

Simultaneous Determination of Production and Maintenance Schedules Using In-Line Equipment Condition and Yield Information

Track: Operations Planning, Scheduling and Control

Thomas W. Sloan
Department of Management, 414 JB
School of Business Administration
University of Miami
Coral Gables, FL 33124
Tel: 305/284-1086 Fax:305/284-3655
E-mail: tsloan@miami.edu

Abstract

In this paper, we develop a semi-Markov decision process model of a single-machine production system with multiple products and multiple repair actions. The model simultaneously determines maintenance and production schedules, accounting for the fact that equipment condition can affect the yield of different products differently. After examining structural properties, we compare the combined method to the approach often used in practice: a simple threshold maintenance policy and an FCFS dispatching policy. In the more than 11,000 problems studied, the simultaneous approach yielded rewards that were an average of more than eight percent greater than the traditional approach. The results indicate that the improvement increases as the uncertainty of repair and variability of yield increase.

1. Introduction

In many manufacturing environments, equipment condition has a significant impact on product quality. Examples range from the wear of tool bits in a machine shop to the contamination of ultra-clean production environments needed for the production of pharmaceuticals and integrated circuits. This paper explores how information about equipment condition can be used to improve production scheduling decisions and increase yield. Consider, for example, the etch operation in semiconductor wafer fabrication in which chemicals are used to strip materials from the surface of silicon wafers. The inner chambers of the etch equipment become more and more contaminated as wafers are processed, and eventually, one must stop production and clean the machine, returning it to an improved state.

In recent years, the development of *in situ* particle monitors and other devices has made it possible to assess the level of contamination while the equipment is in use. Experts have promoted the use of such devices to improve cleaning schedules and process control (Borden and Larson, 1989). However, little effort has been made to use such information in production scheduling decisions despite the fact that the level of contamination may affect different product types differently. For example, leading-edge computer chips typically have smaller circuit sizes than more mature chips and thus are more susceptible to damage by particles in the equipment.

A similar division exists in the operations management/operations research literature related to this problem. Only recently have researchers begun to explore the interaction between equipment condition and yield, and this research has focused almost exclusively on single-product systems, ignoring situations in which the equipment condition affects different products differently. For a complete review of the relevant literature, refer to Sloan (2001).

In this paper we extend the model proposed by Sloan and Shanthikumar (2000) by developing a semi-Markov decision process model that simultaneously determines the production and maintenance schedules for a single-machine, multiple-product system, accounting for product differences. The equipment condition deteriorates over time, and at each decision epoch, the decision maker can choose to produce one of several products or to take one of several cleaning actions. There is uncertainty associated with each action, and the expected rewards or costs depend on the machine state. In addition, the time between decision epochs is a random variable that depends on which action is taken.

2. Model

We consider the problem of scheduling production and maintenance for a single-machine, multiple-product system. As the equipment is used, the condition deteriorates, and the state of the machine affects the yield of different products to varying degrees. The set of machine states is defined as $S = \{0, 1, \dots, M\}$, indexed by i , where state 0 is the best possible condition and M is the worst condition. At each decision epoch, we observe the machine state and decide whether to produce one of K products, denoted as actions $k = 1, 2, \dots, K$, or to stop production and perform one of several cleaning actions. For notational simplicity, we assume that cleaning actions are limited to “minor” clean or a “major” clean, denoted as actions $K+1$ and $K+2$, respectively, but additional cleaning actions could be easily be incorporated. The set of all actions is defined as $A = \{1, 2, \dots, K+2\}$, and is indexed by a .

Define $X(t)$ as the state of the machine at time t . The time between decision epochs is a random variable, and following the notation of Puterman (1994) we define $F(t | i, a)$ as the probability that the next decision epoch occurs within t time units of the current decision epoch given that the current state is i and action a is taken. The decision process is tied to the production and cleaning cycles, but the machine state may change between decision epochs. Thus, we define $p(j | t, i, a)$ as the probability that the machine occupies state j , t time units after a decision epoch, given that action a was taken in state i . Let X_n denote the machine state and a_n denote the action taken at the n th decision epoch, $n = 1, 2, \dots$. Define $P(j | i, a)$ as the probability that the machine will be state j at the next decision epoch given that the state is i and action a is taken at the current decision epoch. We assume that changes in the machine state depend only on the current state and the action taken.

