Proceedings of the 2002 Winter Simulation Conference
E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.

USING SIMULATION TO EVALUATE BUFFER ADJUSTMENT METHODS IN ORDER PROMISING

Hank Grant

Industrial Engineering
University of Oklahoma
100 E. Boyd Suite R-208
Norman, OK 73019-1016 U.S.A.

Scott Moses

Industrial Engineering
University of Oklahoma
202 West Boyd, Room 124
Norman, OK 73019-0631 U.S.A.

Dave Goldsman

Industrial and Systems Engineering
Georgia Institute of Technology
765 Ferst Drive
Atlanta, GA 30332-0205 U.S.A.

ABSTRACT

Much literature exists for scheduling production, but there is little work on establishing the due dates that serve as the inputs to developing a production schedule. We call this Order Promising. This paper explores a simulation-based approach for evaluating methods for promising the delivery of orders based on dynamic buffer adjustment coupled with various methods to forecast the amount of buffer required. The primary objective of the paper is to frame the problem and suggest methods of analysis. Preliminary computational results are also presented.

1 INTRODUCTION

Order promising is similar to, but different from, due date assignment. It is designed to consider the current status of the system and fold that information into an estimate of when an order could be promised to a customer. A scheduled due date is first calculated using traditional flow-time estimation methods. Then, that scheduled date is modified by a "buffer" amount. The buffer is updated periodically to reflect current system conditions. The overall objective is to have the lateness to the promised date be close to zero. As we will see, we are also concerned with the variation of the lateness-to-promise estimator. Promise dates with lateness expectation of zero but with large variance will not be acceptable to customers.

The basic concept is that, if we can do a good job on buffer adjustment, we can more effectively use the current behavior of the system to actually establish promise dates. While we will still need good scheduling and control to ensure shop performance to plan, better promising will help ensure that is achieved. In theory, this should yield better performance to promise and more accurately reflect actual operations.

2 BACKGROUND

A vast literature exists on scheduling to meet pre-specified due dates. Thus, some may infer that the problem of assigning due dates (promising) has been thoroughly investigated. However, very little effective research has been done on due date assignment itself. Various survey papers on due date assignment include (Cheng 1989; Smith 1983; Ragatz 1984a).

Furthermore, this paucity of research on due date assignment is evident in both the production systems and real-time database (RTDB) systems fields. Both fields aim to schedule tasks such that they complete by their deadlines, yet they have given very little attention to procedures for actually setting deadlines! In an RTDB, transactions must not only maintain consistency constraints of the database but also must satisfy deadlines. The main goal of an RTDB is to meet the deadlines of data transactions regardless of system or transaction failures (Gruenwald 1997; Gruenwald 1999). Research on how to schedule real-time transactions so that they can meet their deadlines (Chen 1996a) assumes that a deadline is given as a property of a transaction at the time when it is submitted to the system. Among the popular scheduling techniques are earliest deadline and earliest slack time (Abbott 1992), which are analogous to simple priority rules for task sequencing.

Computational limitations of the 1960's motivated Conway (1965) and other researchers to consider highly simplified parametric approaches such as CON (Conway 1965), TWK (Conway 1965), JIQ (Eilon 1976), and WIQ (Ragatz 1984b) (also (Vig 1993; Raghu 1995; Tsai 1997; Udo 1993)). Several studies have confirmed the basic intuition that consideration of current system congestion improves the accuracy of order lead time estimation, including (Eilon 1976; Heard 1976; Weeks 1979; Bertrand 1983; Ragatz 1984b; Cheng 1988; Rajasekera 1991; Wein 1992). Parametric policies require tuning, so a few procedures have been developed to automatically adjust parameters in a closed loop fashion (Baker 1981; Seidmann 1981; Cheng 1986; Cheng 1987; Shanthikumar 1988), although these procedures all require assumptions about arrival processes or processing time distributions.

