

A MARKOV DECISION PROCESS MODEL FOR OPTIMAL POLICY MAKING IN THE MAINTENANCE OF A SINGLE-MACHINE SINGLE-PRODUCT TOOLSET

Juan Sebastián Borrero
Raha Akhavan-Tabatabaei

Universidad de los Andes
Cra. 1 No. 18A-10
Bogotá D.C, Colombia

ABSTRACT

An aspect of great importance in Semiconductor Manufacturing Systems (SMS) is that machines are subject to unpredictable failures. Since Semiconductor Manufacturing is a highly capital intensive industry, it is crucial to optimize the usage of the resources. Performing preventive maintenance (PM), if done optimally, can reduce the risk of unpredicted failures and hence minimize the cost of outages. However, performing frequent PM results in higher cycle time and WIP accumulation at the toolset. In this paper we present a method to create optimal policies for a single-server single-product workstation using a Markov Decision Process model. The optimal policy determines whether or not to perform the PM based on the WIP level and the time since last repair. We present some numerical examples to illustrate the behavior of the optimal policy under different scenarios and compare the results with some common policies such as fixed frequency PM.

1 INTRODUCTION

SMS fabrication facilities (fabs) are usually arranged as job-shop like manufacturing systems with multiple workstations that perform several operations on different types of products. These fabs usually operate under a different set of conditions such as batching of jobs, failures of machines, multiple products and re-entrant lines, among others. Existence of such components complicates the analysis of SMS fabs in terms of estimating the cycle time or waiting time as explained by Shanthikumar, Ding, and Zhang (2007) and Connors, Geigin, and Yao (1996).

One of the most important aspects in SMS fabs is that the machines are subject to unpredictable failures. To prevent these failures the line managers perform periodic preventive maintenance (PM) procedures with certain frequency. This is mainly done to avoid the variability of machine availability due to unpredictable failures, since highly variable processing times dramatically increase the performance measures such as mean cycle time of the manufacturing systems. Inflated cycle times have negative consequences in the operations of the fabs, which translate to a sharp increase in costs as evidenced by Hopp and Spearman (2008). On the other hand performing too frequent and unnecessary PMs results in excessive cost of outage, WIP accumulation and PM cost itself. This implies the importance of controlling the failures and the maintenance activities of machines in a fab, with the purpose of optimizing the fab performance.

Shanthikumar, Ding, and Zhang (2007) mention that the two common approaches to model and analyze fab behavior are simulation models and queueing theory based models. Due to the complexity of the operations in this industry, simulation models consume much resource and thus result in inefficient outcomes when applied for quick what-if analysis and spontaneous decision making. On the other hand, analytical queueing theory based models provide tools, that although not always exact, can approximate the behavior of toolsets with high performance and can provide fast answers regarding the performance of the toolsets to the managers.

In the past decades, a number of queueing models have been formulated including some customized models for SMS. These models allow the managers estimate the performance measures of the fabs,

and can give fast responses to what if questions in order to improve the efficiency of their operations. Moreover, based on such models the researchers have proposed various policies aimed at different aspects of operations management that provide optimal or near optimal results in some predetermined performance measures of the fabs.

Extensive literature exist on approaching the problem of performing maintenance in general failure prone systems. Wang (2002) performs a survey that discusses and categorizes the different approaches made in order to find optimal maintenance policies subject to different cost functions. The majority of models that have been studied are based solely on the age of the system, ignoring relevant aspects of manufacturing systems such as WIP.

In works by Van der Duyn Schouten and Vannestee (1995) and Meller and Kim (1996) preventive maintenance is considered in a manufacturing system. Both papers explicitly consider the inventory level as well as the age of the system. The inventory information is incorporated by using a buffer that feeds a machine and the model is built in such a way that the optimal policy gives information regarding the level in which the buffer should be maintained. One of the main aspects taken into account is the service level and their approach is to find a maintenance policy that minimizes lost demand. The main disadvantage of these models is that, although the system fails and is repaired in a random fashion, they consider the demand and the service time as deterministic variables. Therefore the variability that is usually observed in demand or service time, although crucial in the dynamics of manufacturing systems, is not taken into account.

In the work of Yao et. al (2004) and Yao et. al (2005), an optimal maintenance model is built using Mixed Integer Programming and Markov Decision Processes. The former considers a complex SMS while the latter focuses only on a single machine. The inventory is taken into account in a similar way that in Van der Duyn Schouten and Vannestee (1995) and Meller and Kim (1996). Although these models incorporate the inventory and the age of system as the main factors, the service times and demand process are again considered deterministic; thus they are not capable of examining the effects of variability that is created in the system by sources of demand and service time.

Das and Sarkar (1999) also address the maintenance problem in a general manufacturing system, with the difference that the behavior of the demand and the service times are ruled by random variables. They obtain optimal maintenance policies by minimizing a cost function that implicitly depends on the inventory and on the age of the system. However, they consider inventory only through finished inventory. Finished inventory is maintained using an (s, S) policy which begins to process jobs whenever an order is made such that jobs are always available in an initial buffer. This implies that they do not consider WIP at the machine, so the effect of inventory on cycle time is not directly taken into account.

In this work we formulate a Markov Decision Process model for a single machine-single product workstation subject to random failures. The purpose of the model is to obtain an optimal policy on when to perform PMs, using a cost function that includes the cost of performing the PM, the cost of holding WIP when the machine is not available and the cost associated with the risk of an unpredicted failure. Although the manufacturing system that we propose is simple, the mathematical derivation is straight forward and we believe that the approach can be extended to multi-machine workstations and to more general cases by fairly easy derivations.