Under these conditions, $\{X(t), t > 0\}$ is a semi-Markov process, and we call $\{X_n, n = 1, 2, \dots\}$ the embedded Markov chain of the semi-Markov process. Our objective is to determine a policy that maximizes the expected average reward per unit time. Suppose that the process begins at time 0, and let $Z(t)$ denote the total reward earned by time t . The average reward earned under policy π is

$$g_{\pi}(i) = \lim_{t \rightarrow \infty} E_{\pi} [Z(t)/t | X(0) = i]. \quad (1)$$

The expected time until the next decision epoch, given that action a is taken in state i at the current decision epoch is defined as

$$\tau(i,a) = \sum_j P(j | i,a) \int dF(t | i,a) \quad (2)$$

If action a is taken at the current decision epoch while the machine is in state i , a reward of $c(i,a)$ is earned immediately, and a reward of $v(j,a)$ is earned per unit time while the machine is in state j . We use the index j to indicate the possibility that the machine state may change between decision epochs. The expected total reward between this decision epoch and the next is defined as

$$r(i,a) = c(i,a) + \int \sum_j [v(j,a) p(j | u,i,a) du] dF(t | i,a) \quad (3)$$

In this particular application, the reward for producing a product depends on the yield and the yield depends on the machine state. Let R_k be the profit earned for a unit of product k and β_{jk} be the yield of product k when the machine is in state j . When the equipment is cleaned, we assume that the fixed cost and the variable cost rate depend only on the machine state at the time the decision is made. However, the time to perform the cleaning is still a random variable.

We make the following assumptions:

- (A1) There exist $\varepsilon > 0$ and $\delta > 0$ such that $F(\delta | i,a) \leq 1 - \varepsilon$ for all i and all a .
- (A2) $\tau(i,a)$ is bounded.
- (A3) $r(i,a)$ is bounded.

Proposition 1: If Assumptions 1 through 3 hold, then the following optimality equation is satisfied:

$$h(i) = \max_a \{r(i,a) - g\tau(i,a) + \sum_j P(j | i,a) h(j)\}, \quad (4)$$

and an average optimal stationary policy exists.

To further characterize the optimal policy we make several additional assumptions:

- (A4) $r(i,a)$ is nonincreasing in i .
- (A5) $\tau(i,a)$ is nondecreasing in i .
- (A6) $\sum_j P(j | i,a)$ is nonincreasing in i .

A stationary policy is a decision rule that is non-randomized and does not depend on time, i.e., a policy that depends only on the current machine state. Define $a(i)$ as the optimal action taken in state i . The following proposition states that the optimal policy be of the control limit type.

Proposition 2: If Assumptions 1 through 6 hold, then there exists a threshold state, \hat{i} , such that $a(i) > K$ for $i \geq \hat{i}$, and $a(i) \leq K$ for $i < \hat{i}$.

Thus far we have only made assumptions about the ordering of values with respect to the machine state. Now we consider the implications of imposing an order on values with respect to the actions. Specifically, we make the following assumptions:

- (A7) $r(i, a_1) - r(i, a_2)$ is nonincreasing in i for $a_1 \geq a_2$.
- (A8) $\sum_j P(j | i, a_1) - \sum_j P(j | i, a_2)$ is nondecreasing in i for $a_1 \geq a_2$.

We can now state the following proposition, based on a theorem from Kao (1973).

Proposition 3: If Assumptions 1 through 8 hold, then the optimal action, $a(i)$, is nondecreasing in i .

We can find the optimal stationary policy using a linear program (LP), as described by Osaki and Mine (1968), Heyman and Sobel (1984), and Puterman (1994). For economy, we omit the formulation here. In many situations, it is necessary to meet some pre-specified production requirements. Suppose that we must meet some long-term production schedule in the sense that a certain proportion of total production must consist of a particular product type. Specifically, let γ_k be the long-run proportion of product k required. We can then modify the LP to include additional constraints to ensure that the long-run average production requirements are met. We refer to the expanded linear program as LP'.

The combined production and maintenance scheduling model above solves two problems simultaneously. Traditionally, the problems of maintenance scheduling and production scheduling have been solved sequentially. First, one determines the maintenance schedule, ignoring how the machine condition may affect different products differently. Then, given a maintenance schedule, one determines a production schedule. This is often the approach used in industry. We wish to determine how much of a difference it makes to solve the problems simultaneously rather than sequentially and also how much of a difference the dispatching policy makes. We examine three different approaches, each of which solves the two problems in a slightly different way:

Approach 1 (sequential, FCFS dispatch): This is the traditional approach. A cleaning threshold is determined by solving the original LP, and an FCFS dispatching policy is used. The production schedule is determined by solving LP' with the additional constraint that the cleaning threshold be enforced.

Approach 2 (sequential, yield-based dispatch): As with Approach 1, a cleaning threshold is determined by solving the original LP. Next, a production schedule is determined by solving LP' with the additional constraint that the cleaning threshold be enforced.

Approach 3 (simultaneous, yield-based dispatch): The cleaning and production schedules are determined simultaneously by solving LP'.