Several analytical works consider the combined problem of due date setting and job sequencing using either queueing theory (Shanthikumar 1988; Rajasekera 1991; Wein 1991; Wein 1992; Duenyas 1995a; Duenyas 1995b; Palaka 1998; Weng 1999) or deterministic optimization (Heard 1976; Seidmann 1981; Cheng 1986; Cheng 1987; Yano 1987; De 1992; Luss 1993; Bagchi 1994; De 1994a; Li 1999). The latter sometimes employ simple parametric policies for due date assignment with optimal parameter values (Cheng 1991; De 1994b; Chen 1996b; Cheng 1996). Almost all of these results are for single-machine systems.

For multi-resource systems, researchers recently have begun to perform statistical estimation of the distribution of flow-time to set due dates for future arrivals (Lawrence 1995; Vig 1991; Kaplan 1993). Hopp (2000) combines factory physics, statistical estimation, and control charting to create a lead-time estimator that is generally applicable and also adaptive to changes in the system over time. Lawrence (1995) assigns due dates by forecasting order flow-times (using one of six different estimators) and then adding some function of the forecast error distribution to this estimate so as to achieve a certain performance objective (the function is selected based on the particular performance criteria).

3 PROMISING SYSTEM ARCHITECTURE

The Promising System Architecture used in this research is shown in Figure 1. Orders arriving are submitted when they arrive to a Promising Service. This service develops a scheduled completion date for the order and adjusts the scheduled due date using the buffer to calculate a promise date. The buffer reflects the current performance of the system, regarding progress to promised dates.

Once a promise date has been established for an order, the Promising Service passes the order on to the Planning Service. The Planning Service is responsible for phasing the order into the existing orders and releasing those orders to the factory floor.

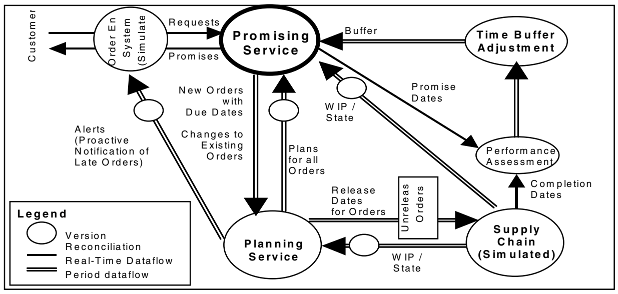

Figure 1: Promising System Architecture

Information is fed back to the Promising Service and the Planning Service from the Supply Chain regarding the status of orders, including notification of any late orders.

4 SIMULATION MODEL STRUCTURE

Figure 2 provides the structure for the simulation model. Independent sub-models were constructed for each of the components described and were based on the Promising System Architecture. The flow is shown in the figure and is described below.

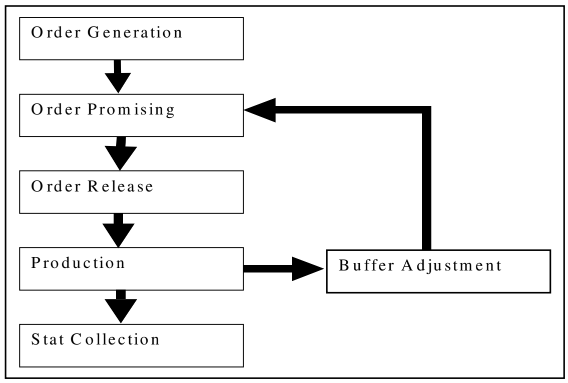

Figure 2: Model Components and Structure

4.1 Order Generation

Order Generation is simply the creation of orders to be processed. This module selects the type of part to be produced and controls the frequency of the arrivals of orders. The parameters of Order Generation available for modification are: arrival rates and distributions, part types, and part routings.

4.2 Order Promising

Order Promising calculates the promise date for each order using the following formula:

$$DATE_{Promise} = DATE_{Schedule} + BUFFER_t \tag{1}$$

where $BUFFER_t$ is the current value of the buffer at time t. The buffer is recalculated periodically in the buffer adjustment module and reflects the current performance of the system. Several methods have been explored to update the buffer value and are discussed later in the paper. The parameters of Order Promising include: calculation method for schedule and calculation method for buffer incorporation.

4.3 Order Release

Order Release controls the release of orders to the production system. They are accumulated prior to release, and then the complete set of active released orders is rescheduled based on the priority scheme selected. The parameters of Order Release include: time between releases, time to reschedule, and time to distribute to the factory floor.

