

SET MANAGEMENT: MINIMIZING SYNCHRONIZATION DELAYS OF PREFABRICATED PARTS BEFORE ASSEMBLY

Manfred Mittler
Michael Purm

University of Würzburg
Institute of Computer Science
Am Hubland, 97074 Würzburg, GERMANY

Ottmar Gühr

IBM Germany
German Manufacturing Technology Center
Max-Eyth-Str. 6, 71065 Sindelfingen, GERMANY

ABSTRACT

This paper addresses the problem of synchronizing prefabricated parts for assembly into modules. We examine production environments arising out of semiconductor manufacturing and multi layer ceramics where several prefabricated parts, so-called sets, are necessary to complete a module. Frequently, some of these components are manufactured by the same production system.

In order to minimize the inventory of the entire production system and simultaneously to achieve short cycle times, we propose sophisticated sequencing rules (named *set management policies*) to control the production of components. Based on global status information on the inventory at each service station these policies determine the component which is to be processed next.

We evaluate the performance of the set management policies proposed in this paper by means of simulation results for several production environments which feature rework and re-entrant flow. In terms of several performance measures we show that our sequencing rules achieve better results than standard rules such as FCFS or EDD. The mean of the performance measures and also their variances are considered.

1 INTRODUCTION

Commonly, prefabricated microelectronic parts of different type from either semiconductor or multi layer ceramics (MLC) manufacturing lines are assembled into modules. These modules which may consist of several prefabricated parts are integrated into final products later on in the manufacturing process. Often some of the prefabricated parts are manufactured by the same production line which is, in general, fed with production orders for complete modules. Ex-

amples for this manufacturing process can be found, among many others, in mainframe computer fabrication. Here, at the first integration level semiconductor chips are mounted on MLC parts. These ceramic parts are identical except for the internal wiring and can therefore be produced by the same production line. At the second integration level, a number of ceramic parts are assembled to a board. Nevertheless, a complete set of MLC parts is needed to form a board.

As companies are aiming at lean production, cf. Womack, Jones, and Roos (1990), the inventory levels prior to assembly stages have to be minimal. Therefore, from an operations manager's point of view, it is desirable to deliver complete sets to assembly machines as far as possible. Otherwise, parts to be assembled to a board have to wait for each other. Since the production process itself is highly complex and prone to malfunctions, manufactured parts have to be reworked or might even be scrapped. Thus, replenishment of incomplete sets waiting at the end of the line takes a very long time (between 40 and 60 days) since needed parts have to go through the whole line from the very beginning. In addition, not all part types might encounter the same problems in manufacturing. For example the structural complexity may vary considerably resulting in lower or higher probability for a prefabricated product to be defect. For brief introductions to semiconductor and MLC manufacturing we refer to Gise and Blanchard (1986), Hogg, Fowler, and Ibrahim (1991), and Ahmed and Tummala (1992), respectively.

By these considerations it is obvious that waiting times at assembly stations occur. Our approach minimizes these waiting times. This leads to a reduced inventory at the assembly station such that the goal of lean production (minimum inventory level) is achieved. We consider both the *mean* waiting times and the *variance* of waiting times, because on-time-

delivery of goods is required nowadays. However, the ability to meet due dates increases with less variable waiting and service times. Several publications address the problem of minimizing the deviation of completion times from a common due date (see for example De, Gosh, and Wells (1992), Hall, Kubiak, and Sethi (1991), and De, Gosh, and Wells (1989)). Since by Little's law the mean inventory is proportional to the mean cycle time we also keep an eye on the cycle time of individual parts. Further, the variability of cycle times also has a direct impact on the inventory at the assembly station. Therefore, the variance of cycle times has also to be taken into account.

In this paper, we develop and analyze several queuing disciplines which try to control the production of sets of parts in unreliable manufacturing lines which feature rework and re-entrant flow. All parts of a module are manufactured by the same production line. We refer to these queuing disciplines as *set management policies* since they try to balance the inventory of every service station and the assembly station in a manner that each time a part is to be selected for service an integer number of sets can be formed. Although there is a large number of publications on sequencing rules and queuing disciplines (see for example Johri (1992), Lee (1992), or Uzsoy, Martin-Vega, Lee, and Leonard (1991)), to the authors' knowledge there is currently no publication considering set management and regarding the problem of synchronizing parts before assembly.

