Proceedings of the 2018 Winter Simulation Conference
M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, eds.

SIMULATION OPTIMIZATION FOR PLANNING PRODUCT TRANSITIONS IN SEMICONDUCTOR MANUFACTURING FACILITIES

Atchyuta Bharadwaj Manda
Reha Uzsoy

Edward P. Fitts Department of Industrial and Systems Engineering
North Carolina State University
111 Lampe Dr
Raleigh, 27695-7906, NC, USA

ABSTRACT

The introduction of a new product into a semiconductor manufacturing facility can have significant adverse effects on both current and new products. We combine a simulation model of a simplified production system and a learning model to optimize releases of old and new products using simulation optimization. Our results show that the simulation optimization model yields significant improvements in performance over simple alternatives, and provides insights into the structure of optimal release policies.

1 INTRODUCTION

Managing new product introductions is critical for companies in high technology industries like semiconductor manufacturing to attract and retain customers. A new product with enhanced features and/or improved performance commands higher prices early in its life cycle when the competition is still catching up, the new product is in high demand, and its price is high (Leachman and Ding 2007; Nemoto et al. 2000). Hence it is important for the firm to rapidly ramp up production of good devices when both demand and prices are high. However, introduction of a new product into an operating factory where it shares capacity with other products adversely impacts all products produced by the factory (Leachman 1993) as problems are identified and require engineering intervention to fix. The time required for these interventions decreases as experience is gained in producing the new product, requiring the modeling of learning effects in production planning.

Effective management of new product introductions requires consideration of both the adverse impact of the new product on other products in the factory, and the fact that the new product will consume more resources than a mature product early in its life cycle. To this end, we study a simplified system that manages the impact of the new product by controlling the releases of both new and old products. We first create a simulation model to capture the impact of a new product introduction on a production line in which an stable product, which will be replaced over some time horizon by the new product, is currently being produced. Klein and Kalir (2006) describe such a situation in an Intel fab which was required to transition from producing logic products to flash memory products. Introduction of the new product on this production line causes increased disruptions affecting both old and new products. A simulation optimization procedure is used to find the releases of both old and new products to maximize the expected contribution over the duration of the product transition.

The rest of the paper is organized as follows. Section 2 provides a brief review of previous related work. Section 3 summarizes a mathematical model of new product introduction developed in previous work which we use in our simulation model. Section 4 presents the simulation model, while the simulation optimization procedure is presented in Section 5. Section 6 describes the numerical experiments and results, while Section 7 concludes the paper with a summary of our principal findings and some directions for future work.

2 LITERATURE REVIEW

The work in this paper lies at the intersection of two research streams: that on learning in manufacturing systems, and the body of queuing and simulation studies that examine the impact of different sources of variability on the performance of manufacturing systems. Our focus in this work is on planning the releases of work into an operating factory in the presence of learning effects as the difficulties encountered with a newly introduced product are identified and remediated over time.

There is an extensive body of research on learning models in manufacturing systems, which mainly deals with measuring and modeling the improvements in productivity obtained as experience of performing a new task is gained. Anzanello and Fogliatto (2011) and Yelle (1979) review the classical learning curve literature and present various models of learning developed across many industries. Liao (1979) and Reeves and Sweigart (1981) assume that learning is driven by cumulative production, and leads to a decrease in capacity consumption by the product.

Later research, such as Chand et al. (1996), Fine (1986; 1988) and Fine and Porteus (1989), distinguishes between learning that takes place through repetition of the task and learning obtained from experimentation that consumes production capacity, and thus competes with revenue generating production. Terwiesch and Bohn (2001) distinguish learning through experimentation from learning by cumulative production. They formulate a dynamic programming model to decide when to experiment and when to produce but do not consider manufacturing lead times or the effects of congestion. Similar problems have been addressed by Tirkel et al. (2016) and Gilenson et al. (2015), among others. Haller et al. (2003) present a methodology to manage the conflict between the requirement of short cycle times and rapid increase in throughput while bringing a new factory up to full production volume. They suggest the use of maximum limits on the WIP in different segments of the line to maintain an acceptable tradeoff between WIP and cycle time, similar to the SLIM approach of Leachman et al. (2002). The WIP limits are updated weekly over the duration of the ramp, and simulation experiments show beneficial results. However, this work considers a single product and does not address interactions between different products.

Several authors have used simulation models to study the behavior of wafer fabs during product transitions or under changing product mix. Nemoto et al. (2000) examine the financial benefit of cycle time reduction in ramping up a new process using stochastic simulation, also considering the impact of decreasing device prices over time. Dümmler (2000) presents a series of interesting simulation studies examining the behavior of a wafer fab in the presence of temporary surges in wafer starts, changing product mixes and different release and dispatching policies. His findings are consistent with the simple mathematical model we develop in this paper, and suggest that the impact of new product introductions on the variability of the effective processing time in the fab is a major contributor to adverse performance. Klein and Kalir (2006) describe a simulation model used to study the performance of an Intel wafer fab transitioning from logic to flash memory products, and found that the simulation successfully identified potential bottlenecks that were not identified by a static capacity analysis. The simulation model was used to suggest managerial actions that avoided performance degradation during the product transition.

A number of authors have focused on yield learning, the improvement in the fraction of good devices over time as the firm gains experience with the product. Bohn (1995), Gruber (1992; 1994), Macher and Mowery (2003) and others examine empirical data on yield learning over time. Tirkel (2013) reviews various yield learning models discussed in the literature. In contrast, our work in this paper focuses on the improvement in factory productivity, as measured by output, over time in the presence of learning by fab engineering personnel remediating production problems on the shop floor.

