x)) = (e_1, \ldots, e_m)$, we see that the bad elements of E maps to a set of size at most $\frac{1}{2}\prod_{i=1}^m \phi(q_i) = \phi(n)/2$. This proves our claim that there are at most $\phi(n)/2$ bad elements.

¶59. Finally, the following result is the essence of the assertion that integer factorization is reducible to order finding:

Lemma 12. *The probability p that a random element $x \in \mathbb{Z}_n^*$ has even order and $s(x)$ is a nontrivial squareroot of unity modulo n is at least $p \geq 1/4$.*

Proof. We have $|E| \geq 3\phi(n)/4$ since at most $1/2^m \leq 1/4$ of the elements in $\mathbb{Z}_n^*$ have even order. At most $\phi(n)/2$ of the elements in E are bad, so at least $|E| - \phi(n)/2 \geq \phi(n)/4$ are not bad. **Q.E.D.**

¶60. Reducing factorization to order finding. Here is the algorithm to do simple factorization via order finding: given N to be factored, we may assume N is a non-power odd number. First we choose a random value $x \in \mathbb{Z}_N$. We must check if $x \in \mathbb{Z}_N^*$, by computing $\texttt{GCD}(x, N)$. If this GCD is not 1, we have in fact found a factor! Hence we may assume $x \in \mathbb{Z}_N^*$ and proceed to find its N-order r. If N is composite, with probability $p \geq 1/4$, r would be even and $s(x) = x^{r/2}$ a non-trivial squareroot of 1 modulo N. Thus with probability $p \geq 1/4$ we can factor N. In the contrary case, we can repeat this test $k \geq 2$ times. If we fail to factor N for k times, we declare N to be prime.

What is the probability of error? Error can only occur if N is composite and we declare it prime. But this happens only if we fail the test for k times. This probability is at most $(1-p)^k \leq (3/4)^k$, which can be as small as we like by making k large enough. For instance $k = 3$, the probability of error would already be less that 50 percent. With $k = 17$, the error probability is less than 1 percent.

But what happens when the order finding algorithm itself is randomized? That is, given $N, x, \varepsilon > 0$, the algorithm returns an r such that $\Pr\{r \neq \texttt{ord}_N(x)\} \leq \varepsilon$. We leave as an exercise to show the following:

Theorem 13. *If there is a randomized polynomial time algorithm for order finding, then there is a randomized polynomial time algorithm for integer factorization.*

Example (contd). Continue with $n = 45 = 9 \times 5$. Let $r = \texttt{LCM}(\phi(9), \phi(5)) = 12$. Consider x such that $\overline{h}(x) = (2, 2)$. We have $2^6 \equiv 5^6 \equiv 1 \pmod 9$ and also $2^6 \equiv 3^6 \equiv -1 \pmod 5$. Hence $\overline{h}(x^6) = (1, -1)$. But what is x^6? Well, $x^6 \equiv 1 \pmod 9$ means $(x^6 \bmod 45) \in \{1, 10, 19, 28, 37\}$. Similarly, $x^6 \equiv -1 \pmod 5$ implies $(x^6 \bmod 45) \in \{4, 9, 14, 19, 24, 29, 34, 39, 44\}$. This means $x^6 = 19$. Check: Modulo 45, we have $19^2 = (20-1)^2 = 400 - 40 + 1 \equiv -50 + 5 + 1 \equiv 1$. Thus $y = 19$ is a nontrivial square root of unity. Our reduction tells us that $\texttt{GCD}(45, y-1)$ or $\texttt{GCD}(45, y+1)$ must be nontrivial. Indeed, $\texttt{GCD}(45, y-1) = 9$ and $\texttt{GCD}(45, 20) = 5$.

The preceding development shows why assuming n is odd in our reduction of the factorization problem is not an arbitrary decision: it ensures that each $\mathbb{Z}_{q_i}^*$ is cyclic.

¶61. Improved Estimate. Let the number of bad elements in E be $\alpha\phi(n)$. Above we show that $\alpha \leq 1/2$. Let us improve this to

$$\alpha \leq \frac{3}{7}. \tag{53}$$

For $m, n, k \in \mathbb{N}$, define $\mathbb{Z}_{m,n}^k := \{(d, e) \in \mathbb{Z}_m \times \mathbb{Z}_n : d > 0, e > 0, \mathrm{v}(d) - \mathrm{v}(e) = k\}$, and $f_{m,n}(k) = |\mathbb{Z}_{m,n}^k|/mn$. It is easy to verify the following:

$$\begin{aligned} f_{m,n}(k) &= \sum_{\ell \geq 0} f_m(k+\ell) f_n(\ell) \\ f_{m,n}(k+1) &\leq f_{m,n}(k) \end{aligned}$$

Clearly, (53) follows from the following bound:

$$f_{m,n}(k) < \frac{1518}{3600} < \frac{3}{7}. \tag{54}$$

It is enough to show that $f_{m,n}(0) \leq \frac{1518}{3600}$. From (48), we have

$$f_m(0) \leq 1/2, \quad f_m(1) \leq 1/3, \quad f_m(2) \leq 1/5, \quad f_m(\ell) < 2^{-\ell}.$$

Then

$$\begin{aligned} f_{m,n}(0) &= \sum_{\ell \geq 0} f_m(\ell) f_n(\ell) \\ &< f_m(0)f_n(0) + f_m(1)f_n(1) + f_m(2)f_n(2) + \sum_{\ell \geq 3} (2^{-\ell})^2 \\ &\leq \frac{1}{4} + \frac{1}{9} + \frac{1}{25} + \frac{1}{48} \\ &= \frac{1518}{3600}. \end{aligned}$$

Exercises

Exercise 14.1: Prove that the function $h^* : \mathbb{Z}_n^* \to G^*$ in (43) is an isomorphism. ♢

Exercise 14.2: Consider $n = 1991 = 11 \times 181$. What is the probability that if you pick a random number $x \in \mathbb{Z}_n$, you can factor n using x, using our reduction to order finding? ♢

Exercise 14.3: Carry out the steps of the factorization algorithm outlined in the text for the input 225. ♢

Exercise 14.4: Show that integer factorization and order finding are polynomially equivalent problems: we have shown that integer factorization is polynomial time reducible to order finding. Show the converse reduction. ♢

Exercise 14.5: Give an estimate of the complexity of factoring prime powers (the second reduction). ♢

Exercise 14.6: Give an upper bound $T(L)$ on the overall complexity of factorization of integers, where $L = \lg N$ in the bit size of the input integer N. You are to use the four reductions described in the text, and to express $T(L)$ relative to the complexity, $t(L)$, of finding orders of odd non-powers. ♢

Exercise 14.7: Prove Theorem 13. ♢

Exercise 14.8: Improve the improved estimate. ♢

End Exercises

§15. Quantum Order Finding

With our preparation in basic techniques of quantum algorithms, and in number theoretic tools, we are ready for the final assault on factorization. Now Theorem 13 assures us that polynomial-time integer factorization can be reduced to polynomial-time order finding. This reduction (as worked out in the previous section) is best done on a classical computer: we only use quantum power for the finding orders.

¶62. To see concretely what is needed, let $N = 2^n$ be fixed. Given $3 \leq m < N$ and $x \in \mathbb{Z}_m^*$, we want to find the m-order of x. For instance, $N = 16 = 2^4$, $m = 15$ and $x = 7$. Then the 15-order of x is $r = 4$. But how can we produce r from m, x using a quantum algorithm? Note that the order problem on m is embedded in N. The trick is to define a unitary operator $U_{m,x}$ and an eigenvector $|v\rangle$ such that $U_{m,x}(|v\rangle) = \omega^{f(r)}|v\rangle$ where $f(r)$ is some easily inverted function of r. Then by phase estimation, we can approximate $f(r)$ and then invert $f(r)$ to get r.

¶63. The operator $U_{m,x}$. The unitary operator $U_{m,x}$ will act on the usual state space of n qubits, with the eigenstates $|0\rangle, \ldots, |N-1\rangle$. This operator is completely described by its actions on the eigenstates:

$$U_{m,x}(|j\rangle) := \begin{cases} |jx \bmod m\rangle & \text{if } j \in \mathbb{Z}_m, \\ |j\rangle & \text{else.} \end{cases} \tag{55}$$

Since x has an inverse modulo m, $jx \equiv j'x \pmod m$ implies $j \equiv jxx^{-1} \equiv j'xx^{-1} \equiv j' \pmod m$. Thus $U_{m,x}$ is a permutation matrix, and *a fortiori*, a unitary matrix.

¶64. Eigenvectors of $U_{m,x}$. We next find eigenvectors of $U_{m,x}$. For each $s \in \mathbb{Z}_r$, let

$$|v_s\rangle := \frac{1}{\sqrt{r}} \sum_{\ell=0}^{r-1} \omega^{-s\ell/r} |x^\ell \bmod m\rangle, \tag{56}$$

where the "ω" notation of ¶38 is used. Then

$$\begin{aligned} U_{m,x}(|v_s\rangle) &= \frac{1}{\sqrt{r}} \sum_{\ell=0}^{r-1} \omega^{-s\ell/r} |x^{1+\ell} \bmod m\rangle \\ &= \frac{1}{\sqrt{r}} \omega^{s/r} \sum_{\ell=0}^{r-1} \omega^{-s(1+\ell)/r} |x^{1+\ell} \bmod m\rangle \\ &= \omega^{s/r} |v_s\rangle. \end{aligned}$$

The last equality follows from the fact that $x^r \equiv 1 \pmod m$ and $\omega^{-s} = 1$ since s is an integer. Thus, $|v_s\rangle$ is an eigenvector of $U_{m,x}$ with phase $\phi = s/r$. If we can estimate ϕ, and assuming we know s, we can trivially recover r, provided $s \neq 0$. Unfortunately, we do not know how to prepare the state $|v_s\rangle$ for a single s. To circumvent this problem, we use another observation: for any k,

$$\begin{aligned} \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \omega^{sk/r} |v_s\rangle &= \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \omega^{sk/r} \frac{1}{\sqrt{r}} \sum_{\ell=0}^{r-1} \omega^{-s\ell/r} |x^\ell \bmod m\rangle \\ &= \frac{1}{r} \sum_{\ell=0}^{r-1} \sum_{s=0}^{r-1} \omega^{(k-\ell)s/r} |x^\ell \bmod m\rangle \\ &= |x^k \bmod m\rangle \end{aligned} \tag{57}$$

where the last equation follows from the fact that (see (25)) $\sum_{s=0}^{r-1} \omega^{(k-\ell)s/r}$ vanishes for $\ell \neq k$, and otherwise equals r. A special case of this identity is when $k = 0$,

$$|1\rangle = \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} |v_s\rangle. \tag{58}$$

¶65. Exponentiation circuit. To apply the result of (57), we need to produce the state $|x^k \bmod m\rangle$ for each k. Hence, we want to construct the exponentiation transformation,

$$|k\rangle \otimes |y\rangle \mapsto |k\rangle \otimes |x^k y \bmod m\rangle.$$

The circuit V is shown in Figure 14, and it uses as subcircuits denoted $U = U_{m,x}$ (to be more precise, control-$U_{m,x}$). Also, U^i denotes i-fold application of U.

Note that this circuit is similar to the first stage of phase estimation (¶47, see Figure 11). What we must remember, however, is that the U^{2^i}-gates must be implemented efficiently (polynomial in i, not in 2^i). This uses the well-known successive squaring trick of classical exponentiation. We leave this as an Exercise. As in Figure 11, we actually use an control-U gates, and the justification is also similar: the control-U^{2^i}, on input $|k_i\rangle \otimes |yx^{(k_1 \cdots k_{i-1})_2}\rangle$, produces the output $|k_i\rangle \otimes |yx^{(k_1 \cdots k_{i-1} k_i)_2}\rangle$.

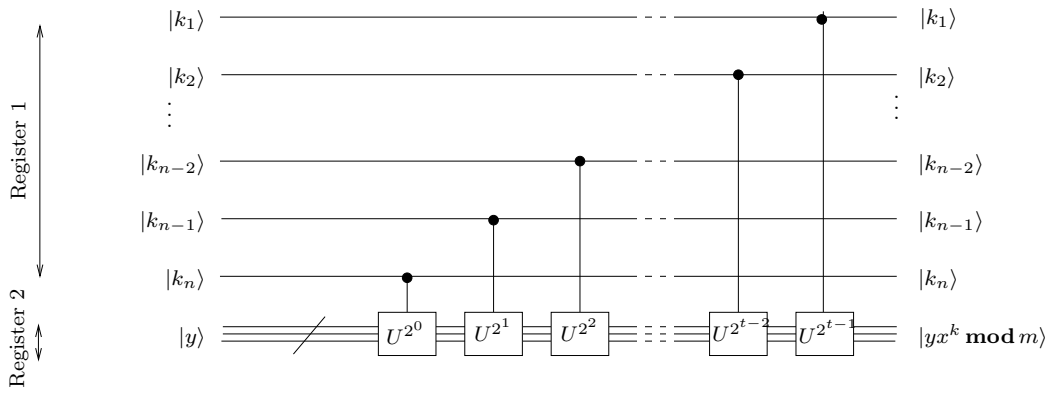

Figure 14: Exponentiation Circuit for $V = V_{m,x} : |k\rangle \otimes |y\rangle \mapsto |k\rangle \otimes |x^k y \bmod m\rangle$.

¶66. Combining. We combine the previous two ideas. Let us start with two quwords, each with n bits. Call these the first and second registers, respectively. In bit notation ¶7, we prepare them as follows:

$$|0^n\rangle \otimes |0^{n-1}1\rangle.$$

These state can also be written as $|0\rangle \otimes |1\rangle$ in the indexing notation ¶7. Applying the Hadamard transformation H_n on the first register, we obtain

$$\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle \otimes |1\rangle. \tag{59}$$

Note that $|0^{n-1}1\rangle = |1\rangle$ where we intepret $0^{n-1}1$ as the binary number 1. Next apply the exponentiation operator $V = V_{m,x}$ above to the two registers to yield

$$\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} |k\rangle \otimes |x^k \bmod m\rangle.$$

From (57), this last result can be expressed as

$$\frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \left(|k\rangle \otimes \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} \omega^{sk/r} |v_s\rangle \right) \quad = \quad \frac{1}{\sqrt{rN}} \sum_{s=0}^{r-1} \left(\sum_{k=0}^{N-1} |k\rangle \otimes \omega^{sk/r} |v_s\rangle \right). \tag{60}$$

The expression in the final pair of parentheses is similar to a Fourier transform (¶39), but with phase $\omega^{s/r}$. We next need to estimate the phase angle s/r.

¶67. Phase Estimation. Let $V = V_{m,x}$ denote the exponentiation circuit of Figure 14. This circuit transforms the input (59) to (60). To see more clearly the actions of V, let us rewrite

$$\begin{aligned} V(|k\rangle \otimes |1\rangle) &= V\left(|k\rangle \otimes \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} |v_s\rangle \right), \qquad \text{(by (58))} \\ &= \frac{1}{\sqrt{r}} \sum_{s=0}^{r-1} |k\rangle \otimes \omega^{sk/r} |v_s\rangle \\ &= \frac{1}{\sqrt{r}} |k\rangle \otimes \sum_{s=0}^{r-1} \omega^{sk/r} |v_s\rangle \end{aligned}$$

Even more simply,

$$V(|k\rangle \otimes |v_s\rangle) = \omega^{sk/r} |k\rangle \otimes |v_s\rangle \tag{61}$$

Thus $|k\rangle \otimes |v_s\rangle$ is an eigenvector of V with phase $\omega^{sk/r}$. Now view V as a blackbox, illustrated in Figure 15(a). To estimate the phase angle sk/r, we use our previous phase estimation technique, by introducing the third register with t qubits, prepared as $|0^t\rangle$. Putting all these together, we have the circuit given in Figure 15(b).

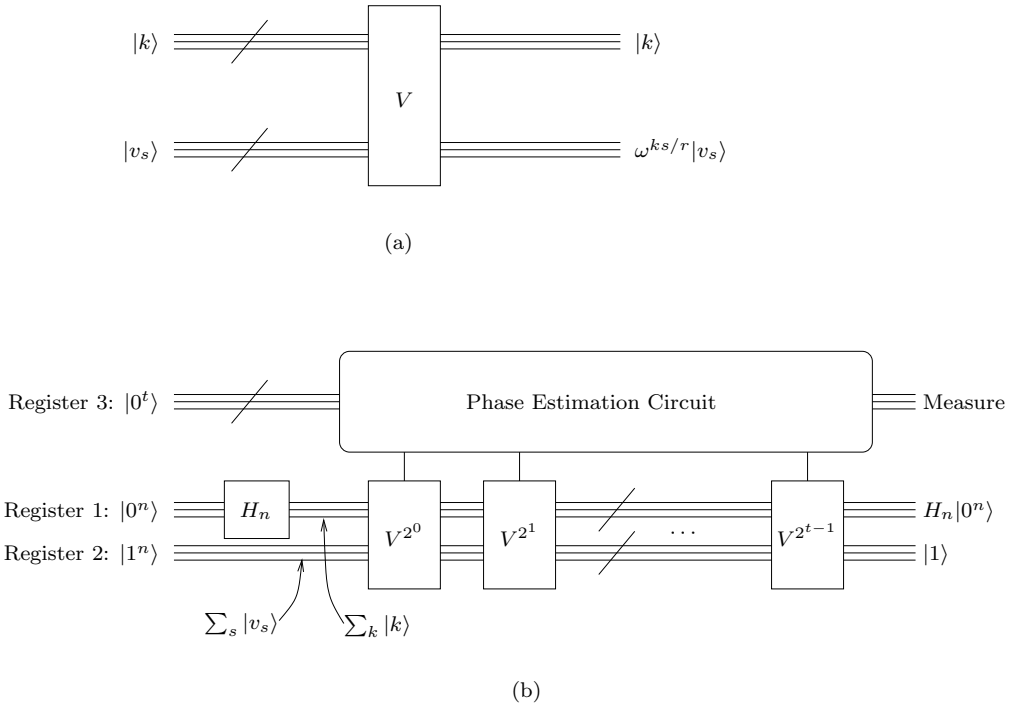

Figure 15: (a) Blackbox V. (b) Circuit for Order Finding.

¶68. Measurement. Suppose we perform the measurement on register three, as shown in Figure 15(b). Since the input to each control-V^{2^i} gate is a linear superposition of the eigenstates $|k\rangle \otimes |v_s\rangle$ for all $k = 0, \ldots, N-1$ and $s = 0, \ldots, r-1$, the result of this measurement an approximate value of $|sk/r\rangle$ for some k and s.

We can recover k by measuring register one. This gives an estimate for s/r unless $k = 0$. Furthermore, if $s = 0$ or $\texttt{GCD}(s, r) > 1$, the estimate is also useless. Otherwise, we are indifferent as to which s or k was measured. Since at least $Cr/\ln(r)$ numbers less r are prime, the probability of obtaining an s that is coprime to r is at least $C/\ln(r)$. So if we repeat at least $2C/\ln(r)$ times, we have a strong chance of obtaining a good s. The reader may device other ways to correctly estimate r despite the presense of bad choices of s.

¶69. Final Recovery. We did not specify t above because it depends on the accuracy we desire for estimating s/r. Here is where we use a well-known fact about rational approximation: if $|w - s/r| \leq 1/(2r^2)$, then we can recover s/r uniquely from w, using the simple continued fraction algorithm. Of course, we do not know r, but only need at upper bound on r (e.g., N will do). So we just make sure that, with high probability, our estimate has at least $\lg(2N^2)$ bits of precision. The continued fraction algorithm does not need any quantum power.

This concludes our description of Shor's algorithm for order finding.

Exercises

Exercise 15.1: Describe an algorithm for the complete factorization of an integer $N \in \mathbb{N}$ that runs in time polynomial in N. ♢

Exercise 15.2: Describe a polynomial time algorithm which, given $N \in \mathbb{N}$, either detects that N is not a power or else completely factorize N, *i.e.*, finds $M, e \in \mathbb{N}$ such that $N = M^e$. HINT: how would you detect if N is a square, $N = M^2$? ♢

Exercise 15.3: (Extended Euclidean Algorithm) Modify the Euclidean algorithm to compute s, t, d for any input numbers m, n, such that $d = \texttt{GCD}(m, n)$ and $d = sm + tn$. ♢

Exercise 15.4: If n is an odd non-power, then $\mathbb{Z}_n^*$ is non-cyclic. ♢

Exercise 15.5: Verify the bounds in (45)-(49). ♢

Exercise 15.6: Improve the lower bound on α in Equation (53). HINT: We can just expand more terms in the proof of (54). ♢

Exercise 15.7: Let $x \in \mathbb{Z}_n^*$ and $n = q_1 \cdots q_m$ ($m \geq 2$) where the q_i's are prime powers, and coprime to each other. Assume $\overline{h}(x) = (g_1, \ldots, g_m)$ where each g_i is a generator of $\mathbb{Z}_{q_i}^*$. Characterize the conditions where $x^\ell \equiv -1 \pmod n$ where $2\ell = \texttt{LCM}(\phi(q_1), \ldots, \phi(q_m))$. ♢

Exercise 15.8: Recently, it was announced that a quantum computer was able to factor the number 15. Infer what probably happened – what was, and what was not done by the quantum computer. ♢

Exercise 15.9: Let q be a prime power and $d = \mathrm{v}(\phi(q))$. Then exactly half of the elements in $\mathbb{Z}_q^*$ has q-order that is divisible by 2^d. ♢

Exercise 15.10: Let U be a unitary transformation and $|v\rangle$ an eigenvector such that $U|v\rangle = \omega^\phi|v\rangle$ where $0 \leq \phi < 1$ and, as usual, we write ω for $e^{\mathbf{i}2\pi}$. See Figure 13.

(i) Consider the circuit in Figure 13 that has two Hadamard gates and a control-U gate. What is the output of this circuit on input $|0\rangle \otimes |v\rangle$?
(ii) Show that the probability of $|0\rangle$ on the control line is $p_0 = (1 + \cos 2\pi\phi)/2$.
(iii) Suppose X is the number of heads in n tosses of a coin. If the coin has probability p ($0 < p < 1$) of showing up heads, then

$$\Pr\{|p - (X/n)| > \varepsilon\} \leq 2\exp(-2n\varepsilon^2).$$

This is known as the "Hoeffding bound". Suppose $\cos 2\pi\phi = \pm 0.b_1b_2b_3\cdots$ in binary notation. Using the Hoeffding bound as well as the quantum circuit, describe an experimental procedure to estimate the sign $\sigma \in \{\pm 1\}$ and first two bits b_1, b_2 so that

$$\Pr\{|\cos 2\pi\phi - \sigma 0.b_1b_2| > 1/8\} \leq \delta \tag{62}$$

where $0 < \delta < 1$ is given.
(iv) Outline a generalization of (iii), so that we efficiently estimate the first m bits of $\cos 2\pi\phi$ with error probability δ. That is, $\Pr\{|\cos 2\pi\phi - 0.b_1b_2\cdots b_m| > 2^{m+1}\} \leq \delta$. HINT: use the control-$U^{2^i}$ circuits ($i \geq 1$) to estimate the bits in parallel. ♢

End Exercises

§16. Grover's Search Algorithm

¶70. We consider another application where quantum power has a provable advantage over classical computation. This is the problem of searching a set B for an item. Assume there are n items in B, and we are looking for a particular element x_0. To do this, assume we have an oracle which, given any $x \in B$, can tell us whether or not $x = x_0$. Clearly, this requires n comparisons in the worst case. On average, we still need $n/2$ comparisons. Grover (1996) showed that a quantum algorithm can do this in $\sqrt{n}$ queries with high probability.

¶71. Oracle operators. Let $f : \mathbb{B}^n \to \mathbb{B}$ be any function. By an **f-oracle** we mean the unitary operator $O_f : \mathbb{B}^{\otimes n} \otimes \mathbb{B} \to \mathbb{B}^{\otimes n} \otimes \mathbb{B}$ where

$$O_f(|\mathbf{x}\rangle \otimes |b\rangle) = |\mathbf{x}\rangle \otimes |f(\mathbf{x}) \oplus b\rangle.$$

That is, the $n+1$st qubit b is flipped iff $f(\mathbf{x}) = 1$. Setting $|b\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$, then as in Deutsch's problem,

$$\begin{aligned} O_f(|\mathbf{x}\rangle \otimes |b\rangle) &= |\mathbf{x}\rangle \otimes \frac{1}{\sqrt{2}}(|f(\mathbf{x}) \oplus 0\rangle - |f(\mathbf{x}) \oplus 1\rangle) \\ &= (-1)^{f(\mathbf{x})}|\mathbf{x}\rangle \otimes \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \\ &= (-1)^{f(\mathbf{x})}|\mathbf{x}\rangle \otimes |b\rangle. \end{aligned}$$

In other words, $|\mathbf{x}\rangle \otimes |b\rangle$ is an eigenvector of O_f with eigenvalue $(-1)^{f(\mathbf{x})}$.

¶72. Example. Fix $1 < m < 2^n$, and let $f : \mathbb{B}^n \to \mathbb{B}$ be the function $f(j) = 1$ iff $j|m$ (where $j \in \mathbb{B}^n$ is interpreted as a binary number). Thus, the oracle O_f amounts to searching for a factor of m. We leave it as an exercise to construct a quantum circuit for O_f.

¶73. The Search Problem. Let $N = 2^n$ and $f : \mathbb{B}^n \to \mathbb{B}$. Viewing $\mathbb{B}^n$ as the set $\mathbb{Z}_{2^n} = \{0, 1, \ldots, 2^n - 1\}$, we define the sets

$$P := \{j \in \mathbb{B}^n : f(j) = 1\}, \quad Q := \{j \in \mathbb{B}^n : f(j) = 0\}.$$

Also, $p := |P|$ and $q := |Q|$. We want to construct a quantum circuit using the oracles O_f so that, if we measure the first n output lines, we obtain some value $j \in P$ with high probability.

This is interpreted as the problem of searching for an element of $j \in \mathbb{B}^n$ with property P.

The idea is easy to see geometrically: consider the states $|P\rangle := \frac{1}{\sqrt{p}} \sum_{j \in P} |j\rangle$ and $|Q\rangle := \frac{1}{\sqrt{q}} \sum_{j \in Q} |j\rangle$. Clearly, $|P\rangle$ and $|Q\rangle$ are orthogonal states:

$$\langle P|Q\rangle = 0.$$

Let the initial state of the system be

$$|\phi\rangle = \frac{1}{\sqrt{N}} \sum_{j=0}^{2^n-1} |j\rangle.$$

We know that $|\phi\rangle$ can be prepared as

$$H^{\otimes n}|0\rangle := \frac{1}{\sqrt{N}} \sum_{j=0}^{2^n-1} |j\rangle.$$

We may verify that

$$|\phi\rangle = \sqrt{\frac{q}{N}}|P\rangle + \sqrt{\frac{p}{N}}|Q\rangle.$$

This shows that $|\phi\rangle$ lies in the plane spanned by $|P\rangle$ and $|Q\rangle$.

We next consider the effect of O_f on $|\phi\rangle$: this amounts to a reflection of $|\phi\rangle$ about the $|Q\rangle$:

$$O_f(|\phi\rangle) = -\sqrt{\frac{q}{N}}|P\rangle + \sqrt{\frac{p}{N}}|Q\rangle.$$

§17. Quantum Turing Machines

¶74. A circuit computes a finite function. We now address general computational models that takes inputs of arbitrarily large size. In complexity theory, we can take one of two paths. One way is to define a general computing model such as Turing machines. Another way is to start from circuits, and to define[14] a "uniform circuit family". These two approaches can also be taken to define a more general computational model for quantum computing. A basic result here is the existence of a universal Turing machine. Deutsch [15] proved the analogous result for quantum computers. Bernstein and Vazirani [12] describes a similar model, usually known as the quantum Turing machine (QTM). Benioff [6] has argued for a different basis for such models.

Need to give the results of Yao??

§A. APPENDIX: Review of Linear Algebra

We review of some facts from linear algebra needed in quantum computing. We work exclusively with finite-dimensional linear spaces over the complex field $\mathbb{C}$.

[14] There is a bit of circularity in the conventional definition of "uniformity" via some Turing machine (say). We can avoid this problem by using some logical circuit description language, say.