We investigate the performance of these set management policies by means of simulation for several basic manufacturing systems and for a complex semiconductor manufacturing process. In order to justify the application of our algorithms in real production environments we compare our strategies to common queuing disciplines such as FCFS (First-Come First-Served) and EDD (Earliest Due Date). For evaluation purposes the so-called *set synchronization times* are taken into account, as well as the cycle time of manufactured parts and the variance thereof.

2 FABRICATION ENVIRONMENT AND SET SYNCHRONIZATION TIMES

Prior to the definition of set synchronization times, we describe the production environment to be examined. As we will see in the following, for set management purposes this particular environment is the basic model incorporating rework and re-entrant flow.

The manufacturing system comprises N service stations, each of which consists of a number of servers and an infinite queue. The mean service time at station i is μ_i^{-1} . The final product to be manufactured

is made up out of several components (*prefabricated parts*), where each component is of one of M different types. All part types are manufactured by the same production line. Nevertheless, the process flows of different part types may be different from each other. The process flow of part type j has w_j processing steps (operations) which have to be performed on the N service stations. Therefore a job of type j is worked on at stations $s_{j,1}$ through s_{j,w_j} . A group of prefabricated parts which form a complete final product is called a *set*. A set contains exactly g_j parts of type j ($j = 1, \dots, M$) such that a final product comprises $G = \sum_{j=1}^M g_j$ components. We assume that *entire* sets arrive at the system with a mean interarrival time of λ^{-1} . If the process flows of different parts are identical components belonging to the same set arrive at the same station in random order to ensure that all part types have identical cycle times. Otherwise the first (last) part would have the smallest (longest) cycle time. This effect might not be desirable from an operations manager's point of view.

Individual parts start their manufacturing process at the first station indicated by the process flow of their type. Immediately prior to departure from some specific stations parts are subject to inspection. If a part fails to pass the quality check it is sent back to an earlier processing step in its process flow for rework purposes. After having gone through all processing steps prefabricated parts arrive at a *storage* and wait to be assembled to final products. Arrivals to the storage are used on a FCFS basis to complete entire sets. We assume that the assembly time is negligible such that sets leave the storage as final products immediately at the instant of arrival of the last component needed to complete a set. Depending on the state of the storage an arriving part may be used for further completion of a final product or may alternatively initiate the assembly of a new final product. There is no individualization of parts at the assembly. Therefore, all prefabricated parts of a particular type are interchangeable in the storage.

The queuing disciplines proposed in this paper are based on the concept of balanced service stations. We refer to a service station as a *balanced* station if the amount of parts currently present at that service center (waiting in the queue or receiving service) is exactly the amount of prefabricated parts needed to form a certain number of final products. If counter $x_{i,j}$ denotes the number of parts of type j currently residing at station i then station i is balanced if the equation $x_{i,j} = c \cdot g_j$ holds for $j = 1, \dots, M$, with an arbitrary constant c . Note, that the number of parts of type j currently held in the storage is denoted by $x_{N+1,j}$.

Taking into account the production environment described above we define the *set synchronization time* S as the time which elapses between the arrival of the first component of a set at the storage and the departure of the corresponding final product. Besides set synchronization times we look at cycle times of different part types. We define the *cycle time* T of an individual component as the time elapsing from the instant of arrival at the manufacturing system until the instant of arrival at the assembly station.

3 QUEUEING DISCIPLINES FOR SET MANAGEMENT

We propose two advanced queuing disciplines for the purpose of reducing set synchronization times. These disciplines will rely on globally available status information for either each *service station* or for each *processing step*.