4.4 Production

Production represents the actual production of orders. In the system under study here, three machining stations are available for production. The parts move through the stations in various orders, depending upon the part. Each part goes through one of the three stations once, and may have different processing times at each station. Various probability distributions are incorporated for the processing times. The layout of the production system is shown in Figure 3. The parameters of the Production module include: processing time and distribution for each step, resources available, and product routing.

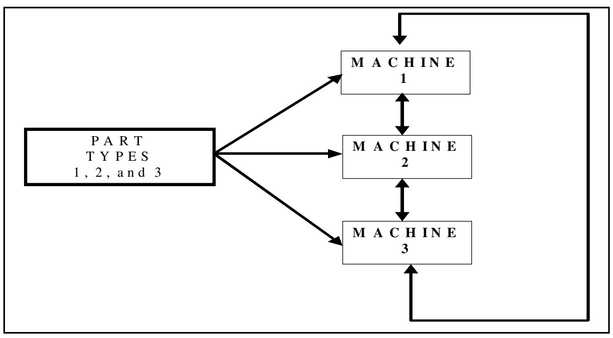

Figure 3: Production System Components

4.5 Statistical Collection

Statistical Collection performs the collection of information regarding various system performance parameters. These include:

- Lateness to Schedule;
- Lateness to Promise;
- k, the lag, in number of orders, between the arrival and completion of an order;
- Lateness as a function of k; and
- Time in the system

4.6 Buffer Adjustment

Buffer Adjustment is performed periodically and performs the recalculation of the buffer using various algorithms discussed below.

4.7 Software

The AWESIM simulation system (Pritsker 1999) was used to perform the experiments in conjunction with SAS and EXCEL.

5 METHODOLOGY

In this research, we have experimented with various methods of establishing promise dates using simple methods for establishing the scheduled date of completion and more complex methods for buffer adjustment. The relationship was shown in equation (1). The idea is that the value used for $BUFFER_t$ will characterize the variation in the system such that the resulting $DATE_{Promise}$ will be close to the actual completion date.

5.1 Buffer Adjustment

The buffer adjustment process consists primarily of characterizing the current system performance and feeding back that information in the form of a forecast of the lateness to scheduled completion to be used in the calculation of a promise date. First, we calculate an estimate of the scheduled completion time for parts. In this study, we have simply used the sum of each task's expected processing times. More complex methods could also be used.

We then observe the system for a period of time, and collect information concerning the lateness to schedule. We then use that data to estimate the future lateness to schedule for orders and include this estimate in establishing the promise date to provide to customers.

The process is complicated as follows. In Figure 4, we illustrate two processes, the arrivals of orders and the completion of orders. Orders arrive according to some probability distribution. The i^{th} arrival arrives at time t.

In the completion process, the j^{th} completion happens at some time before t, and the $(j+1)^{st}$ completion happens after time t. The i^{th} arrival becomes the $(j+k)^{th}$ completion, where k is a random variable. We would, of course, like to estimate the time of the $(j+k)^{th}$ completion as accurately as possible in order to establish a promise date to provide to customers.

We can define k as the number of orders completing between the arrival of an order and its completion, including that order.

Forecasting tools like exponential smoothing would use the lateness of the j^{th} completion and use it to forecast the lateness of the $(j+1)^{st}$ completion. However, in this application we must forecast the lateness of the $(j+k)^{th}$ completion where k is a random variable. We can estimate the distribution of k, as well as the lateness using simulation, as discussed below.

The lateness forecast is assigned to the variable BUFFER in equation (1), using one of the methods described below.

