Proceedings of the 2011 Winter Simulation Conference
S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, eds.

PRODUCTION PLANNING FOR SEMICONDUCTOR MANUFACTURING VIA SIMULATION OPTIMIZATION

Jingang Liu
Chihui Li
Feng Yang

West Virginia University
P.O.Box 6070
Morgantown, WV 26505, USA

Hong Wan

School of Industrial Engineering
Purdue University
West Lafayette, IN 47907, USA

Reha Uzsoy

North Carolina State University
Campus Box 7906
Raleigh, NC 27695-7906, USA

ABSTRACT

This paper is concerned with production planning in manufacturing, which can be loosely defined as the problem of finding a release plan for jobs that minimizes the total cost (or maximizes the total profit). Production planning is a challenging optimization problem due to the variability in manufacturing systems and uncertainty in future demand, both of which have not been adequately addressed by existing production planning models. To address both these issues, this paper formulates the production planning problem as a simulation-based multi-objective optimization problem, and adapts a genetic algorithm to search for a set of release plans that are near-Pareto optimal. The solutions from the simulation optimization approach can serve as a useful benchmark for existing and new production planning methods.

1 INTRODUCTION

This paper is concerned with production planning in manufacturing, which can be loosely defined as the problem of finding a release schedule of jobs into the facility over time so that the actual outputs over time satisfy, as closely as possible, the predetermined requirements (Missbauer and Uzsoy 2010). Clearly, production planning is an optimization problem with the decision variables being the quantity of jobs released into the system for processing over a planning horizon, and the objective is usually to minimize the total cost (or sometimes profit) associated with a release plan. The total cost typically includes the holding cost for finished goods and work in process inventories (WIP), penalty costs for failing to satisfy customer demand in time, and raw material cost.

Evaluating the quality of a given release plan is challenging due to the variability (or uncertainty) involved in the manufacturing planning environment. As pointed out by Cheng (1987), two sources of variability are inherent in any production planning problem: fluctuation in product demand, and variation in operation work contents due to causes such as random machine failures. The presence of such variability leads to two major difficulties in evaluating a release plan. First, for a given release plan, the total cost incurred over the planning horizon is a random variable, and hence, a thorough evaluation of a plan needs to be made based on the distribution of the total cost, or at least on the mean and variance of the cost which are considered as the two fundamental distribution characteristics. Note that it is important to take into account not only the mean but also the variance of the cost, since the latter is a measure of the risk associated with a release plan.

Second, it is difficult to quantify the dependence of the mean/variance of the total cost upon the release plan. The cost depends on the output performance of the manufacturing system over the planning horizon: the quantity of products produced and the WIP (i.e., the number of jobs in the system). Due to the queueing effects resulting from the system variability, these outputs are stochastic processes whose distribution evolves over time depending on the ongoing release of jobs and the initial status of the system (e.g., the WIP and inventory levels at the beginning of the planning period). For a realistic manufacturing system, the relationships between release plans and stochastic outputs are nonlinear and time-dependent, and cannot as yet be described accurately by analytical methods. At present simulation appears to be the only approach that is able to provide a good approximation of the stochastic input-output relationships based on numerical instances.

The existing research on production planning has not adequately addressed these two difficulties, i.e., the risk assessment of a production plan and the quantification of the input-output relationships for manufacturing systems. Most production planning models, including the widely used Material Requirements Planning (MRP) procedure (Baker 1993; Vollmann 1988) and most mathematical programming models (e.g., Hackman and Leachman 1989; Johnson and Montgomery 1974; Leachman 1994), disregard the stochastic nature of the problem and do not even consider the dependence of a system's outputs upon the input release (or the system workload). The relationship between releases and output is defined by lead times that are specified as exogenous parameters independent of workload, contradicting results from queuing models and industrial practice. Ignoring this fundamental relationship between workload, as determined by releases, and output may well lead to inferior production plans (Byrne and Bakir 1999; Pahl et al. 2005). In recent years, there has been a growing interest in exploring this input-output relationship in production planning (Pahl et al. 2005). One approach within this new stream of work is to integrate deterministic mathematical programming with computer simulation. A number of researchers (Hung and Leachman 1996; Byrne and Bakir 1999; Kim and Kim 2001; and Byrne and Hossain 2005) have combined simulation and mathematical programming models in an iterative scheme to evaluate the effects of the release decisions upon the performance of a manufacturing system. Other researchers (Asmundsson et al. 2006; Asmundsson et al. 2009) have proposed the clearing function (CF) methods, in which the CFs are first estimated from simulation and then incorporated in the optimization model as constraints describing the dependence of output upon expected WIP, which, in turn, is defined by the release rates of jobs. In the linear/nonlinear programming models of this more recent work, simulation is used to refine or estimate static constraints, providing a deterministic approximation to the system's time-dependent input-output relationships. A number of authors (Peters et al. 1977; Sen and Higle 1999; Aouam and Uzsoy 2011) have proposed stochastic programming models with recourse, while Higle and Kempf (2010) propose chance-constrained programming models. These models explicitly address the stochastic nature of the problem, but involve different levels of approximations whose effect on solution quality is still under investigation. Also, to the best of our knowledge, no work in the literature has considered the risk factor (or variance of cost) in production planning.

