Proceedings of the 1992 Winter Simulation Conference
ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson

REAL-TIME SIMULATION FOR DECISION SUPPORT IN CONTINUOUS FLOW MANUFACTURING SYSTEMS

Paul G. McConnell

RWD Technologies Inc.
10480 Little Patuxent Parkway, Suite 1200
Columbia, Maryland 21044, U.S.A.

D. J. Medeiros

Department of Industrial Engineering
Penn State University
University Park, Pennsylvania 16802, U.S.A.

ABSTRACT

Conventional, hierarchical, and optimizing approaches to real-time decision support using simulation are developed and compared. Issues related to the use of simulation for real time decision support are considered. The decision support systems are tested using an emulation of a continuous flow manufacturing system.

1 INTRODUCTION

The dynamic control of manufacturing systems is an issue of concern to both the manufacturing and the academic communities (Harmonosky and Robohn 1991). Without an effective means of control, a manufacturing system will not survive in the long term, especially as markets become more competitive.

One idea for the effective control of manufacturing systems is to make control decisions in real time by using simulation. Real-time control by simulation involves initializing a simulation model with the current status of the system, and using the non-steady state results of the model to aid in making an immediate decision. Real-time control by simulation is a relatively new idea, made possible by the introduction of computers on the factory floor, advances in computer memory capacity and processing speed, and improvements in simulation software.

Figure 1 illustrates the use of real-time simulation for factory floor decision making. Any discrete change in a manufacturing system's status should be detected by sensors feeding information to a control system. The control system will determine if there is a pre-programmed response to this change of state. If there is, the control system retrieves the response and feeds it back to the physical system. If there is not, an external decision is required. The real-time simulation model provides some form of decision support when an external decision is required. The model is initialized with the manufacturing system's current status, and used to either evaluate a decision or to generate a decision. The decision is fed back to the control system, then in turn to the physical system.

For real-time decisions, the correct decision depends not only on the results of the computation, but also on the time it takes to produce the decision. A real-time decision is not necessarily instantaneous; instead, a decision is considered a real-time decision as long as it is made within a certain time increment. That time increment is defined by a later event which will invalidate or at least degrade the decision (Stankovic 1988). Depending on the nature of the system, a decision made in, for example, five minutes may still be considered real-time.

2 MANUFACTURING APPLICATIONS OF REAL-TIME SIMULATION

Currently, there is no consistent perspective on how a decision tool based on real-time simulation should be developed and used. Proposed methods include:

- Using a set of stand-alone models on the shop floor,
- Using one model as a real-time process monitoring tool with look-ahead system assessment capabilities,
- Using a simplex optimization algorithm such as the Nelder-Mead method in the simulation model, or,
- Using the real-time simulation with an expert system.

Gaafer and Cochran (1989) perceive a real-time simulation model as a shop floor tool that consists of a set of models, any one of which can be invoked by the user. The user enters changes to the experimental framework of the chosen model, and selects a purpose, either estimation or optimization. The advantage of using a set of models is that only those factors relative to a specific decision need be in a model, thereby decreasing run time.