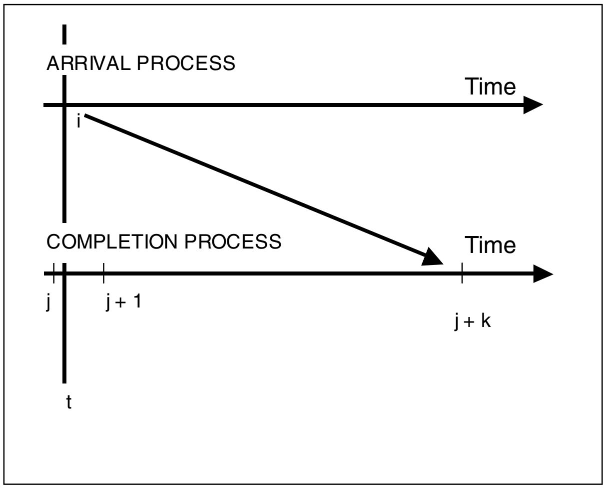

Figure 4: Lateness to Promise

5.2 Methods of Forecasting Lateness to Scheduled Completion

The following forecasting methods were explored in the estimation of the buffer:

Accumulated average lateness to schedule: Figure 5 shows the average lateness to schedule as time progresses. The cumulative average approaches a constant value. This lateness to schedule is the value used for the buffer value.

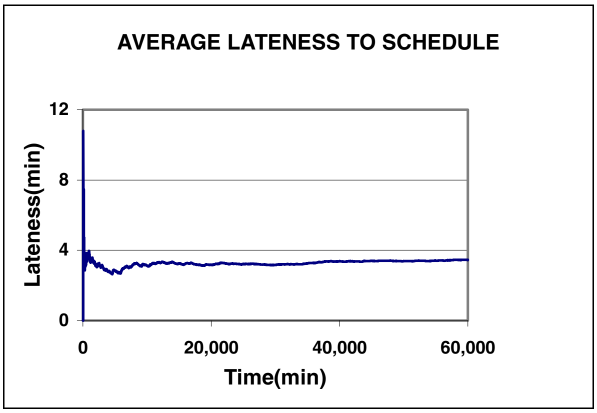

Figure 5: Average Lateness to Schedule

Exponential Smoothing: The exponential smoothing for lateness to schedule is based on periodically calculating the smoothing coefficients and then applying them to the forecast of the lateness to schedule, or the buffer.

ARIMA Time Series: The ARIMA models are updated in much the same way. Periodically, a subset of the data is fit to an ARIMA time series model and this model is used to forecast the future value of the buffer.

Regression: In using regression, several system parameters that might be correlated with lateness to schedule are observed. Periodically, a regression model is fit to these parameters, and the resulting model is used to forecast the lateness to schedule.

Time-shift due to k:

In most forecasting techniques, we are attempting to forecast the next observation of some variable. In this case, as was discussed above, we are trying to forecast k product completion events in the future, where k is a random variable that can be estimated by observation, but not (thus far) solved analytically. The incorporation of the time shift k into the methodology is still being examined and is discussed below.

6 OBSERVATIONS

6.1 Effects of the Variance of the Production Process

Extremely complex behavior can be exhibited by even relatively simple systems. In the system studied, we examined three machining systems, with unlimited queues. Three part types flow through all three machines only once, but in any order, and with different service times for each step. We examined constant, exponentially distributed, and normally distributed machining times. We also examined several levels of system congestion, from utilizations around 25% up to utilizations approaching 100%.

Figure 6 shows a plot of the lateness to promise for 60,000 minutes of production. As can be seen from this plot, the lateness to schedule is moderately variable throughout the process even thought the average is near zero. The system was configured to have approximately 25% utilization, and the process times were assumed to be constant.

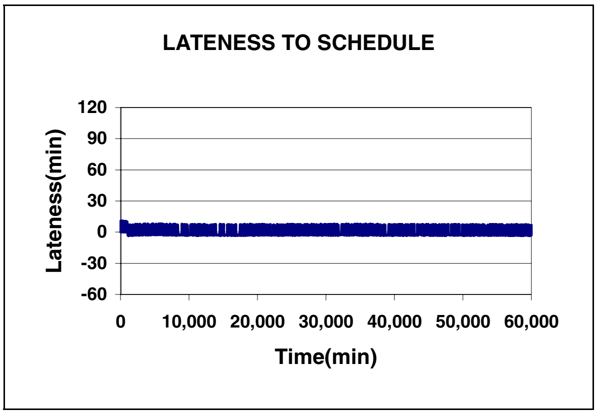

Figure 6: Lateness to Schedule - Constant Processing Times

In this case, the time units are in minutes. The variability in lateness to schedule is quite acceptable, even without buffer adjustment. However, when the system congestion increases, this is no longer the case.

