

A NUMERICAL STUDY ON THE STRUCTURE OF OPTIMAL PREVENTIVE MAINTENANCE POLICIES IN PROTOTYPE TANDEM QUEUES

Taehyung Kim
James R. Morrison

Department of Industrial and System Engineering
Korea Advanced Institute of Science and Technology
291 Daehak-ro, Yuseong-gu
Daejeon, 34141, SOUTH KOREA

ABSTRACT

While high levels of automation in modern manufacturing systems increase the reliability of production, tool failure and preventive maintenance (PM) events remain a significant source of production variability. It is well known for production systems, such as the M/G/1 queue, that optimal PM policies possess a threshold structure. Much less is known for networks of queues. Here we consider the prototypical tandem queue consisting of two exponential servers in series subject to health deterioration leading to failure and repair. We model the PM decision problem as a Markov decision process (MDP) with a discounted infinite-horizon cost. We conduct numerical studies to assess the structure of optimal policies. Simulation is used to assess the value of the optimal PM policy relative to the use of a PM policy derived by considering each queue in isolation. Our simulation studies demonstrate that the mean cycle time and discounted operating costs are 10% superior.

1 INTRODUCTION

1.1 Overview

While high levels of automation in modern manufacturing systems increase the reliability of production, tool failure and preventive maintenance (PM) events remain a significant source of production variability. It is well known for production systems, such as the M/G/1 queue, under certain assumptions and subject to tool failures, repairs, and PMs, that optimal PM policies possess a threshold structure. Much less is known for queues in tandem, or more generally, networks of queues. Here we consider the prototypical tandem queue consisting of two exponential servers in series subject to deterioration in health, failure, and repair and numerically explore the structure of optimal policies in this system.

1.2 Literature Review

Much of the research on PM policies has concentrated on non-production systems or has taken an isolated view of the production equipment. When production systems are considered, there are two general categories of work: production systems that build into an inventory and queueing systems. We will briefly discuss equipment, inventory, and queueing systems next, with an eye toward motivating our work.

PMs are essential for efficient machine operation and much work has been devoted to the study of PM policies, c.f. (Sim and Endrenyi 1988; Nicolai and Dekker 2008; Moghaddan and Usher 2011). Applications include manufacturing equipment (cf. Laggoune et al. (2009)), windmills (cf. Krishna (2012)), and trucks (cf. Barde et al. (2016)). Often, a PM policy to maximize the mean availability is sought. Alternately, the cost of maintaining the machine may be pursued by considering failure costs, repair costs, and maintenance costs. Recent examples in this vein of work include Barde et al. (2016) and Barde et al. (2019) which focus

on PMs for a military vehicle. There, a Markov decision process (MDP) model was developed that assumed age-based deterioration of the components of the vehicle. Numerical approaches for obtaining near optimal policies were used.

In production systems, inventory holding costs and production rewards may also factor into the PM decisions. Some production system models assume a single stage or tool which produces into an inventory of finished goods. The tool production rate can be controlled when the tool is functional; otherwise the production rate is zero. Yao et al. (2005) study such a single tool production-inventory system. They modeled their system as an MDP and proved structural properties for the optimal policy as a function of system age and inventory level. A reinforcement learning approach was used in Das and Sarkar (1999) with a single tool, single/multiproduct inventory system as a case study.

Queueing systems can be considered as another class of production systems. Work arrives to a queue ahead of the tool and waits for service. After service, the finished work exits the system. The tool cannot work if there is nothing in the queue and if the tool has suffered a failure or PM event. For an M/G/1 queue, Yao et al. (2003) proved structural properties of optimal maintenance policies using a Semi-Markov Decision Process (SMDP) model. In Borrero (2010), a further deeper numerical investigation into a similar system was conducted. Cai et al. (2013) derived structural properties for PM policies in single machine multiproduct systems. Similar structural properties were observed. In those studies, for fixed queue length, there was an equipment wear or age threshold after which maintenance was dictated. In Kaufman and Lewis (2007), a theoretical characterization of a replacement and repair policy for a single stage, single server system showed similar results. In their numerical studies, they also demonstrated that the threshold is not necessarily monotonically non-decreasing in the queue length. This non-monotonicity was only observed to occur when the queue length is zero. These theoretical and numerical observations on the structure of optimal policies have thus far been for systems with a single stage (and single server at that stage).