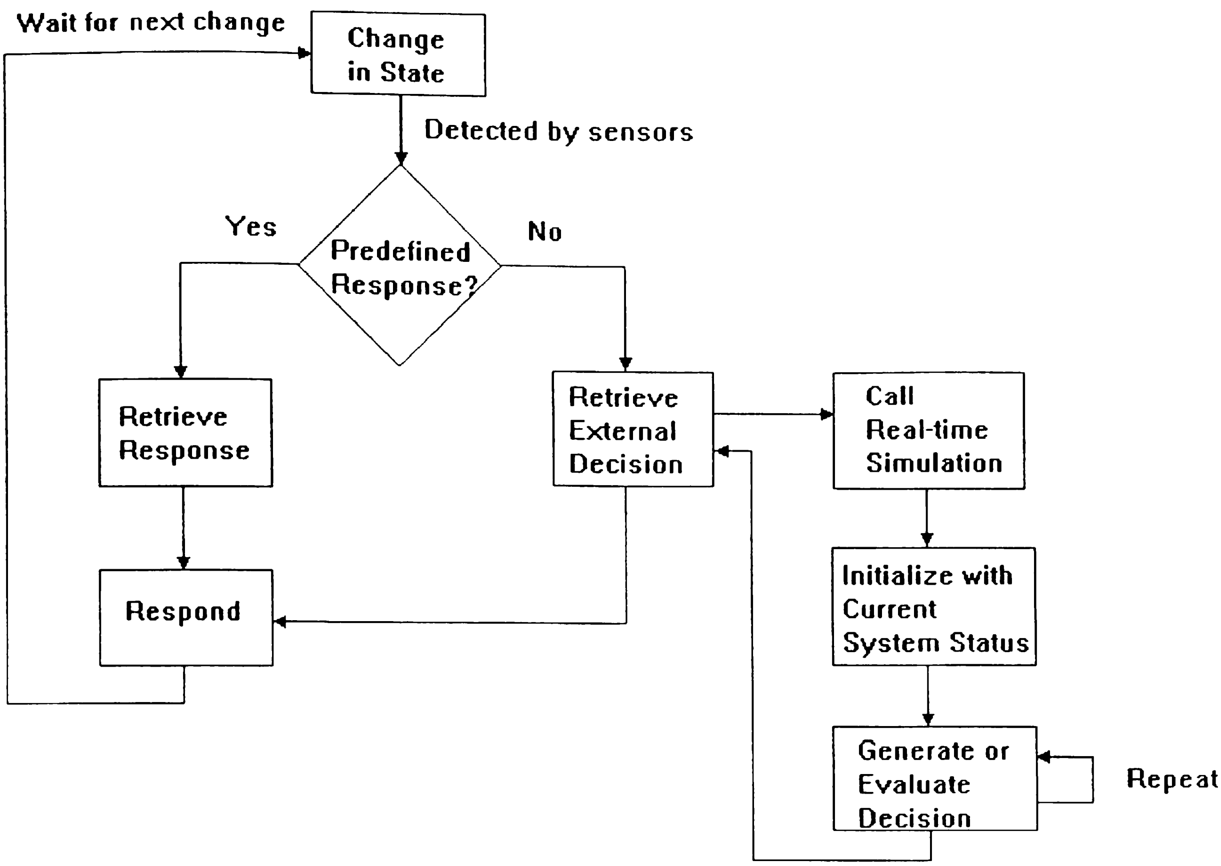

Figure 1: Real Time Simulation for Decision Support

Rather than using a set of models, Harmonosky (1990) proposes linking the computer simulation with the physical system, and having the simulation logic controlled by the actual system communication signals. When a decision is required, the model switches from a monitoring mode to a look-ahead mode. Some knowledge of the possible decisions that will be evaluated should exist; therefore, the user need not spend time making model modifications and re-compiling code when evaluating possible decisions. Having a priori knowledge of decisions and having the simulation always initialized with the current status of the real system speeds the process of evaluating a decision. However, switching the system from a look-ahead mode back to a monitoring mode may be difficult. Also, if different alternatives are to be compared using the real-time simulation, or if multiple replications of the simulation are required, then the initial system status must be saved and re-initialized for each alternative, thereby eliminating any savings in run time due to having the model already initialized.

In both these approaches, the user may have to evaluate an extensive series of decisions, or may simply evaluate poor decisions. If so, then the tool has lost both its real-time characteristics and its effectiveness. Using either an optimization method or an expert system shell is one possible way to alleviate these problems.

The combination of simulation and optimization can yield an effective decision support system. Beveridge and Schechter (1970) observe that the Box Complex Method and Nelder-Mead Method perform best over a wide range of problems. Pegden and Gately (1980) note that an optimization technique should attempt to minimize the number of simulation runs rather than the computational effort involved in the optimization algorithm. The number of iterations to optimal will directly effect the ability to maintain a real-time decision environment. Filip, Neagu, and Donciulescu (1983) present an optimization algorithm for real-time production control in a job shop. The algorithm is a two level hierarchical optimization method based on parametric decomposition.