We can quickly see the effects of increased process time variation. By simply replacing the constant times in the process with exponentially distributed variables, there is a dramatic increase in the variation of the lateness.

Figure 7 illustrates this effect where the only change in the model is in the processing time distributions.

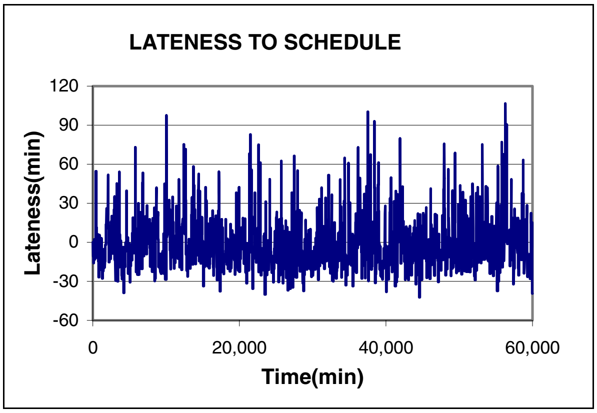

Figure 7: Lateness to Schedule - Exponential Processing Times

One of the challenges in this methodology is concerned with the underlying variance of the production process. If the underlying variance is large, this tends to be reflected in the buffer and thus, in the promise date estimation. While it is not too difficult to adjust the buffer such that the average lateness to promise is close to zero, this is not sufficient. If the resulting time series of lateness to promise observations are highly variable, customers will not be happy.

6.2 Distribution of k

As was discussed above, k is the lag between the current completion and the completion of the current arrival. It is a random variable and can be estimated by experimentation.

For example, simulating our test system with low utilization (25%) and exponential service times, Figure 8 shows the distribution of k. The average is 1.8, meaning that a current arrival completes approximately 2 orders later than the most recent completion before the current arrival, with the range from 1 to 6.

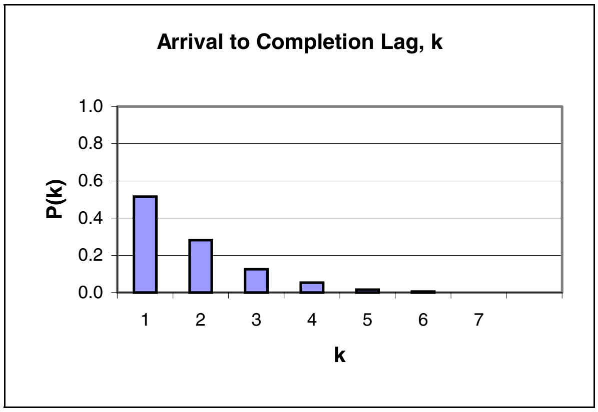

Figure 8: Distribution of k

We were also able to collect observations regarding the lateness to promise for various values of k. For instance, when k was 1, meaning that the order arriving was the next to complete, the lateness was observed to be -14.290, which means that the promise date was over-optimistic for this example. These are shown in Table 1.

The fact that these values exhibit the behavior that they do is no surprise. More orders completing between the arrival of an order and its completion provides more time for interference and delay. But it is interesting to be able to characterize it and perhaps use it in the development of a better characterization of the buffer.

Table 1: Lateness to Promise as a function of k

k	Average Lateness to Promise
1	-14.3
2	-1.8
3	15.7
4	24.1
5	46.2
6	56.8

7 CONCLUSIONS

Several conclusions can be made based on the work to date.

There are various, perhaps conflicting, performance measures that one would be interested in concerning this methodology. Those include:

- Average lateness to promise. We would like this to be zero.
- Variance of the lateness to promise. We would like this to be small, or within defined limits.
- Probability that the lateness to promise is less than or equal to some specified value. We would like to be able to specify this performance in terms

such as “95% of the orders completed with lateness to promise less than one day.”

The variance of the lateness to schedule and lateness to promise dates is a critical performance measure in the effective application of this methodology. For it to be useful, the variance of the estimates must be reduced or better characterized.