This paper addresses the drawbacks of existing production planning models by adapting and applying the simulation optimization approach, which incorporates discrete-event simulation as an objective evaluator in an optimization routine (Chapters 19-21 in Henderson and Nelson 2006). The simulation optimization (SO) approach allows for the evaluation of a release plan based on not only the mean but also the variance of the total cost incurred by that plan. By using simulation to directly estimate the system outputs (and subsequently the cost objective) for a given release plan, the time-dependent input-output relationship is quantified in detail, subject to modelling error in the simulation model.

Admittedly, the SO method may be very time consuming due to the possibly large amount of discrete-event simulation required. Thus, the use of SO in the practice of production planning may require parallel machines to meet the computational loads and deliver a timely decision. However, the SO method is of substantial research interest in that it provides "true" optimum production plans, which serve as a useful benchmark to evaluate other production planning methods. Currently, the production plans resulting from the existing models can only be evaluated for feasibility or compared with each other for relative superiority. Due to the advantages discussed above, the SO method is expected to lead to better plans than the existing methods, and hence can provide a benchmark evaluation of the existing and future production planning models.

As our first step in approaching the production planning problem with the SO method, we have completed two tasks in this paper.

- Build a C++ discrete-event simulation model representing a scaled-down semiconductor wafer fabrication (fab) system. The initial condition of the simulation can be specified to reflect the status of the real system at the time when planning decisions are to be made. The C++ model is integrated into the optimization routinee as a sampling tool.
- Adapt a deterministic multi-objective heuristic to solve the SO problem for production planning. The heuristic was applied to the scaled-down wafer fab model, and approximately Pareto "optimal" production plans were obtained under two different demand scenarios.

2 OVERVIEW OF THE SIMULATION OPTIMIZATION APPROACH

The problem of optimum production planning is stochastic by nature due to the variability inherent in the decision environment. For a release (or input) plan over a time horizon, the outputs of a manufacturing system are stochastic processes that evolve with time. The total cost associated with a plan is a function of the system outputs and the pre-specified random demand over the planning horizon, and hence is a random variable. Therefore, to evaluate a release plan, the performance measures of interest should include both the expectation and variance of the total cost. Since analytical expressions for the relationships between the input decision and the performance measures are unavailable, SO is a promising alternative to approach the optimum planning of production activities.

SO provides an optimization framework within which candidate solutions are systematically generated and evaluated in terms of performance measures estimated by simulation experiments. With the wide use of simulation in industry, SO is gaining increased popularity, and nowadays most discrete-event simulation packages include some form of "optimization" routine. Compared to deterministic optimization, SO involves an additional complication due to the stochastic nature of the problem that it intends to solve: In SO, the performance measures of a particular decision setting (e.g., a production plan) cannot be evaluated exactly, but instead must be estimated and are thus subject to stochastic noise/errors. Such uncertainty makes it important to address the stochastic convergence of solutions (Hong and Nelson 2006). However, most algorithms used in commercial software are heuristics, such as tabu search and scatter search in OptQuest (OptTek Systems, Inc.); they do not address the stochastic nature of the performance estimates obtained from simulation, and provide no convergence or "correct selection" guarantees. Such heuristics typically evaluate the objective function by averaging over a small (often fixed) number of replications, and then treat the average as deterministic (Xu et al. 2010).

