Improving Manufacturing System Performance through Rescheduling

Jeffrey W. Herrmann
Computer Integrated Manufacturing Laboratory
University of Maryland
College Park, Maryland 20742
E-mail: jwh2@eng.umd.edu

June 21, 2001

Abstract

Many production facilities generate and update production schedules. Rescheduling is necessary because the manufacturing system is dynamic and unexpected events occur. This paper discusses rescheduling policies and how they affect the performance of manufacturing systems. The benefits of proper rescheduling include improved manufacturing system performance with reduced cost due to computational effort (human or computer) and disruptions to existing plans (nervousness). The paper summarizes previous work in this area and also discusses how rescheduling requires methods to solve production scheduling problems. The goal is to help production planners, engineers, and researchers understand the importance of rescheduling and the proper role of production scheduling in dynamic manufacturing systems.

1. Introduction

Manufacturing facilities are complex, dynamic, stochastic systems. From the beginning of organized manufacturing, workers, supervisors, engineers, and managers have developed many clever and practical methods for controlling production activities.

Although dispatching rules, kanban cards, and other decentralized production control policies are in use, many manufacturing facilities generate and update production schedules, which are plans that state when certain controllable activities (e.g., processing of jobs by resources) should take place. Dispatching rules are usually quick but myopic because they typically they do not use global information. Production schedules can enable better coordination to increase productivity and minimize operating costs. A production schedule can identify resource conflicts, control the release of jobs to the job shop, ensure that required raw materials are ordered in time, determine whether delivery promises can be met, and identify time periods available for preventive maintenance.

Production scheduling is a very difficult combinatorial optimization problem, so manual solution of any reasonably large problem is impossible. Research scientists, software companies, and manufacturing consultants have developed and implemented advanced scheduling systems that can perform production scheduling for complex manufacturing systems. These scheduling systems exploit results from production scheduling theory and advanced optimization techniques.

Moreover, because the manufacturing system is dynamic and unexpected events occur, rescheduling is necessary because the production schedule is a plan that must be updated when the state of the manufacturing system makes the current production schedule infeasible. There are many types of disturbances that can upset the plan, including machine failures, processing time delays, rush orders, quality problems, and unavailable material. As Bean *et al.* (1991) state, rescheduling is a dynamic approach that responds to disruptions, yet it considers future information (by creating a plan for the future).

Rescheduling is related to order release. Rescheduling can occur without order release when revising an existing schedule (to respond to a disruption). Order release makes an order available for processing and is thus the same as a job arrival. If order releases (job arrivals) are not predictable, processing cannot begin until the job is scheduled, so either rescheduling must be done when the job arrives or the job must wait until the next rescheduling point. If order releases are known ahead of time, then the arriving jobs can be scheduled along with the existing backlog.

Most rescheduling research has studied methods for updating a schedule or creating that a schedule that is robust with respect to disturbances. There have also been a few studies that provide guidance on when rescheduling should be done. Traditionally the rescheduling period is based on management measurement periods (one week or one day or one shift). But rescheduling also occurs during this period as unexpected events occur. Task start times are delayed, jobs are reassigned to different resources, and other adjustments take place. Intuitively, one can see that while small disruptions may (or should) be ignored (or handled in some simple way), larger disruptions may require significant changes to the production schedule to maintain good system performance. Huge disruptions will require many changes.

This paper discusses rescheduling policies and how they affect the performance of manufacturing systems. The paper summarizes previous work in this area and also discusses how rescheduling requires methods to solve production scheduling problems. The goal is to help production planners, engineers, and researchers understand the importance of rescheduling and the proper role of production scheduling in dynamic manufacturing systems.

This paper does not address the sequencing of parts processed in high-volume, repetitive manufacturing systems. In such settings, one can look to JIT and lean manufacturing principles for how to control production. These approaches generally do not need the same type of production schedules discussed here.

The remainder of this paper is organized as follows: Section 2 provides definitions for some terms to be used. Section 3 discusses related work on the control of manufacturing systems. Section 4 discusses how rescheduling affects manufacturing system performance. Section 5 discusses the need to solve production scheduling problems. Section 6 concludes the paper and identifies some promising research directions.

