Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

LOOK-AHEAD STRATEGIES FOR CONTROLLING BATCH OPERATIONS IN INDUSTRY – AN OVERVIEW

Durk-Jouke van der Zee

Production Systems Design Group
Faculty of Management and Organization
University of Groningen
P.O. Box 800
9700 AV, Groningen, THE NETHERLANDS

ABSTRACT

Batching jobs in a manufacturing system is a very common policy in most industries. Main reasons for batching are avoidance of set ups and/or facilitation of material handling. Examples of batch-wise production systems are ovens found in aircraft industry and in semiconductor manufacturing. Starting from the early nineties much research efforts have been put in constructing strategies for the dynamic control of these systems in order to reduce cycle times. Typically, these so-called "look-ahead strategies" base their scheduling decision on the information on a few near future product arrivals. In this paper we give a literature overview of the developed strategies, consider basic insights in their construction and highlight issues for further research.

1 INTRODUCTION

In many manufacturing and transportation systems batch servers are used for efficient processing. Main reasons for batching, i.e., the grouping of a number of jobs which may be processed simultaneously, are the avoidance of setups or the facilitation of material handling. Examples of batch servers are ovens used for hardening of synthetic aircraft parts (Hodes et al. 1992), the diffusion or oxidation tubes in semiconductor wafer fabrication and the burn-in ovens in semiconductor testing (Fowler et al. 1992, 2000, Uzsoy et al. 1992, 1994). For both types of industry addressed competition is severe and management attention is focused on shortening lead times from the perspective of both cost reduction and customer service.

In this paper we consider a particular model of a batch-processing machine motivated by the oven systems found in aircraft industry and semiconductor manufacturing, see (Fowler et al. 1992, Glassey and Weng 1991, Hodes et al. 1992, Weng and Leachman 1993). The oven processes share similarities with respect to the need for specific settings for e.g. temperatures, pressures and service times, which relate to different products. Consequently, different types of product cannot be batched together. Batch sizes are restricted by e.g. physical sizes of the oven and products, or process constraints. Service times are considered to be constant, depending on product and/or oven characteristics. Service may not be interrupted, i.e., jobs may not be preempted, because this would make products worthless for any further use. This is due to strict quality constraints.

In this article we survey strategies for on-line scheduling of these batch servers. Typically, these rules should be computationally efficient, and responsive. The need for a responsive strategy follows from shop dynamics that restricts planning information to queue lengths and forecast data on a few near-future arrivals. Therefore, there is a need for continuous updating of the schedule. Main objectives of our survey of these so-called look-ahead strategies are to give *an overview of their current fields of application* in terms of shop configurations and to consider *basic insights in rule construction.* Given these observations we highlight *issues for future research.*

The remainder of the paper is organized as follows: in the next section we review literature. We relate look-ahead strategies to alternative types of rules developed for controlling batch processes. Next existing look-ahead strategies will be classified according to basic shop characteristics. In this way we get an insight in their field of application. In Section 3 we consider a decision framework for describing look-ahead strategies. It supplies us with a general format for describing their construction. In Section 4 we use this framework to obtain basic insights in rule construction as they follow from the strategies developed so far. Finally, conclusions and directions for future research are summarized in Section 5.

2 LITERATURE REVIEW

In this section we relate look-ahead strategies to alternative types of strategies for planning batch operations. Next we classify look-ahead strategies according to the shop configurations addressed.

Control strategies for scheduling batch processes may be classified according to the amount and quality of information that is known on future arrivals. In queueing theory threshold strategies are studied which relate the decision to schedule a batch to a certain minimum queue length. A well-known example of such a strategy is the Minimum Batch Size rule (MBS), which was introduced by Neuts (1967). According to this strategy a batch starts service as soon as at least a certain fixed number of customers is present.

While the above strategies base their decision on local information only, full knowledge of future arrivals is assumed to be known in the field of deterministic machine scheduling. Much of the research in this field is related to oven systems found in semiconductor manufacturing. Uzsoy et al. (1992, 1994) summarize scheduling models for this industry. Other surveys are supplied by Webster and Baker (1995) and Potts and Kovalyov (2000).

