DATA COMMUNICATION

The massive proliferation of computers and the explosive growth of the Internet have led to a great demand for data communication equipment and services. Just in the past few years, we have experienced tremendous progress in the speed at which computers are able to communicate with each other. Data communications technology has made remarkable strides from supporting low-bit-rate text-based communications during the previous decade to being able to support high-speed, low-latency multimedia-based applications of the future. There has also been significant progress made toward extending all these services to the mobile user by means of wireless technology.

There are a number of applications that require the use of data communication networks. Some of these applications operate at low data rates and require less bandwidth. These include applications such as point-of-sale, e-mail transfer, short messaging, low-speed Internet access, and telemetry. Other applications such as high-speed Internet and Intranet access, video conferencing, local area network (LAN) interconnections, full-motion video, image transfer, telemedicine, virtual private networks (VPNs), distance learning, residential multimedia, and TV/video distribution require large bandwidths and hence need to be supported by broadband networks.

A data communications network is, by definition, a collection of applications hosted on separate computers that are interconnected by means of a communications infrastructure. This article provides a concise introduction to the subject of data communications with an emphasis on describing the basic technical concepts by examples of existing standards and services. A list of well-recognized books and articles in the literature available on this subject is provided at the end of the article for the interested reader.

We begin our discussion with a brief overview of the units that make up a data communications network. We then introduce some general networking concepts and discuss the relative merits of different networking technologies. We next describe the layered architecture for data network design and the function of each layer. A more detailed discussion of the principles and concepts associated with the design of the lower three layers follows. Toward the end of this article, we provide a brief overview of some of the existing and emerging data communications standards.

GENERAL NETWORKING CONCEPTS

Figure 1 shows a representation of a generic data communication network. The basic conceptual units that make up the

J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*. Copyright © 1999 John Wiley & Sons, Inc.

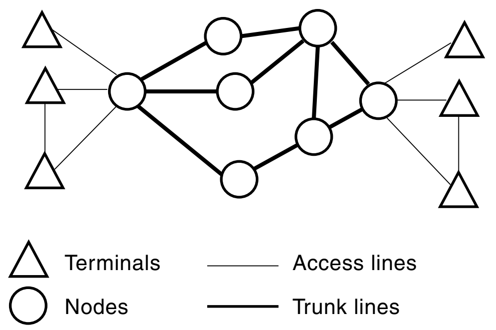

Figure 1. A simple representation of a data communication network showing the interconnection between different network components.

network are the *terminals,* the *access lines,* the *trunks,* and the *switching nodes.* Data communication is achieved when the users of the network and the terminals exchange information or messages over the network. The terminals represent the end-user equipment that is connected to the network. This equipment could be general purpose computers, database systems, dumb terminals, or any other communications devices. The terminals provide, either through software or human interaction, the functions required for information exchange between pairs of application programs or between application programs and people. These functions include call setup, session management, and message transmission control. Examples of applications include electronic mail transfer, terminal to computer connection for time sharing or other purposes, and terminal to database connections.

Access lines provide for data transmission between the terminals and the switching nodes. These connections may be set up on a permanent or switched basis. The access lines provide a channel for data communication at a certain rate (bits per second or bps) called data rate or channel capacity. The access line capacity may range from a few hundred bps to a few million bps, and they are usually not the same for all terminals of a given network. The actual information-carrying capacity or throughput of the link is typically less than the raw bit rate that the channel provides, and depends on the overhead associated with the protocols employed.

The trunks are the transmission facilities that provide data communication between pairs of switching nodes, and typically have much larger capacity than the access lines. They carry data for multiple connections. The trunk lines are typically multiplexed using frequency division multiplexing (FDM), time division multiplexing (TDM), or asynchronous statistical multiplexing (ASM).

When there are a large number of terminals connected to a network, it is impossible to connect every station in the network to every other station by means of a direct dedicated line. Therefore communication between a source and destination device is achieved by traversing through a series of intermediate switching nodes. The function of these switching nodes is to provide a switching or routing facility that will move the data from one node to another node until it is delivered to the destination. The network shown in Fig. 2 is a simple example of a switched communication network. As seen from Fig. 2, if station 1 wishes to communicate with station 5, it may go through switching nodes A, C, and F. Every switching node in the network may not have a direct link to every other node, but there will be at least one or more possible paths between any two nodes.

There are two conceptually different approaches to switching that are in use today: *circuit switching* and *packet switching.* In a circuit-switched environment there is a dedicated path established between two stations that wish to communicate. This path is a sequence of links between network nodes, and in each internode link there is a physical channel dedicated to the particular connection. Circuit switching involves setting up an end-to-end circuit before any data are transmitted. For example, if station 1 wishes to communicate to station 5, it sends a connection request to node A. Node A then establishes the next link in the path to F, and allocates a free channel on that link. This continues until an end-to-end connection is established. Once the connection is established, data transmission occurs through the allocated channels. The allocated channels are not available to other users until the current call is completed. Circuit switching is a technology designed for voice communications where end-to-end connectivity has to be guaranteed with virtually no transmission delay throughout the conversation. Though designed for voice communication, there is a large volume of data traffic that is carried over circuit-switched networks even today. For example, the circuit-switched global public switched telephone network (PSTN) is used for data communication using dial-up modems. Integrated services digital network (ISDN) is another example of a circuit-switched network that provides data communications. There are a number of private networks that use dedicated leased lines such as T1/EI or T3/E3 for data communications.

In a circuit-switched network, call establishment, maintenance, and termination are done by means of control signals. Until the late 1970s signaling for circuit establishment was performed using the same channel used by the data traffic. That is, in order to set up a call through the network, the call setup information was sent sequentially from node to node using the same circuit that would eventually become the circuit used for connecting the end users. This signaling system, called *in-band signaling,* has two major disadvantages. First, the rate of signaling is limited to the circuit speed, and second, the circuits that could have been used for sending traffic is consumed simply to find a path between the end points. These limitations resulted in unnecessary bottlenecks.