Optimum production planning for semiconductor manufacturing typically involves a relatively large number of decision variables, usuallt into the hundreds if not thousands. We attempted to use theoretically sound algorithms such as that developed by Xu et al. (2010) to solve the problem, but were not able to obtain solutions within a reasonable amount of time (say, a week or so). Hence in this preliminary study we have adapted and applied a multi-objective optimization (MOO) heuristic to solve the planning problem. The two objectives of the optimization problem are minimizing the mean and variance of the total cost.

The remainder of the paper is organized as follows. Section 3 describes the simulation model of the scale-down wafer fab investigated in this work. In Section 4, the formulation of the optimization problem for production planning is given. In Section 5, the MOO heuristic is adapted and applied to solve the production planning problem . Section 6 presents the empirical results obtained from applying the SO method for optimum production planning. Section 7 gives a brief summary.

3 SIMULATION MODEL

We consider the scale-down semiconductor wafer fab system described in Kayton et al. (1997), which has been used in the literature (e.g., Asmundsson et al. 2009; Irdem et al. 2010) to evaluate a variety of production planning and scheduling methods. The system consists of 11 workstations, and involves the representative features of real semiconductor fabs such as machine failures, re-entrant flows and batch processing.

The system is designed to process three types of wafers (products) with each type following a distinct route. A wafer is considered as an individual unit of processing materials that passes through a wafer fab. Product 1 requires twenty-two operations, Product 2 fourteen operations, and Product 3 fourteen operations. All the products are released into the fab in a fixed lot size of 50. A lot is a number of wafers of the same type that are processed and transported between workstations as a unit. Following the common practice of production planning literature, the system is treated as a push system (Hopp 2007), and thus the release rate of jobs (e.g., the weekly wafer starts) can be controlled by production managers.

Using Microsoft Visual Studio C++, we built the simulation model that represents the scale-down wafer fab. The C++ model was verified using techniques recommended in Law and Kelton (2000), such as running the model with simplified assumptions to detect logical mistakes and testing the model outputs under a variety of input settings. Also, the C++ model was validated against the simulation model built in Arena by Irdem et al. (2010). The outputs from the C++ and Arena models are compared for a wide set of inputs. We chose to construct the simulation model in C++ as opposed to using the existing Arena model so that the simulation can be integrated as a sampling tool into the MOO algorithm, which is developed in Matlab.

4 FORMULATION OF THE PRODUCTION PLANNING PROBLEM

We formulate the optimum production planning as a multi-objective optimization problem. Denote the planning time horizon as $(0,H]$. Following the existing production planning models, we divide the entire horizon into T equal-length time periods. For convenience of discussion, the following notations are given.
Indices:

k: Product index.
t: Planning period index.

Parameters:

K: Number of different types of products/wafers in the manufacturing system.
T: Number of time periods within the planning horizon $(0,H]$.
ω_{kt}: Unit WIP holding cost per period for product of type k in period t.
h_{kt}: Unit inventory holding cost per period for product of type k in period t.
b_{kt}: Unit backlogging cost for product type k in period t.
D_{kt}: Demand for type k product in period t; demand is considered as a random variable that follows a pre-specified distribution obtained from forecasting models which are outside the scope of this work. We assume that all demands are realized at the end of each planning period.

Independent decision variables:

x_{kt}: Number of lots of type k products released in period t. We assume that lots are released into the system with constant inter-release times. Denote $\mathbf{x} = \{x_{kt}; k = 1,2,\dots,K; t = 1,2,\dots,T\}$ as the decision vector to be determined in the planning optimization.

Dependent random variables:

Z_{kt}: Number of lots of product type k products produced in period t.

W_{kt}: Cumulative WIP (in number of lots) of type k products in period t. Denote $Q_k(\tau)$ as the WIP level of type k products at a time instant τ ($-\infty < \tau < \infty$). Since the WIP level jumps up or down upon the arrival or departure of products, $Q_k(\tau)$ is a piece-wise constant function of the continuous time τ. Let $(s_t, e_t]$ be the interval of the t^{th} time period, then we have

$$W_{kt} = \int_{s_t}^{e_t} Q_k(\tau) d\tau.$$

I_{kt}: Cumulative inventory level of type k end products at the end of period t. Let I_{k0} represent the initial inventory of type k finished goods, and $V_k(\tau)$ the inventory of product k at time instant τ. Note that $V_k(\tau)$ is also a piece-wise constant function of time τ. Then the cumulative inventory I_{kt} can be written as:

$$I_{kt} = \int_{s_t}^{e_t} V_k(\tau) d\tau.$$

The inventory level at the beginning of each period is given as:

$$\begin{aligned} V_k(s_1) &= I_{k0} \\ V_k(s_t) &= \max\{0, V_k(s_{t-1}) + Z_{k,t-1} - B_{k,t-1} - D_{k,t-1}\}; \;\; t = 1, 2, \ldots, T \end{aligned}$$

where $B_{k,t-1}$ is defined as follows.