Kim and Uzsoy (2008;2013) propose an integrated production planning model using clearing functions that incorporate learning effects by considering production and engineering lots separately. Production lots are sold to generate revenue, whereas engineering lots help increase capacity in future periods. They perform a marginal cost analysis to provide insights on managing the system. However, they do not explore the problem of ramp up and new product introduction. The idea of distinguishing production and engineering lots is applied to a full fab model by Ziarnetzky et al. (2017), who give three production planning formulations with fixed lead times that incorporate learning and different capacity allocation scenarios for the engineering and production lots.

Manda et al. (2016) formulate a mathematical model using clearing functions that captures the impact of a new product introduction on the effective processing time distribution of a production line on which an existing product is already being produced, while incorporating the learning effects gained by experience in producing the new product. The impact of the new product introduction on factory performance is captured through the mean and variance of the effective processing time, the time required for a job to be processed accounting for all detractors such as setups and machine failures. Specifically, the new product randomly encounters production difficulties that require the production equipment to be taken offline for the problem to be resolved. This creates a stochastic series of non-preemptive disruptions whose frequency decreases over time as experience with the new product accumulates. In this paper, we embed this mathematical model into a simulation model of a simplified production system, and use simulation optimization to obtain product release schedules that maximize the total contribution (revenue minus variable costs) from both products over the duration of the product transition.

3 MODELING AND ANALYSIS

We consider a simplified production system represented as a single resource which is currently producing a mature product, Product 1, to meet a known, stable demand. We assume that the manufacturing process for this product has been debugged to the point that it has achieved its steady-state effective processing time distribution. At a specified point in time demand for this product begins to decrease and is replaced by demand for a new product, Product 2, whose demand eventually replaces that of Product 1 entirely.

The mature product, Product 1, has a natural processing time with mean t_0 and standard deviation σ_0. A process event requiring engineering activity occurs following a Poisson process on average after every Q_1 units are processed, with mean duration P time units and standard deviation of σ_p.

Product 2 has the same natural processing time and engineering activity duration distributions as Product 1. However, this product requires engineering activities more frequently, on average, after every Q_{20} units, again following a Poisson process, when it is first introduced. As process issues are identified and resolved, the average number of units processed between disruptions increases from Q_{20} to Q_{2s} following the exponential learning model

$$Q_{2t} = Q_{20} + (Q_{2s} - Q_{20})(1 - e^{-\alpha X_2(t)}) \tag{1}$$

where $X_2(t)$ denotes the cumulative production of Product 2 in the time interval $[0,t)$ and α is a parameter governing the rate of learning. The mean and standard deviation of the effective processing time, representing the contributions of both products, is given by

$$t_e(t) = t_0 + \frac{P}{Q_t} \tag{2}$$

$$\sigma_e^2(t) = \sigma_0^2 + \frac{\sigma_p^2}{Q_t} + \frac{(Q_t - 1)}{Q_t^2} P^2 \tag{3}$$

where

$$Q_t = \left(Q_{10} \left(\frac{W_{1t}}{W_{1t} + W_{2t}} \right) + \left(Q_{20} + (Q_{2s} - Q_{20})(1 - e^{-\alpha X_2(t)}) \right) \left(\frac{W_{2t}}{W_{1t} + W_{2t}} \right) \right) \tag{4}$$

and W_{it} denotes the workload of product i at the start of period t.

In ongoing work (Manda and Uzsoy 2018) we have extended the approach of Manda et al. (2016) by deriving a nonlinear clearing function describing the expected throughput of the system in each planning period as a function of the workload of each product available for processing in each period. This clearing function can then be incorporated into a deterministic nonlinear optimization model that determines the optimal release schedule of each product in each period. However, the presence of the learning model (4) in the clearing function results in a non-convex feasible region for which even locally optimal solutions are difficult to obtain in reasonable CPU times. In addition, the practical problem is subject to many stochastic elements, including the demands for each period, the amount of learning in each period and line yield. The use of simulation optimization allows us to progressively implement simulation models of more complex multi-stage production systems for which deterministic optimization models are computationally burdensome. The work in this paper constitutes a first step in this direction by using a simplified, aggregate model of a wafer fab in implementing a product transition.

4 SIMULATION MODEL

In order to explore the performance of simulation optimization for the problem of interest, we consider a simplified production system where the wafer fab is modeled as a single server queuing system with two customer classes over a planning horizon consisting of J discrete periods. The discrete periods are of the order of three months in duration, representing the common time scale used in the semiconductor industry when considering production transitions. The structure of the queuing system is illustrated in Figure 1. Releases of each product into the fab during each planning period are assumed to be uniformly distributed

