Proceedings of the 1993 Winter Simulation Conference
G.W. Evans, M. Mollaghasemi, E.C. Russell, W.E. Biles (eds.)

SIMULATION FOR REAL-TIME DECISION MAKING IN MANUFACTURING SYSTEMS

Paul Rogers
Rosalee J. Gordon

Division of Manufacturing Engineering
Department of Mechanical Engineering
The University of Calgary
Calgary, Alberta, T2N 1N4
CANADA

ABSTRACT

This paper addresses the use of simulation for the support of real-time decision-making in manufacturing systems, focusing on dynamic scheduling decisions. Previous work is critically surveyed and the most salient unresolved issues both for future academic research on the approach and for its successful industrial application are highlighted.

1 INTRODUCTION

This paper is concerned with the use of discrete-event simulation as a component in real-time decision support tools for the control of manufacturing systems. The focus will be on those types of "control" decision dealing with the dynamic scheduling (or rescheduling) of a manufacturing facility in response to a wide range of unexpected events (or "stimuli").

Increasing competitive pressures on manufacturing organisations have prompted increasing interest in scheduling problems with the goal being to find ways of reducing lead time and increasing productivity (without reducing quality). Significant performance improvements can accrue from making scheduling decisions better, and a wide range of techniques, from the disciplines of control theory, operations research (O.R.), and artificial intelligence (A.I.), have been investigated as candidates for the construction of decision support tools. Discrete-event simulation is one such technique and simulation-based approaches to dynamic, finite capacity scheduling have been investigated with much interest over the last decade. The purpose of this paper is to review the work carried out to date, critically assess the current state and future potential of the approach, and highlight some of the problems which remain to be solved if good solutions to real-time scheduling problems are to be found.

The structure of the remainder of the paper is as follows. First, the dynamic scheduling problem is discussed in the context of wider production planning and control issues, and the justification for pursuing simulation-based approaches is established. Next the literature on the technique is critically surveyed moving from simple models to more sophisticated approaches including those incorporating elements of other decision-making tools (from the O.R. and A.I. domains). Following this, some of the major unresolved issues concerning the future development of simulation-based scheduling are discussed highlighting both academic/research and industrial/application concerns.

2 DYNAMIC SCHEDULING IN PRODUCTION PLANNING AND CONTROL

Figure 1 shows a simplified diagram of the production planning and control system for a discrete component manufacturer. Planning and control decisions are made in a three level hierarchy moving from the aggregate, long term planning level at the top (defining the Master Production Schedule), through medium term requirements analysis (probably using a variant of materials requirements planning), to short term production activity control at the bottom. To support the decision process, the modules in the hierarchy need access to a wide range of information including: product design, bill of materials and process plans; machine and labour availability and capability; system demand in the form of forecasts and orders.

The main concern of this paper is with the lowest level of the hierarchy, production activity control (PAC), which is responsible for executing the plans devised by higher levels. At this level, the finite capacity of resources is a key concern (higher levels may assume capacity is infinite or use simplistic models of the true capacity constraints). The central decision of PAC is scheduling, i.e. the allocation of a set of resources over time to perform a set of tasks to meet desired performance objectives. This scheduling task, however, is a very dynamic one since changes in any of the planning input information noted above, or any difference between the actual and expected execution of a plan on the shop floor, may necessitate some form of rescheduling at one or more levels in the hierarchy.

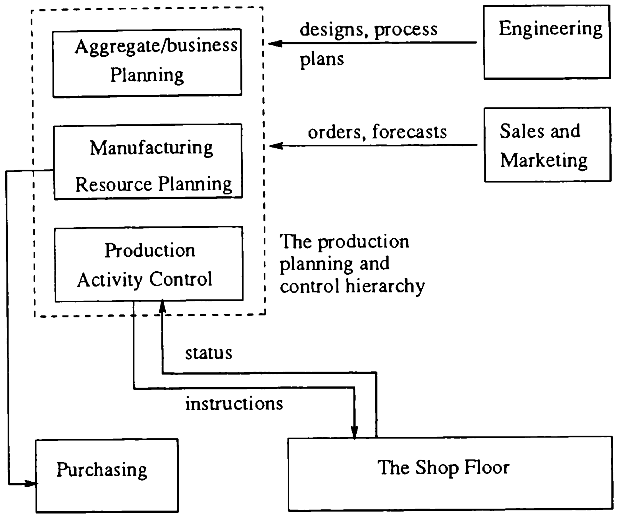

Figure 1: Basic Production Planning and Control System for Discrete Product Manufacture