When using an expert system shell with a real-time simulation model, Moore (1985) states that there are three basic considerations to be addressed in any attempt to apply expert system technology to real-time processes. First, an efficient means of real-time data access must be established. Second, the inference engine must be designed to operate in a real-time context -- its computational efficiency must be enhanced by every means possible. Finally, a simple and easy to use means of knowledge acquisition must be devised.

Wu and Wysk (1989) use a discrete event simulation to evaluate a set of dispatching rules for a short planning horizon. The rule with the best simulated performance for the time period is applied to the physical system. Wu and Wysk found that in most cases, their system performed better when compared to a single pass, static scheduling mechanism.

In conclusion, the literature presents a number of very different approaches to developing a real-time simulation tool for decision making. Possibly different types of applications, or even the specific characteristics of a system, will dictate which approach to use when constructing a real-time simulation model.

3 OBJECTIVE AND SYSTEM DESCRIPTION

The objective of this paper is to demonstrate the feasibility and utility of a real-time simulation model used for factory floor decision-support.

Simulation based decision support will be demonstrated for a continuous flow manufacturing system. A continuous flow manufacturing system is arguably better suited for control with a real-time simulation than a discrete part manufacturing system because significantly fewer events occur in a continuous manufacturing system. Real-time simulation has been successfully used to control the flow of gas through a pipeline system (Sjoen 1987).

In this paper, a simplified system that still contains the fundamental characteristics of continuous flow coating or pumping systems is used to evaluate simulation-based decision support. This type of system is used in industries as diverse as the food industry, paper industry, photographic industry, and the chemical industry.

The simplified system involves coating a flexible base with three perishable solutions, as shown in the schematic of Figure 2. There are two tanks for each solution; when a tank is empty, the flow switches to the other tank. The solutions are prepared in the tank; this task requires a significant amount of time. The solution flows out of a tank, past the tank switch, through a filter, into a surge tank, and into a hopper. Flow control meters before the filter and between the surge tank and the hopper control the rate of flow of the solution. The base is conveyed continuously, without interruption, to the hopper. The solutions flow out of the hopper onto the base, and the base is conveyed to a finishing section. The base can run at variable speeds.

4 MODELING ISSUES

To build a real-time simulation model for decision support in a continuous flow manufacturing system, a number of issues related to the construction and use of the real-time simulation model had to be resolved. They are:

- Type of decision (open or predefined).
- Modeling approach.
- Treatment of random events.
- Treatment of discrete operations in process at the decision point.
- Number of replications of the simulation.
- Horizon length.
- Initializing the system.

4.1 Type of Decision

The first issue is choosing between a model designed to analyze specific decisions or a flexible model which can handle any system decision that may arise. The latter approach is difficult to implement as it requires both a very flexible model (or family of models) and a well-designed user interface. We utilize the former approach, and characterize the decisions required as follows:

1. When should the refilling of an empty tank begin?
   A solution should be prepared soon enough so it is always available when needed, yet not so soon that it may perish if the system is unexpectedly interrupted.
2. What is the optimal speed for the base?
   The objective is to maximize throughput, while avoiding stopping the production line (stopping usually results in a large amount of waste for these kinds of systems). If a solution will soon perish, running faster may be in order. If a tank is not ready, and the paired tank will soon run out, running slower may be in order.
3. At what time should a product change occur?
   Usually in continuous flow systems there is a goal production level, but actual production need only be within some range around this goal. Therefore, production of a certain product is terminated at some point in the range where the quantity of solutions left in the tanks (which will be scrapped) is at a minimum.