The rest of the paper is structured as follows: in Section 2 we present the formulation of the model, in Section 3 some numerical results and analysis are presented, also comparisons with common policies are presented in this section. Finally section 4 concludes the paper.

2 MODEL FORMULATION

We consider a single tool workstation in a semiconductor manufacturing fab that works in shifts of fixed duration denoted by $\varepsilon > 0$. Lots arrive to this workstation according to a Poisson process with rate λ , we denote this process by $A = \{A(t), t \geq 0\}$, where $A(t)$ represents the number of arrivals up to time $t \geq 0$. The processing time for every lot is a random variable B , all the processing times are independent and identically distributed. We assume that the arrival process and the processing times are independent. The workstation has a maximum capacity of keeping K lots.

This workstation is also subject to unforeseen failures with random time between failures. We assume that every time the machine fails it takes until the end of the shift to fix it and that at the beginning of each shift the machine is always ready to work. The machine is also required to undergo Preventive Maintenance (PM) periodically. It is assumed that performing a PM reduces the risk of having a failure in the following shifts and once a PM is scheduled it is always performed at the beginning of the shift.

Each PM is assumed to take exactly $\delta > 0$ units of time and costs c_m units. We define m as the number of shifts that have elapsed since the last PM or the last failure, where $0 < m < M$, and M is the maximum number of shift allowed to go without any repair or PM.

At the beginning of each shift the Operation Manager (OM) faces the decision of whether or not to perform the PM on that shift. Although performing the PM reduces the risk of failure in the near future it incurs a cost associated with holding the WIP during the time of PM. This cost is denoted by c_w per unit WIP per unit time. On the other hand facing a failure incurs the repair cost of c_r per unit time plus the cost of holding the WIP.

Finally we define Q as the random variable that denotes the lifetime of a machine that has just been repaired or maintained. For each $b \in \{0, 1\}$, $T_b(m)$ denotes the lifetime of the machine given that it has been working for m consecutive shifts since the last outage where the reason for the last outage is b . If $b = 0$ the last outage has been due to a PM and if $b = 1$ the last outage has been due to a failure. The random variable $T_b(m)$ can be defined in terms of Q in the following way,

$$P[T_b(m) \leq t] = \begin{cases} P[Q \leq t | Q > m\epsilon], & \text{if } t \geq m\epsilon, 0 \leq m < M, b = 1 \\ P[Q \leq t | Q > m\epsilon - \delta], & \text{if } t \geq m\epsilon - \delta, 1 \leq m < M, b = 0 \end{cases}$$

The b has to be introduced, because, for the same m , the distribution of the failure time is different depending on the type of the last outage. If it was due to a PM, then if at the beginning of any shift the machine has been working for m shifts, the time elapsed since the outage is $(m - 1)\epsilon + \epsilon - \delta$, and if the outage was due to a failure, then the time elapsed is $m\epsilon$ (so for the same m there is a difference of δ units of time).

2.1 The MDP Model Formulation

To provide an optimal solution for this problem we develop the following MDP model.

2.1.1 Decision epochs

The decision epoch for such a model is the beginning of each shift when the OM has to make the decision of whether or not to perform the PM. We define the set of decision epochs as $E = \{1, 2, 3, \dots, N\}$, $N \leq \infty$, where $i \in E$ is the beginning of the i -th shift.

2.1.2 State Space and State Variables

We define the state of the system with three variables denoting the existing WIP, n , the number of shifts elapsed since the last PM or repair, m , and the reason for the last outage of the machine b ($b = 0$ indicates that the last stop was due to a PM and $b = 1$ indicates that the last stop was due to a repair). Let S_1 and S_2 be the following sets:

$$S_1 = \{(n, m, b) : 0 \leq n \leq K, 0 \leq m \leq M, b \in \{0, 1\}\}$$

$$S_2 = \{(n, 0, 0) : 0 \leq n \leq K\}$$

Then the state space S is defined as

$$S = S_1 - S_2 \tag{1}$$

We denote $U(t) \in S$ as the state of the system at decision epoch $t \in E$.

2.1.3 Decision Set

Independent of the time and the state of the machine, only two decisions can be made: do a preventive maintenance at the start of the shift (Y), or not (N). Let \mathcal{D}_s be the set of decisions possible decisions in state s . If S_3 is the following set

$$S_3 = \{(n, M, b) : 0 \leq n \leq K, b \in \{0, 1\}\} \tag{2}$$

then:

$$\mathcal{D}_s = \begin{cases} \{Y, N\}, & \text{for } s \in S - S_3 \\ \{Y\}, & \text{for } s \in S_3 \end{cases} \quad (3)$$

For each $t \in E$, $D(t)$ will denote the decision that is made at time epoch t .

2.1.4 Transition probabilities

When the system begins a shift in state (n, m, b) the decision that is made at that moment determines the set of possible states of the system at the beginning of the next shift.

If decision N is made then the system can either transition to $(l_N, 0, 1)$, $0 \leq l_N \leq K$, if the machine fails within the current shift or if no failure happens it transitions to $(l_N, m + 1, b)$, $0 \leq l_N \leq K$. The probability of occurrence of each case depends on the state of the system and the probability of a failure. If decision Y is made then the system transitions to states $(l_Y, 1, 0)$, for some $0 \leq l_Y \leq K$ if there is no failure during the shift. If there is a failure, then the system transitions to states $(l_Y, 0, 1)$ for some $0 \leq l_Y \leq K$.