3.1 Most Requested Part Type per Service Station

According to the definition of a balanced service station from Sec. 2 *Most Requested part type per Service Station* (MRSS) is a queuing discipline which tries to balance downstream stations by selecting the appropriate part type to be processed next. Concerning the process flow of a specific part type, station k has a location *downstream* of station i if station i is encountered prior to station k during the manufacturing process of this part type. Note, that in the case of re-entrant flow a station can be located down- or upstream of itself. For every part type j visiting station i the counter $z_{i,j}$ indicates the number of parts currently needed at this station in order to supplement the inventory of this station in a way that a certain number of complete sets is obtained. A variable with the same meaning is also needed for the storage, where prefabricated parts are assembled to final products. Thus, $z_{N+1,j}$ denotes the number of required parts to balance the storage. Formally, $z_{i,j}$ is defined as

$$z_{i,j} = c \cdot g_j - x_{i,j},$$

with $c = \max_{1 \leq j \leq M} \left\lceil \frac{x_{i,j}}{g_j} \right\rceil$.

We now calculate the number of parts of each type needed in order to balance all stations located downstream from a particular station. Let $y_{i,j}$ be the sum of all parts of type j needed to replenish the inventory of all service stations downstream from station i

($i = 1, \dots, N$), i.e.

$$y_{i,j} = \begin{cases} \sum_{l = s_{j,a_{i,j}+1}}^{s_{j,w_j}} z_{l,j} + z_{N+1,j} & \text{if } a_{i,j} \neq 0, \\ 0 & \text{otherwise,} \end{cases}$$

where $a_{i,j}$ is the first processing step of the process flow of part type j that is performed by station i . The relation $a_{i,j} = 0$ implies that service station i never occurs in the process flow of part type j . Thus, $y_{i,j}$ is the exact amount of parts of type j needed in the downstream section of the fabrication facility (from the point of view of station i) to balance the inventory of all stations located there.

A server working according to MRSS at station i will now select the most required part type j^* for service, i.e. the part type for which

$$y_{i,j^*} = \max_{1 \leq j \leq M} \{ y_{i,j} \mid y_{i,j} \neq 0 \}, \quad (1)$$

holds. If there are several types with exactly the same amount of parts needed, a random selection is made among these types. If more than one part of the selected type is waiting in the queue the part with the earliest arrival time at the system is loaded into the server. The advantage of regarding time stamps when selecting the job to be processed next is obvious. Parts that were subject to rework will receive priority at stations which they now have to pass once again. Thus, the variance of cycle times can be reduced. The average sojourn time itself will stay unchanged, however, since this queuing discipline is work-conserving. For a definition of work conserving queuing disciplines we refer to the literature, e.g. Heyman and Sobel (1982).

3.2 Most Requested Part Type per Processing Step

With MRSS the counters of requested parts $y_{i,j}$ depend on the service station. Because of the re-entrant process flow several processing steps on a particular part type may be performed by the same station. Therefore, it is appropriate to distinguish at each station among parts of the same type waiting for different processing steps to be performed. The maintenance of part requirement counters for each processing step yields a more detailed knowledge of how many parts are needed to balance the downstream production system, in particular the assembly station. In this case, we try to balance the inventory at the processing step level instead of at the service station level. This goal is achieved by using the sequencing rule *Most Requested part type per Processing Step*

(MRPS) which employs requirement counters for all processing steps.

For this sequencing rule let $x_{k,j}$ be the number of parts of type j which are currently involved in step k of their manufacturing process. Then, $z_{k,j}$ is the number of parts of type j required to balance processing step k , i.e.

$$z_{k,j} = c \cdot g_j - x_{k,j},$$

$$\text{with } c = \max_{1 \leq j \leq m} \left\lceil \frac{x_{k,j}}{g_j} \right\rceil.$$

We will now calculate the amount of requested components in the downstream part of the manufacturing process as seen from the last processing step that involves station i . Calculating such a value makes only sense if all part types adhere to the same process flow. Otherwise, it is impossible to define the requested material of a specific processing step in the manufacturing process. Therefore, we assume that the process flows of all part types are identical. Under this assumption the number of parts of type j needed to balance the downstream section of the manufacturing process is given by:

$$y_{i,j} = \sum_{k = a_{i,j} + 1}^{w_j} z_{k,j} + z_{N+1,j},$$

where $a_{i,j}$ is the last processing step of the process flow executed at station i . Notice that $a_{i,j}$ cannot be zero since all parts follow the same process flow. Otherwise, not a single processing step would be carried out by service station i .