In this article we will focus on so-called look-ahead strategies. These strategies have been developed based on the observation that in many practical cases the assumptions underlying deterministic machine scheduling are not met. In those cases the amount and quality of data on future arrivals does not allow for a deterministic approach, cf. (Duenyas and Neale 1997, Fowler et al. 1992, Glassey and Weng 1991). Typically, look-ahead strategies assume that only a limited number of near future arrivals are known and/or predicted. In Table 1 we give an overview of the look-ahead strategies developed so far.

Glassey and Weng (1991) were among the first to introduce look-ahead strategies for (semi-conductor) batch processing systems. They discuss the practical usability of a dynamic programming approach to find a sequence of loading times of given lots, in such a way that total delay is

Table 1: Overview of Developed Look-Ahead Strategies

Rule	*Machines*		*Products*		*Forecast Data*	*Criterion*
	Number	*Char.*	*Number*	*Char.*		
MBS	≥1	I	1	I	No	F,C
DBH	≥1*	*	1	*	No	F
NACH	≥1	I	≥1	NI	Yes	F
MCR	1	I	≥1	NI	No	F,C**
RHCR	1	I	≥1	NI	Yes	F,C**
HA	1	I	≥1	NI	Yes	F
DJAH	≥1	I	≥1	NI	Yes	F,C
DSH	≥1	NI	≥1	NI	Yes	F,C
RHCR-S	1***	I	≥1	NI	Yes	F
DJAH-F	1***	I	≥1	NI	Yes	F

MBS = Minimum Batch Size rule (Neuts 1967)
DBH = Dynamic Batching Heuristic (Glassey and Weng 1991, Glassey et al. 1993)
NACH = Next Arrival Control Heuristic (Fowler et al. 1992, 2000)
MCR = Minimum Cost Rate heuristic (Weng and Leachman 1993)
RHCR = Rolling Horizon Cost Rate heuristic (Robinson et al. 1995)
DJAH = Dynamic Job Assignment Heuristic (van der Zee et al. 1997)
HA = Heuristic with next Arrival information (Duenyas and Neale 1997)
DSH = Dynamic Scheduling Heuristic (van der Zee et al. 2001)
RHCR-S = Rolling Horizon Cost Rate heuristic for Batch-Serial system (Robinson et al. 1995)
DJAH-F = Dynamic Job Assignment Heuristic for Flow shops (van der Zee 2002)
I = Identical machine(product) characteristics only (service time, batch size)
NI = Non-identical machine(product) characteristics allowed (service time, batch size)
F = Average flow time
C = Logistic costs
* = No explicit formulation available in literature
** = See van der Zee et al. 1997
*** = Network of batch machine and serial machine

minimized. They argue that this approach fails for reasons of computational feasibility, and availability and quality of data on future arrivals. Therefore they present a Dynamic Batching Heuristic (DBH). This heuristic decides when to start the next machine cycle thereby aiming for a minimal average flow time. The planning horizon in DBH is just one service time. DBH proves to perform better than MBS, based upon the knowledge of just a few arrivals. Starting from the single product single machine shop discussed by Glassey and Weng other authors proposed new look-ahead strategies in order to deal with several extensions. The first extension of the DBH rule concerned the multiple products case, which was considered by Fowler, Hogg and Phillips (1992). Differences between products concern the required service time and/or maximum allowed batch size. Their Next Arrival Control Heuristic (NACH) proves to be a robust heuristic in case forecast data on future arrivals are used, i.e., estimated arrival moments for new lots. Weng and Leachman (1993) show how performance can be improved for the multiple product single machine case by their Minimum Cost Rate heuristic (MCR), which shows an analogy with the Silver and Meal heuristic (1973). However, a disadvantage of MCR is the relatively large amount of data needed to realize the improvement in system performance. Also robustness of the heuristic with respect to forecast data is weaker than for NACH. For that reason Robinson, Fowler and Bard (1995) propose a slightly altered and more robust version of the MCR heuristic, named Rolling Horizon Cost Rate heuristic (RHCR). Finally, Duenyas and Neale (1997) have obtained structural results and developed an effective heuristic (HA).