The derivation of the transition probabilities now follow. There are four main scenarios that have to be considered depending on the decision made and whether or not there is a failure during the shift. By $P_{n,l}^Y(v)$ we denote the probability that the WIP in the machine moves from n to l in v units of time, given that the machine is working, and denote $P_{n,l}^A(v)$ the same probability but given that the machine is not working. Note that these transition probabilities do not depend on the past because the random variables involved are exponential. The transition probabilities if decision is N will be the first that will be examined. Suppose that at the beginning of a shift the system is in state (n, m, b) and assume that the machine does not fail during the shift. This means that at the beginning of the following shift the system must be in a state of the form $(l_N, m + 1, b)$ for some $0 \leq l_N \leq K$. The probability of being in that state is the probability that the WIP moves from n to l_N in ε units of time, times the probability of not having a failure. In a more concise way this is

$$P_{n,l_N}^Y(\varepsilon) \times P[\text{not having a failure during the shift}]$$

Now suppose decision Y is made. The only difference in the analysis with the previous cases is that the PM downtime has to be taken into account. If there is no failure during the shift after the PM then the system can only transition to states $(l_Y, 1, 0)$, for some $0 \leq l_Y \leq K$. In this case it is just necessary to condition on the WIP there will be at the time the PM is done. Suppose that this WIP is r , then the probability of being in state $(l_Y, 1, 0)$ at the beginning of the next shift is the probability that the WIP moves from n to r during the PM, times the probability that the WIP moves from r to l_Y during the rest of the shift while the machine is working times the probability of not having a failure. Since the machine is not working while the PM is being done, this is precisely

$$P_{n,r}^A(\delta) \times P[\text{not having a failure during the rest of the shift}] \times P_{r,l_Y}^Y(\varepsilon - \delta) \quad (4)$$

In this case, the total probability is obtained by summing over all the possible units of WIP there are at the time when the PM is done.

The analysis for the case where decisions N or Y are made and a failure occurs follow a similar fashion. A detailed mathematical derivation of these probabilities is presented in Appendix A.

2.1.5 Cost Function

We define $r(s, d)$ as the cost associated with residing in state s and making decision d , for every $s \in S$, $d \in \mathcal{D}_s$. This cost includes the holding cost of WIP, c_m , cost of performing a PM, c_w , and cost of repair c_r .

If decision N is made, then the expected cost during the shift will be

$$r((n, m, b), N) = c_w E[XI_\varepsilon^N(n, m, b)]\varepsilon + c_r E[[s_b - T_b(m)]^+] \quad (5)$$

where $XI_\varepsilon(n, m, b)$ is the average units of WIP there are in the machine during a shift when the machine is not working, given that at the start of the shift the system was in state (n, m, b) and decision N is

made, $[x]^+$ is defined as $[x]^+ := \max\{x, 0\}$. Finally, s_b takes the following value

$$s_b = \begin{cases} (m + 1)\varepsilon - \delta, & \text{if } b = 0 \\ (m + 1)\varepsilon, & \text{if } b = 1 \end{cases} \tag{6}$$

Note that according to this $E[[s_b - T_b(m)]^+]$ is just the expected value of time that the machine will be under repair during the shift.

If decision Y is made, then the expected cost will be of the form

$$r((n, m, b), Y) = c_m + c_w E[XI_\varepsilon^Y(n, m, b)]\varepsilon + c_r E[[\varepsilon - \delta - Q]^+] \tag{7}$$

where $XI_\varepsilon^Y(n, m, b)$ is the average extra units waiting in the machine during a down time of the machine, given that at the start of the shift there system was in state (n, m, b) and decision Y is made. The computation of $E[[s_b - T_b(m)]^+]$ follows using straightforward integration. For the expected value of WIP the detailed mathematical derivation is presented in Appendix B.

2.2 The Optimization Model

In order to find the optimal policy a linear program is formulated, with the purpose of minimizing the discounted total cost in the long run. Let β be the discount factor. Let $V_\phi(s)$ be the cost that the system incurs in the long run following a stationary policy ϕ , given that at the start is in the state s . If $\phi(s)$ is defined as the decision that is made under the stationary policy, given that the system is in state s , it follows that

$$V_\phi(s) = r(s, \phi(s)) + \beta \sum_{s' \in S} p(s'|s, \phi(s))V_\phi(s') \tag{8}$$

We define V_s as the decision variable of the linear program, where $s \in S$, and $z = \sum_{s \in S} V_s$ as the function to be optimized. Then the linear program is given by:

$$\max z = \sum_{s \in S} V_s \tag{9}$$

subject to the constraints

$$V_s \leq r(s, d) + \beta \sum_{s' \in S} p(s'|s, d)V_{s'} \quad \forall s \in S, d \in \mathcal{D}_s \tag{10}$$

and where $V_s \geq 0 \forall s \in S$.

3 NUMERICAL RESULTS

In this section the proposed model is examined under different sets of parameters. Some base cases are first presented and analyzed. Then a sensitivity analysis is made in which certain parameters of the base cases are changed one by one to observe the effect of this changes on the behavior of the optimal policy. Then a comparison is performed between the value of the optimal policy and the values resulting with the use of some fixed policies that are commonly used.

3.1 Base Cases

Ten base cases are built using a model with fixed parameters, the only difference between various base cases is that the Failure Time Distribution (FTD) parameters are changed. This is done in order to get FTDs with different variances but the same mean; the rest of the parameters of the model are fixed. The parameters of each base case are presented in Table 1. The information regarding the FTD for each of the ten base cases is presented in Table 2. We assume that all the FTDs are Weibull and we denote k for the shape parameter of the distribution and θ for its scale parameter.