¶75. Basic notations. The (complex) conjugate $\overline{z}$ of a complex number $z \in \mathbb{C}$ is denoted $\overline{z} = x - \mathbf{i}y$ where $z = x + \mathbf{i}y$ with $x, y \in \mathbb{R}$ and $\mathbf{i} = \sqrt{-1}$. We will shortly see that the complex conjugate of z can also be denoted z^*, and this form is common in quantum physics. The **absolute value** $|z|$ of z is equal to $\sqrt{z\overline{z}} = \sqrt{x^2 + y^2}$. Complex numbers with absolute value 1 are called **complex signs**, a generalization of the real signs ± 1. A complex sign z can be written in the form $e^{\mathbf{i}\theta}$ for some real θ. The value θ is also called the **phase**. If $|z| \leq 1$, we call z a (probability) **amplitude**.

A matrix $A \in \mathbb{C}^{m\times n}$ is viewed as a transformation $t_A : \mathbb{C}^n \to \mathbb{C}^m$, where $t_A(x) = Ax$ (a matrix-vector multiplication). Thus, $range(A)$ is simply $\{Ax : x \in \mathbb{C}^n\} \subseteq \mathbb{C}^m$. Let us assume $m = n$ unless otherwise noted. The (i, j)-th entry of a matrix A is denoted $(A)_{ij}$. The **transpose** A^T and **conjugate** $\overline{A}$ of a matrix A is given by $(A^T)_{ij} = (A)_{ji}$ and $(\overline{A})_{ij} = \overline{(A)_{ij}}$, respectively. Then the **conjugate transpose** A^* is given by $A^* = \overline{A^T} = \overline{A}^T$. Note that $(AB)^T = B^T A^T$ and $(AB)^* = B^* A^*$. A matrix $A \in \mathbb{C}^{n\times n}$ is **Hermitian** if $A^* = A$, **unitary** if $A^* A = I$ (identity), and **orthogonal** if $A^T A = I$. So if A is unitary then $A^{-1} = A^*$, $A^* A = AA^*$. Hermitian matrices are also known as "self-adjoint" matrices (as A^* is sometimes called the "adjoint" of A). In case A is a 1×1 matrix, A^* is just another way of writing complex conjugation, since $A^* = \overline{A}$. Unitary matrices are fundamental in quantum computing. For a unitary U, it is clear that $\det(U) = 1$ and hence its eigenvalues λ_i are complex signs, $|\lambda_i| = 1$.

Define the **scalar product** for two n-vectors x, y such that $\langle x, y\rangle := \sum_{i=1}^n x_i^* y_i$. Thus, $\langle x, x\rangle$ is real and nonnegative. Moreover, $\langle x, x\rangle = 0$ iff $x = \mathbf{0}$. Define $\|x\| := \sqrt{\langle x, x\rangle}$, called the (Euclidean) **norm** of x. Two vectors x, y are **orthogonal** if $\langle x, y\rangle = 0$.

A matrix A is **normal** if $A^* A = AA^*$. Note that unitary matrices, Hermitian matrices, skew-Hermitian matrices ($A^* = -A$) are all normal. In quantum mechanics, Hermitian and unitary matrices are of paramount importance.

¶76. Orthogonalization and QR-Factorization. A useful tool for investigating the structure of linear spaces is based on a certain factorization of matrices. Let $A = [a_1|a_2|\cdots|a_m] \in \mathbb{C}^{n\times m}$ where a_i is the ith column. Let $S_i \subseteq \mathbb{C}^n$ be the subspace spanned by the first i columns.

Assume A has rank m (so $m \leq n$). Let the sequence $(q_1, \ldots, q_m)$ of vectors form an orthonormal basis for $S_m \subseteq \mathbb{C}^n$. If we form the matrix

$$Q := [q_1|q_2|\cdots|q_m],$$

then there is some $m \times m$ matrix R such that

$$A = QR. \tag{63}$$

To see this, note that the ith column r_i of R represents the vector a_i relative to the basis $(q_1, \ldots, q_m)$. Let us call $(q_1, \ldots, q_m)$ (or Q) a **Gram-Schmidt basis** for A if for each $i = 1, \ldots, m$, the prefix $(q_1, \ldots, q_i)$ forms an ordered basis for S_i. An ordered basis is just a basis whose elements are linearly ordered. In this case, the matrix R is upper triangular in (63). The well-known Gram-Schmid orthogonalization procedure that compute such a basis Q from any A.

The factorization (63) of A is known as a **reduced QR-factorization** when Q is an Gram-Schmidt basis of A. Each vector q_i in this basis is unique up to some scalar multiple of modulus 1. Equivalently, we say q_i is determined up to "complex signs". To make the factorization unique, we choose the complex signs to make the diagonal elements of R real and non-negative. The **full QR-factorization** of A is the following variant,

$$A = \widehat{Q}\widehat{R} \tag{64}$$

where $\widehat{Q}$ is $n \times n$ and $\widehat{R}$ is $n \times m$. It is obtained from (63) by augmenting Q and R: the matrix $\widehat{Q}$ is obtained by appending $n - m$ additional columns so that the columns of $\widehat{Q}$ form an orthonormal basis for $\mathbb{C}^n$; the matrix $\widehat{R}$ is obtained by appending $n - m$ additional rows of 0's. Note that when $m = n$, the full QR-factorization is just the reduced QR-factorization. Since $\widehat{Q}^*\widehat{Q} = I$, the matrix Q is unitary in the full QR-factorization.

So far, we have assumed that A has rank m. Suppose m, n are arbitrary and A has has rank k (so $k \leq \min\{m, n\}$. We first apply a permutation P to the columns of A so that the first k columns are linearly independent. Then we have $AP = QR$ where $Q \in \mathbb{C}^{n\times n}$ is unitary, $R \in \mathbb{C}^{n\times m}$ is upper triangular, and the first k columns of Q forms a orthonormal basis for S_m,

¶77. Eigenvalues and Eigenvectors. A non-zero vector $x \in \mathbb{C}^n$ is called an **eigenvector** of A if $Ax = \lambda x$ for some $\lambda \in \mathbb{C}$. In this case, we call λ an **eigenvalue** of A that is associated with the eigenvector x. Note that while eigenvectors must be non-zero, we have no such restriction on eigenvalues. In particular, A is singular iff $\lambda = 0$ is an eigenvalue: $Ax = 0x = 0$ ($x \neq 0$) iff A is singular. The set of eigenvalues of A, denoted $\Lambda(A)$, is called the **spectrum** of A. The **characteristic polynomial** of A is $p_A(z) = \det(zI - A)$.

For example, if $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$ then $p_A(z) = \det \begin{bmatrix} z-a & b \\ c & z-d \end{bmatrix} = (z-a)(z-d) - bc = z^2 - (a+d)z + ad - bc$. If

$$A = \begin{bmatrix} 0 & 1 & & & \\ & 0 & 1 & & \\ & & \ddots & & \\ & & & 0 & 1 \\ -a_0 & -a_1 & \cdots & -a_{n-2} & -a_{n-1} \end{bmatrix}$$

then $p_A(z) = z^n + \sum_{i=0}^{n-1} a_i z^i$.

In general, it is easily seen that $p(z)$ is monic of degree n with the constant term equal to $\det(A)$ and the coefficient of z^{n-1} equal to $-trace(A) = -\sum_{i=1}^{n} a_{ii}$. Also, $\lambda \in \Lambda(A)$ iff λ is a zero of $p_A(z)$. [In proof, $Ax = \lambda x$ iff $(\lambda I - A)x = 0$ iff $\lambda I - A$ is singular iff $\det(\lambda I - A) = 0$.] The multiplicity of λ as a root of $p_A(z)$ is called the **algebraic multiplicity** of λ. It follows that the cardinality of $\Lambda(A)$ is between 1 and n.

Example. If $A_1 = \begin{bmatrix} 4 & 3 \\ 3 & 4 \end{bmatrix}$ then $p_{A_1}(z) = z^2 - 8z + 7$ and so $\Lambda(A_1) = \{1, 7\}$. To compute the eigenvectors associated with a given eigenvalue $\lambda \in \Lambda(A)$, we find a basis for the nullspace of the matrix $\lambda I - A$ (or $A - \lambda I$). Recall that the null space of a matrix B is the set of vectors x such that $Bx = 0$. We compute the LU decomposition of B, $B = LU$ and then solve for U by backward substitution. For instance, the eigenvectors associated with $\lambda = 1$ for A_1 above is the nullspace of $A_1 - I = \begin{bmatrix} 3 & 3 \\ 3 & 3 \end{bmatrix}$. Clearly the nullspace in this case is spanned by $x_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$. So x_1 is associated with $\lambda = 1$. Similarly, the nullspace of $A - 7I = \begin{bmatrix} -3 & 3 \\ 3 & -3 \end{bmatrix}$ is spanned by $x_2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

Lemma 14. *If A can be written as a block matrix of the form*

$$A = \begin{bmatrix} B & C \\ \mathbf{0} & D \end{bmatrix}$$

where B and D are square blocks, then

$$p_A(z) = p_B(z) p_D(z). \tag{65}$$

Proof. The lemma is equivalent to $\det(zI - A) = \det(zI - B)\det(zI - D)$. We only have to show that any of the $n!$ terms of the determinant $\det(zI - A)$ that involves an entry of C is zero. Suppose B is $k \times k$ and the term $t = \prod_{i=1}^{n} a_{i,j(i)}$ contains an entry $a_{i_0,j(i_0)}$ of C. Then the i_0th row of B does not contribute to t. If B contributes ℓ entries to t, this implies $\ell \leq k-1$. So for some $1 \leq c \leq k$, the c-th column of B does not contribute to t. Consider the index i_1 where $j(i_1) = c$: clearly, the factor $a_{i_1,j(i_1)}$ of t is 0 and so $t = 0$, concluding our proof. **Q.E.D.**

As corollary, we have $\Lambda(A) = \Lambda(B) \cup \Lambda(D)$.

¶78. Invariant subspaces. A subspace $E \subseteq \mathbb{C}^n$ is **A-invariant** if $AE = \{Ax : x \in E\}$ is contained in E. If $E = \{0\}$ then it is clearly A-invariant. We call this the trivial case. If $\lambda \in \Lambda(A)$, the set $E_\lambda = \{x \in \mathbb{C}^n : Ax = \lambda x\}$ is easily seen to be a non-trivial A-invariant subspace; we call E_λ the **eigenspace** of A associated with λ. In particular, the eigenspace of A associated with $\lambda = 0$ is the nullspace of A. If $\lambda \neq \lambda'$ then clearly $E_\lambda \cap E_{\lambda'} = \{0\}$. Further, $x \in E_\lambda$ and $y \in E_{\lambda'}$ are linearly dependent: for, if $c = ax + by = 0$ for some $a, b \in \mathbb{C}$ then $Ac = \lambda ax + \lambda' by = 0$, which easily implies $a = b = 0$. In E_λ, transformation by A amounts to scaling by a factor of λ. The dimension of E_λ is called the **geometric multiplicity** of λ. We will show below that the geometric multiplicity of λ is at most the algebraic multiplicity.

¶79. Similarity. Invariant subspaces are intimately connected to the notion of "similarity". Two matrices $A, B \in \mathbb{C}^{n \times n}$ are **similar** if

$$A = XBX^{-1}$$

for some non-singular matrix X. Intuitively, A, B are similar means that the transformation (represented by) A in one coordinate system S is the same as transformation of B in another coordinate system S'. To see this, suppose $X : S' \to S$ is the coordinate transformation from S' to S, and let $A : S \to S$. Then the map $x \mapsto Ax$ is the same as first transforming $x \mapsto X^{-1}x \in S'$, then applying B to get $BX^{-1}x \in S'$, and finally transforming the result back to $XBX^{-1}x \in S$. Thus, similarity classifies transformations of isomorphic space S.

In case X is unitary, $X^*X = I$, we say A and B are **unitarily similar**: $A = XBX^*$. Similar matrices A, B have the same characteristic polynomial since

$$\det(zI - X^{-1}BX) = \det(X^{-1}(zI - B)X) = \det(X^{-1})\det(zI - B)\det(X) = \det(zI - B).$$

Thus similar matrices have the same spectrum,

$$\Lambda(A) = \Lambda(B), \tag{66}$$

with the same algebraic multiplicities. Next,

$$Ax = \lambda x \Leftrightarrow X^{-1}BXx = \lambda x \Leftrightarrow B(Xx) = \lambda(Xx). \tag{67}$$

This shows that the eigenspace E_λ of A with λ and the eigenspace E'_λ of B are related as follows: $E_\lambda = XE'_\lambda$. It follows that λ has the same geometric multiplicity relative to A and B.

¶80. Defective Matrices. Geometric multiplicity can be different from algebraic multiplicity. An eigenvalue λ is **defective** if its geometric multiplicity is different from its algebraic multiplicity. A matrix is **defective** if any of its eigenvalue is defective. For instance, if

$$A = \begin{bmatrix} \lambda & 0 \\ 0 & \lambda \end{bmatrix}, \qquad A' = \begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}, \qquad (\lambda \neq 0) \tag{68}$$

then $\Lambda(A) = \Lambda(A') = \{\lambda\}$ and the algebraic multiplicity of λ is 2 for both A and A'. Let E_λ and E'_λ be the eigenspaces associated with λ for A and A', respectively. Clearly, $E_\lambda = \mathbb{C}^2$ so that the geometric multiplicity of λ relative to A is 2. But if $x = (a, b)^T \in E'_\lambda$ then $A'x = \lambda x$. This means $\lambda x = (\lambda a + b, \lambda b)^T$, and thus $b = 0$. So E'_λ has dimension 1, not 2. Hence λ is defective for A'.

Lemma 15. *For all A, the geometric multiplicity of any $\lambda \in \Lambda(A)$ is at most the algebraic multiplicity of λ.*

Proof. To see this, suppose λ has geometric multiplicity m and $x_1, \ldots, x_m$ are m linearly independent eigenvectors all associated with λ. We may assume that the x_i's are unit vectors ($x_i^* x_i = 1$). Let $X = [x_1|x_2|\cdots|x_m|\cdots|x_n]$ where the columns $x_{m+1}, \ldots, x_n$ are additional unit vectors that span the complement of the eigenspace E_λ. Thus X is unitary ($X^*X = I$) and $X^{-1} = X^*$. So $AX = [\lambda x_1|\cdots|\lambda x_m|x'_{m+1}|\cdots|x'_n]$, where $x'_j = Ax_j$ for $j = m+1, \ldots, n$. Then a simple calculation shows

$$B = X^*AX = \begin{bmatrix} \lambda I & C \\ \mathbf{0} & D \end{bmatrix}$$

for some C and D. This shows that the characteristic polynomial of B is divisible by $(z-\lambda)^m$, and so the algebraic multiplicity of B is at least m. Since A and B are similar, the algebraic multiplicities λ in A and B are equal. **Q.E.D.**

¶81. Diagonalizability. The diagonal matrix whose (i,i)th element is d_i (for $i = 1, \ldots, n$) is denoted $D = \mathtt{diag}(d_1, \ldots, d_n)$. A matrix is **diagonalizable** if it is similar to a diagonal matrix. To see why this concept is useful, suppose A is diagonalizable:

$$A = XDX^{-1}, \quad D = \mathtt{diag}(d_1, \ldots, d_n) \tag{69}$$

for some X. Then observe that the columns of X are eigenvectors for A. To see this, (69) implies $AX = XD$ and hence $Ax_i = d_i x_i$ where x_i is the ith column of X. Furthermore, this set of eigenvectors is **complete**, i.e., they span the whole space.

We restate this observation as a theorem. Let e_i denote the ith **elementary vector**, with 1 in the ith position and 0 elsewhere. Thus $x_i = Xe_i$. and $d_i = De_i$.

Theorem 16. *If $A = XDX^{-1}$ where $D = \mathtt{diag}(d_1, \ldots, d_n)$ is a $n \times n$ diagonal matrix, then the set $\{Xe_1, \ldots, Xe_n\}$ is a complete set of eigenvectors of A. Moreover, d_i is the associated eigenvalue of Xe_i.*

It follows that a *diagonal matrix $D = \mathtt{diag}(d_1, \ldots, d_n)$ is always non-defective.* To see this, note that the characteristic polynomial of D is $\prod_{i=1}^n (x - d_i)$ and so each $\lambda \in \{d_1, \ldots, d_n\}$ is an eigenvalue of D and it appears in D as many times as its algebraic multiplicity. Since similarity transformations preserve eigenvalues and their multiplicities, we conclude:

Theorem 17. *A matrix A is diagonalizable iff it is non-defective.*

For instance, the matrix A' in (68) is not diagonalizable.

¶82. Unitary Similarity and Schur Form. Canonical form for matrices that are equivalent under various notions of equivalence is an extremely powerful tool in linear algebra. We will consider matrices that are equivalent under unitary similarity transformations: $A \equiv B$ iff $A = UBU^*$ for some unitary U. Unitary operators are basically isomorphisms of the inner product space: if $y = Ux$ then $y^*y = (Ux)^*(Ux) = x^*(U^*U)x = x^*x$.

The invariant subspace relationship can be captured by a matrix equation: let X be a $n \times m$ matrix whose m columns span some A-invariant subspace E. Then there exists some $B \in \mathbb{C}^{m\times m}$ such that

$$AX = XB. \tag{70}$$

Conversely, every such equation (70) shows that the space spanned by the columns of X is A-invariant.

Next, assume that $E \subseteq E_\lambda$ for some eigenvalue λ. Then B has the form λI in (70). Let $X = QR$ be a full QR-factorization of X (see (64)) where

$$R = \begin{bmatrix} T \\ \mathbf{0} \end{bmatrix}$$

for some upper triangular $T \in \mathbb{C}^{m\times m}$. To say that a matrix A is upper triangular means that $(A)_{ij} = 0$ for $j > i$. Thus (70) becomes $AQR = QRB$, or $Q^*AQR = RB$. Since $RB = \lambda R$, we have

$$(Q^*AQ)R = \lambda R. \tag{71}$$

Let us write

$$Q^*AQ = \begin{bmatrix} C & D \\ E & F \end{bmatrix}$$

where $C \in \mathbb{C}^{m\times m}$ and $F \in \mathbb{C}^{(n-m)\times(n-m)}$. Then the block version of (71) implies $ET = \mathbf{0}$. Since T is non-singular, this means $E = \mathbf{0}$:

$$Q^*AQ = \begin{bmatrix} C & D \\ \mathbf{0} & F \end{bmatrix}. \tag{72}$$

By repeated application of this transformation of A, we obtain the **Schur Decomposition** of a matrix:

Theorem 18 (Schur Decomposition). *Every matrix A is unitarily similar to a upper diagonal matrix T.*

Proof. We use induction on n where $A \in \mathbb{C}^{n\times n}$. The result is trivial for $n = 1$. So assume $n =\geq 2$ and let $Ax = \lambda x$ for some eigenvector $x \neq 0$. Then using (72) with $m = 1$, there is a unitary Q such that

$$Q^*AQ = \begin{bmatrix} c & y^T \\ \mathbf{0} & F \end{bmatrix}$$

for some $c \in \mathbb{C}$, some vector y and square matrix F. By induction, there is unitary V such that V^*FV is upper triangular. If

$$U = \begin{bmatrix} 1 & \mathbf{0} \\ \mathbf{0} & V \end{bmatrix}$$

then $T = U^*(Q^*AQ)U$ is upper triangular. As QU is unitary, this shows that A is unitarily similar to T. **Q.E.D.**

¶83. Unitary Diagonalizability. We have introduced two concepts: unitary similarity and diagonalizability. Combining them, we say a matrix A is **unitarily diagonalizable** iff it has the form

$$A = U\Lambda U^*$$

for some unitary U and $\Lambda = \mathtt{diag}(\lambda_1, \ldots, \lambda_n)$.

Clearly, unitarily diagonalizability implies diagonalizability. We show the converse as well: let A be diagonalizable, so $A = X\Lambda X^{-1}$ for some X. Here X is not necessarily unitary. By theorem 16, the columns $\{Xe_1, \ldots, Xe_n\}$ of X forms a complete set of eigenvectors for A. In order for X to be unitary, we need the eigenvectors to be normalized, i.e., $(Xe_i)^*(Xe_i) = 1$. Let $w_i := \sqrt{(Xe_i)^*(Xe_i)}$ and $u_i := (Xe_i)/w_i$. Next define the matrices U and W via

$$X = U \cdot W = [u_1|u_2|\cdots|u_n] \cdot \mathtt{diag}(w_1, \ldots, w_n).$$

Clearly, U is unitary and we have $X^{-1} = W^{-1}U^{-1} = W^{-1}U^*$. Of course, $W^{-1} = \mathtt{diag}(1/w_1, \ldots, 1/w_n)$. We then have $A = (UW)\Lambda W^{-1}U^* = U\Lambda U^*$. So A is unitarily diagonalizable. So we have proved:

Lemma 19. *A matrix is diagonalizable iff it is unitarily diagonalizable.*

Recall that a matrix A is normal if $A^*A = AA^*$. We note a useful lemma.

Lemma 20. *If A is upper diagonal, then A is normal iff A is diagonal.*

Proof. One direction is easy: if A is diagonal, then clearly $A^*A = AA^*$. Conversely, suppose $A^*A = AA^*$. We claim that A must be diagonal. Let c be the first column of A and r be the first row of A. Then top-left corner entry of A^*A is c^*c, and the corresponding entry of AA^* is r^*r. Thus A is normal implies $c^*c = r^*r$. But the first entry in c and in r are equal to some $\alpha \in \mathbb{C}$, and $c^*c = |\alpha|^2$. Hence $r^*r = |\alpha|^2$, which implies that all the remaining entries in r are zero. Continuing in this fashion, we argue that all the off-diagonal entries in the ith row must be zero. **Q.E.D.**

Theorem 21. *A matrix is normal iff it is diagonalizable.*

Proof. Let $A = UTU^*$ be the Schur decomposition of A given by the previous theorem. Since $AA^* = (UT)(T^*U^*)$ and $A^*A = (UT^*)(TU^*)$, we have

$$AA^* = A^*A \Leftrightarrow T^*T = TT^*.$$

Thus A is normal iff T is normal. But we had just shown that T is normal iff T is diagonal. Thus A is normal iff T is diagonal. But T is diagonal means A is diagonalizable. **Q.E.D.**

As a corollary, we have

- Unitary and Hermitian matrices are diagonalizable. This is because such matrices are normal. It follows that such matrices have complete sets of eigenvectors.
- A matrix is non-defective iff it is normal. This follows from theorems 21 and 17.

¶84. Projection and spectral decomposition. A **projection operator** P is a matrix satisfying the equation $P^2 = P$. Note that the operator $I - P$ is orthogonal to P since $(I - P)P = 0$. It is easy to find projections: for any unit length vector x, the matrix xx^* is a projection operator since $(xx^*)(xx^*) = x(x^*x)x^* = xx^*$. Note that xx^* is Hermitian, since $(xx^*)^* = xx^*$. For any y, $(xx^*)y = (x^*y)x = \alpha x$ where $\alpha = x^*y$ is a scalar. Thus the range of the projection (xx^*) is the linear subspace spanned by x. Generalizing this, if $\{x_i : i = 1, \ldots, k\}$ is a set of orthonormal vectors, then $P = \sum_{i=1}^k x_i x_i^T$ is a projection operator.

Theorem 22. *If A is normal, then it has a* **spectral decomposition**, *i.e., an expression of the form*

$$A = \sum_{i=1}^n \lambda_i P_i \tag{73}$$

where each P_i is a projection operator. Moreover, the $\lambda_1, \ldots, \lambda_n$ are the eigenvalues of A, and there is a unitary matrix $U = [u_1|u_2|\cdots|u_n]$ such that each P_i has the form $P_i = u_i u_i^*$.

Proof. Since A is normal, we may write $A = U\Lambda U^*$ for some unitary matrix $U = [u_1|u_2|\cdots|u_n]$, and $\Lambda = \mathtt{diag}(\lambda_1, \ldots, \lambda_n)$. It suffices to prove that

$$A = \sum_{i=1}^n \lambda_i u_i u_i^*. \tag{74}$$

By definition of $A = U\Lambda U^*$, we have $A_{jk} = r_j \Lambda c_k$ where r_j is the jth row of U and c_k is the kth column of U^*. But $r_j = (u_{1j}, u_{2j}, \ldots, u_{nj})$ and $c_k = (u_{1k}, u_{2k}, \ldots, u_{nk})^*$. Hence

$$A_{jk} = \sum_i u_{ij} \lambda_i u_{ik}^*.$$

This proves (74). We just set $P_i = u_i u_i^*$ in the theorem. **Q.E.D.**

¶85. Hermitian Matrices and Real Numbers. We show a remarkable family analogies between Hermitian matrices and real numbers. First, real numbers can be positive, negative or zero. We have analogous but more subtle classification of Hermitian matrices. A matrix A is **positive semidefinite** (resp., **positive definite**) if $x^*Ax \geq 0$ (resp., $x^*Ax > 0$) for all nonzero vector x. Analogous definitions can be given by replacing "positive" by "negative".

Theorem 23. *Let H be Hermitian.*
(i) All its eigenvalues are real.
(ii) H is positive definite iff all its eigenvalues are positive.
(iii) H is positive semidefinite iff all its eigenvalues are non-negative.

Proof. Let $H = U\Lambda U^*$ for some unitary U and diagonal Λ.
(i) Since H is Hermitian, we have $H = H^* = U\Lambda^*U^*$. Hence $\Lambda = \Lambda^*$, i.e., Λ is real.
(ii) Let the columns of U be $x_1, \ldots, x_n$. Each x_i is an eigenvector of H with associated eigenvalue λ_i. Since $x_i^* H x_i = \lambda_i |x_i|^2$, the positive definiteness of H implies $\lambda_i > 0$. Conversely, if each $\lambda_i > 0$ we can show that H is positive definite: any non-zero vector $x \in \mathbb{C}$ can be expressed as $\sum_{i=1}^n c_i x_i$ ($c_i \in \mathbb{C}$). Then $x^*Hx = \sum_{i=1}^n \lambda_i |c_i|^2 |x_i|^2 > 0$
(iii) The proof is similar to (ii). **Q.E.D.**

¶86. Analogy between Hermitian Matrices and Real Numbers. The above connection between Hermitian matrices with real numbers goes much deeper. The special role of real numbers in the complex field $\mathbb{C}$ is mirrored in many ways by the Hermitian matrices in the context of complex matrices. This analogy can be extended as follows:

real number	$\leftrightarrow$	Hermitian
pure complex number	$\leftrightarrow$	anti-Hermitian
complex sign, $\|z\| = 1$	$\leftrightarrow$	unitary
positive real	$\leftrightarrow$	positive definite Hermitian
non-negative real	$\leftrightarrow$	positive semidefinite Hermitian

In the following, z is a complex number. The complex conjugate of $z = x + \mathbf{i}y$ is $\overline{z} = x - \mathbf{i}y$. We point out the matrix analogues of the following properties:

1. z is real iff $z = \overline{z}$; z is pure complex iff $\overline{z} = -z$.
2. z is real iff $\mathbf{i}z$ is pure complex; z is pure complex iff $\mathbf{i}z$ is real.
3. $z + \overline{z}$ is real and $z - \overline{z}$ is pure complex.
4. $\overline{z}z = z\overline{z}$ is real.
5. z can be uniquely written as $z = x + \mathbf{i}y$ where x, y are real.
6. z can be uniquely written as $z = x + w$ where w is real and w is pure complex.
7. A real number r is non-negative iff there is a real number s such that $r = s^2$.
8. A z has the **polar form**, $z = rs$ where r is non-negative real, and s is a complex sign, $|s| = 1$. This form is unique if z is non-zero.
9. A complex sign s can be uniquely written as $s = e^{\mathbf{i}\theta}$ for some real θ.

In the following, let $A \in \mathbb{C}^{n\times n}$.

1. A is Hermitian iff $A = A^*$; A is anti-Hermitian iff $A^* = -A$.
   Thus, the matrix analogue of complex conjugation, $z \mapsto \overline{z}$, is conjugate transpose, $A \mapsto A^*$.
2. A is Hermitian iff $\mathbf{i}A$ is anti-Hermitian; A is anti-Hermitian iff $\mathbf{i}A$ is Hermitian.
3. $A + A^*$ is Hermitian and $A - A^*$ is anti-Hermitian.
4. A^*A and AA^* are both Hermitian.
5. A can be uniquely written as $G + \mathbf{i}H$ where G, H are Hermitian.