The parallel machine case is addressed by Fowler et al. (2000) and van der Zee et al. (2001). Fowler, Hogg and Phillips (2000) show how their NACH heuristic may be extended to the multiple machine case. van der Zee et al. introduce the Dynamic Job Assignment Heuristic (DJAH). The criterion for optimization for DJAH is the minimization of logistic costs per part (customer) on the long term. Logistic costs associated with a job consist of linear waiting costs and a fixed amount of setup costs (e.g. energy costs). The definition of this cost function also covers an important special case: if setup costs are zero, minimization of logistic costs comes down to minimization of average flow time, cf. (Fowler et al. 1992). While both DJAH and NACH address the control of identical machines, the DSH heuristic (van der Zee et al. 2001) is intended to assist the planner in scheduling non-identical machines.

Network configurations are addressed by Robinson et al. and van der Zee (2002). Both they consider a Batch-Serial system. Here the production system consists of a batch machine followed by a serial machine, which processes piece-wise.

Above we related existing look-ahead strategies to their assumed fields of application using general shop characteristics, i.e., available machines, product types, availability of data on new jobs, and the criterion for optimization adopted. In Section 5 we come back to this overview as we highlight alternative shop configurations and systems as candidates for further research in this field.

3 SYSTEM DESCRIPTION

In this section the batch shop under study is described in detail. Given this description a general framework for decision-making in batch shops is introduced. Next, in Section 4, the framework is used to support our discussion on the construction of look-ahead strategies.

3.1 Batch Shop

We describe the shop using figure 1 as a starting point. Next to a controller, the shop consists of a batch server and a buffer. Buffers are used to store lots that queue at a production stage. We assume buffers to have an unlimited storage capacity and lot sizes to equal one product. Multiple types (j) of products are considered, with $j \in J = \{1,2,..,N\}$. Each product type sets different requirements to processing conditions for the batch machine, like e.g. pressure and temperature. Consequently, batches have to be made up of the same type of products. Next to processing conditions, essential type differences concern service times for the batch machine (T_j), and the maximum batch sizes allowed for the batch machine (C_j). The latter characteristic may e.g. be related to volume restrictions. Service times per product type are fixed and include setup and transport times. Hence, setup activities are sequence independent. Also, service times for the batch machine (T_j) are independent of batch size.

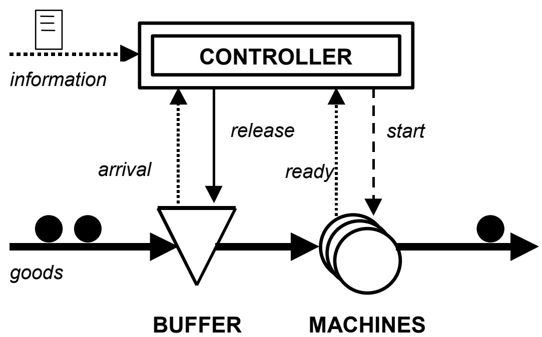

Figure 1: Batch Shop

3.2 A Framework for On-Line Decision Making in Batch Shops

The above description of the shop floor sets the context for the decision problem. Let us now consider this problem in some more detail. Therefore we use the general framework for decision-making in batch shops developed by van der Zee at al. (2001). They used the framework for defining look-ahead strategies that address shops consisting of multiple parallel batch machines (see Section 2). We will now relate the framework (see Figure 2) to the description we gave of the batch shop. We will do so by characterizing control strategies, i.e., scheduling routines, in terms of their triggers, information availability and usage, and decision structure.

