

SURVEY OF JOB SHOP SCHEDULING TECHNIQUES

INTRODUCTION

In the United States today, there are approximately 40,000 factories producing metal-fabricated parts. These parts end up in a wide variety of products sold here and abroad. These factories employ roughly 2 million people and ship close to \$3 billion worth of products every year. The vast majority of these factories are what we call “job shops”, meaning that the flow of raw and unfinished goods through them is completely random. Over the years, the behavior and performance of these job shops have been the focus of considerable attention in the Operations Research (OR) literature. Research papers on topics such as factory layout, inventory control, process control, production scheduling, and resource utilization can be found in almost every issue of every OR journal on the market today. The most popular of these topics is production (often referred to as job shop) scheduling. Job shop scheduling can be thought of as the allocation of resources over a specified time to perform a predetermined collection of tasks. Job shop scheduling has received this large amount of attention, because it has the potential to dramatically decrease costs and increase throughput, thereby, profits.

A large number of approaches to the modeling and solution of these job shop scheduling problems have been reported in the OR literature, with varying degrees of success. These approaches revolve around a series of technological advances that have occurred over that last 35 years. These include mathematical programming, dispatching rules, expert systems, neural networks, genetic algorithms, and inductive learning. In this article, we take an evolutionary view in describing how these technologies have been applied to job shop scheduling problems. To do this, we discuss a few of the most important contributions in each of these technology areas and the most recent trends.

MATHEMATICAL TECHNIQUES

Mathematical programming has been applied extensively to job shop scheduling problems. Problems have been formulated using integer programming (Balas 1965, 1967), mixed-integer programming (Balas 1969, 1970), and dynamic programming (Srinivasan 1971). Until recently, the use of these approaches has been limited because scheduling problems belong to the class of NP-complete problems. To overcome these deficiencies, a group of researchers began to decompose the scheduling problem into a number of subproblems, proposing a number of techniques to solve them. In addition, new solution techniques, more powerful heuristics, and the computational power of modern computers have enabled these approaches to be used on larger problems. Still, difficulties in the formulation of material flow constraints as mathematical inequalities and the development of generalized software solutions have limited the use of these approaches.

Decomposition strategies

Davis and Jones 1988 proposed a methodology based on the decomposition of mathematical programming problems that used both Benders-type (Benders 1960) and Dantzig/Wolfe-type (Dantzig and Wolfe, 1960) decompositions. The methodology was part of closed-loop, real-time, two-level hierarchical shop floor control system. The top-level scheduler (i.e., the supremal) specified the earliest start time and the latest finish time for each job. The lower level scheduling modules (i.e., the infimals) would refine these limit times for each job by detailed sequencing of all operations. A multicriteria objective function was specified that included tardiness, throughput, and process utilization costs. The decomposition was achieved by first reordering the constraints of the original problem to generate a block angular form, then transforming that block angular form into a hierarchical tree structure. In general, N subproblems would result plus a constraint set that contained partial members of each of the subproblems. The latter was termed the “coupling” constraints, and included precedence relations and material handling. The supremal unit explicitly considered the coupling constraints, while the infimal units considered their individual decoupled constraint sets. The authors pointed out that the inherent stochastic nature of job shops and the presence of multiple, but often conflicting, objectives made it difficult to express the coupling constraints using exact mathematical relationships. This made it almost impossible to develop a general solution methodology. To overcome this, a new real-time simulation methodology was proposed in (Davis and Jones, 1988) to solve the supremal and infimal problems.

Gershwin 1989 used the notion of temporal decomposition to propose a mathematical programming framework for analysis of production planning and scheduling. This framework can be characterized as hierarchical and multi-layer. The problem formulations to control events at higher layers ignored the details of the variations of events occurring at lower layers. The problem formulations at the lower layers view the events at the higher layers as static, discrete events. Scheduling is actually carried out in bottom three layers so that the production requirements imposed by the planning layers can be met. First, a hedging point is found by solving a dynamic programming problem. This hedging point is the number of excess goods that should be produced to compensate for future equipment failures. This hedging point is used to formulate a linear programming problem to determine instantaneous production rates. These rates are then used to determine the actual schedule (which parts to make and when). A variety of approaches are under investigation for generating schedules.

Enumerative techniques and Lagrangian relaxation

Two popular solution techniques for integer-programming problems are branch-and-bound and Lagrangian relaxation. Branch-and-bound is an enumerative technique (Agin 1966, Lawler and Wood 1966). Summarizing Morton and Pentico 1993, “The basic idea of branching is to conceptualize the problem as a decision tree. Each decision choice point—a node—corresponds to a partial solution.

From each node, there grow a number of new branches, one for each possible decision. This branching process continues until leaf nodes, that cannot branch any further, are reached. These leaf nodes are solutions to the scheduling problem". Although efficient bounding and pruning procedures have been developed to speed up the search, this is still a very computational intensive procedure for solving large scheduling problems. If the integer constraint is the main problem, then why not remove that constraint. A technique called Lagrangian relaxation, which has been used for more than 30 years, does just that (Shapiro 1979). Lagrangian relaxation solves integer-programming problems by omitting specific integer-valued constraints and adding the corresponding costs (due to these omissions and/or relaxations) to the objective function. As with branch and bound, Lagrangian relaxation is computationally expensive for large scheduling problems.

Recent trends

Model-Based Optimization (MBO) is an optimization approach that uses mathematical expressions (e.g., constraints and inequalities) to model scheduling problems as mixed integer (non) linear programs (MINLP's) (Zentner et al., 1994). A set of methods such as linear programming, branch-and-bound, and decomposition techniques are used to search the scenario space of solutions. Due to the advances in computer technologies, the computation times are becoming very practical. According to Subrahmanyam et al. 1996 "For problems of moderate size, solutions of type D are given." Solutions of type D are optimal solutions of the maximum desirability possible within the constraints of operation. These approaches are being enhanced by the development of English-like "scheduling languages" and high-level graphical interfaces. The scheduling languages support the developing of the mathematical formulations with minimum intervention from the user.

DISPATCHING RULES

Dispatching rules have been applied consistently to scheduling problems. They are procedures designed to provide good solutions to complex problems in real-time. The term dispatching rule, scheduling rule, sequencing rule, or heuristic are often used synonymously (Panwalker and Islander 1977, Blackstone et al., 1982, Baker 1974). Dispatching rules have been classified mainly according to the performance criteria for which they have been developed. Wu (1987) categorized dispatching rules into several classes. Class 1 contains simple priority rules, which are based on information related to the jobs. Sub-classes are based on the particular piece of information used. Example classes include those based on processing times (such as shortest processing time (SPT)), due dates (such as earliest due date (EDD)), slack (such as minimum slack (MINSLACK)), and arrival times (such as first-in first-out (FIFO)). Class 2 consists of combinations of rules from class one. The particular rule that is implemented can now depend on the situation that exists on the shop floor. A typical example of a rule in this class is, for example, SPT until the queue length exceeds 5, then switch to FIFO. This prohibits

jobs with large processing times from staying in the queue for long periods. Class 3 contains rules that are commonly referred to as Weight Priority Indexes. The idea here is to use more than one piece of information about the jobs to determine the schedule. Pieces of information are assigned weights to reflect their relative importance. Usually, an objective function $f(x)$ is defined. For example, $f(x) = \text{weight}_1 \times \text{Processing Time of Job}(x) + \text{weight}_2 \times (\text{Current Time} - \text{Due Date of Job}(x))$. Then, any time new sequence is needed, the function $f(x)$ is evaluated for each job x in the queue. The jobs are ranked based on this evaluation.