Invoked at station i MRPS will now select a part of type j^* to be processed next according to Equation (1). In this equation, the index i indicating the service station has to be replaced by the index k which corresponds to processing step k . In case that Equation (1) holds for more than one type, the result of a random selection indicates the part type to be serviced next. If the number of waiting jobs of this particular type is greater than one, the arrival time to the system determines the part that is to enter the server. Parts with earlier time stamps are favored over recent arrivals to the system since it is commonly known that this will reduce the variance of cycle time of manufactured parts. If no part of the chosen type is available at processing step $a_{i,j}$ those types with the next higher number of requested parts are subject to the selection process just described. If continued usage of this selection scheme does not yield an appropriate part for processing step $a_{i,j}$, another processing step performed by this station has to be taken into account. We then look further upstream

in the process flow for the next processing step involving station i . For this step the whole selection process is started once again. If at least a single part is waiting in the queue — and we assume that selection according to a queuing discipline is only invoked if the queue is not empty — this approach leads to the component to be processed next.

MRPS proves to be very flexible concerning reentrant flow. Regardless of how many processing steps involve a particular station and independent of the current stock of parts at this station, MRPS will always make a reasonable choice on which job to be processed next. However, it has the disadvantage that all part types manufactured have to have the same process flow through the system. In manufacturing systems involving different process flows for a variety of prefabricated parts one has to switch to other queuing disciplines, for instance MRSS. However, it turns out that in this case MRSS is outperformed by a sequencing rule which is based on heuristics, cf. Purm (1994).

4 SIMULATION ENVIRONMENTS

We evaluate the performance of the set management policies described in Sec. 3 for several production environments by means of simulation. These environments are introduced briefly. We assume that a set consists of two parts, one of type A and one of type B . Further, we assume that both arrival and service times are exponentially distributed. The arrival rate and the mean service time are adjusted so as to achieve a utilization of $\rho = 0.9$ at each service station.

4.1 Basic Model with Rework

Since we focus mainly on rework we start with a single $M/M/1$ Bernoulli feedback station (cf. Figure 1).




Figure 1: Basic Model with Rework

Customers of both types arrive according to a Poisson process with rate λ ; service times are exponentially distributed for both types with mean μ^{-1} . In this model rework is incorporated via the feedback probability p . After service time completion parts are fed back to the tail of the queue with probability p . For an arbitrary customer this probability is independent of its history, and therefore independent of the number of feedbacks which this particular customer already has experienced (Bernoulli feedback,

see for example van den Berg and Boxma 1987). Finally, parts successfully processed have to wait for assembly at the storage. We examine this model for a rework probability of 10%.

4.2 Basic Model with Re-entrant Flow

The basic model incorporating re-entrant flow consists of three service stations as shown in Figure 2. Customers take the following route during their visit to the system. First, they have to be processed at station 1, then they have to pass three times the sequence consisting of station 2 and station 3. Before entering the storage, parts have to receive service one more time at station 2. Additionally, there is a rework probability of 0.1 at each station.




Figure 2: Basic Model with Re-entrant Flow

Simulation results for these basic models were obtained by 20 independent replications. The length of a replication was set to 100,000 time units whereas the initial transient phase was assumed to be finished after 10,000 time units.

4.3 Semiconductor Fabrication Facility

We evaluate the performance of the queuing disciplines in realistic production environments by means of the semiconductor manufacturing facility investigated by Resende (1984). In this system, 74 processing steps have to be carried out on 28 different tools to produce DRAMs. The process flow of this system is highly re-entrant, service times depend on the processing step, and rework is carried out at the same tools as it occurs (local feedback). In opposite to the original model published by Resende, we assume that there are no tool breakdowns, the rest of the model remaining unchanged. For further details, the interested reader is referred to the original paper.

In order to achieve better confidence intervals for this fabrication facility, we examined 20 independent replications, each with an initial transient phase of 1000 final products manufactured (equivalent of approx. 160,000 time units) and a steady state length of 10,000 assembled sets (equivalent of approx. 1,600,000 time units). Station 5 was the bottleneck

tool and the arrival rate was tailored in a way to achieve a load of $\rho = 0.88$ at this station. As before a set was made up of two components with different type.