It is worth noting that Ramirez-Hernandez and Fernandez (2010) used approximate dynamic programming to study the Intel mini-fab which is a small scale queueing network. They controlled the PM decisions at the bottleneck toolset alone. Their numerical study did not investigate policy structure but rather focused on assessing improvement in mean cycle time relative to a fixed PM schedule.

1.3 Contribution and Organization

In this work, we focus on a prototype queueing network consisting of a series of two exponential servers, a Poisson arrival process, and infinite buffers preceding each server. Servers experience deterioration, failure, repair and PMs. We explore the structure of optimal policies. The contributions follow. We

- Provide an MDP model for this two-stage tandem line;
- Numerically demonstrate the existence of optimal threshold policies for PMs; and
- Numerically demonstrate various monotonicity properties of these thresholds.

Note that because the system now has two stages and two buffers, the thresholds are more complicated. They depend on the state of the other machine and there are interesting cross effects.

The remainder of the paper is organized as followed. Our formal MDP model for the two-stage tandem queue is presented in Section 2. In Section 3, we review the numerical studies conducted and state our numerical observations. We assess the value of using a joint PM control relative to a PM control derived from consideration of the queues in isolation. Concluding remarks, including comments on future directions, are provided in Section 4.

2 SYSTEM DESCRIPTION

In this section, we first introduce the notations that are used for the description.

2.1 Notation

S :	Finite set of states
w_t^i :	WIP at buffer i at time t (including any in the machine)
s_t^i :	Deterioration status of machine i at time t
A :	finite set of actions
a_t^i :	action chosen for machine i at time t
λ :	arrival rate of incoming jobs
μ^i :	service rate of processing at machine i
σ^i :	deterioration rate of machine i
θ_{pm}^i :	repair rate for preventive maintenance of machine i
θ_f^i :	repair rate for corrective maintenance of machine i
D :	total rate of events in the system
T :	Random variable for the time between events in the sampled system
C_{pm} :	cost of preventive maintenance event
C_f :	cost of corrective maintenance event
C_h :	holding cost per job per time unit
γ :	Discount factor
K_{max} :	Deterioration status indicating that a machine is receiving preventive maintenance

2.2 Assumptions

We make the following assumptions:

- 1) The cost of failure and repair (C_f) is greater than the cost of a PM (C_{pm}), that is $C_f > C_{pm}$.
- 2) The mean time to conduct a repair is greater than the mean time to conduct at PM. That is, in terms of the repair rates, $\theta_f^i < \theta_{pm}^i$.
- 3) All interarrival times, service times, inter-deterioration times, repair times, and PM times are independent of each other and each has its own IID exponentially distributed duration. We can thus resort to uniformization (Lippmann (1975)) via the inclusion of virtual transitions and formulate the problem as an MDP.
- 4) PM events are non-preemptive. When a machine is operating, a transition to a PM state (that is, deterioration status of K_{max}), can only occur if the machine is idle or after a service completion event. PM events return the machine to a like new state, that is, to a deterioration status of 0.
- 5) Failure events occur when the deterioration state reaches $K_{max} - 1$. They are preemptive. Repair events then start immediately. When the repair event is complete, the machine is returned to a like new state, that is, to a deterioration status of 0.
- 6) Although the model can allow infinite queue lengths, for numerical convenience later, we truncate the state space and allow at most L_{max} jobs at each machine. If an arrival event occurs when machine 1's buffer is full, the arrival is lost. When machine 2's buffer is full, service at machine 1 is not allowed.

2.3 MDP Modeling of a Prototypical Tandem Queue with Failures and PMs

We study a queueing network consisting of two service stations. Jobs arrive as a Poisson process with rate λ to an infinite queue serving machine 1. They are served in FIFO manner when machine 1 is available. After service they immediately proceed to the infinite queue serving machine 2 and are similarly served in

FIFO manner. Jobs that have received service from machine 2 exit the system. Service times are IID exponential random variables at each machine (independent of each other and with possibly different rates μ_1, μ_2). Machines experience independent degradation events incrementing their deterioration status by one. Such events occur after IID exponential durations of time with rates σ_i , for machines $i = 1, 2$. Deterioration may lead to failure and is alleviated by repair events and PM events; see assumptions 4 and 5 above. Repair events and PM events are IID exponentially distributed with rates θ_f^i and θ_{pm}^i , respectively, for machines $i \in \{1, 2\}$.