2.1 The Dynamic Scheduling Problem

Scheduling problems are notoriously difficult even in situations where the production requirements and resource characteristics are completely known in advance (French 1982). Manufacturing systems typically involve considerable complexity in terms of inter-task relationships, resource requirements and constraints making it rare that any "optimal" algorithm can be formulated. In situations where there is significant stochasticity and dynamism (the norm for manufacturing organisations where new orders arrive, machines break down, raw material goes missing, and many other unexpected events occur), the probability of defining an algorithm to derive a schedule which is optimal in any real sense is vanishingly small. What is required is some way of determining a new schedule in response to any change very rapidly, in "real-time", before the results of the change can adversely affect system performance.

2.2 A Basic Simulation-Based Finite Capacity Scheduler

Simulation offers one method of dynamically generating a new schedule in response to some change in requirements, capability, or objectives. This can be viewed as a non-traditional application of discrete-event simulation which has typically been used as an "off-line" tool for system design to answer questions of the type "If I choose system design X what will the performance be?" (e.g. X might be one of a set of candidate shop dispatching rules). In the off-line mode, alternative choices are evaluated once for each decision made during the design process and the selected option is implemented. In this on-line, real-time mode, as described by Rogers and Flanagan (1991), simulation is used to support repetitive operational decisions, e.g. those concerned with dynamic scheduling. Whenever a decision is to be made, in response to some change in the state of the controlled system, the candidate options are simulated, their performance compared, and the best is chosen.

The basic operation of a simulation-based finite capacity scheduler (Pritsker, Grant, and Duket 1990) is illustrated in Figure 2. The simulation model is linked to both the production planning and control system and to the shop floor. The model can be initialised to the current state of the shop floor and a human scheduler can use it to predict the performance of alternative decision options for responding to a change and thus to help make the most appropriate choice. The user should be able to easily evaluate likely alternatives, thus rapidly planning around the unpredictable event. A secondary potential use for the simulation model would be to allow marketing/sales personnel to use it to predict order completion times for customers on-line. This would ensure that promises made to customers were based on valid estimates of completion times, allowing rush orders to be bid on whenever the delivery dates can definitely be met and refused if insufficient capacity is available.

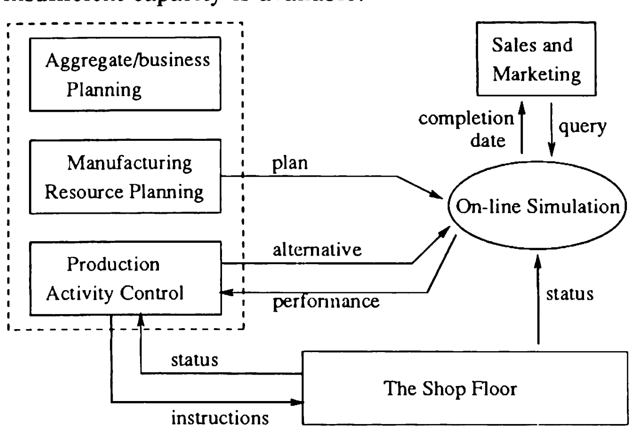

Figure 2: Role of a Simulation-based Finite Capacity Scheduler

The basic model in Figure 2 will be enhanced below, to better take into account more of the real constraints of the dynamic scheduling problem and to illustrate more sophisticated approaches to deriving better schedules.

3 A REVIEW OF THE LITERATURE ON SIMULATION-BASED APPROACHES

Simulation-based tools represent only one approach to the generation of schedules in real-time. We do not attempt a rigorous definition of "real-time" here but will take it to mean that a decision is classed as being one which must be made in real-time if failure to make the decision within a specified time limit will result in deterioration in the system performance (Harmonosky and Robohn 1991). In laypersons terms, a real-time decision support system is one which is "fast enough to be useful". Depending on the time characteristics of the manufacturing system being scheduled, a dynamic scheduler may still be considered real-time if it takes up to as much as an hour to derive a new schedule.