In the following section, we compare simulation results obtained by set management policies MRSS and MRPS to those of two widely used sequencing rules, *First-Come First-Served* (FCFS) and *Earliest Due Date* (EDD). The employment of FCFS implies the service of waiting parts according to the order of their arrival. EDD requires the assignment of due dates to parts entering the system and selects that part with the closest due date for service. Instead of setting due dates explicitly we assign the arrival time to each part and select that part which has the longest life time in the manufacturing system to be processed next. This approach corresponds exactly to the EDD sequencing rule if due dates are set equally for all parts to the arrival time plus a fixed amount of time. In case that there are several queued parts with the same due date, FCFS is used among these parts to dissolve the contention for service.

Additionally to FCFS and EDD one might suggest to consider also the sequencing rule *Shortest Remaining Processing Time* (SRPT) in order to evaluate the performance of MRSS and MRPS. SRPT is common both in theory and in practice. However, SRPT turns out to be not well suited for our purposes. Obviously, given that two lots waiting for the same processing step to be performed at a particular machine and having experienced a different number of rework steps, SRPT is incapable to distinguish among these lots since the remaining processing times are the same for both lots.

5 NUMERICAL RESULTS

In the following tables simulation results for the set synchronization time S and the cycle time T of parts A and B are shown. For the performance measure X we denote by $E[X]$ the mean of X and by $\text{Var}[X]$ the variance of X . In addition to the mean and the variance we show confidence intervals for a confidence level of 95% (within tables denoted by “Conf.Int.”).

5.1 Basic Model with Rework

Table 1 contains the simulation results for the basic production system featuring rework. Concerning the quality of the flow management there is no major discernible difference between the queuing disciplines except for FCFS. Since FCFS does not take into account the part requirements at the storage the employment of FCFS results in a higher mean and a higher variance of the synchronization time S . The mean and

variance of cycle times are roughly the same for both part types and all sequencing rules. The slight differences occurring are not significant since the 95 % confidence intervals overlap each other. Since entire sets arrive at the system with single parts being arranged at random and since the sequencing rules applied are work conserving, it is clear that cycle times are independent from the discipline used.

5.2 Basic Model with Re-entrant Flow

As observed for the previously examined system EDD outperforms FCFS substantially in all characteristics (cf. Table 2). Obviously, MRSS might not be used in an environment where re-entrant flow occurs. While the mean and the variance of cycle times are not affected by this discipline, the corresponding results for the set synchronization time are worse. This is due to the fact that MRSS can not distinguish between parts of a particular type waiting for different processing steps to be performed by the same service station. This is an immediate consequence of the fact that the selection scheme used by sequencing rule MRSS is based upon request counters attached to each station. On the other hand, the impressive results achieved by MRPS prove that request counters maintained at each processing step are very effective since they assure the independence of the material flow control from re-entrant flow. Despite the fact that the mean and the variance of cycle times are slightly higher than compared to EDD, the use of MRPS cuts the $E[S]$ and $\text{Var}[S]$ in half.

Additional simulation experiments have shown that the performance advantage of MRPS over EDD does not depend on the degree of re-entrant flow. Simulation results for the basic system with re-entrant flow where the number of processing cycles leading through stations 2 and 3 was increased are very similar to those shown in Table 2. Irrespective of the number of processing cycles EDD never outperformed MRPS in terms of $E[S]$ and $\text{Var}[S]$.

We already examined systems without re-entrant flow consisting of 5 or 15 service stations. The results for various rework settings show that MRSS is well suited to control the material flow of production lines without re-entrant flow. In this case, MRSS achieves the best synchronization times among the sequencing rules considered in this paper, whereas in general, cycle times are slightly higher and more variable at the same time. However, due to the lack of space the results are not reported explicitly.