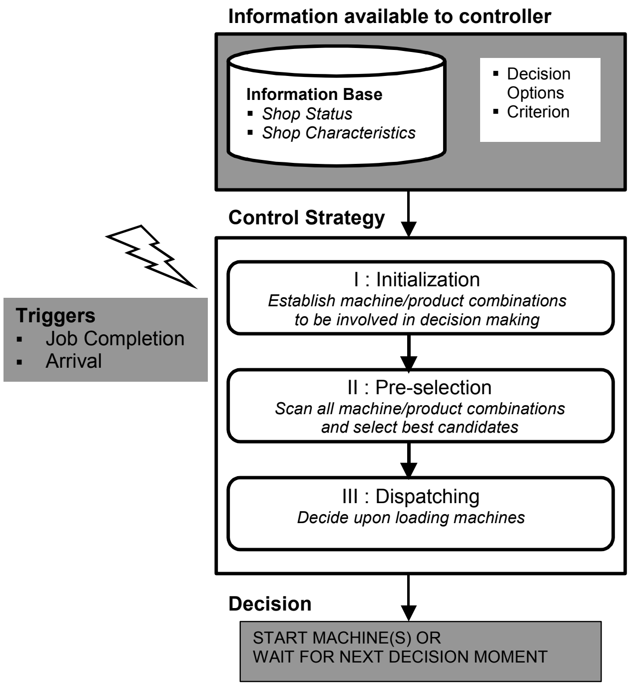

Figure 2: A Framework for Decision-Making

3.2.1 Triggers

Three types of events govern shop dynamics: product arrivals, job completion and information on future product arrivals. Each of these types of events may trigger the controller. As such these events correspond to *decision moments*. Obviously, new operations are only released if both machine and products are available.

3.2.2 Information

Next to the static shop characteristics, as discussed in Subsection 3.1, the planner has the following information at his disposal at the decision moment (t_0):

- Local information on shop status at t_0:
  - Queue lengths for each product j (q_j), with $j \in J = \{1,2,..,N\}$.
  - The moments t' the batch machine is available (again).
- Information on future arrivals:
  - For each product j the present and successive future arrivals $t_{k,j}$ ordered through the index k according to the moment of arrival, up to some specified look-ahead horizon LH.

Note that there is no general agreement in the literature with respect to the definition of the look-ahead horizon (see Section 4). However, in all cases it is assumed that the amount of look-ahead information available only allows for scheduling the next machine cycle.

3.2.3 Decision Structure - Decision Options, Criterion for Optimization, and Decision Procedure

The task of the planner boils down to scheduling the batch machine, i.e., making a decision on batch contents and scheduling moment for the next machine cycle. Allowed decision options are to:

- Release the job characterized by batch contents and scheduling moment to the shop floor.
- Postpone decision making to a later decision moment.

The criterion for optimization specifies the long term goals for controlling the batch shop. Given the length of the look-ahead horizon the rules will typically apply a *reduced criterion*, where the optimization is related to system performance for the next machine cycle.

The procedure for deciding among the alternative decision options is subdivided in three sequential steps:

1. *Initialization* is meant to establish the set of machine/product combinations, which is to be involved in decision making. Typically, the set of all possible combinations is reduced on a basis of general exclusion principles. Exclusion, because it is a priori clear, that certain norms cannot be met. Hence, even in the best case (potentially), it is induced from the information base that some prespecified upper or lower bound for some criterion will be violated. The principles may follow e.g. basic insights with respect to problem structure or from company standards. Clearly, benefits of the initialization step lie in reduction of the problem in terms of candidate solutions.
2. In the second step, *pre-selection*, the aim is to (further) reduce the combinatorial problem by selecting the most promising machine/product combinations.
3. In this final step no further selection of product machine combinations is foreseen. This implies that product machine combinations should be unique with respect to the machine chosen. *Dispatching* concerns the question whether machines available at the decision moment should be loaded right now, or whether it is better to wait for a later moment. Such moments typically correspond to future product arrivals. The trade-off involves a comparison of logistic costs for both possibilities for each of the selected machine/product combinations.

In the next section we will use the framework to discuss construction of existing look-ahead strategies.

4 RULE CONSTRUCTION

In the previous section we considered a general format for describing the construction of look-ahead strategies. In this section we will apply this framework to describe alternative choices made in the rule construction. Building on this description we highlight basic insights obtained in rule construction so far. Please note that references for the rules mentioned can be found in Section 2 (Table 1).

4.1 Triggers

In literature only product arrivals and the completion of jobs are considered as triggers for activating a control strategy. The third type of trigger – the receipt of new information on future arrivals - is not considered in literature. While this is true, its inclusion in existing rules would not require significant modifications in their construction. The question remains, however, what the effects of its inclusion would be for system performance.