Figure 3 shows in more detail the operation of a dynamic rescheduling system, highlighting the volume of information required in order for valid decisions to be made. As can be seen in the figure, both relatively static information (e.g. products, components, process plans, and resource capabilities) and dynamic information (e.g. current shop floor state, and current and future known orders) must be accessible and accurate else the well known "garbage-in-garbage-out" principle will lead to the generation of useless schedules. A large appropriately structured database is needed to

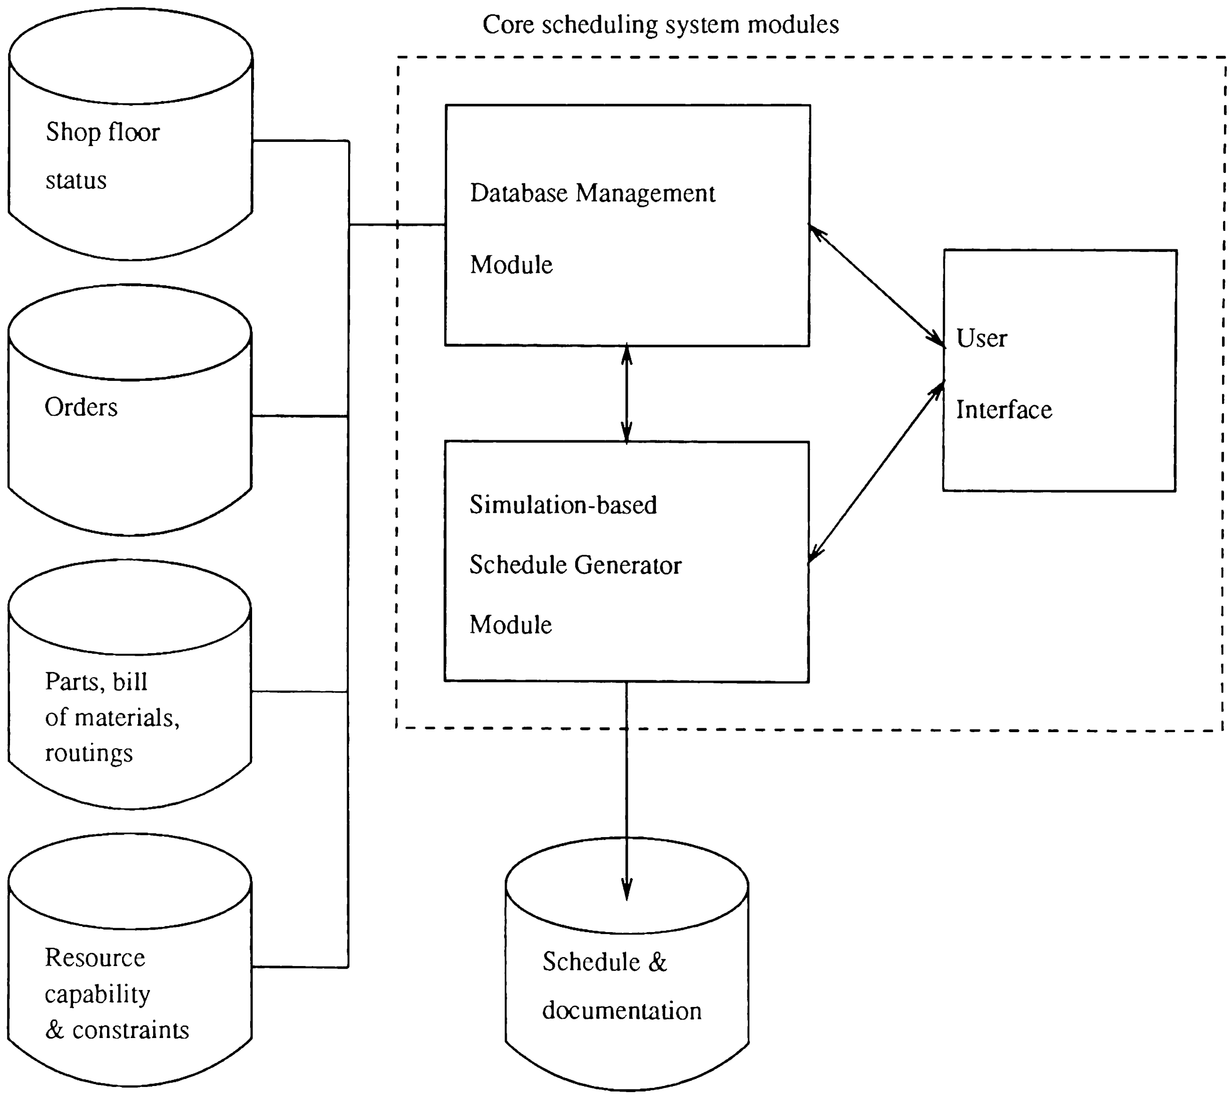

Figure 3: Structure of a Dynamic Rescheduling System

represent all this information to the level of detail required for finite capacity scheduling. The three essential components that this type of scheduling system must possess, as identified by Pinedo (1992), are: database management (for storing and accessing all of the required information); schedule generator; and user-interface.

The main assumption in Figure 3 is that the schedule generator module is based on discrete-event simulation, while further assumptions (some of which will be discussed later) concern the level of automation of the schedule evaluation and editing activity. The figure assumes that the system is not completely automated and that some human validation of the schedule is required before its release to the shop floor. As we will discuss later, the framework in the figure can be augmented with additional modules that increase the level of automation of the scheduling process, replacing some or all of the user's "intelligence" with knowledge-based modules to perform certain functions.

3.1 Pure Non-Simulation-Based Approaches

Before concentrating on simulation-based approaches, it should be noted that there has been much research reported on two alternative approaches to schedule generation, the first using O.R. tools and the second based on A.I. concepts. Space does not allow a review of this work here but the interested reader should consult Dempster, Lenstra, and Rinnooy Kan (1982) for more on the O.R. work and Kusiak (1990) for more on the A.I. approaches.

3.2 Pure Simulation-Based Approaches

A number of researchers have become interested in using a simulation-based schedule generator as the heart of the schedule generator module in systems similar to that shown in Figure 3. As noted earlier, the concept dates back at least ten years (Pritsker, Grant, and Duket 1990) and the potential benefits in terms of schedule performance have been well described by Grant (1988).