During the last 30 years, the performance of a large number of these rules has been studied extensively using simulation techniques (Montazer and Van Wassenhove, 1990). These studies have been aimed at answering the question: If you want to optimize a particular performance criterion, which rule should you choose? Most of the early work concentrated on the shortest processing time rule (SPT). Conway and Maxwell (1967) were the first to study the SPT rule and its variations. They found that, although some individual jobs could experience prohibitively long flow times, the SPT rule minimized the mean flow time for all jobs. They also showed that SPT was the best choice for optimizing the mean value of other basic measures such as waiting time and system utilization. Many similar investigations have been carried out to determine the dispatching rule which optimizes a wide range of job-related (such as due date and tardiness) and shop-related (such as throughput and utilization) performance measures. This problem of selecting the best dispatching rule for a given performance measure continues to be a very active area of research. However, the research has been expanded to include the possibility of switching rules to address an important problem: error recovery. Two early efforts to address error recovery were conducted by Bean and Birge (1986) and Saleh (1988). Both developed heuristic rules to smooth-out disruptions to the original schedule, thereby creating a match-up with that schedule. Bean and Birge (1986) based their heuristic on Turnpike Theory (McKenzie 1976) to optimize a generalized cost function. Saleh showed that he could minimize duration of the disruption by switching the objective function from mean flow time to makespan based on disjunctive graphs (Adams et al., 1988).

ARTIFICIAL INTELLIGENCE (AI) TECHNIQUES

Starting in the early 80s, a series of new technologies were applied to job shop scheduling problems. They fall under the general title of artificial intelligence (AI) techniques and include expert systems, knowledge-based systems, and several search techniques. Expert and knowledge-based systems were quite prevalent in the early and mid 1980s. They have four main advantages. First, and perhaps most important, they use both quantitative and qualitative knowledge in the decision-making process. Second, they are capable of generating heuristics that are significantly more complex than the simple dispatching rules described above. Third, the selection of the best heuristic can be based on information about the entire job shop including the current jobs, expected new jobs, and the current status of re-

sources, material transporters, inventory, and personnel. Fourth, they capture complex relationships in elegant new data structures and contain special techniques for powerful manipulation of the information in these data structures. There are, however, serious disadvantages. They can be time consuming to build and verify, as well as difficult to maintain and change. Moreover, since they generate only feasible solutions, it is rarely possible to tell how close that solution is to the optimal solution. Finally, since they are tied directly to the system they were built to manage, there is no such thing as a generic AI system.

Expert/knowledge-based systems

Expert and knowledge-based systems consist of two parts: a knowledge base, and inference engine to operate on that knowledge base. Formalizations of the “knowledge” that human experts use—rules, procedures, heuristics, and other types of abstractions—are captured in the knowledge base. Three types of knowledge are usually included: procedural, declarative, and meta. Procedural knowledge is domain-specific problem solving knowledge. Declarative knowledge provides the input data defining the problem domain. Meta knowledge is knowledge about how to use the procedural and declarative knowledge to actually solve the problem. Several data structures have been utilized to represent the knowledge in the knowledge base including semantic nets, frames, scripts, predicate calculus, and production rules. The inference engine selects a strategy to apply to the knowledge bases to solve the problem at hand. It can be forward chaining (data driven) or backward chaining (goal driven).

ISIS (Fox 1983) was the first major expert system aimed specifically at job shop scheduling problems. ISIS used a constraint-directed reasoning approach with three constraint categories: organizational goals, physical limitations and causal restrictions. Organizational goals considered objective functions based on due-date and work-in-progress. Physical limitations referred to situations where a resource had limited processing capability. Procedural constraints and resource requirements were typical examples of the third category. Several issues with respect to constraints were considered such as constraints in conflict, importance of a constraint, interactions of constraints, constraint generation and constraint obligation. ISIS used a three level, hierarchical, constraint-directed search. Orders were selected at level 1. A capacity analysis was performed at level 2 to determine the availability of the resources required by the order. Detailed scheduling was performed at level 3. ISIS also provided for the capability to interactively construct and alter schedules. In this capacity, ISIS utilized its constraint knowledge to maintain the consistency of the schedule and to identify scheduling decisions that would result in poorly satisfied constraints.

Wysk et al. (1986) developed an integrated expert system/simulation scheduler called MPECS. The expert system used both forward and backward chaining to select a small set of potentially good rules from predefined set of dispatching rules and other heuristics in the knowledge base. These rules optimized a single performance measure, although that measure could change from one scheduling

period to the next. The selected rules were then evaluated one at a time using a deterministic simulation of a laboratory manufacturing system. After all of the rules were evaluated, the best rule was implemented on the laboratory system. Data could be gathered about how the rule actually performed and used to update the knowledge base off-line. They were able to show that periodic rescheduling makes the system more responsive and adaptive to a changing environment. MPECS was important for several reasons. It was the first hybrid system to make decisions based on the actual feedback from the shop floor. It incorporated some learning into its knowledge base to improve future decisions. The same systems could be used to optimize several different performance measures. Finally, it utilized a new multi-step approach to shop floor scheduling.

Other examples of expert/knowledge-based scheduling systems developed OPIS (Opportunistic Intelligent Scheduler) (Smith 1995), and SONIA (Le Pape 1995).

Distributed AI: agents

Due to the limited knowledge and the problem solving ability of a single expert or knowledge based system, these AI approaches have difficulty solving large scheduling problems as well. To address this, AI researchers have also begun to develop distributed scheduling system approaches (Parunak et al., 1985). They have done this by an application of their well-known “divide and conquer” approach. This requires a problem decomposition technique, such as those described above, and the development of different expert/knowledge-based systems that can cooperate to solve the overall problem (Zhang and Zhang, 1995). The AI community’s answer is the “agent” paradigm. An agent is a unique software process operating asynchronously with other agents. Agents are complete knowledge-based systems by themselves. The set of agents in a system may be heterogeneous with respect to long-term knowledge, solution-evaluation criteria, or goals, as well as languages, algorithms, hardware requirements. Integrating agents selected from a “library” creates a multi-agent system.

For example, one such multi-agent system could involve two types of agents: tasks and resources. Each task agent might be responsible for scheduling a certain class of tasks such as material handling, machining, or inspection, on those resources capable of performing those tasks. This can be done using any performance measure related to tasks, such as minimize tardiness, and any solution technique. Each resource agent might be responsible for a single resource or a class of resources. Task agents must send their resource requests to the appropriate resource agent, along with the set of operations to be performed by that resource (Daouas et al., 1995). Upon receipt of such a request, the resource agent must generate a new schedule using its own performance measures, such as maximize utilization, which includes this request. The resource agent will use the results to decide whether to accept this new request or not. To avoid the situation where no resource will accept a request, coordination mechanisms must be developed. There are, now, no general guidelines for the design and implementation of this coordination. Therefore, the debates about centralized vs. decentralized approaches to job

shop scheduling go on. The agents' formalism may provide an answer to these debates.

ARTIFICIAL NEURAL NETWORKS

Neural networks, also called connectionist or distributed parallel processing models, have been studied for many years in an attempt to mirror the learning and prediction abilities of human beings. Neural network models are distinguished by network topology, node characteristics, and training or learning rules. An example of a three-layer, feed-forward neural network is shown in Figure 1.