On another note, the system studied was not only extremely sensitive to the distribution of the process times but also the congestion of the system. It also exhibited cyclic renewal behavior at extremely long time intervals, on the order of 16,000 minutes. This long interval would not normally be investigated in analysis studies but would be significant in actual practice.

8 FUTURE WORK

We have identified several activities to pursue next in this project.

First, the incorporation of k lags in product completion into the model is necessary to better represent the lag between the forecast of the completion that will occur k completions in the future. Forecasting techniques will be used which will consider this lag.

Further, more information is available for buffer calculation than is currently being used. While the regression models explored included quantities such as queue size and waiting time, this should be expanded as a part of the buffer adjustment methodology.

Next, the variation of the production process is extremely significant regarding the use of this technique. We will investigate reducing the variability in calculating the buffer and the related promised date lateness.

Additionally, a complete set of analysis studies should be completed which characterize the performance of the various methods of buffer calculation and adjustment. We have completed preliminary studies that are promising, concerning the performance of the techniques posed. Basically, if we are able to fit a model to the time series data on the lateness to schedule, then we can use that model in buffer estimation and promise date calculation. Current studies have attained promising R-squared values on the order of 80%.

Finally, we will also investigate the incorporation of other methods of forecasting which are not as sensitive to the variance of the process. These would include support vector machines and neural networks.

ACKNOWLEDGMENTS

We wish to acknowledge the following students who supported this work: Jennifer Pittman and Andres Lucas. This research is funded by the National Science Foundation, Grant No. DMI0111082, “Real Time Promising for Authority Domains Operating in a Build to Order Mode,” Phase II, S. Moses, H. Grant, S. Pulat, and L. Gruenwald, 2001.

REFERENCES

Abbott, R. and H. Garcia-Molina. 1992. Scheduling Real-Time Transactions: A Performance Evaluation, *ACM Transaction on Database Systems*, 17 (3), 513-560.

Bagchi, U., F. M. Julien, and M.J. Magazine. 1994. Note: Due-Date Assignment to Multi-Job Customer Orders, *Management Science*, 40 (10), 1389-1392.

Baker, K.R. and J. W. M. Bertrand. 1981. An Investigation of Due-Date Assignment Rules with Constrained Tightness, *Journal of Operations Management*, 1 (3), 109-120.

Bertrand, J.W.M. 1983. The Effect of Workload Dependent Due-Dates on Job Shop Performance, *Management Science*, 29 (7), 799-816.

Chen, Y., and L. Gruenwald. 1996. Effects of Deadline Propagation on Nested Transactions in Real-Time Database Systems, *Information Systems,* Special Issue on Real-Time Database Systems, 21 (1), 103-124.

Chen, Z.L. 1996. Scheduling and Common Due Date Assignment with Earliness-Tardiness Penalties and Batch Delivery Costs, *European Journal of Operational Research*, 93, 49-60.

Cheng, T.C.E. 1986. Optimal Due-Date Assignment in a Job Shop, *International Journal of Production Research*, 24 (3), 503-515.

Cheng, T.C.E. 1987. Optimal TWK-Power Due-Date Determination and Sequencing, *International Journal of Systems Sciences*, 18 (1), 1-7.

Cheng, T.C.E. 1988. Integration of Priority Dispatching and Due-Date Assignment in a Job Shop, *International Journal of Systems Science*, 19 (9), 1813-1825.

Cheng, T.C.E. and M. C. Gupta. 1989. Survey of Scheduling Research Involving Due Date Determination Decisions, *European Journal of Operational Research*, 38, 156-166.

Cheng, T.C.E. 1991. Optimal Assignment of Total-Work-Content Due-Dates and Sequencing in a Single-Machine Shop, *Journal of the Operational Research Society*, 42, 177-181.

Cheng, T.C.E., C. Oguz, and X. D. Qi. 1996. Due-Date Assignment and Single Machine Scheduling with Compressible Processing Times, *International Journal of Production Economics*, 43, 107-113.

Conway, R.W. 1965. Priority Dispatching and Job Lateness in a Job Shop, *Journal of Industrial Engineering*, 16 (4), 228-237.

De, P., J. B. Ghosh, and C. E. Wells 1992. Optimal Due-Date Assignment and Sequencing, *European Journal of Operational Research*, 57, 323-331.