An alternate *out-of-band signaling* system, usually called *common channel signaling,* was developed to solve this problem. Under this setup, signaling takes place over a separate signaling network, which is partitioned from the network that carries the user traffic. The use of dedicated separate signal-

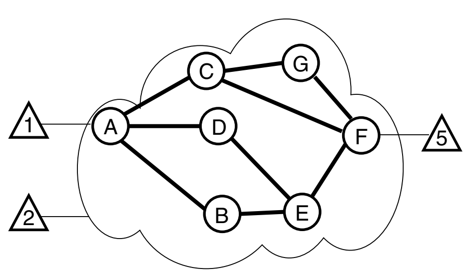

Figure 2. Illustration of a simple switching network showing how the end terminals are interconnected.

ing channels reduces the call setup time, and it is more adaptable to the evolving functional needs of advanced intelligent networking.

Circuit-switched networks have a lot of deficiencies which motivated the development of packet switching networks for transmission of data. In many data communications scenarios such as surfing the Internet, the communication link between the host and user computer is idle most of the time. It is only occasionally that the user downloads a web page or sends an e-mail. Most of the time is spent on browsing the downloaded material or composing the e-mail. It is a terrible waste to tie up a communication link for the entire surfing session when a user is only actually using it for a fraction of the time that he or she is connected. Apart from inefficient link usage, circuit-switched networks also have the disadvantage of supporting only a constant data rate transmission. This makes it mandatory that all the devices connected to the network transmit and receive at the same constant data rate, and hence interconnecting a variety of computer and network devices becomes more cumbersome. These inherent disadvantages of circuit-switched networks are overcome in packet-switched networks. Examples of packet switching technologies include X.25, frame relay, TCP/IP, and ATM in the wide area networking scenarios, and Ethernet, Token Ring, and ATM in the local area networking scenarios.

A packet switching network is a distributed collection of interconnected packet switching nodes. Each of these nodes is capable of receiving, storing, and forwarding small packets of data. Data are routed from the source to destination through a series of these nodes. To send data over the packet network, the user message is broken up into a series of small packets, which are transmitted one at a time. A typical packet varies in size from a few tens of octets to a few thousand octets. Each packet is appended with a header, which contains control information that helps the network route the packet to the correct destination. At each node along the route from source to destination, the packet is received, stored briefly, and passed on to the next node.

There are many advantages to this approach. First, the links between the source and destination may be shared by multiple users, each of them capturing the links only when required. This allows for much greater efficiency in multiplexing. Also each node is capable of data rate conversion, and hence devices with different data rates can communicate with each other.

Within the packet switching paradigm, there are fundamentally two different approaches to handling the stream of packets through the network. One is the *connectionless* or *datagram* approach, and the other is the *connection-oriented virtual circuit* approach. In the datagram approach, each packet is treated independently without any reference to the other packets that make up the message. Each packet contains the destination address in its header but no information on the route to be followed. These packets, called datagrams, are similar to the letters that we send using our postal system. Each letter is put in an envelope with a destination address and dropped in the mailbox. On receiving a packet, the switching node determines which is the best route to follow to get to the destination based on information available to it about the network conditions. After finding the best route, the node forwards the packet to the next node along the route. Since each packet is handled independently, the route taken by each packet may be different depending on the network conditions at the time it arrives. It is therefore possible that some packets take longer routes than others, and hence arrive out of sequence at the destination. It is up to the destination station to figure out how to reorder them. The destination stations are also responsible for detecting lost and corrupted packets, and ensuring that the source retransmits them until correctly received.

The other packet switching approach is called the virtual circuit (VC) approach. In this case, before any packet is sent across the network, the network establishes a route from the source to the destination. This is established by transferring call-request and call-accept packets between the two stations that wish to communicate. Once a route is established, all subsequent data packets in that session use the same route, and in this sense it is similar to circuit switching. This preassigned route used for the entire duration of the call is called a virtual circuit, and it is identified by means of a virtual circuit identifier (VCI). Each packet that uses the preassigned route has the VCI information in its header. Having a VC established prior to sending data does not imply that there is a dedicated path as in circuit switching. Packets still go through the buffering and queuing processes, and can experience delays. However, the routing decisions are made at call setup, and each packet follows the same route, ensuring that they arrive at the destination in the right sequence. Virtual circuits can also use error control mechanisms to ensure that the packets arrive at the destination without errors. Virtual circuits are well suited for transferring larger amounts of data such as in a file transfer, since the network does not have to make routing decisions on every individual packet. For small bursts of data though, the overhead associated with setting up the call a priori may not be justified. Also datagrams are more flexible in adapting to congestion and node failures, since packets can quickly be rerouted through alternate routes without tearing down the virtual connection. Another important distinction between the datagram and virtual circuit approach is that it is easier to design VC-based systems for guaranteed quality of service, whereas datagram based systems are often engineered for best effort service.

SEVEN-LAYER ARCHITECTURE

It is very instructive to look at a data communications network as a hierarchy of nested modules or layers. Each layer performs a function in support of other layers, and this function is called a service. Due to this modular design, a user of a given layer needs only to know what its inputs, outputs, and service are and does not need to be aware of its inner workings. Each given layer in the hierarchy regards the next lower layer as one or more black boxes that provide a specified service to the given higher layer. Each layer at a given node is able to communicate with the layer above and below it, and also to the corresponding layer (peer layer) at another node. Associated with each layer is a protocol designed to enable it to communicate with its peer layers. The collection of these layered protocols is referred to as a *protocol stack*.