4.2 Information

While all look-ahead strategies assume full information on queue lengths and machine availability, differences exist with respect to the length of the look-ahead horizon. Some rules relate the length of the horizon to a fixed number of arrivals. For example the NACH rule only considers the first arrival. Alternatively, DBH considers all arrivals within a fixed period, which is set equal to one processing time. Also superposition of these types is possible - the horizon for RHCR covers k arrivals plus the arrivals that take place during processing if the machine would be loaded at the time of the k-th arrival. Simulation experiments indicate that the marginal yield of more information (arrivals) in terms of performance is (strongly) decreasing, cf. Glassey and Weng (1991), Duenyas and Neale (1997). Also the larger the information horizon is set the higher the data collection costs. In this sense there is a trade-off determining an optimum for the information horizon.

4.3 Decision Options

The first look-ahead strategies developed (DBH, MCR) allowed for job release at the time of future arrivals. Research by Fowler et al. (1992) made clear that control strategies adopting such a policy prove to be less robust in case of forecast data. Clearly, delaying the decision is better under uncertainty since there is never an advantage to making it earlier (and possibly a disadvantage). Therefore more recent rules modified the first decision option mentioned in Subsection 3.2 by restricting the loading of products in the batch machine to the decision moment (t_0).

4.4 Criterion

So far, two criterions for optimization have been considered, cf. van der Zee et al. (1997):

- The minimization of average flow time per part in the long run.
- The minimization of average cost price per part in the long run.

Note that for the shop configuration as described above the minimization of average flow time implies minimization of average waiting time. This logically follows from the fact that service times are assumed to be fixed. Costs considered are fixed set up costs and linear waiting costs. Typically look-ahead strategies try to reach these goals by adopting a reduced criterion for optimization – costs are to be minimized for the next machine cycle.

4.5 Decision Procedure

In this subsection we will consider the decision procedure.

4.5.1 Initialization

So far exclusion principles only have been used for the case of parallel identical machines. Both NACH and DJAH reduce the set of machines to be scheduled to those available at the decision moment. This reduction is based on the idea that in suchlike situation machines that are currently not available will never be better candidates for loading.

4.5.2 Pre-Selection

Two perspectives are taken in literature in defining alternative rules for pre-selection:

- A short term perspective that relates shop status and system performance.
- A long term perspective that relates machine capacity and system performance.

4.5.2.1 Short-Term Perspective

As long as the batch shop consists of identical machines only, the need for pre-selection is not great, as computation times are within reasonable bounds. It should be remarked, however, that these observations are related to simulation studies of small batch shops in terms of products and number of machines. Consequently, there may be a need for additional rules in case of larger shops.

The consideration of alternative machine types leads to a substantial increase in the complexity of the combinatorial problem faced by the planner. van der Zee et al. reduce the problem by using a throughput related rule. As a definition of throughput is used:

$$TH_{m,j} = \frac{\min(q_j, C_j)}{t_m^{'} + T_{m,j} - t_0}\,. \qquad (1)$$

Note how throughput is influenced by queue length (q_j), machine availability (t_m') and capacity ($C_{m,j}$) and service times ($T_{m,j}$). Using this formula as basis a schedule of unique machine product combinations is being built. For a full description of the rule see (van der Zee et al. 2001).

4.5.2.2 Long-Term Perspective

The above perspective focuses at the best candidate product machine combinations within the planning horizon. The system description makes clear that this horizon typically is short being bounded by the look-ahead horizon. Consequently, long-term effects of a specific schedule on system performance may be neglected. For batch shops this is especially clear if scheduling decisions have to be made in the presence of a full loads, i.e., the number of a specific type products in queue meets or exceeds machine capacity. In such cases the serving of a partial load of an alternative type of products may have severe impact on long-term performance. It may even lead to an unstable system (Duenyas and Neale 1997) due to "capacity loss". To avoid these effects on system performance a priority rule was proposed by Fowler et al. (1992) According to this rule product machine combinations for which a full load is available at the decision moment are preferred over alternative combinations. The beneficial effects of this rule are confirmed by simulation experiments (van der Zee et al. 2001).

The pre-selection step concludes with a reduction of candidate product machine combinations to those for which a machine is available at the decision moment. In this way it is allowed for a rolling horizon approach in scheduling batch jobs.