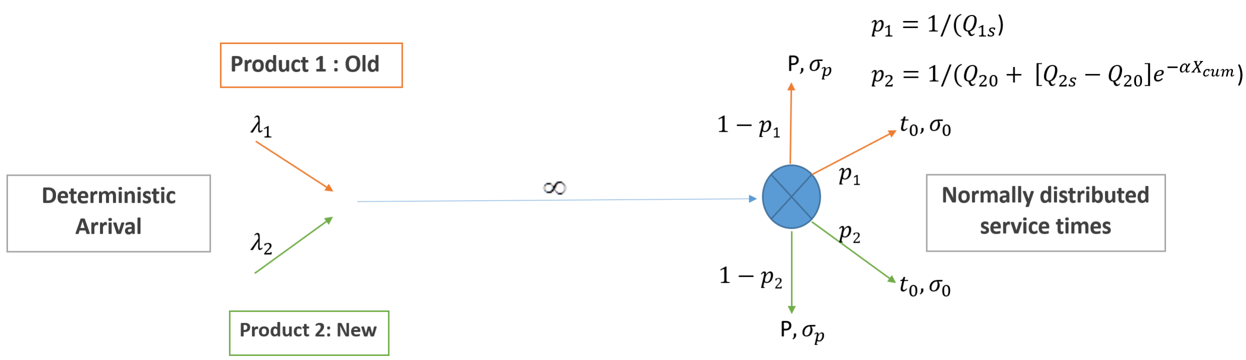

Figure 1: Schematic of queuing system.

over the period with negligible variability, and the demand D_{it} for each product i in period t is assumed to be known with certainty. The quantities R_{it} of each product i released in each planning period t are the decision variables. We shall denote the mean arrival rate of product i in planning period t by $\lambda_{it} = \frac{R_{it}}{T}$.

Each arriving unit (lot) of product i will encounter an engineering disruption with probability p_{it}. The disruption probability of the stable product, Product 1, remains constant at $p_1 = \frac{1}{Q_{1s}}$ for the entire planning horizon, while that of Product 2 decreases over the planning horizon following the learning model, with disruption rate

$$p_{2t} = \frac{1}{Q_{20} + \left(Q_{2s} - Q_{20}e^{-\alpha X_2(t)}\right)} \tag{5}$$

at time t. The duration of an engineering activity is same for both the products and has a mean and standard deviation given by P and σ_p. With probability $1-p_{2t}$ for Product 2 and probability $1-p_{1t}$ for Product 1, the service time distributions for these products follow the natural processing time distribution.

The simulation operates in continuous time given a vector of the decision variables R_{it} specifying the number of lots of each product i released in period t. When a lot of product i arrives at the server at some time τ it generates a disruption with probability $p_{i\tau}$. Otherwise the service time is generated from the natural processing time distribution. At each departure epoch t of a lot of Product 2 from the system, the cumulative production $X_2(t)$ and the disruption probability p_{2t} of a disruption for the next arriving lot of Product 2 are updated based on the learning model (5). At the end of each period, the total throughput of each product obtained in the period and the number of each product remaining in the system are stored. Using the demand for the period, the backorders and finished goods inventory of each product is computed at the end of each period. Finally, the total contribution (revenue - variable production costs) is computed over the entire time horizon of J periods. We obtain n independent replications of the simulation for each release vector and calculate the expected total contribution (revenue minus variable costs) of the system over the planning horizon, and report the 95% confidence interval for the total contribution.

5 SIMULATION OPTIMIZATION PROCEDURE

The simulation optimization procedure seeks a release vector R_{it} for each time period t that maximizes the expected total contribution over the planning horizon of J periods. The problem can be stated as follows:

$$\min \quad E\left[\sum_{\forall t,i}[\pi_{it}(D_{it}+S_{it-1}-S_{it})-(r_{it}R_{it}+w_{it}W_{it}+i_{it}I_{it}+s_{it}S_{it})]\right]$$

$$\text{s.t.} \tag{6}$$

$$R_{it} \geq 0 \tag{7}$$

S_{it} denotes the number of units of product i backlogged at the end period t, W_{it} the amount of product i in work in progress (WIP) inventory at the end of period t, I_{it} the amount of product i in finished goods inventory at the end of period t. r_{it}, w_{it}, i_{it} and s_{it} are the associated cost coefficients, and π_{it} the market price of product i in period t.

We use a Genetic Algorithm (GA) for our simulation optimization engine in this exploratory work due to its simplicity of implementation. This requires careful selection of the algorithm parameters, particularly population size, mutation and crossover probabilities, elite and immigration fractions to ensure efficient convergence to a global optimum in reasonable CPU times. Extensive initial experiments explored the impact of these parameters on performance, leading to our final choices. However, despite these settings we experienced extremely slow convergence when starting the GA from a completely random initial population. To address this issue, we seeded the initial population with several release vectors obtained from a relaxed version of the nonlinear programming model of Manda and Uzsoy (2018), which resulted in significantly more rapid convergence to good solutions. The specific parameter settings and results of the computational experiments are described in the following section. This approach of seeding a search heuristic with different initial solutions is widely used in the literature, as discussed by Kalir and Sarin (1999) for simulated annealing. The specific approach we use is to obtain optimal solutions to a relaxed problem by perturbing the problem data, based on the problem space approaches of Storer et al. (1992).

We use Matlab to implement the Genetic Algorithm. By randomly perturbing the demand data, the nonlinear programming model was used to create 6 good initial solutions that were used in initial population. The total population size was 100 and the rest of the population is generated randomly. The elite count is 10. There are 90 individuals other than elite children and the cross over fraction is set to 0.6. This means that of the 90 remaining individuals, on average $(0.6)(90) = 54$ are crossover children and the remaining 36 are mutation children. The mutation, and scale and shrink and migration options are left at their default values. Four independent replications of the GA are run in each experiment.

6 NUMERICAL EXPERIMENTS

In this section, we examine the performance of the model and the structure of the near-optimal solutions yielded by the simulation optimization procedure. We consider a product transition in which demand

Table 1: Demand data.