De, P., J. B. Ghosh, and C. E. Wells, C.E. 1994a. Due-Date Assignment and Early/Tardy Scheduling on Identical Parallel Machines, *Naval Research Logistics*, 41 (1), 17-32.

De, P., Ghosh, J.B., and Wells, C.E. 1994b. Solving a Generalized Model for CON Due Date Assignment and Sequencing, *International Journal of Production Economics*, 34, 179-185.

Duenyas, I. (1995a) Single Facility Due Date Setting with Multiple Customer Classes, *Management Science*, 41 (4), 608-619.

Duenyas, I. and W. J. Hopp. 1995b. Quoting Customer Lead Times, *Management Science*, 41 (1), 43-57.

Eilon, S.E. and I. G. Chowdhury, 1976. Due Dates in Job Shop Scheduling, *International Journal of Production Research*, 14 (2), 223-237.

Gruenwald, L., and J. Cheng. 1997. A Logging Technique Based on Transaction Types for Real-Time Databases, *Real-Time Database and Information Systems-Research Advances*, Editors Azer Bestavros and Victor Fay-Wolfe, Boston: Kluwer Academic Publishers, 379-392,

Gruenwald, L, L. Brown, R. Dirckze, S. Guinepain, C. Sanchez, B. Summers, and S. Vanichayobon. 1999. Database Research at The University of Oklahoma, *ACM SIGMOD RECORD*, 28(3), 73-78.

Heard, E.L. 1976. Due-Dates and Instantaneous Load in the One-Machine Shop, *Management Science*, 23 (4), 444-450.

Hopp, W.J. and M. L. Roof Sturgis. 2000. Quoting Manufacturing Due Dates Subject to a Service Level Constraint, *IIE Transactions*, 32 (9), 771-784.

Kaplan, A.C. and A. T. Unal. 1993. A Probabilistic Cost-Based Due Date Assignment Model for Job Shops, *International Journal of Production Research*, 31 (12), 2817-2834.

Lawrence, S.R. 1995. Estimating Flowtimes and Setting Due-Dates in Complex Production Systems, *IIE Transactions*, 27, 657-668.

Li, C.L and T. C. E. Cheng. 1999. Due Date Determination with Resequencing, *IIE Transactions*, 31 (2), 183-188.

Luss, H. and M. B. Rosenwein. 1993. A Due Date Assignment Algorithm for Multiproduct Manufacturing Facilities, *European Journal of Operational Research*, 65, 187-198.

Palaka, K., S. Erlebacher, and D. H. Kropp. 1998. Lead-Time Setting, Capacity Utilization, and Pricing Decisions Under Lead-Time Dependent Demand, *IIE Transactions*, 30, 151-163.

Pritsker, A and J. O'Reilly. 1999. *Simulation with Visual SLAM and AweSim*, John Wiley.

Ragatz, G.L. and V. A. Mabert. 1984a. A Framework for the Study of Due Date Management in Job Shops, *International Journal of Production Research*, 22 (4), 685-695.

Ragatz, G.L. and V. A. Mabert. 1984b. A Simulation Analysis of Due Date Assignment Rules, *Journal of Operations Management*, 5 (1), 27-39.

Raghu, T.S. and V. Rajendran. 1995. Due-Date Setting Methodologies Based on Simulated Annealing—An Experimental Study in a Real-Life Job Shop, *International Journal of Production Research*, 33 (9), 2535-2554.

Rajasekera, J.R., M. R. Murr, and K. C. So. 1991. A Due-Date Assignment Model for a Flow Shop with Application in a Lightguide Cable Shop, *Journal of Manufacturing Systems*, 10 (1), 1-7.

Seidmann, A. and M. L. Smith. 1981. Due Date Assignment for Production Systems, *Management Science*, 27 (5), 571-581.

Shanthikumar, J.G. and U. Sumita. 1988. Approximations for the Time Spent in a Dynamic Job Shop with Applications to Due-Date Assignment, *International Journal of Production Research*, 26 (8), 1329-1352.

Smith, M.L. and A. Seidmann. 1983. Due Date Selection Procedures for Job-Shop Simulation, *Computers and Industrial Engineering*, 7 (3), 199-207.