2. Definitions

To avoid the possibility of confusion, let us define some terms before continuing. A *manufacturing system* organizes equipment, people, and information to fabricate and assemble finished goods that are shipped to a customer. This system may be as large as a factory or as small as a manufacturing cell. According to Black (2000), a manufacturing system is “the collection of operations and processes used to produce a desired product.” A manufacturing system does not include finance, design engineering, research and development, production and inventory planning, purchasing, or distribution. Note that it does include order release, shop floor control, and material handling.

Order release controls a manufacturing system’s input by determining which orders (jobs) should be moved into production. It may be known as job release, order review/release, input/output control, or just input control.

Shop floor control determines which operation each person and piece of equipment should do and when they should do it. In general this process controls all production and material handling resources. Design decisions include order release policies (including WIP levels for pull systems), dispatching rules, batch sizes, and preventive maintenance policies.

A *production schedule* specifies, for each resource required for production, the planned start time and end time of each job assigned to that resource.

Scheduling is the process of creating a production schedule for a given set of jobs and resources.

Rescheduling is the process of updating an existing production schedule in response to disruptions or other changes. This includes the arrival of new jobs, machine failures, and machine repairs.

A *rescheduling policy* specifies how rescheduling is done. The policy specifies the events that trigger rescheduling. These events may be predictable (even regular) or unpredictable. The policy specifies the method used to revise the existing schedule. Note that the policy may specify different methods for different situations. If these policies have any parameters (for instance, the length of the rescheduling period), the policy specifies these parameters.

Dynamic scheduling schemes do not create production schedules. Instead, these decentralized production control methods dispatch jobs when necessary and use all information at the moment of dispatching. Such schemes use dispatching rules to prioritize jobs waiting for processing at a resource.

3. Literature Review

There are a large number of sources that describe approaches for shop floor control. This section will highlight those that are most relevant to the topic of rescheduling.

Production Planning. Vollmann *et al.* (1997) describe manufacturing planning and control systems, which include not only shop floor control methods but also materials requirements planning and aggregate production planning procedures. In traditional production planning approaches, aggregate production plans cover long periods of time (e.g., quarters and years) and are updated infrequently (monthly), master production schedules cover shorter periods (weeks and months) and are updated frequently (weekly), and production schedules cover the shorter periods (shifts and days) and are updated very frequently (daily).

Order Release. Herrmann (2000) gives an overview of order release procedures. Although, as mentioned above, order release and rescheduling are closely related (since order release controls the arrival of new orders that need to be included in the schedule), few researchers have described integrated order release and rescheduling decisions. Church and Uzsoy (1992) and Vieira *et al.* (2000a, b) do explicitly include order release as part of the rescheduling policy.

Production Scheduling. Pinedo (1995) covers mathematical formulations of many important machine scheduling problems and describes the design and implementation of scheduling systems. Most research on scheduling problems has focused on the complexity of the combinatorial optimization problems that occur and exact and approximate solution techniques. For example, Herrmann *et al.* (1993, 1995a, 1995b, 1997) study production scheduling problems for a variety of problem settings and present a range of solution approaches.

Rescheduling. A number of authors have proposed rescheduling approaches for a variety of scheduling environments. In general there are three primary types of studies: one, methods for repairing a schedule that has been disrupted; two, methods for creating a schedule that is robust with respect to disruptions; and three, studies of how rescheduling policies (which specify when rescheduling is done) affect the performance of the dynamic manufacturing system.

Bean *et al.* (1991) discuss a matchup scheduling procedure that repairs a production schedule when a disruption occurs. Their results show that matchup scheduling is an optimal approach when disruptions are infrequent enough to allow the system to get back on schedule before the next disruption. However, the approach does not address how an initial schedule should be created. Zweben et al. (1993) described rescheduling problems that occur in space shuttle ground processing, which are project scheduling problems. Bierwirth and Mattfeld (1999) present a genetic algorithm that reuses the previous solution to solve a job shop scheduling problem every time a new job arrives.