Supervised learning neural networks

Through exposure to historical data, supervised learning neural networks attempt to capture the desired relationships between inputs and the outputs. Back-propagation is the most popular and widely used supervised training procedure. Back-propagation (Rumelhart et al., 1986, Werbos 1995) applies the gradient-descent technique in the feed-forward network to change a collection of weights so that some cost function can be minimized. The cost function, which is only dependent on weights (W) and training patterns, is defined by:

$$C(W) = \frac{1}{2} \sum (T_{ij} - O_{ij}) \quad (1)$$

where the T is the target value, O is the output of the network, i represents the output nodes, and j represents the training patterns.

After the network propagates the input values to the output layer, the error between the desired output and actual output will be "back-propagated" to the previous layer. In the hidden layers, the error for each node is computed by the weighted sum of errors in the next layer's nodes. In a three-layered network, the next layer means the output layer. The activation function is usually a sigmoid function with the weights modified according to

$$\Delta W_{ij} = \eta X_j (1 - X_j) (T_j - X_j) X_i \quad (2)$$

or

$$\Delta W_{ij} = \eta X_j (1 - X_j) \left(\sum \delta_k W_{jk} \right) X_i \quad (3)$$

where W_{jk} is weight from node i to node (e.g., neuron) j , η is the learning rate, X_j is the output of node j , T_j is the target value of node j , and δ_k is the error function of node k .

If j is in the output layer, Eq. (2) is used. If j is the hidden layers, Eq. (3) is used. The weights are updated to reduce the cost function at each step. The process continues until the error between the predicted and the actual outputs is smaller than some predetermined tolerance.

Rabelo (1990) was the first to use back-propagation neural nets to solve job shop scheduling problems with several job types, exhibiting different arrival patterns, process plans, precedence sequences and batch sizes. Training examples were generated to train the neural network to select the correct characterization of the manufacturing environments suitable for various scheduling policies and the chosen performance criteria. In order to generate training samples, a performance simulation of the dispatching rules

available for the manufacturing system was carried out. The neural networks were trained for problems involving 3, 4, 5, 8, 10, and 20 machines. To carry out this training, a special, input-feature space was developed. This space contained both job characteristics (such as types, number of jobs in each type, routings, due dates, and processing times) and shop characteristics (such as number of machines and their capacities). The output of the neural network represented the relative ranking of the available dispatching rules for that specific scheduling problem and the selected performance criteria. The neural networks were tested in numerous problems and their performance (in terms of minimizing Mean Tardiness) was always better than each single dispatching rule (25% to 50%).

Relaxation models

Neural networks based on relaxation models are defined by energy functions. They are pre-assembled systems that relax from input to output along a predefined energy contour. Hopfield neural networks (Hopfield and Tank 1985) are a classical example of a relaxation model that has been used to solve some classic, textbook scheduling problems (Foo and Takefuji, 1988). Two-dimensional Hopfield networks were used to solve 4-job, 3-machine problems and 10-job, 10-machine problems (Zhou et al., 1990). They were extended in (Lo and Bavarian, 1991) to 3 dimensions to represent jobs ($i=1, \dots, J$), machines $j=1, \dots, J$, and time ($m=1, \dots, M$). In each case, the objective was to minimize the makespan, total time to complete all jobs, which is defined as

$$E = \sum_{j=1} \sum_{i=1} \sum_{m=1} (v_{ijm})(m + T_{ij} - 1) \quad (4)$$

where v_{ijm} is the output (1 or 0) of neuron ijm , and T_{ij} is the time required by j^{th} resource (e.g., machine) to complete the i^{th} job.

Due to a large number of variables involved in generating a feasible schedule, these approaches tend to be computationally inefficient and frequently generate infeasible solutions. Consequently, they have not been used to solve realistic scheduling problems.

Temporal reinforcement learning

It was noted above that supervised learning neural networks attempt to capture the desired relationships between inputs and the outputs through exposure to training patterns. However, for some problems, the desired response may not always be available during the time of learning. When, the desired response is obtained, changes to the neural network are performed by assessing penalties for the scheduling actions previously decided by the neural network. As summarized by Tesauro (1992), "In the simplest form of this paradigm, the learning system passively observes a temporal sequence of input states that eventually leads to a final reinforcement or reward signal (usually a scalar). The learning system's task in this case is to predict expected reward given an observation of an input state or sequence of input states. The system may also be set up so that it can generate control signals that influence the sequence of states." For scheduling, the learning task is to

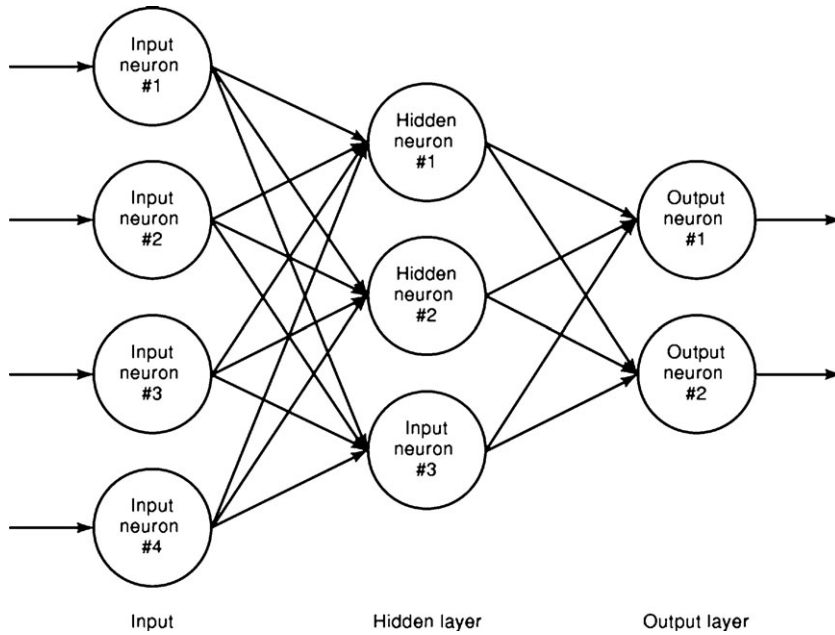


Figure 1. An example of a three-layer, feed-forward neural network.

produce an scheduling action that will lead to minimizing (or maximizing) the performance measure (e.g., makespan, tardiness) based on the state of the system (e.g., inventories, machine status, routings, due dates, layouts). Several procedures have been developed to train neural networks when the desired response is not available during the time of learning. Rabelo et al. (1994) utilized a procedure developed by Watkins (1989), denominated Q-learning, to implement a scheduling system to solve dynamic job shop scheduling problems. The scheduling system was able to follow trends in the shop floor and select a dispatching rule that provided the maximum reward according to performance measures based on tardiness and flow time. On the other hand, Zhang and Dietterich (1996) utilized a procedure developed by Sutton (1988) called TD(λ) to schedule payload processing of NASA's space shuttle program.