In its most basic form the following steps are involved in making use of a real-time simulation-based dynamic rescheduling system whenever a decision needs to be made:

1) The user initialises the simulation model by uploading information on the current state of released orders and production resources from the shop floor and downloading complete production planning and order requirements from higher level information systems.
2) The user identifies the set of actions which can be taken at this decision point as candidates for evaluation (i.e. possible responses to the need to reschedule).
3) The user makes one deterministic simulation run for each of the candidates and stores the results (in terms of predicted performance) of each option for later analysis.
4) After all options have been evaluated the one with the best performance given system objectives is selected and implemented.

The candidate actions identified in 2) could, for example, take the form of different dispatching rules to try (if the manufacturing system is a job shop) or a vector of dispatching rules (possibly a different one for each machine).

Various researchers have attempted to develop simulation-based schedulers according to this basic scheme, or variants of it. In the process, a number of possible enhancements to the basic approach outlined above have been proposed and include the following:

- In step 1) make the simulation model track the execution of the real system thereby saving model initialisation time when a new decision is to be made (Harmonosky 1990). This does introduce additional complexity in that some technique must be devised for switching the simulation between "tracking" and "prediction" modes.
- In step 1) there might be a range of possible models to choose from (instead of one fixed model) which differ in terms of the level of detail modelled. For certain decisions it might be possible to select a simpler (and hence faster) model to allow a decision to be reached more quickly and/or more options to be evaluated before the time when a decision must be made (McConnell and Medeiros 1992).
- In step 2) add a sub-module to automatically select a subset of candidate actions from the full set of all possible actions thereby saving time. This choice can be made automatically based on the scheduling objectives of the user (Wichmann 1990).
- In step 2) make use of a fixed, preselected set of alternative actions to evaluate, again saving time (Harmonosky 1990).
- In steps 3) and 4) keep evaluating more alternatives (keeping track of the current best option) until the time is reached when a decision must be made immediately. At this point choose the current best option (Harmonosky 1990).