Leon *et al.* (1994) analyze how a single disruption delays a job shop schedule and present surrogate measures for estimating that delay in more general cases. They present a genetic algorithm to find robust schedules that minimize expected delay and expected makespan. Byeon *et al.* (1998) and Wu *et al.* (1999) present approaches to create a robust partial schedule for a job shop that is subject to disturbances. The incomplete portions of the schedule are resolved at the appropriate time, giving the shop some flexibility to handle disruptions. Their results show that, in a range of situations, such a schedule leads to better system performance than dispatching rules. However, as the amount of processing time variability increases, dispatching rules led to better performance. This does not address the dynamic aspects of the manufacturing system. Similarly, Mehta and Uzsoy (1998) present an approach to create predictive schedules that include inserted idle time as a means to reduce the impact of disruptions.

A number of papers present the results of simulation studies to show how the length of the rescheduling period influences the performance of scheduling heuristics for dynamic manufacturing systems. Farn and Muhlemann (1979) study a single-machine system with sequence-dependent setup times. Arriving jobs are included in the schedule at the next rescheduling point, and the schedule is created using a heuristic. Muhlemann *et al.* (1982) study the dynamic job shop scheduling problem and experimentally compare different scheduling heuristics across a range of scenarios, including rescheduling period length, the number of jobs in the backlog, and the amount of uncertainty in processing times and machine failures.

Church and Uzsoy (1992) developed a hybrid event-driven rescheduling policy for single- and parallel-machine models with dynamic job arrivals. Their system reschedules the facility, periodically taking into account work that is already in the system. Regular events occurring between routine rescheduling are ignored until the next rescheduling moment. However, when an event is classified as an exception, immediate action should be taken, with the entire facility being rescheduled and resulting schedule implemented until the next schedule generation point. To create a schedule, the system uses the Earliest Due Date rule to minimize maximum lateness. The paper also presents analytical models to bound the maximum completion time. Their models do not consider other system performance measures. This work also did not consider different types of jobs or setups.

Vieira *et al.* (2000a, b) have also studied rescheduling policies. Vieira *et al.* (2000a) have studied a single-machine system and developed analytical models to estimate system performance. That work considered two rescheduling policies: periodic and event driven based on queue size. Their results show that the analytical models can accurately predict the performance of a single-machine system operating under those rescheduling strategies. Vieira *et al.* (2000b) extended that study by investigating parallel machine systems, which have more complex rescheduling strategies.

Currently, John Fowler and collaborators at Arizona State University are investigating production scheduling problems for unrelated parallel machines that have processing time variability and equipment breakdowns and rescheduling strategies for semiconductor wafer fabrication facilities.

Dynamic Scheduling. Dispatching rules and pull mechanisms are used to control production without a production schedule. When a machine becomes available it chooses from among the jobs in its queue by using a dispatching rule that sorts the jobs by some criteria. Common dispatching rules employ processing times and due dates in simple rules and complex combinations. Some dispatching rules are extensions of policies that work well on simple machine scheduling problems (e.g. Shortest Processing Time and Earliest Due Date).

Panwalkar and Iskander (1977) provide an extensive list of dispatching rules. They categorize these rules into five classes: simple dispatching rules, combination of simple rules, weighted priority indexes, heuristic scheduling rules, and other rules.

Green and Appel (1981) examine the problem of job shop scheduling by asking the following questions: What traditional dispatching rules do experienced schedulers select? Would dispatch rule selection be influenced by urgency? Would schedulers select a dispatch order based on organizational influence or peer pressure? The authors asked schedulers in a number of plants to denote which of the following rules they used: due date, slack, operation due date, slack per operation, SPT, FCFS, COVERT, Program in Greatest Trouble (PGT), or friend needs a favor (FNF). The authors report that influence systems affect scheduling. The PGT rule (a coalition rule) was highly valued, but FNF (an individual rule) was rejected. Traditional and theoretical rules were not highly valued.

The computational effort of dispatching rules is low when simple rules (like SPT or EDD) are used. However, some dispatching rules require a large amount of information, and the job priorities must be recalculated at every dispatching decision.