Figure 3 illustrates an example of a layered architecture for data networks. This seven-layered architecture, called the open systems interconnection (OSI), serves a reference model and has a relatively clean structure. Data networks in exis-

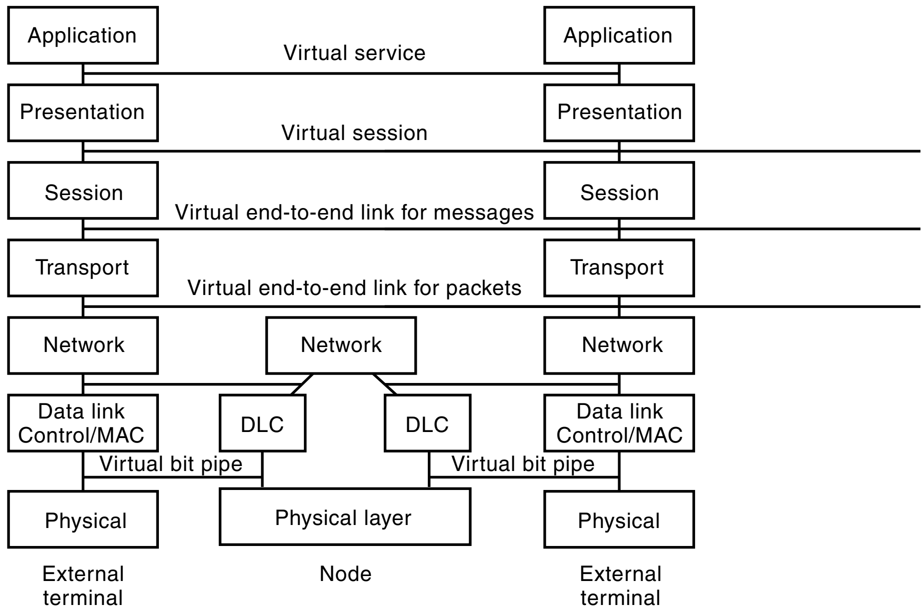

Figure 3. Illustration of the interaction among the seven layers of the OSI network architecture. Notice only the lower layers are involved at the intermediate nodes.

tence today may or may not use all the seven layers of the reference model, and may have additional layers. As an example, Fig. 4 shows how the TCP/IP protocol stack matches the OSI reference model.

The seven layers in the reference model are the physical layer, data link control (DLC) layer, network layer, transport layer, session layer, presentation layer, and the applications layer. The physical layer is concerned with the transmission of a sequence of bits over the physical medium. It deals with the mechanical, electrical, and functional characteristics to access the physical medium. The physical layer is a virtual bit pipe that maps the bits coming from the next higher layer (DLC layer) into signals appropriate for the physical channel at the transmitting end, and maps the signals back into bits at the receiving end.

The DLC layer ensures that the bit stream it sends across the physical link is received without any errors. Each point-to-point link has a DLC layer at each end of the link that ensures that the data transferred across the physical link is reliable. The DLC layer may introduce some framing and synchronization to the bit stream, but it does not concern itself about the details of packetization that is done at the network layer. While the DLC ensures reliable transfer of bits over a point-to-point link, there is a need for an intermediate layer to manage the multiaccess link so that frames can be sent by each node without constant interference from the other nodes. This layer is called the medium access control (MAC) layer. It is usually considered as the lower sublayer of layer 2 with the conventional DLC considered as the higher sublayer.

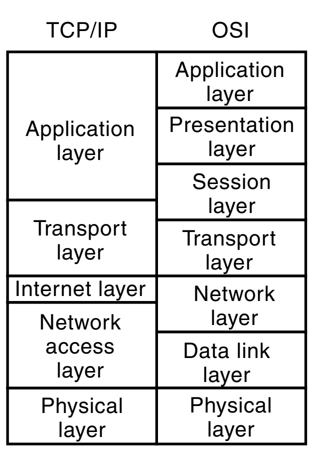

Figure 4. Comparison of TCP/IP and OSI layering to show how data networks deviate from the reference model.

The network layer is responsible for establishing, maintaining, and terminating connections. It makes the upper-layers independent of the data transmission and switching technologies used to connect systems. There is one network layer process associated with each node and the terminals of the network. All these processes are peers, and they work together in implementing routing and flow control for the network. When a frame enters a node from a communication link, the bits in that frame pass through the physical layer to the DLC layer. The start and end points of the frame are determined at the DLC layer, and if the frame is accepted as correct, the DLC strips off the DLC header and trailer from the frame and passes the packet up to the network layer. A packet consists of two parts: a packet header followed by the packet body and the DLC trailer. The network layer module uses the packet header along with stored information at the module to carry out its routing and flow control functions.

The transport layer provides reliable end-to-end transfer of data. Its functions include breaking messages into packets and reassembling packets into messages, performing error recovery if the lower layers are not error-free, doing flow control if not done adequately at the network layer, and multiplexing and demultiplexing sessions together. Breaking messages into packets is rather simple, with packet size being the only decision to be made. There is a significant relationship between packet size and transmission time. Increasing the packet size means less overhead, but at the same time it increases the transmission delay or latency. The reassembling function is relatively simple as long as the transport layer process has plenty of buffer space available, but it can be quite tricky if there is limited buffer space. If the packets arrive out of order at the transport layer, the problem of reassembling becomes even more difficult. Multiplexing of sessions is done to reduce the overhead when there are several low-rate sessions from the same source site to the same destination site. Conversely, one high-rate session may be split into multiple sessions to improve flow control.

The layer above the transport layer is called the sessions layer. This layer provides the framework for communication between applications. It is responsible for establishing, managing, and terminating sessions between cooperating applications. The session layer deals with the access rights of a service or application. Above the session layer is the presentation layer. The major functions of the presentation layer include data encryption, data compression, and code conversion. Encryption is required for security purposes, compression reduces the bandwidth requirements, and code conversion is necessary because of incompatible terminal devices.

The highest layer in the OSI model is called the application layer. Each application requires its own software to perform the required task. While all the lower layers work the tasks required for many different applications, the application layer performs tasks that are specific to the particular application such as file transfer or remote log in.

The merits of the layered approach will become clearer as we discuss the lower three layers in more detail. While layering provides great modularity and simplifies the entire design process, there is an enormous amount of overhead and delay generated due to the seven-layer process. Protocol designers continually strive to improve the efficiency of the protocol stack by attempting to minimize the overheads and decrease the delay, while maintaining the simplicity and modularity.

PHYSICAL LAYER—TRANSMISSION MEDIA AND DATA ENCODING

The physical layer is the lowest layer in the seven-layer model. It is in this layer that the actual transmission of bits takes place. This layer deals with the interface between the terminal and the communication medium. The function of the physical layer is to convert the digital bit stream that it receives from the higher layer into signals that can be transmitted over the physical medium of transmission.