Despite its apparent conceptual simplicity, there are a number of thorny issues to be resolved in developing this framework into a system capable of implementation, as highlighted by Harmonosky (1990) and Rogers and Flanagan (1991). These are described in greater detail in the next section but are summarised below.

- The simulation model must have fast and reliable links to other manufacturing information systems such as production planning and control and shop floor data collection.
- It has to be decided how the system will be used in terms of at least the following parameters:
  - rescheduling frequency
  - simulation run length
  - schedule evaluation criteria
  - the range of possible responses
  - the candidate selection process
  - user interface functionality.

3.3 Hybrid Approaches

Two of the deficiencies of a pure simulation-based approach to dynamic rescheduling are: i) only a very small number of possible schedules can be evaluated since each time-consuming simulation run generates only one scenario; ii) significant time delays occur due to the human user activities such as candidate action selection, performance evaluation, and schedule choice. Hybrid approaches, combining a simulation engine with a heuristic optimisation capability or with A.I. expertise are aimed at solving these problems.

To make a better selection of the limited number of schedules to generate and evaluate by simulation, analytical models (using hierarchical decomposition approaches to scheduling) can be used (Davis and Jones 1988). If the selected candidate actions can then be simulated concurrently on some parallel computing device (or network) then even faster real-time schedule generation is possible.

Artificial intelligence tools can be used to either replace or assist the human user in selecting candidate actions to evaluate or in interpreting the results of simulation to make a choice of which action to implement. Systems have been reported which assist the user in improving a schedule (Ben-Arieh 1988), or which provide assistance in choosing alternatives to simulate (Wu and Wysk 1989). Beyond these capabilities, it has been proposed (Wu and Wysk 1988) to include additional modules of expertise to improve the selection process over time by learning (i.e. monitoring how well previous selections performed and using this information to modify the selection criteria).

3.4 Commercial Systems and Industrial Implementations

There are now a wide range of companies supplying dynamic finite capacity scheduling tools, based on some combination of O.R. algorithms, knowledge-based systems, and discrete-event simulation. Although there has not been much published on the performance of these systems, examples reporting encouraging results can be found in Bauer (1991) and Jain, Barner, and Osterfield (1989).

4 UNRESOLVED PROBLEMS IN SIMULATION-BASED SCHEDULING

In section 3.2, a number of unresolved problems were briefly raised concerning the implementation of the dynamic rescheduling framework illustrated in Figure 3. These issues are described in greater detail below:

- How to interface the simulation to other systems.
  The generation of good schedules requires that complete and valid input data be provided to the model in terms of both the current state of the shop floor (i.e. every in-process order, machine, person, and other limited resource) and the production plan over the time window to be simulated. Without fast and robust communication links between the scheduler and these other systems, the approach is unviable. If a comprehensive shop floor data collection system is not in place then manual data input would be required which would be impracticable for all but the smallest manufacturing systems.
- What triggers the need to reschedule?
  Should rescheduling be done on a regular basis (e.g. at the start of every shift or every day) or should it be triggered by the occurrence and detection of certain "events" (e.g. rush order arrival, unexpectedly low process yield for some operation, or machine breakdown). One possibility would be to specify certain tolerance ranges on the operation completion times in the existing schedule so that no rescheduling is carried out provided discrepancies between the schedule and its execution are within the tolerance range.
- Is a purely deterministic simulation valid?
  Typically (Grant 1988), deterministic simulation is used to evaluate the alternative actions due to the need to make a real-time decision and since the multiple replications required by a stochastic model would take much more time. There are questions as to the validity of such simulations for systems where there is known and quantifiable stochasticity. It might be desirable to represent only the critical sources of uncertainty as stochastic (McConnell and

Medeiros 1992) and approximate all others deterministically.