B_{kt}: Quantity of type k products that cannot be satisfied on time at the end of planning period t. The unmet demand is considered as backlog, and will be fulfilled at the end of nearest planning period(s). We have:

$$B_{kt} = -\min\{0, V_k(s_t) + Z_{k,t} - B_{k,t-1} - D_{k,t}\}; \;\; t = 1, 2, \ldots, T$$

Given pre-specified product demand and initial system status (including initial WIP, inventory levels, backlog quantities, etc.), all the dependent variables can be evaluated through simulation with a specified release plan.

It is worth pointing out that for the sake of solving the optimization problem in a realistic setting, it is necessary to consider time periods beyond the planning horizon. In this work, we follow the approach of Hung and Leachman (1996), and assume that after the planning horizon demands for each product type will continue forever following the rate of the last planning period. An extra post-planning period is added, during which the demands are satisfied by the products that are released but not finished during the planning horizon. This extra period lasts until all the products released during the planning horizon have been completed.

The total cost incurred by the production plan $\mathbf{x}$ can be written as

$$L(\mathbf{x}) = \sum_{k=1}^{K} \sum_{t=1}^{T+1} \omega_{kt} W_{kt} + \sum_{k=1}^{K} \sum_{t=1}^{T+1} h_{kt} I_{kt} + \sum_{k=1}^{K} \sum_{t=1}^{T} b_{kt} B_{kt},$$

which consists of three parts: the WIP holding cost, the inventory holding cost, and the backlog cost. The total cost $L(\mathbf{x})$ is a random variable, and the purpose of production planning is to find a release plan that excels in terms of both the expectation and variance of $L(\mathbf{x})$. Hence, we formulate the production planning problem as the following MOO problem:

$$\begin{aligned} \min \quad & \mathrm{E}[L(\mathbf{x})] \\ \min \quad & \mathrm{Var}[L(\mathbf{x})] \\ \text{Subject to:} \quad & x_k^L \le x_{kt} \le x_k^U; \;\; k = 1, 2, \ldots, K; \; t = 1, 2, \ldots, T \end{aligned} \tag{1}$$

The lower and upper bounds, x_k^L and x_k^U, are parameters specified by the user.

For an MOO problem, there generally does not exist any single solution which can provide the optimal value on all the objectives, and thus it is of interest to generate a set of non-dominated solutions, where no objective can be improved without worsening at least one other objective. The set of all non-dominated solutions is referred to as the Pareto optimal front (Deb 2009), and our goal is to obtain a set of solutions for $\mathbf{x}$ as close as possible to the Pareto optimal front.

5 THE MULTI-OBJECTIVE OPTIMIZATION PROCEDURE

5.1 The Multi-Objective Optimization Algorithm

To solve the problem (1), we adapted the MOO function "gamultiobj" provided by Matlab. The function gamultiobj uses a controlled elitist genetic algorithm, which is a variant of the elitism non-dominated sorting GA (NSGA-II). Due to its various appealing features, the NSGA-II is a most widely used algorithm for MOO problems (see Deb 2009 for details). Since the gamultiobj function is developed for deterministic optimization, we adapted it in a straightforward way to accommodate the stochastic simulation involved in the simulation optimization procedure.