The state at time t consists of the WIP at each buffer w_t^i (including any job in service) and the deterioration status of each machine s_t^i , $i \in \{1, 2\}$: $S_t = (w_t^1, s_t^1, w_t^2, s_t^2)$. Here, $w_t^i \in \{0, 1, \dots, L_{max}\}$, $i \in \{1, 2\}$ and $s_t^i \in \{0, 1, \dots, K_{max}\}$, $i \in \{1, 2\}$. The action at time t is the vector of machine actions $A_t = (a_t^1, a_t^2)$, where a_t^i is the indicator of the intent to start a PM event, $i \in \{1, 2\}$. The PM event indicated occurs on the next transition unless the machine is working, in which case the PM event begins after a service event at that machine (on account of our non-preemptive assumption). The PM event transitions the machine deterioration status to $s_t^i = K_{max}$, after which the machine can only transition to $s_t^i = 0$ and may not serve jobs until that time. Note that, if $A_t = (1, 1)$ and the queues are empty, then both machines will transition into a PM when the next event occurs.

We resort to uniformization, see Lippmann (1975). We associate an exponential clock to all random variables and sample the system state when there are real or virtual events. Let D denote the sum of all rates for all the exponential random variables, that is,

$$D = \lambda + \sum_i \mu^i + \sum_i \sigma^i + \sum_i \sum_k \theta_k^i.$$

We hereafter use t as our index for the discrete time system. Thus, t and $t + 1$ represent subsequent discrete time steps and not actual instants on the timeline. Let T denote the actual time interval between events in the sampled system, so that $E[T] = 1/D$.

For each sampling epoch, an operating cost is incurred. The operating cost is the sum of job holding costs and machine repair/PM costs. Note that the repair/PM cost is incurred one time when experiencing a failure or entering a PM as indicated by the conditions: $s_t^i < K_{max} - 1$ and $s_{t+1}^i \geq K_{max} - 1$. Our cost function is as follows:

$$g(w_t^1, s_t^1, s_{t+1}^1, w_t^2, s_t^2, s_{t+1}^2) = -c_h \cdot (w_t^1 + w_t^2) \cdot E(T) - I(s_t^i < K_{max} - 1) \cdot \{c_f \cdot \sum_i I(s_{t+1}^i = K_{max} - 1) - c_{pm} \cdot \sum_i I(s_{t+1}^i = K_{max})\},$$

where $I(A)$ is the indicator of the event A .

We consider the infinite-horizon discounted-cost version of the problem. With initial state $S = (w^1, s^1, w^2, s^2)$ and policy π , the expected total discounted cost is

$$V^\pi(w^1, s^1, w^2, s^2) = \lim_{n \rightarrow \infty} E\left\{ \sum_0^{t_n} \gamma^t \cdot g^\pi(w_t^1, s_t^1, s_{t+1}^1, w_t^2, s_t^2, s_{t+1}^2) \right\},$$

where g^π denotes the costs incurred via cost function g operating under policy π . We define

$$V(w^1, s^1, w^2, s^2) = \min \{Q^{a^1, a^2}(w^1, s^1, w^2, s^2) | (a^1, a^2) \in A\}$$

where $Q^{a^1, a^2}(w^1, s^1, w^2, s^2)$ is the cost-to-go under actions $(a^1, a^2) \in A$.

We require some helpful notation. For $j \in \{arr, ser, det, rep, tot\}$, we use sp^{ij} as the deterioration status of the machine as follows:

$$sp^{ij} = \begin{cases} k_{max}, & w^i = 0 \\ s^i + 1, & w^i > 0, a^i = 1, j = det \\ s^i, & otherwise \end{cases}$$

If $w^i \geq 1$, only non-preemptive maintenance is allowed, so that $sp^{ij} = s^i$. However, if j indicates a deterioration event ($j = det$), then $sp^{ij} = s^i + 1$. When there are no jobs at both machines, preventive maintenance is allowed on any next transition.

An optimal policy π^* can be obtained by solving the optimality equations detailed below.