NEIGHBORHOOD SEARCH METHODS

Neighborhood search methods are very popular. Neighborhood search methods provide good solutions and offer possibilities to be enhanced when combined with other heuristics. Wilkerson and Irwin (1971) developed one of the first neighborhood procedures. This method iteratively added small changes ("perturbations") to an initial schedule, which is obtained by any heuristic. Conceptually similar to hill climbing, these techniques continue to perturb and evaluate schedules until there is no improvement in the objective function. When this happens, the procedure is ended. Popular techniques that belong to this family include Tabu search, simulated annealing, and genetic algorithms. Each of these has its own perturbation methods, stopping rules, and methods for avoiding local optimum.

Tabu search

The basic idea of Tabu search (Glover 1989, 1990) is to explore the search space of all feasible scheduling solutions

by a sequence of moves. A move from one schedule to another schedule is made by evaluating all candidates and choosing the best available, just like gradient-based techniques. Some moves are classified as tabu (i.e., they are forbidden) because they either trap the search at a local optimum, or they lead to cycling (repeating part of the search). These moves are put onto something called the Tabu List, which is built up from the history of moves used during the search. These tabu moves force exploration of the search space until the old solution area (e.g., local optimum) is left behind. Another key element is that of freeing the search by a short term memory function that provides "strategic forgetting". Tabu search methods have been evolving to more advanced frameworks that includes longer term memory mechanisms. These advanced frameworks are sometimes referred as Adaptive Memory Programming (AMP, Glover 1996).

Tabu search methods have been applied successfully to scheduling problems and as solvers of mixed integer programming problems. Nowicki and Smutnicki (Glover 1996) implemented tabu search methods for job shop and flow shop scheduling problems. Vaessens (Glover 1996) showed that tabu search methods (in specific job shop scheduling cases) are superior over other approaches such as simulated annealing, genetic algorithms, and neural networks.

Simulated annealing

Simulated annealing is based on the analogy to the physical process of cooling and recrystallization of metals. The current state of the thermodynamic system is analogous to the current scheduling solution, the energy equation for the thermodynamic system is analogous to the objective function, and the ground state is analogous to the global optimum. In addition to the global energy J , there is a global temperature T , which is lowered as the iterations progress. Using this analogy, the technique randomly generates new schedules by sampling the probability distribution of the

system (Kirkpatrick et al., 1983):

$$P_j \propto \exp(-T(\Delta J_{\text{best}} - \Delta J_j)/K) \quad (5)$$

where P_j represents the probability of making move j from among the neighborhood choices. ΔJ_{best} represents the improvement of the objective function for the best choice, and ΔJ_j represents the improvement for choice j . K is a normalization factor. Since increases of energy can be accepted, the algorithm is able to escape local minima.

Simulated annealing has been applied effectively to job shop scheduling problems. Vakharia and Chang (1990) developed a scheduling system based on simulated annealing for manufacturing cells. Jeffcoat and Bulfin (1993) applied simulated annealing to a resource-constrained scheduling problem. Their computational results indicated that the simulated annealing procedure provided the best results in comparison with other neighborhood search procedures.

Genetic algorithms

Genetic algorithms (GA) are an optimization methodology based on a direct analogy to Darwinian natural selection and mutations in biological reproduction. In principle, genetic algorithms encode a parallel search through concept space, with each process attempting coarse-grain hill climbing (Goldberg 1988). Instances of a concept correspond to individuals of a species. Induced changes and recombinations of these concepts are tested against an evaluation function to see which ones will survive to the next generation. The use of genetic algorithms requires five components:

1. A way of encoding solutions to the problem—fixed length string of symbols.
2. An evaluation function that returns a rating for each solution.
3. A way of initializing the population of solutions.
4. Operators that may be applied to parents when they reproduce to alter their genetic composition such as crossover (i.e., exchanging a randomly selected segment between parents), mutation (i.e., gene modification), and other domain specific operators.
5. Parameter setting for the algorithm, the operators, and so forth.

A number of approaches have been utilized in the application of genetic algorithms (GA) to job shop scheduling problems (Davis 1985, Goldberg and Lingle 1985, Starkweather et al., 1992):

1. Genetic algorithms with blind recombination operators have been utilized in job shop scheduling. Their emphasis on relative ordering schema, absolute ordering schema, cycles, and edges in the offsprings will arise differences in such blind recombination operators.
2. Sequencing problems have been addressed by mapping their constraints to a Boolean satisfiability problem using partial payoff schemes. This scheme has

produced good results for very simple problems.

3. Heuristic genetic algorithms have been applied to job shop scheduling. In these genetic schemes, problem specific heuristics are incorporated in the recombination operators (such as optimization operators based).

Starkweather et al. (1993) were the first to use genetic algorithms to solve a dual-criteria job shop scheduling problem in a real production facility. Those criteria were the minimization of average inventory in the plant and the minimization of the average waiting time for an order to be selected. These criteria are negatively correlated (The larger the inventory, the shorter the wait; the smaller the inventory, the longer the wait.). To represent the production/shipping optimization problem, a symbolic coding was used for each member (chromosome) of the population. In this scheme, customer orders are represented by discrete integers. Therefore, each member of the population is a permutation of customer orders. The Genetic Algorithm used to solve this problem was based on blind recombinant operators. This recombination operator emphasizes information about the relative order of the elements in the permutation, because this impacts both inventory and waiting time. A single evaluation function (a weighted sum of the two criteria) was utilized to rank each member of the population. That ranking was based on an on-line simulation of the plant operations. This approach generated schedules that produced inventory levels and waiting times that were acceptable to the plant manager. In addition, the integration of the genetic algorithm with the on-line simulation made it possible to react to system dynamics.

These applications have emphasized the utilization of genetic algorithms as a “solo” technique. This has limited the level of complexity of the problems solved and their success. Recent research publications have demonstrated the sensitivity of genetic algorithms to the initial population. When the initial population is generated randomly, genetic algorithms are shown to be less efficient than the annealing-type algorithms, but better than the heuristic methods alone. However, if the initial population is generated by a heuristic, the genetic algorithms become as good as, or better than the annealing-type algorithms. In addition, integration with other search procedures (e.g., tabu search) has enhanced the capabilities of both. This result is not surprising, as it is consistent with results from nonlinear optimization. Simply stated, if you begin the search close to the optimal solution you are much more likely to get the optimum than if you begin the search far away.

FUZZY LOGIC

Fuzzy set theory has been utilized to develop hybrid scheduling approaches. Fuzzy set theory can be useful in modeling and solving job shop scheduling problems with uncertain processing times, constraints, and set-up times. These uncertainties can be represented by fuzzy numbers that are described by using the concept of an interval of confidence. These approaches usually are integrated with other methodologies (e.g., search procedures, constraint re-

laxation). For example, Slany (1994) stresses the imprecision of straight-forward methods presented in the mathematical approaches and introduces a method known as fuzzy constraint relaxation, which is integrated with a knowledge-based scheduling system. His system was applied to a steel manufacturing plant. Grabot and Geneste (1994) use fuzzy logic principles to combine dispatching rules for multi-criteria problems. On the other hand, Krucky (1994) addresses the problem of minimizing setup times of a medium-to-high product mix production line using fuzzy logic. The heuristic, fuzzy logic based algorithm described helps determine how to minimize setup time by clustering assemblies into families of products that share the same setup by balancing a product's placement time between multiple-high-speed placement process steps. Tsujimura et al. (1993) presented a hybrid system, which uses fuzzy set theory to model the processing times of a flow shop scheduling facility. Triangular Fuzzy Numbers (TFNs) are used to represent these processing times. Each job is defined by two TFNs, a lower bound and an upper bound. A branch and bound procedure is utilized to minimize makespan.