The major change made to the gamultiobj function is to include a two-stage process which determines for each candidate solution (release plan) the number of simulation replications required to evaluate the cost objectives associated with that candidate to a desired degree of statistical precision. As opposed to deterministic optimization where only one replication is needed to evaluate the objective values for a candidate solution, the number of simulation replications required to solve (1) may well vary on a candidate-by-candidate basis. In the adapted algorithm, for each newly generated solution $\mathbf{x}^*$, the following two-stage process is used to determine the number of replications required to obtain an expected cost estimate $\widehat{\mathrm{E}}[L(\mathbf{x}^*)]$ of the desired precision. Suppose that the precision is measured by the variance of $\widehat{\mathrm{E}}[L(\mathbf{x}^*)]$, and the desired variance is pre-specified as σ^2. At the first stage, a conservative number, say n_0, independent replications are generated. From the n_0 cost realizations, the sample variance of the cost is estimated as σ_0^2. The number of replications that is likely to provide a desired variance σ^2 for $\widehat{\mathrm{E}}[L(\mathbf{x}^*)]$ is then set as $\lceil \sigma_0^2/\sigma^2 \rceil$. At the second stage, $\lceil \sigma_0^2/\sigma^2 \rceil - n_0$ additional replications are performed. The data obtained from both stages are then used to evaluate the objectives, the mean and variance of the total cost.

5.2 Simulation-Based Objective Evaluation

As mentioned above, for a candidate release plan $\mathbf{x}^*$, multiple simulation replications will be performed with each one leading to an instance of the associated total cost $L(\mathbf{x})$. Averaging across those replications will lead to an estimate for $\mathrm{E}[L(\mathbf{x}^*)]$, denoted as $\widehat{\mathrm{E}}[L(\mathbf{x}^*)]$; an estimate of the variance $\mathrm{Var}[L(\mathbf{x})]$ can also be obtained, which is denoted as $\widehat{\mathrm{Var}}[L(\mathbf{x})]$.

Next, we detail how a simulation experiment is performed under a plan $\mathbf{x}^*$ to obtain an instance of the total cost. (i) First, the initial status of the simulation model is specified to reflect the status of the manufacturing system at the beginning of the planning horizon, when the planning is being performed. For each replication, the simulation model is initialized with certain WIP levels at different workstations, the completion time for the lots currently in process at all the stations, the last failure time of each machine, etc. Hence each simulation starts from a state consistent with the real system being considered. (ii) Second, under the plan $\mathbf{x}^*$, the model is simulated for the time length equal to the planning horizon, which is 12 weeks in the experiments considered in Section 6, plus the extra period mentioned in Section 4. The $(T+1)^{th}$ additional period is determined in such a way that the products released during the planning horizon will complete processing by the end of the extra period. Based on the range of the total time it takes for a wafer lot to go through the system, the length of the extra period is set to three weeks in our experiments. (iii) Third, during the discrete-event simulation process, the system outputs $\{Y_{kt}; W_{kt}; k = 1,\ldots,K; t = 1,\ldots,T+1\}$ are collected. For each replication, by generating a demand realization $\{D_{kt}; k = 1,\ldots,K; t = 1,\ldots,T+1\}$

Table 1: The mean demand $\{d_{kt}; k = 1,2,3; t = 1,2,\ldots,12\}$. (The demands are given in terms of lots per week.)

Period	1	2	3	4	5	6	7	8	9	10	11	12
Prod1	23	26	29	31	34	37	40	43	45	48	51	54
Prod2	66	62	58	54	49	45	41	37	33	28	24	20
Prod3	16	16	15	15	15	15	14	14	14	14	13	13

from the given random demand distributions (Section 6.1), $\{I_{kt}; B_{kt}; k = 1,\ldots,K; t = 1,\ldots,T+1\}$ can also be obtained.

Therefore, for a certain release plan $\mathbf{x}^*$, each simulation run leads to a possible value of the total cost. With multiple replications, the mean and variance (and possibly the distribution) of the cost can be estimated, and these statistics represent the variability in the manufacturing system as well as the demand uncertainty.

6 EMPIRICAL RESULTS

The multi-objective simulation optimization procedure was applied to solve the production planning problem for the scale-down semiconductor wafer fab described in Section 3. The parameters and variables defined in Section 4 are specified as follows. The planning horizon considered is 12 weeks, with each time period being one week long and $T = 12$ periods within the planning horizon. The extra planning period is set as 3 weeks. It is assumed that the cost parameters are time-independent and product-independent, and hence, for any time period and any product type, we have $\omega_{kt} = 7$ for the unit WIP holding cost per time unit, $h_{kt} = 15$ for the unit inventory holding cost per time unit, and $b_{kt} = 35$ for the backlog cost.