- How long should each simulation run?
  This issue is related to the preceding one. Time constraints will place an upper bound on the maximum length of simulation but even this may be too long in terms of model validity. For example, if there is a 90% chance that a rush order will arrive in the next 2 hours is it valid to deterministically simulate for a 4 hour time window?
- What schedule evaluation criteria should be used?
  There is a wide range of criteria that could be used to evaluate schedule performance and a number of questions concerning the evaluation process. Which of the performance measures should be supported by the system? Should they all be built in, automatically calculated, and presented to the user for each option or should the ones to calculate be configurable by the user? Should the system do some automatic filtering so that only certain ones are presented to the user?
- How is the generated schedule to be implemented?
  It might be desirable to use this type of system both for a largely manual manufacturing system, where the system will need to generate hard copies of individual machine schedules for operators, and for fully automated systems where an electronic representation of the schedule could be passed to the shop floor control computer(s).
- What range of candidate responses to support?
  In a typical manufacturing environment there is a daunting array of possible actions that might be taken to attempt to produce a better schedule. The problem this leads to is what range of these actions are to be supported by the simulation system? It may not be enough to simply support different dispatching rules since other issues, for example the order release decision (Melnyk 1988), may have greater control over schedule performance.
- What particular candidate responses to evaluate?
  This issue is closely related to the previous one. Irrespective of the breadth of possible options, only a limited number of options can actually be simulated and compared due to time constraints. The problem is in rapidly identifying the best candidates for a particular decision, either fully automatically or providing guidance to the human user. Rather than identifying a subset of possible actions it might be better to choose one and then incrementally modify it in some kind of neighbourhood search approach.
- What features should the user interface have?
  In systems intended to support a human scheduler there are a whole range of user interface issues to resolve. In particular, what guidance should be provided to the user in selecting options to evaluate and how can a powerful schedule editing facility (Kanet and Sridharan 1990) be used, allowing the user to force certain operations to take place at defined times and locations.

Although a clear distinction between academic and industrial issues cannot (and should not) be made, some of the above points seem to be of more concern to one or other of these two groups. Broadly speaking, the industrial group is concerned with the development of systems which are usable in tackling existing scheduling problems within manufacturing facilities. The concern is with easy-to-use techniques which can generate feasible and robust schedules and can fit into the exising control scheme without too much upheaval. The academic group, on the other hand, is concerned with fundamental knowledge and understanding of this problem domain, for example the investigation of systems which will allow new knowledge to be gained about dynamic rescheduling and the determination of quantitative results on the performance of alternative solution approaches. Clearly the concerns of both groups need to be addressed since a balance is required between developing working solutions and gaining a fuller understanding of why those solutions work (and hence whether better responses exist).

4.1 Academic/Research Issues

A number of the issues described above fall squarely in the academic domain. Clearly, working simulation-based scheduling systems exist but there is a lack of scientific knowledge on their range of applicability and on their performance compared to alternative approaches. Studies are needed here to expand our understanding concerning:

- the range of manufacturing environments where the approach is beneficial
- quantitative effects of variation of the simulation window
- quantitative effects of variation of run frequency
- quantitative comparisons of simulation and other approaches
- quantitative analysis of the performance of deterministic simulation
- the need for fully/partially stochastic simulations
- appropriate designs for hybrid simulation/A.I./O.R. systems
- structuring the simulation to allow interactive partial rescheduling

- the most appropriate set of actions to be evaluated.

To give more detail on some of the above, studies on the characteristics of manufacturing environments where the approach is beneficial are required. What combination of manufacturing characteristics (e.g. complexity, uncertainty, shop structure, and congestion) make the approach feasible. If there is little uncertainty then arguably algorithmic approaches should be capable of generating better schedules. If there is too much uncertainty then perhaps on-line simulation has nothing to offer and traditional off-line studies are best.

It should be noted that experimentation with this sort of approach raises a number of issues new to simulation researchers over and above the usual ones such as level of detail, verification and validation, and statistical concerns (Law and Kelton 1991). One of these is that new mechanisms need to be developed for model initialisation since the usual "empty and idle" approach is clearly unrealistic. A second is that to evaluate the performance of the approaches, the model must be linked to a second, more detailed simulation (or really "emulation") model (McConnell and Medeiros 1992) which can mimic the operation of the system to be scheduled, incorporating all the real sources of uncertainty.