3. Numerical Results

In this section, we report the results of numerical test problems designed to explore the effectiveness of incorporating more equipment information in production and maintenance scheduling decisions and to study the sensitivity to changes in different parameters of the model.

3.1 Experimental Design

We study two problem sets, one for each set of assumptions. Problem Set 1 incorporates Assumptions 1 through 6, and Set 2 incorporates Assumptions 1 through 8. By examining the two cases, we can get a sense of the importance of the assumptions regarding product ordering.

We examine a single-machine system that produces four products and has five machine states. Two cleaning actions are available: “minor” clean and “major” clean. Thus, we can choose one of six actions at each decision epoch. Since this problem was motivated by an application in semiconductor manufacturing, we make several additional assumptions for our test problems which are appropriate for this context. Below we describe how each of the model parameters is treated and summarize the values used in the test problems. Each set has a total 5,632 problems.

Product mix: The long-run proportion of product k needed is γ_k , where $\sum \gamma_k = 1$. We test 11 different product mixes, ranging from a uniform mix of all products to a mix that requires mostly one product.

Unit profits: Each product has a different unit profit, but we would expect the relative rewards rather than the absolute rewards to make a difference in the decision problem. Thus, a level describes the percent difference between the rewards of adjacent products. For example, a level of 100 percent means that $R_k = 2R_{k+1}$ for $k = 1, 2, 3$. We set $R_4 = \$100$ throughout the experiments, and we test two levels: 100 and 200 percent.

Yield matrix: This matrix indicates the expected yield when producing each product in each state. Probabilities were generated from a beta distribution and sorted to ensure that the assumptions were met. We test two levels: low and high variance. This is the only factor for which the two problem sets differ.

Cleaning costs: We assume that the fixed cleaning cost does not depend on the machine state, and we assume that variable cleaning cost value is constant. Recall that the time to clean the machine is a random variable and does depend on the machine state. We assume that the mean cleaning time increases in i and increases in a . In other words, it takes longer to perform a major clean than a minor clean, and it takes longer to clean as the machine state gets worse. We assume that the variable cleaning cost is fixed at \$10 for both minor and major cleaning for all problems. We test two levels of fixed cleaning costs.

Transition probabilities: We assume that if production is picked, the machine state transitions are not affected by which product is produced, i.e., $P(j | i, k) = P(j | i)$ for all k . This assumption is valid in the context of the etch operation in semiconductor manufacturing because the process itself does not depend on the product. We must also specify the transition probabilities for the

cleaning actions, so there are a total of three matrices for each problem. When production is picked, the level describes the “rate” at which the equipment deteriorates, and we test two levels: slow and fast. When cleaning is picked, the level describes the effectiveness of clean action. We test two levels: low and high, where “low” means that the probability of moving to a better machine state is low, and “high” means that the probability of moving to a better machine state is high.

Transition times: Processing times at the etch operation typically do not vary by product nor are they affected by the equipment condition. So we assume that all products have the same mean processing times and that this mean does not depend on the state. We allow two cleaning actions, minor and major clean. We assume that the cleaning time increases in i and increases in a . In other words, it takes longer to perform a major clean than a minor clean, and it takes longer to clean as the machine state gets worse. We test two levels: low and high.

3.2 Results

Table 1 reports the results of the experiments. The table reports the percentage improvement provided by Approach 3 over Approaches 1 and 2 for each experiment. The table lists summary results for the entire set of 5,632 test problems in each set as well as averages for the subset of problems corresponding to a particular factor level.

See Figure 1

The average reward earned using Approach 3 is an average of 5.0 percent higher than Approach 1 and an average of 2.0 percent higher than Approach 2 for Set 1. For Set 2, the average reward earned using Approach 3 is an average of 12.3 percent higher than Approach 1 and an average of 3.8 percent higher than Approach 2 for Set 1. The average Approach 3/Approach 1 difference for both sets is 8.7 percent, and the average Approach 3/Approach 2 difference is 2.9 percent. These results demonstrate that using equipment condition and yield information to simultaneously determine production and maintenance schedules can substantially improve performance. Even when production schedules are fine-tuned using Approach 2, Approach 3 is still significantly better. The average improvement is bigger for Set 2 than for Set 1, which seems reasonable given the additional structure imposed by Assumptions 7 and 8.