Table 1: Parameters of the base cases

Parameter	Description	Value
λ	Arrival rate	1
μ	Service rate	3
K	Capacity	18
M	Max. Shifts	10
ε	Shift length	12
δ	PM length	2
β	Discount factor	.99
c_m	Cost of performing a PM	100
c_w	Cost of extra WIP per unit per unit time	40
c_r	Cost of repair per unit time	150

Table 2: FTD parameters of the base cases

Parameter/Case	1	2	3	4	5	6	7	8	9	10
θ	80	81.567	83.391	85.047	86.619	88.056	89.228	90.154	90.276	88.914
k	1	1.05	1.12	1.2	1.3	1.43	1.6	1.9	2.35	3.5
$E[Q]$	80	80	80	80	80	80	80	80	80	80
$CV^2(Q)$	1	.9	.8	.7	.6	.5	.4	.3	.2	.1

Since the FTDs are Weibull, the hazard (or failure) rate, $h(t)$, in each case is given by

$$h(t) = \frac{kt^{k-1}}{\theta^k} \quad t \geq 0 \tag{11}$$

Since $k \geq 1$ in all of the cases, it follows that the distributions of the lifetime of the machine has the Increasing Failure Rate property. This means that the likelihood of a failure increases as the time that the machine has been working increases. Figure 1 shows the behavior of the hazard rate of some of the FTD used in the analysis.

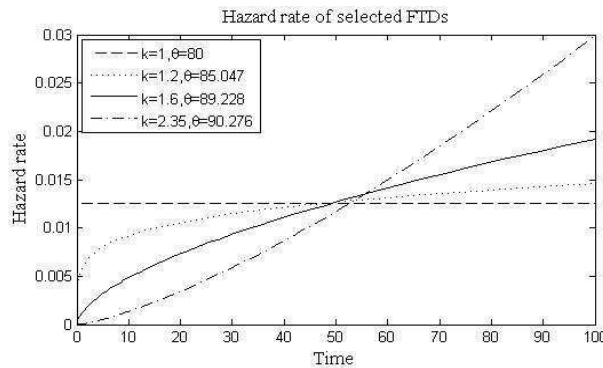


Figure 1: Hazard rate of some of the FTDs used in the analysis

As it is clear from Figure 1 the hazard rates become more convex and have a faster increase over time as k increases, and this increase makes the coefficient of variation (CV) of the FTD decrease.

In Figure 2 the results of the optimal policy for the base case parameters of four different FTDs are shown.

This graphs should be interpreted as follows: in the case #2, the optimal policy indicates that independent of WIP at the machine at the beginning of a shift, if the number of shifts that the machine has elapsed without an outage, m , is less than or equal to 8 then the PM should not be performed. If $m = 9$ then performing a PM is recommended only if there is no WIP at the machine at the beginning

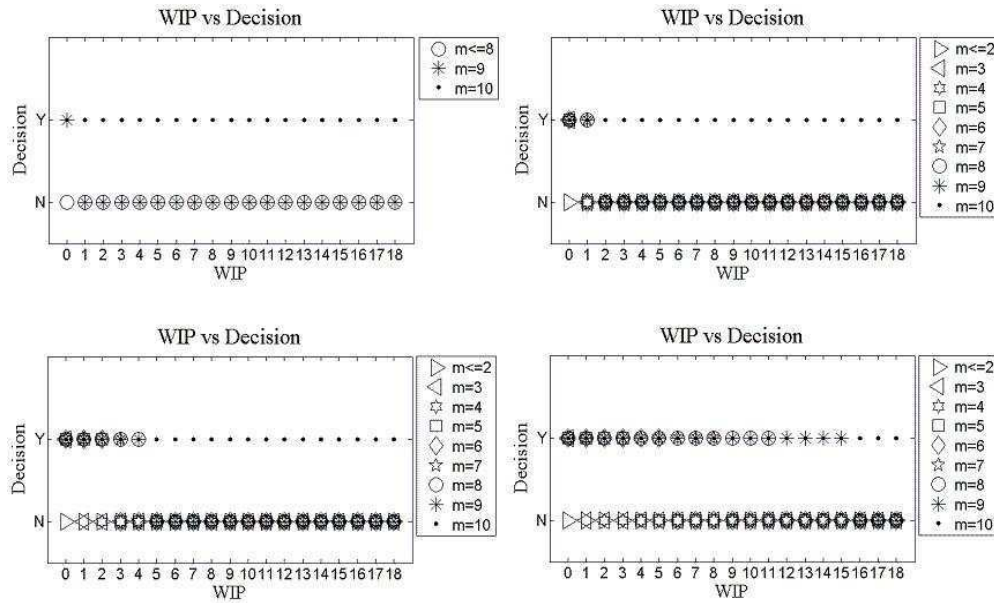


Figure 2: Optimal policies when the base case parameters are used, using different FTDs. From top left to bottom right respectively: FTD case #2, #5, #8 and #10.

of the shift. Note that when $m = 10$ the PM is always performed. This is because the maximum number of shifts that the machine can go without a PM is 10.

In a similar fashion, in the graph corresponding to case #5, the optimal policy is telling us that independent of the WIP, when $m \leq 2$ PM should never be performed. If $m = 3$ PM should be performed only if the WIP at the machine at the beginning of any shift is zero. The same applies to values of m between 4 and 7. If $m = 8$ or $m = 9$ then the PM should be performed only if the WIP is 1 or 0. The results for the cases #8 and #10 are analyzed in a similar fashion.