HLA Scheduling

Linn, et. al (2002) developed a research tool to facilitate the development of distributed simulations that will then be used to analyze and solve various manufacturing related problems. They use the High Level Architecture/Run Time Infrastructure (HLA/RTI) software (developed by the Defense Modeling and Simulation Office), the Distributed Manufacturing Simulation (DMS) Adapter (developed by the National Institute of Standards and Technology) and ARENA simulation tool.

Their goal is to develop a transporter simulation model that can be easily integrated with other simulators through Internet using the DMS Adaptor that will handle the transmission, receipt and internal updates to all the federation object model (FOM) objects used by a federate. They aim at describing the development of a distributed simulation model for a transportation system that includes multiple transporters. DMS Adapter is imbedded in order to establish information flow exchange between simulation models.

The basic problem in Production Planning is to determine the type and quantity of the products to produce, to meet uncertain demand in the future time periods. Hierarchical production planning (HPP) provides a formal bridge between long-term plans and short-term schedules. HPP was proposed first by Hax and Meal (1975) to solve this problem. Venkateswaran et. al (2005) confirm that the basic advantages of the hierarchical approach to production planning (Vicens et al. 2001) include reduction in complexity, gradual absorption of random events, increased insight due to the use of aggregated figures, reduced need for detailed information, and better forecasting.

The solution techniques depend on the scope and the specific manufacturing scenario (Venkateswaran et al 2005). They include heuristics based on linear programming (LP), stochastic programming, Enterprise Resource Planning (ERP) tools, and optimization coupled with simulation-based evaluation. Those techniques have

their drawbacks; using deterministic data at the aggregate level, for example, does not account for the stochastic nature of the system. A major drawback of these techniques is that they require reruns in the case of unexpected external or internal events so any exception (such as machine failures, new order arrivals) leads to the regeneration of the entire plan.

Venkateswaran, et al (2004) looked at the impacts of planning level decisions on the scheduling function and scheduling level decisions on the planning function. They focused on developing an integrated production plan and schedule for an enterprise. HPP as proposed by Hax and Meal (1975) separates the planning problem into different sub-problems based on the length of planning horizon, time and cost. Hence, Venkateswaran, et al (2004) proposed a two level HPP architecture in which Systems Dynamics (SD) is used to model the behavior of the system while Discrete Event Simulation (DES) is used to study its design and operations. In their approach, SD provides a framework to understand the operations of complex dynamic systems and view the impact of decisions on the entire enterprise. DES models the uncertainty in the manufacturing systems. The decision maker outputs two weights: the weight for the WIP factor (α), and the weight for the inventory factor (β), that is then supplied to the SD model to calculate the weekly production order quantities.

Venkateswaran, et al (2004) used High Level Architecture (HLA) RunTime Infrastructure (RTI) to enable the interfacing between the SD (PowerSim®) and DES (Arena®) models. The distributed Manufacturing Simulation (DMS) adapter developed by NIST has been employed to interface the simulation models with the HLA/RTI.

In the past decade, there has been a renewed interest in the development of a feasible yet optimal production plan and schedule, given the uncertainties of the production environment (Venkateswaran et al. 2005). Those, however, lacked the interaction between planning and scheduling.

Venkateswaran et al. (2005) build on the work of Venkateswaran, et al (2004). They provide a novel approach in solving the hierarchical production planning problem using system dynamic and discrete event simulations coupled with optimization.

They use Meta-heuristics coupled with simulation to determine the optimal queue rule to be used within the shop. Meta-heuristics are employed to find near-optimal solutions for complex optimization problems, such as scheduling. The decision variables of the meta-heuristics are the queue rules that control the part release to each machine. The best set of queue rules is determined by evaluating the rules using a simulation model. Rule-based heuristic is employed to measure the actual performance against the expected performance.

DES is used to describe the most complex manufacturing systems, include stochastic elements which cannot be described easily by mathematical or analytical models, track the status of individual entities and resources in the facility and estimate numerous performance measures associated with those entities. These are highly important for the detailed scheduling level. Moreover, the models can be changed easily and run quickly to reflect changes that occur in the real shop.

The HLA establishes common simulation architecture to facilitate the interoperability of all types of models and simulations. The RTI software implements the specification and represents one of the tangible products of the HLA. RTI provides Object data exchange and time management services.

SWARM

Veeramachaneni et al. (2004) presented a swarm intelligence based approach for optimal scheduling in sensor networks. Sensors are characterized by their transaction times and interdependencies. They used particle swarm optimization (PSO) to solve the problem and reach an optimal schedule. Given a set of sensors for a particular mission, their goal was to find the optimal sensor schedule in sensor networks, constrained by many practical issues such as communication concurrency constraints, interdependencies, and limited computational power at sink. The PSO generates the optimal schedule off-line, which is used by the sensor manager to schedule the sensors in real time. The optimal schedule can then be determined by finding the maximum clique and partitioning the graph into cliques. A clique is a complete subgraph of the graph, where each vertex pair is joined by an edge. Veeramachaneni et al. (2004) designed a cost function to evaluate the maximum cliques and hence select the optimal clique for the specific application. The evaluation is closely attached to reducing the overall transaction time.

The particle swarm optimization algorithm has come to be widely used as a problem solving method in engineering and computer science and has since proven to be a powerful competitor to genetic algorithms (Veeramachaneni et al. 2004).

The particles are entered into a multidimensional search space where each particle represents a possible solution to the multidimensional problem. The particles move based on two factors: its own best solution and any particle's best solution. As the particles traverse through the search space, each particle stores its own best solution in memory and experiences a pull towards this position (called pbest). Each particle also stores the global best solution in memory and experiences a pull towards that position (called gbest). The pbest and gbest are updated after each iteration if a more dominating solution, in terms of fitness, is found by the particle and by the population, respectively. The process continues until either the desired result is achieved or the computational power is exhausted.

According to Veeramachaneni et al. (2004), the PSO formulae define each particle in the D-dimensional space as $X_i = (x_{i1}, x_{i2}, \dots, x_{iD})$ where the subscript i represents the particle number, and the second subscript is the dimension. The memory of the previous best position is represented as $P_i = (p_{i1}, p_{i2}, \dots, p_{iD})$ and a velocity along each dimension as $V_i = (v_{i1}, v_{i2}, \dots, v_{iD})$. After each iteration, the velocity term is updated, and the particle is pulled in the direction of its own best position, P_i , and the global best position, P_g , found so far. The PSO algorithm takes the dependencies and the transaction times as inputs.

Veeramachaneni et al. (2004) report that the algorithm has successfully found the optimal solutions leading to minimum overall transaction times for the system.

Schill, Zimmer, and Trumppf (2005) Schill et al. (2005) study a new TDMA scheduling problem that tries to minimize the duration of total information exchange throughout a multi-hop wireless network. Their goal is to achieve an algorithm for a robust, decentralized ad hoc network that minimizes latency, maximizes usable bandwidth, and allows exchanging control parameters throughout the swarm in minimal time. Communication in swarms of robots is a continuous flow of messages and updates, to distribute control parameters, environmental information, and to fuse sensor data.

Schill et al. (2005) state that one of the biggest advantages of swarms is that they do not require collision detection, which allows the use of simpler hardware. Swarms also maximize the usage of the available bandwidth. A very important problem especially in swarm control and formation control is the exchange of certain parameters throughout the entire network.