Demand	1	2	3	4	5	6	7	8	9	10
Product 1	1250	1250	1250	937.5	625	312.5	0	0	0	0
Product 2	0	0	0	312.5	625	937.5	1250	1250	1250	1250

Table 2: Parameter values.

Parameters	J	T	t_0	c_0	P	c_p	α	$Q_{1s} = Q_{2s}$	Q_{20}
Values	10	129600	80	0.25	800	0.5	0.0019/0.0030	50	10

Table 3: Cost Parameters.

Parameters	π_1	π_2	s_1	s_2	w_1	w_2	i_1	i_2
Values	5	7.5	10	15	1	1	4	4

for a mature product is replaced by that for a new product over a 10 period planning horizon. As each period represents a quarter, the total planning horizon is two and half years. Table 1 gives the demand for the two products over the planning horizon. Only Product 1 has demand during the initial three periods. From periods 4 to 6, the demand for Product 1 decreases to zero and that for Product 2 increases. From period 7 onwards we only have demand for Product 2. The average utilization in period 1 for Product 1 is approximately 91%. This utilization changes over the planning horizon as the new product is introduced, altering the effective processing time distribution following the learning model.

We conduct two numerical experiments with two different learning rates α in the learning model (1). The values of α and the other parameters used in these experiments are summarized in Tables 2 and 3. Note that the learning model is completely deterministic except in its dependence on the cumulative production of Product 2 which is, of course, a random variable. More complex, stochastic versions of this model, such as that suggested by Filho and Uzsoy (2013), have been proposed and will be included in future work.

Each lot has a natural processing time t_0 of 80 minutes with coefficient of variation c_0 = 0.25. The new product, Product 2, causes a disruption on average every Q_{20} = 10 lots which, with experience gained by producing the product, increases to Q_{2s} = 50 lots. The mature product, Product 1, causes a disruption on average every Q_{1s} = 50 lots. Each disruption lasts an average of P = 800 minutes with coefficient of variation c_p = 0.75.

The cost parameters are given in Table 3. Each lot of Product 1 generates a revenue of 5 while Product 2 generates a revenue of 7.5. The unit backorder cost is assumed to be 10 for Product 1 and 15 for Product 2. The finished goods inventory holding cost is assumed to be 4, in line with the rapid price decrease and obsolescence seen in the the semiconductor industry. There is no cost for releasing a product into the fab, but a WIP holding cost of 1 per unit per period is assessed for both products.

We compare the performance of the simulation optimization procedure to that of a simple procedure suggested by practitioners in which $R_{it} = D_{it}$ for both products. The demand data in our experiments was chosen such that the total demand for both products remains constant across the planning horizon. This essentially replaces releases of one product with those of the other, which is quite reasonable if the workload imposed on the factory by each product is the same. However, due to the increased frequency of disruptions associated with product 2, this is not the case in our experiments, with severe consequences for factory performance.

6.1 Experiment 1: Slow Learning

In this experiment we use a alpha value of 0.0019, representing a relatively slow rate of learning. The results summarized in Figures 3 and 4, are quite striking: the baseline scheme performs extremely poorly, generally yielding negative expected contribution. Early in its life cycle, the high rate of disruptions experiences by Product 2 adversely affects the effective processing time distribution, raising utilization above 1 in some periods. This results in long cycle times, high backlogs and lower revenues. Figure 2illustrates the evolution of the mean effective processing time over the planning horizon in one particular simulation replication. The mean effective processing time in period 7 is more than 25% higher than that at the start of period 1, which represents the effective processing time of the stable Product 1. It is thus easy to see how simply replacing releases of Product 1 with an equal number of releases of Product 2 will yield poor performance.

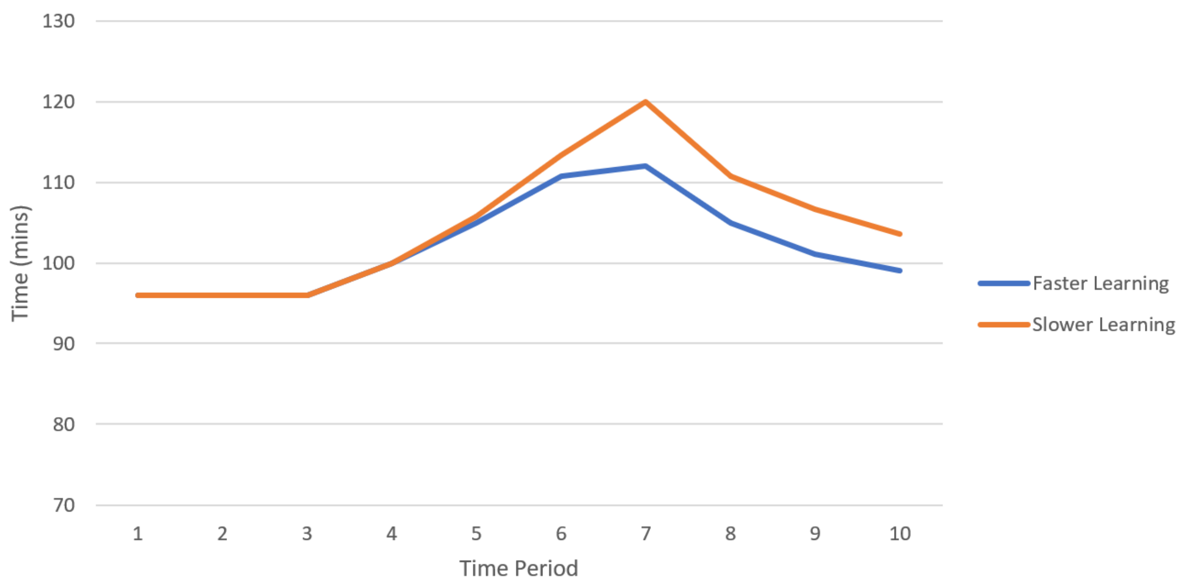

Figure 2: A Sample effective processing time plot for fast and slow Learning.