Before we get into the details of how this conversion is achieved, let us briefly look at the various types of media that can be used for data communication. The available media may be broadly classified as being either *wired* or *wireless.* A wired medium is typically made of untwisted or twisted pair wires, coaxial cables, or optical fibers. A wireless medium makes use of electromagnetic waves in the radio frequency (RF), microwave, millimeter wave, or infrared bands, which are capable of propagating in the atmosphere or outer space.

The characteristics and quality of data transmission are determined both by the characteristics of the medium and the characteristics of the signal. For any data communication system, the two most important performance parameters are data rate and range. Data rate defined in bits per second quantifies the volume of data that can be transmitted in a given time. The range specifies the maximum distance over which the data transmission can be reliably achieved. The maximum data rate that can be achieved over a given medium depends on the bandwidth of the medium. In a very general sense, the higher the data rate associated with a signal, the higher will be the bandwidth of the signal, and hence transmission of such a signal will require a medium that can support a higher bandwidth. For instance, toll quality voice signals have a bandwidth of only 3 kHz and can be transmitted over a relatively narrowband medium, whereas video signals have a bandwidth of a few megahertz and hence require a broadband medium for transmission.

While the bandwidth of the transmission media dictates the data rate that can be achieved, the distance over which the signals can be transported depend on the transmission impairments such as attenuation, distortion, noise, and interference in the medium. Attenuation is a measure of the decrease in signal amplitude as it travels over a distance. The greater the attenuation and noise in the medium, the shorter will be the communication range. Among the wired media, twisted pair suffers greater attenuation than coaxial cables, which in turn suffer more than optical fiber. In wireless systems, attenuation depends on the frequency band of operation, with higher frequencies causing greater attenuation. Attenuation in a radio system also depends heavily on the nature of the radio environment such as terrain conditions, presence of obstruction, atmospheric absorption, fading, and so on.

Apart from the attenuation and bandwidth characteristics, the choice of transmission medium is also determined by factors such as cost, availability, ease of installation, need for mobility, etc. The appropriate medium should be selected depending on the need of the particular application.

The most common transmission medium is the twisted pair wire used in the access link of the telephone network. This link is now being increasingly used for data. Using voice band modems, these lines provide a data rate of 33 kbps, and with asymmetric digital subscriber line (ADSL) modems they can provide up to a few megabits per second at distances less than 1500 feet. Coaxial cables are widely used in local area networks and cable TV networks. These cables can provide data rates from ten to several hundred Mbps. Optical fiber allows data rates from several thousand Mbps to terabits per second.

After selecting an appropriate medium for transmission, the next step is to design the appropriate interface to convert the digital bit stream to signals suitable for transmission over the medium. This process is called data encoding. In many cases, such as the telephone line, the transmission medium is more suited for carrying analog signals. Since the original bit stream is digital in nature, it is necessary to undertake a digital-to-analog conversion before digital data can be sent over such a transmission media. The process of encoding digital data into an analog signal suitable for transmission is called *modulation,* and the reverse process of decoding the analog signal into digital data is called *demodulation.* Devices that perform the twin operations of modulation and demodulation are called *modems.*

Modulation is the process of impressing the information contained in a digital data stream onto a carrier signal. This is achieved by varying the amplitude, phase, frequency, or a combination of amplitude and phase of the carrier signal in accordance with the value of the digital data. The type of modulation used has a significant effect on the achievable data rate and range. Some modulation schemes achieve greater data rates at the expense of a shorter range, and others achieve higher ranges at the expense of a slower data rate. There is always a trade-off between data rate and range, and it should be the endeavor of the designer to make the trade-off in a way that is favorable to the design considerations. Continual advances in modulation techniques have led to the development of sophisticated schemes that permit this trade-off to be made at points approaching the theoretical limits established by Shannon in his classical channel capacity theorem.

DATA LINK LAYER—ERROR DETECTION AND RECOVERY

As discussed in the previous section, the physical layer suffers from various transmission impairments that can cause errors in the bit stream. It is the function of the data link control layer to ensure that those errors are detected and corrected before the frames received from the physical layer are passed on to the upper layers. Many data networks, especially those with very reliable transmission links such as fiber-optic channels, sometimes do not employ a data link layer. In such cases, since errors occur so rarely, corrective action is left to the higher-layer modules. However, some data links, especially those that use wireless links, require a data link layer to ensure that network layer receives the data with minimum errors, thereby reducing the load on the higher layers.

Error detection requires extra bits to be appended to the frame. The simplest way to do this is to append a single bit, called a parity bit, to the string of bits that make up a frame. The parity check bit is assigned a value of 1 if the number of 1's in the bit string is odd, and a value of 0 otherwise. This forces the total number of 1's in an encoded string to be always even. Now, if one bit gets corrupted in the channel, the total number of 1's in the string will no longer be even, and the receiver will be able to detect it. This error detection scheme is remarkably simple, and works very well for single bit errors. However, when more numbers of bits get corrupted in a frame, a single-bit parity check will not suffice. Parity checking codes can be expanded to detect multiple errors by using more than one bit for parity checking. The underlying idea is to map a string of K bits to another string of $K + L$ bits called codewords, where L is the number of parity bits. Since there are 2^K possible strings of size K, and 2^{K+L} possible codewords of size $K + L$, we only need a subset of the 2^{K+L} possible codewords to represent all the 2^K strings. If we can cleverly select this subset such that the chosen codewords are as dissimilar to one another as possible, we can maximize the number of errors that can be detected. The dissimilarity between two codes is quantified by the distance between the two codes and is computed as sum of the modulo-two sums of corresponding bits of the two codewords. While maximizing the distance between two codes improves the error detection capability, it does not guarantee detection of burst errors. Other techniques such as interleaving are used to combat burst errors. The parity check code that is most commonly used today is called the cyclic redundancy check (CRC) codes.