Pull mechanisms such as kanban cards and CONWIP order release policies add production authorization cards to the system so that a resource can work only when both material and cards are available. Hopp and Spearman (1996) provide a good introduction to these topics. Buzacott and Shanthikumar (1993) analyze a generalized production authorization policy.

Control Theoretic Approaches. Gershwin (1994) reviews literature of control theoretic models of manufacturing systems. The models are used to develop rules for deciding what action to take and when to take it in response to random disruptions. For instance, these result in control policies implemented as dispatching rules or hedging-point policies.

In a number of papers, Kumar and others (e.g., Kumar, 1994; Perkins and Kumar, 1989; Chase and Ramadge, 1992) have studied the control of dynamic manufacturing systems. Specifically, they have described classes of dispatching rules that identify which waiting job a resource should process next. For machines without setup times, the proposed dispatching rules are a class of least slack policies that prioritize each job by the difference between its due date (or some surrogate) and the expected amount of time until the job is completed. For resources with setup times, the proposed dispatching rules focus on completing all waiting jobs of one type before performing a setup and processing jobs of another type. All of the rules studied keep a machine working if there are any jobs waiting for processing. (That is, the machine cannot ignore waiting jobs.) Kumar (1994) summarizes the results of work on the stability and performance of these policies. This important work demonstrates why certain classes of dispatching rules work well and provides guidance when selecting dispatching rules. However, there exist dynamic manufacturing systems for which these types of dispatching rules are inappropriate or suboptimal. For example, Chase and Ramadge (1989) demonstrated that there exist idling policies that have superior performance.

For a single machine operating in a dynamic, stochastic environment, Markowitz and Wein (2001) classify scheduling problems based on three attributes: the presence of setups, the presence of due dates, and the type of products (standardized or customized). They present dynamic cyclic policies that minimize the long-run expected average costs of earliness, tardiness, holding, and setups.

4. The Impact of Rescheduling Policies

Determining the impact of a rescheduling policy on a dynamic manufacturing system requires careful study, modeling, and analysis of the specific manufacturing system. Unfortunately, at this time there exist few models for understanding the impact of rescheduling policies. The following discussion summarizes the limited results that do exist.

Intuitively, it seems natural that rescheduling more often yields better performance. A number of experimental studies support this hypothesis. See, for example, Farn and Muhlemann (1979) and Muhlemann *et al.* (1982). The latter paper also suggests that the rescheduling period affects system performance more when there is greater uncertainty and that managers need to explore the tradeoff between the cost of scheduling and the benefits of more frequent scheduling.

The cost of rescheduling includes computational effort (human or computer) and disruptions to existing plans (nervousness). The rescheduling period affects the number of jobs being considered for scheduling. A longer rescheduling period means that more jobs (and tasks) will be considered in the scheduling problem. This will increase the computational effort needed to create the production schedule. Moving jobs from one scheduled machine to another may require additional material handling work. For instance, Bean *et al.* (1991) use the number of jobs reassigned as a measure of solution quality.

Bean *et al.* (1991) show that the matchup algorithm (which requires more job reassignments) leads to better performance (less total tardiness) than a simple pushback strategy that simply delays tasks.

According to Wu *et al.* (1999) a robust, partial schedule leads to better system performance (less weighted tardiness) than dispatching rules. However, as processing time variability increases, dispatching rules lead to better performance. Leon *et al.* (1994) state that, as processing time variability increases, the improvement (in expected makespan and expected delay) due to robust schedules increases.

Mehta and Uzsoy (1998) state that predictive schedules (with inserted idle time) increase predictability (reduce nervousness) but do not significantly degrade system performance (maximum lateness), compared to schedules generated by ignoring possible breakdowns.

Church and Uzsoy (1992) state that periodic rescheduling procedures leads to near optimal performance (minimal maximum lateness) when order release is periodic. In addition, rescheduling at the arrival of a “rush” job (one with a tight due date) is useful, but more frequent rescheduling does not improve system performance significantly. Thus, if done carefully, good system performance can be maintained while reducing the rescheduling effort (the number of rescheduling events).