The reasons for the improvements obtained by the simulation optimization can be seen in Figure 4. While there are some differences between the different GA replications, a common theme emerges. Early in the planning horizon releases of Product 1 are increased to build inventory against future demand. Small quantities of Product 2 are released to take advantage of the decreasing learning rate of the exponential learning model, under which low production quantities early in the life cycle result in large improvements in Q. Inventory of Product 2 is accumulated during periods 4 through 6. Due to the build up of inventory of Product 1 in the first 3 periods, the system is now free and more units of Product 2 can be released, increasing the rate of learning and building inventory of Product 2 during periods 4 through 6 which is then used to meet demand in periods 7 through 9.

The important observations are the front-loading of releases of both products to build inventory of Product 1 before operations are disrupted by the introduction of large quantities of Product 2, and the introduction of small quantities of Product 2 early in the horizon to take advantage of the high learning benefits without unduly disrupting production of Product 1. Clearly, careful planning of releases provides a significant performance advantage in this situation.

Another striking feature of the results in Figure 3 is the width of the confidence intervals obtained. Even in this very simplified production system, the performance of the system over the planning horizon is subject to a high degree of uncertainty, which can only be increased by the presence of multiple work centers with different processing and learning characteristics, shifting production bottlenecks and multiple products whose demand ramps up and down at different times.

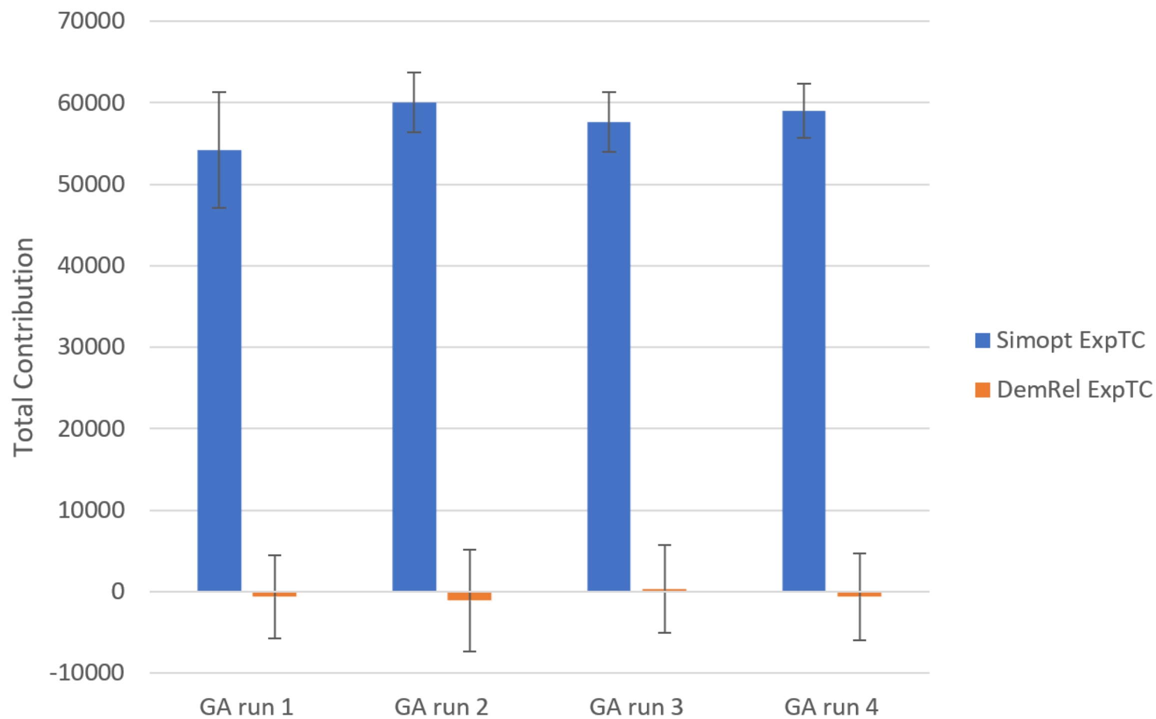

Figure 3: Expected Total Contribution for Slow Learning showing 95% Confidence Intervals.

6.2 Experiment 2: Faster Learning

This experiment uses a value of $\alpha = 0.003$ which increases the rate of learning obtained by producing Product 2. The results, shown in Figures 2, 5 and 6, are qualitatively similar to those for the slow learning experiment, although because there is less variability between the GA replications the figure is easier to read. Once again releases, of Product 1 are raised early in the horizon, while small quantities of Product 2 are released to start the learning process without undue disruption of Product 1. The accumulated inventory of Product 1 allows its releases to be reduced in periods 5 and 6, increasing releases of Product 2 and allowing inventory to be built up as the effective processing time recovers towards the original distribution. The baseline release schedule now performs substantially better than in the slow learning experiment, but is still dominated by the simulation optimization procedure.