After examining the whole set of base cases, having into account that the only difference between these cases is a decrease in the CV of the FTDs that imply an increase in the hazard rate, one can observe the great impact that the decrease in CV has on the behavior of the optimal policy. When the CV of the FTD decreases then the optimal policy calls for performing the PM more often. This result has an intuitive explanation in that when the CV is less, then the probability of having a failure increases faster as m increases, compared to higher values of CV, and that whenever the CV decreases the hazard rate increases in a faster way with time. This is true when the FTDs are modeled by Weibull distributions with shape parameters greater than 1.

Another conclusion of great importance after examining the base cases is that the general behavior of the optimal policy follows some rules. In order to study this behavior it is convenient to introduce some notations. Let b be fixed, denote by m^* the first level of m at which decision Y begins to be made and n_m as the largest value of WIP at which decision Y is made for the last time in level m . If in level m decision Y is never made, define n_m to be -1 . For example, from Figure 2 case #2 $m^* = 9$, $n_{m^*} = n_9 = 0$, and $n_m = -1$ if $0 \leq m \leq 8$. In case #5 $m^* = 3$, $n_{m^*} = n_3 = 0$, in case #8 $m^* = 3$, $n_{m^*} = 0$, $n_7 = 3$, $n_8 = n_9 = 4$, and so on. Using this notation it is observed that the behavior of the optimal policy follows the following pattern; for any m if at a level of WIP of n decision Y is made, then for values of WIP less than n the decision Y is always made (for the same value of m), and that $n_m \leq n_{\hat{m}}$ for $m^* \leq m \leq \hat{m} \leq M - 1$.

According to this observed pattern, for each b one can define a subset of the state space S , in which decision Y is always made. We call this subset the Y region, and denote it by $S_Y(b)$. Then $S_Y(b)$ is given by

$$S_Y(b) = \{s \in S : s = (n, m, b), m^* \leq m \text{ and } n \leq n_m\}. \tag{12}$$

This means that the Y region can be completely characterized by the vector $(m^*, n_{m^*}, n_{m^*+1}, \dots, n_{M-1})$. Although this vector can be obtained for every instance, using the *second rule*?? one can get a great amount of information of $S_Y(b)$ by just using m^* , n_{m^*} and n_{M-1} and observing that n_m increases

uniformly from n_{m^*} to n_{M-1} as m increases. For this reason, in what follows we define (m^*, n_{m^*}, n_{M-1}) to be the Y information vector.

Using these previous observations, the question arises if this is true in a general analytic way, i.e., if it is possible to find an structure for the optimal policy. In order to get conditions that assure that this is the case, an ordering has to be imposed into the state space in such a way that the cost function and the complementary transition probabilities are non-decreasing (or non-increasing) in the the state space. However in this model it is not evident how to order the states such that these two conditions hold despite numerous attempts. However, from the numerical results, it is quite plausible that a structure for the optimal policy exists in this model.

3.2 Sensitivity analysis

In Table 3, the Y information vector is shown for different instances obtained by changing some parameter in the base cases using $b = 1$. The results when $b = 0$ are similar. For example row 3 of the table shows the Y information vector when the repair cost c_r is changed to 300 and the rest of the base cases parameters remain fixed. Row 4 shows the Y information vector when c_r is changed to 75 and the rest of the base cases parameters remain fixed. The idea applies for the rest of the row of the table. In the first column is indicated what parameter changed, an by what amount, with respect to the base cases parameters.

Table 3: Y information vector (m^*, n_{m^*}, n_{M-1}) for different instances based on the base cases. In the first column the parameter that is changed is shown, in the first row the case # is shown.

	1	2	3	4	5	6	7	8	9	10
Base Par.	(10,18,18)	(9,0,0)	(7,0,0)	(4,0,0)	(3,0,1)	(3,0,2)	(3,0,3)	(3,0,4)	(3,0,6)	(3,0,15)
$c_r = 300$	(10,18,18)	(4,0,0)	(2,0,1)	(2,0,2)	(1,0,3)	(1,0,5)	(2,1,6)	(2,1,9)	(2,0,14)	(3,1,18)
$c_r = 50$	(10,18,18)	(9,0,0)	(9,0,0)	(8,0,0)	(6,0,0)	(5,0,0)	(4,0,1)	(4,0,2)	(4,0,3)	(4,0,7)
$c_w = 80$	(10,18,18)	(9,0,0)	(8,0,0)	(5,0,0)	(4,0,0)	(3,0,1)	(3,0,1)	(3,0,2)	(3,0,3)	(3,0,8)
$c_w = 20$	(10,18,18)	(9,0,0)	(6,0,0)	(4,0,1)	(3,0,2)	(3,0,4)	(2,0,5)	(2,0,8)	(3,1,12)	(3,0,18)
$\delta = 4$	(10,18,18)	(9,0,0)	(8,0,0)	(5,0,0)	(4,0,1)	(3,0,2)	(3,0,3)	(3,0,4)	(3,0,6)	(3,0,12)
$\delta = 1$	(10,18,18)	(9,0,0)	(6,0,0)	(4,0,1)	(3,0,1)	(3,0,2)	(2,0,3)	(2,0,4)	(2,0,7)	(2,0,16)
$\mu = 6$	(10,18,18)	(9,0,0)	(5,0,0)	(3,0,0)	(3,0,1)	(3,0,2)	(3,0,3)	(3,0,4)	(3,0,6)	(3,0,13)
$\mu = 1.5$	(10,18,18)	(7,0,1)	(5,0,1)	(4,0,1)	(3,0,2)	(2,0,3)	(2,0,3)	(2,0,5)	(2,0,9)	(3,1,18)
$\beta = .9$	(10,18,18)	(9,0,0)	(8,0,0)	(5,0,0)	(4,0,1)	(3,0,2)	(3,0,2)	(3,0,4)	(3,0,6)	(3,0,16)