Schill et al. (2005) claim that the results for the omnicast problem, a new problem that implements many-to-many communication and is similar to a concurrent multiple broadcast from every node to every other node, can be applied to TDMA scheduling, in order to achieve information dissemination in minimal time. The most important next step is to apply the results of the global analysis of optimal omnicast to distributed, local TDMA scheduling algorithms.

REACTIVE SCHEDULING

Reactive scheduling is generally defined as the ability to revise or repair a complete schedule that has been "overtaken" by events on the shop floor (Zweben et al., 1995). Such events include rush orders, excessive delays, and broken resources. There are two approaches: reactive repair and the proactive adjustment. In reactive repair, the scheduling system waits until an event has occurred before it attempts to recover from that event. The match-up techniques described in section 3 fall into this category. Proactive adjustment requires a capability to monitor the system continuously, predict the future evolution of the system, do contingency planning for likely events, and generate new schedules, all during the execution time of the current schedule. The work of Wysk et al. (1986) and Davis and Jones (1988) fall into this category. Approaches that are more recent utilize artificial intelligence and knowledge-based methodologies (Smith 1995). Still most of the AI approaches propose a quasi-deterministic view of the system, i.e., a stochastic system featuring implicit and/or explicit causal rules. The problem formulation used does not recognize the physical environment of the shop floor domain where interference not only leads to readjustment of schedules but also imposes physical actions to minimize them.

LEARNING IN SCHEDULING

The first step in developing a knowledge base is knowledge acquisition. This in itself is a two step process: get the knowledge from knowledge sources and store that knowledge in digital form. Much work has been done in the area of knowledge acquisition, such as protocol analysis and interactive editing (Shaw et al., 1992). Knowledge sources may be human experts, simulation data, experimental data, databases, and text. In scheduling problems, the knowledge sources are likely to be human experts or simulation data. To extract knowledge from these two sources, the machine learning technique that learns from examples (data) becomes a promising tool. Inductive learning is a state classification process. If we view the state space as a hyperplane, the training data (consisting of conditions and decisions) can be represented as points on the hyperplane. The inductive learning algorithm seeks to draw lines on the hyperplane based on the training data to divide the plane into several areas within which the same decision (conclusion) will be made.

One algorithm that has been implemented in inductive aids and expert system shells is that developed by Quinlan (1986), called Iterative Dichotomiser 3 or ID3. ID3 uses examples to induce production rules (e.g. IF ... THEN ...), which form a simple decision tree. Decision trees are one way to represent knowledge for the purpose of classification. The nodes in a decision tree correspond to attributes of the objects to be classified, and the arcs are alternative values for these attributes. The end nodes of the tree (leaves) indicate classes to which groups of objects belong. Each example is described by attributes and a resulting decision. To determine a good attribute to partition the objects into classes, entropy is employed to measure the information content of each attribute, and then rules are derived through a repetitive decomposition process that minimizes the overall entropy. The entropy value of attribute A_k can be defined as

$$H(A_k) = \sum_{j=1}^{M_k} P(a_{kj}) \left\{ - \sum_{i=1}^N P(c_i | a_{kj}) \log_2 P(c_i | a_{kj}) \right\} \quad (6)$$

where $H(A_k)$ is the entropy value of attribute A_k , $P(a_{kj})$ is the probability of attribute k being at its j^{th} value, $P(c_i | a_{kj})$ is the probability that the class value is c_i when attribute k is at its j^{th} value, M_k is the total number of values for attribute A_k , and N is the total number of different classes (outcomes).

The attribute with the minimum entropy value will be selected as a node in the decision tree to partition the objects. The arcs out of this node represent different values of this attribute. If all the objects in an arc belong to one class, the partition process stops. Otherwise, another attribute will be identified using entropy values to further partition the objects that belong to this arc. This partition process continues until all the objects in an arc are in the same class. Before applying this algorithm, all attributes that have continuous values need to be transformed to discrete values.

In the context of job shop scheduling, the attributes represent system status and the classes represent the dis-

patching rules. Very often, the attribute values are continuous. Yih (1988) proposed a trace-driven knowledge acquisition (TDKA) methodology to deal with continuous data and to avoid the problems occurring in verbally interviewing human experts. TDKA learns scheduling knowledge from expert schedulers without a dialogue with them. There are three steps in this approach. In Step 1, an interactive simulator is developed to mimic the system of interest. The expert will interact with this simulator and make decisions. The entire decision making process will be recorded in the simulator and can be repeated for later analysis. The series of system information and the corresponding decision collected is called a "trace." Step 2 analyzes the "trace" and forms classification rules to partition the trace into groups. The partition process stops when most of the cases in each group use the same dispatching rule (error rate is below the threshold defined by the knowledge engineer). Then, the decision rules are formed. The last step is to verify the generated rules. The resulting rule base is used to schedule jobs in the simulator. If it performs as well as or better than the expert, the process stops. Otherwise, the threshold value is increased, and the process returns to Step 2.

As the job shop operates over time, it is important to be able to modify the knowledge contained in these rule bases. Chiu (1994) looks at knowledge modification for job shop scheduling problems by a framework of dynamic scheduling schemes that explores routing flexibility and handles uncertainties. The proposed methodology includes three modules: discrete-event simulation, instance generation, and incremental induction. First, a simulation module is developed to implement the dynamic scheduling scheme, to generate training examples, and to evaluate the methodology. Second, in an instance-generation module, the searching of good training examples is successfully fulfilled by a genetic algorithm. Finally, in an incremental-induction module, a tolerance-based incremental learning algorithm is proposed to allow continuous learning and facilitate knowledge modification. This algorithm uses entropy values to select attributes to partition the examples where the attribute values are continuous. The tolerance is used to maintain the stability of the existing knowledge while the new example is introduced. The decision tree will not be reconstructed unless there is enough momentum from the new data, that is, the change of the entropy value becomes significant. The experimental results showed that the tolerance-based incremental learning algorithm cannot only reduce the frequency of modifications, but also enhances the generalization ability of the resulting decision tree in a distributed job shop environment.

THEORY OF CONSTRAINTS

The Theory of Constraints (TOC) developed by Eliyahu Goldratt (1990, 1992) is the underlying philosophy for synchronized manufacturing. Goldratt (1990) defined synchronized manufacturing as any systematic method that attempts to move material quickly and smoothly through the production process in concert with market demand. A core concept to TOC is the idea that a few critical constraints exist. Goldratt contends that there is only one constraint in

a system at any given time. As defined by Dettmer (1997), a constraint is “any element of a system or its environment that limits the output of the system”. A constraint will prevent increases in throughput regardless of improvements made to the system. The best schedule is obtained by focusing on the planning and scheduling of these constraint operations. In essence, the constraint operations become the basis from which the entire schedule is derived. TOC has several important concepts and principles. Among them (Goldratt 1990, 1992):

1. Systems function like chains.
2. The system optimum is not the sum of the local optima.
3. The effect-cause-effect method identifies constraints.
4. System constraints can be either physically or policy.
5. Inertia is the worst enemy of a process of ongoing improvement.
6. Throughput is the rate at which the entire system generates money through sales.
7. Inventory is all the money the system invests in things it intends to sell.
8. Operating expense is all the money the system spends turning inventory into throughput.