4.5.3 Dispatching

The trade-off involves a comparison of logistic costs for both possibilities for each of the selected machine/product combinations. In Figure 3 this trade off is displayed for the single product single machine case. It shows how starting the batch machine at t_0 causes waiting costs for the items that arrive during processing (T). On the other hand, if the batch machine were to be loaded at t_1, next to these costs also the waiting costs for items in queue (q) at the decision moment (t_0) should be taken into account. Next to waiting costs the trade-off may involve other costs, like e.g. set up costs, see van der Zee et al. (1997). Of course dispatching decisions only make sense in case of non-full loads.

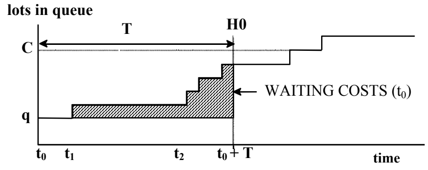

(a)

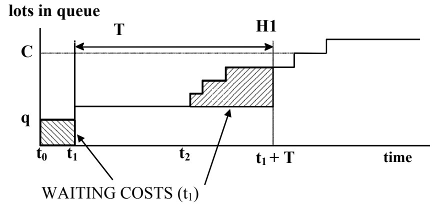

(b)

Figure 3 a, b: Trading-Off Waiting Costs in Case of Delaying the Loading Decision.

Differences between the look-ahead strategies developed concentrate on two issues:

- The number of arrivals to be considered in the trade-off.
- The choice of a weight factor in comparing costs associated with alternative schedules.

The number of arrivals included in the trade-off varies significantly among the rules. While NACH only considers the next arrival as a possible candidate for postponing the decision, MCR and RHCR consider arrivals up to the moment they make up a full load. Note how the choice of alternative scheduling moments corresponds with the assumptions with respect to the length of the look-ahead horizon.

The choice to include more information on future arrivals in decision-making corresponds with the addition of a weight factor. Where NACH and DBH do not consider a weight factor in comparing costs for alternative decision options, the more information intensive strategies DJAH, DSH, MCR and RHCR do. The latter two strategies weight waiting costs for the associated planning horizon, i.e., the period the machine is not available for any alternative use. For example, following Figure 3, if the machine would be loaded at t_1 this period would be [t_0, t_1 + T]. Alternatively, DJAH and DSH, adopt batch size as a weight factor.

Recently, look-ahead strategies have been developed that address network configurations Robinson et al. (1995) and van der Zee (2002) focus at a batch-serial system, i.e., a batch machine followed by a discrete processor. On the other hand Neale and Duenyas (2000) focus at a serial-batch system, where the discrete processor precedes the batch machine. Basically, these strategies are extensions of existing strategies. New elements are in the computation of waiting costs and the choice of a weight factor. The computation of waiting costs includes waiting at the serial stage (batch-serial system). Further, it was found that the planning horizon is not very well suited as a weight factor in network configurations.

5 CONCLUSIONS AND RESEARCH AGENDA

In the preceeding sections we surveyed existing look-ahead strategies for the control of batch processes. We studied their assumed field of application and the progress with respect to rule construction. In this section we will summarize main conclusions on both subjects and relate them to suggestions for future research.

5.1 Field of Application

- At this moment the assumed field of application for look-ahead strategies is rather limited. The focus is mainly towards oven systems in semiconductor manufacturing. Similar systems in other industries are hardly studied.
- Within the context of semiconductor manufacturing only specific configurations are studied. Alternative configurations that assume the availability of compatible product families (cf. Duenyas and Neale 1997), compound arrivals, multiple processing steps (Robinson et al. 1995) or re-entry flows (Glassey et al. 1993) have not or hardly been studied yet.
- Almost all look-ahead strategies adopt a flow time criterion. Practice may require the use of alternative performance criteria based on due date settings or possibilities to prioritize the processing of certain products (for example because they are needed urgently elsewhere).

5.2 Rule Construction

- Significant progress has been made with regard to rule construction in the past years. It would be worthwhile to support this progress by mare structural analysis like that of e.g. Duenyas and Neale (1997). They relate rule construction to queueing theory. In this way they do not only try to improve rule construction but integrate it with insights from other fields.
- There are no uniform assumptions with respect to the availability of information on future arrivals underlying existing look-ahead strategies. An interesting question is this respect is whether the application of a specific rule should be related to the amount of data available on future arrivals.
- Activation of look-ahead strategies is related to shop status and not to the receipt of data on future arrivals. It would be interesting to know how performance would be influenced if it would be included as a trigger.