From the information in Table 3 some conclusions follow. First, it is evident that the variability of the FTD plays a major role in the behavior of the optimal policy, and that in general, whenever the variability of the FTD decreases the PM's are performed more often. Second, independent of the parameter that is changed, the overall behavior of the Y information vector has a recognizable pattern as the variability of the FTD decreases. Note that as the variability decreases the m^* is the value that decreases first leaving n_{m^*} and n_{M-1} static, then after it has decreased to a certain level (for example 5 or 4) the value of n_{M-1} begins to increase leaving m^* and n_{m^*} static, and finally in some extreme cases (like when c_r is doubled or c_w is half the base case) c_{m^*} begins to increase. This leads to another conclusion that n_{m^*} is in general 0.

On the other hand, in general the behavior of the optimal policies follow the intuition while in some other cases it helps to clarify the situation. For the former consider the instance where c_r is doubled. One would expect to perform PM more often, and of course the optimal policy cofirms intuition. For the latter, consider the case where c_w is doubled. Intuition would say that PMs have to be performed more often since is more costly to hold WIP if the machine fails. However, c_w also directly affects the time when the PM's are being performed, so the behavior of the optimal policy with respect to c_w is more complex because it also depends on δ , the amount time that the PM lasts. The importance of δ is also evident when that value is doubled (when we perform less frequent PMs) and when is half (performing PMs more often). Also note that the optimal policies behave quite similar for different values of μ . The reason of that is the cost function's independence from μ , the transition probabilities being the only part of the MDP changed by changing μ . Finally note that the model is practically insensitive to changes in the discount factor β .

3.3 Comparison with Fixed Policies

In this section the optimal policy obtained by the MDP is compared against fixed policies that are commonly used. For this purpose base case #6 of the previous section is compared against two types of fixed policies. The first policy performs a PM whenever the WIP level at the beginning of the shift is below or equal to some prescribed level n_0 . The second policy only performs a PM whenever the number of shifts since the last outage is at a certain level m_0 . The results of these comparisons are shown in Table 4 and in Table 5.

Table 4: Average values for the first type of fixed policies where a PM is always performed if the WIP at the beginning of any shift is less or equal than n_0 . The average value for the optimal policy in this case is 12041.95.

n_0	Value	Increase
0	16524.36	37.22%
1	18364.92	52.51%
2	19280.78	60.11%
3	19697.72	63.58%
4	19878.39	65.08%
5	19956.53	65.73%
6	19993.27	66.03%
7	20013.29	66.20%
8	20027.16	66.31%
9	20038.87	66.41%
10	20050.11	66.50%

Table 5: Average values for the second type of fixed policies where a PM is always performed whenever the number of shifts since the last outage is m_0 . If the number of shifts is less than m_0 the PM is never performed. The average value for the optimal policy in this case is 12041.95.

m_0	Value	Increase
2	13566.95	12.66%
3	12534.32	4.1%
4	12413.85	3.1%
5	12568.17	4.37%
6	12893.43	7.07%
7	13230.59	9.87%
8	13544.76	12.48%
9	13849.85	15.01%
10	14115.76	17.22%

Finally the optimal policy is compared against two extreme policies, the policy that makes a PM at every shift and the policy that never performs a PM, except when $m = M$. When a PM is always performed, the average value of the cost function is 20160.15756 which is an increase of 67.42% with respect to the optimal policy. If a PM is never performed until the limit is reached then the average value of the cost function is 14115.76885 which is an increase of 17.22% with respect to the optimal policy value.

4 CONCLUSIONS

The proposed model provides an optimal policy for performing preventive maintenance in single-machine single-product systems. The main contribution of the model is that it considers random machine deterioration and random service and arrival processes and it explicitly includes the age of the machine and the WIP accumulation during machine outage as the criteria for optimality.

The current model cannot directly be applied to toolset in High Volume Manufacturing fabs, since they consist of multi-machine workstations. However, it can give insight on how WIP and machine age interact and the trade-off between the two so that an optimal maintenance policy is adopted minimizing the overall costs. The current model can also be applied to find the optimal PM policy at Research and Development fabs with single-tool workstations and also test equipment such as reticle inspection or metrology tools that do not typically exist in large numbers within a fab.

As the next step to this research a multi-machine workstation model is being considered. In this framework, aspects such as dispatching the lots to parallel servers and maintaining the overall availability of the toolset at a certain level are to be considered. This extension further complicates the MDP model, specially in the definition of the state space that cannot be extended in a straight-forward fashion.

APPENDIX A

In this appendix the mathematical derivation of the transition probabilities for the MDP are shown. For notation simplicity denote

$$p_t[s'|s, d]: = P[U(t + 1) = s'|U(t) = s, D(t) = d] \tag{13}$$

Let F_m^b and f_m^b denote the distribution and the density functions of $T_b(m)$ respectively. Define $X(t)$ as the WIP there is in the machine at time $t \geq 0$ and set $X: = \{X(t), t \geq 0\}$ as the associated stochastic process. The behavior of X will depend on whether the machine is working or not. If it is assumed that the service time of each job B is exponential with service rate μ , then it follows that when the machine is working, the evolution of X through time will be governed by a Continuous time Markov Chain $Y: = \{Y(t), t \geq 0\}$, and when the machine is not working, the evolution of X will be governed by the arrival process A .