7 CONCLUSIONS AND FUTURE WORK

The work in this paper constitutes a very early stage in the development of a simulation optimization environment to study the effective management of product transitions in semiconductor wafer fabrication. While the production system is stylized to a single-stage queuing model, this paper is the first to our knowledge to consider two interacting products during the product transition, and to propose a specific mathematical model of their interactions. Most previous work, except the simulation studies of Dümmler (2000) and Klein and Kalir (2006) has generally focused on a single product, ignoring the impact of product introductions on other products with which they share capacity.

A wide range of future research directions remain open in what appears to be a complex but under-researched area. The development of deterministic optimization models for planning releases in the presence of learning and product interactions during ramp-up is an important direction for future work. Manda and Uzsoy (2018) have developed such models, which are rendered non-convex by the presence of the nonlinear

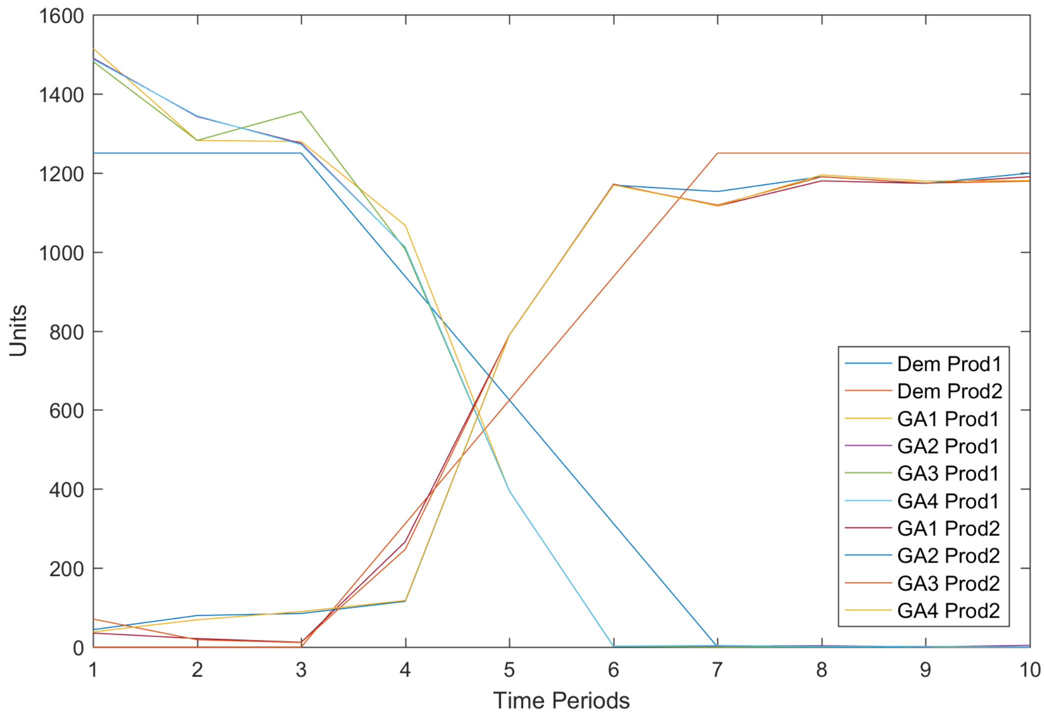

Figure 4: Product Releases for Slow Learning Experiment.

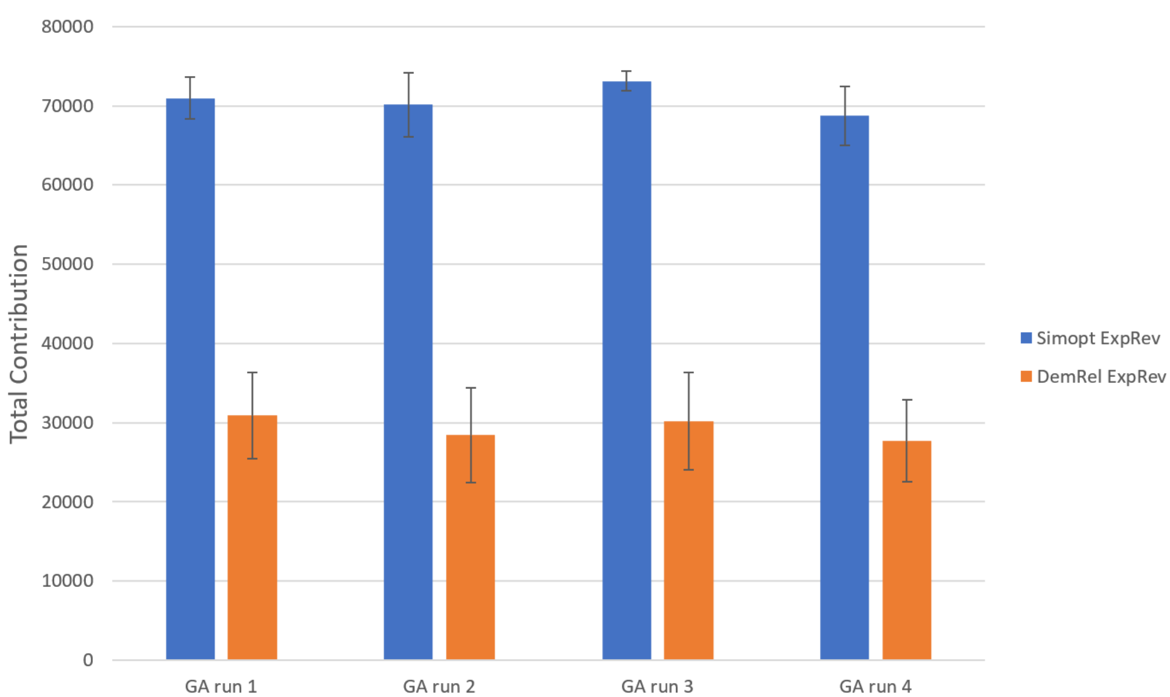

Figure 5: Expected Total Contribution for Fast Learning showing 95% Confidence Intervals.