Vieira *et al.* (2000a, b) have shown that rescheduling frequency can significantly affect the system performance (average flow time). A lower rescheduling frequency (which causes longer rescheduling periods) lowers the number of setups (and reduce time wasted on setups) but increases manufacturing cycle time and WIP. Event-driven and periodic strategies exhibit similar performance. Rescheduling when a machine fails or becomes available after a repair decreases manufacturing cycle time slightly but increases the frequency of rescheduling.

Although it can increase computational effort, a longer rescheduling period (which increases the number of jobs that are considered simultaneously) can improve system performance through better coordination. For example, Herrmann and Delalio (2001) consider the affect of rescheduling period on decisions regarding batching and scheduling of sheet metal punch press operations. Their results indicate that, when material is inexpensive, decreasing the scheduling frequency can significantly reduce costs because fewer setups occur and more parts are produced from inexpensive unsheared sheets. However, when material is expensive, changing the scheduling frequency does not affect costs as much.

5. The Need to Solve Scheduling Problems

Rescheduling provides a perspective that can put into proper context the need to solve production scheduling problems in dynamic, stochastic manufacturing systems. Rescheduling policies identify not only when rescheduling should be done but also the objectives and constraints of the resulting scheduling problem. For example, Bean *et al.* (1991) present the matchup scheduling problem, which attempts to recover the original schedule as soon as possible while satisfying a constraint on allowable tardiness cost. Vieira *et al.* (2000a, b) study rescheduling policies that require the production schedule to minimize the number of setups and the job flow time.

Portougal and Robb (2000) discuss the gap between production scheduling theory and practice and emphasize the importance of the planning period. Their paper argues that, if job cycle times are greater than the planning period, then careful scheduling is needed to coordinate activities in multiple planning periods, and complex models are appropriate. If the cycle time is smaller, then scheduling is seldom important. The paper states that, in the latter case, the only important objective is that the resource (or production unit) completes all of the desired work in the planning period.

Despite that claim, however, one can easily see that scheduling is indeed important if careless scheduling would prevent the resource (or production unit) from accomplishing this goal. In the presence of sequence-dependent setup times, for instance, scheduling significantly affects the total time required. A poor schedule would waste valuable time doing setups.

Thus, it may be more appropriate to state that, when job cycle times are shorter than the planning period, satisfying the production target should set the constraints and objectives of the production scheduling problem. The resulting production scheduling problems may emphasize finding feasible solutions over optimization, but such problems can be extremely difficult in realistic settings.

6. Summary and Conclusions

This paper has described the role of rescheduling in dynamic manufacturing systems and discussed previous work related to rescheduling. In general there are three primary types of studies: one, methods for repairing a schedule that has been disrupted; two, methods for creating a schedule that is robust with respect to disruptions; and three, studies of how rescheduling policies affect the performance of the dynamic manufacturing system.

The paper discussed previous results on how rescheduling policies affect manufacturing system performance. Sophisticated repair algorithms can yield better performance than simple delay methods. Robust schedules can increase predictability and improve system performance. The impact of rescheduling frequency is mixed. In some cases, more frequent rescheduling yields better performance because the system reacts more quickly to unplanned events. However, more frequent rescheduling requires more effort, and too-often rescheduling can yield poor performance by requiring too many setups.

Studying rescheduling helps bridge the gap between theory and practice of production scheduling. Most scheduling results do not consider important characteristics of the dynamic environment in which scheduling occurs, which limits their usefulness. Rescheduling provides a perspective that can put into proper context the need to solve production scheduling problems in dynamic, stochastic manufacturing systems.

Modeling rescheduling. Mathematical models of dynamic, stochastic manufacturing systems can provide useful information to analysts and managers trying to design manufacturing systems. There are a wide variety of models available, including queueing network models and discrete event simulation models. Typically, however, these types of models do not explicitly represent the production control policies (e.g., rescheduling policies) that will control the system. Consequently, because these policies significantly affect system performance, the resulting system models will be inaccurate, which can lead to poor design decisions.