There are $K = 3$ different types of wafers processed in the system, so the release plan $\mathbf{x}$ is a 36×1 decision vector. The constraints in (1), i.e., the upper and lower bounds on the release rates are given as: $x_L^k = 10$ lots/week and $x_U^k = 80$ lots/week. If all three types of products are released at their lower bound rates, the system utilization is 0.3. If one product is released at its upper bound rate while the other two products are released at their lower bound rates, the system utilization is about 1.2. Hence, the system is allowed to be temporarily overloaded in our production planning problem.

The forecasted demand over the planning horizon is specified in Section 6.1.

6.1 Prespecified Demand

It has been explained in Section 5.2 that for each simulation replication, an instance of the total cost is obtained with a realization of the random demand. In the simulation optimization procedure, the random demand is allowed to follow any type of distribution, among which the most sophisticated distributions may be those represented by multivariate time-series models. To generate demands that can be modeled as general multivariate time-series, please refer to the algorithms developed by Biller and Nelson (2003). In this paper, two different demand scenarios are considered.

Case 1

In the first case, it is assumed that the customer demand over the planning horizon is deterministic and given with certainty as shown in Table 1. As can be seen from Table 1, over the planning horizon, the demand for product 1 increases gradually, the demand for product 2 decreases, and the demand for product 3 remains at a low, stable level. If the wafer fab is operated at a release rate equal to the the demand rate in Table 1, then the system utilization is about 80% throughout the planning horizon.

Case 2

In the second case, the product demand is random and assumed to follow a normal distribution. Specifically, denote the random demand vector as $\{D_{kt}; k = 1,2,3; t = 1,2,\ldots,12\}$, and for product k in period t, D_{kt} is normally distributed with a mean of d_{kt} as given in Table 1, and a standard deviation of $0.1 \times d_{kt}$.

The demands across different products and time periods are assumed independent. In each simulation replication, a realization of the random vector $\{D_{kt}; k = 1,2,3; t = 1,2,\ldots,12\}$ is generated and used to evaluate the total cost of a release plan.

6.2 Optimization Results

Figure 1 shows the objective performance (in terms of the mean and standard deviation of the total cost) of the two set of solutions obtained for the two demand scenarios. The fourteen crosses represent the performance estimates of the solutions for case 1, and are labelled as 1.1 to 1.14. The nine circles represent the performance estimates of the solutions for case 2, and are labelled as 2.1 to 2.9. As can be seen from the Figure 1, a set of non-dominated solutions were obtained for both cases. For each case, the solutions (or release plans) vary over a fairly wide range in terms of the mean and standard deviation of the cost. Compared to the solutions obtained from single-objective optimization, the results (as those presented in Figure 1) from solving our MOO problem provide more complete information to decision makers who have to weigh the trade-off between average cost and the risk associated with that cost.