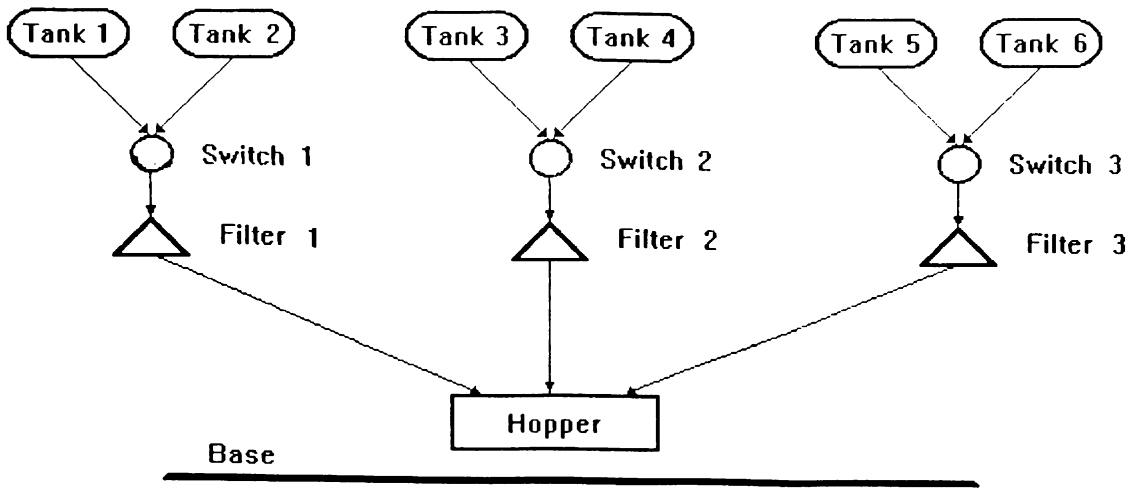

Figure 2: Three Solution Coating Process

4. When should a filter change occur?
   A sensor will detect when a filter is nearing its limit. Ideally, if a downtime event will occur, the filter will be changed during this event.

4.2 Modeling Approach

Several modeling approaches can be used including:

- A conventional simulation model.
- A set of simulation models, each using different bounds and depth of the system. (The set of models will be called a hierarchical model.)
- A simulation model using an optimization algorithm.
- A simulation model with a knowledge-based shell.
- A simulation model that serves as a process monitoring tool with look-ahead capabilities.

Simulation models with knowledge-based shells appear to be slow to yield a decision, and expensive to develop. A simulation model that functions as a process monitoring tool with look-ahead capabilities offers no advantage over a conventional simulation model for continuous manufacturing systems.

A conventional simulation model, a hierarchical model, and an optimization model were implemented. These simulation models include differential equations describing the solution flow and the flow of the base. The filter life is modeled based on time (in reality, filter life is a function of the amount and quality of solution flowing through the filter). Logic for detecting filters requiring change and being changed is included. Logic for detecting when a tank is empty, switching tanks, and refilling a tank if appropriate is included. Likewise, logic is included for testing the quality due to age of a solution and scrapping that solution if appropriate. Each model is initialized with the current status of the system, which is available from the process control system.

4.2.1 Conventional Model

The conventional simulation model serves as a decision evaluation tool for the decisions regarding when to refill a tank, at what speed to run the base, and when to change filters. The conventional simulation model serves as a step-wise optimization tool for evaluating when a product change should occur (that is, it evaluates all possibilities and provides the user with the best changeover point based on current system status).

The criterion for evaluating the decisions of when to refill a tank, when to change a filter, and at what speed to run the base is maximum output (i.e., decisions are made by the user based on obtaining the maximum production over time). The criterion for evaluating when to end a production run is to minimize the sum of the solutions left in the tanks subject to production being within the acceptable range (this ending time is reported to the user for his consideration when making a decision).

4.2.2 Hierarchical Model

The hierarchical model is really a set of three models, each evaluating different decisions. One model evaluates the optimal base speed, another model evaluates when to refill a tank, and a third model evaluates when to change a filter and when to terminate production for a specific product.