Because the rescheduling policy affects the performance of the manufacturing system, it needs to be considered in manufacturing system design. Rarely is the dynamic behavior of the manufacturing system considered during the design phase. When it is, more effort is spent modeling the resources in the factory and the flow of parts through the system. Little effort is spent modeling the production control scheme. This occurs because existing analytical and simulation models provide little support for rescheduling. Often, they are limited to predefined sets of dispatching rules. Although modern software for building discrete event simulation models allows an analyst to create complex models and sophisticated production control policies, building such models and conducting the necessary experiments can require a large amount of time and effort.

Research directions. Rescheduling research needs to proceed from in-depth studies of manufacturers with rescheduling problems. This will allow investigators to identify important issues, validate assumptions, and illustrate insights.

More research is needed to compare the performance of manufacturing systems under rescheduling policies to their performance under dynamic scheduling (such as dispatching rules). This will yield additional insight into the advantages and disadvantages of rescheduling in different problem settings. This study could be done by examining analytical models (for those systems where such models exist or can be constructed) or conducting simulation studies (for more complex systems). Although there have been some studies, a comprehensive campaign is still needed.

Rescheduling provides a systems view of manufacturing that includes not only material flow and resource availability but also order release and control systems. This perspective will be useful for a wide variety of manufacturing system control problems and should lead to developments in the design and optimization of manufacturing planning and control systems. It will bring closer the vision of a comprehensive set of models that can describe the complete set of dynamics within a manufacturing system.

Acknowledgements

The author wishes to thank the students, colleagues, and industrial collaborators who have contributed many valuable insights about the nature of production scheduling.

References

Bean, James C., John R. Birge, John Mittenthal, and Charles E. Noon, "Matchup scheduling with multiple resources, release dates, and disruptions," *Operations Research*, Volume 39, Number 3, pages 470-483, 1991.

Bierwirth, Christian, and Dirk C. Mattfeld, "Production scheduling and rescheduling with genetic algorithms," *Evolutionary Computation*, Volume 7, Number 1, pages 1-17, 1999.

Black, J.T., "Manufacturing systems," in *Encyclopedia of Production and Manufacturing Management*, Paul M. Swamidass, editor, Kluwer, Norwell, Massachusetts, 2000.

Buzacott, John A., and J.George Shanthikumar, *Stochastic Models of Manufacturing Systems*, Prentice Hall, Englewood Cliffs, New Jersey, 1993.

Byeon, Eui-Seok, S. David Wu, and Robert H. Storer, "Decomposition heuristics for robust job shop scheduling," *IEEE Transactions on Robotics and Automation*, Volume 14, Number 2, pages 303-313, 1998.

Chase, C., and P.J. Ramadge, "On the real-time scheduling policies for flexible manufacturing systems," *IEEE Transactions on Automatic Control*, Volume 37, Number 4, pages 491-496, 1992.

Church, L.K., and R. Uzsoy, "Analysis of periodic and event-driven rescheduling policies in dynamic shops," *International Journal of Computer Integrated Manufacturing*, Volume 5, pages 153-163, 1992.

Farn, C.K., and A.P. Muhlemann, "The dynamic aspects of a production scheduling problem," *International Journal of Production Research*, Volume 17, pages 15-21, 1979.

Gershwin, Stanley B., *Manufacturing Systems Engineering*, PTR Prentice Hall, Englewood Cliffs, New Jersey, 1994.

Green, G.I., and L.B. Appel, "An empirical analysis of job shop dispatch rule selection," *Journal of Operations Management*, Volume 1, pages 197-203, 1981.

Herrmann, Jeffrey W., "Order release," in *Encyclopedia of Production and Manufacturing Management*, Paul M. Swamidass, editor, Kluwer, Norwell, Massachusetts, 2000.

Herrmann, Jeffrey W., and David R. Delalio, "Algorithms for Sheet Metal Nesting," *IEEE Transactions on Robotics and Automation*, Vol. 17, No. 2, pp. 183-190, 2001.