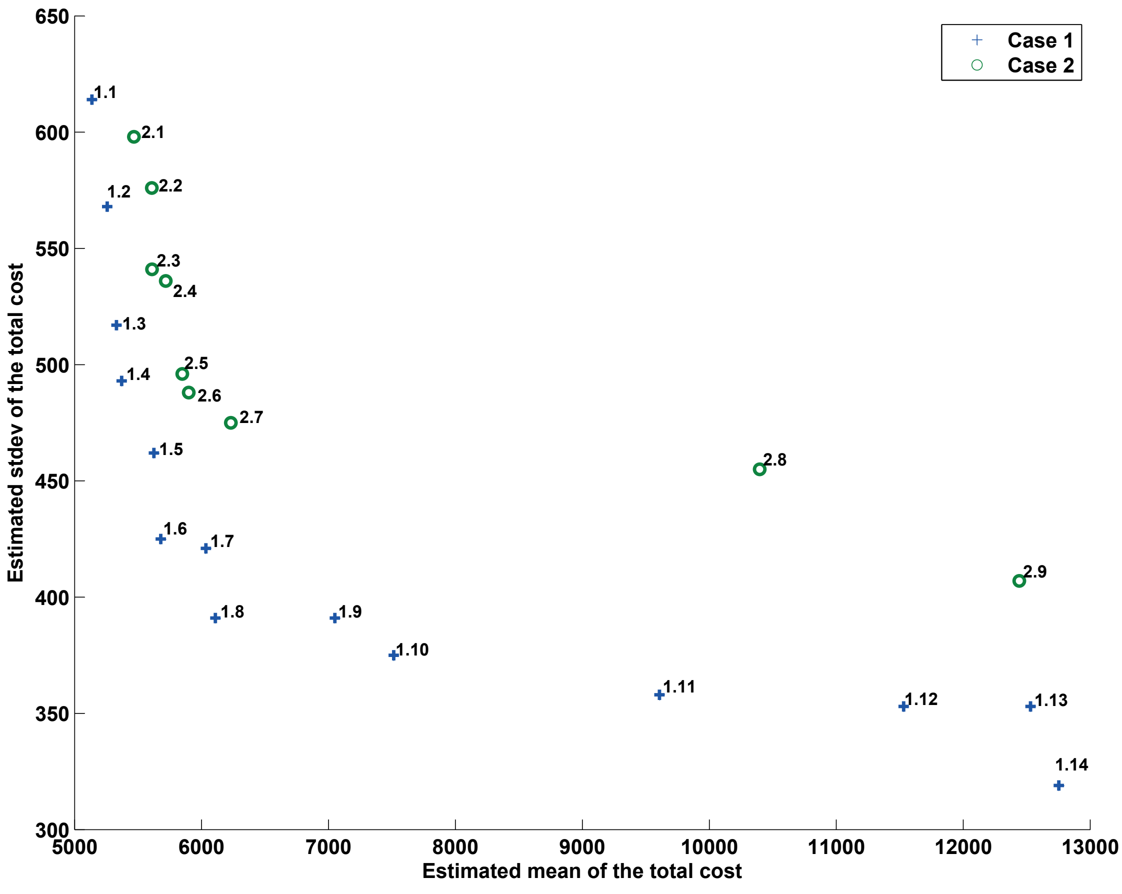

Figure 1: Objective performance of the solutions obtained for the two demand scenarios.

As expected, the solutions in Case 2 result in a higher variability in the total cost compared to those in Case 1. For instance, solutions 1.5 and 2.3 have almost the same mean cost, but the standard deviation of solution 2.3 is substantially higher than that of solution 1.5.

7 SUMMARY

This work is a preliminary step toward approaching the production planning problem with simulation optimization methods. The optimum production planning is formulated as a multi-objective optimization (MOO) problem, and the two objectives considered simultaneously are minimizing the mean and variance of the total cost associated with a production plan. A heuristic algorithm is adapted to solve the MOO problem due to the high-dimensional decision variables and high computational time required by simulation. The solutions resulting from the heuristics can serve as the starting points for asymptotically locally convergent algorithms, or can be further investigated through ranking and selection procedures.

ACKNOWLEDGMENTS

This research was supported by National Science Foundation Grant CMMI-1068131.

REFERENCES

Aouam, T. and R. Uzsoy. 2011. "An Exploratory Analysis of Production Planning in the face of Stochastic Demand and Workload-Dependent Lead Times". In *Decision Policies for Production Networks*, edited by K. G. Kempf and D. Armbruster. Boston: Springer.

Asmundsson, J. M., R. L. Rardin and R. Uzsoy. 2006. "Tractable Nonlinear Production Planning Models for Semiconductor Wafer Fabrication Facilities". *IEEE Transactions on Semiconductor Manufacturing* 19 (1): 95–111.

Asmundsson, J., R. L. Rardin, and R. Uzsoy. 2009. "Production Planning Models for Resources Subject to Congestion". *Naval Research Logistics* 56: 142–157.

Baker, K. R. 1993. "Requirements Planning". In: *Handbooks in Operations Research and Management Science: Logistics of Production and Inventory*, edited by S. C. Graves, A. H. G. Rinnooy Kan and P. H. Zipkin. Amsterdam, Netherlands: Elsevier Science.

Byrne, M. D. and M. A. Bakir. 1999. "Production Planning Using a Hybrid Simulation–Analytical Approach". *International Journal of Production Economics* 59 (1-3): 305–311.

Byrne, M. D. and M. M. Hossain. 2005. "Production Planning: An Improved Hybrid Approach". *International Journal of Production Economics* 93-94 (1): 225–229.

Cheng, T. C. E. 1987. "A Simulation Study of MRP Capacity Planning with Uncertain Operation Times". *International Journal of Production Research.* 25 (2): 245–258.

Deb, K. 2009. *Multi-Objective Optimization using Evolutionary Algorithms*. New York: John Wiley & Sons.

Hackman, S. T. and R. C. Leachman. 1989. "A General Framework for Modeling Production". *Management Science* 35 (4): 478–495.