Under the assumptions that the service times are i.i.d exponential with rate μ and that A is a Poisson process with rate λ , it follows that Y is a finite birth and death process with $K + 1$ states, birth rates given by $\lambda_i = \lambda$ for all $0 \leq i \leq K$, and with death rates given by $\mu_i = \mu$ for all $1 \leq i \leq K$. Define $P_{i,j}^Y(t): = P[Y(t) = j|Y(0) = i]$ for $0 \leq i, j \leq K$, $P_{i,j}^A(t): = P[A(t) = j|A(0) = i]$ and $\bar{P}_{i,j}^A(t): = P[A(t) \geq j|A(0) = i]$, for $0 \leq i, j$.

Finally, note that since Q resets to zero every time a PM or a repair is performed, if $U(t) = (n, m, 0)$, the t can be written as $m\epsilon - \delta$ for $m \geq 1$, and if $U(t) = (n, m, 1)$, then t can be written as $m\epsilon$, $m \geq 0$.

Suppose the decision Y is made, and that there is a failure during the shift. Then the system has to be conditioned on the WIP at the time when the PM is done, the WIP at the time of the failure and the time of the failure. If $U(t) = (n, m, 0)$ then the system should be in a state of the form $(l, 0, 1)$ at the beginning of the next shift, for some $0 \leq l \leq K - 1$. It follows that

$$\begin{aligned} p_t[(l, 0, 1)|(n, m, 0), Y] &= \sum_{r=0}^K \left\{ \sum_{j=n}^{K-1} P[A(t + \delta) = j|A(t) = n] \int_0^{\epsilon - \delta} (P[Y(t + \delta + u) = r|Y(t + \delta) = j] \right. \\ &\quad \times P[A(t + \epsilon) = l|A(t + \delta + u) = r] f_0(u)) du + P[A(t + \delta) \geq K|A(t) = n] \\ &\quad \times \left. \int_0^{\epsilon - \delta} (P[Y(t + \delta + u) = r|Y(t + \delta) = K] P[A(t + \epsilon) = l|A(t + \delta + u) = r] f_0(u)) du \right\} \\ &= \sum_{r=0}^K \left\{ \sum_{j=n}^{K-1} P_{n,j}^A(\delta) \int_0^{\epsilon - \delta} P_{j,r}^Y(u) P_{r,l}^A(\epsilon - \delta - u) f_0(u) du + \bar{P}_{n,K}^A(\delta) \int_0^{\epsilon - \delta} P_{K,r}^Y(u) P_{r,l}^A(\epsilon - \delta - u) f_0(u) du \right\} \end{aligned}$$

If $l = K$, then it can be shown that $p_t[(K, 0, 1)|(n, m, 0), Y]$ is:

$$\sum_{r=0}^K \left\{ \sum_{j=n}^{K-1} P_{n,j}^A(\delta) \int_0^{\epsilon - \delta} P_{j,r}^Y(u) \bar{P}_{r,K}^A(\epsilon - \delta - u) f_0(u) du + P_{n,K}^A(\delta) \int_0^{\epsilon - \delta} P_{K,r}^Y(u) \bar{P}_{r,K}^A(\epsilon - \delta - u) f_0(u) du \right\} \tag{14}$$

Finally note that if $U(t) = (n, m, 1)$, then the transition probabilities can be calculated in the same for that above give that $D = Y$, because the PM is made. This means that $p_t[(l, 1, 0)|(n, m, 1), Y] = p_t[(l, 1, 0)|(n, m, 0), Y]$ and that $p_t[(l, 0, 1)|(n, m, 1), Y] = p_t[(l, 0, 1)|(n, m, 0), Y]$. The case when $D = N$ and there is a failure is derived in a similar way.

APPENDIX B

In this appendix explicit formulas for the cost function of the MDP model are obtained. First expressions for $E[XI_{\epsilon}^d(n, m, b)]$, for $d = N$ and $d = Y$ are derived. In order to get them some previous results are needed. Define

$$I_{x,y}^*(t, u) = \frac{1}{u-t} \int_t^u 1_{\{A(z)-A(t)=y-x\}} dz \tag{15}$$

where 1_C is the indicator function of the set C . We set y to be of the form $y = x + k$, for some non negative integer k , and $0 \leq t < u$. (If $u < t$ then $I_{x,y}^*(t, u)$ is defined to be 0).

Note that $I_{x,y}^*(t, u)$ is a random variable that give the proportion of time there are y units of WIP during the time interval $[t, u]$ given that at t there were x units of WIP. Using I^* , XI_{ϵ}^d can be defined as

$$XI_{\epsilon}^N(n, m, b) = \sum_{j=r_b}^{K-1} (j - r_b) I_{r_b, j}^*(w_b, z_b) + K \left[1 - \sum_{j=r_b}^{K-1} I_{r_b, j}^*(w_b, z_b) \right] \tag{16}$$

for $b \in \{0, 1\}$, where in this case $w_0 = T_0(m)$, $w_1 = T_1(m)$, $z_0 = (m + 1)\epsilon - \delta$, $z_1 = (m + 1)\epsilon$ and

$$r_0 = E[Y_{m\epsilon-\delta}^{T_0(m)}(n)] \text{ and } r_1 = E[Y_{m\epsilon}^{T_1(m)}(n)] \tag{17}$$

where $Y_t^u(n)$, $0 \leq t \leq u$, is defined to be the WIP there is in the machine at time u given that at t there were n units of WIP and that from t to u the machine is working.