Each model simulates only the aspects of the system critical for the particular decision, thereby either decreasing model execution time or improving the quality of the decision. The model to evaluate speed assumes infinite filter life. The model to evaluate when to refill a tank also assumes infinite filter life, and uses stochastic values for critical refill parameters. The model to evaluate when to change a filter and when to terminate a production run is identical to the conventional model.

4.2.3 Optimization Model

The optimization model uses an objective function of maximizing throughput for the decisions of when to refill a tank, when to change a filter, and at what speed to run. The objective function for when to change products is the same as for the previous models. This model differs in that the optimal solution is computed and applied to the system without user intervention.

The Nelder-Mead Simplex method for simulation optimization was used because of its generality, efficiency, and robustness (Barton and Ivey 1991).

4.3 Treatment of Random Events

Examination of recent literature reveals that the authors often assume that the simulated systems are deterministic during the look-ahead horizon. Although the assumption of a deterministic look-ahead horizon will allow the use of data from a single replication of the real-time simulation model, the results of the simulation may not present an accurate representation of the physical system.

Goldsman, Swain, and Withers (1990) discuss the specialized problem of analyzing simulations in a short, single replication. Our treatment of random variables is based on their work.

Table 1 gives a summary of random variables of the manufacturing system, and how they are handled in each real-time simulation model. In the conventional model, all stochastic characteristics of the system are described with deterministic values. The filter change time and the time to refill a tank are replaced by a deterministic value describing the time when eighty percent of all refills or filter changes will have occurred. The quantity in a full tank, the filter life, and the solutions' life are described by a deterministic value equal to the mean of the respective distribution.

In the hierarchical model used to evaluate speed, all operation times are constant. The model used to evaluate when to refill a tank assumes constant operation times except for a truncated distribution used for time to refill the tank; thus multiple replications must be run.

The optimization model uses the sample distributions for the time to refill a tank and the time to change a filter. All other random events are replaced by a deterministic value equal to the mean of the distribution.

4.4 In-process Operations

This issue occurs when the system contains random events, and thus may have an operation with a random delay time in process at the decision point. In this study, for operations already in process, all models use a deterministic evaluation of time remaining for the operation. If the actual time of an operation exceeds this deterministic value, then the model assumes that completion of the operation will require one additional standard deviation of time (based on the sample distribution).

4.5 Number of Replications

If the real-time simulation is deterministic, as in the conventional model, then obviously only one replication is required for each decision alternative that is being evaluated. For simulations that include randomness, the number of replications is set to allow for acceptable estimates of the parameters of interest. Setting the number of replications in advance is possible because the decisions to be made are known in advance. For the hierarchical model, six replications were used for each alternative evaluated. The optimization model used one replication for each combination of parameters, with a maximum of 30 replications allowed to converge on an optimal solution.

4.6 Horizon Length

The horizon length will depend on the decision that is to be made. The look-ahead period should end when an expected event in the system will change the status of the parameter of concern. For example, when the real-time simulation model is used to determine when to schedule a filter change, the look-ahead horizon will end (automatically) when either the filter fails or the system goes down for some other reason allowing the filter to be changed.

4.7 Initializing the System

Initialization issues include what data to pass to the real-time simulation model, how to pass the data, and when to pass the data. For this manufacturing system, only the state of continuous variables (a relatively small amount of data compared to the amount that would be

Table 1: Representation of Random Variables in Real-time Simulation Models

Variable	Conventional	Hierarchical (speed)	Hierarchical (refill)	Optimizing
Speed of base	constant	constant	constant	constant
Acceleration of base	constant	constant	constant	constant
Solution flow rates	constant	constant	constant	constant
Solution life	constant	constant	constant	constant
Solution remaining in "empty" tank	constant	constant	constant	constant
Refill time	constant	constant	truncated distribution	distribution
Refill level	constant	constant	constant	constant
Filter life	constant	constant	constant	constant
Filter change time	constant	constant	constant	distribution
System downtime	eliminated	eliminated	eliminated	eliminated
Time between failures	eliminated	eliminated	eliminated	eliminated

required for a discrete-parts system), and the status of discrete events in process, if any, need to be passed. Therefore, data is passed to the real-time simulation models only when a decision is required. This data base is also temporarily stored so the model can be accurately re-initialized for each replication.