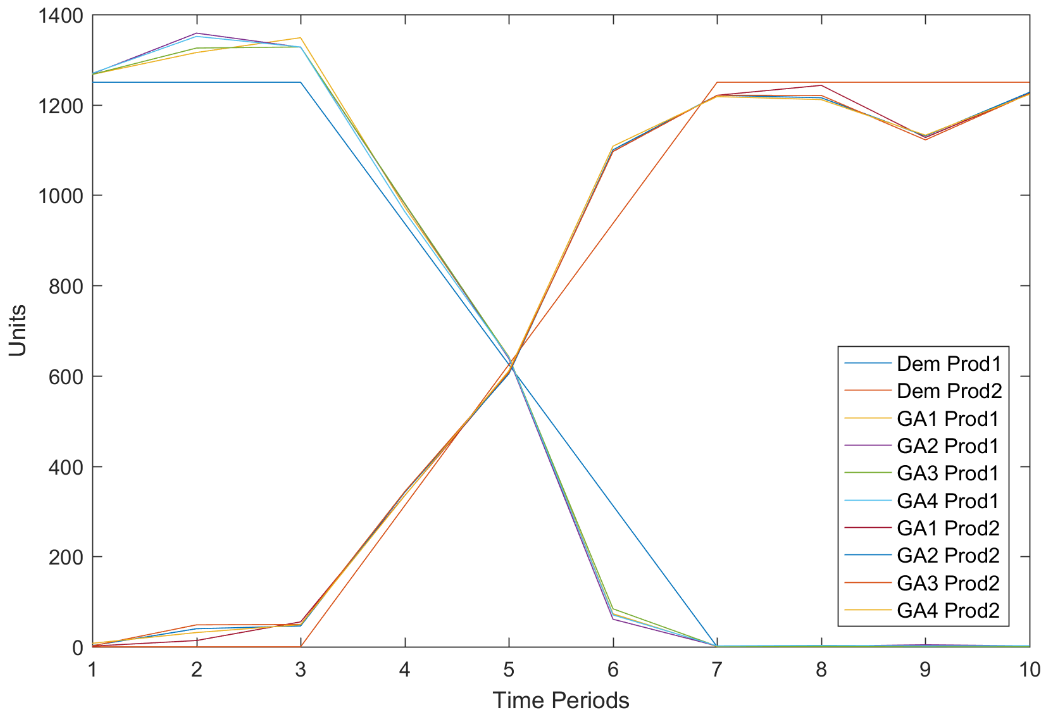

Figure 6: Product Releases for Fast Learning Experiment.

learning effects. However, initial computational results show that the solutions obtained from these models exhibit a structure very similar to those reported for the simulation optimization approach in this paper.

Another important direction for future work is to develop a clear understanding of the different types of learning mechanisms at work in an operating wafer fab. Some learning takes place by repetition, as in the classical learning curve models, while some is the result of systematic process improvement activity based on experiments that consume production capacity. This paper has focused on improvements in processing efficiency resulting from learning, while ignoring improvements in yield which have been a major concern of other authors. The development of a unified framework in which to view these issues, as well as mathematical and computational models that can provide practitioners with useful insights, constitutes an important direction for future work.

REFERENCES

Anzanello, M. J., and F. S. Fogliatto. 2011. "Learning Curve Models and Applications: Literature review and Research Directions". *International Journal of Industrial Ergonomics* 41(5):573–583.

Bohn, R. E. 1995. "Noise and Learning in Semiconductor Manufacturing". *Management Science* 41(1):31–42.

Chand, S., H. Moskowitz, A. Novak, I. Rekhi, and G. Sorger. 1996. "Capacity Allocation for Dynamic Process Improvement with Quality and Demand Considerations". *Operations Research* 44(6):964–975.

Dummler, M. A. 2000. "Analysis of the Nonstationary Behavior of a Wafer Fab During Product Mix Changes". In *Proceedings of the 2000 Winter Simulation Conference*, edited by J. A. Joines et al., Volume 2, 1436–1442. Orlando, Florida: IEEE.

Filho, M. G., and R. Uzsoy. 2013. "The Impact of Simultaneous Continuous Improvement in Setup Time and Repair Time on Manufacturing Cycle Times Under Uncertain Conditions". *International Journal of Production Research* 51(2):447–464.

Fine, C. H. 1986. "Quality Improvement and Learning in Productive Systems". *Management Science* 32(10):1301–1315.

Fine, C. H. 1988. "A Quality Control Model with Learning Effects". *Operations Research* 36(3):437–444.

Fine, C. H., and E. L. Porteus. 1989. "Dynamic Process Improvement". *Operations Research* 37(4):580–591.

Gilenson, M., M. Hassoun, and L. Yedidsion. 2015. "Setting Defect Charts Control Limits to Balance Cycle Time and Yield for a Tandem Production Line". *Computers & Operations Research* 53:301–308.