Case1: $s^i < k_{max} - 1, i \in \{1,2\}$

$$Q^{0,0}(w^1, s^1, w^2, s^2) = C_h(w^1, w^2) + \frac{\gamma(\sigma^1 + \sigma^2)}{D} C_f(s^1, s^2) + \frac{\gamma}{D} \{ \lambda \cdot V(w^1 + 1, s^1, w^2, s^2) + \mu^1 \cdot V((w^1 - 1)^+ + I(w^1 = L_{max}), s^1, w^2 + I(w^1 \geq 1), s^2) + \mu^2 \cdot (V(w^1, s^1, (w^2 - 1)^+, s^2) + \sigma^1 \cdot V(w^1, s^1 + 1, w^2, s^2) + \sigma^2 \cdot V(w^1, s^1, w^2, s^2 + 1) + (D - (\lambda + \sum_i \mu^i + \sum_i \sigma^i)) \cdot V(w^1, s^1, w^2, s^2)) \}, \quad (1)$$

$$Q^{0,1}(w^1, s^1, w^2, s^2) = C_h(w^1, w^2) + \frac{\gamma(\sigma^1 + \sigma^2)}{D} C_f(s^1, s^2) + \frac{\gamma \mu^2}{D} C_{pm}(s^1, s^2) + \frac{\gamma}{D} \{ \lambda \cdot V(w^1 + 1, s^1, w^2, sp^{2,arr}) + \mu^1 \cdot V((w^1 - 1)^+ + I(w^1 = L_{max}), s^1, w^2 + I(w^1 \geq 1), sp^{2,ser}) + \mu^2 \cdot V(w^1, s^1, (w^2 - 1)^+, k_{max}) + \sigma^1 \cdot V(w^1, s^1 + 1, w^2, sp^{2,det}) + \sigma^2 \cdot V(w^1, s^1, w^2, sp^{2,det}) + (D - (\lambda + \sum_i \mu^i + \sum_i \sigma^i)) \cdot V(w^1, s^1, w^2, sp^{2,tot}) \}, \quad (2)$$

$$Q^{1,0}(w^1, s^1, w^2, s^2) = C_h(w^1, w^2) + \frac{\gamma(\sigma^1 + \sigma^2)}{D} C_f(s^1, s^2) + \frac{\gamma \mu^1}{D} C_{pm}(s^1, s^2) + \frac{\gamma}{D} \{ \lambda \cdot V(w^1 + 1, sp^{1,arr}, w^2, s^2) + \mu^1 \cdot V((w^1 - 1)^+ + I(w^1 = L_{max}), k_{max}, w^2 + I(w^1 \geq 1), s^2) + \mu^2 \cdot V(w^1, sp^{1,ser}, (w^2 - 1)^+, s^2) + \sigma^1 \cdot V(w^1, sp^{1,det}, w^2, s^2) + \sigma^2 \cdot V(w^1, sp^{1,det}, w^2, s^2 + 1) + (D - (\lambda + \sum_i \mu^i + \sum_i \sigma^i)) \cdot V(w^1, sp^{1,tot}, w^2, s^2) \}, \quad (3)$$

$$Q^{1,1}(w^1, s^1, w^2, s^2) = C_h(w^1, w^2) + \frac{\gamma(\sigma^1 + \sigma^2)}{D} C_f(s^1, s^2) + \frac{\gamma(\mu^1 + \mu^2)}{D} C_{pm}(s^1, s^2) + \frac{\gamma}{D} \{ \lambda \cdot V(w^1 + 1, sp^{1,arr}, w^2, sp^{2,arr}) + \mu^1 \cdot V((w^1 - 1)^+ + I(w^1 = L_{max}), k_{max}, w^2 + I(w^1 \geq 1), sp^{2,ser}) + \mu^2 \cdot V(w^1, sp^{1,ser}, (w^2 - 1)^+, k_{max}) + \sigma^1 \cdot V(w^1, sp^{1,det}, w^2, sp^{2,det}) + \sigma^2 \cdot V(w^1, sp^{1,det}, w^2, sp^{2,det}) + (D - (\lambda + \sum_i \mu^i + \sum_i \sigma^i)) \cdot V(w^1, sp^{1,tot}, w^2, sp^{2,tot}) \}, \quad (4)$$

$$C_h(s^1, s^2) = -c_h \cdot (w^1 + w^2),$$

$$C_f(s^1, s^2) = \begin{cases} -c_f & s^i + 1 = K_{max} - 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$C_{pm}(s^1, s^2) = \begin{cases} -c_{pm} & s^i + 1 = K_{max}, \\ 0 & \text{otherwise.} \end{cases}$$

Equations (1-4) consider the cases when both machines are operating normally.