4.2 Industrial/Application Issues

A number of researchers have highlighted the considerable gap between much academic research on production control and scheduling problems and industrial needs for working systems (Melnyk 1988). The main problem with academic work on simulation-based scheduling approaches is a failure to adequately consider the intended operation of these systems within manufacturing companies and the key role that the human user must play. Some of the concerns here have been well described (McKay, Buzacott, and Safayeni 1988), for example the development of models of the shop floor which are capable of representing it to the level of detail required. There is a vast range of constraints which is difficult to incorporate into computer models but which experienced human schedulers (with intimate knowledge of the the real system) can utilise to produce good schedules.

One factor here is the characteristics of the manufacturing system to be scheduled. If the system is fully automated (e.g. a flexible machining system with an automatic materials handling system) then it should be easier to develop a model which can describe the system in the necessary detail and which can be kept up-to-date, ready to carry out rescheduling as required. As the proportion of manual activity in the system increases so the probability that critical constraints are not represented or that certain changes in state are not made known to the model increases. In this case, it is important to have a rescheduling system which can make use of the expertise of a knowledgeable human scheduler so that the two can work in close concert, each complementing the other's capabilities.

For non-automated scheduling systems which must rely on a human user (and for non-automated manufacturing shop floors where people implement the generated schedule), how the system interfaces to all humans with which it must interact is critical. The problem of system acceptance (whereby people will refuse to use a system in which they have no confidence or which makes their tasks more onerous), which occurs in traditional (off-line) applications of simulation too, is one which must be dealt with here.

A further key issue for industrial installations is flexibility and modularity. Any system must be flexible to change since change is certain to occur in terms of updates to the scheduling problem database, and new links to other computing systems.

One final industrial issue that must be raised is that systems must fit within emerging production planning and control paradigms. There is some debate as to what the decomposition should be between off-line "planning" and on-line "control". An emerging popular approach is to use more sophisticated off-line techniques to constrain the uncertainty that the on-line systems must face, thereby simplifying the shop floor plan execution environment. Simulation-based scheduling systems will be needed that do more than simply allow experimentation with different dispatching rules (Melnyk 1988), for example allowing different order release strategies to be evaluated which pay close attention to the setting of due dates. Further, systems which can fit within the framework of emerging distributed shop floor control models (Bauer *et al.* 1991) are required.

5 CONCLUSIONS

Simulation-based real-time scheduling is clearly a technique that will continue to be implemented in manufacturing plants over the coming years. As for many other industrial problems, it might be argued that discrete-event simulation is the only tool capable of modelling in sufficient detail to carry out an analysis.

This paper has critically surveyed reported work in this area and highlighted some of the major problems which remain to be addressed. The following summarises the main areas where future development is needed and will take place:

- improved scientific understanding of the technique
- tools more targeted at supporting an expert human user in a highly uncertain environment
- hybrid scheduling approaches integrating knowledge-based, optimisation, and simulation-based elements
- tools fitting with emerging production planning and control paradigms.

In closing, a possible spectre hanging over the long term applicability of simulation-based real-time scheduling should be noted. Arguably, simulation is a tool of last resort for analysing complex, poorly understood systems. The following two trends may make simulation-based tools obsolete in tackling scheduling problems in the future:

i) Once dynamic scheduling is better understood, it may be possible to create hybrid state-driven A.I./O.R. tools that can search for a schedule generation algorithm and rapidly generate a provably good schedule. Provided the necessary knowledge or experience can be accumulated, it might be possible to move directly from data describing characteristics of the scheduling problem to the knowledge-based choice of a fast algorithm for schedule generation.

ii) The movement towards production planning and control systems with significantly simplified production activity control elements (and correspondingly more sophisticated planning elements) may make the schedule execution environment so simple that complex highly integrated simulation-based tools are not required.

ACKNOWLEDGEMENTS

The authors wish to thank the Natural Sciences and Engineering Research Council of Canada and the University of Calgary for their generous research support under grants OGP-012-1522 and URGC-911526 respectively.

REFERENCES

Bauer, A., R. Bowden, J. Browne, J. Duggan, and G. Lyons. 1991. *Shop Floor Control Systems: From Design to Implementation.* Chapman & Hall.

Bauer, M. 1991. Keeping commitments with CIM. *Manufacturing Systems* 9:18-23.