6. A can be uniquely written as $A = G + F$ where G is Hermitian and F is anti-Hermitian.
   This is just a restatement of the previous property.

7. A Hermition matrix H is positive semidefinite iff $H = G^*G$ for some positive semidefinite G.

8. A has two polar forms, $A = HU$ and $A = U'H'$, where H, H' are positive semidefinite Hermitian, and U, U' are unitary. If A is non-singular, then these polar forms are unique.

9. A unitary A can be uniquely written as $e^{\mathbf{i}H}$ for some Hermitian H.

Let us prove the non-obvious cases of the properties.
Property 5: Let $G = (A + A^*)/2$ and $H = (A - A^*)/2\mathbf{i}$. Then clearly $A = G + \mathbf{i}H$ and from the preceding, we conclude that G and H are Hermitian. Conversely, if $A = G + \mathbf{i}H$ for some Hermitian G and H, then $A^* = G^* + (\mathbf{i}H)^* = G^* - \mathbf{i}H^* = G - \mathbf{i}H$. It follows that $G = (A + A^*)/2$ and $H = (A - A^*)/2\mathbf{i}$.
Property 7: If A is positive semidefinite, we know that $A = U\Lambda U^*$ for some unitary U and diagonal $\Lambda = \mathtt{diag}(\lambda_1, \ldots, \lambda_n)$, where $\lambda_i > 0$. If $B = U\mathtt{diag}(\sqrt{\lambda_1}, \ldots, \mathtt{diag}\sqrt{\lambda_n}U^*$, then $BB^* = A$.
Property 8:
Property 9:

¶87. Operator functions. If $A \in \mathbb{C}^{n\times n}$, how can we define $\exp(A)$ or $\sqrt{A}$? More generally, if $f : \mathbb{C} \to \mathbb{C}$ is a function, we want to extend f to a function on matrices $f : \mathbb{C}^{n\times n} \to \mathbb{C}^{n\times n}$. When extended in this way, we call f an **operator function** (since matrices represent operators). We do not define $f(A)$ for arbitrary $A \in \mathbb{C}^{n\times n}$, but only when A is normal. The definition is rather simple using the spectral decomposition (Theorem 22) of A. If $A = \sum_i \lambda_i P_i$, we just define $f(A) = \sum_i f(\lambda_i)P_i$.

¶88. Jordan Form. The ultimate canonical form under general similarity transformation is the Jordan form.

¶89. Hilbert Space. Let S be a complex vector space (or linear space), endowed with an inner product $\langle x, y\rangle \in \mathbb{C}$ such that for all $x, y, z \in S$ and $a \in \mathbb{C}$,

- (linearity) $\langle x, y + z\rangle = \langle x, y\rangle + \langle x, z\rangle$
- (homogeneity) $\langle x, ay\rangle = a\langle x, y\rangle$, $a \in CC$.
- (skew symmetry) $\langle x, y\rangle = \overline{\langle y, x\rangle}$
- (positivity) $\langle x, x\rangle$ is real and $\langle x, x\rangle \geq 0$ with equality iff $x = 0$

There is an asymmetry in the two arguments of the inner product: it follows from the above axioms that

$$\langle ax, y\rangle = \overline{\langle y, ax\rangle} = \overline{a}\overline{\langle y, x\rangle} = \overline{a}\langle x, y\rangle.$$

How does such inner products arise? Suppose H is a matrix and we define $\langle x, y\rangle$ to be x^*Hy. Linearity and homogeneity are obvious: $x^*H(y + z) = (x^*Hy) + (x^*Hz)$ and $x^*H(ay) = a(x^*Hy)$. If H is Hermitian then skew symmetry holds: $\overline{\langle y, x\rangle} = (y^*Hx)^* = x^*H^*y = x^*Hy = \langle x, y\rangle$. It follows from theorem 23 that if H is positive definite then positivity also holds. The simplest case is to choose H to be the identity matrix I.

The **norm** $\|x\|$ of $x \in S$ is defined to be $\sqrt{\langle x, x\rangle}$. An infinite sequence $\overline{x} = (x_1, x_2, \ldots, x_k, \ldots)$ is **Cauchy** if for every $\varepsilon > 0$ there is some k such that for all $i, j \geq k$, $\|x_i - x_j\| < \varepsilon$. The limit of $\overline{x}$ is the element $x_0 \in S$ such that $\|x_k - x_0\| \to 0$ as $k \to \infty$ (clearly, x_0 is unique). The space S is **complete** if every Cauchy sequence has a limit. A **Hilbert space** is a complex vector space S with an inner product and norm as defined, and which is complete. The literature sometimes require Hilbert space to be infinite dimensional; for our limited purposes, we will actually assume S is finite dimensional. The infinite dimensional setting is somewhat more complicated. For instance, in the finite dimensional setting, if an operator A satisfies $A^*A = I$ (i.e., it is unitary), then $AA^* = A^*A$. This property may fail in the infinite dimensional setting.

¶90. Duality. For each $x \in S$ we obtain a linear function $f_x : S \to \mathbb{C}$ where $f_x(y) = \langle x|y\rangle$:

$$f_x(cy) = cf_x(y), \qquad f_x(y + z) = f_x(y) + f_x(z).$$

The function f_x is also continuous where the underlying topology on S is given by the metric $d(x, y) = \|x - y\|$. Conversely, if $\phi : S \to \mathbb{C}$ is linear and continuous, there exists $y \in S$ such that $\phi = f_y$. This duality between

elements of S and continuous linear functions on S gives rise to the $|x\rangle$ (this is just x) and $\langle y|$ notation (this is f_y). Moreover, $\langle y||x\rangle = f_y(x) = \langle y, x\rangle$.

Exercises

Exercise A.1: Show that $\det(AB) = \det(A)\det(B)$. ♢

Exercise A.2: Give the spectral decomposition of $\begin{bmatrix} 4 & 3 \\ 3 & 4 \end{bmatrix}$. ♢

Exercise A.3: Show if $AB = I$ then $BA = I$ and B is unique. HINT: let the (i,j)**-cofactor** of A be $(-1)^{i+j}$ times the determinant of the matrix A with the ith row and jth column deleted. Consider the matrix C whose (i,j)th entry is the (j,i)-cofactor. How close is C to the inverse of A? Show that $AC = CA$. ♢

Exercise A.4: Show a positive definite matrix that is not Hermitian. ♢

Exercise A.5: Let C_{ij} denote the (i,j)-cofactor of A, defined to be $(-1)^{i+j}$ times the determinant of A after the i-th row and j-th column is deleted. Assume the following fact: for all i, $\det(A) = \sum_{j=1}^{n} a_{ij}C_{ij}$ where $a_{ij} = (A)_{ij}$.
(i) Prove that if $\det(A) \neq 0$ then there exists B such that $AB = BA = I$. HINT: Consider the "adjoint" $adj(A)$ of a matrix A where the (i,j)-th entry of $adj(A)$ is the (j,i)-cofactor C_{ji} (note the transposed subscripts).
(ii) Assume $AB = BA = I$. Suppose $AB' = I$ or $B'A = I$ for some other B'. Prove that $B = B'$. [From (i) and (ii), we conclude that inverses are defined and unique whenever $\det A \neq 0$. ♢

Exercise A.6: Consider $n \times n$ matrices with complex entries which are either orthogonal or unitary.
(i) If $n = 1$, what do these matrices look like?
(ii) If $n = 2$, what do these matrices look like? ♢

Exercise A.7: Let $\lambda_1, \ldots, \lambda_k$ be distinct eigenvalues of A and for $i = 1, \ldots, k$, B_i is a set of linearly independent vectors of the invariant subspace E_{λ_i}. Then the set $B = \cup_{i=1}^{k} B_i$ is linearly independent. ♢

Exercise A.8: (SVD) The singular value decomposition (SVD) of an $m \times n$ matrix A is $A = U\Sigma V$ where $U \in \mathbb{C}^{m\times m}$ and $V \in \mathbb{C}^{n\times n}$ are both unitary, and Σ is diagonal. The diagonal entries of Σ are called the **singular values** of A. Show that every A has a SVD. Further, up to complex signs and ordering of the singular values in Σ, the columns of U and rows of V are unique. ♢

Exercise A.9: A polynomial $q(z)$ is a **minimal polynomial** for a matrix A if $q(A) = 0$ and q has minimal degree.
(i) Show that $q(z)$ divides the characteristic polynomial of A.
(ii) Characterize the matrices $A \in \mathbb{C}^{n\times n}$ such that the minimal polynomial of A is z^n ♢

End Exercises

References

[1] A. Barenco, C. Bennett, R. Cleve, D. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52(5):3457–3467, 1995.

[2] D. Beckman, A. N. Chari, S. Devabhakturi, and J. Preskill. Efficient networks for quantum factoring. *Phy. Rev. A*, 54(2):1034–1063, 1996.

[3] P. Benioff. The computer as a physical system: A macroscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *J. Stat. Phys.*, 22(5):563–590, 1980.

[4] P. Benioff. Quantum mechanical Hamiltonian models of Turing machines. *J. Statist. Phys.*, 29:515–546, 1982.

[5] P. Benioff. Quantum mechanical Hamiltonian models of Turing machines that dissipate no energy. *Phys. Review Letters*, 48:1581–1585, 1982.

[6] P. Benioff. Quantum ballistic evolution in quantum mechanics: Applications to quantum computers. *Phy. Rev. A*, 54(2):1106–1123, 1996.

[7] C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Develop.*, 17:525ff, 1973.

[8] C. H. Bennett. Notes on the history of reversible computation. *IBM J. Research and Develop.*, 32:16–23, 1988.

[9] C. H. Bennett. Time/space trade-offs for reversible computation. *SIAM J. Computing*, 18:766–776, 1989.

[10] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Physical Review Letters*, 70:1895–1899, 1993.

[11] C. H. Bennett and P. W. Shor. Quantum information theory. *IEEE Trans. Info. Theory*, 44:2724–2742, 1998.

[12] E. Bernstein and U. Vazirani. Quantum complexity theory. *Proc. ACM Symposium on Theory of Computing*, 25:11–20, 1993.

[13] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20):4091ff, 1995.

[14] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proc. Royal Soc. London, A*, 454:339–354, 1998.

[15] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proc. Royal Soc. London, A*, 400(97–117), 1985.

[16] D. Deutsch. Quantum computational networks. *Proc. Royal Soc. London, A*, 439:553–558, 1992.

[17] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proc. Royal Soc. London, A*, 439(553–558), 1992.

[18] R. P. Feynman. Quantum mechanical computers. *Foundations of Physics*, 16(6):507–531, 1986. Reprinted from Opt.New Vol.11, 11(1985).

[19] E. Fredkin and T. Toffoli. Conservative logic. *Int. J. Theor. Phys.*, 21:219–253, 1982.

[20] Y. Lecerf. Machines de Turing réversibles. récursive insolubilté en $n \in \mathbb{N}$ de l'équation $u = \theta^n u$, oú θ est un isomorphisme de codes. *C. R. Acad. Fran caise Sci.*, 257:2597–2600, 1963.

[21] C. Monroe, D. Meekhof, B. King, W. Itano, and D.J.Wineland. Demonstration of a universal quantum logic gate. *Physical Review Letters*, 75:4714ff, 1995.

[22] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[23] M. Raizen, J. Gilligan, J. Bengquist, W. Itano, and D. Wineland. Ionic crystals in a linear paul trap. *Phy. Rev. A*, 45:6493ff, 1992.

[24] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Computing*, 26(5):1484–1509, 1997.

[25] T. Toffoli. Reversible computing. In J. de Bakker and J. van Leeuwen, editors, *Proc. 7th Int. Colloquium on Automata, Languages and Programming*, pages 632–644, New York, 1980. Springer. Lecture Notes in Computer Science, vol.84.

Lecture XXIII
THEORY OF REAL COMPUTATION

The "computable" numbers may be described briefly as the real numbers whose expressions as a decimal are calculable by finite means. Although the subject of this paper is ostensibly the computable numbers, it is almost equally easy to define and investigate computable functions of an integral variable or a real or computable variable, computable predicates, and so forth. The fundamental problems involved are, however, the same in each case, and I have chosen the computable numbers for explicit treatment as involving the least cumbrous technique. I hope shortly to give an account of the relations of the computable numbers, functions, and so forth to one another. This will include a development of the theory of functions of a real variable expressed in terms of computable numbers. According to my definition, a number is computable if its decimal can be written down by a machine.

—Alan Turing (1936)

April 13, 2009

¶1. REMARK: Rewrite this chapter based on my paper "Theory of Real Computation according to EGC". This requires generalization to partial functions on nominal domains, introducing explicit set theory, explicit algebraic structures, etc.

Until now, we have treated computation over a discrete, countable domain such as natural numbers ($\mathbb{N}$) or finite strings (Σ^*). We now treat computations over an uncountable domain $\mathbb{R}$. Although computation over $\mathbb{R}$ is one of the earliest motivations for computability theory, it has remained relatively underdeveloped. Nevertheless, it is an extremely important aspect of computability and complexity since it forms a large part of numerical analysis and scientific computation. For some current views on this, see Smale [5], Blum [1], and Braverman and Cook [3].

While theory of discrete computability has a widely accepted foundation, the foundation for "continuous" computability over uncountable domains has a less settled status. We shall return to these questions at the end of this chapter. For the main part, we shall follow the development based on the school of "computable analysis", also known as the Polish School or type-2 Theory of Computability. Our basic references are Ko [4] and Weihrauch [7]. The strong influence of [4] should be evident in these notes.

§1. Introduction

¶2. Turing's paper [6] in 1936 is usually noted for its introduction of the computing device that now bear his name. But a major part of his paper addresses the computability of real numbers and their functions. But subsequent development of computability focused mainly on computations over countable domains such as natural numbers or finite strings. Turing promised in the paper to further develp the computability of real numbers, although it was not forthcoming. Grzegorczyk (1955) and Lacombe (1955) introduced the model of real computation that is today spearheaded by the "TTE school" [7]. The starting point of this approach is the representations of real numbers as Cauchy sequences. We can directly manipulate the infinite Cauchy sequence, or we can use an oracle view of such sequences [4]. Another approach, the "Russian School" of Ceitin (1959), focus on representing real numbers as finite objects (programs) that are intepreted by universal machines. In recent years, Smale and co-workers have promoted the algebraic model as the correct model for real computation. The algebraic model is the de facto model used in theoretical algorithms. For a survey of these and related approaches, see [7, Chap. 9].

¶3. The foundation of the real number system $\mathbb{R}$ was clarified by the analysts of the 19th century, and comes down to us in several equivalent forms:

- (i) The Dedekind Approach: real numbers are defined as (Dedekind) **cuts**, defined to be any proper non-empty subset C of $\mathbb{Q}$ with the properties that C has no maximum value, and if $p \in C$ and $q < p$ then $q \in C$. For instance, if $p \in \mathbb{Q}$ then the set of all rational numbers less than p is a cut which we denote by p^*. We then identify real numbers with these cuts.
- (ii) The Cauchy Approach: This is the idea of real numbers as a **Cauchy sequence**, i.e., an infinite sequence
$$\mathbf{a} = (a_0, a_1, a_2, \ldots) \tag{1}$$
of rational numbers with the property that for all $\varepsilon > 0$, there exists $k = k(\varepsilon)$ such that for all $m > n > k$, $|a_m - a_n| \leq \varepsilon$. We call a_i the i**th convergent** of the Cauchy sequence. It is easy to see that each Cauchy sequence corresponds to a unique real number $x \in \mathbb{R}$. To compute with such sequences, we need some additional assumptions about the rate of convergence: we say (1) is **rapidly converging** if $|a_n - x| \leq 2^{-n}$ for all n. In the following, *all Cauchy sequences are rapidly converging.*
- (iii) Binary Expansion Approach: Perhaps the view of real numbers is as an ω-sequence of bits $(0, 1)$, preceded by a sign $(+, -)$, and containing a single binary point (.) somewhere in this sequence. E.g., -3 is represented by -11.0^ω, and $1/3$ is represented by
$$+.01010101\cdots = +.(01)^\omega. \tag{2}$$
We assume the reader is familiar with such representations, but recall here some basic properties. Leading zeros up to the first non-zero symbol (either 1 or binary point) are insignificant. A representation with no such leading zeros is said to be **normalized**. Thus, $1/3$ has infinitely many non-normalized representations, given by the regular expression $+0^+.(01)^\omega$. Two normalized sequences represent the same real number iff they have the form $w01^\omega$ and $w10^\omega$ respectively.

Among these 3 ways to representation real numbers, it turns out that Cauchy sequences is the most appropriate. Turing's original paper explicitly suggested the use of decimal expansion, which is just a variant of binary expansion. But such expansions turned out to be rather limiting, as shown in the Exercise.

¶4. Dyadic Numbers. The Dedekind Approach and the Cauchy Approach both rely on the rational numbers $\mathbb{Q}$. The main property we need is that $\mathbb{Q}$ is dense in $\mathbb{R}$. In the following development, we prefer to replace $\mathbb{Q}$ by the **dyadic rationals** (or **dyadic numbers**),

$$\mathbb{D} := \{m2^n : m, n \in \mathbb{Z}\}.$$

Alternatively, we could define $\mathbb{D}$ as the ring $\mathbb{Z}[\frac{1}{2}]$. Each $d \in \mathbb{D}$ can be represented as a finite sequence

$$\pm b_0 b_1 \cdots b_i . b_{i+1} \cdots b_n = \sum_{j=0}^{n} b_j 2^{i-j},$$

with a binary point between b_i and b_{i+1}. This is just the finite version of Binary Expansion of real numbers above. This set $\mathbb{D}$ is countable and dense in the reals. It has nicer complexity properties than $\mathbb{Q}$ in practical algorithms.

We now (re)define Dedekind cuts as non-empty, proper subsets of $\mathbb{D}$ that has no maximum value, and such that if $p \in C$ and $q < p$ then $q \in C$. Similarly, Cauchy sequences is (re)defined as in (1), but with a_i's in $\mathbb{D}$.

¶5. Sets such as $\mathbb{Q}$ and $\mathbb{D}$ will be called **base reals** because they will be our computational surrogate for approaching real numbers and real functions. In practice, $\mathbb{D}$ will be the preferred set of base reals because of efficient algorithms in this setting. We introduce several notations related to $\mathbb{D}$:

- $\mathbb{D}_n := \{m2^{-n} : m \in \mathbb{Z}\}$ is the set of all dyadic rationals with at most n bits *after* the binary point. For instance, if $n < 0$, then $\mathbb{D}_n$ comprise those integers that are divisible of $2^{-n} > 1$.
- $\mathbb{D}[a, b] := \mathbb{D}[a, b]$ is the restriction of $\mathbb{D}$ to an interval $[a, b] \subseteq \mathbb{R}$. Combining these two notations, we could write $\mathbb{D}_n[a, b]$ for $\mathbb{D}_n \cap [a, b]$.
- We generalize the standard **floor** and **ceiling** functions of number theory. For any $x \in \mathbb{R}$, let $\lfloor x \rfloor_n$ denote the largest dyadic in $\mathbb{D}_n$ that is $\leq x$. Similarly, $\lceil x \rceil_n$ is the smallest dyadic in $\mathbb{D}_n$ that is $\geq y$. Thus $\lfloor x \rfloor_n \leq x \leq \lceil x \rceil_n$. The usual floor and ceiling functions are $\lfloor x \rfloor = \lfloor x \rfloor_0$ and $\lceil x \rceil_0 = \lceil x \rceil$.

¶6. Basic Operations on Dyadics. We gather some well-known facts about operations on dyadic numbers. Because of our emphasis on bit-sizes, it is useful to approximate $\lg |x| = \log_2 |x|$ where $x \in \mathbb{R}$. Define

$$c\ell g\, \mathbb{R} \to \mathbb{Z}$$

where $c\ell g\, x) = \lceil \lg |x| \rceil$ if $x \neq 0$, and $c\ell g\, 0) = \uparrow$. Similarly, $f\ell g\, \mathbb{R} \to \mathbb{Z}$ is defined using floor instead of ceiling.

Lemma 1. *Assume that dyadic numbers are represented by their binary expansions.*
(i) The restrictions of $c\ell g\, x)$ and $f\ell g\, x)$ to $x \in \mathbb{D}$ can be computed in polynomial time.
(ii) The restrictions of $\lceil x \rceil$ and $\lfloor x \rfloor$ functions to $x \in \mathbb{D}$ can be computed in polynomial time.
(iii) Comparision of two dyadic numbers can be computed in linear time.
(iv) The ring operations $\pm, \times$ and division by powers of 2 on $\mathbb{D}$ is computable in polynomial time.

Proof. Assume $d \neq 0$, and

$$d = \pm b_n b_{n-1} \ldots b_1 b_0 . b_{-1} \ldots b_{-m}$$

for some $m, n \geq 0$ and each $b_i \in \{0, 1\}$.
(i) Let $k \in \{n, \ldots, -m\}$ be the largest index such that $b_k = 1$, and $\ell \in \{n, \ldots, -m\}$ be the smallest index such that $b_\ell = 1$. Then $f\ell g\, d) = k$ and $c\ell g\, d) = -k - \delta(k \neq \ell)$ where $\delta(P) = 1$ if the predicate P is true and $\delta(P) = 0$ else.
(ii) and (iii): This is clear.
(iv) Division by powers of two amounts to shifting the binary point. Addition is easy. Multiplication can be done by the high-school algorithm in quadratic time. **Q.E.D.**

REMARK: An alternative to binary expansions is to represent the dyadic number $m2^n$ $(m, n \in \mathbb{Z})$ by the pair (m, n), where m, n are represented in binary. The size of this representation is $\lg(2|mn|)$. We may call this the **floating point** representation. The complexity bounds under this representation is unchanged as far as the operations in this lemma is concerned. But for other operations, the two complexity bounds can be rather different.

¶7. Partial Functions. Following [7], we distinguish partial from total functions in our notations. Suppose R, S are sets. We write

$$f :\subseteq R \to S$$

to indicate that f a partial function, with **nominal domain** R and range S. So for all $x \in R$, $f(x)$ is either an element of S or it is undefined. We write $f(x) =\uparrow$ if it is undefined, and $f(x) =\downarrow$ when $f(x) \in S$. Call $\texttt{domain}(f) := \{x \in R : f(x) =\downarrow\}$ the **proper domain** of f. Also $\texttt{range}(f) := \{f(x) : x \in \texttt{domain}(f)\}$ is the **proper range** of f. We say f is **onto** or **surjective** iff its nominal and proper ranges agree. In case the nominal domain and proper domain of f agree, we call f a **total function** and indicate this in the standard way, $f : R \to S$.

We remark that the distinction between nominal domain and proper domain has import when we discuss computation of partial functions: an algorithm to compute the given function is expected to receive inputs from its nominal domain, and the algorithm should be able to recognize improper inputs. Thus, we will see an example of a computable function which becomes uncomputable simply by enlarging the nominal domain.

¶8. Computing real functions with Type-2 Turing Machines. How do we compute a real function $f : \mathbb{R} \to \mathbb{R}$? It is understood that we intend to use one of the above representations of real numbers. But there are still some other decisions.

We can take any standard Turing machine M that, besides the standard work tape(s), has two extra tapes for input and output. The dedicated input tape is read-only, and store the input real number x, represented by an infinite string in one of the three formats (Dedekind, Cauchy, Binary). This input string can be *any* representation of x. The dedicated output tape is write-only, and will "compute in the limit" *some* representation of $f(x)$ in the chosen format (Dedekind, Cauchy, Binary). Computing in the limit means, for any n, if we wait long enough, the output tape of M will contain the first n symbols of the representation of $f(x)$. Note that M is necessarily non-halting.

A Turing machine that computes in this manner is called a **type-2 Turing Machine**. Thus, f is **type-2 computable** iff there is a type-2 Turing machine that computes f in the above sense. See [7] for this approach. We will shortly introduce an alternative machine formalism for computing real functions, using the concept of "oracle Turing machines".

Exercises

Exercise 1.1: Show that (2) represents 1/3. NOTE: It is not considered a proof if you show that multiplying the string of (2) by the binary expansion of 3 gives the binary expansion of 1. This procedure assumes that the multiplication algorithm for binary numbers is correct. ♢

Exercise 1.2: Show that if we assume the binary expansion of real numbers, then the function $x \mapsto 3x$ is not computable by a type-2 Turing machine. HINT: consider $x = 1/3$. ♢

Exercise 1.3: Our definition of type-2 machine forces the the Turing machine to compute forever even when both the input output could have some finite representation.

(i) Let us extend these representations by giving them some finite means to indicate an infinite sequence. **Finite Cauchy Representation**: the sequence (1) can be finite. It just means that the last entry is repeated forever. **Finite Binary Representation**: we allow a finite string provided it has the form $x(y)$ where x denotes an element of $\mathbb{D}$ and $y \in \{0,1\}^+$. This is just a finite way of writing the binary expansion xy^ω. E.g., 1/3 has the finite representation $+.(01)$, and $-1/2$ has the finite representations $-.111(1)$ or $-.(1)$. What kinds of real numbers hve finite representations?

(ii) An extended type-2 Turing machine now accepts the usual representation as well as the finite variant. Its output can be the usual infinite representation, or if it halts, it produces the finite variant. Show that extended type-2 Turing machines do not compute any larger class of real functions.

(iii) Characterize the real functions that are computed by extended type-2 Turing machines that always halt. ♢

End Exercises

§2. Representation of Real Numbers

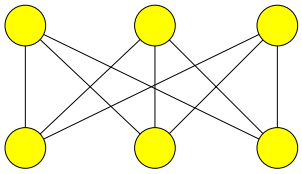

Figure 1: $K_{3,3}$

¶9. Representation of graphs The concept of representations is central to computing. This is the starting point of Turing's analysis of computability. We illustrate these ideas using the the concept of finite graphs (both the directed and undirected varieties). Graphs are relatively abstract objects, and to manipulate them in computers, we need concrete representations. For instance, consider the "utilities graph" $K_{3,3}$ illustrated in Figure 1.

The main point about $K_{3,3}$ is that the vertices of $K_{3,3}$ are unlabeled (anonymous). We shall introduce labeled graphs shortly. A standard representation of graphs is via Boolean matrices. The following Boolean matrices M_1 and M_2 are both representations of $K_{3,3}$:

$$M_1 = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad M_2 = \begin{bmatrix} 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix}.$$

Let $\mathcal{G}$ and $\mathcal{M}$ be the set of (unlabeled) undirected graphs and Boolean matrices, respectively. We view $K_{3,3}$ as an element of $\mathcal{G}$. The representation of undirected graphs by Boolean matrices can be formalized as a partial injective function,

$$\rho_{GM} :\subseteq \mathcal{M} \to \mathcal{G}.$$

It is a partial function because not all Boolean matrices represent an element of $\mathcal{G}$, only the symmetric ones do. It is surjective because every element of $\mathcal{G}$ has at least one representation, but an n-vertex graph has at most $n!$ representations.

We are not done yet: Boolean matrices are still too abstract for computers (Turing machines). We need a representation by strings. But here, there is the obvious representation of Boolean matrices by strings:

$$\rho_{MS} :\subseteq \{0,1\}^* \to \mathcal{M}$$

where each $M \in \mathcal{M}$ is represented by listing its rows in order. For instance,

$$\rho_{MS}(000111000111000111111000111000111000) = M_1, \qquad \rho_{MS}(010101101010010101101010010101101010) = M_2.$$

Again, ρ_{MS} is a partial surjective function, but now it happens to also be injective (every Boolean matrix has a unique representation). By composing these two representations, we get a representation of graphs that is suitable for computing:

$$\rho_{GS} :\subseteq \{0,1\}^* \to \mathcal{G}$$

where $\rho_{GS}(w) = \rho_{GM}(\rho_{MS}(w))$. In other words, *the representation of a representation of an object G is a representation of G.*

Another well-known representation of graphs is by a pair (V, E) where V is a finite set and $E \subseteq V^2$. We call (V, E) a a **labeled directed graph**. It is a "labeled" graph where the identity of each node is relevant. Two labeled directed graphs (V, E) and (V', E') are isomorphic if there is renaming of the elements of V that turns E to E'. There are also standard representations of undirected graphs $\mathcal{G}$ as such labeled directed graphs. To ultimately get a representation by strings, we need to compose this with some representation for finite sets, and also representations for the vertices (elements of V). This is developed in the Exercise.