5 IMPLEMENTATION

An actual manufacturing system was not available for evaluating the feasibility and utility of real-time simulation. Instead, a real-time emulation of this system has been constructed on an engineering workstation. The emulation generates data representing the dynamic behavior of the system and also models the process control system. Since information about the actual system comes not from observation of the system, but from data from the process sensors, we concluded that an emulation would be an acceptable alternative to using an actual manufacturing system. An advantage of using an emulation is that the emulation produces pseudorandom sequences that are dynamic but repeatable. Different real-time simulation models can be tested against known, repeatable scenarios.

The emulation was constructed using the combined discrete-continuous simulation modeling capability of SIMAN; it runs on a VAXStation 3100. The emulation was run in real-time by fixing the step size of the continuous simulation and suspending execution of the model on the computer until the specified increment of real-time has passed.

Some discrete events can occur between steps. These events are generated by the process monitoring system detecting conditions such as a tank reaching a minimum level or a filter nearing its limit. Discrete events could be handled in a similar manner by suspending the process until the event time. However, since the step size is small relative to the frequency of decisions (15 second step size verses several minutes between decisions) the difference is insignificant and was ignored in this study.

The emulation differs significantly from the simulation running on the control machine. All system variables that can possibly be random are treated as such in the emulation. The emulation also includes functions generating natural variation, whereas the control system describes this behavior as it would commonly be simulated. For example, filter life is primarily a function of the quality and quantity of the solution flowing through the system. In the emulation, filter life is described by a differential equation that models the number of particles that the filter has trapped; each batch of each solution has a randomly generated number of particles per unit volume. In contrast, the decision support simulation models calculate filter life based on the time since the filter was installed.

The emulation runs as a single process on the engineering workstation. The workstation is networked to another workstation which runs the decision-support simulation. This provides an ideal environment for testing decision-support systems. The emulation is an independent system which sends data similar to real process data via the network to the decision support system. After an appropriate configuration for the decision support system has been determined, the emulation can be discarded and code can be written on the process control computer to transfer the required information to the decision-support system.

6 RESULTS

The real-time simulation models were tested on two different production runs of the system, each producing 100,000 linear feet of coated base. The maximum production speed was set to 400 linear feet per minute. Production could be terminated anywhere between 90,000 and 110,000 linear feet.

Table 2 summarizes the decisions made using the real-time simulation models for the first production run, as well as the simulated production and waste solution remaining in the tanks. The second production run gave similar results for the speed, refill, and filter decisions and so is not included in the table. For each controllable production event shown in Table 2, the clock time to make the decision and the decision itself is listed. Some of these decisions are discussed further below.

6.1 Refill Decision

All models were capable of quickly making a decision related to when to refill the raw materials. Because the hierarchical model required both operator input and multiple replications, it took longer than the other two models to evaluate an alternative. However, the hierarchical model's performance was acceptable because the shortest recommended delay until refill was 5 minutes and the longest time for running the hierarchical model was 4.5 minutes.

In one case (event 11), the optimization model recommended no delay; clearly this is infeasible. Since time to execute the optimization model does not vary significantly, a lower bound could be placed on delay time to avoid this problem.

The conventional model used a deterministic value for refill time equal to the upper limit of the truncated sample distribution for refill time. Using this upper limit provided a safety margin. The other two models used a distribution to represent refill time. The optimization model produced results that were generally consistent with those from the conventional model. In 7 out of 12 decisions, the hierarchical model suggested the same or a longer delay until refill than the other two models; in all except one of these (event 17), the longer delay was appropriate. In event 17, the hierarchical model caused 2 minutes of unscheduled downtime.