5.3 Semiconductor Fabrication Facility

The results shown in Table 3 exhibit that both mean cycle times and the variance of cycle times of the semiconductor fabrication facility are independent of the sequencing rule employed while the variance thereof is slightly improved by EDD. The synchronization time for MRSS is by far the worst in terms of the mean and the variance. As already shown by the results for the basic model with re-entrant flow this again proves that MRSS is not appropriate for manufacturing environments featuring re-entrant flow. It is therefore necessary to attach part requirement counters to processing steps themselves. This notion is supported by the results achieved by the set management policy MRPS. Here, the mean and the variance of set synchronization times can be reduced to a very low level, topping even the results of EDD. The only disadvantage of the reduction of set synchronization times is a slight increase in the variance of cycle times for both components.

Additional experiments have shown that the improvements achieved by MRPS depend on the system load. Simulations for the semiconductor fabrication line where the load ranged between 0.5 and 0.7 at the bottleneck station showed that $E[S]$ and $\text{Var}[S]$ varied only slightly for different queueing disciplines. However, for $\rho \geq 0.7$ the use of MRPS outperforms EDD in terms of these system characteristics. At the same time, for MRPS, the variance of cycle times increases roughly proportional to the system load, while staying on an even level for EDD. However, this increase is only of secondary importance since the mean cycle time stays the same for all policies under consideration and for varying loads.

6 CONCLUSION

We conclude that even under realistic production conditions as provided by the semiconductor fabrication facility, set management policy MRPS outperforms EDD in terms of lower set synchronization times at the assembly station. The slight increase in the variance of cycle times compared to EDD is inconceivable. On the other hand, MRSS is not equipped to handle re-entrant flow and — as expected — the results achieved grow worse with system complexity. However, in manufacturing systems without re-entrant flow MRSS performs very well. The commonly used queueing discipline FCFS is outperformed by MRPS in every production environment we examined.

One might call into question the applicability of the sequencing rules MRSS and MRPS since they

are based on global status information of the entire production facility. However, in modern automated production environments this information should be available and accessible for online decisions like which job is to be processed next.

It remains to future research to investigate more realistic production environments including not only rework and re-entrant flow but also, among many others, the availability of operators, machine break-

downs, setup times, batching effects, and scrap. We already examined the set management policies MRSS and MRPS for our basic models with machine breakdowns. It turned out that also in this case our new sequencing rules outperform standard sequencing rules. Therefore, we are quite confident that these new sequencing rules will prove to be very useful even under real production conditions.

Table 1: Simulation Results for Basic Model with Rework, $\rho = 0.9$

Q.Dis.	E[S]	Conf.Int.	Var[S]	E[T]	Conf.Int.	Var[T]	E[T]	Conf.Int.	Var[T]
				part type A			part type B		
FCFS	2.07	± 0.0383	5.01	16.7	± 0.731	316	16.7	± 0.727	313
EDD	1.11	± 0.00234	1.23	17.5	± 0.770	315	17.5	± 0.770	315
MRSS	1.11	± 0.00267	1.24	16.8	± 0.705	277	16.8	± 0.704	277
MRPS	1.11	± 0.00227	1.23	16.5	± 0.497	264	16.5	± 0.498	264

Table 2: Simulation Results for Basic Model with Re-entrant Flow, $\rho = 0.9$

Q.Dis.	E[S]	Conf.Int.	Var[S]	E[T]	Conf.Int.	Var[T]	E[T]	Conf.Int.	Var[T]
				part type A			part type B		
FCFS	2.94	± 0.0402	7.49	42.7	± 1.24	786	42.7	± 1.24	787
EDD	0.796	± 0.00238	0.562	39.3	± 0.969	549	39.3	± 0.968	549
MRSS	26.4	± 1.54	1030	40.2	± 0.988	1010	40.9	± 1.38	1170
MRPS	0.382	± 0.00238	0.300	44.4	± 1.03	803	44.4	± 1.03	803

Table 3: Simulation Results for the Semiconductor Fab, $\rho = 0.88$

Q.Dis.	E[S]	Conf.Int.	Var[S]	E[T]	Conf.	Var[T]	E[T]	Conf.	Var[T]
				part type A			part type B		
FCFS	506	± 4.92	157,000	10,400	± 101	4,240,000	10,400	± 98	4,260,000
EDD	448	± 6.00	128,000	10,700	± 162	3,690,000	10,700	± 162	3,680,000
MRSS	1,360	± 132	3,200,000	10,500	± 147	4,640,000	10,400	± 155	4,100,000
MRPS	393	± 3.83	10,9000	10,900	± 145	4,250,000	10,900	± 144	4,250,000

ACKNOWLEDGMENTS

The authors would like to thank Prof. Dr. P. Tran-Gia for the support of this work and Prof. Dr. H. Fromm for stimulating discussions.