For clarity in this first case, we describe some of the terms. The first two terms in (1) is the operating cost. The subsequent terms arise due to the arrival process, two service processes, two deterioration processes, respectively. Following a service event at machine 1, a job enters the buffer at machine 2. The next state is thus $((w^1 - 1)^+ + I(w^1 = L_{max}), s^1, w^2 + I(w^1 \geq 1), s^2)$. If no jobs are in a machine buffer when a service event occurs, the state remains the same. The remaining term in (1) details the virtual transitions with total probability $1 - (\lambda + \sum_i \mu^i + \sum_i \sigma^i)/D$. Equation (2) details the case when there is a preventive maintenance intention at machine 2. Note that since we assume non-preemptive maintenance, a transition to a deterioration status of k_{max} is only allowed when there is service event at machine 2 (μ^2). The others are similar.

Case2: $s^1 \geq k_{max} - 1, s^2 < k_{max} - 1$

$$Q^{0,0}(w^1, s^1, w^2, s^2) = C_h(w^1, w^2) + \frac{\gamma(\sigma^1 + \sigma^2)}{D} C_f(s^1, s^2) + \frac{\gamma}{D} \{ \lambda \cdot V(w^1 + 1, s^1, w^2, s^2) + \mu^2 \cdot V(w^1, s^1, (w^2 - 1)^+, s^2) + \sigma^2 \cdot V(w^1, s^1, w^2, s^2 + 1) + \theta_k^1 \cdot V(w^1, 0, w^2, s^2) + (D - (\lambda + \mu^2 + \sigma^2)) \cdot V(w^1, s^1, w^2, s^2) \}, \text{ for } k \in \{f, pm\},$$

$$Q^{0,1}(w^1, s^1, w^2, s^2) = C_h(w^1, w^2) + \frac{\gamma(\sigma^1 + \sigma^2)}{D} C_f(s^1, s^2) + \frac{\gamma}{D} \{ \lambda \cdot V(w^1 + 1, s^1, w^2, sp^{2,arr}) + \mu^2 \cdot V(w^1, s^1, (w^2 - 1)^+, k_{max}) + \sigma^2 \cdot V(w^1, s^1, w^2, sp^{2,det}) + \theta_k^1 \cdot V(w^1, 0, w^2, sp^{2,rep}) + (D - (\lambda + \mu^2 + \sigma^2 + \theta_k^2)) \cdot V(w^1, s^1, w^2, sp^{2,tot}) \}, \text{ for } k \in \{f, pm\},$$

The case $s^1 \geq k_{max} - 1, s^2 < k_{max} - 1$ occurs when the machine 1 is in repair and the machine 2 is available. Once machine is in repair, service and deterioration are not allowed. Machine 1 returns to a zero deterioration state with probability $\frac{\theta_k^1}{D}$. A PM decision is not allowed in the repair state. The state will transit to $(w^1, s^1, (w^2 - 1)^+, k_{max})$ if a PM is indicated and next event is a service.

Case 3 ($s^1 < k_{max} - 1, s^2 \geq k_{max} - 1$) and Case 4 ($s^i \geq k_{max} - 1, i \in \{1,2\}$) behave similarly.

3 NUMERICAL STUDIES

In this section, we first provide the parameter settings for our experiments. Numerically based observations on the structure of optimal policies are reviewed next. Finally, we report the results of simulation studies to assess the performance of the optimal policy for the discounted operating cost and mean cycle time relative to a policy that treats each machine in isolation. The results demonstrate that considering the joint control of PMs at two machines in tandem is superior to treating them as if they were isolated.

3.1 Baseline Setting

Recall that we truncated the maximum buffer space at each machine to L_{max} . For our studies we set $L_{max} = 100$ and $K_{max} = 9$. We use value iteration to obtain solutions. Each experiment thus takes roughly 50 mins on an i5-6600 computer with 8GB RAM.