Henderson, S. G. and B. L. Nelson. 2006. *Handbooks in Operations Research and Management Science: Simulation*. Amsterdam, Netherlands: Elsevier Science.

Higle, J. L. and K. G. Kempf. 2010. "Production Planning under Supply and Demand Uncertainty: A Stochastic Programming Approach". In: *Stochastic Programming: The State of the Art*, edited by G. Infanger. Berlin: Springer.

Hong, L. J. and B. L. Nelson. 2006. "Discrete Optimization via Simulation using COMPASS". *Operations Research* 54: 115–129.

Hung, Y. F. and R. C. Leachman. 1996. "A Production Planning Methodology for Semiconductor Manufacturing Based on Iterative Simulation and Linear Programming Calculations". *IEEE Transactions on Semiconductor Manufacturing* 9 (2): 257–269.

Johnson, L. A. and D. C. Montgomery. 1974. *Operations Research in Production Planning, Scheduling and Inventory Control*. New York: John Wiley & Sons.

Irdem, D. F., N. B. Kacar and R. Uzsoy. 2010. "An Exploratory Analysis of Two Iterative Linear Programming-Simulation Approaches for Production Planning". *IEEE Transactions on Semiconductor Manufacturing* 23 (3): 442–455.

Kayton, D., T. Teyner, C. Schwartz and R. Uzsoy. 1997. "Focusing Maintenance Improvement Efforts in a Wafer Fabrication Facility Operating under Theory of Constraints". *Production and Inventory Management Journal* 38 (4): 51–57.

Kim, B. and S. Kim. 2001. "Extended Model for a Hybrid Production Planning Approach". *International Journal of Production Economics* 73 (2): 165–173.

Leachman, R. C. 1994. "Modeling Techniques for Automated Production Planning in the Semiconductor Industry". In: *Optimization in Industry*, edited by T. Ciriani and R. C. Leachman. New York: John Wiley & Sons.

Missbauer, H. and R. Uzsoy. 2010. "Optimization Formulations of Production Planning Problems". In: *Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook*, edited by K. G. Kempf, P. Keskinocak and R. Uzsoy. New York: Springer.

Pahl, J., S. Voss and D. L. Woodruff. 2005. "Production Planning with Load Dependent Lead Times". *4OR: A Quarterly Journal of Operations Research* 3 (4): 257–302.

Peters, R. J., K. Boskma and H. A. E. Kupper. 1977. "Stochastic Programming in Production Planning: A Case with Non-Simple Recourse". *Statistica Neerlandica* 31: 113–126.

Sen, S. and J. L. Higle. 1999. "An introductory tutorial on stochastic linear programming models". *Interfaces* 29 (2): 33–61.

Vollmann, T. E., W. L. Berry and D. C. Whybark. 1988. *Manufacturing Planning and Control Systems*. Massachusetts: Irwin/McGraw-Hill.

Xu, J., L. J. Hong and B. L. Nelson. 2010. "Industrial Strength COMPASS: A Comprehensive Algorithm and Software for Optimization via Simulation". *ACM TOMACS* 20: 1–29.

AUTHOR BIOGRAPHIES

JINGANG LIU is with Cummins, Inc. He received his Ph.D. degree in Industrial and Management Systems Engineering from West Virginia University. He is interested in developing simulation models and statistical methods for manufacturing and energy applications. His e-mail is jingangliu@gmail.com.

CHIHUI LI is a PhD student in the Industrial and Management Systems Engineering Department at West Virginia University. His research interest lies in computer integrated manufacturing and quality and reliability engineering. His e-mail is simonleewin@hotmail.com.

FENG YANG is an Assistant Professor in the Industrial and Management Systems Engineering Department at West Virginia University. Her research interests include simulation and metamodeling, design of experiments, and applied statistics. Her email address is feng.yang@mail.wvu.edu.

HONG WAN is an Associate Professor in the School of Industrial Engineering at Purdue University. Her research interests include design and analysis of simulation experiments, simulation optimization, applied statistics, quality management, and healthcare engineering. She has taught a variety of courses and is a member of INFORMS and ASA. Her email address is hwan@purdue.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. His research has been supported by the National Science Foundation, Intel Corporation,

Hitachi Semiconductor, Harris Corporation, Kimberly Clark, Union Pacific, Ascension Health and General Motors. He was named a Fellow of the Institute of Industrial Engineers in 2005 and 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.