The above expression for XI_{ϵ}^N is a reasonable approximation, since the true random variable for the average inventory has r_0 replaced with $Y_{m\epsilon-\delta}^{T_0(m)}(n)$ and r_1 for $Y_{m\epsilon}^{T_1(m)}(n)$. The reason for doing this approximation will be explained below.

In order to compute $E[XI_{\epsilon}^N(n, m, b)]$ it will be useful to first compute $E[I_{x,y}^*(t, u)] =: M_{x,y}^*(t, u)$ for any positive numbers $0 \leq t \leq u$. It can be shown that

$$M_{x,y}^*(t, u) = \begin{cases} \frac{(u-t)^{-1}}{\lambda(y-x)!} \{ (y-x)\Gamma(\lambda(u-t), y-x) - e^{-\lambda(u-t)}(\lambda(t-u))^{y-x} \}, & y > x \\ \frac{1-e^{-\lambda(u-t)}}{(u-t)\lambda}, & y = x \end{cases} \tag{18}$$

where $\Gamma(\cdot, \cdot)$ is the upper incomplete gamma function:

$$\Gamma(s, i) = \int_0^s e^{-t} t^{i-1} dt \tag{19}$$

Conditioning on $T_b(m)$, and given that $T_b(m)$ and $A(t)$ are independent for all $0 \leq m \leq M$, $b \in \{0, 1\}$ and $t \geq 0$, it follows that $E[XI_{\epsilon}^N(n, m, b)]$ if given by

$$\sum_{j=r_b}^{K-1} (j - r_b) \int_{z_b-\epsilon}^{z_b} M_{r_b, j}^*(v, z_b) f_m^b(v) dv + K \left[1 - \sum_{j=r_b}^{K-1} \int_{z_b-\epsilon}^{z_b} M_{r_b, j}^*(v, z_b) f_m^b(v) dv \right] \tag{20}$$

for $b \in \{0, 1\}$. The reason for doing the approximation in (16) is because the computation of the above expected value can get very complicated if it is not used. Although an expression in terms of f_m^0 and f_m^1 can be obtained if the approximation is not used (by conditioning again), the numerical computation can get very time-consuming.

If at the start of the shift decision Y is made, it must be taken into account the extra inventory during the time the PM is being performed. If the state of the system is (n, m, b) , then the expected

value of extra inventory during the PM is given by

$$PME(n) = \sum_{j=n}^{K-1} (j-n)M_{n,j}^*(0, \delta) + K \left[1 - \sum_{j=n}^{K-1} M_{n,j}^*(0, \delta) \right] \quad (21)$$

for all $0 \leq m \leq M$, $b \in \{0, 1\}$. Because the machine can fail the rest of the shift, using similar arguments as before, it can be shown that $E[XI^Y(n, m, b)]$ is given by

$$PME(n) + \sum_{j=q}^{K-1} (j-q) \int_0^{\varepsilon-\delta} M_{q,j}^*(v, \varepsilon - \delta) f_0(v) dv + K \left[1 - \sum_{j=q}^{K-1} \int_0^{\varepsilon-\delta} M_{q,j}^*(v, \varepsilon - \delta) f_0(v) dv \right] \quad (22)$$

for all $0 \leq m \leq M$, $b \in \{0, 1\}$, where in this case

$$q = E[Y_0^Q(n)] \quad (23)$$

REFERENCES

- Connors D.P, Geigin G.E. and Yao D.D. 1996. A queueing model for semiconductor manufacturing. *IEEE Transactions On Semiconductor Manufacturing* 9(3): 412-427.
- Das T.K. and Sarkar S. 1999. Optimal preventive maintenance in a production inventory system. *IIE Transactions* 31: 537-551.
- Hopp W.J and M. L. Spearman, 2008, *Factory Physics: Foundations of Manufacturing Management*, New York: McGraw - Hill.
- Meller R.D. and Kim D.S. 1996. The impact of preventive maintenance on system cost and buffer size. *European Journal of Operational Research* 95: 577-591.
- Shanthikumar J. G., Ding S. and Zhang, M. 2007. Queueing theory for semiconductor manufacturing systems: a survey and open problems. *IEEE Transactions on Automation Science and Engineering* 4: 321-335.
- Van der Duyn Schouten F.A. and Vanneste S.G. 1995. Maintenance optimization of a production system with buffer capacity. *European Journal of Operational Research* 82: 328-338.
- Wang. H. 2002. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research* 139: 469-489.
- Yao X., E. Fernández-Gaucherand, M. Fu and S. I. Marcus. 2004. Optimal preventive maintenance scheduling in semiconductor manufacturing. *IEEE Transactions on Semiconductor Manufacturing* 17 (3): 345-355.
- Yao X., X. Xiaolan, M. Fu and S. I. Marcus. 2005. Optimal joint preventive maintenance and production policies. *Naval Reseach Logistics* 52:669-681.

AUTHOR BIOGRAPHIES

JUAN SEBASTIAN BORRERO is a master student of Industrial Engineering at Universidad de los Andes and holds a BS.c in mathematics from the same university. The title of his master thesis is "An MDP model for Optimal Maintenance Policy in multi-machine SMS Environment".

RAHA AKHAVAN-TABATABAEI is currently an assistant professor at Universidad de los Andes, Bogota, Colombia. Prior to this she worked for Intel Corporation, AZ, USA, as a senior industrial Engineer for HVM fabs. She has a PhD in Industrial and Systems Engineering from North Carolina State University.