Once an error is detected at the receiving DLC layer, the next step is to simply inform the transmitting DLC that the frame has been received in error, and hence request a repeat of the bit frame. In principle, this procedure may be repeated as many times as required until the frame is received correctly. Positive or negative acknowledgments are sent back by the receiver to the sender over a reliable feedback channel in order to report whether a previously transmitted frame has been received error free or with errors. A positive acknowledgment (ACK) signals the transmitter to send the next packet, and a negative acknowledgment (NAK) is a request for frame retransmission. This strategy for controlling errors is called automatic repeat request (ARQ). There are three basic versions of ARQ: *stop-and-wait, go-back-N,* and *selective repeat.*

The stop-and-wait ARQ is the simplest ARQ scheme. After transmission of each frame, the sender waits for an acknowledgment from the receiver before sending the next frame or repeating the previous frame. Stop-and-wait ARQ is not very suitable for high-speed transmission, since it makes very inefficient use of the channel. A lot of time is wasted in simply waiting for an acknowledgment, and this can cause significant inefficiencies especially in long delay channels such as satellite links.

Sliding window protocols were invented to overcome the problems associated with stop-and-wait ARQ. Here the data link is connected via a full-duplex channel, and multiple frames are transmitted successively at a given time. A particular form of sliding window protocol is called the go-back-N protocol, and has been implemented in standards such as the high-level data link control (HDLC) and transmission control protocol/Internet protocol (TCP/IP). Here N frames are sent before receiving an acknowledgment. Each frame is numbered by $l = 0, 1, \ldots, 2^{k-1}$, where k is the number of bits used to represent the sequence number of a frame. Anytime a frame is received in error, all frames succeeding it are discarded. This provides for easy sequencing at the receiver but is wasteful of bandwidth. This is particularly true when the value of N is very high as in satellite links where it is typically set to 127. The bandwidth inefficiency can be improved if only the erroneous frames are retransmitted. This is what is done in selective repeat ARQ. The only disadvantage here is that the receiver needs to hold all the frames in the buffer to reorganize the sequence of frames in the event of an error.

Table 1 provides the expressions for the throughput efficiency for the three ARQ schemes discussed above. As seen from the expressions, the efficiency is a strong function of α, which is the ratio between the propagation delay time and frame duration. The efficiency also depends on the probability of frame error, P.

The DLC schemes that are implemented today in data networks are all based on enhancements on variations and com-

Table 1. Throughput Efficiency of Various ARQ Schemes[a]

	Throughput Efficiency
Stop-and-wait	$\eta_{sw} = \dfrac{1 - P}{1 + 2a}$
Go-back-N	$\eta_{GB} = \dfrac{1 - P}{1 + 2aP}, N \geq 2a + 1$
	$\eta_{GB} = \dfrac{N(1 - P)}{(1 + 2a)(1 - P + PN)}, N < 2a + 1$
Selective repeat	$\eta_{SR} = 1 - P, N \geq 2a + 1$
	$\eta_{SR} = \dfrac{N(1 - P)}{1 + 2a}, N < 2a + 1$

[a] Ref. 5.

binations of the basic techniques described in this section. There are a number of standards developed for data link control. They include the HDLC protocol standardized by ISO, and the link access procedure on the D channel (LAPD) standardized by the ITU-T. HDLC is the data link control used in X.25 networks, and LAPD was developed for ISDN.

MEDIUM ACCESS CONTROL LAYER—PROTOCOLS FOR SHARING RESOURCES

A medium access control (MAC) protocol is a set of rules to control access of distributed clients or terminals to a shared communication medium. The function of the MAC protocol is to schedule, control, and coordinate the use of the shared resource. There are various multiple access schemes available, and they may be broadly classified into three categories: *fixed assignment techniques, demand-based assignment,* and *random access techniques*.

Fixed assignment schemes include techniques such as frequency division multiple access (FDMA), time division multiple access (TDMA), and code division multiple access (CDMA). In an FDMA scheme the available channel bandwidth is divided into a fixed number of bands, each of which are allocated to a requesting terminal based on availability. Once a frequency band is allocated to a particular user, it is not available to other users until the existing user terminates the connection. In a TDMA scheme each terminal requiring to use the channel is allocated a time slot during which it is allowed to transmit or receive data. The total number of time slots in a frame is fixed, and once a time slot is allocated to a particular user, other users are blocked from using that time slot. In an FDMA scheme, the number of frequency bands determines the maximum number of simultaneous connections possible, whereas in a TDMA scheme the number of time slots is the determining factor. In a CDMA system all the users are allocated the same frequency, and they all transmit at the same time. They are distinguished from one another using unique pseudorandom spreading codes. While FDMA and TDMA provide completely contention-free access, the same cannot be said about CDMA under all conditions. In CDMA systems, as the number of users increase, there is performance degradation due to interference. Under such conditions, even though there is a fixed channel exclusively for the allocated user, interference from other users can be thought of as a form of contention. Since FDMA, TDMA, and CDMA schemes allocate a fixed bandwidth to a user, they are best suited for constant bit rate applications.

Many applications in data communications require a variable bandwidth. A fixed bandwidth assignment is very inefficient for data traffic that arrives in random bursts. In order to overcome this inefficiency, many data networks employ demand-based assignment. Demand-based assignments are implemented using either a reservation scheme or a polling mechanism. In a reservation-based system a terminal specifies the required bandwidth, and the system allocates an appropriate number of time slots or frequency bands to the requesting user. These schemes are called demand-assigned TDMA (DA-TDMA) and demand-assigned FDMA (DA-FDMA), respectively. Polling schemes may have either a centralized or distributed structure. In the centralized polling scheme there is a central controller that polls every terminal in a sequential manner. Those terminals that have data to transmit respond to the poll and get bandwidths allocated to them. In a distributed polling scheme, a token is passed from one terminal to the other in turn. If a particular terminal is idle upon receiving the token, it lets the token pass by. If a terminal is active while receiving the token, it seizes the token, transmits its packets, regenerates the token, and puts it on the medium when its transmission is complete or upon reaching its time limit. Examples of networks that use token passing include IBM's Token Ring network, IEEE 802.5, and fiber data distribution interface (FDDI).