Table 1 shows the baseline settings for our experiments. The rates for the events are given as are the costs. Since the maximum deterioration status was set at 9, values 8 and 9 indicate that a machine is in repair or PM, respectively. Note that the mean duration of a repair following failure is twice the mean of a PM event. Also note that we set $c_{pm} = 0$ so that the structure of optimal policies are clearly visible. (Using other values provides similar behavior but not as distinctly visible on graphs.)

Table 1: Parameter settings for the baseline experiment.

Parameter	λ	μ^1	μ^2	δ^1	δ^2	θ_{PM}^1	θ_{PM}^2	θ_F^1	θ_F^2
Rate	0.2	0.32	0.32	0.04	0.04	0.08	0.08	0.04	0.04
Mean	5	3.125	3.125	25	25	12.5	12.5	25	25
Parameter	Q1	Q2	M1	M2	c_f	c_{PM}	c_h		
Constant	100	100	10	10	20000	0	3		

From the baseline settings in Table 1, we varied one parameter in each experiment with the other settings held fixed. Table 2 shows the parameters varied, the range over which they were varied, and the number of cases. There were 50 resulting experiments.

Table 2: Variation of experiment from baseline settings.

Varying Parameter	Values			No. Cases
	Min	Max	Increment	
Arrival rate	0.21	0.3	0.01	10
Deterioration rate	0.01	0.1	0.01	10
Repair rate from PM	0.02	0.2	0.02	10
Holding cost	0.01	0.1	0.01	10
Failure cost	20	200	20	10
Total				50

3.2 Structural Behavior

In our 50 experiments, the following four properties were always observed.

Numerical observation 1: Threshold behavior as a function of w^1, s^1 . Consider fixed values for w^2 and s^2 . For all $w^1 \in \{0, 1, \dots, L_{max}\}$, there exists an $\alpha(w^1, s^2, w^2)$ such that

$$a^1(w^1, s^1, s^2, w^2) = \begin{cases} 1, & s^1 \geq \alpha(w^1, s^2, w^2), \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, for $w^1 \geq 1$, $\alpha(w^1, s^2, w^2)$ is non-decreasing with respect to w^1 .

Figure 1 depicts this behavior. There is a clear threshold below which a PM is not indicated and above which a PM is intended. For example, as in Figure 1, when $w^1 = 5$, $\alpha(w^1, s^2, w^2) = 2$. If the deterioration status of machine 1 is less than or equal to two, no PM is indicated. On the other hand, if the deterioration status of machine 1 is greater or equal to three, a PM is indicated. This behavior mimics that of Yao et al (2003) for a single machine system.

8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
6	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
4	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
3	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20

Figure 1: Structural behavior of the optimal policy at (w^1, s^1) .

Similar behavior was observed at machine 2 with w^1 and s^1 held constant.

Numerical observation 2: Threshold behavior as a function of w^2, s^2 . Consider fixed values for w^1 and s^1 . For all $w^2 \in \{0, 1, \dots, L_{max}\}$, there exists an $\alpha(w^1, s^1, w^2)$ such that

$$a^1(w^1, s^1, s^2, w^2) = \begin{cases} 1, & s^2 \geq \alpha(w^1, s^1, w^2), \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, for $w^2 \geq 1$, $\alpha(w^1, s^1, w^2)$ is non-decreasing with respect to w^2 .

Numerical observation 3: Threshold behavior as a function of w^2, s^1 . Consider fixed values for w^1 and s^2 . For all $w^2 \in \{0, 1, \dots, L_{max}\}$, there exists an $\alpha(w^1, s^1, w^2)$ such that

$$a^1(w^1, s^1, s^2, w^2) = \begin{cases} 1, & s^1 \geq \alpha(w^1, s^2, w^2), \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, for $w^2 \geq 1$, $\alpha(w^1, s^1, w^2)$ is non-decreasing with respect to w^2 .

8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
6	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
5	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	
4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20

Figure 2: Structural behavior of the optimal policy at (w^2, s^1) .

With w^1 and s^2 fixed, the s^1 threshold for PM intention at machine 1 is non-increasing in the number of jobs at machine 2. As the number of jobs at machine 2 grows large, machine 1 prefers to initiate PMs at

comparatively lower deterioration status. Intuitively, when machine 2 has much work, there is less risk for initiating a PM at machine 1.