Tsai, C.H., G. T. Chang, and R. T. Li. 1997. Integrating Order Release Control with Due-Date Assignment Rules, *International Journal of Production Research*, 35 (12), 3379-3392.

Udo, G. 1993. An Investigation of Due-Date Assignment Using Workload Information of a Dynamic Shop, *International Journal of Production Economics*, 29, 89-101.

Vig, M.M. and K. J. Dooley. 1991. Dynamic Rules for Due-Date Assignment, *International Journal of Production Research*, 29 (7), 1361-1377.

Vig, M.M. and K. J. Dooley. 1993. Mixing Static and Dynamic Flowtime Estimates for Due-Date Assignment, *Journal of Operations Management*, 11, 67-79.

Weeks, J.K. 1979. A Simulation Study of Predictable Due-Dates, *Management Science*, 25 (4), 363-373.

Wein, L.M. 1991. Due-Date Setting and Priority Sequencing in a Multiclass M/G/1 Queue, *Management Science*, 37 (7), 834-850.

Wein, L.M. and P. B. Chevalier. 1992. A Broader View of the Job-Shop Scheduling Problem, *Management Science*, 38 (7), 1018-1033.

Weng, Z.K. .1999. Strategies for Integrating Lead Time and Customer-Order Decisions, *IIE Transactions*, 31 (2), 161-171.

Yano, C.A. 1987. Setting Planned Leadtimes in Serial Production Systems with Tardiness Costs, *Management Science*, 33 (1), 95-106.

AUTHOR BIOGRAPHIES

HANK GRANT joined the faculty at the University of Oklahoma in December of 1993 as Director of the School of Industrial Engineering and Southwestern Bell Professor. Since his arrival, he has been actively involved in all aspects of the growth of School as well as in the startup of the Center for the Study of Wireless Electromagnetic Compatibility. He is currently Director of the Center and Dugan Professor of Industrial Engineering. Prior to joining the University of Oklahoma, Dr. Grant was with the National Science Foundation in Washington, DC, where he directed programs in Production Systems, Engineering Design, and Operations Research. Before that, he was Director of the Measurement and Manufacturing Systems Laboratory with Hewlett-Packard in HP Labs. Before joining HP, Dr. Grant was involved in the startup and development of two Industrial Engineering software companies: Pritsker Corporation and FACTROL. Dr. Grant is a Fellow of the Institute of Industrial Engineers and is a member of the following societies: INFORMS, the Institute of Industrial Engineers, Tau Beta Pi and The Institute of American Entrepreneurs. His email address is <hgrant@ou.edu>.

SCOTT MOSES has been Assistant Professor of Industrial Engineering at the University of Oklahoma since 1999. His current research activity focuses on scalable algorithms for real-time order promising and on high-speed tactical-level planning of production tasks and material flow in large discrete systems. Emphasis in research is given to computationally oriented approaches that work well for industrial sized systems and that have the requisite flexibility to yield solutions meaningful to industry. Prior to joining the OU faculty, he was with i2 Technologies for six years, where he was highly active in the design, development, implementation, and integration of advanced decision support software for supply chain planning and management. Previously he worked at Pritsker Corporation where he constructed discrete-event simulation models of manufacturing operations. He received his Ph.D. degree in industrial engineering from Purdue University in 1995 and holds an M.S. degree in industrial engineering and a B.S. degree in mechanical engineering, both from Oklahoma State University. He is a member of IIE, INFORMS, AAAI, IEEE and ASEE. His email address is <moses@ou.edu>.

DAVID GOLDSMAN is a Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. His research interests include simulation output analysis and ranking and selection. He also studies applications arising in the healthcare field. Dave has been an active participant in the Winter Simulation Conference — he is currently on the Board of Directors, and was the 1995 Program Chair and 1992 Associate *Proceedings* Editor. He is a member of IIE, INFORMS, and SCS. Dave received his undergraduate education at Syracuse University, and his Ph.D. in Operations Research and Industrial Engineering from Cornell in 1984. His email and web addresses are <sman@isye.gatech.edu> and <www.isye.gatech.edu/~sman>.