The third category of medium access scheme is the random assignment scheme. Under this scheme there is no preassigned time for each terminal. Any terminal that requires access to the medium can seek access in a random fashion. There are many algorithms that fall under the category of random assignment. The most basic one is the unslotted or *pure ALOHA* scheme. Under this scheme a terminal that has data to transmit simply sends the data frames across the medium. If the channel is free at that time, the frames get through successfully. If the transmission of one terminal coincides with a transmission from another terminal, the frames collide and get corrupted in the channel. Under such a situation the terminal retransmits its frame after a random delay. When the network is heavily loaded, the unslotted ALOHA scheme generates a large number of collisions and hence requires a lot of retransmissions, which decreases the effective utilization of the channel. The maximum channel utilization in a pure ALOHA scheme is only 18%. The efficiency of this random access scheme can be considerably improved by modifying the ALOHA scheme to allow transmission to begin only at the start of predefined time slots. This scheme is called *slotted ALOHA,* and achieves an efficiency of 37%.

Both ALOHA schemes are rather simplistic; they do not take advantage of the fact that the propagation delay between terminals is very small when compared to the packet transmission time. A *carrier sense multiple access* (CSMA) scheme was developed to overcome some of the limitations of the ALOHA scheme and improve channel utilization. Under this scheme a terminal wishing to transmit first listens to the medium to determine if another transmission is in progress. If it senses that the medium is in use, the terminal waits and listens again later. If the medium is idle, the terminal transmits the packet. Even under this scheme it may happen that two or more stations attempt to transmit at about the same time. If this happens, there could be a collision leading to packet loss. Under this condition the terminal reschedules its next sensing of the channel to take place after a random back-off time and then begins the transmission process again. The randomness of the back-off time ensures that two stations do not get locked continuously into collisions. The maximum utilization of the channel that is achievable through CSMA far exceeds that of the ALOHA schemes. There are variations of CSMA that can achieve greater than 80% efficiency.

Even though CSMA minimizes the probability of collisions, when two frames collide, the medium remains unusable throughout the duration of the frame. For long frames the amount of capacity wasted can be very high. This wastage can be reduced if a terminal continues to listen to the medium while transmitting, and aborts transmission the moment it senses a busy medium. This modified algorithm is called

CSMA with *collision detection* (CSMA-CD), and forms the basis of the IEEE 802.3 Ethernet MAC layer.

CSMA-CD, while being very efficient for wired media, suffers from a particular problem called the hidden terminal problem when used in radio channels. In a radio system it is quite possible that two terminals are in radio range of a third terminal while the link between the two terminals themselves is opaque. For instance, it is possible that both terminals A and C can talk to B while A and C cannot hear each other's transmission. Under such a scenario it is impossible to detect collisions. This problem is overcome by adopting a collision avoidance mechanism that requires that every interterminal transmission go through a central hub or access port that can hear the transmission of every other terminal. This architecture is incorporated as part of the MAC process in the IEEE 802.11 Wireless LAN standard.

There has been tremendous research in the area of medium access control, and there are a number of algorithms developed to deal with the needs of different applications. Medium access schemes to support integrated services that combine data, voice, and video are still an active area of research. MAC layer design is a challenging topic especially for systems such as wireless ATM LANs, which aim to provide multimedia services in a mobile radio environment.

NETWORK LAYER—ROUTING, CONGESTION, AND FLOW CONTROL

The two major issues that are handled by the network layer are routing and flow control. Routing is one of the most complex and critical functions performed in a packet switching network, and there are scores of schemes developed to handle this function. A good routing scheme should be able to find the best possible route with the least overhead. It should be fairly simple, flexible, and accurate. It should be very robust to localized network problems and make a judicious tradeoff between fairness to all users and maximizing the average throughput by ranking packet exchange between nearby nodes.

All routing schemes use some performance criterion to select a particular route. This criterion could be to minimize the number of hops required to get from source to destination, minimize some cost metric, minimize network delay, or maximize throughput. The measurement of these performance criteria could be based on information obtained from the local node, adjacent nodes, all nodes, or nodes along the route. And the information could be updated on a continuous or periodic basis or as major load or topology changes occur in the network.

The basic routing strategies that are implemented in packet networks may be classified into four broad categories: *fixed routing, flooding, random routing,* and *adaptive routing.* Fixed routing is the most primitive of these techniques. In this case a fixed route is selected for each source destination pair of nodes in the network using some least-cost routing algorithm. Once selected, the routes are fixed and do not dynamically vary based on traffic conditions. Since the routes between source and destination are fixed, there is no difference between datagram and virtual circuit routing in this particular case. Flooding is the other simple routing technique that requires no information about the network traffic. Every node that receives a packet forwards it to every one of its neighboring nodes. This leads to multiple copies of the same packet arriving at the destination via different routes, and there has to be a mechanism to discard all but the first copy. Flooding can lead to incessant retransmissions unless some additional intelligence is built in to prevent repetitious transmission. Flooding is extremely wasteful of link resources but has tremendous robustness against network node failures. Random routing is another simple scheme where a packet received at any node is forwarded to an adjacent node that is randomly picked. This avoids the unnecessary loading associated with flooding, but often the randomly selected route will not be the least-cost one. Assigning a certain probability value to each adjacent node, and forwarding the packets to the one that has the highest probability of success may improve this technique.

Virtually every routing scheme that is in use today employs some form of adaptive routing technique. In adaptive routing, routing decisions change with changes in network conditions. For example, node failures and congestion are constantly monitored and taken into account while making routing decisions. In order to implement adaptive routing, there has to be a mechanism to communicate network condition information between the various nodes in the network. The routing algorithm has to make a trade-off between the amount and precision of the information that is communicated between the nodes, and the overhead required to do so. Adaptive routing is very complex and requires massive processing power at the nodes. While this scheme is very powerful in controlling congestion, it is possible that at times the algorithm may react too quickly and cause congestion-producing oscillations in the network.