In normal usage, we move between the concepts of labeled graphs and unlabeled graphs without much fanfare (often there is not even a name to distinguish them). Any concrete representation has invariably a great deal of superfluous details, but in practice we learned to automatically filter out these details. For instance, even the sets requires a representation which is some arbitrary listing of its elements (perhaps with duplicates!). But we learned quite early to filter out duplicates in set representations, and how to re-order elements (by sorting) so that we can quickly check equivalent representations of sets.

¶10. Representations. All computations over abstract mathematical domains must ultimately be reduced to manipulation of more concrete representations. We saw this with finite graphs above. In particular, computating over real numbers require such representations. Representations[1] of real numbers and their properties are treated in depth by Weihrauch [7].

Let S and T be sets. A **representation** for S (with **representig set** T) is a partial, onto function of the form

$$\rho :\subseteq T \to S. \tag{3}$$

If $\rho(x) =\downarrow$, we say x is **proper** (or ρ-proper), otherwise **improper**. We may call a proper x a **representating element** (or ρ-representation) of $\rho(x) \in S$. Each $s \in S$ has at least one, but possibly many more, representations. If $\rho(x) = \rho(y)$ then we say that x, y are **ρ-equivalent**.

There is an important special case already used in Chapter 0: when the representing set T is Σ^* for some alphabet Σ, we call ρ a **notation**. Without loss of generality, we assume $\Sigma = \{0, 1\}$ unless otherwise noted.

¶11. Discussion. The set S is normally more abstract than its representing set T. We deduce properties of $s \in S$ by manipulating its representating elements. If a representation is meant to used for computation, then it must ultimately be a notation. Nevertheless it is useful to introduce intermediate representations. We consider $T = \mathbb{N}$ or $T = \Sigma^*$ to be adequate for computation.

We emphasize that the function ρ itself is not an object of computation. In particular, the issue of "effective" ρ does not arise. Rather, ρ lies in the realm of mathematical analysis; they are used in discussion and in proving properties of algorithms that use this representation. We will discuss effectivity concepts related to ρ later (under "decidable sets").

¶12. Convention for representations. Invariably, there is a gap between abstract mathematical objects and their representations. The example of graphs illustrate this gap. Although we must use representations in computation, at other times, we prefer to talk directly about the ideal or abstract objects. To make the gap between objects and representations as nonobtrusive as possible, we use a convention.

> Let S be a set with a representation system ρ. In most discussions, ρ may be assumed to be held fixed, and we could effectively suppress ρ. This can be achieved by writing
>
> $$\overline{x}, \overline{y}, \overline{z}, etc. \tag{4}$$
>
> for the **ρ-representations** of $x, y, z, \ldots \in S$.
> Thus "$\overline{x}$" denotes an arbitrary but fixed element of $\rho^{-1}(x)$. Within the scope of discussion, the choice of this element will be held fixed.

This convention is similar to the asymptotic notations such as $O(f)$ or $\Omega(f)$ in that it is a way to hide details (like ρ) that could be obtrusive.

¶13. Higher type classes and objects. We want to view natural numbers as type-0 objects, real numbers as type-1 objects, and real functions as type-2 objects. This explains the Weihrauch's terminology of type-2 Turing machines, as they compute real functions. Here we give a simplified account of types.

For sets R and S, let the set of all partial functions from R to S be denoted

$$[R \Rightarrow S]. \tag{5}$$

Let us define a **type-0 class** to be a countable set S. Note that S can be finite or denumerable. If R is a type-i class and S is a type-j class, then $R \times S$ is a type-$\max\{i, j\}$ class and $[R \Rightarrow S]$ is **type-k class** where $k = 1 + \max\{i, j\}$.

If the representation $\overline{x}$ of an object x comes from a type-i class, we call x a **type-i object**.

Natural numbers, dyadic numbers, finite graphs, etc, can be viewed as type-0 objects. Rational numbers are also type-0 objects because they are represented by a pair of type-0 objects. Real numbers are represented by elements of $[\mathbb{N} \Rightarrow \mathbb{D}]$, so they are type-1 objects. Real functions $f :\subseteq \mathbb{R} \to \mathbb{R}$ are represented by elements of

$$[[\mathbb{N} \Rightarrow \mathbb{D}] \Rightarrow [\mathbb{N} \Rightarrow \mathbb{D}]]$$

and so are type-2 objects.

[1] The "representations" in Weihrauch [7] are representations in our sense, but with representing set restricted to Σ^ω. Our notation terminology follows Weihrauch. By a "naming system", Weihrauch refers to either a "representation" or a notation.

¶14. Induced Representations. The representation of higher-type objects is normally induced from representations of the underlying type-0 objects. Let $\rho :\subseteq T \to R$ be a notation of R. Consider a k-ary function $f :\subseteq R^k \to R$. We call $\overline{f} :\subseteq (T)^k \to T$ a **(ρ-induced) representation** of f if for all $(\overline{x}_1, \ldots, \overline{x}_k) \in T^k$,

$$\rho(\overline{f}(\overline{x}_1, \ldots, \overline{x}_k)) = f(\rho(\overline{x}_1), \ldots, \rho(\overline{x}_k)). \tag{6}$$

Equation (6) is interpreted in the strong sense that the LHS is undefined iff the RHS is undefined.

We say f is **computable** if there exists a notation ρ such that the ρ-induced notation $\overline{f}$ is computable.

¶15. We now build up the representation of real numbers and their operations.

- **Dyadic numbers and ring operations.** We had described a notation system for the dyadic numbers $\mathbb{D}$. This is the function
$$\rho :\subseteq \Sigma^* \to \mathbb{D}$$
where $\Sigma = \{0, 1, \bullet, -, +\}$ and the proper strings have the form $\pm b_0 b_1 \cdots b_i \bullet b_{i+1} \cdots b_n$, which is a notation for the number $\pm \sum_{j=0}^{n} b_j 2^{i-j}$. E.g., $\rho(-01 \bullet 01) = -1.25$. Since $\mathbb{N} \subseteq \mathbb{Z} \subseteq \mathbb{D}$, this notation for dyadics is an extension of standard notation for natural numbers and integers.

  Consider the representation of ring operations over $\mathbb{D}$. Relative to an induced ρ, the multiplication function $\times : \mathbb{D}^2 \to \mathbb{D}$ has an induced representation
$$Mul : (\Sigma^*)^2 \to \Sigma^*$$
such that $Mul(\overline{m}, \overline{n}) = \overline{m \times n}$ for all $m, n \in \mathbb{Z}$. Similarly, the other ring operations has induced representations Add, Sub and Sub with the properties $Add(\overline{m}, \overline{n}) = \overline{m+n}$ and $Sub(\overline{m}, \overline{n}) = \overline{m-n}$. There are well-known algorithms for implementing Mul, Add and Sub.

- **Real Numbers.** A **Cauchy representation** of real number $x \subseteq \mathbb{R}$ is a total function $\overline{x} : \mathbb{N} \to \mathbb{D}$ such that $|\overline{x}(n) - x| \leq 2^{-n}$ for all $n \in \mathbb{N}$. Let
$$\rho :\subseteq [\mathbb{N} \Rightarrow \mathbb{D}] \to \mathbb{R}$$
where ρ maps each Cauchy sequence to the corresponding real numbers. Thus ρ is a representing system for real numbers.

¶16. Standard Cauchy function. Every real number x can be written as

$$n + 0.b_1 b_2 \cdots$$

where $n \in \mathbb{Z}$ and $b_i \in \{0, 1\}$. The b_i's are uniquely determined by x when $x \notin \mathbb{D}$. Otherwise, there are two possibilities (either b_i's are eventually 0 or eventually 1). We obtain uniqueness by requiring that b_i's are eventually 0. Using this unique sequence, we define the **standard Cauchy function** of x to be

$$\beta_x[p] = n + \sum_{i=1}^{p} b_i 2^{-i}.$$

For instance, $-5/3$ is written $-2 + 0.01010101 \cdots$. This function has nice monotonicity properties:

Lemma 2. *Let $x \in \mathbb{R}$ and $p \in \mathbb{N}$.*
(i) $\beta_x[p] \leq \beta_x[p+1] \leq x$.
(ii) $x - \beta_x[p] < 2^{-p}$.
(iii) If $y \in \mathbb{R}$ and $|y - \beta_x[p]| \leq 2^{-p}$, then for all $n \leq p$, we also have $|y - \beta_x[n]| \leq 2^{-n}$. In particular, there is a representation $\overline{y}$ such $\overline{y}[n] = \beta_x[n]$ for all $n \leq p$.

Proof. We only prove (iii). We use induction on $p - n$. The result is true for $p - n = 0$. If $p - n \geq 1$, then we have either $\beta_x[n] = \beta_x[n+1]$ or $\beta_x[n] = \beta_x[n+1] - 2^{-n-1}$. In the former case, $|y - \beta_x[n]| = |y - \beta_x[n+1]| \leq 2^{-n-1} < 2^{-n}$, so the result holds. Otherwise, $|y - \beta_x[n]| = |y - \beta_x[n+1] + 2^{-n-1}| \leq |y - \beta_x[n+1| + 2^{-n-1} \leq 2^{-n}$ (by induction hypothesis). **Q.E.D.**

¶17. Gödel numbers of real numbers. If x is a type-1 object, it is represented as a function $\overline{x} : S \to T$. This is generally an infinite object. Now, if $\overline{x}$ can be computed by the ith Turing machine M_i, then we say that i (the Gödel number of M_i) is a representation of $\overline{x}$. Extending the convention of ¶12, for any real number x, let $\overline{\overline{x}}$ refer to the Gödel number of a Turing machine that computes any representation $\overline{x}$ of x.

¶18. Convention for approximation of real numbers. The following convention facilitates the manipulation of expressions involving approximate numbers.

> The expression "$x \pm \delta$" refers to a real value of the form $x + \theta\delta$, for some $|\theta| \leq 1$. Thus θ is an anonymous variable in this notation (like the big-Oh notation).
> In general, any occurrence of the symbol "$\pm$" in an approximate expression calls for a textual substitution by "$+\theta$" for some anonymous variable θ satisfying $|\theta| \leq 1$. Different occurences of "$\pm$" may call for different θ's depending on the context.

Thus, we write "$\widetilde{x} = x \pm 2^{-p}$" to indicate that $\widetilde{x} = x + \theta 2^{-p}$ for some $\theta \in [-1, 1]$. We call $\widetilde{x}$ a **p-bit absolute approximation** of x in this case.

Examples of manipulations using this notation: We have $x = y \pm 2^{-p}$ iff $y = x \pm 2^{-p}$. Also, $x = (y \pm 2^{-p}) \pm 2^{-p} = y \pm 2^{-p+1}$.

Exercises

Exercise 2.1: For each of the following sets S, provide a string representation. Use Σ^* for the representing set, where you can choose the alphabet Σ. For complex representations, we want you to use composition of simpler representations whenever possible. E.g., the representation in each part is supposed to use the representation constructed in the earlier part.
(i) S is the set of finite subsets of $\{0, 1\}^*$.
(ii) S is the set of labeled directed graphs whose vertices are elements of $\{0, 1\}^*$. We want you to represent each graph as a pair (V, E) as described in the text.
(iii) S is the set of (unlabeled) directed graphs.
(iv) S is the set of (unlabeled) binary trees. Note that nodes in a binary tree distinguishes between its left and right children. Even when it has one child, this child is either a left or a right child. ♢

Exercise 2.2: For each of the representations in the previous exercise, describe any reasonable algorithm for checking equivalence. Polynomial time complexity may be hard to achieve in some cases. ♢

Exercise 2.3: Let us explore the notion of types. Write $\tau(R) = i$ if R is a type-i class.
(a) What is the cardinality of the sets $[\mathbb{N} \Rightarrow \{0, 1\}]$ and $[\{0, 1\} \Rightarrow \mathbb{N}]$? This suggests that if $T = [R \Rightarrow S]$ then $\tau(T) = \max\{\tau(R) + 1, \tau(S)\}$. What are some properties of this definition of type?
(b) What is the cardinality of the sets $[R \times S \Rightarrow T]$, $[R \Rightarrow S \times T]$? What does this suggest about extending the notion of types to sets which are constructed recursively by using the two operations of forming functions $[\cdot \Rightarrow \cdot]$ and Cartesian product $(\cdot \times \cdot)$?
(c) What is the cardinality of $[R \Rightarrow [S \Rightarrow T]]$ and $[[R \Rightarrow S] \Rightarrow T]$? ♢

End Exercises

§3. Computable Real Functions

¶19. We introduce **oracle Turing machines** (OTM). Such machines are alternatives to type-2 Turing machines (¶8) as devices for computing real functions $f :\subseteq \mathbb{R} \to \mathbb{R}$.

Let M denote such an OTM. Intuitively, the input real number x to M is viewed[2] as an "oracle". The oracle is just a Cauchy representation $\overline{x} = (a_0, a_1, \ldots)$ for x. For any given $n \in \mathbb{N}$, the oracle $\overline{x}$ returns the nth convergent a_n. We must provide M with a special tape to query the oracle, and to receive answers. What is the output of this OTM on oracle $\overline{x}$? Well, we want it to compute an oracle for $f(x)$. An oracle for $f(x)$ must take an input $p \in \mathbb{N}$ and output the pth convergent of some Cauchy representation of $f(x)$. Hence, the machine M needs another input tape (the **precision tape**) to hold p. We often write

$$M^{\overline{x}}[p] \tag{7}$$

to denote the output of M, or the entire computation, on $\overline{x}$ and p,

An advantage of OTM over type-2 Turing machines is that we avoid the concept of "computing in the limit". It also allows a more natural definition of complexity. Similar oracles could be defined if x were represented by a binary expansion or a Dedekind cut.

[2] An oracle is a black box which will supply answers to any (valid) query. The temples of ancient Greece often have such oracles. The most famous of these temple oracles is at Delphi.

¶20. Oracle view of real functions. We generalize OTM's to allow any number of oracles. For instance, to compute a function $f :\subseteq \mathbb{R}^k \to \mathbb{R}$, we need an OTM M that takes as input k oracles, $\overline{x}_1, \ldots, \overline{x}_k$ corresponding to an input $(x_1, \ldots, x_k) = \mathbf{x} \in \mathbb{R}^k$. A query to the oracle amounts to placing a pair (i, n) on the oracle tape, and the output is the nth convergent in $\overline{x}_i$.

Let $f : \mathbb{R}^k \to \mathbb{R}$ be partial k-ary real function. An **oracle function** for f is any function

$$\overline{f} :\subseteq \mathbb{R}^k \times \mathbb{N} \to \mathbb{D} \tag{8}$$

such that for all $\mathbf{x} \in \mathbb{R}^k$ and $p \in \mathbb{N}$, $\overline{f}(\mathbf{x}, p)$ is defined iff $\mathbf{x} \in \mathtt{domain}(f)$. In case it is defined, we have $|\overline{f}(\mathbf{x}, p) - f(\mathbf{x})| \leq 2^{-p}$. Alternatively, $\overline{f}(\mathbf{x}, p) = f(\mathbf{x}) \pm 2^{-p}$.

An OTM M can be viewed as computing oracle functions (cf. (7)). For any $\mathbf{x} \in \mathbb{R}^k$, and for all Cauchy representations $\overline{\mathbf{x}}$, and $p \in \mathbb{N}$, the OTM computes $M^{\overline{\mathbf{x}}}[p] = f(\mathbf{x}) \pm 2^{-p}$. The output depends not just on $\mathbf{x}$, but on the specific representation $\overline{\mathbf{x}}$. In case $k = 0$, f is just a real number x and the oracle function $\overline{f}$ is just a Cauchy representation for x. The OTM reduces to an ordinary Turing machine.

¶21. Absolute approximation of real functions. An **(absolute) approximation** $\widetilde{f}$ for f is just the restriction of an oracle function $\overline{f}$ to dyadic numbers. Instead of (8), we have

$$\widetilde{f} :\subseteq \mathbb{D}^k \times \mathbb{N} \to \mathbb{D}$$

such that $(\mathtt{domain}(f) \cap \mathbb{D}^k) \times \mathbb{N} = \mathtt{domain}(\widetilde{f})$, and for all $\mathbf{x} \in \mathbb{D}^k$ and $p \in \mathbb{N}$, we have

$$\widetilde{f}(\mathbf{x}, p) = \begin{cases} f(\mathbf{x}) \pm 2^{-p} & \text{if } \mathbf{x} \in \mathtt{domain}(f) \\ \uparrow & \text{else.} \end{cases}$$

Let $\mathcal{A}_f$ denote the set of all absolute approximations of f.

We drop the qualifier "absolute" if this is clear from context; the reason for this qualification is that we will later introduce "relative approximation". Since the input and output of approximations are dyadic numbers, they can be computed by ordinary Turing machines. We say that f is **(absolutely) approximable** if some $\widetilde{f} \in \mathcal{A}_f$ is computable in the standard sense, by an ordinary Turing machine.

A real number x can be viewed as 0-ary real function; in this case, there is no difference between oracle functions $\overline{x}$ for x or absolute approximations $\widetilde{x}$ of x. We may write $\mathcal{A}_x$ instead of $\mathcal{A}_f$. Thus oracle functions or absolute approximations of real numbers are just another form of rapidly convergent Cauchy sequences.

¶22. Special treatment of precision parameter. The precision parameter p in oracle functions and in absolute approximations has different status from the other arguments. This fact will be seen later, when we discuss complexity of computing such functions. To visually distinguish this parameter from the others, we may write either "$f(\mathbf{x})[p]$" or "$\overline{f}(\mathbf{x}; p)$" instead of $\overline{f}(\mathbf{x}, p)$. Note that "$f(\mathbf{x})[p]$" refers to any p-bit approximation of $f(\mathbf{x})$, and this value depends on the context. In particular, if x is a real number, then we may write "$x[p]$" or "$\overline{x}(p)$" for the pth convergent of $\overline{x}$.

¶23. Oracle machines. Since oracle machines are our official computational model for real functions, we will provide the details in full.

An oracle Turing machine (OTM) is a Turing machine M that has, in addition to the standard work tape(s), also three special tapes, called the **oracle tape**, the **precision tape** and the **output tape**. It also has two special states,

$$q_?, q_! \tag{9}$$

called the **query state** and the **answer state**. We view M as computing a function $\overline{f} : \mathbb{R} \times \mathbb{N} \to \mathbb{R}$. The input is $\overline{x} \in \mathcal{A}_x$, and a precision $p \in \mathbb{N}$ placed on the precision tape. The computation of M begins with the When it enters the query state $q_?$, we assume the oracle tape contains a binary number $\overline{n}$. In the next instant, M will enter the answer state $q_!$, and simultaneously the string $\overline{n}$ is replaced by a representation of the nth convergent $\overline{x}[n]$. Then M continues computing as usual. Finally, M halts and on the output tape should contain $\overline{d}$ where $d \in \mathbb{D}$ such that $|d - f(x)| \leq 2^{-p}$. We shall denote the output d by $M^{\overline{x}}[p]$.

¶24. Computing partial real functions. Let $S \subseteq \mathbb{R}^k$ for some $k \geq 0$. Typically, we have $S = \mathbb{R}^k$ or $S = [a, b] \subseteq \mathbb{R}$. We even allow $k = 0$ (in which case $S = \emptyset$). To compute a partial function

$$f :\subseteq \mathbb{R}^k \to \mathbb{R} \tag{10}$$

we need an OTM that accepts k oracles $\overline{x_i}$ $(i = 1, \ldots, k)$ with the following properties. A query on the oracle tape consists of a pair $\langle n, i\rangle \in \mathbb{N} \times \{1, \ldots, k\}$, and the response will be a representation of $\overline{x_i}[n]$ on the oracle tape. If $\mathbf{x} = (x_1, \ldots, x_k)$, we write $\overline{\mathbf{x}}$ for $(\overline{x_1}, \ldots, \overline{x_k})$.

We say that the OTM M **computes** the partial function f if for all $\mathbf{x} \in S$, representations $\overline{\mathbf{x}}$, and $p \in \mathbb{N}$,

$$M^{\overline{\mathbf{x}}}[p] = \begin{cases} f(\mathbf{x}) \pm 2^{-p} & \text{if } \mathbf{x} \in \mathtt{domain}(f), \\ \uparrow & \text{if } \mathbf{x} \in S \setminus \mathtt{domain}(f). \end{cases} \tag{11}$$

Such an f is said to be **computable**. Moreover, we may define $\mathtt{domain}(M)$ as $\mathtt{domain}(f)$.

In case f is a 0-ary function, f is a real number x, and we call x a **computable real**; this means x has a recursive Cauchy function $\widetilde{x} : \mathbb{N} \to \mathbb{D}$.

¶25. Remarks.
0. In (11), the notation

$$M^{\overline{\mathbf{x}}}[p] = \uparrow$$

means that the OTM goes into an infinite loop, i.e., does not halt. If we omit this requirement, but have no requirements at all when $\mathbf{x} \notin \mathtt{domain}(f)$, then we say that M **conditionally computes** the function f. Equivalently, it means we promise not to feed improper inputs to the OTM.
1. Instead of weakening the notion of computability, we can strengthen it: assume the OTM has two special states,

$$q_\uparrow, \quad q_\downarrow \tag{12}$$

called the **improper** and **proper** states. We require that, for all inputs $\mathbf{x}$, the OTM must halt in the state $q_\uparrow$ if $\mathbf{x} \in \mathtt{domain}(f)$, and otherwise it halts in the state $q_\downarrow$ (with the correct output). In this case, we say the OTM **strongly computes** f.
2. Note that we have define the "computability" of partial functions, not the "partial computability" of functions. We will avoid the terminology of "partially computable", which is commonly used in discrete computability theory.
3. To distinguish between computability and partial computability in discrete computability from real computability, we will use the terms **recursive** and **partially recursive** for computability over a discrete domain.
4. We extend this distinction when complexity is taken into account. E.g., we say **polynomial-time recursive** to refer to discrete functions that can be computed in polynomial-time. This is the functional analogue of the well-known class P comprising languages that are recognized in deterministic polynomial time. Write FP for the class of polynomial-time recursive functions.

¶26. Example: OTM to multiply by three. Let us construct an OTM M_3 to compute the function $f(x) = 3x$. With an oracle $\overline{x}$ and $p \in \mathbb{N}$, M_3 first place $\overline{p+2}$ on the oracle tape and enter $q_?$. Let the response be $d \in \mathbb{D}$. Then M_3 computes $\overline{3d}$ on the output tape and halt, using any algorithm for multiplying binary numbers. To see that M_3 computes $f(x) = 3x$, note that $|d - x| \leq 2^{-p-2}$ and so $|3d - f(x)| \leq 3 \cdot 2^{-p-2} < 2^{-p}$, as desired.

¶27. Example: Division is a computable partial function. Is the partial function $\div : (x, y) \mapsto x/y$ computable? To show this is so, we must make sure that $y = 0$ iff our OTM loops. This is not hard to do: for $p = 0, 1, 2, \cdots$, we query the y-oracle for the pth convergent and check if $|\overline{y}[p]| > 2^{-p}$. If $y = 0$, then the condition $|\overline{y}[p]| > 2^{-p}$ will never hold and we loop. Otherwise, we know that $y \neq 0$, and we proceed to compute $x/y \pm 2^{-p}$.

Here is the procedure: first we find a $k \in \mathbb{N}$ such that $|y| \geq 2^{-k}$ and $|x| \leq 2^k$. In fact, if $|\overline{y}[p]| > 2^{-p}$ then we may choose k greater than $\lg(|\overline{y}[p]| - 2^{-p})$. We now query the oracle for the values $x' = \overline{x}[p + 3k + 3]$ and $y' = \overline{y}[p + 3k + 3]$. Thus $|y'| \geq |y| - 2^{-p-3k-3} > 2^{-k-1}$. Note that

$$\begin{aligned} \left|\frac{x'}{y'} - \frac{x}{y}\right| &= \left|\frac{x'y - xy'}{yy'}\right| \\ &< \frac{2^{k+1}}{y}\,|x'y - xy'| \\ &= \frac{2^{k+1}}{y}\,\left|(x \pm 2^{-p-3k-3})y - x(y \pm 2^{-p-3k-3})\right| \\ &\leq 2^{-p-2k-2}((x/y) + 1) \\ &\leq 2^{-p-2k-2}(2^{2k} + 1) \\ &\leq 2^{-p-1}. \end{aligned}$$

We then compute z' as a $(p+1)$-bit approximation of x'/y' (there are standard algorithms for this). Then $|z' - (x/y)| \leq |z' - (x'/y')| + |(x'/y') - (x/y)| \leq 2^{-p}$.

¶28. Computable reals. We discuss two special cases of computing a real function f. The case when the range of f is a finite set is discussed under "real predicates" below. Here we discuss the case where f is a constant function:

When $k = 0$ in (10), the OTM M that computes f is just computing a real number $x \in \mathbb{R}$. We may regard M is just an ordinary Turing machine (no oracles). For any $p \in \mathbb{N}$, we have $|M[p] - x| \leq 2^{-p}$. We call $x \in \mathbb{R}$ a **computable real number** if such an M exists.

¶29. Uncomputable reals. Any rational number is computable. So is any real algebraic number. To construct an uncomputable real, we use an idea of Specker (1949). For $A \subseteq \mathbb{N}$, define the number $x_A := \sum_{n \in A} 2^{-n}$. We can also write $x_A = \sum_{n=0}^{\infty} \chi_A(n) 2^{-n}$ where $\chi_A : \mathbb{N} \to \{0, 1\}$ is the characteristic function of A.

Lemma 3. *x_A is computable iff A is recursive.*

Proof. If A is recursive, then the standard Cauchy representation $\beta = \beta_{x_A}$ is easily seen to be computable: $\beta[p] = \sum_{i=0}^{p} \chi_A(i) 2^{-i}$.

Conversely, suppose $\overline{x} = \overline{x_A}$ is a computable representation of x_A. If $x_A \in \mathbb{D}$, then clearly A is recursive. So assume otherwise. It is sufficient to show that the standard Cauchy representation β of x_A is computable. This is because $p \in A$ iff $\beta[p] > \beta[p-1]$. To compute $\beta[p]$, we compute the first $n = n(p)$ such that

$$[\overline{x}[n] - 2^{-n}, \overline{x}[n] + 2^{-n}] \cap \mathbb{D}_p = \emptyset. \tag{13}$$

Such an n exists since $x \notin \mathbb{D}$. The predicate (13) is also computable. Then $\beta[p] = \max\{d \in \mathbb{D}_p : d < \overline{x}[n]\}$. **Q.E.D.**

Corollary 4. *If A is nonrecursive, x_A is uncomputable. In particular, if A is the set* DIAG $= \{i \in \mathbb{N} : \phi_i(i) = \downarrow\}$ *then x_A is uncomputable.*

¶30. The field of computable reals.

Lemma 5. *The field operations of $\pm, \times, \div$ are computable functions.*

Proof. The case of $\div$ has been shown in ¶27. Let us consider addition and subtraction: we can define

$$\overline{x \pm y}[n] := \overline{x}[n+1] \pm \overline{y}[n+1].$$

Correctness is clear from the fact that

$$\overline{x+y}[n] = (x \pm 2^{-n-1}) + (y \pm 2^{-n-1}) = (x+y) \pm 2^{-n}.$$

Similarly for subtraction, $\overline{x-y}$.

The remaining operation is multiplication: first compute some $k \in \mathbb{N}$ such that x, y satisfy $|x| + |y| < 2^{k-1}$. Then define

$$\overline{xy}[n] := \overline{x}[n+k] \times \overline{y}[n+k].$$

Correctness follows from:

$$\begin{array}{rcl} \overline{xy}[n] & = & (x \pm 2^{-n-k})(y \pm 2^{-n-k}) \\ & = & xy \pm x2^{-n-k} \pm y2^{-n-k} \pm 2^{-2(n+k)} \\ |\overline{xy}[n] - xy| & \leq & (|x| + |y|)2^{-n-k} + 2^{-2(n+k)} \\ & = & 2^{-n-1} + 2^{2n-1} = 2^{-n}. \end{array}$$

Q.E.D.