4. Conclusions

In many manufacturing environments, equipment condition has a significant impact on product quality. This paper explored the effects of incorporating equipment condition information into decisions about production and maintenance scheduling. We developed a semi-Markov decision process model that simultaneously determines the maintenance and production schedules for a single-machine, multiple-product production system, accounting for the fact that the equipment condition affects the yield of the different products differently. Production times and cleaning times are random variables, and the effectiveness of cleaning is subject to some uncertainty. Furthermore, constraints are imposed on the (long-run) product mix. Under reasonable assumptions about the product yield and machine state transition probabilities, we showed that a stationary optimal policy would exist. Furthermore, we found sufficient conditions for a

monotone policy, i.e., a policy for which the optimal action is nondecreasing in the machine state.

Motivated by an application in semiconductor manufacturing, we compared the simultaneous approach to the approach often used in practice: a simple threshold maintenance policy and a first-come, first-serve dispatching policy. The results of more than 11,000 test problems indicate that the improvement provided by the simultaneous approach increases as the uncertainty increases — uncertainty in terms of yield variability and effectiveness of the maintenance actions. Based on the cases studied, we can conclude that much of the improvement comes from improving the maintenance schedule. However, simply changing the production schedule to account for yield differences can substantially improve performance. In the problems studied, the simultaneous approach yielded rewards that were an average of more than eight percent greater than the traditional approach. In highly competitive industries such as semiconductor manufacturing, such improvements could make the difference between long-term success and failure.

References

- Borden, P. G. and L. A. Larson. “Benefits of Real-Time, In Situ Particle Monitoring in Production Medium Current Implantation”. *IEEE Transactions on Semiconductor Manufacturing*. Vol. 2, No. 4 (1989). pp. 141–145.
- Heyman, D. P. and M. J. Sobel. *Stochastic Models in Operations Research, Volume II: Stochastic Optimization*. New York, NY: McGraw-Hill, 1984.
- Kao, E. P. C. “Optimal Replacement Rules When Changes of State Are Semi-Markovian”. *Operations Research*. Vol. 21, No. 6 (1973). pp. 1231–1249.
- Osaki, S. and H. Mine. “Linear Programming Algorithms for Semi-Markovian Decision Processes”. *Journal of Mathematical Analysis and Applications*. Vol. 22 (1968). pp. 356–381.
- Puterman, M. L. *Markov Decision Processes: Discrete Dynamic Programming*. New York, NY: John Wiley and Sons, 1994.
- Sloan, T. W. “Simultaneous Determination of Production and Maintenance Schedules Using In-Line Equipment Condition and Yield Information”. Working Paper, School of Business, University of Miami, 2001.
- Sloan, T. W. and J. G. Shanthikumar. “Combined Production and Maintenance Scheduling for a Multiple-Product, Single-Machine Production System”. *Production and Operations Management*. Vol. 9, No. 4 (2000). pp. 379–399.

TABLE 1: Summary of Numerical Results

	All Problems	Set 1		Set 2	
		A3/A1	A3/A2	A3/A1	A3/A2
	Average	5.0	2.0	12.3	3.8
	Maximum	325.4	45.9	140.6	45.4
	Minimum	0.0	0.0	0.0	0.0
Factor	Level				
Product Mix	0	0.3	0.2	1.0	0.8
	1	0.9	0.5	4.7	0.5
	2	3.7	0.7	18.8	1.6
	3	12.2	1.7	31.8	4.1
	4	3.5	2.7	9.7	6.3
	5	4.8	2.7	8.4	5.2
	6	2.1	1.4	6.8	3.8
	7	5.3	2.7	13.8	4.6
	8	6.2	3.0	12.2	3.5
	9	9.8	3.3	19.6	5.6
	10	6.3	3.0	8.9	6.1
Unit Profits	100	6.7	2.0	16.8	3.9
	200	3.4	1.9	7.8	3.8
Yield Matrix	low	3.3	1.2	8.1	2.6
	high	6.7	2.8	16.5	5.1
Minor Clean Cost	100	2.2	1.4	9.1	3.6
	200	7.9	2.6	15.6	4.1
Major Clean Cost	300	5.2	2.1	11.6	3.8
	400	4.8	1.9	13.1	3.9
Equipment Deterioration	slow	1.5	1.1	3.3	2.3
	fast	8.5	2.9	21.4	5.4
Minor Clean Effectiveness	low	6.4	2.4	14.1	4.1
	high	3.6	1.6	10.6	3.5
Major Clean Effectiveness	low	4.5	1.9	12.0	3.8
	high	5.5	2.1	12.6	3.9
Minor Clean Time	short	4.3	1.6	11.1	3.5
	long	5.7	2.3	13.6	4.1
Major Clean Time	short	7.5	3.2	17.5	5.4
	long	2.5	0.8	7.2	2.3

Notes:

A3/A1 is the average percentage improvement of Approach 3 over Approach 1 for the given factor level.

A3/A2 is the average percentage improvement of Approach 3 over Approach 2 for the given factor level.