The original routing algorithm designed for the first generation Internet in 1969 was a distributed adaptive algorithm called the Bellman–Ford algorithm. This algorithm used estimated delay as the performance criterion, and selected the route that minimized the delay. This algorithm was inaccurate, since the estimation of delay was wholly based on queue length and did not take into account the line speed. In 1979 this algorithm was modified to measure delay directly based on timestamps made on every packet on arrival and departure at every node. With this modification, when a positive acknowledgment was received, delay was measured as the departure time minus the arrival time plus transmission time and propagation delay. Every 10 s each node computed the average delay on each outgoing link, and any significant changes in delay were communicated to other nodes using flooding. As every time delay information was updated, the routing tables were recomputed. This routing scheme worked very well as long as the load on the Internet was light. As the load grew, this strategy created congestion oscillations. Since all nodes were attempting to find the least-delay route at the same time, all traffic shifted to the least-delay route at the same time. This caused congestion at the best route and freed up any previously congested route. This would increase the delay value on the new route and decrease the delay value on the previously congested route, thereby causing a shift back. This oscillation is primarily caused due to the fact that every node was attempting to find the best route. In 1987 a new algorithm was designed that attempted to give the average route a good path instead of attempting to give all routes the best path. This was done by modifying the cost function used to calculate the effect of delay.

DATA COMMUNICATION STANDARDS

We now very briefly cover a few of the important data communications standards. There are a number of standards in existence today, and what is discussed here is only a small representative sample. We only provide some high-level details about the standards such as service objectives and basic functions and do not get into the details of their operation. The interested reader is referred to the books in the bibliography for more details on each of these standards. All the standards described here use the concepts and techniques detailed in the previous sections to accomplish the operations that are specified.

ISDN

Integrated services digital network (ISDN) was developed to provide the user with a single interface that supported a range of different devices simultaneously. The basic ISDN connection consists of two B channels (2B) of 64 kbps each and a single D channel of 16 kbps. The B channels are designed to carry user data, and the D channel is meant for carrying control and signal information. The 2B + D format is known as the *basic rate interface* (BRI). With frame control and other overheads, an ISDN BRI provides a capacity of 192 kbps. A higher rate interface called the *primary rate interface* (PRI) is also available. PRI offers 23B channels and one D channel at 64 kbps giving a total of 1.544 Mbps (TI rate) in North America, and 30B channels and one D channel at 64 kbps giving a total of 2.048 Mbps (E1 rate) in Europe and other parts of the world.

ISDN uses the existing telecommunications dial-up infrastructure, though special ISDN connection interface boxes are required at the user premise. Since its inception in the early 1980s, ISDN has not been very successful, especially in North America. One of the reasons for its poor success was the lack of a "killer application" during its early days. Now, with the growing demand for Internet access, the regional Bell operating companies (RBOCs) are having some success in selling ISDN services. Notwithstanding its poor success in the marketplace, ISDN has provided the industry with many landmark standards such as the LAPD data link control, and the Q.931 messaging protocol. For example, LAPD has been the founding technology for frame relay, and Q.931 is used extensively in other signaling systems such as mobile radio, frame relay, and ATM.

X.25

The X.25 was designed to perform a function similar to that of ISDN in terms of providing an interface between an end-user device and a network. Unlike ISDN, which connects to a circuit switched network, the X.25 connects to a packet switched network. The X.25 was designed by the ITU-T in the early 1970s to define how a public packet data network would handle a user's payload and accommodate various quality of service (QoS) features that are requested by the user. There have been many revisions to the original X.25 standard with the last major revision made in 1988.

The X.25 is not a switching standard specification. It only specifies the interface between a packet network and a user data terminal. It does not concern itself with the internal operations of the network. The X.25 recommendation encompasses the lower three layers of the OSI model. The physical layer uses a V-series, X.21, or X.21bis interface, and the DLC layer uses the LAPB protocol, which is the subset of HDLC.

The X.25 is the oldest packet data standard, and it was designed to be implemented over noisy analog phone lines. It therefore has a lot of built-in error control that is an unnecessary overhead when used in today's relatively low-error links. An X.25 connection supports a number of virtual circuits both in the form of *permanent virtual circuits* (PVCs) and *switched virtual circuits* (SVCs) at data rates from 19.2 kbps to 64 kbps. Even though X.25 has enjoyed considerable success in the past, given that it is an old standard, it will likely see decreasing usage as other technologies evolve.

Frame Relay

The purpose of a frame relay network is to provide a high-speed virtual private network (VPN) capable of supporting high-bit-rate applications. It is designed for modern networks that do not require a lot of error correction. Typical frame relay connections range from 56 kbps to 2 Mbps.

Frame relay, like X.25, implements multiple virtual circuits over a single connection but does so using statistical multiplexing techniques that make efficient use of the available bandwidth and provide flexibility. Frame relay includes a CRC for detecting corrupted bits but does not have any mechanism for error correction. In addition, because many higher-level protocols include their own flow control algorithms, frame relay implements a simple congestion notification mechanism to notify the user when the network is nearing saturation. Frame relays have relatively high initial cost, and are most commonly used for interconnecting remote LANs together.

ATM

Asynchronous transfer mode (ATM) is a high-performance packet switching and multiplexing technology that utilizes fixed length packets (cells) to carry different types of traffic. ATM was envisioned as the technology of future public networks to deliver broadband ISDN services. The fundamental difference between an ATM and other packet switching technologies is that ATM is designed for high-performance multimedia networking with quality of service (QoS) guarantees. ATM has been implemented in a very broad range of networking devices in both LAN and WAN environments. With most public carriers deploying ATM in their backbones, ATM appears to be poised for great success in the WAN environment. In the LAN marketplace, however, it faces strong competition from high-speed Ethernet-based technologies.