Next to the above suggestions for further research it is also important to direct more efforts to applied research. In this way the practical validity of several extensions can be tested and benefits of the new rules may be exploited to a greater extent.

REFERENCES

Duenyas, I. and J.J. Neale. 1997. Stochastic scheduling of a batch processing machine with incompatible job families. *Annals of the Operations Research* 70:191-220.

Fowler, J.W., Hogg, G.L. and D.T. Phillips. 1992. Control of multiproduct bulk service diffusion/oxidation processes. *IIE Transactions* 24:84-96.

Fowler, J.W., Hogg, G.L and D.T. Phillips. 2000. Control of multiproduct bulk service diffusion/oxidation processes. Part 2: Multiple servers. *IIE Transactions* 32:167-176.

Glassey, C.R. and W.W. Weng,. 1991. Dynamic batching heuristic for simultaneous processing. *IEEE Transactions on Semiconductor Manufacturing* 4:77-82.

Glassey, C.R., Markgraf, F. and H. Fromm. 1993. Real time scheduling of batch operations. In *Optimization in Industry: mathematical programming and modeling techniques in practice*, ed. T.A. Ciriani and R.C. Leachman, 113-137. Chichester: Wiley.

Hodes, B., Schoonhoven, B. and R. Swart. 1992. On line planning van ovens. Technical Report, University of Twente, The Netherlands (in Dutch).

Neale, J.J. and I. Duenyas. 2000. Control of manufacturing networks which contain a batch processing machine. *IIE Transactions* 32:1027-1041.

Neuts, M.F. 1967. A general class of bulk queues with Poisson input. *Annals of Mathematical Statistics*, 38:759–770.

Potts, C.N. and M.Y. Kovalyov. 2000. Scheduling with batching: a review, *European Journal of Operational Research* 120:228-249.

Robinson, J.K., Fowler, J.W. and J.F. Bard. 1995. The use of upstream and downstream information in scheduling semiconductor batch operations. *International Journal of Production Research* 33:1849-1869.

Silver, E.A. and H.C. Meal. 1973. A heuristic for selecting lot size quantities for the case of a deterministic time varying demand rate and discrete opportunities for replenishment. *Production and Inventory Management* 14, 64-74.

Uzsoy, R., Lee, C.Y. and L.A. Martin-Vega. 1992. A review of production planning and scheduling models in the semiconductor industry, Part I: System characteristics, performance evaluation and production planning. *IIE Transactions* 24:47-60

Uzsoy, R., Lee, C.Y. and L.A. Martin-Vega. 1994. A review of production planning and scheduling models in the semiconductor industry, Part II: Shop-floor control. *IIE Transactions* 26:44-55.

Webster, S. and K.R. Baker. 1995. Scheduling groups of jobs on a single machine. *Operations Research* 43:692-703.

Weng, W.W. and R.C. Leachman,. 1993. An improved methodology for real-time production decisions at batch-process work stations. *IEEE Transactions on Semiconductor Manufacturing* 6:219-225.

Zee, D.J. van der, Harten, A. van and P.C. Schuur. 1997. Dynamic job assignment heuristics for multi-server batch operations - A Cost-Based Approach. *International Journal of Production Research* 35:3063-3093.

Zee, D.J. van der, Harten, A. van and P.C. Schuur. 2001. On-line scheduling of multi-server batch operations. *IIE Transactions* 33:569-586.

Zee, D.J. van der. 2002. Adaptive scheduling of batch servers in flow shops. *International Journal of Production Research* 40: 2811-2833.

AUTHOR BIOGRAPHY

DURK-JOUKE VAN DER ZEE is Assistant Professor of Production Systems Design at the Faculty of Management and Organization, University of Groningen, The Netherlands. Dr. van der Zee received his M.Sc. and Ph.D. in Industrial Engineering at the University of Twente, The Netherlands. His research interests include simulation methodology and applications, shop floor control systems and decision making related to the introduction and use of flexible manufacturing systems. He can be contacted by e-mail at <d.j.van.der.zee@bdk.rug.nl>.