Numerical observation 4: Threshold behavior as a function of w^1, s^2 . Consider fixed values for w^2 and s^1 . For all $w^1 \in \{0,1, \dots, L_{max}\}$, there exists an $\alpha(w^1, s^2, w^2)$ such that

$$a^1(w^1, s^1, s^2, w^2) = \begin{cases} 1, & s^2 \geq \alpha(w^1, s^1, w^2), \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, for $w^1 \geq 1$, $\alpha(w^1, s^2, w^2)$ is non-decreasing with respect to w^2 .

8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
7	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20

w^1

Figure 3: Structural behavior of the optimal policy at (w^1, s^2) .

In all of our studies, the PM policy at machine 2 was not influenced by the number of jobs at machine 1.

3.3 Evaluation of the optimal policy

We assess the performance of the optimal PM policy relative to the use of the optimal PM policy for such a machine in isolation (and applied separately to each machine). The discrete event simulator is utilized for the experiments which is developed based on Choi and Kang(2014). We consider mean cycle time and operating cost. Fifty different simulation experiments were conducted for each policy (consisting of 30 replications each) with the parameter settings from Tables 1 and 2. Each simulation runs for 100,000 time units and warm-up periods (10% of arrivals from the beginning) are excluded from the results. The summary results are provided in Table 3. There, the results are averaged over all 50 simulation runs.

The discounted operating costs averaged over 47 simulations are provided in Table 3. We disregarded the data from arrivals rate 0.28, 0.29, 0.3 as the number of jobs in the buffer was near 100 and we did not trust the results due to our queue length truncation. There was a 10.55% reduction in total discounted operating costs. There were fewer failures in the joint PM policy and 3% more PM events.

Cycle time is an important manufacturing metric. Although the objective of the current model is discounted operating costs, it is worth investigating the mean of the cycle time. The MDP based joint PM policy exhibited a 11.4% reduction in mean cycle time compared to the isolated machine PM policy

Table 3: Performance of the joint PM policy relative to isolated PM policies

		PM policy of the joint MDP model	Isolated PM policy applied to both machines	Difference (%)
Discounted operating costs		21863.27	24443.07	10.55
CT	Mean	25.02	28.23	11.36
Mean queue length		5.568433	6.44	11.55
Number of failures per simulation run		3.52	5.89	40.2
Number of PMs per simulation run		1319.89	1360.1	2.95

4. CONCLUDING REMARKS

Observing the absence of studies addressing the structure of optimal PM policies in networks of failure prone queues, we investigated a prototype tandem system. Two exponential machines work on jobs in tandem. For computational tractability, we considered a truncated state space. The machines were subject to deterioration based failure and repair. PM events could be initiated to return a machine to a like new status. We developed an MDP model using uniformization seeking a PM policy minimizing the long-term discounted costs.

For this prototype tandem system, we observed the following numerical behaviors. If the queue length and deterioration of the first stage was held fixed, the structure of the optimal PM policy at the second stage behaved similarly to such a queue in isolation. There was a deterioration threshold above which a PM was dictated and this threshold was non-decreasing in the queue length. Similar structure existed at the first stage when holding the state of the second stage fixed. However, these thresholds can be a function of the state of the other queue and may exhibit their own non-increasing tendency.

We further assessed the performance of the optimal PM policy obtained from the MDP model that considers both machines jointly with a myopic policy that is optimal for a single machine in isolation. The optimal joint PM policy was superior by about 10% in discounted operating costs as well as mean cycle time.

There are several directions for future work. First, as this paper explores numerical tendencies of optimal policies, it would be very interesting to theoretically prove the existence of such structural behaviors. Second, longer tandem lines should be studied. The exact analysis conducted here is limited by computational tractability issues. As such, exploiting the observations of this work to develop heuristics could allow for the treatment of larger systems. Third, the queueing system with general arrival and service process should be considered. We assume all interarrival times, service times, inter-deterioration times, repair times, and PM times are IID exponentially distributed. Investigating general distribution will make the problem realistic.

REFERENCES

- Barde, S., H. Shin, and S. Yacout. 2016. "Opportunistic preventive maintenance strategy of a multi-component system with hierarchical structure by simulation and evaluation". In 2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA), September 6th-9th, Berlin, Germany, 1-8.
- Barde, S., S. Yacout, and H. Shin. 2019. "Optimal preventive maintenance policy based on reinforcement learning of a fleet of military trucks". *Journal of Intelligent Manufacturing* 30(1):147-161.
- Borrero, J. S. and R. Akhavan-Tabatabaei. 2010. "A Markov decision process model for optimal policy making in the maintenance of a single-machine single-product toolset". In *Proceedings of the 2010 Winter Simulation Conference*, edited by J. Hagan, E.