Theorem 6. *The class of computable real numbers is a subfield of $\mathbb{R}$.*

Proof. Clearly, 0 and 1 are computable reals. If x, y are computable reals, then we claim that so are

$$x \pm y, xy, x/y$$

provided x/y is defined (i.e., $y \neq 0$). According to the previous lemma, there are OTM's that compute each of the four field operations. Let us consider the case of $(x, y) \mapsto x + y$, as the other cases are similar. Let the OTM M_+ with oracles $\overline{x}$ and $\overline{y}$ and precision parameter p, produce a p-bit approximation of $x + y$, denoted $M_+^{\overline{x},\overline{y}}[p]$. Since x, y are computable, there exists Turing machines M_x and M_y that computes some $\overline{x}$ and $\overline{y}$. We must show that $\overline{x+y}$ is computable. Let us now construct a Turing machine M_{x+y} which computes $\overline{x+y}$ from $\overline{x}$ and $\overline{y}$. On input precision p, M_{x+y} will simulates M_+ in a step-by-step manner. But whenever M_+ queries its oracles for $\overline{x}$ or $\overline{y}$, we simulate M_x or M_y to obtain the answer. Thus, $M_{x+y}[p]$ computes the pth convergent of $\overline{x+y}$. Thus $x + y$ is a computable real. In a similar way, the numbers $x - y$, xy and x/y are computable. **Q.E.D.**

¶31. Function composition.

Lemma 7. *Let $f, g :\subseteq \mathbb{R} \to \mathbb{R}$ be partial functions. If f and g are computable, so is their function composition, $f \circ g : \mathbb{R} \to \mathbb{R}$*

Proof. Suppose M_f and M_g are OTMs to compute f and g. To compute $h = f \circ g$, we construct an OTM M_h that, on input $\overline{x} \in \mathcal{A}_x$ and $p \in \mathbb{N}$, begins to simulate M_f on input $g(x)$ and p. Whenever M_f attempts to query its oracle $g(x)$ with some precision p', M_h will first place p' in the precision tape of M_g in order to compute $M_g^{\overline{x}}[p']$. Upon getting this value, it can resume the simulation of M_f. Eventually, M_h will halt with the correct output $f(g(x))$. It is also clear that if $\overline{x} \notin \mathtt{domain}(g)$, then $M_h^{\overline{x}}[p] \uparrow$. **Q.E.D.**

In contrast to this result, suppose f and g are (absolutely) approximable. There is no obvious way to approximate $f \circ g$.

Exercises

Exercise 3.1: Show that the following partial functions are computable: $\ln x$ and $f(x) = p(x)/q(x)$ where f is a rational function defined by integer polynomials p, q. Note: $\ln x$ is undefined iff $x \leq 0$, and $f(x)$ is undefine when $q(x) = 0$. ♢

Exercise 3.2: Recall a function $f : \mathbb{R} \to \mathbb{R}$ is type-2 computable if it is computed by a Type-2 Turing machine (see ¶8). What is the relationship between type-2 computability and our notion of computability using OTM's?
(i) Show that f is type-2 computable iff f is computable.
(ii) In type-2 computation, if the machine halts or produces only a finite output sequence, it's output is considered undefined. How does the result of (i) extend to partial functions $f :\subseteq \mathbb{R} \to \mathbb{R}$? ♢

Exercise 3.3: Show that the real number x_A is uncomputable if A is a r.e., but not recursive. ♢

Exercise 3.4: If $f :\subseteq \mathbb{R} \to \mathbb{R}$ is a computable partial function and $x \in \mathtt{domain}(f)$ is computable, then $f(x)$ is computable. ♢

Exercise 3.5: What does it mean for an OTM M to compute a real function $f : \mathbb{R} \to \mathbb{R}$ using the Dedekind cut representation? The oracle for $x \in \mathbb{R}$ is just the Dedekind cut $\overline{x} \subseteq \mathbb{D}$. Our oracle, given any $d \in \mathbb{D}$, will return with 1 if $d \in \overline{x}$ and return 0 otherwise. Instead of the precision tape, our input tape now contains a dyadic d'. Then $M^{\overline{x}}(d')$ will output 1 iff $d' \in \overline{f(x)}$. Prove that a function f is computable using Cauchy oracles iff it is computable usind Dedekind oracles. ♢

Exercise 3.6: Complete the proof of Lemma 5. ♢

End Exercises

§4. Real Predicates

¶32. Real functions with countable range. Consider the computation of partial real functions whose range $\mathtt{range}(f) := \{f(\mathbf{x}) : \mathbf{x} \in \mathtt{domain}(f)\}$ is a countable set. Without loss of generality, assume

$$f :\subseteq \mathbb{R}^k \to \mathbb{D}. \tag{14}$$

In the special case where $\mathtt{range}(f)$ is a finite set, we call f a **real predicate**. In computability theory, one often speaks of "deciding" predicates instead of "computing" predicates.

In typical predicates, we have $\mathtt{range}(f) = \{0, 1\}$; call these "logical predicates" of standard two-valued logic. But in the field of computational geometry, predicates with $\mathtt{range}(f) = \{-1, 0, +1\}$ is more common (call these "geometric predicates"). An example of a geometric predicate is $f(p, C) = 0$ if point p lies on circle C, $f(p, C) = -1$ if p lies outside C, or $f(p, C) = +1$ if p lies inside C.

The computation of a function f such as in (14) presents some difficulties: *suppose the domain* $\mathtt{domain}(f)$ *contains an open set S, and f restricted to S is not a constant function, then f cannot be continous.*

Since all computable functions are continuous (see below), this implies that f is uncomputable. One solution is to relax the notion of computability for such functions, by considering conditional computability. Recall that this means we do not restrict the behavior of the OTM for inputs outside of $\mathtt{domain}(f)$.

For example, consider the function

$$f_0 :\subseteq \mathbb{R} \to \mathbb{D} \tag{15}$$

where $f_0(x) =\uparrow$ if $x \notin \mathbb{N}$; otherwise $f_0(x) := 2^{-n}$ if $\phi_x(x)$ (i.e., the j-th STM on input j) halts in n steps.

We can conditionally compute f_0 as follows: on input oracle $\overline{x}$ and $p \in \mathbb{N}$, we first compute a $j \in \mathbb{N}$ such that $x = j$ iff $x \in \mathbb{N}$. Note that we cannot be sure if $x \in \mathbb{N}$, that is why our algorithm is only a "conditional" computation of f_0. Nevertheless, we may simulate $\phi_j(j)$ for p steps. If it halts in $k \leq p$ steps, we output $\overline{2^{-k}}$. Otherwise, we output 0. It is not hard to prove that the this procedure is correct.

¶33. The equality and comparison predicates. Consider the following real predicates:

- (Equality) $f_= : \mathbb{R}^2 \to \{0, 1\}$ where $f_=(x, y) = 1$ if $x = y$ and $f_=(x, y) = 0$ otherwise.
- (Comparison) $f_{\neq} : \mathbb{R}^2 \to \{0, 1\}$ where $f_{\neq}(x, y) = 1$ if $x > y$, $f_{\neq}(x, y) = 0$ if $x < y$ and $f_{\neq}(x, y) =\uparrow$ if $x = y$.

Theorem 8.
(i) The predicate $f_=$ is uncomputable.
(ii) The predicate $f_{\neq}$ is computable.

Proof. (i) This holds for the simple reason noted above: after a finite number of steps, an OTM does not have enough information to determine equality of the input numbers. By way of contradiction, assume there is an OTM $M_=$ to decide $f_=$. Consider the operation of $M_=$ on the input $\overline{x} = (0, 0, 0, \ldots)$ comprising all zeros, and the input $\overline{y} = (1, 1/2, 1/3, \ldots)$ whose pth convergent is $1/(p + 1)$. Clearly, $M_=$ outputs 1, and there is some $k \in \mathbb{N}$ such that the oracles only queries $\overline{x}[p]$ and $\overline{y}[p]$ where $p \leq k$. We modify $\overline{y}$ so that for all $p > k$, the pth convergent is now $1/(k + 1)$. So the output on the modified input ought to have been 0. But $M_=$'s behavior on the modified input is exactly as before, and will wrongly output 1.

(ii) We construct an OTM $M_{\neq}$ with oracles $\overline{x}$ and $\overline{y}$ operates as follows: for $p = 0, 1, 2, \ldots$, it checks if

$$|\overline{x}[p] - \overline{y}[p]| > 2^{-p+1}.$$

If so, it outputs 1 iff $\overline{x}[p] > \overline{y}[p]$ and otherwise outputs 0. If the input satisfies $x \neq y$, our $M_{\neq}$ will eventually produce an output. If $x = y$, then $M_{\neq} \uparrow$.

Q.E.D.

¶34. The Russian approach. We try to give a deeper reason for why the function $f_=$ is uncomputable, by avoiding arguments based on an OTM's inability to process the entire input before producing an output. To do this, we use the **Russian Approach** to real computation [7, Chapter 9]. Here, a real number x is represented by its Gödel number $\overline{\overline{x}} \in \mathbb{N}$ (see ¶17).

First, we connect the Russian Approach to the Polish Approach. Consider an enumeration $M_0, M_1, \ldots$ of all STM's, where each M_i is interpreted to compute a partial function of the form

$$\phi_i :\subseteq \mathbb{N} \to \mathbb{D}.$$

Moreover, there is a deterministic UTM U such that $U(i, n) = \phi_i(n)$ for all i, n. If $|\phi_i[p] - x| \leq 2^{-p}$ for all $p \in \mathbb{N}$, then we say i is a Gödel number of the real number x. The following is immediate:

Lemma 9. *A real number x is computable iff it has a Gödel number $\overline{\overline{x}}$.*

This shows that the Russian approach is necessarily confined to computable reals.

¶35. Russian computability. A function $f :\subseteq \mathbb{R}^2 \to \mathbb{R}$ is **Russian computable** if there exists a STM M that for all $x, y \in \mathbb{R}$, if $f(x, y) =\downarrow$, then on any input of the form $(\overline{\overline{x}}, \overline{\overline{y}})$ M will output some Gödel number of the form $\overline{\overline{f(x, y)}}$.

Remarks:
(a) This definition uses a regular Turing machine, as opposed to oracle Turing machines.
(b) Only the behavior of f and M on computable real inputs are relevant. In particular, the behavior of M is irrelevant if the input are not valid Gödel numbers.
(c) A necessary condition for f to be computable is that for all computable x, y, the real number $f(x, y)$ must be computable.

¶36. Equality is not Russian computable.

Lemma 10. *The predicate $f_=$ is not Russian computable.*

Proof. For any n, let

$$h_n(k) = \begin{cases} 1 & \text{if } M_n(n) \text{ halts in} \leq k \text{ steps,} \\ 0 & \text{else.} \end{cases}$$

Consider the real number $x_n := \sum_{i=0}^{\infty} h_n(i)2^{-i}$. Note that $n \in$ co-DIAG iff $x_n = 0$. We can construct the Gödel number $\overline{\overline{x_n}}$ of the Turing machine which, on input k, outputs $\sum_{i=0}^{k} h_n(i)2^{-i}$. Hence $\overline{\overline{x_n}}$ is the Gödel number of x_n.

Let Z be the set of $n \in \mathbb{N}$ such that $x_n = 0$. Since the function $t : n \mapsto \overline{\overline{x_n}}$ is computable, we have just shown that

$$\text{co-DIAG} \leq_m Z \text{ (via } t).$$

This proves that Z is not recursive. If $f_=$ is Russian computable, then Z is recursive, which is a contradiction. **Q.E.D.**

The general conclusion is that we cannot decide equality of real numbers. This is equivalent to the undecidability of checking if a given real number is zero (Exercise).

Later, we will return to the question of computing predicates: then we will introduce a way to approximate predicates which is motivated by recognizing subsets of $\mathbb{R}^k$.

¶37. Is square root computable? The answer depends on how we define the nominal domain of the square root function. Consider the following three definitions of the square root function:

- $sqrt_1 :\subseteq \mathbb{R} \to \mathbb{R}$, with $sqrt(x) = \uparrow$ if $x < 0$, and $sqrt(x) = \sqrt{x}$ for $x \geq 0$.
- $sqrt_0 : (0, \infty) \to \mathbb{R}$. This is the restriction of $sqrt(x)$ to $(0, \infty)$.
- $sqrt : [0, \infty) \to \mathbb{R}$. This is the restriction of $sqrt(x)$ to $[0, \infty)$.

We regard $sqrt$ as the main definition among these three variants. Note that 0 is in the proper domains of $sqrt_1$ and $sqrt$, but not in the proper domain of $sqrt_0$. One might suggest that both $sqrt_1$ and $sqrt$ should not be computable since it is undecidable if an input real number is 0 or not (¶33). But note that $sqrt_1$ and $sqrt$ only has to "approximate 0 in its the output", not to "decide if the input is 0". There is an important difference — in approximating, the oracle machine can output a representation of zero (as precision $p \to \infty$) without ever knowing (for any particular p) whether the input is 0. Indeed, the next result shows that we *can* approximate zero in case of $sqrt$.

Theorem 11. *(i) The function $sqrt_0$ is computable.*
(ii) The function $sqrt_1 :\subseteq \mathbb{R} \to \mathbb{R}$ is uncomputable.
(iii) The function $sqrt : [0, \infty) \to \mathbb{R}$ is computable.

Of course (iii) implies (i). However we will prove (iii) under the assumption (i).

Proof of (ii): Suppose M is an OTM to compute $sqrt_1$. To obtain a contradiction, consider the oracle function $\overline{0^-}$ for zero where $\overline{0^-}[p] = -2^{-p}$ for all p. Consider the output of M on input $\overline{0^-}$ and $p = 1$. Clearly, M must halt after some t steps. Now, consider another oracle function $\overline{x}$ where $\overline{x}[p] = \overline{0^-}[p]$ if $p \leq t$, and $\overline{x}[p] = -2^{-t}$ for all $p > t$. Clearly, $\overline{x}$ represents the negative number 2^{-t}, but M on input $\overline{x}$ and $p = 1$ will halt with an output as before, contradiction.

Proof of (iii): This proof assume (i) holds. To show the computability of $sqrt$, we must ensure that, in case the input represents 0, then we output a representation of 0. We do this as follows. On input $\overline{x}$ and p where $x \geq 0$, we first compute $\overline{x}[2p+2]$ and compare this to 2^{-2p-1}. If $\overline{x}[2p+2] \geq 2^{-2p-1}$ then $x \geq 2^{-2p-2}$, and we can compute a p-bit approximation of $\sqrt{x}$ using an algorithm for $sqrt_0$ (from part(i)). Otherwise, we simply output 0. We show that this output of 0 is correct: since $\overline{x}[2p+2] < 2^{-2p-1}$, we have $x < 2^{-2p}$, or $\sqrt{x} < 2^{-p}$. Hence 0 is a p-bit approximation to $\sqrt{x}$.

¶38. Proof that $sqrt_0$ is computable. We now prove (i).

1. *The function $sqrt_0$ is approximable.* This means that for each $x \in \mathbb{D} \cap (0, \infty)$, and $p \in \mathbb{N}$, we can compute some $\sqrt{}(x)[p] \in D$ which is a p-bit approximation to $\sqrt{x}$.

   There are well-known efficient Newton-type methods for doing this. However, since we do not care about time bound, we can naively compute z by binary search among the elements of $\mathbb{D}_{p+1}[0, \lceil y \rceil]$. More precisely,

suppose we know an interval $[a,b]$ containing $\sqrt{x}$ where $a,b \in \mathbb{D}_{p+1}[0,\lceil y \rceil]$. let $c = \lfloor (a+b)/2 \rfloor_{p+1}$ is the nearest element in $\mathbb{D}_{p+1}$ to that is $\leq (a+b)/2$. By computing c^2 and computing it to x, we can decide whether $c \geq \sqrt{x}$ or $c \leq \sqrt{x}$. If $c^2 \geq x$, then we narrow our interval to $[a,c]$; otherwise we narrow it to $[c,b]$.

2. We need a basic estimate: *If* $|\rho| \leq 1/2$ *then* $\sqrt{1+\rho} = 1 \pm \frac{5\rho}{8}$.

   This amounts to a linearized estimate of the squareroot of $(1+\rho)$, when ρ is small. For instance, this implies that if $|\rho| \leq 1/2$ then $\sqrt{1+\rho} = 1 \pm \rho$. The proof is standard, using Taylor's expansion of $\sqrt{1+\rho}$. We leave it as an Exercise.

3. We are ready to present the OTM for computing $sqrt_0$, assuming an input oracle $\overline{x}$ for $x > 0$, and a precision $p \in \mathbb{N}$. We want to compute some $z \in \mathbb{D}$ such that $z = \sqrt{x} \pm 2^{-p}$.

   - STEP 1. Iteratively compute $\overline{x}[2n]$ for $n = 1, 2, 4, 8, \ldots$, until $\overline{x}[2n] \geq 2^{-2n+1}$ holds. This iteration must terminate because $x > 0$.
   - STEP 2. Determine the unique ℓ such that $2^{-2\ell+1} \leq \overline{x}[2n] < 2^{-2\ell+2}$. This implies that $\ell \leq n$. It follows that
   $$x \leq \overline{x}[2n] + 2^{-2n} < 2^{-2\ell+3}$$
   and
   $$x \geq \overline{x}[2n] - 2^{-2n} \geq 2^{-2\ell}.$$
   Hence $2^{-\ell} \leq \sqrt{x} < 2^{-\ell+1}$.
   - STEP 3. Compute $y := \overline{x}[\ell + p + 4]$. Then $y = x \pm 2^{-\ell-p-4} = x(1 \pm 2^{\ell-p-2})$.
   - STEP 4. Using our approximation function $\widetilde{\sqrt{}}$ above, compute an approximation $z := sqrt(y)[p+1] = \sqrt{y} \pm 2^{-p-1}$. Output this z. To see that this output is correct, it suffices to show that
   $$\sqrt{y} = \sqrt{x}\sqrt{1 \pm 2^{\ell-p-2}} = \sqrt{x}(1 \pm \frac{5}{8}2^{\ell-p-2}) = \sqrt{x} \pm 2^{-p-1}.$$

EXERCISES

Exercise 4.1: Let $f_{=0} : \mathbb{R} \to \mathbb{R}$ be the predicate such that $f_{=0}(x) = 1$ iff $x = 0$.
(i) Show that $f_{=0}$ is computable iff $f_=$ is computable.
(ii) Show that $f_{=0}$ is Russian computable iff $f_=$ is Russian computable. ♢

Exercise 4.2: The predicate $f_=$ is total and undecidable. Let us show that it is undecidable for a much stronger reason. Consider a partial function analogue: define $g_= :\subseteq \mathbb{R}^2 \to \mathbb{R}$ where $g_=(x,y) = \uparrow$ if $x \neq y$ and $g_=(x,x) = 1$. Show that $g_=$ is conditionally undecidable. ♢

Exercise 4.3: Let $f : \mathbb{R} \to \mathbb{R}$ have a finite range, i.e., $\{f(x) : x \in \texttt{domain}(f)\}$ is a finite set. ♢

Exercise 4.4: Give an alternative proof that $sqrt_0$ is approximable, by using Newton iteration. This should be more efficient than the binary search method of ¶38. Analyze the complexity of your algorithm. ♢

Exercise 4.5: We want to bound $\sqrt{1+\rho}$ (¶38(b)). Let $n!!$ denote the **double factorial** defined by $n!! = (n-2)!! \cdot n$ when $n \geq 1$, and $(-1)!! = 0!! = 1$.
(i) Show that
$$(1+\rho)^{1/2} = 1 + \frac{\rho}{2} - \sum_{k \geq 1} \rho^{2k} \frac{(4k-3)!!}{4^k (2k)!} \left(1 - \rho \frac{4k-1}{4k+2}\right).$$
HINT: use the Binomial Theorem to expand $(1+\rho)^{1/2} = 1 + \frac{\rho}{2} + \frac{(1/2)(-1/2)}{2!}\rho^2 + \cdots$.
(ii) If $|\rho| \leq \frac{1}{2}$ then $\sqrt{1+\rho} = 1 + \rho/2 - K$ where $0 \leq K \leq \frac{3\rho^2}{16(1-\rho)}$.
(iii) If $|\rho| \leq \frac{1}{2}$ then $\sqrt{1+\rho} = 1 \pm 5\rho/8$. ♢

Exercise 4.6: Prove that in the Taylor expansion,
$$\sqrt{1+\rho} == 1 + \frac{\rho}{2} - \frac{1}{8}\rho^2 + \frac{1}{16}\rho^3 - \frac{5}{128}\rho^4 + \frac{7}{256}\rho^5 + \cdots$$
every coefficient is a dyadic number. ♢

END EXERCISES

§5. Continuity of Computable Functions

¶39. Compact Domains. It is often simpler to work with real functions whose domains are real intervals. For $[a,b] \subseteq \mathbb{R}$, let $C[a,b]$ be the set of real functions $f : [a,b] \to \mathbb{R}$ which are continuous.

¶40. Modulus of continuity. Let $f : [a,b] \to \mathbb{R}$. In analysis, we say that f is **continuous** if for all $\varepsilon > 0$ there is a $\delta > 0$ such that $|x-y| \leq \delta$ implies $|f(x) - f(y)| \leq \varepsilon$. We will work with a quantitative version of this concept: a function $m : \mathbb{N} \to \mathbb{N}$ is a **modulus of continuity** (or modulus function) for f if for all $n \in \mathbb{N}$, and $x, y \in [a,b]$,

$$|x-y| \leq 2^{-m(n)} \quad \Rightarrow \quad |f(x) - f(y)| \leq 2^{-n}.$$

The following is immediate:

Lemma 12. *$f : [a,b] \to \mathbb{R}$ is continuous iff it has a modulus function.*

Proof. If f is continuous, then for every $n \in \mathbb{N}$, there is a $\delta(n)$ such that $|x-y| \leq \delta(n)$ implies $|f(x)-f(y)| \leq 2^{-n}$. We may choose $m(n) = \lceil \lg(1/\delta(n)) \rceil$. Conversely, if m is a modulus function, then for every $\epsilon > 0$ we choose $\delta = 2^{-m(n)}$ where $n = \lceil \lg(1/\epsilon) \rceil$. **Q.E.D.**

¶41. Heine-Borel Theorem. Let $S \subseteq \mathbb{R}$ and I be any index set. A collection $\{C_i : i \in I\}$ where each $C_i \in \mathbb{R}$ is an open set is called an **open cover** of S if $S \subseteq \bigcup_{i \in I} C_i$. The Heine-Borel Theorem is a fundamental fact about the real numbers: it says that if S is compact, i.e., S is closed and bounded, and $\{C_i : i \in I\}$ is an open cover of S, then there exists a finite subset $J \subseteq I$ such that $\{C_j : j \in I\}$ is also an open cover of S.

¶42. Suppose M is an OTM that computes a partial function $g :\subseteq \mathbb{R} \to \mathbb{R}$. Define the function

$$k_M(x,p) := \text{the largest } k \text{ such that the computation of } M^{\beta_x}[p] \text{ queries } \beta_x[k].$$

Note that this definition assumes that M uses the standard Cauchy representation β_x as oracle. If this oracle was never queried, define $k_M(x,p) = 0$.

Write $k(x,p) = k_M(x,p)$ when M is understood. For each $x \in [a,b]$, define the open interval $I(x,p) := (\ell(x,p), r(x,p))$ where

$$\begin{aligned} \ell(x,p) &= \beta_x[k(x,p)] - 2^{-k(x,p)}, \\ r(x,p) &= \beta_x[k(x,p)] + 2^{-k(x,p)}. \end{aligned}$$

Lemma 13. *Suppose $y \in I(x,p) \cap [a,b]$ then $|f(y) - f(x)| \leq 2^{-p+1}$.*

Proof. Consider the representation $\overline{y}$ with the property

$$\overline{y}[n] = \begin{cases} \beta_x[n] & \text{if } n \leq k(x,p), \\ \beta_y[n] & \text{if } n > k(x,p). \end{cases}$$

By Lemma 2(iii), we know that $\overline{y}$ is a representation of y. Our construction of $\overline{y}$ ensures that the computation of $M^{\beta_x}[p]$ is indistinguishable from $M^{\overline{y}}[p]$. Therefore

$$\begin{aligned} |f(y) - f(x)| &\leq |f(y) - M^{\overline{y}}[p]| + |M^{\beta_x}[p] - f(x)| \\ &\leq 2^{-p} + 2^{-p} = 2^{-p+1}. \end{aligned}$$

Q.E.D.

¶43. Computability implies continuity. The key lemma is the following:

Lemma 14 (Continuity). *If $f : [a,b] \to \mathbb{R}$ is computable then f has a recursive modulus function $m : \mathbb{N} \to \mathbb{N}$.*

Proof. Let M be an OTM that computes f, and let $k(x,p) = k_M(x,p)$. We will fix $p \in \mathbb{N}$ for most of this argument.

Consider the set $\mathcal{C} = \{I(x, p+3)) : x \in [a,b]\}$. Note that $x \in I(x,p+3)$ because the center of $I(x,p+3)$ is $\beta_x[k(x,p+3)]$ and $x - \beta_x[k(x,p+3)] < 2^{-k(x,p+3)}$ (by Lemma 2(ii)). Hence $\mathcal{C}$ is an open cover of $[a,b]$.

Hence $\mathcal{C}$ contains a finite subcover by Heine-Borel. Let the finite subcover be $\{I(x_i, p+3) : i = 1, 2, \ldots, t\}$. Next note that if we replace x by $x' = \beta_x[k(x,p+3)]$, the interval $I(x', p+3)$ is identical to $I(x, p+3)$. Hence wlog, assume that each x_i is equal to the dyadic number $\beta_{x_i}[k(x_i, p+3)]$.

Writing $I(x_i, p+3) = (\ell_i, r_i)$ and we may assume that $\ell_1 < \ell_2 < \cdots < \ell_t$ and $r_1 < r_2 < \cdots < r_t$. Define

$$m(p) = 2 + \max\{k(x_i, p+3) : i = 1, \ldots, t\}.$$

We will show that $m(p)$ acts as a modulus function.

CLAIM 1: If $y, z \in [a, b]$, and $0 < z - y \leq 2^{-m(n)}$ then there is some $i, j = 1, \ldots, t$ such that either (A)

$$\ell_i < y < z < r_i$$

or (B)

$$\ell_i < y \leq \ell_j < r_i < z < r_j.$$

To see this, let i be the maximum index such that $\ell_i < y$ and j the minimum index such that $z < r_j$. If $i = j$, then we have the situation of (A). If $i = j - 1$, we have the situation of (B). If $i < j - 1$, then $y \leq \ell_{i+1} < r_{i+1} \leq z$. Thus $z - y \geq r_{i+1} - \ell_{i+1} = 2^{1-k(x_{i+1}, p+3)} > 2^{-m(n)}$, contradicting our assumption about $z - y$.

CLAIM 2: The function $m(p)$ is a modulus function.

To see this, let y, z as in CLAIM 1. We want to show $|f(y) - f(z)| \leq 2^{-p}$. There are two cases in CLAIM 1. In case (A), we have

$$|f(y) - f(z)| \leq |f(y) - f(x_i)| + |f(x_i) - f(z)|$$

where $|f(y) - f(x_i)|$ and $|f(x_i) - f(z)|$ are each upper bounded by 2^{-p-2} (Lemma 13). Hence $|f(y) - f(z)| \leq 2^{-p-1}$. In case (B), choosing any $u \in (\ell(x_j), r(x_i))$, we have

$$|f(y) - f(z)| \leq |f(y) - f(x_i)| + |f(x_i) - f(u)| + |f(u) - f(x_j)| + |f(x_j) - f(z)|$$

Again Lemma 13 that each absolute value on the right hand side is at most 2^{-p-2} and so $|f(y) - f(z)| \leq 2^{-p}$. This proves our CLAIM.

Although $m(p)$ is a modulus function, its recursiveness is unclear. But here is a recursive version m^*: on input $p \in \mathbb{N}$, search for the smallest $s^* = 0, 1, 2, \cdots$ such that

$$[a, b] \subseteq \bigcup_{d \in \mathbb{D}_{s^*}} I(d, p+3). \tag{16}$$

Note that this search will terminate since there exists an s such that $\{x_1, \ldots, x_t\} \subseteq \mathbb{D}_s$ and so $[a, b] \subseteq \bigcup_{i=1}^t I(x_i, p+3)$. If (16) holds, we may redefine the modulus function as $m^*(p) = 2 + \max\{k(d, p+3) : d \in \mathbb{D}_s^* \cap [a, b]\}$. It is clear that m^* is recursive. Since $m^*(p) \geq m(p)$, we conclude that it is also a modulus for f. **Q.E.D.**

¶44. We conclude that discontinous functions such as the step function $x \mapsto \delta x > 0$, or floor function $x \mapsto \lfloor x \rfloor$, are not computable. One view is that this conclusion is "to be expected". Some people in computing hold the position that input data in real computation is inherently uncertain, and this conclusion is consistent with such a view. Nevertheless, there are many problems in real computation where data is not at all uncertain. Examples include the standard problems in the field of computational geometry, in automatic theorem proving for elementary geometry, or intersections of algebraic surfaces within a CAD model. For such applications, discontinuity is essential, and so such a conclusion is unacceptable. We shall return to this important issue.