Herrmann, Jeffrey W., and Chung-Yee Lee, "On scheduling to minimize earliness-tardiness and batch delivery costs with a common due date," *European Journal of Operational Research*, Vol. 70, No. 3, pp. 272-288, 1993.

Herrmann, Jeffrey W., and Chung-Yee Lee, "Solving a class scheduling problem with a genetic algorithm," *ORSA Journal of Computing*, Vol. 7, No. 4, pp. 443-452, 1995a.

Herrmann, Jeffrey W., Chung-Yee Lee, and Jim Hinchman, "Global job shop scheduling with a genetic algorithm," *Production and Operations Management*, Vol. 4, No. 1, pp. 30-45, 1995b.

Herrmann, Jeffrey W., Jean-Marie Proth, and N. Sauer, "Heuristics for unrelated machine scheduling with precedence constraints," *European Journal of Operational Research*, Vol. 102, No. 3, pp. 528-537, 1997.

Hopp, Wallace J., and Mark L. Spearman, *Factory Physics*, Irwin/McGraw-Hill, Boston, 1996.

Kumar, P.R., "Scheduling manufacturing systems of re-entrant lines," in *Stochastic Modeling and Analysis of Manufacturing Systems*, David D. Yao, ed., pages 325-360, Springer-Verlag, New York, 1994.

Leon, V. Jorge, S. David Wu, and Robert H. Storer, "Robustness measures and robust scheduling for job shops," *IIE Transactions*, Volume 26, Number 5, pages 32-43, 1994.

Markowitz, David M., and Lawrence M. Wein, "Heavy traffic analysis of dynamic cyclic policies: a unified treatment of the single machine scheduling problem," *Operations Research*, Volume 49, Number 2, pages 246-270, 2001.

Mehta, S.V., and R.M. Uzsoy, "Predictable scheduling of a job shop subject to breakdowns," *IEEE Transactions on Robotics and Automation*, Volume 14, pages 365-378, 1998.

Muhlemann, A.P., A.G. Lockett, and C.K. Farn, "Job shop scheduling heuristics and frequencies of scheduling," *International Journal of Production Research*, Volume 20, pages 227-241, 1982.

Panwalkar, S.S., and W. Iskander, "A survey of scheduling rules," *Operations Research*, Volume 25, pages 45-61, 1977.

Perkins, J.R., and P.R. Kumar, "Stable, distributed real-time scheduling of flexible manufacturing/assembly/disassembly systems," *IEEE Transactions on Automatic Control*, Volume 34, Number 2, pages 139-148, 1989.

Pinedo, Michael, *Scheduling: Theory, Algorithms, and Systems*, Prentice Hall, Englewood Cliffs, New Jersey, 1995.

Portougal, Victor, and Robb, David J., "Production scheduling theory: just where is it applicable?" *INTERFACES*, Volume 30, Number 6, pages 64-76, 2000.

Vieira, Guilherme E., Jeffrey W. Herrmann, and Edward Lin, "Analytical models to predict the performance of a single-machine system under periodic and event-driven rescheduling strategies," *International Journal of Production Research*, Volume 38, Number 8, pages 1899-1915, 2000a.

Vieira, Guilherme E., Jeffrey W. Herrmann, and Edward Lin, "Predicting the performance of rescheduling strategies for parallel machine systems," *Journal of Manufacturing Systems*, Volume 19, Number 4, pages 256-266, 2000b.

Vollmann, Thomas E., William L. Berry, and D. Clay Whybark, *Manufacturing Planning and Control Systems*, fourth edition, Irwin/McGraw-Hill, New York, 1997.

Wu, S. David, Byeon, Eui-Seok, and Storer, Robert H., "A graph-theoretic decomposition of the job shop scheduling problem to achieve scheduling robustness," *Operations Research*, Volume 47, Number 1, pages 113-124, 1999.

Zweben, M., E. Davis, B. Daun, and M.J. Deale, "Scheduling and rescheduling with iterative repair," *IEEE Transactions on Systems, Man, and Cybernetics*, Volume 23, pages 1588-1596, 1993.