Ben-Arieh, D. 1988. A knowledge-based simulation and control system. In *Artificial Intelligence: Implications for Computer Integrated Manufacturing*, ed. A. Kusiak, 461-472. IFS (Publications) and Springer-Verlag.

Davis, W.J. and A.T. Jones. 1988. A real-time production scheduler for a stochastic manufacturing environment. *International Journal of Computer Integrated Manufacturing* 1:101-112.

French, S. 1982. *Sequencing and Scheduling: An Introduction to the Mathematics of the Job-Shop.* Ellis Horwood Limited.

Grant, F. Hank. 1988. Simulation in designing and scheduling manufacturing systems. In *Design and Analysis of Integrated Manufacturing Systems*, ed. W. Dale Compton, 134-147. Washington D.C.:National Academy Press.

Harmonosky, C.M. 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.

Harmonosky, C.M. and S.F. Robohn. 1991. Real-time scheduling in computer integrated manufacturing: a review of recent research. *International Journal of Computer Integrated Manufacturing* 4:331-340.

Jain, S., K. Barber, and D. Osterfield. 1989. Expert simulation for on-line scheduling. In *Proceedings of the 1989 Winter Simulation Conference*, ed. E.A. MacNair, K.J. Musselman, and P. Heidelberger, 930-935.

Kanet, J.J. and V. Sridharan. 1990. The electronic leitstand: a new tool for shop scheduling. *Manufacturing Review* 3:161-170.

Kusiak, A. 1990. *Intelligent Manufacturing Systems.* Prentice-Hall.

Law, A.M. and W.D. Kelton. 1991. *Simulation Modeling and Analysis, Second Edition.* McGraw-Hill.

McConnell, P.G. and D.J. Medeiros. 1992. Real-time simulation for decision support in continuous flow manufacturing systems. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J.J. Swain, D. Goldsman, R.C. Crain, and J.R. Wilson, 936-944.

McKay, K.N., J.A. Buzacott, and F.R. Safayeni. 1988. The scheduler's knowledge of uncertainty: the missing link. *Proceedings of the IFIP Working Conference on Knowledge Based Production Management Systems*, 223-241.

Melnyk, S.A. 1988. Production control: issues and challenges. In *Intelligent Manufacturing: Proceedings from the First International Conference on Expert Systems and the Leading Edge in Production Planning and Control*, ed. M.D. Oliff, 199-232. The Benjamin/Cummings Publishing Company, Inc..

Pinedo, M. 1992. Scheduling. In *The Handbook of Industrial Engineering, 2nd Edition*, ed. G. Salvendy, 2131-2153.

Pritsker, A.A.B., F. Hank Grant, and S.D. Duket. 1990. Simulation and real-time factory control. In *Papers*

Experiences Perspectives, ed. A.A.B. Pritsker, 472-483. Systems Publishing Corporation.

Rogers, P. and M.T. Flanagan. 1991. On-line simulation for the real-time scheduling of manufacturing systems. *Industrial Engineering* 23:37-40.

Wichmann, K.E. 1990. Simulation-based production schedule generation. *Production Planning and Control* 1:179-189.

Wu, S.-Y.D. and R.A. Wysk. 1988. Multi-pass expert system control system - a control/scheduling structure for flexible manufacturing cells. *Journal of Manufacturing Systems* 7:107-120.

Wu, S.-Y.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* 27:1603-1623.

AUTHOR BIOGRAPHIES

PAUL ROGERS is an Associate Professor in the Department of Mechanical Engineering (Division of Manufacturing Engineering) at the University of Calgary. His research interests include discrete-event simulation, production planning and control systems, object-oriented modelling for intelligent manufacturing, and models for analysis of manufacturing systems. He is a member of IIE, SCS, and TIMS and serves on the Editorial Board of the International Journal of Computer Integrated Manufacturing. He holds Ph.D. and M.Eng. degrees from Cambridge University in England.

ROSALEE J. GORDON is a graduate student in the Department of Mechanical Engineering (Division of Manufacturing Engineering) at the University of Calgary, working towards her M.Sc. degree. She also has a B.Comm. degree in Operations Management from the University of Calgary. Her research interests are in the areas of simulation and scheduling, and she is a member of IIE, SME, and APICS.