All refill decisions (except the two noted above) were good decisions. It is difficult to determine which of the models is generating the best delay response because system conditions are not identical due to differences in previous decisions.

6.2 Filter Decision

All three models quickly generated a filter change decision. The models yielded the same response for filters 1 and 3. For filter 2 (event 13), the conventional and hierarchical models recommended against a filter change while the optimization model scheduled a change. Thus, the optimization model caused two shutdowns for filter changes while the other models accomplished this in one shutdown. Again, because of differences in previous decisions, there is insufficient information to conclude that the optimization model made an incorrect decision, although the other runs suggest that two shutdowns were unnecessary.

6.3 End Production Decision

Neither the conventional nor the hierarchical model performed consistently well in advising when to end production. In the first production run, both models suggested the best of the three reasonable production end times. (Production should end when both tanks of a pair are empty within the allowable range of linear feet produced). In the second run, however, both models advised the worst of three reasonable end times. The optimization model advised the best end time in both production runs.

To illustrate this, the hierarchical model's end time decisions are described. Table 3 contains the three reasonable production end points for each of the two runs. For the first run, the hierarchical model advised ending production at 94,709 feet, which minimizes solution waste. For the second run, the model advised ending at 99,943 feet, which maximizes waste, and did not recognize the third point as a possible solution.

7 CONCLUSION

Simulation models can be effective real-time decision support tools for controlling manufacturing systems similar to the one described in this paper. The conventional, hierarchical, and optimization models

Table 2: Results of Real Time Simulation for Production Run 1

	Conventional		Hierarchical		Optimization	
Event	Time needed (min.)	Decision	Time needed (min.)	Decision	Time needed (min.)	Decision
1. Increase speed	.75	full speed	.75	full speed	1.	full speed
2. Tank 1 empty	1.25	delay 5 min	3.25	delay 10 min	1.25	delay 5 min
3. Tank 3 empty	2.	delay 12 min	2.5	delay 12 min	1.25	delay 12 min
4. Increase speed	1.	full speed	.75	full speed	1.25	full speed
5. Tank 5 empty	1.5	delay 26 min	2.25	delay 26 min	1.25	delay 17 min
6. Tank 2 empty	1.75	delay 5 min	4.5	delay 10 min	1.25	delay 3 min
7. Tank 4 empty	1.25	delay 17 min	2.5	delay 12 min	1.5	delay 10 min
8. Increase speed	1.25	full speed	.75	full speed	1.25	full speed
9. Tank 6 empty	2.	delay 31 min	1.75	delay 21 min	1.25	delay 30 min
10. Tank 1 empty	3.	delay 10 min	3.5	delay 5 min	1.5	delay 14 min
11. Tank 3 empty	2.	delay 12 min	2.5	delay 17 min	1.	delay 0 min
12. Filter 3 warning	1.5	do not change	1.	do not change	1.	do not change
13. Filter 2 warning	1.	do not change	.75	do not change	1.	change now
14. Tank 2 empty	1.75	delay 10 min	2.25	delay 10 min	1.5	delay 7 min
15. Filter 1 warning	.75	change now	1.	change now	1.	change now
16. Tank 5 empty	2.	delay 26 min	2.25	delay 12 min	1.25	delay 26 min
17. Tank 4 empty	2.5	delay 7 min	2.5	delay 26 min	1.	delay 7 min
18. Tank 1 empty	1.5	delay 5 min	2.25	delay 10 min	1.25	delay 13 min
19. Production can end	2.	do not refill any tanks	2.	do not refill any tanks	2.5	do not refill any tanks
Total produced		96,326 ft.		94,709 ft.		96,316 ft.
Production time		311 minutes		307 minutes		310 minutes
Production rate (average)		309.4 ft/min		308.5 ft/min		310.7 ft/min
End Waste		642 L		794 L		550 L

generated good results for most decision alternatives, but the optimization model was more successful in evaluating the production end time decision. All models operate quickly enough to be used in a real time mode for the sample production system, although timing may become critical with more complex systems.