¶45. Characterization of computability. The previous lemma leads to an important characterization of computable real functions.

Theorem 15. *A function $f : [a, b] \to \mathbb{R}$ is computable iff f is absolutely approximable and has a recursive modulus function.*

Proof. Suppose f has a recursive modulus function $m : \mathbb{N} \to \mathbb{N}$ and a computable absolute approximation $\widetilde{f}$. Consider the following OTM M to compute f: give an oracle function $\overline{x}$ and precision $p \in \mathbb{N}$, M first computes $m(p+1)$ and then obtains from the oracle $y = \overline{x}[m(p+1)]$. Finally, M computes and outputs $z = f(y)[p+1]$. Note that

$$z = f(y) \pm 2^{-p-1} = (f(x) \pm 2^{-p-1}) \pm 2^{-p-1} = f(x) \pm 2^{-p},$$

as desired.

Conversely, suppose f is computable. By Lemma 14, f has a computable modulus m, so it remains to show that f is absolutely approximable. Let M be an OTM that computes f. It is easy to construct an ordinary Turing machine M' that, on input $d \in \mathbb{D}$ and $p \in \mathbb{N}$, simulates M on the oracle β_d and p. M will output $M^{\beta_d}[p]$ which is clearly a p-bit approximation of $f(d)$. **Q.E.D.**

¶46. We seek examples of continuous functions are not computable. Let $\beta \in \mathbb{R}$. Consider the function $g_\beta :\subseteq \mathbb{R} \to \mathbb{R}$ where $g_\beta(x) =\uparrow$ iff $x \geq \beta$. We may assume g_β is continuous within its proper domain $[\beta, \infty)$. Nevertheless, g_β is uncomputable since we can reduce the problem of checking if an input x is equal to β to g_β. The reader will rightfully object to such examples because they say nothing about the nature of the function restricted to its proper domain. We will provide more acceptable examples later.

EXERCISES

Exercise 5.1: Show that g_β is uncomputable. ♢

Exercise 5.2: Generalize the characterization of computable f to the case where $\texttt{domain}(f) = \mathbb{R}$. HINT: use localized modulus function as described in the text. ♢

Exercise 5.3: Modify the proof of Lemma 14, to show that if t is the time complexity of an OTM which a real function f, then t is also a modulus function for f. See below for definition of time complexity. ♢

Exercise 5.4: Generalize the continuity characterization of computable functions to multidimensional real functions, $f : [0, 1]^k \to \mathbb{R}$. ♢

END EXERCISES

§6. Complexity of Real Functions

¶47. Time complexity. A **complexity function** T is any partial real function $T :\subseteq \mathbb{R}^m \to \mathbb{R}$.
For example, $T(x) = \sqrt{x}$ and $T(x) = x \lg x$ are complexity functions. So is $T(x, y) = x^2y + y^x \lg x$.
Let $f :\subseteq \mathbb{R}^k \to \mathbb{R}$ be a partial real function.

- (i) A **local time complexity** of f is $T :\subseteq \mathbb{R}^{k+1} \to \mathbb{R}$ such that for some OTM M that computes f, and for all $\overline{\mathbf{x}}$, $\mathbf{x} \in \texttt{domain}(f)$ and $p \in \mathbb{N}$, $M^{\overline{\mathbf{x}}}[p]$ halts in $\leq T(\mathbf{x}, p)$ steps.
- (ii) A **global time complexity** of f is $T :\subseteq \mathbb{R} \to \mathbb{R}$ such that for some OTM M that computes f, and for all $\overline{\mathbf{x}}$ where $\mathbf{x} \in \texttt{domain}(f)$, and $p \in \mathbb{N}$, $M^{\overline{\mathbf{x}}}[p]$ halts $\leq T(p)$ steps.

¶48. Remarks.
1. If T is a local complexity function for f, then $\texttt{domain}(f) \times \mathbb{N} \subseteq \texttt{domain}(T)$.
2. Note that when $\mathbf{x} \notin \texttt{domain}(f)$, then we do not care how much time the computation $M^{\overline{\mathbf{x}}}[p]$ takes.
3. The bit size of p is normally take to be $O(\log p)$ and complexity T implies the machine uses $\leq T(\log p)$ steps. However, that is not our definition; the precision parameter p is not treated like ordinary input parameters. Instead, the precision parameter p has **(precision) size** p. One way to ensure this special treatment of precision is to require that an OTM can only write in unary notation on the query tape.
4. It is desirable not to assume unary notation for representing p. So, in practice, we expect p to be represented in binary. But we still regard the precision size to be "p" and not "$\lg p$".

¶49. Complexity of approximations. In the local time complexity $T(\mathbf{x}, p)$, the parameter $(x_1, \ldots, x_k) = \mathbf{x} \in \mathbb{R}^k$ is not really an "input size" measure for complexity: we could have $\sum_{i=1}^k |x_k|$ that are arbitrarily small but representing arbitrarily complex input values. For each fixed $\mathbf{x}$, the function $T_{\mathbf{x}}(p) := T(\mathbf{x}, p)$ is closer to our usual idea of a complexity function. However, we have no idea of its dependence on $\mathbf{x}$. On the other hand, global time complexity $T(p)$ explicitly exclude any dependence on $\mathbf{x}$. In either case, it is counter-intuitive to our usual ideas of complexity functions.

We propose to look at another complexity function which looks at *some* choices of $\mathbf{x}$. Let

$$\widetilde{f} :\subseteq \mathbb{D}^k \times \mathbb{N} \to \mathbb{D}$$

be an absolute approximation of f. Define **input size** of $\mathbf{x} \in \mathbb{D}^k$ to be the length of the string representation $\overline{\mathbf{x}}$ (this depends on the convention that the representation of dyadic numbers are held fixed). Write $size(\mathbf{x})$ for this input size. The **precision size** of p is declared to be p (as before).

We say $\widetilde{f}$ has **time complexity** $T : \mathbb{R}^2 \to \mathbb{R}$ if there is a Turing machine M such that on input $\mathbf{x} \in \texttt{domain}(\widetilde{f})$ and $p \in \mathbb{N}$, M takes $\leq T(size(\mathbf{x}), p)$ steps. We say f is **polynomial-time approximable** if it has a approximating function $\widetilde{f}$ whose time complexity is polynomial.

¶50. Polynomial time computability. We focus on real functions that are efficiently computable.

- Let $T(\mathbf{x}, p)$ be a local time complexity of an OTM M. We say that M is **locally polynomial-time** if for each $\mathbf{x} \in \mathtt{domain}(M)$, the function $T_{\mathbf{x}}(p) = T(\mathbf{x}, p)$ is bounded by a polynomial in the parameter p.
- Let $T(p)$ be the global time complexity of an OTM M. We say that M is **globally polynomial-time** if $T(p)$ is bounded by a polynomial in the parameter p.
- Define
$$P_{C[a,b]}$$
to be the set of real functions $f : [a, b] \to \mathbb{R}$ with global time complexity $T(p)$ that is a polynomial.
- We say the function f is **polynomial-time approximable** if some absolute approximation $\widetilde{f} :\subseteq \mathbb{D}^k \times \mathbb{N} \to \mathbb{D}$ can be computed in time that is polynomial in $size(\mathbf{d}) + p$ where the input is $\mathbf{d} \in \mathbb{D}^k$ and $p \in \mathbb{N}$.
- Alternatively, f is polynomial-time approximable if $\widetilde{f} \in FP$ where FP denotes the class of discrete functions $h : \Sigma^* \to \Sigma^*$ that are computed in polynomial-time by ordinary Turing machines.

Theorem 16. *The field operations $(\pm, \times, \div)$ are locally polynomial-time computable functions.*

Proof. We just verify that the algorithms for the field operations as presented in (¶27,¶30) are locally polynomial-time. **Q.E.D.**

¶51. We extend the continuity characterization of computable function to $P_{C[0,1]}$. We say $f : [0, 1] \to \mathbb{R}$ has a **polynomial modulus** $m : \mathbb{N} \to \mathbb{N}$ if m is a modulus and $m(n) = O(n^k)$ for some constant $k > 0$.

Theorem 17. *$f \subseteq P_{C[0,1]}$ iff f is polynomial-time approximable and f has a polynomial modulus.*

Proof. ($\Rightarrow$) Suppose $f \in P_{C[0,1]}$ is computable by an OTM M in global polynomial time $T(p)$. In the proof of Lemma 14, we show that function $m(p) := 2 + k_M(x, p + 3)$ serves as a modulus function. Since $T(p) \geq k_M(x, p)$ for all x, p, it follows that $m(p) := 2 + T(p + 3)$ is a polynomial modulus function.

It remains to show that the approximation function $f(d)[p]$ can be computed in polynomial-time. To do this, we simulate M on input $\overline{d}$ and p. Whenever M queries the oracle $\overline{d}[n]$ on some precision n, we can easily compute the answer in $O(size(d) + n)$ steps. So the overall computation is polynomial in $size(d) + p$.

($\Leftarrow$) Conversely, from m is a polynomial modulus and $\widetilde{f} \in FP$. We construct an OTM M whose operation for $M^{\widetilde{x}}[p]$ goes as follows: it computes $m(p+1)$ and asks the oracle for $y = x[m(p+1)]$. Finally it outputs $f(y)[p+1]$. The proof of Theorem 15 shows the correctness of this output. We now observe that this computation is polynomial time in p. **Q.E.D.**

Exercises

Exercise 6.1: Modify our proof of Lemma 14 to show that if an OTM M computes f in uniform time $T : \mathbb{R} \to \mathbb{R}$ then $m(n) := 2 + T(n + 2)$ is a modulus function for f. ♢

End Exercises

§7. Recursively Open Sets

In the previous section we characterized computable functions with a compact domain, of the form $f : [a, b] \to \mathbb{R}$. In this section we want to extend this characterization to two other classes of functions: (1) Total functions with non-compact domain such as $f : \mathbb{R} \to \mathbb{R}$ or $f : (a, b) \to \mathbb{R}$. (2) Partial functions $f :\subseteq \mathbb{R} \to \mathbb{R}$. We will need the concept of recursively open sets.

¶52. Let $\phi : \mathbb{N} \to \mathbb{D}$. This function defines a sequence $(I_0, I_1, I_2, \ldots)$ of open intervals where $I_n = (\phi(2n), \phi(2n+1))$. We say ϕ is an **interval representation** of the open set $S = \bigcup_{n\geq 0} I_n$.

A set $S \subseteq \mathbb{R}$ is **recursively open** if there is an interval representation of S that is recursive. We say S is **recursively closed** if its complement $\mathbb{R} \setminus S$ is recursively open.

Note that if $\phi(2n) \geq \phi(2n+1)$ then the open interval $(\phi(2n), \phi(2n+1)$ is empty. In particular, the empty set $S = \emptyset$ is recursively open. So is $S = \mathbb{R}$. On the other hand, the sets $[0, 1]$ and $(0, 1]$ are not recursively open, since recursively open sets are open sets in the standard topology of $\mathbb{R}$.

Theorem 18. *The following statements are equivalent:*
(i) $S \subseteq \mathbb{R}$ is recursively open.
(ii) S is equal to $\mathtt{domain}(f)$ *for some computable $f :\subseteq \mathbb{R} \to \mathbb{R}$.*

Proof. Let $f :\subseteq \mathbb{R} \to \mathbb{R}$ be computed by an OTM M. We want to show that $S := \mathtt{domain}(f)$ is recursively open. Wlog, assume $\mathtt{domain}(f) \neq \emptyset$. For $k \in \mathbb{N}$, define

$$A_k := \left\{ d \in \mathbb{D} : M^{\beta_d}[1] \text{ halts within } k \text{ steps} \right\}.$$

Note that $A_k \subseteq A_{k+1}$. Also, A_k is recursive.

CLAIM 1: $d \in A_k$ implies that if $x = d \pm 2^{-k}$ then $x \in S$. Pf: if $x = d \pm 2^{-k}$ then there is a representation $\overline{x}$ such that $\overline{x}[n] = \beta_d(n)$ for all $n \leq k$. Then $M^{\overline{x}}[1] = M^{\beta_d}[1]$. So $M^{\overline{x}}[1]$ halts, and $x \subseteq S$. QED

CLAIM 2: if $x \in S$ then there is a k such that $d \in A_k$ and $x = d \pm 2^{-k}$. Pf: Consider $M^{\beta_x}[1]$. It halts in some k steps. Let $d = \beta_x[k]$. Then $M^{\beta_d}[1]$ halts in k steps. So $d \in A_k$. Moreover, $d = x \pm 2^{-k}$. QED.

From CLAIMS 1 and 2, we conclude that $S = \bigcup_{k\geq 1} \bigcup_{d\in A_k} (d - 2^{-k}, d + 2^{-k})$. This shows that S is recursively open.

Conversely, suppose $S = \bigcup_{n\geq 0}(\phi(2n), \phi(2n+1))$ where ϕ is recursive. We construct a OTM M such that M on input $\overline{x}$ and p will halt iff $x \in S$. We just use dovetailing to check the following list of conditions:

$$\phi(2n) < x < \phi(2n+1), \qquad (n = 0, 1, 2, \ldots).$$

If any of these conditions hold, we halt. Otherwise we loop. **Q.E.D.**

¶53. Squareroots again. Does this result contradict our result in ¶37, where the function $sqrt : [0, \infty) \to \mathbb{R}$ was shown computable. The set $\mathtt{domain}(sqrt) = [0, \infty)$ is clearly not open. This is not a contradiction because the nominal domain of $sqrt$ is not $\mathbb{R}$ as required by the theorem. Indeed, ¶37 show that the related $sqrt_1 :\subseteq \mathbb{R} \to \mathbb{R}$ with $\mathtt{domain}(sqrt_1) = [0, \infty)$ is not computable, consistent with the present more general result. The generalization of the previous theorem to such situations such as $sqrt$ is treated in the Exercise.

¶54. Local modulus function We observe that Lemma 14 which shows the existence of computable modulus functions can be slightly generalized as follows: if $[a, b] \subseteq \mathtt{domain}(f)$, the original proof still leads to a modulus function m, but for the restriction of f to $[a, b]$. This suggests the following generalization of modulus functions to arbitrary functions.

Let $f :\subseteq \mathbb{R} \to \mathbb{R}$. A function

$$m :\subseteq \mathbb{D}^2 \times \mathbb{N} \to \mathbb{N}$$

is called a **local modulus function** for f if for all $d_1, d_2 \in \mathbb{D}$ and $p \in \mathbb{N}$, if $[d_1, d_2]$ is not contained in $\mathtt{domain}(f)$, then $m(d_1, d_2, p) =\uparrow$. Otherwise, for all $x, y \in [d_1, d_2]$, if $|x - y| \leq 2^{-m(d_1,d_2,p)}$ then $|f(x) - f(y)| \leq 2^{-p}$.

Note m is now a partial function.

Lemma 19. *Let $f :\subseteq \mathbb{R} \to \mathbb{R}$ with recursively open* $\mathtt{domain}(f)$*. If f is computable, then f has a computable local modulus function.*

Proof. Let OTM M compute f, and ϕ a recursive interval representation for $\mathtt{domain}(f)$. We construct a computable local modulus function m for f as follows: on input $d_1, d_2 \in \mathbb{D}$ and $p \in \mathbb{N}$, we first check that d_1, d_2 belong to the domain of f by computing $M^{\beta_d}[1]$ where $d = d_1$ and d_2. Assuming both computations halt, we next compute $\Phi(i) = \bigcup_{j=0}^{i}(\phi(2j), \phi(2j+1))$ for $i = 0, 1, 2, \ldots$. Note that $\Phi(i)$ is just a union of disjoint intervals, which we can maintain as a list of their endpoints. We stop when we verify that $(d_1, d_2) \subseteq \Phi(i)$ for some i.

We proceed to compute a value $m(d_1, d_2, p)$ as in the proof of Lemma 14. For $s = 0, 1, 2, \ldots$, as in (16), we check if

$$[d_1, d_2] \subseteq \bigcup_{d\in\mathbb{D}_s} I(d, p+3),$$

adopting notations from that proof. This process is guaranteed to halt for some s^*. Then we output $m(d_1, d_2, p) = 2 + \max\{k(d, p+3) : d \in \mathbb{D}_s^* \cap [a, b]\}$. **Q.E.D.**

From this, we can conclude:

Theorem 20. *Let $f :\subseteq \mathbb{R} \to \mathbb{R}$ with recursively open* $\texttt{domain}(f)$. *Then f is computable iff f has a local modulus function and f is approximable.*

Exercises

Exercise 7.1:
(i) Give the proof of Theorem 20.
(ii) Show by a counter example that the hypothesis that $\texttt{domain}(f)$ is recursively open cannot be dropped. ♢

Exercise 7.2: Let $S \subseteq \mathbb{R}$. Suppose $f :\subseteq S \to \mathbb{R}$ is computable. Show that $\texttt{domain}(f)$ is relatively open in S. ♢

End Exercises

§8. Recursively Closed Sets

¶55. Let $f : [0,1] \to \mathbb{R}$ be a computable function with modulus function m and approximation $\widetilde{f}$. Let $y = y_0$ have the computable Cauchy representation $\overline{y}$. We aim to show that $f^{-1}(y_0)$ is recursively closed. The result is trivial if $f^{-1}(y_0)$ is empty. Hence we assume there exists an $x_0 \in [0,1]$ such that $f(x_0) = y_0$.

For $p \in \mathbb{N}$, define

$$y_* = y_*(p) := \max\left\{f(d)[p+2] : d \in \mathbb{D}_{m(p+2)}[0,1] \text{ and } f(d)[p+2] \leq \overline{y}[p+2]\right\}, \tag{17}$$

$$y^* = y^*(p) := \min\left\{f(d)[p+2] : d \in \mathbb{D}_{m(p+2)}[0,1] \text{ and } f(d)[p+2] \geq \overline{y}[p+2]\right\}. \tag{18}$$

The definitions of y_* and y^* involve maximizing and minimizing over sets that may be empty. When the corresponding set is empty, y^* (resp., y_*) is defined to be $+\infty$ (resp., $-\infty$).

Lemma 21.
(i) Consider the set $Y = \left\{f(d)[p+2] : d \in \mathbb{D}_{m(p+2)}[0,1]\right\}$. If $y, z \in Y$, $y < z$ then $z - y > 3 \cdot 2^{-p-2}$. Then there exists $w \in Y$ such that $y < w < z$.
(ii) Suppose there is an $x \in [0,1]$ such that $f(x) < y_0 - 2^{-p+1}$ (resp., $f(x) > y_0 + 2^{-p+1}$). Then y_* (resp., y^*) is finite and $y_* \geq y_0 - 2^{-p}$ (resp., $y^* \leq y_0 + 2^{-p}$).

Proof. (i) This follows from the fact that if $d_1, d_2 \in \mathbb{D}_{m(p+2)}[0,1]$ and $d_1 = d_2 \pm 2^{-m(p+2)}$ then $\widetilde{f}(d_1) = \widetilde{f}(d_2) \pm 3 \cdot 2^{-p-2}$:

$$\widetilde{f}(d_1) = f(d_1) \pm 2^{-p-2} = f(d_2) \pm 2 \cdot 2^{-p-2} = \widetilde{f}(d_2) \pm 3 \cdot 2^{-p-2}.$$

(ii) Let $d = \lfloor x \rfloor_{m(p+2)}$ and $d_0 = \lfloor x_0 \rfloor_{m(p+2)}$. We assume $f(x) < y_0 - 2^{-p+1}$ (the case $f(x) > y_0 + 2^{-p+1}$ is symmetric). If $f(d_0)[p+2] \leq \overline{y}[p+2]$ then it is immediate that $y_* \geq f(d_0)[p+2] > y_0 - 2^{-p}$. Hence we may now assume $f(d_0)[p+2] > \overline{y}[p+2]$.

Note that $f(d)[p+2] \leq \overline{y}[p+2]$ since

$$\begin{aligned} f(d)[p+2] &\leq f(d) + 2^{-p-2} \\ &\leq f(x) + 2^{-p-1} \\ &\leq y_0 - 3 \cdot 2^{-p-1} \\ &< \overline{y}[p+2]. \end{aligned}$$

This shows that y_* is finite. But from $f(d_0)[p+2] > y[p+2] \geq f(d)[p+2]$ and part (i), we know that there is a $w \in Y$ such that $f(d_0)[p+2] > y[p+2] \geq w$ and $f(d_0)[p+2] - w \leq 3 \cdot 2^{-p-2}$. Hence $f(d_0)[p+2] \leq y_* + 3 \cdot 2^{-p-2}$. Finally, we derive $y_0 \leq f(d_2)[p+2] + 2^{-p-2} \leq (y_* + 3 \cdot 2^{-p-2}) + 2^{-p-2} = y_* + 2^{-p}$. **Q.E.D.** .

¶56. Inverse of a computable value is recursively closed. Using y^*, y_*, define the sets

$$A_p := \left\{d \in \mathbb{D}_{m(p+2)}[0,1] : f(d)[p+2] > y^* + 2^{-p}\right\}, \tag{19}$$

$$B_p := \left\{d \in \mathbb{D}_{m(p+2)}[0,1] : f(d)[p+2] < y_* - 2^{-p}\right\}. \tag{20}$$

For $d \in A_p \cup B_p$, define the interval $I_d := (d - 2^{-m(p+2)}, d + 2^{-m(p+2)})$.

Theorem 22. *If $f : [0,1] \to \mathbb{R}$ is computable and y_0 is a computable real then the set $f^{-1}(y_0)$ is recursively closed. In fact, $f^{-1}(y_0) = [0,1] \setminus S$ where*

$$S := \bigcup_{p \in \mathbb{N}} \bigcup_{d \in A_p \cup B_p} I_d$$

is recursively open.

Proof. We claim that S is recursively open: for each $p = 0, 1, 2, \ldots$, and each $d \in \mathbb{D}_{m(p+2)}[0,1]$, we can decide if $d \in A_p \cup B_p$. This is because we can compute $y^*(p)$ and $y_*(p)$, and can thus decide whether $f(d)[p+2] > y^* + 2^{-p}$ or $f(d)[p+2] < y_* - 2^{-p}$. If $d \in A_p \cup B_p$, and if d is the ith value so determined by this procedure, we define ϕ by the equation $I_d = (\phi(2i), \phi(2i+1))$. Thus ϕ is a recursive interval representation of S, establishing our claim.

First we show that $f^{-1}(y_0) \subseteq [0,1] \setminus S$. If $x \in I_d$ where $d \in B_p$, then

$$\begin{array}{rcll} f(x) & \leq & f(d) + 2^{-p-2} & (|x-d| \leq 2^{-m(p+2)}) \\ & \leq & f(d)[p+2] + 2^{-p-1} & (f(d) \leq f(d)[p+2] + 2^{-p-2}) \\ & < & y_* - 2^{-p-1} & (f(d)[p+2] < y_* - 2^{-p}) \\ & < & y_0 & (y_* \leq \overline{y}[p+2] \leq y_0 + 2^{-p-2}). \end{array}$$

Similarly, if $x \in I_d$ where $d \in A_p$, then $f(x) > y_0$. This proves that $I_d \cap f^{-1}(y_0) = \emptyset$, and so $f^{-1}(y_0) \subseteq [0,1] \setminus S$.

We now show that $[0,1] \setminus S \subseteq f^{-1}(y_0)$. This is equivalent, $[0,1] \setminus f^{-1}(y_0) \subseteq S$. Suppose $x \in [0,1] \setminus f^{-1}(y_0)$. Choose $p \in \mathbb{N}$ such that either $f(x) < y_0 - 5 \cdot 2^{-p-1}$ or $f(x) > y_0 + 5 \cdot 2^{-p-1}$. By symmetry, it suffices to consider the case

$$f(x) < y_0 - 5 \cdot 2^{-p-1}. \tag{21}$$

Let $d := \lfloor x \rfloor_{m(p+2)}$. We show that $d \in B_p$:

$$\begin{array}{rcll} y_* - 2^{-p} & \geq & (y_0 - 2^{-p}) - 2^{-p} & (\textit{Lemma } 21(ii)) \\ & > & f(x) + 2^{-p-1} & (y_0 - 5 \cdot 2^{-p+1} > f(x)) \\ & \geq & f(d) + 2^{-p-2} & (f(x) \geq f(d) - 2^{-p-2}) \\ & \geq & f(d)[p+2]. & \end{array}$$

Furthermore, $|d - x| < 2^{-m(p+2)}$, so that $x \in I_d$. Thus $x \in S$, as desired. **Q.E.D.**

Exercises

Exercise 8.1: Let y_0 be a computable real and $T \subseteq [0,1]$ a recursively closed set. Construct a real function $f : [0,1] \to \mathbb{R}$ such that $T = f^{-1}(y_0)$. ♢

End Exercises

§9. Linear Approximations

¶57. Let $f : [0,1] \to \mathbb{R}$. Sometimes, it is easier to approximate f by constructing a sequence of functions that converges to f.

A sequence $(f_n : n \in \mathbb{N})$ of real functions is said to **rapidly converge** to f if for each n, we have $\texttt{domain}(f_n) = \texttt{domain}(f)$ and $f(x) = f_n(x) \pm 2^{-n}$.

Assume each f_n in the sequence $(f_n : n \in \mathbb{N})$ is piecewise linear, with breakpoints contained in $\mathbb{D}$, and let $f_n(\mathbb{D}) \subseteq \mathbb{D}$. It follows that for all $d \in \mathbb{D}$, $f_n(d) \in \mathbb{D}$. Hence the sequence $(f_n : n \in \mathbb{N})$ may be represented by a function

$$F : \mathbb{N} \times \mathbb{D}[0,1] \to \mathbb{D}$$

where $F(n,d) = f_n(d)$. Call F a **uniform linear approximation** of f in case F represents a sequence $(f_n : n \in \mathbb{N})$ that rapidly converges to f.

¶58.

Theorem 23. *A function $f : [0,1] \to \mathbb{R}$ is computable iff f has a computable modulus function m and a computable uniform linear approximation F.*

Proof. In view of the characterization of Theorem 15, we may assume that f has a computable modulus m. We only have to show that the existence of of a computable approximation $\widetilde{f} : \mathbb{D} \times \mathbb{N} \to \mathbb{D}$ is equivalent to the existence of a computable uniform linear approximation F.

Given a computable F, we define $\widetilde{f}$ by $f(d)[n] = F(n,d)$. Clearly, $\widetilde{f}$ is computable and is an approximation of f.

Conversely, given $\widetilde{f}$, we define $f_n(d) = f(d)[n+3]$ for all $d \in \mathbb{D}_{m(n+3)}$, and $f_n(x)$ is obtained by linear interpolation if $x \notin \mathbb{D}_{m(n+3)}$. Clearly, $F(n,d) = f_n(x)$ is computable.

To show that $(f_n : n \in \mathbb{N})$ rapidly converges to f, recall the generalized floor $\lfloor x \rfloor_n$ and ceiling $\lceil x \rceil_n$ functions. Note that $f_n(x)$ lies between $f_n(\lfloor x \rfloor_{m(n+3)})$ and $f_n(\lceil x \rceil_{m(n+3)})$. Moreover,

$$\begin{aligned} &|f_n(\lceil x \rceil_{m(n+3)}) - f_n(\lfloor x \rfloor_{m(n+3)})| \\ &\quad \leq |f_n(\lceil x \rceil_{m(n+3)}) - f(\lceil x \rceil_{m(n+3)})| + |f(\lceil x \rceil_{m(n+3)}) - f(\lfloor x \rfloor_{m(n+3)})| + |f(\lfloor x \rfloor_{m(n+3)}) - f_n(\lfloor x \rfloor_{m(n+3)})| \\ &\quad \leq 3 \cdot 2^{-n-3} < 2^{-n-1} \end{aligned}$$

Hence $f_n(x) = f_n(\lfloor x \rfloor_{m(n+3)}) \pm 2^{-n-1}$. Finally,

$$\begin{aligned} |f(x) - f_n(x)| &\leq |f(x) - f(\lfloor x \rfloor_{m(n+3)})| + |f(\lfloor x \rfloor_{m(n+3)}) - f_n(\lfloor x \rfloor_{m(n+3)})| + |f_n(\lfloor x \rfloor_{m(n+3)}) - f_n(x)| \\ &\leq 2^{n-3} + 2^{-n-1} + 2^{n-3} < 2^{-n}. \end{aligned}$$

Q.E.D.