The general process of TOC is as follows (Goldratt 1990):

1. Identify the systems’ constraints.
2. Decide how to exploit the system’s constraints.
3. Subordinate everything else to the above decision.
4. Elevate the system’s constraints.
5. If in the previous steps a constraint have been broken, go back to Step 1, but do not allow inertia to cause a system constraint.

TOC has been successfully applied to scheduling problems (Academic and Industrial) (Advanced Manufacturing Research, Inc. 1996). Its tools that comprised five distinct logic trees (explained extensively in (Dettmer 1997)) are the Current Reality Tree, the Evaporating Cloud Diagram, the Future Reality Tree, the Prerequisite Tree, and the Transition Tree. These trees are tied to the Categories of Legitimate Reservation (that provide the logic to guide the construction of the trees). These tools have not only been used in production scheduling but also in other enterprise functions such as marketing and sales.

SUMMARY AND CONCLUSIONS

Since job shop scheduling problems fall into the class of NP-complete problems, they are among the most difficult to formulate and solve. Operations Research analysts and engineers have been pursuing solutions to these problems for more than 35 years, with varying degrees of success.

While they are difficult to solve, job shop scheduling problems are among the most important because they impact the ability of manufacturers to meet customer demands and make a profit. They also impact the ability of

autonomous systems to optimize their operations, the deployment of intelligent systems, and the optimizations of communications systems. For this reason, operations research analysts and engineers will continue this pursuit well into the next century.

BIBLIOGRAPHY

- Adams, J., E. Balas and D. Zawack (1988), “The shifting bottleneck procedure for job shop scheduling,” *Management Science*, **34** (3), 391–401.
- Advanced Manufacturing Research, Inc. (1996), “Advanced planning and scheduling systems: just a fad or a breakthrough in manufacturing and supply chain management?,” *The Report on Manufacturing*, December 1996.
- Agin, N. (1966), “Optimum seeking with branch and bound,” *Management Science*, **13**, 176–185.
- Baker, K. (1974), *Introduction to Sequencing and Scheduling*, New York: John Wiley & Sons.
- Balas, E. (1965), “An additive algorithm for solving linear programs with zero-one variables,” *Operations Research*, **13**: 517–546.
- Balas, E. (1967), “Discrete programming by the filter method,” *Operations Research*, **15**, 915–957.
- Balas, E. (1969), “Machine sequencing via disjunctive graphs: An implicit enumeration algorithm,” *Operations Research*, **17**: 1–10.
- Balas, E. (1970), “Machine sequencing: disjunctive graphs and degree-constrained subgraphs,” *Naval Research Logistics Quarterly*, **17**, 941–957.
- Bean, J. and J. Birge (1986), “Match-up real-time scheduling,” *NBS Special Publication*, **724**: 197–212.
- Benders, J. (1960), “Partitioning procedures for solving mixed-variables mathematical programming problems,” *Numerische Mathematik*, **4** (3): 238–252.
- Blackstone, J., D. Phillips and G. Hogg (1982), “A state-of-the-art survey of dispatching rules for manufacturing job shop operations,” *International Journal of Production Research*, **20** (1): 27–45.
- Chiu, C. (1994), “A Learning-Based Methodology for Dynamic Scheduling in Distributed Manufacturing Systems,” Ph.D. Dissertation, Purdue University.
- Conway, R. and W. Maxwell (1967), *Theory of Scheduling*, Reading, Massachusetts: Addison-Wesley.
- Dantzig, G. and P. Wolfe (1960), “Decomposition principles for linear programs,” *Naval Research Logistics Quarterly*, **8** (1): 101–111.
- Daouas, T., K. Ghedira and J. Muller (1995), “Distributed flow shop scheduling problem versus local optimization,” *Proceedings of the First International Conference on Multi-Agent Systems*, Cambridge, Massachusetts: MIT Press, 1995
- Davis, L. (1985), “Job shop scheduling with genetic algorithms,” *Proceedings of an International Conference on Genetic Algorithms and Their Applications*, Carnegie Mellon University, 136–140.
- Davis, W. and A. Jones (1988), “A real-time production scheduler for a stochastic manufacturing environment. *International Journal of Computer Integrated Manufacturing*,” **1** (2): 101–112.
- Dettmer, W. (1997), *Goldratt’s Theory of Constraints: A Systems Approach to Continuous Improvement*, Milwaukee, Wisconsin: Quality Press.