Gruber, H. 1992. "The Learning Curve in the Production of Semiconductor Memory Chips". *Applied Economics* 24(8):885–894.

Gruber, H. 1994. "The Yield Factor and the Learning Curve in Semiconductor Production". *Applied Economics* 26(8):837–843.

Haller, M., A. Peikert, and J. Thoma. 2003. "Cycle Time Management During Production Ramp-up". *Robotics and Computer-Integrated Manufacturing* 19(1-2):183–188.

Kalir, A. A., and S. C. Sarin. 1999. "The Role of Advanced Start and Dominance Rules in Simulated Annealing for Parallel Processor Scheduling Problems". *Production Planning & Control* 10(8):757–766.

Kim, S., and R. Uzsoy. 2013. "Modeling and Analysis of Integrated Planning of Production and Engineering Process Improvement". *IEEE Transactions on Semiconductor Manufacturing* 26(3):414–422.

Kim, S., and R. M. Uzsoy. 2008. "Integrated Planning of Production and Engineering Process Improvement". *IEEE Transactions on Semiconductor Manufacturing* 21(3):390–398.

Klein, M., and A. Kalir. 2006. "A Full Factory Transient Simulation Model for the Analysis of Expected Performance in a Transition Period". In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. F. Perrone et al., 1836–1839. Monterey, California: IEEE.

Leachman, R. C. 1993. "The Competitive Semiconductor Manufacturing Survey". In *International Symposium on Semiconductor Manufacturing*, 359–381. IEEE.

Leachman, R. C., and S. Ding. 2007. "Integration of Speed Economics into Decision-Making for Manufacturing Management". *International Journal of Production Economics* 107(1):39–55.

Leachman, R. C., J. Kang, and V. Lin. 2002. "SLIM: Short cycle Time and Low Inventory in Manufacturing at Samsung Electronics". *Interfaces* 32(1):61–77.

Liao, W. M. 1979. "Effects of Learning on Resource Allocation Decisions". *Decision Sciences* 10(1):116–125.

Macher, J. T., and D. C. Mowery. 2003. "Managing Learning by Doing: An Empirical Study in Semiconductor Manufacturing". *Journal of Product Innovation Management* 20(5):391–410.

Manda, A. B., and R. Uzsoy. 2018. "Optimizing Releases in New Product Introductions". Technical report, Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, Raleigh, North Carolina.

Manda, A. B., R. Uzsoy, K. G. Kempf, and S. Kim. 2016. "Modeling the Impact of New Product Introduction on the Output of Semiconductor Wafer Fabrication Facilities". In *Proceedings of the 2016 Winter Simulation Conference*, edited by P. I. Frazier et al., 2547–2558. Arlington, Virginia: IEEE Press.

Nemoto, K., E. Akcali, and R. M. Uzsoy. 2000. "Quantifying the Benefits of Cycle Time Reduction in Semiconductor Wafer Fabrication". *IEEE Transactions on Electronics Packaging Manufacturing* 23(1):39–47.

Reeves, G. R., and J. R. Sweigart. 1981. "Product-mix Models When Learning Effects are Present". *Management Science* 27(2):204–212.

Storer, R. H., S. D. Wu, and R. Vaccari. 1992. "New Search Spaces for Sequencing Problems with Application to Job Shop Scheduling". *Management Science* 38(10):1495–1509.

Terwiesch, C., and R. E. Bohn. 2001. "Learning and Process Improvement During Production Ramp-up". *International Journal of Production Economics* 70(1):1–19.

Tirkel, I. 2013. "Yield Learning Curve Models in Semiconductor Manufacturing". *IEEE Transactions On Semiconductor Manufacturing* 26(4):564–571.

Tirkel, I., G. Rabinowitz, D. Price, and D. Sutherland. 2016. "Wafer Fabrication Yield Learning and Cost Analysis Based on In-line Inspection". *International Journal of Production Research* 54(12):3578–3590.

Yelle, L. E. 1979. "The Learning Curve: Historical Review and Comprehensive Survey". *Decision Sciences* 10(2):302–328.

Ziarnetzky, T., L. Mönch, T. Ponsignon, and H. Ehm. 2017. "Rolling Horizon Planning with Engineering Activities in Semiconductor Supply Chains". In *Automation Science and Engineering (CASE), 2017 13th IEEE Conference on*, 1024–1025. IEEE.

AUTHOR BIOGRAPHIES

ATCHYUTA BHARADWAJ MANDA is a Doctoral student in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. He holds a Masters in Industrial and Systems Engineering from North Carolina State University, and a Bachelor of Technology degree in Mechanical Engineering from GITAM University. His research interests are in production systems, supply chain management and new product introduction, focusing on data oriented decision analysis using stochastic simulation. His e-mail address is amanda@ncsu.edu.

REHA UZSOY is Clifton A. Anderson Distinguished Professor in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. He holds BS degrees in Industrial Engineering and Mathematics and an MS in Industrial Engineering from Bogazici University, Istanbul, Turkey. He received his Ph.D. in Industrial and Systems Engineering in 1990 from the University of Florida. His teaching and research interests are in production planning, scheduling, and supply chain management. He was named a Fellow of the Institute of Industrial Engineers in 2005, Outstanding Young Industrial Engineer in Education in 1997, and has received awards for both undergraduate and graduate teaching. His email address is ruzsoy@ncsu.edu.