§10. Numerical Functionals and Operators

¶59. Higher-type Computations. Consider the following problems:

- Evaluation Problem:
$$EVAL : C[0,1] \times [0,1] \to \mathbb{R}$$
where $EVAL(f,x) = f(x)$.
- Maximization Problem:
$$MAX : C[0,1] \to C[0,1]$$
where $MAX(f)$ is the function $g \in C[0,1]$ given by $g(x) = \max\{f(y) : 0 \leq y \leq x\}$. We can specialize MAX by defining
$$MAX_1 : C[0,1] \to \mathbb{R}$$
where $MAX_1(f) = \max\{f(y) : 0 \leq y \leq 1\}$. Alternatively, $MAX_1(f) = MAX(f)(1)$.
- Integration Problem:
$$INT \ :\subseteq \ C[0,1] \to C[0,1]$$
where $INT(f)$ is the function $g \in C[0,1]$ given by $g(x) = \int_0^x f$. We can specialize INT by defining
$$INT_1 \ :\subseteq \ C[0,1] \to \mathbb{R}$$
where $INT_1(f) = \int_0^1 f$.

A **(numerical) functional** F is a function from real functions to $\mathbb{R}$. A **(numerical) operator** H is a function from real functions to real functions.

Thus, $EVAL, INT_1, MAX_1$ are functionals, while MAX and INT are operators. A functional F belongs to $[\mathbb{R} \Rightarrow \mathbb{R}] \Rightarrow \mathbb{R}$, and an operator H belongs $[\mathbb{R} \Rightarrow \mathbb{R}] \Rightarrow [\mathbb{R} \Rightarrow \mathbb{R}]$. The input of functionals and operators include real functions (type-2 objects). Thus they are type-3 objects by our classification.

¶60. How should we compute such objects? We need to introduce suitable input and output representations.

- Let $f : [0,1] \to \mathbb{R}$, viewed as input to an OTM for computing a functional F or operator H. We say f is **represented** by a pair $(m, \widetilde{f})$ where
$$m : \mathbb{D}[0,1] \times \mathbb{N} \to \mathbb{D}$$
is a modulus function for f and
$$\widetilde{f} : \mathbb{D}[0,1] \times \mathbb{N} \to \mathbb{D}$$
is an absolute approximation (¶21) for f. If m and f are both recursive, then we call $(m, \widetilde{f})$ a **recursive representation** of f. This definition is justified by Theorem 15.

- For functional F, the output is $\mathbb{R}$. We use the same trick as for computing functions: instead of the OTM outputing a representation $\overline{x}$ of real numbers (i.e., a Cauchy representation $\overline{x} : \mathbb{N} \to \mathbb{D}$), we incorporate the precision argument p of $\overline{x}$ into the OTM, so that the OTM only outputs $\overline{x}[p] \in \mathbb{D}$ (cf. ¶20).

  We say an OTM M computes the functional F if, on oracle inputs $m, \widetilde{f}$ and precision input p, computes a dyadic number
$$M^{m,\widetilde{f}}[p]$$
such that $|M^{m,\widetilde{f}}[p] - F(f)| \leq 2^{-p}$.

- For operator H, instead of the OTM outputing a representation of a function $g : [0,1] \to \mathbb{R}$, we incorporate the argument x for g and the precision desired for $g(x)$ into the input of the OTM. We say an OTM M computes the operator H if, on oracle inputs $m, \widetilde{f}, \overline{x}$ and precision input p, computes a dyadic number
$$M^{m,\widetilde{f},\overline{x}}[p]$$
such that $|M^{m,\widetilde{f},\overline{x}}[p] - I(f)(x)| \leq 2^{-p}$.

The technique used to define computation of functionals and operators amounts to reducing the type of the range: by moving arguments from the range into the domain of functions. For instance, instead of computing $f \in [A \Rightarrow [B \Rightarrow C]]$, we compute $f' \in [[A \times B] \Rightarrow C]$. This is called "Currying" (after the logician Curry).

¶61. Non-polynomial computability of INT_1 Let $F : D \to \mathbb{R}$ be a functional for some $D \subseteq C[0,1]$. We say F is **polynomial time** if there is an OTM M and polynomials q, q' such that for all representations $(m, \widetilde{f})$ of $f \in D$, $M^{m,\widetilde{f}}[n]$ runs in time polynomial in n.

We will now prove that the INT_1 operator cannot be computed in polynomial time.

Suppose OTM M computes INT_1 in polynomial time $t(n)$. Choose n sufficiently large so that $t(n) < 2^{n-3}$. We now run $M^{m,\widetilde{f}}[n]$ for suitably chosen m and $\widetilde{f}$. The modulus function is the constant 1: $m(n) = 1$ for all n. First, let $\widetilde{f}$ be the constant 0: $\widetilde{f}(d; n) = 0$ for all d, n. Let $M^{m,\widetilde{f}}[n]$ query $\widetilde{f}$ on the $t \leq t(n)$ inputs
$$(d_1, n_1), \ldots, (d_t, n_t).$$
Wlog, assume $d_1 < d_2 < \cdots < d_t$. For each query, the oracle returns 0, $f(d_i)[k_i] = 0$.

Next, we replace $\widetilde{f}$ by another function $\widetilde{g}$ which is an approximation to the piecewise linear g defined as follows: the break points of g are $d_1, \ldots, d_t$ and also $(d_i + d_{i+1})/2$ for $i = 1, \ldots, t-1$. Then g is determined by these values: $g(d_i) = 0$ for $i = 1, \ldots, t$, and $g((d_i + d_{i+1})/2) = (d_{i+1} - d_i)/2$. See Figure 2 for an illustration. Our approximation $\widetilde{g}$ also has the property that $g(d_i)[n] = 0$ for all n ($i = 1, \ldots, t$). It follows that $M^{m,\widetilde{f}}[n]$ and $M^{m,\widetilde{g}}[n]$ has identical behavior and must output the same answer.

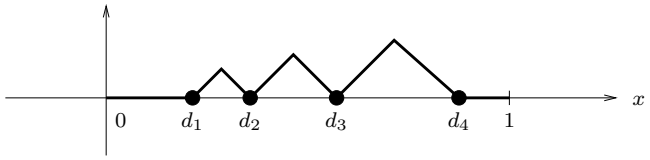

Figure 2: Function g to fool the integrator OTM

We note that

$$\begin{aligned}\int_0^1 g &= \sum_{i=1}^{t-1}\left(\tfrac{d_{i+1}-d_i}{2}\right)^2\\ &\geq \tfrac{1}{t}\left(\sum_{i=1}^{t-1}\tfrac{d_{i+1}-d_i}{2}\right)^2 \quad (Schwartz'sinequality)\\ &= \tfrac{1}{4t}.\end{aligned}$$

Since $t \leq t(n) < 2^{n-3}$ (by choice of n), we conclude that $\int_0^1 g > 2^{-n+1}$. If $d = M^{m,\widetilde{f}}[n] = M^{m,\widetilde{g}}[n]$, then $|d - \int_0^1 f|$ or $|d - \int_0^1 g| > 2^n$, contradiction.

Exercises

Exercise 10.1: Consider an alternative to representing functions. Suppose our representation of f does not include a modulus function. Instead, we represent f by a pair $(\widetilde{f}, \widetilde{f'})$ where $\widetilde{f}$ is an absolute approximation of f, and $\widetilde{f'}$ an absolute approximation of f'. We say an OTM M "computes" the $EVAL$ functional if, on oracle inputs $\widetilde{f}, \overline{x}$ and precision input p, the machine $M^{\widetilde{f},\overline{x}}[p]$ eventually outputs $EVAL(f,x) \pm 2^{-p}$. Prove or disprove: $EVAL$ cannot be computed in this sense. ♢

End Exercises

§11. Maximization of Functions

¶62. Maximum value and maximum points. For $f : [0,1] \to \mathbb{R}$, define $\max(f) := \max\{f(x) : x \in [0,1]\}$, and $\text{maxDom}(f) := \{x \in [0,1] : f(x) = \max(f)\}$.

Call $\max(f)$ the **maximum value** of f and each $x \in \text{maxDom}(f)$ is a **maximum point** of f.

We study the maximization of real functions. Finding maximum points efficiently is connected to the $P = NP$ question. Extending this analogy, there is a parallel between the polynomial-time hierarchy and problems defined by alternating maximization- and minimization- of real functions.

¶63. Computability of maximum value.

Lemma 24. *Let f be computable and $f \in C[0,1]$. Then $\max(f)$ is a computable real.*

Proof. Let m be a recursive modulus function and $\widetilde{f}$ a recursive approximation of f. We compute a p-bit approximation of $\max(f)$ as follows: we compute $n = m(p+1)$ and then compute

$$y^* = \max\{f(d)[p+1] : d \in \mathbb{D}_n\}.$$

Let x^* be a maximum point, i.e., $f(x^*) = \max(f)$. If $d^* = \lfloor x^* \rfloor_n$, then

$$\begin{aligned}y^* &\geq f(d^*)[p+1]\\ &\geq f(d^*) - 2^{-p-1}\\ &\geq (f(x^*) - 2^{-p-1}) - 2^{-p-1} = f(x^*) - 2^{-p}.\end{aligned}$$

If $y^* = f(d)[p+1]$ then $f(x^*) \geq f(d) \geq y^* - 2^{-p-1}$. Thus y^* is an p-bit approximation of $\max(f)$. **Q.E.D.**

Corollary 25. $\text{maxDom}(f)$ *is recursive closed.*

Proof. Since $y_0 = \max(f)$ is computable and $\text{maxDom}(f) = f^{-1}(y_0)$, this follows from Theorem 22. **Q.E.D.**

¶64. Uncomputable continuous functions.

Corollary 26. *Let $f : [0,1] \to \mathbb{R}$. If $\max(f)$ is an uncomputable real, then f is uncomputable function.*

This furnishes us with simple examples of continuous but uncomputable functions. For $\alpha \in \mathbb{R}$, define the piecewise linear function $f_\alpha : [0,1] \to \mathbb{R}$ with a breakpoint at $1/2$, and $f_\alpha(1/2) = \alpha$ and $f_\alpha(0) = f_\alpha(1) = 0$.

Corollary 27. *The function f_α is computable iff α is computable.*

¶65. Characterization of maximum points. The proof is similar to Theorem 28.

Theorem 28. *Let $T \subseteq [0,1]$. The following statements are equivalent:*
(1) T is recursively closed.
(2) $T = \text{maxDom}(f)$ where f is computable and $f \in C[0,1]$.

Proof. (2)⇒(1) is just Corollary 25. To show (1)(⇒)(2), suppose T is recursively closed. We may write $T = [0,1] \setminus \bigcup_{n\geq 0} I_n$, for some recursive $\phi : \mathbb{N} \to \mathbb{D}$ such that each $I_n = (\phi(2n), \phi(2n+1))$.

We construct a function f such that $T = \text{maxDom}(f)$. For each interval $I_n = (\phi(2n), \phi(2n+1)$, we define the piecewise linear function $g_n : [0,1] \to \mathbb{R}$ with breakpoints are $\phi(2n) + \phi(2n+1))/2$ and also the endpoints of $I_n \cap [0,1]$. Then g_n is completely specified if we define

$$g_n\left(\frac{\phi(2n) + \phi(2n+1))}{2}\right) = -1$$

and $g_n(x) = 0$ for $x \notin I_n$. Note that $g_n(x) \leq 0$ for all $x \in [0,1]$. In case $I_n \in [0,1]$, the graph of g_n has a "V" shape on the interval I_n.

Now define

$$f(x) := \sum_{n=0}^{\infty} 2^{-n} g_n(x).$$

Clearly, $\max(f) = 0$. We claim that $T = \text{maxDom}(f)$. To see this, if $x \in T$ then $f(x) = 0$, and so $T \subseteq \text{maxDom}(f)$. Conversely, if $x \notin T$, then $x \in I_n$ for some n and $g_n(x) < 0$. Thus $f(x) < 0$ and $x \notin \text{maxDom}(f)$.

It remains to prove that f is computable. Let $f_n(x) = \sum_{i=0}^{n} 2^{-i} g_i(x)$. Clearly, the sequence (f_n) rapidly converges to f since $|f(x) - f_n(x)| \leq \sum_{i\geq n+1} 2^{-i} \leq 2^{-n}$. The uniform linear approximation function $F(n,x) = f_n(x)$ is clearly computable. By our characterization (¶58) of computable functions by linear approximations, it remains to show that f has a computable modulus function m. We define $m(n)$ to be any k such that

$$|f_{n+2}(x) - f_{n+2}(y)| \leq 2^{-n-2}. \tag{22}$$

whenever $|x - y| \leq 2^{-k}$. This would imply

$$|f(x) - f(y)| \leq 3 \cdot 2^{-n-2} < 2^{-n},$$

so that $m(n) = k$ is a modulus function. Without loss of generality, assume $\phi(2i) \leq \phi(2i+1)$ for all i, and let $k' := -\lg \min\{\phi(2i+1) - \phi(2i) : i = 0, \ldots, n+2\}$. We choose $k = n + 2 + \max\{0, \lfloor k' \rfloor\}$ For $i \leq n+2$, we have $|g_i(x) - g_i(y)| \leq 2^{k'-1}|x - y|$. Hence if $|x - y| \leq 2^{-k}$, then $|g_i(x) - g_i(y) \leq 2^{-n-3}$. Thus $|f_{n+2}(x) - f_{n+2}(y)| \leq \sum_{i=0}^{n+2} 2^{-i} 2^{-n-3} \leq 2^{-n-2}$, satisfying (22). **Q.E.D.**

We can even assume in this result that f is polynomial-time computable (Exercise).

¶66. Uncomputable maximum points. This result is from Specker.

Lemma 29. *There is a computable function $f_0 : [0,1] \to \mathbb{R}$ such that* maxDom(f_0) *is an uncountable set of all uncomputable reals.*

Proof. By the previous theorem, it suffices to construct a closed set containing only uncomputable reals. Equivalently, we construct an open set $S \subseteq [0,1]$ which contains all computable reals in $[0,1]$.

Let $\phi_0, \phi_1, \ldots$ be an enumeration of the partial recursive functions. Define the open interval

$$I_i := (\phi_i(i+4) - 2^{-i-3}, \phi_i(i+4) + 2^{-i-3})$$

provided $\phi_i(i+4)$ represents a dyadic number in $[0,1]$. The set $S := \bigcup\{I_i : \phi_i(i+4) \in \mathbb{D}\}$ is recursively open since we can dovetail all the computations of $\phi_i(i+4)$ $(i \in \mathbb{N})$, and whenever $\phi_i(i+4)$ halts and returns an element of $\mathbb{D}[0,1]$, we output the corresponding I_i.

Note that if $x \in \mathbb{R}$ is recursive, then there is some i such that ϕ_i computes a Cauchy function of x. This means that $|\phi_i(i+4) - x| \leq 2^{-i-4}$ and hence $x \in I_i$. Thus $x \in S$. Finally, to ensure that S is a proper subset of $[0,1]$, we see that the measure of S is at most $\sum_{i\geq 0} 2^{-i-2} = 1/2$. This concludes the construction. **Q.E.D.**

¶67. Computability of the MAX operator. By examining the proof of Lemma 24, we see that it also shows that the operator $MAX_1 : C[0,1] \to \mathbb{R}$ is computable. Recall (¶60) that it means there is an OTM M taking a pair of oracles $(m, \widetilde{f})$ representing $f \in C[0,1]$ and a precision $p \in \mathbb{N}$, and the output is a p-bit approximation of $MAX_1(f)$. We can further generalize this argument:

Lemma 30. *The functional $MAX : C[0,1] \to C[0,1]$ is computable.*

Proof. We describe an OTM M that takes as input a pair $(m, \widetilde{f})$ of modulus and approximation that represents $f \in C[0,1]$. It takes a third oracle $\overline{x}$ representing $x \in \mathbb{R}$, and finally a precision $p \in \mathbb{N}$. We want to output a value $MAX(f)(x)[p]$: We first compute $n = m(p+1)$. Then we compute the set

$$\{f(d)[p+1] : d \in \mathbb{D}_n, d[p+1] > x[p+2]\}.$$

Let y^* be the maximum of this set. ... **Q.E.D.**

¶68. Smooth bump function. Consider the "sigmoid" function

$$h(x) := \begin{cases} e^{-1/x^2} & \text{if } x > 0 \\ 0 & \text{if } x \leq 0. \end{cases}$$

Note that $0 \leq h(x) \leq 1$ and $h'(x) = e^{-1/x^2}(2x^{-3}) > 0$ for $x > 0$. Then

$$\begin{aligned} h'(x) &= \frac{2}{x^3} \cdot \frac{1}{1 + x^{-2} + x^{-4}2! + x^{-6}3! + \cdots} \\ &= \frac{2}{x^3 + x + x^{-1}2 + x^{-3}6 + \cdots} \end{aligned}$$

and thus $h'(x) \to 0$ as $x \to 0^+$. Hence $h'(0) = 0$. We can similarly show that $h^{(n)}(0) = 0$ for all $n \geq 2$.

Figure 3: Sigmoid function h

This function is shown in Figure 3. We now modify $h(x)$ to $H(x)$ to create a more pronounced step-like function:

$$H(x) := \frac{h(x)}{h(1-x) + h(x)}$$

Observe that

$$H(x) = \begin{cases} 0 & \text{if } x \leq 1, \\ 1 & \text{if } x \geq 1, \\ \frac{1}{2} & \text{if } x = \frac{1}{2}, \\ \text{strictly increasing} & 0 < x < 1. \end{cases}$$

Finally, we create a "bump function"

$$B(x) := \begin{cases} H(2x) & x \leq 1/2, \\ H(2-2x) & x \geq 1/2. \end{cases}$$

Observe that

$$B(x) = \begin{cases} 0 & x \leq 0 \text{ or } x \geq 1, \\ 1 & x = 1/2, \\ \text{strictly increasing} & 0 < x < 1/2, \\ \text{strictly decreasing} & 1/2 < x < 1. \end{cases}$$

For all $n \geq 1$, $H^{(n)}(0) = H^{(n)}(1) = 0$.

FIGURE: functions h and H and B

¶69. A smooth function f_1 such that $MAX(f_1)$ is polynomial time iff $P = NP$. Suppose $A \in NP$. Assume that $w \in A$ iff $(\exists y)[|y| \leq p(|w|) \ \wedge \ R(y, w)]$ where $p(n)$ is a polynomial function and $R(y, w)$ is a polynomial-time predicate. We construct f_1 as follows.

We will partition $[0, 1)$ into countably many pairwise disjoint intervals of the form $I_s = [a_s, b_s)$ where $s \in \{0, 1\}^*$. The width $w(I_s)$ of the interval is $2^{-1-2|s|}$, and the $2^{|s|}$ strings of length $|s|$ will have their intervals laid out consecutively, followed by all the strings of length $|s| + 1$, etc. Thus the total width of all intervals with width $2^{-1-2|s|}$ is $2^{-1-|s|}$.

FIGURE: Layout of $[0, 1)$

The half-interval $[a_s, (a_s + b_s)/2]$ is subdivided into $2^{-p(|s|)}$ subintervals, each corresponding to a string of length $p(|s|)$.

INCOMPLETE.

Exercises

Exercise 11.1: Show that in Theorem 28, the function f can be assumed to be polynomial-time. ♢

Exercise 11.2: Here is a simpler proof of Lemma 24: Let $f \in C[0, 1]$ be computable by an OTM M. Then $\max(f) = f(x^*)$ for some $x^* \in \mathbb{R}$. Hence $M^{\overline{x}^*}[p]$ yields a p-bit absolute approximation of $\max(f)$. Since p is arbitrary, we have shown the approximability of $\max(f)$. Why is the error? ♢

Exercise 11.3: State and prove the analogue of Lemma 24 for partial functions $f :\subseteq \mathbb{R} \to \mathbb{R}$. ♢

Exercise 11.4: Show that f_α is computable iff α is computable. ♢

Exercise 11.5: Show that the function f in ¶65 may be assumed to be polynomial-time computable. HINT: Define $f(x) := \sum_{n \geq \infty} g_n(x) 2^{-t(n)}$ where $t(n)$ bounds the number of steps it takes to compute $\phi(i)$ for $i = 0, \ldots, 2n + 1$. ♢

End Exercises

§12. Roots of Functions

¶70. Zero sets of computable functions are recursively closed. The **zero set** of a function $f :\subseteq \mathbb{R} \to \mathbb{R}$ is $f^{-1}(0) = \{x \in \mathbb{R} : f(x) = 0\}$, denoted $\text{Zero}(f)$. Elements of $\text{Zero}(f)$ is called a **zero** of f.

We have the another characterization of recursively closed subsets of $[0, 1]$, with basically the same proof as Theorem 28.

Theorem 31. *Let $T \subseteq [0, 1]$. The following statements are equivalent:*
(1) T is recursively closed.
(2) $T = \text{Zero}(f)$ where f is computable and $f \in C[0, 1]$.

Proof. (2)⇒(1): this follows from Theorem 22 by choosing $y_0 = 0$.
(1)⇒(2): the function f constructed in Theorem 28 has the property that $\text{Zero}(f) = \text{maxDom}(f)$. **Q.E.D.**

¶71. Isolated Zeros. Zero sets (like recursively closed sets) of computable functions can be very complex: a zero set may be an uncountable set of uncomputable reals (¶66). The zeros in this example involve non-isolated zeros. Nice functions (e.g., analytic functions) have only isolated zeros. The following result speaks to such functions:

Lemma 32. *Let $f : [0, 1] \to \mathbb{R}$ be computable. Then isolated zeros of f are computable.*

Proof. Let x_0 be an isolated zero of f. If $x_0 \in \mathbb{D}$, then it is clearly computable. Assuming otherwise, we show how to compute an n-bit approximation of x_0: we may assume that $\text{Zero}(f) \cap [a, b] = \{x_0\}$ for some $a, b \in \mathbb{D}[0, 1]$. Then, for each $d \in \mathbb{D}_n[a, b]$, we compute $f(d)[p]$ for $p = 0, 1, 2, \ldots$, until $|f(d)[p]| > 2^{-p}$. When this happens, we know the sign of $f(d)$, since $\texttt{sign}(f(d)) = \texttt{sign}(f(d)[p])$. We also assume that the sign of $f(a)$ and $f(b)$ is computed in this way. Since x_0 is unique zero in $[a, b]$, there is a unique $d_1, d_2 \in \mathbb{D}_n[a, b] \cup \{a, b\}$ such that $f(d_1)f(d_2) < 0$ and $|d_1 - d_2| \leq 2^{-n}$. Now either d_1 or d_2 can be used as an n-bit approximation of x_0. **Q.E.D.**

¶72. Complexity of isolated zeros. Although isolated zeros are computable, they can have arbitrarily high complexity:

Lemma 33. *Let $x_0 \in [0,1]$ be a computable real. Then there is a strictly increasing computable function $f_0 : [0,1] \to \mathbb{R}$ such that* $\text{Zero}(f_0) = \{x_0\}$.

Proof. Suppose $\overline{x} : \mathbb{N} \to \mathbb{D}$ is a computable approximation of x_0. Without loss of generality, assume $a, b \in \mathbb{D}$ such as $0 < a < x_0 < b < 1$. We construct a uniform linear approximation $F : \mathbb{N} \times \mathbb{D}[0,1] \to \mathbb{D}$ for f_0 as follows: for each $n \in \mathbb{N}$ and $d \in \mathbb{D}[0,1]$, define

$$x_*(n) = \max\left\{a, \overline{x}[i] - 2^{-i} : i = 0, \ldots, n\right\}.$$

and

$$x^*(n) = \min\left\{b, \overline{x}[i] + 2^{-i} : i = 0, \ldots, n\right\}.$$

Thus,

$$x_*(n) \le x_*(n+1) \le x_0 \le x^*(n+1) \le x^*(n)$$

for all n. We define $g_n : [0,1] \to \mathbb{R}$ to be the piecewise linear function with breakpoints at $x_*(n)$ and $x^*(n)$. For $d \in \{0, 1, x_*(n), x^*(n)\}$, we have

$$g_n(d) := \begin{cases} -2^{-n} & \text{if } d = 0, \\ 0 & \text{if } d = x_*(n) \text{ or } d = x^*(n), \\ +2^{-n} & \text{if } d = 1. \end{cases}$$

SEE FIGURE ILLUSTRATING g_n

Note that for each t and for i large enough ($i > i(t)$) we have

$$g_i(t) \begin{cases} = 0 & \text{if } t = x_0, \\ < 0 & \text{if } t < x_0, \\ > 0 & \text{if } t > x_0. \end{cases}$$

Now define the function F_n where $F_n(t) = \sum_{i=0}^{n} g_i(t)$. Note that as $n \to \infty$, $F_n(t)$ converges to

$$f_0(t) = \sum_{i=0}^{\infty} g_i(t) \begin{cases} = 0 & \text{if } t = x_0, \\ < 0 & \text{if } t < x_0, \\ > 0 & \text{if } t > x_0. \end{cases}$$

Thus $\text{Zero}(f_0) = \{x_0\}$.

It is easy to see that each g_n is strictly increasing, and hence f_0 is strictly increasing.

The $\{F_n(d) = F(n,d) : n \in \mathbb{N}\}$ rapidly converges to f_0 since for all $t \in [0,1]$,

$$|F_n(t) - f_0(t)| \le \sum_{i=n+1}^{\infty} 2^{-i} \le 2^{-n}.$$

Clearly, the uniform linear approximation function $F(n,d) := F_n(d)$ is clearly recursive.

Finally, to show that f_0 is computable, it remains to show that f_0 has a recursive modulus function. Note that the slope of each g_i is non-negative but at most

$$\frac{1}{2^i \min\{a, 1-b\}} \le 2^{-i-k_0}$$

for some integer constant k_0. That means that if $|x - y| \le 2^{-n}$ then $|g_i(x) - g_i(y)| \le 2^{-n-i-k_0}$. Summing, $|f_0(x) - f_0(y)| \le \sum_i^n 2^{-n-i-k_0} < 2^{-n-k_0+2}$. Hence, we may define $m(n) = n + k_0 + 2$. This is clearly recursive. **Q.E.D.**

§13. Derivatives

¶73. Computability of Derivatives

Lemma 34 (Pour-El and Richards). *Let $f : [0,1] \to \mathbb{R}$ be computable and f' is defined and continuous. Then f' is computable iff f' has a recursive modulus.*

Proof. One direction is immediate from a characterization of computable functions. In the other direction, suppose f' has a recursive modulus m. To show that f' is computable, it is sufficient to show that it is approximable. Given $d \in \mathbb{D}$ and $p \in \mathbb{N}$, we compute a p-bit approximation $f'(d)[p]$ as the following value

$$y = (f(d')[m(p+1)+p+1] - f(d)[m(p+1)+p+1])2^{m(p+1)} \tag{23}$$

where $d' = d + 2^{-m(p+1)}$. To show that y is a correct value, note that the mean value theorem implies that $f'(x) = (f(d') - f(d))2^{m(p-1)}$ for some $x \in [d, d']$. Hence

$$\begin{aligned} f'(d) &= f'(x) \pm 2^{-p-1} \\ &= (f(d') - f(d))2^{m(p+1)} \pm 2^{-p-1} \\ &= \left(\widetilde{f}(d') \pm 2^{-m(p+1)-p-1} - \widetilde{f}(d) \pm 2^{-m(p+1)-p-1}\right) 2^{m(p+1)} \pm 2^{-p-1} \\ &= \left(\widetilde{f}(d') - \widetilde{f}(d)\right) 2^{m(p+1)} \pm 2^{-p-1} \pm 2^{-p-1} \\ &= y \pm 2^{-p}. \end{aligned}$$

Q.E.D.