Table 3: Production End Times for Hierarchical Model

	Total linear feet	Total waste (L)
Run 1	94,709	794
	99,109 (estimated)	1034 (estimated)
	105,709 (estimated)	896 (estimated)
Run 2	95,741	946
	99,943	1233
	108,609 (estimated)	768 (estimated)

The average time to make a decision was shortest for the optimization model and longest for the hierarchical model. Timing for the optimization model was consistent because iterations were not stopped when the system converged. Timing for the other models varied depending on the number of alternatives considered. In more complex systems, it may be necessary to restrict the number of alternatives or iterations to achieve a solution in an acceptable time. Choice of model type, model execution time, and need to implement restrictions on run time depend on the characteristics of the manufacturing system.

REFERENCES

Barton, R. R., and J. S. Ivey, Jr. 1991. Modifications of Nelder-Mead simplex method for stochastic simulation response optimization. In *Proceedings of the 1991 Winter Simulation Conference*, ed. B. C. Nelson, D. Kelton, and G. M. Clark, 945-953. Institute of Electrical and Electronics Engineers, Phoenix, Arizona.

Beveridge, G. S. G., and R. S. Schechter. 1970. *Optimization Theory and Practice*. New York: McGraw-Hill.

Filip, F. G., G. Neagu, and D. A. Donciulescu. 1983. Job shop scheduling optimization in real-time production control. *Computers in Industry* 4:395-403.

Gaafer, L. K., and J. K. Cochran. 1989. Developing a real-time simulation tool for shop-floor decision making. In *Proceedings of the 1989 Summer Computer Simulation Conference*, ed. J. Clema, 79-85. Institute of Electrical and Electronics Engineers, Austin, Texas.

Goldsman, D., J. J. Swain, and D. Withers. 1990. Single replication simulation. In *Proceedings of the 1990 Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, and R. E. Nance, 387-391. Institute of Electrical and Electronics Engineers, New Orleans, Louisiana.

Harmonosky, C. H. 1990. Implementation issues using simulation for real-time scheduling, control, and monitoring. In *Proceedings of the 1990 Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, and R. E. Nance, 595-598. Institute of Electrical and Electronics Engineers, New Orleans, Louisiana.

Harmonosky, C. H., and S. F. Robohn. 1991. Real-time scheduling in computer integrated manufacturing: A review of recent research. IMSE Working Paper 90-108, Industrial Engineering Department, Penn State University, University Park, Pennsylvania.

Moore, R. L. 1985. Adding real-time expert system capabilities to large distributed control systems. *Control Engineering* 32:118-121.

Pegden, C. P., and M. P. Gately. 1980. A decision optimization module for SLAM. *Simulation* 37:18-25.

Sjoen, K. 1987. Real-time simulation on the Statpipe system. *Pipe Line Industry* 53:46-51.

Stankovic, J. A. and K. Ramamritham, ed. 1988. *Hard real-time systems: A tutorial.* Washington: Computer Society Press of the Institute of Electrical and Electronics Engineers.

Wu, S. D., and R. A. Wysk. 1989. An application of discrete-event simulation to on-line control and scheduling in flexible manufacturing. *International Journal of Production Research* 29:1603-1623.

AUTHOR BIOGRAPHIES

PAUL G. MCCONNELL is an Associate Engineer with RWD Technologies, Inc. He consults in Total Quality Systems, Systems Modeling, and Safety Systems primarily for the chemical and petroleum industries. He holds a BS in Industrial Engineering and Operations Research from Virginia Polytechnic Institute and State University, and an MS in Industrial Engineering and Operations Research from Penn State.

D. J. MEDEIROS is an Associate Professor in the Department of Industrial Engineering at Penn State University. Her research interests include design, analysis and control of manufacturing and material handling systems, and manufacturing applications of coordinate measuring machines. She holds a BSIE from the University of Massachusetts, and an MSIE and PhD from Purdue University. She is Manufacturing Applications Track Coordinator for the 1992 Winter Simulation Conference.