REFERENCES

- Ahmed, S. and R. R. Tummala. 1992. Overview of packaging for the IBM Enterprise System/9000 based on the glass-ceramic copper/thin film thermal conduction module. *IEEE Transactions on Components, Hybrids and Manufacturing Technology* 15, 426–431.
- De, P., J. B. Gosh, and C. E. Wells. 1989. A note on the minimization of mean squared deviation of completion times about a common due date. *Management Science* 35(6), 1143–1147.
- De, P., J. B. Gosh, and C. E. Wells. 1992. On the minimization of completion time variance with a bicriteria extension. *Operations Research* 40(6), 1148–1155.
- Gise, P. and R. Blanchard. 1986. *Modern semiconductor fabrication technology*. Englewood Cliffs, NJ: Prentice-Hall.
- Hall, N. G., W. Kubiak, and S. P. Sethi. 1991. Earliness tardiness scheduling problems II: Deviation of completion times about a restrictive common due dates. *Operations Research* 39(5), 847–856.
- Heyman, D. P. and M. J. Sobel. 1982. *Stochastic Models in Operations Research, Volume I: Stochastic Processes and Operating Characteristics*. New York: McGraw-Hill.
- Hogg, G. L., J. Fowler, and M. Ibrahim. 1991. Flow control in semiconductor manufacturing: A survey and projection of needs. Technical report 91110757A-GEN, SEMATECH, Austin, TX.
- Johri, P. K. 1992. Practical issues in scheduling and dispatching in semiconductor wafer fabrication. *Journal of Manufacturing Systems* 12(6), 474–485.
- Lee, C.-Y. 1992. Efficient algorithms for scheduling semiconductor burn-in operations. *Operations Research* 40(4), 764–775.
- Purm, M. 1994. Reduzierung von Gruppenbearbeitungszeiten in Produktionssystemen. Master's thesis, Institute of Computer Science, University of Würzburg, Würzburg, Germany. In German.
- Resende, M. G. C. 1984. Computer simulation of semiconductor wafer fabrication. Technical Report ORC 86-14, Operations Research Center, University of California at Berkeley, Berkeley, CA.
- Uzsoy, R., L. A. Martin-Vega, C.-Y. Lee, and P. A. Leonard. 1991. Production scheduling algorithms for a semiconductor test facility. *IEEE Transactions on Semiconductor Manufacturing* 4(4), 270–280.
- Womack, J. P., D. T. Jones, and D. Roos. 1990. *The machine that changed the world*. New York, NY: Rawson Associates.

AUTHOR BIOGRAPHIES

MANFRED MITTLER is an Assistant Professor at the Institute of Computer Science of the University of Würzburg, Germany, since 1992. He received the Master's degree in Mathematics from the same university and is currently pursuing his Ph.D. His research interests include the application of queueing theory, statistics, simulation, and probability theory to the performance analysis of manufacturing systems and telecommunication systems. Mr. Mittler is member of INFORMS, GI (German Chapter of the ACM), DGOR (German Operations Research Society), and VDI (German Society of Engineers).

MICHAEL PURM received the Master's degree in Computer Science from the University of Würzburg, Germany, in 1995. He wrote his Master thesis on the performance evaluation of set management sequencing rules. Currently, Mr. Purm is with Nortel/DASA, Friedrichshafen, Germany.

OTTMAR GIHR received the Master degree in Electrical Engineering from the University of Stuttgart, Germany, in 1984 and the Ph.D. degree in 1989 from the same university. From 1989 to 1990 Dr. Gihl was a visiting scientist at the Centre for Teletraffic Research of Bond University, Queensland, Australia. Since 1990 Dr. Gihl is with the German Manufacturing Technology Center of IBM Germany in Sindelfingen, where he is involved in the performance analysis of semiconductor manufacturing systems and the application of queueing analysis in business process reengineering.