- Felix Schill, Uwe R. Zimmer and Jochen Trumpp (2005), "Towards Optimal TDMA Scheduling for Robotic Swarm Communication," distributed to *TAROS '05*
- Foo, Y. and Y. Takefuji (1988), "Stochastic neural networks for solving job-shop scheduling: Part 2 Architecture and simulations," *Proceedings of the IEEE International Conference on Neural Networks*, published by IEEE TAB: II283-II290.
- Fox, M. (1983), "Constraint-Directed Search: A case study of Job Shop Scheduling," Ph.D. Dissertation, Carnegie-Mellon University.
- Gershwin, S. (1989), "Hierarchical flow control: a framework for scheduling and planning discrete events in manufacturing systems," *Proceedings of IEEE Special Issue on Discrete Event Systems*, 77: 195-209.
- Glover, F. (1989), "Tabu search—Part I," *ORSA Journal on Computing*, 1 (3): 190-206.
- Glover, F. (1990), "Tabu search—Part II," *ORSA Journal on Computing*, 2 (1): 4-32.
- Glover, F. (1996), "Tabu search and adaptive memory programming—advances, applications and challenges," To appear in: *Interfaces in Computer Science and Operations Research*, The Netherlands: Kluwer Academic Publishers.
- Goldberg, D. and R. Lingle (1985), "Alleles, loci, and the traveling salesman problem," *Proceedings of the of the International Conference on Genetic Algorithms and Their Applications*, Carnegie Mellon University, 162-164.
- Goldberg, D. (1988), *Genetic Algorithms in Search Optimization and Machine Learning*, Menlo Park: California: Addison-Wesley.
- Goldratt, E. (1990), *Theory of Constraints*, Great Barrington, Massachusetts: North River Press.
- Goldratt, E. (1992), *The Goal*, Great Barrington, Massachusetts: North River Press.
- Grabot, B. and L. Geneste (1994), "Dispatching rules in scheduling: a fuzzy approach," *International Journal of Production Research*, 32 (4): 903-915.
- Hopfield, J. and D. Tank (1985), "Neural computation of decisions in optimization problems," *Biological Cybernetics*, 52: 141-152.
- J. Venkateswaran and Y.-J. Son (2005), "Hybrid System Dynamic - Discrete Event Simulation-Based Architecture for Hierarchical Production Planning," *International Journal of Production Planning Research*, 43 (20): 4397-4429
- Jayendran Venkateswaran, Young-Jun Son and Albert Jones (2004), "Hierarchical Production Planning Using a Hybrid System Dynamic-Discrete Event Simulation Architecture," *Proceedings of the 2004 Winter Simulation Conference*
- Jeffcoat, D. and R. Bulfin (1993), "Simulated annealing for resource-constrained scheduling," *European Journal of Operational Research*, 70: 43-51.
- Kalyan Veeramachaneni and Lisa Ann Osadciw (2004), "Optimal Scheduling in Sensor Networks Using Swarm Intelligence"
- Kirkpatrick, S., C. Gelatt and M. Vecchi (1983), "Optimization by simulated annealing," *Science*, 220 (4598) : 671-680.
- Krucky, J. (1994), "Fuzzy family setup assignment and machine balancing," *Hewlett-Packard Journal*, June: 51-64.
- Lawler, E. and D. Wood (1966), "Branch and bound methods: a survey," *Operations Research*, 14, 699-719.
- Le Pape, C. (1995), "Scheduling as intelligent control of decision-making and constraint propagation," *Intelligent Scheduling*, " edited by M. Zweben and M. Fox, San Francisco, California: Morgan Kaufman, 67-98.
- Lo, Z. and B. Bavarian (1991), "Scheduling with neural networks for flexible manufacturing systems," *Proceedings of the IEEE International Conference on Robotics and Automation*, Sacramento, California, 818-823.
- McKenzie, L. (1976), "Turnpike theory," *Econometrics*, 44: 841-864.
- Montazer, M. and L. Van Wassenhove (1990), "Analysis of scheduling rules for an FMS," *International Journal of Production Research*, 28: 785-802.
- Morton, E. and D. Pentico (1993), *Heuristic Scheduling Systems*, New York: John Wiley & Sons.
- Panwalker, S. and W. Iskander (1977), "A survey of scheduling rules," *Operations Research*, 25 (1): 45-61.
- Parunak, H., B. Irish, J. Kindrick and P. Lozo (1985), "Fractal actors for distributed manufacturing control," *Proceedings of the Second IEEE Conference on Artificial Intelligence Applications*, 653-660.
- Quinlan, J. (1986), "Induction of decision trees," *Machine Learning*, 1: 81-106.
- Rabelo, L. (1990), "Hybrid Artificial Neural Networks and Knowledge-Based Expert Systems Approach to Flexible Manufacturing System Scheduling," Ph.D. Dissertation, University of Missouri-Rolla.
- Rabelo, L., M. Sahinoglu and X. Avula (1994), "Flexible manufacturing systems scheduling using Q-Learning," *Proceedings of the World Congress on Neural Networks*, San Diego, California: 1378-1385.
- Richard J. Linn, Chin-Sheng Chen and Jorge A. Lozan (2002), "Development of Distributed Simulation Model for the Transporter Entity in a Supply Chain Process," *Proceedings of the 2002 Winter Simulation Conference*
- Rumelhart, D., J. McClelland and the PDP Research Group (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, 1: Foundations*, Cambridge, Massachusetts: MIT Press.
- Saleh, A. (1988), "Real-Time Control of a Flexible Manufacturing Cell," Ph.D. Dissertation, Lehigh University.
- Shapiro, J. (1979), "A survey of Lagrangian techniques for discrete optimization," *Annals of Discrete Mathematics*, 5: 113-138.
- Shaw, M., S. Park and N. Raman (1992), "Intelligent scheduling with machine learning capabilities: The induction of scheduling knowledge," *IEE Transactions on Design and Manufacturing*, 24: 156-168.
- Slany, W. (1994), "Scheduling as a fuzzy multiple criteria optimization problem." CD-Technical Report 94/62, Technical University of Vienna.
- Smith, S. (1995), "OPIS: A methodology and architecture for reactive scheduling," *Intelligent Scheduling*, edited by M. Zweben and M. Fox, San Francisco, California: Morgan Kaufman, 29-66.
- Srinivasan, V. (1971), "A hybrid algorithm for the one machine sequencing problem to minimize total tardiness," *Naval Research Logistics Quarterly*, 18: 317-327.
- Starkweather, T., D. Whitley, K. Mathias and S. McDaniel (1992), "Sequence scheduling with genetic algorithms," *Proceedings of the US/German Conference on New Directions for OR in Manufacturing*, 130-148.
- Starkweather, T., D. Whitley and B. Cookson (1993), "A Genetic Algorithm for scheduling with resource consumption," *Proceedings of the Joint German/US Conference on Operations Research in Production Planning and Control*, 567-583.
- Subrahmanyam, S., M. Zentner and J. Pekny (1996), "Making the most out of corporate information assets: the next generation

- of process scheduling, planning, and design tool," *Proceedings of the Process Industry Technical Conference: Looking Toward the 21st Century*, June 26–27, Erie, Pennsylvania.
- Sutton, R. (1988), "Learning to predict by the methods of temporal differences," *Machine Learning*, **3**: 9–44.
- Tesauro, G. (1992), "Practical issues in temporal difference learning," *Machine Learning*, **8**: 257–277.
- Tsujimura, Y., S. Park, S. Chang and M. Gen (1993), "An effective method for solving flow shop scheduling problems with fuzzy processing times," *Computers and Industrial Engineering*, **25**: 239–242.
- Vakharia A. and Y. Chang (1990), "A simulated annealing approach to scheduling a manufacturing cell," *Naval Research Logistics*, **37**: 559–577.
- Watkins, C. (1989), "Learning from Delayed Rewards," Ph.D. Dissertation, King's College, Cambridge.
- Werbos, P. (1995), "Neurocontrol and supervised learning: An overview and evaluation," *Handbook of Intelligent Control: Neural, Fuzzy, and Adaptive Approaches*, New York: Van Nostrand Reinhold Publication, 65–89.
- Wilkerson, L. and J. Irwin (1971), "An improved algorithm for scheduling independent tasks," *AIIE Transactions*, **3**: 239–245.
- Wu, D. (1987), "An Expert Systems Approach for the Control and Scheduling of Flexible Manufacturing Systems," Ph.D. Dissertation, Pennsylvania State University.
- Wysk, R., D. Wu and R. Yang (1986), "A multi-pass expert control system (MPECS) for flexible manufacturing systems," *NBS Special Publication*, **724**: 251–278.
- Yih, Y. (1990), "Trace-driven knowledge acquisition (TDKA) for rule-based real-time scheduling systems," *Journal of Intelligent Manufacturing*, **1** (4): 217–230.
- Zentner, M., J. Pekny, G. Reklaitis and N. Gupta (1994), "Practical considerations in using model-based optimization for the scheduling and planning of batch/semicontinuous processes," *Journal of Process Control*, **4** (4): 259–280.
- Zhang, M. and C. Zhang (1995), "The consensus of uncertainties in distributed expert systems," *Proceedings of the First International Conference on Multi-Agent Systems*, Cambridge, Massachusetts: MIT Press.
- Zhang, W. and T. Dietterich (1996), "High-performance job-shop scheduling with a time-delay TD(λ) network," *Advances in Neural Information Processing Systems*, Cambridge, Massachusetts: MIT Press. 1025–1030.
- Zhou, D., V. Cherkassky, T. Baldwin and D. Hong (1990), "Scaling neural networks for job shop scheduling," *Proceedings of the International Conference on Neural Networks*, **3**: 889–894.
- Zweben, M., B. Daun, E. Davis and M. Deale (1995), "Scheduling and rescheduling with iterative repair," *Intelligent Scheduling*, edited by M. Zweben and M. Fox, San Francisco, California: Morgan Kaufman, 241–256.

SDM Program Massachusetts
Institute of Technology,
Cambridge, MA 02139-0254
Department of Industrial
Engineering and
Management Systems
University of Central
Florida, Orlando, FL 32816

ALBERT JONES Ph.D.
LUIS C. RABELO Ph.D.
ABEER T. SHARAWI
National Institute of Standards
and Technology, Building
220, Room A127,
Gaithersburg, MD
20899-0001