Ko observed that an analogous lemma is true for polynomial-time computable derivatives:

Lemma 35. *Let $f : [0,1] \to \mathbb{R}$ be polynomial-time computable and f' is defined and continuous. Then f' is polynomial-time computable iff f' has a polynomial-time recursive modulus.*

Proof. We just note that in the previous proof, the value (23) can be computed in polynomial time, since $m(p)$ is polynomial time computable. **Q.E.D.**

§14. Recognition of Sets

¶74. Approximation of predicates. Let $S \subseteq \mathbb{R}^k$ be any set. What does it mean to "compute" S? The standard approach is identify this with the deciding membership in S, i.e., with computability of the **characteristic function** $\chi_S : \mathbb{R}^k \to \{0,1\}$ where $\chi_S(\mathbf{x}) = 1$ iff $\mathbf{x} \in S$. Unfortunately, this means S is uncomputable except for the trivial cases of $S = \emptyset$ or $S = \mathbb{R}^k$. This issue was discussed at length above (¶32) since χ_S is a predicate. One way to circumvent this is to introduce measures on sets. Then we can say that S is decidable (in measure) if there is an OTM M that correctly decides membership in S except for an input set of measure 0. But measure theory may be an overkill for some applications. Here we use an elementary approach, based on metrics $\|\cdot\|$ on $\mathbb{R}^k$ (see Braverman [2]).

To be specific, assume the Euclidean metric $\|\mathbf{x}\| = \sqrt{\sum_{i=1}^k x_i^2}$ where $\mathbf{x} = (x_1, \ldots, x_k)$. Let $B(\mathbf{d}, n)$ denote the corresponding closed Euclidean ball centered at $\mathbf{d} \in \mathbb{R}^k$ with radius 2^{-n}.

Let $f : \mathbb{R}^k \to \{0,1\}$ be a predicate. A function $\widetilde{f} : \mathbb{D}^k \times \mathbb{D} \to \{0,1\}$ is called a **positive approximation** of f if, for all $\mathbf{d} \in \mathbb{D}^k$ and $p \in \mathbb{D}$, we have

$$\widetilde{f}(\mathbf{d}; p) = \begin{cases} 1 & \text{if } f(\mathbf{d}') = 1 \text{ for some } \mathbf{d}' \text{ where } \|\mathbf{d} - \mathbf{d}'\| \leq 2^{-p} \\ 0 & \text{if } f(\mathbf{d}') = 0 \text{ for all } \mathbf{d}' \text{ where } \|\mathbf{d} - \mathbf{d}'\| \leq 2^{1-p} \end{cases} \tag{24}$$

Note that the two clauses in this definition is not exhaustive. In cases not covered by the clauses, $\widetilde{f}(\mathbf{d}; p)$ may be 0 or 1 (we do not care which).

We say $\widetilde{f}$ above is a **negative approximation** of f if, for all $\mathbf{d} \in \mathbb{D}^k$ and $p \in \mathbb{D}$, we have

$$\widetilde{f}(\mathbf{d}; p) = \begin{cases} 0 & \text{if } f(\mathbf{d}') = 0 \text{ for all } \mathbf{d}' \text{ where } \|\mathbf{d} - \mathbf{d}'\| \leq 2^{-p} \\ 1 & \text{if } f(\mathbf{d}') = 1 \text{ for some } \mathbf{d}' \text{ where } \|\mathbf{d} - \mathbf{d}'\| \leq 2^{1-p} \end{cases} \tag{25}$$

As usual, we write $f(\mathbf{d})[p]$ for $\widetilde{f}(\mathbf{d}; p)$.

A set $S \subseteq \mathbb{R}^k$ is **positively recognizable** if its characteristic function χ_S is positively approximable. If f is the characteristic function of a set S, then $\widetilde{f}$ is also called a **positive approximation** of S. Thus, (24) becomes

$$\chi_S(\mathbf{d})[p] = \widetilde{f}(\mathbf{d};p) = \begin{cases} 1 & \text{if } B(\mathbf{d};p) \cap S \neq \emptyset \\ 0 & \text{if } B(\mathbf{d};p-1) \cap S = \emptyset \end{cases}$$

As before, these two clauses is incomplete. To express the requirements in a "complete form", we may write

$$\begin{aligned} \chi_S(\mathbf{d})[p] = 1 &\Rightarrow B(\mathbf{d};p) \cap S \neq \emptyset \\ \chi_S(\mathbf{d})[p] = 0 &\Rightarrow B(\mathbf{d};p-1) \cap S = \emptyset \end{aligned}$$

There is an analogous definition for **negative recognizability**:

$$\chi_S(\mathbf{d})[p] = \widetilde{f}(\mathbf{d};p) = \begin{cases} 0 & \text{if } B(\mathbf{d};p) \cap S = \emptyset \\ 1 & \text{if } B(\mathbf{d};p-1) \cap S \neq \emptyset \end{cases}$$

Finally, we say S is **recognizable** if it is both positively and negatively recognizable.

¶75. Distance function. We shall study recognizability via the distance function. The **distance function** of S is $d_S : \mathbb{R}^k \to \mathbb{R} \cup \{\infty\}$ where $d_S(\mathbf{x}) = \inf\{\|\mathbf{x}-\mathbf{y}\| : \mathbf{y} \in S\}$. In case $S = \emptyset$, we have $d_S(\mathbf{x}) = \infty$ for all $\mathbf{x}$. However, to avoid this trivial case, from now on, we assume S is non-empty. The next lemma shows that d_S is continuous in a very strong sense:

Lemma 36. *The distance function d_S has Lipschitz constant* 1*, i.e., for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^k$,*

$$|d_S(\mathbf{x}) - d_S(\mathbf{y})| \leq \|\mathbf{x} - \mathbf{y}\|.$$

. *Proof.* Wlog, let $d_S(\mathbf{x}) \geq d_S(\mathbf{y})$ and $d_S(\mathbf{y}) = \|\mathbf{y} - \mathbf{z}\|$ for some $\mathbf{z} \in S$. Then

$$\begin{aligned} |d_S(\mathbf{x}) - d_S(\mathbf{y})| &= d_S(\mathbf{x}) - d_S(\mathbf{y}) \\ &\leq \|\mathbf{x} - \mathbf{z}\| - \|\mathbf{y} - \mathbf{z}\| \\ &\leq \|\mathbf{x} - \mathbf{y}\|. \end{aligned}$$

Q.E.D.

Corollary 37. *d_S is computable iff d_S is approximable.*

Proof. The Lipschitz constant of 1 for d_S implies that d_S has a recursive modulus function, namely $m(n) = n$. The result follows from the characterization of computable functions, Theorem 15. **Q.E.D.**

Theorem 38. *Let $S \subseteq \mathbb{R}^k$ be non-empty. Then $S \subseteq \mathbb{R}^k$ is positively recognizable iff its distance function d_S is computable.*

Proof. ($\Leftarrow$) Assume d_S is computable. Given $\mathbf{d}, p$, we want to compute a positive approximation $\widetilde{\chi}_S(\mathbf{d};p)$. The algorithm computes

$$y = d_S(\mathbf{d})[p+1]$$

and then outputs 1 iff $y \leq 3 \cdot 2^{-p-1}$. The correctness of this output is checked in two cases. If the output is 1, this means $d_S(\mathbf{d}) \leq y + 2^{-p-1} \leq 2^{-p+1}$, and this is correct. If the output is 0, this means $d_S(\mathbf{d}) \geq y - 2^{-p-1} > 2^{-p}$. and this is correct too.

($\Rightarrow$) Assume χ_S is positively approximable. To compute $d_S(\mathbf{d})[p]$, consider the uniform grid centered at $\mathbf{d}$ with grid size $2^{-p'}$ where $p' = p + \left\lceil \lg \sqrt{k} \right\rceil$. We "systematically search" the vertices of the grid, starting from the vertices closest to $\mathbf{d}$. At each vertex $\mathbf{y}$, we compute $\chi_S(\mathbf{y})[p]$. If $\mathbf{y}_0$ is the first vertex such that $\chi_S(\mathbf{y}_0)[p] = 1$, we output $z := (\|\mathbf{d} - \mathbf{y}_0\|)[p+1] - 2^{-p-1}$ as our estimate (note that we cannot compute $\|\mathbf{d} - \mathbf{y}\|$ exactly, so we compute a $p+1$-bit approximation).

REMARK: "systematically searching" is a bit nontrivial: it is somewhat reminiscent of Dijkstra's algorithm. However, here is a slightly more brute force solution: we search each concentric square (centered at $\mathbf{d}$) of increasing size. Suppose we found the first vertex which returns a 1 in a square of side length d. We must still continue to search squares whose side length is up to $\lceil d\sqrt{2} \rceil$. Among those vertices that returns a 1, we pick the one whose

distance to the center is minimum. This algorithm (or the Dijkstra-based algorithm) depends of the ability to compare two square roots of integers, but this is clearly possible.

Correctness of this procedure: first we show that this procedure halts. Now, the diameter of a grid cell is $\sqrt{k}2^{-p'} = 2^{-p}$. Hence, if any of the 2^k grid cells adjacent to a vertex $\mathbf{y}$ contains a point of S, then $\chi_S(\mathbf{y})[p]$ will output a 1.

To show that the output is correct, let the point $\mathbf{a} \in S$ be closest to $\mathbf{d}$. Then $d_S(\mathbf{d}) = \|\mathbf{d} - \mathbf{a}\|$, and we have

$$\|\mathbf{d} - \mathbf{y}\| \leq \|\mathbf{d} - \mathbf{a}\| \leq \|\mathbf{d} - \mathbf{y}\| + 2^{-p}.$$

On the other hand, we also have

$$\|\mathbf{d} - \mathbf{y}\| \leq z \leq \|\mathbf{d} - \mathbf{y}\| + 2^{-p}.$$

Thus $|d_S(\mathbf{d}) - z| \leq 2^{-p}$, i.e., z is a p-bit approximation of $d_S(\mathbf{d})$. **Q.E.D.**

The following shows why positive recognizability is a generalization of computability:

Lemma 39. *Let $a \in \mathbb{R}$. Then a is computable iff the singleton $A = \{a\}$ is positively recognizable.*

Proof. Suppose a is computable. Then we out $\chi_A(x)[p]$ to be 1 iff $|a[p+2] - x| \leq 3 \cdot 2^{-p-1}$. To see that this is correct, consider two possibilities: if we output 1 then $|a[p+2]-x| \leq 3{\cdot}2^{-p-1}$, and $|a-x| \leq 3{\cdot}2^{-p-1}+2^{-p-1} = 2^{-p+1}$. This is correct. If we output 0 then $|a[p+2] - x| > 3 \cdot 2^{-p-1}$, and $|a - x| > 3 \cdot 2^{-p-1} - 2^{-p-1} = 2^{-p}$, which is also correct.

Suppose A is positively recognizable. Wlog, $a > 0$. To compute $a[p]$, we find the smallest $i \in \mathbb{N}$ such that $\chi_A(i2^{-p-1})[p-1] = 1$. Note that $|a - i2^{-p-1}| \leq 2^{-p}$, from the definition of positive approximation of χ_A. Furthermore, this procedure must halt because if $i \in \mathbb{N}$ is the largest value such that $i2^{-p-1} \leq a$, then $\chi_A(i2^{-p-1})[p-1] = 1$. Thus, our procedure would have halted by step i. **Q.E.D.**

¶76. Set Operators and Set Oracles. Functions whose input arguments are sets and whose output is also a set are called **set operators** (cf. ¶59). E.g., the "set union operator" is given by

$$UNION : (A, B) \mapsto A \cap B.$$

We use OTM's to compute such operators: *each input set A is viewed as an input oracle*, with its own oracle tape. In fact, the oracle for A will be any positive approximation of A. To query an oracle A, we write on A's oracle tape a pair $(\mathbf{d}, p)$ where $\mathbf{d} \in \mathbb{D}^k$, $p \in \mathbb{D}$, and then enter the special query state $q_?$. In the next instant, the OTM will enter one of the two special states, q_0 or q_1. Entering state q_x $(x = 0, 1)$ amounts to the oracle saying that $\chi_A(\mathbf{d})[p] = x$.

Suppose M is an OTM for computing UNION. Then the input for M is the quadruple $(A, B, \mathbf{d}, p)$ where A, B are sets (represented by some positive approximations), and $\mathbf{d} \in \mathbb{D}^k, p \in \mathbb{D}$. The OTM must eventually halt in one of special states q_x ($x = 0$ or $x = 1$). We **declare** the output of M to be $M^{A,B}(\mathbf{d}; p) = x$. We say M **computes** the UNION operator if $M^{A,B}(\mathbf{d}; p)$ is a positive approximation to $\chi_{A\cap B}$, for all A, B.

Lemma 40. *The set UNION operator, $A, B \subseteq \mathbb{R}^k \mapsto A \cup B$ is computable.*

Proof. We claim that $\chi_{A\cup B}(\mathbf{x})[p]$ can be computed as

$$\chi_A(\mathbf{x})[p] \vee \ \chi_B(\mathbf{x})[p].$$

To see that this is correct, if $\chi_A(\mathbf{x})[p] = 1$ or $\chi_B(\mathbf{x})[p] = 1$ then $d_{A\cup B}(\mathbf{x}) \leq d_A(\mathbf{x}) \leq 2^{1-p}$, and so an output of 1 is warranted. If $\chi_A(\mathbf{x})[p] = 0$ and $\chi_B(\mathbf{x})[p] = 0$ then $d_{A\cup B}(\mathbf{x}) \geq 2^{-p}$, and so an output of 0 is warranted. **Q.E.D.**

Next, we ask whether the intersection operator is computable? The Exercise asks you to give a direct proof that no OTM can compute the operator $INTERSECT : (A, B) \mapsto A \cap B$. However, here is a stronger result:

Lemma 41. *There exist positively recognizable sets A, B such that $A\cap B$ is non-empty and not positively recognizable.*

Proof. Let $\mathbb{D} = \mathbb{D}_0 \uplus \mathbb{D}_1$ be any partition of $\mathbb{D}$ into two disjoint sets, each dense in the reals. To be specific, we first partition the set $(0, 1] \cap \mathbb{D}$ into $D_0 \uplus D_1$ where D_0, D_1 are both dense in $(0, 1]$. Each number in $(0, 1] \cap \mathbb{D}$ has the unique form $0.b_1b_2 \cdots b_n$ where $n \geq 1$ and $b_n = 1$. We put this number into D_0 iff n is even. Then $\mathbb{D}_0 = \mathbb{Z} + D_0 = \{n + d : n \in \mathbb{Z}, d \in D_0\}$ and $\mathbb{D}_1 = \mathbb{Z} + D_1$. Now let $A = \mathbb{D}_0 \cup \{a\}$ and $B = \mathbb{D}_1 \cup \{a\}$ for some uncomputable real number a. Clearly, A and B are both positively recognizable: the approximation function is $\chi(d)[p] = 1$ for all $d, p \in \mathbb{D}$. However, $A \cap B = \{a\}$ is not positively approximable, by lemma 39. **Q.E.D.**

Remark: the proof exploit the fact that positive recognizability of A cannot detect the presence or absence of non-dyadic numbers x, provided the set of dyadic numbers in A is dense in the neighborhood of x. For instance, $A = [0, a)$ and $A' = [0, a]$ are indistinguishable.

¶77. Graphs of functions. We now consider the relationship between a function

$$f : [0,1] \to \mathbb{R}$$

and its graph, $Graph(f) = \{(x, f(x)) : x \in [0,1]\} \subseteq \mathbb{R}^2$. The following lemma is easy to see geometrically:

Lemma 42. *Let S be the graph of $f : [0,1] \to \mathbb{R}$. If $|f(x) - y| > a$ then $d_S(x,y) > a/\sqrt{2}$.*

We leave the proof as an exercise. Now we may show:

Theorem 43 (Braverman). *Suppose $f : [0,1] \to \mathbb{R}$ is continuous. Then f is computable iff $S = graph(f)$ is positively recognizable.*

Proof. ($\Rightarrow$) Suppose f is computable. Then we positively approximate χ_S as follows:

$$\chi_S(x,y)[p] = 1 \quad \text{iff} \quad |f(x)[p+2] - y| \leq 7 \cdot 2^{-p-2}.$$

To show that this is correct, suppose $|f(x)[p+2] - y| \leq 7 \cdot 2^{-p-2}$. Then $|f(x) - y| \leq 8 \cdot 2^{-p-2} = 2^{1-p}$. This implies $d_S(x,y) \leq 2^{1-p}$. Next suppose $|f(x)[p+2] - y| > 7 \cdot 2^{-p-2}$. Then $|f(x) - y| > 6 \cdot 2^{-p-2} = 3 \cdot 2^{-p-1}$. By Lemma 42, this implies $d_S(x,y) \geq 3 \cdot 2^{-p-1}/\sqrt{2} < 2^{-p}$.

($\Leftarrow$) Suppose S is positively recognizable. We compute $f(x)[p]$ by iterating the following loop

$$i \leftarrow 0; \mathsf{while}\ (\chi_S(x, i \cdot 2^{-p-3})[p+3] = 0), i := i + 1$$

It is clear that this halts. At halting, we have $d_S(x, i \cdot 2^{-p-3)}) \leq i \cdot 2^{-p-2}$. By Lemma 42, we conclude that $|f(x) - i \cdot 2^{-p-3)}| \leq \sqrt{2} i \cdot 2^{-p-2} < 2^{-p}$. Hence we can outpot $i \cdot 2^{-p-3}$. **Q.E.D.**

Exercises

Exercise 14.1: Prove Lemma 42 ♢

Exercise 14.2: Extend the above theorem that, for continuous functions $f : [0,1]^k \to \mathbb{R}$, f is computable iff $Graph(f)$ is positively recognizable. ♢

Exercise 14.3: Show that the set COMPLEMENT operator is not computable. ♢

Exercise 14.4: Show that the INTERSECT operator is not computable. ♢

Exercise 14.5: Let $\alpha > 0$ and $f : \mathbb{R}^k \to \{0,1\}$. Call $\widetilde{f} : \mathbb{D}^k \times \mathbb{D} \to \{0,1\}$ an α-approximation if for all $\mathbf{x} \subseteq \mathbb{R}^k$ and $p \in \mathbb{N}$, $\widetilde{f}(\mathbf{x}; p) = 1$ if there is some $\mathbf{y}$ such that $\|\mathbf{x} - \mathbf{y}\| \leq 2^{-p}(1+\alpha)$ and $f(\mathbf{y}) = 1$. Also, $\widetilde{f}(\mathbf{x}; p) = 0$ if there is no $\mathbf{y}$ such that $\|\mathbf{x} - \mathbf{y}\| \leq 2^{-p}$ and $f(\mathbf{y}) = 1$. Thus "positive approximation" is the same as 1-approximation. Is the following true: if f is α-approximable for some α, then it is β-approximable for all $\beta > 0$. ♢

Exercise 14.6: (Braverman) Let $S \subseteq \mathbb{R}^k$ be bounded. Then the following are equivalent:
(1) S is positively recognizable
(3) For any $\varepsilon > 0$, we can compute a finite set T such that the Hausdorff distance $d(S,T) \leq \varepsilon$. ♢

End Exercises

§15. Theory of Real Approximations

¶78. Need for Approximability and Zero Decision. There are two severe restrictions in the theory developed so far: (1) not being able to decide equality $f_=$ (¶33), and (2) only continuous functions are computable (Lemma 14). There are many situations where both of these restrictions are undesirable or even unacceptable:

- Suppose we want to evaluate the polynomial $p(X) = 1 + X + X^2 + \cdots + X^{n-1}$. The application of Horner's rule requires $2n-1$ ring operations. But suppose we use the formula $p(X) = \frac{1-X^n}{1-X}$. When $X \neq 1$, this gives us a method with $2 \lg n + O(1)$ operations, including one division. We must detect the case where the input is $X = 1$, and in this case output n. This is an exponential speed up in terms of the number of operations. But in the current theory, such an algorithm is not possible since we cannot decide if $X = 1$.

- A similar situation arise in evaluating the determinant of an $n \times n$ matrix. We can do this in $O(n^3)$ field operations by using Gaussian elimination. But this requires testing that the pivot is non-zero, and to do pivoting if necessary. But the current computational model cannot support a pivoting algorithm. Therefore, we must compute the determinant using the expansion of determinants with $n!$ terms. We can save a little by exploiting a dynamic programming trick, but we still incur an exponential time complexity.
- Let us consider the problem of automatic theorem proving. Suppose we want to prove the following theorem...
- Practically all the problems of computational geometry are discontinous in a certain sense.

¶79. Decidable sets. We now consider a way out of this dilemma. First, we demand some additional properties from notations.

A set S is **decidable** if it has notation ρ such that the following two sets are decidable: $\{w \in \{0,1\}^* : \rho(w) = \downarrow\}$ and $\{(w, w') \in (\{0,1\}^*)^2 : \rho(w) = \rho(w')\}$. We call ρ a **dedidable notation**. Thus, given w, w' we can decide if they are proper, and if they are, we can decide if they are ρ-equivalent. Necessarily, decidable sets are countable.

Lemma 44. *Let S, T be decidable sets. Then the following sets are decidable:*
(i) Finite sets,
(ii) disjoint union $S \uplus T$,
(iii) Cartesian product $S \times T$,
(iv) Finite sequences S^*,
(v) Finite power set $\overline{2}^S$.

Proof. Let us sketch some of these constructions: suppose $\rho : \{0,1\}^* \to S$ and $\rho' : \{0,1\}^* \to T$ are two decidable notations. Then a decidable notation ρ'' for $S \uplus T$ may be given by $\rho''(0w) = \rho(w)$ and $\rho''(1w) = \rho'(w)$.

In case of Cartesian product, we define $\rho''(w\#w') = (\rho(w), \rho'(w'))$ using the representing set $\Sigma^* = \{0, 1, \#\}^*$.

For finite power set, we just represent each set by a list of its elements (we may allow or disallow duplicates). To test for equality of two sets, we first eliminate duplicates (if they are allowed) to obtain an equal number n of elements in each set (else they are unequal). Then we do at most n^2 equality tests to determine equality of the sets. **Q.E.D.**

Note that set union and set intersection are not among the decidable operations.

¶80. Decidable algebraic structures. An **algebraic structure** is a set R together with a finite number of algebraic operators and predicates. We say R is **decidable** if R has a decidable notation ρ and all its operators and predicates have ρ-notations.

E.g., The integers $R = \mathbb{Z}$ is an ordered ring with operators $\{+, -, \times, 0, 1\}$ and a total ordering $<$ on R. It is easy to see that $\mathbb{Z}$ is a decidable ordered ring.

Lemma 45. *If R is a decidable ring, then so is its quotient field $Q(R)$, and the polynomial ring $R[X]$.*

Proof. (i) The standard quotient construction via pairs $\overline{x, y}$ such that $\overline{x, y} \equiv \overline{x', y'}$ iff $\overline{xy'} \equiv \overline{x'y}$. Since $\overline{xy'}$ and $\overline{x'y}$ are both computable and we can decide if elements of R are equal, we are done.

(ii) We view polynomials in $R[X]$ as finite sequences $\overline{a_0, \ldots, a_n}$ over R, representing the coefficients. Equality of two polynomials $\overline{a_0, \ldots, a_n} \equiv \overline{b_0, \ldots, b_m}$ is easy to decide. All the ring operations are decidable. **Q.E.D.**

Theorem 46. *The algebraic closure $\overline{F}$ of a decidable field is a decidable field.*

Proof. **Q.E.D.**

Theorem 47. *There is a decidable real field that extend the field of algebraic numbers.*

Proof. Use our decidability of shortest path result? **Q.E.D.**

¶81. Real basis. Dedekind has shown how to use rationals as the basis for constructing the reals. It is well-known that the role of rational numbers could be replaced by dyadics, which is to be preferred for computational purposes. In fact, in the world of computing (programming languages), elements of $\mathbb{D}$ are often called "reals". Both $\mathbb{Q}$ and $\mathbb{D}$ illustrates what might be called the "basis" for real computation. Our theory of approximation can begin with any such set.

¶82. Remarks. Despite our interest in approximations, it is important to emphasize that we are not really interested (per se) in functions of the form $\widetilde{f} :\subseteq \mathbb{D} \to \mathbb{D}$, but in functions of the form $f :\subseteq \mathbb{R} \to \mathbb{R}$. The reason is this: there are not enough properties in functions in $\widetilde{f} :\subseteq \mathbb{D} \to \mathbb{D}$. Indeed, all the standard tools of analysis (continuity, differentiability, etc) we use to understand real functions are lost. But by focusing on $f :\subseteq \mathbb{R} \to \mathbb{R}$ we retain all these properties. In short, we insist on discussing real functions even after we decide to substitute approximability for computability. In other words, we must not yield to the intuitionistic impulse to conflate the world of computation with the world of analysis. Both worlds are necessary and have complementary roles: observe that notations such as $\rho :\subseteq \Sigma^* \to S$ serve to connect these two worlds. Moreover, there is no direct computational purpose in ρ – in general, it is meaningless to speak of "computing ρ". Remove the world of analysis, and the soul of our functions is lost.

§16. Transfer Theorem

¶83. A **computational basis** (or simply, **basis**) is any set Ω of partial real functions. So each $f \in \Omega$ has an arity $k \in \mathbb{N}$ and $f :\subseteq \mathbb{R}^k \to \mathbb{R}$. Note that if f has arity 0, then f represents a real constant.

Examples:

1. $\Omega_0 = \{+, -, \times\} \cup \mathbb{Z}$
2. $\Omega_1 = \Omega_0 \cup \{\div\}$
3. $\Omega_2 = \Omega_1 \cup \{\sqrt{\cdot}\}$

¶84. The Real RAM Model is an idealized computing model which has infinitely many registers, and each register can store an arbitrary real number (or could be undefined). Each machine in this model is a finite sequence of instructions which has the usual set of instructions found in modern assembly language, including the ability to index any register by an integer i. We could increment and decrement the contents of registers by 1. The constants 0 and 1 are available. We could compare the contents of any two registers and branch based on the result of the comparison. However, the operations which we allow on real numbers may vary, depending on the choice of a basis Ω.

An **algebraic model** over a basis Ω is a Real RAM Model in which the machines could also perform operations of the form:

$$Z = f(X_1, \ldots, X_k)$$

where $Z, X_1, \ldots, X_k$ are variables and $f \in \Omega$ has arity $k \geq 0$. Note that the result of this operation is stored in register Z; note that the content of Z may be undefined in case f is undefined at its arguments.

An algebraic machine over the basis Ω may also be called an **Ω-machine**.

¶85. Suppose P is a problem that shown to computable in the algebraic model. Let us assume for simplicity that P is a simple numerical problem,

$$f :\subseteq \mathbb{R}^* \to \mathbb{R}. \tag{26}$$

Let M be an algebraic machine. We say M **$\mathcal{A}$-approximates** P if, for $x \in \mathbb{R}^*$ and $p \in \mathbb{R}$, if $f(x) = \uparrow$, then M halts in a special state $q_\uparrow$. Otherwise, M halts in another special state $q_\downarrow$ and a special output register will contain a value z such that $|z - f(x)| \leq 2^{-p}$.

We have two other variations of this definition:

1. If the output of M satisfies $z = f(x)$ whenever $f(x) \downarrow$, we say M **exactly solves** P.
2. If the output of M satisfies $|z - f(x)| \leq 2^{-p}|f(x)|$ whenever $f(x) \downarrow$, we say M **$\mathcal{R}$-approximates** P.

¶86. Transfer Theorem. Now, we say f in (26)

Now, it turns out that many of the algorithms in numerical analysis are not immediately approximation algorithms in our sense. For instance, consider the problem of computing the largest eigenvalue of a matrix.

For this, we need to prove some general theorems.

References

[1] L. Blum. Computing over the reals: Where Turing meets Newton. *Notices of the Amer. Math.Soc.*, 51:1024–1034, 2004.

[2] M. Braverman. On the complexity of real functions. In *46th IEEE Foundations of Computer Sci.*, pages 154–164, 2005.

[3] M. Braverman and S. Cook. Computing over the reals: Foundations for scientific computing. *Notices of the AMS*, 53(3):318–329, Mar. 2006.

[4] K.-I. Ko. *Complexity Theory of Real Functions*. Progress in Theoretical Computer Science. Birkhäuser, Boston, 1991.

[5] S. Smale. Some remarks on the foundations of numerical analysis. *SIAM Review*, 32:211–220, 1990.

[6] A. M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc. Series 2*, 42:230–265, 1936. Corrections: Vol.43 (1937) pp.544-546.

[7] K. Weihrauch. *Computable Analysis*. Springer, Berlin, 2000.

[8] C. K. Yap. On guaranteed accuracy computation. In F. Chen and D. Wang, editors, *Geometric Computation*, chapter 12, pages 322–373. World Scientific Publishing Co., Singapore, 2004.