ATM is designed for handling large amounts of data across long distances over high-speed backbone. Instead of allocating a dedicated virtual circuit for the duration of each call, ATM assembles data into small packets and statistically multiplexes them according to their traffic characteristics. One of the problems with other protocols that implement virtual connections is that some time slots are wasted if no data are being transmitted. ATM avoids this by dynamically allocating bandwidth for traffic on demand. This means greater utilization of bandwidth and better capacity to handle heavy load situations. When an ATM connection is requested, details about the connection are specified, which allow decisions concerning the route and handling of the data to be made. Typical details are the type of traffic, destination, peak and average bandwidth requirements, a cost factor, and other parameters. Using the information provided, the network assigns priorities to the packets and chooses a route that fits within the cost structure.

The major benefits of ATM technology are its high speed, scalability, and ability to support multimedia applications with QoS guarantee. ATM uses a 53-byte fixed packet size that comprises a 48-byte payload and 5-byte header. The 53-byte size was chosen as a compromise between the needs of low-latency voice application and bandwidth-intensive data applications. The fixed packet size simplifies the switching process, enabling the use of hardware switching, which can be implemented at gigabit rates. ATM is immensely scalable, with data rates specified at 155.52 Mbps and 622.08 Mbps. Other data rates, both lower and higher, are also possible. While ATM was originally designed with fiber-optic channels using a SONET physical layer, there is work going on in the standards bodies to develop ATM standards for twisted-pair and wireless media, albeit at lower rates. For example, one version of the emerging ADSL lines are expected to carry ATM cells at 1 Mbps and above to the end user. There are many research and standards activities in progress to develop a 25 Mbps wireless ATM interface. For instance, the European LAN standard HIPERLAN II will carry ATM cells at 25 Mbps.

WIRELESS DATA STANDARDS

Though wireless data has been around for a number of years, it is yet to gain widespread acceptance. There is, however, a good deal of optimism in the industry that the market for wireless data will see tremendous growth sooner than later. The massive proliferation of handheld computing devices has lead to an increase in demand for services that allow for mobile computing applications. There is also a great demand to provide wireless connectivity to the Internet. Even though most of the present cellular radio traffic is voice, it is projected that a large percentage of future traffic will be data.

One of the early standards developed to provide data services over a cellular network is called cellular digital packet data (CDPD). CDPD was designed as an overlay network to the existing analog cellular telephone network, AMPS. It uses existing data communication protocols such as connectionless network layer protocol (CLNP), the Internet protocol (IP), the OSI transport layer, and the transmission control protocol (TCP). Effectively, existing protocols run at layer 3 and above on top of the CDPD physical and MAC layer protocols (6).

The CDPD operates in two modes: the dedicated channel mode and the hopping mode. In the dedicated channel mode, one of the available 30 kHz analog cellular channels is dedicated for CDPD service. In the hopping mode, the CDPD uses one of the free analog channels and hops away to another free channel when a voice call arrives in the channel that is being used for data. In the hopping mode, the CDPD does not take away any of the capacity of the voice system and only uses the idle periods in each channel. The CDPD provides a data rate of 19.2 kbps, and is typically used for applications such as point-of-sale, telemetry, short messaging, public safety, and transportation. Due to its low data rate, and the infant state of mobile data market, the CDPD has been rather slow in gaining acceptance. Standards are under development for higher-speed mobile data.

While CDPD was designed to run over the analog cellular network, standards are under development for packet data services over the new digital cellular systems. For example, a packet radio standard called general packet radio service (GPRS) is being developed as part of the global system for mobile (GSM) communications digital cellular standard. GPRS is designed with the objective of efficiently accommodating bursty data traffic within the GSM framework. GPRS shares GSM frequency bands with voice and circuit switched data, and makes use of many of the physical layer properties of GSM such as the TDMA frame structure and modulation technique (7). GPRS will be able to provide data services to the end users at a maximum rate of 14.4 kbps per time slot when no error recovery mechanisms are required, and a maximum of 13.2 kbps when error recovery is required. GPRS also allows a single user to acquire all the eight time slots in the GSM TDMA frame and thereby provide up to 110 kbps. GPRS is designed to interwork with other public data networks using the IP, CLNP, and X.25. Packet data standards similar to GPRS are being developed for the IS-136 TDMA and IS-95 CDMA digital cellular systems as well.

Given the slow data rates that can be achieved using the first- and second-generation cellular systems, the applications that can be supported by these systems are very limited. Discussions are under way for the development of a third-generation cellular system that can support much higher speed data. The International Telecommunications Union has proposed that the third-generation cellular system be able to support a data rate of 144 kbps for high-speed mobile users, 384 kbps for low-speed pedestrian users, and 2 Mbps for indoor stationary users. There are many proposals under consideration for evolving the current second-generation systems to the proposed third-generation requirements, and these are being heavily debated under the ITU project code-named IMT-2000.

CONCLUSION

Data communications is a vast subject with a tremendous amount of activity going on in recent years. There are a large variety of systems, technologies, and standards, and there is continuing battle between the proponents of different technology. Demand for data has grown remarkably over the past few years, and it is expected that the data market will surpass the voice communications market in the near future. Most of this demand is spawned by the explosive growth of the Internet coupled with the falling prices of computer and communications equipment. Given the rate at which the technology and market demand for data communications is growing, it will not be too far in the future when broadband access at homes and high-speed access to the Internet from anywhere and anytime become a reality.

BIBLIOGRAPHY

1. D. Bertsekas and R. Gallager, *Data Networks,* Englewood Cliffs, NJ: Prentice-Hall, 1992.
2. W. Stallings, *Data and Computer Communications,* Upper Saddle River, NJ: Prentice-Hall, 1997.

3. G. Held, *Understanding Data Communications,* New York: Wiley, 1997.
4. T. Sheldon, *Encyclopedia of Networking,* New York: McGraw-Hill, 1998.
5. D. Haccoun and S. Pierre, Automatic Repeat Request, in Gibson (ed.), *Communications Handbook.* Boca Raton, FL: CRC Press, 1997.
6. M. S. Taylor, W. Waung, and M. Banan, *Internetwork Mobility—the CDPD Approach,* Upper Saddle River, NJ: Prentice-Hall, 1996.
7. J. Cai and D. Goodman, General Packet Radio Service in GSM, *IEEE Commun. Mag.* 122–131, October 1997.

RIAS MUHAMED
SBC Technology Resources