- Yücesan, M. Fu, B. Hohansson, S. Jain, J. Montoya-Torres. 2594-2605. Baltimore, Maryland, Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Cai, Y., J. J. Hasenbein, E. Kutanoglu, and M. Liao. 2013. "Single-machine multiple-recipe predictive maintenance". *Probability in the Engineering and Informational Sciences* 27(2):209-235.
- Cassady, C. R. and E. Kutanoglu. 2005. "Integrating preventive maintenance planning and production scheduling for a single machine". *IEEE Transactions on reliability* 54(2):304-309.
- Choi, B. K. and D. Kang, 2014. "How to develop your own simulators for discrete-event systems". In *Proceedings of the 2014 Winter Simulation Conference*, edited by S. J. Buckley, J. A. Miller, A. Tolk, L. Yilmaz, S. Y. Diallo and I. O. Ryzhov. 147-161. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.
- Das, T. K., A. Gosavi, S. Mahadevan, and N. Marchallick. 1999. "Solving semi-Markov decision problems using average reward reinforcement learning". *Management Science* 45(4):560-574.
- Kaufman, D. L. and M. E. Lewis. 2007. "Machine maintenance with workload considerations". *Naval Research Logistics (NRL)* 54(7):750-766.
- Krishna, D. G. 2012. "Preventive maintenance of wind turbines using Remote Instrument Monitoring System". In *2012 IEEE Fifth Power India Conference*, December 19th-22th, Murthal Delhi-National Capital Region, India, 1-4.
- Lippman, S.A. 1975. "Applying a new device in the optimization of exponential queuing systems". *Operations Research* 23(4):687-710.
- Laggoune, R., A. Chateaufneuf, and D. Aissani. 2009. "Opportunistic policy for optimal preventive maintenance of a multi-component system in continuous operating units". *Computers & Chemical Engineering* 33(9):1499-1510.
- Moghaddam, K. S. and J. S. Usher. 2011. "Preventive maintenance and replacement scheduling for repairable and maintainable systems using dynamic programming". *Computers & Industrial Engineering* 60(4):654-665.
- Nicolai, R. P. and R. Dekker. 2008. "Optimal maintenance of multi-component systems: a review". In *Complex system maintenance handbook*, 263-286. London: Springer
- Ramírez-Hernández, J. A. and E. Fernandez. 2010. "Optimization of preventive maintenance scheduling in semiconductor manufacturing models using a simulation-based approximate dynamic programming approach". In *49th IEEE Conference on Decision and Control (CDC)*. Dec 15th-17th, Atlanta, USA, 3944-3949.
- Sim, S. H. and J. Endrenyi. 1988. "Optimal preventive maintenance with repair". *IEEE Transactions on Reliability* 37(1): 92-96.
- Yao, X., X. Xie, M. C. Fu, S. I. Marcus, E. Fernández-Gaucherand, and S. Krishnaswamy. 2003. "Optimal preventive maintenance policies for unreliable queueing/production systems with applications to semiconductor manufacturing." *TECHCON 2003*:25-27.
- Yao, X., X. Xie, M. C. Fu, and S. I. Marcus. 2005. "Optimal joint preventive maintenance and production policies". *Naval Research Logistics (NRL)* 52(7):668-681.

AUTHOR BIOGRAPHIES

TAEHYUNG KIM is an PhD candidate in Department of Industrial and Systems Engineering, KAIST, South Korea. He hold B.S. degrees in Industrial and Systems Engineering, Business and Technology Management, and M.S. degree in Industrial and System Engineering from KAIST, South Korea. His email address is rlaxo0604@kaist.ac.kr

JAMES R. MORRISON is an Associate Professor in the Department of Industrial and Systems Engineering, KAIST, South Korea. He holds B.S. degrees from the University of Maryland at College Park, USA and M.S. and Ph.D. degrees in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign, USA. He is a co-Chair of the IEEE RAS Technical Committee on Semiconductor Manufacturing Automation. His email address is james.morrison@kaist.edu