

Machine Maintenance with Workload Considerations

David L. Kaufman

*Department of Industrial Engineering
University of Pittsburgh
1048 Benedum Hall, Pittsburgh, PA 15261
davidlk@engr.pitt.edu
412-383-5103*

Mark E. Lewis

*School of Operations Research and Information Engineering
Cornell University
226 Rhodes Hall, Ithaca, NY 14853
mel47@cornell.edu
607-255-0757*

submitted June 19, 2006, revised December 11, 2006,
revised April 20, 2007

Abstract

Machine maintenance is modeled in the setting of a single-server queue. Machine deterioration corresponds to slower service rates and failure. This leads to higher congestion and an increase in customer holding costs. The decision-maker decides when to perform maintenance, which may be done preemptively; before catastrophic failures. Similar to classic maintenance control models, the information available to the decision-maker includes the state of the server. Unlike classic models, the information also includes the number of customers in queue. Considered are both a repair model and a replacement model. In the repair model, with random replacement times, fixed costs are assumed to be constant in the server state. In the replacement model, both constant and variable fixed costs are considered. It is shown in general that the optimal maintenance policies have switching curve structure that is monotone in the server state. However, the switching curve policies for the repair model are not always monotone in the number of customers in the queue. Numerical examples and two heuristics are also presented.

1 Introduction

Traditional maintenance control models include what van Dijkhuizen and van Harten [36] refer to as the “technical” state of the system (e.g., machine failure characteristics). These models do not however typically include other important characteristics associated with the “operating” states of manufacturing systems (e.g., workload, inventory, due dates, etc.). For many systems, in particular controlled queueing systems, it is precisely the operational properties of the system that are crucial when making maintenance decisions. For example, when the system is heavily loaded, is it better to begin preventative maintenance or to postpone repair until some of the congestion subsides? This question is complicated by the fact that as machines deteriorate, so does their ability to produce acceptable products; service rates become slower. The goal of this paper is to gain a better understanding of questions related to maintenance when the decision-maker has access to operating information and the machine’s ability to produce deteriorates over time. The main contributions of the paper are as follows (under various assumptions):

- There exists an optimal maintenance policy that is monotone in the state of the server.
- On the other hand, the optimal maintenance policy need not be monotone in the congestion level.

From a managerial perspective, this implies that the search for (and the storage of) optimal policies can be reduced since only switching curve policies need to be considered. The second result serves as somewhat of a warning since these switching curves need not be monotone in the number of customers in the system. Moreover, we point out that in all of the examples that exhibit the non-monotonicity it is such that repair is optimal when the system is empty, do not repair is optimal when the system has a single customer, and repair is again optimal as the number of customers increases. Intuitively, this means that when the system is empty, it is optimal in some cases to take advantage of the server idleness.

Considered are two models: **repair** and **replacement**. The repair model has repair times that are random with positive mean. The replacement model is similar but has instantaneous replacements. There are fixed costs associated with repairs or replacements and costs are included that explicitly depend on the operational state of the queue: linear customer holding costs.

Several textbooks include the well-known, fundamental maintenance control models of Derman [10, 11] and Ross [24]. The deterioration of a machine is modeled as a discrete-time Markov chain. The machine in state s deteriorates to state j in the next stage with some probability. At each stage, the decision-maker chooses one of two actions: either replace or do not replace (do nothing). Derman [10] considers a finite

number of states with forced replacement at the worst state. Ross [24] considers a countable number of states. Both show sufficient conditions for the optimal policy to be of “control limit” form. This simple form aids in both computation and implementation of the optimal control. The optimality of control limit policies in replacement models has been addressed more recently by So [33]. The reader is referred to the survey papers [22], [29], [9], and [38].

Queueing models incorporating maintenance control decisions include those of Federgruen and So [12, 13]. In their models, a single server is either ‘up’ or ‘down.’ While the server is up, breakdowns occur according to a Poisson process. In [12], when the server goes down, the decision-maker chooses either to start repair immediately or to postpone. Optimal policies are characterized by a single threshold: a repair is initiated if and only if the number of customers in the system exceeds the threshold. In [13], there are two different repair options. Roughly, one is faster and the other is slower but less expensive. Again, optimal policies are monotone: the more expensive, faster repair is chosen if and only if the number of customers in the system exceeds some threshold. Also related to these papers is [14] where a vacation model is considered and again monotone threshold policies are optimal, ‘D-policies.’ These papers are fundamentally different than the models we consider because they do not include multiple functioning server states, and, therefore, preventive maintenance is precluded.

The non-queueing production/maintenance literature can be split into two categories: 1) papers that study how failures or fixed maintenance policies affect traditional production/inventory policies, but do not consider optimal maintenance control; 2) papers that include optimal maintenance control decisions, but that may or may not jointly incorporate optimal production decisions. The present work is related to category 2. The papers in category 2 include [18], [7], and [34]. These models each assume a single functioning state. Other related papers with multiple functioning states (but perhaps not multiple yield rates) include [36], [35], [3], [31, 32], and [30]. As for category 1, Irvani and Duenyas [15] consider a single-machine make-to-stock system. Included in the model are multiple functioning states for the facility where production rates change with the deterioration of the system. By comparing optimal heuristic policies based on the production rates of the facility in different deterioration states to heuristics based on average production rates, Irvani and Duenyas [15] argue that it is important for firms to collect information on how the production rates change as a function of deterioration states. The models we consider also incorporate yield information, but in the setting of a queue. For other non-queueing production/maintenance literature, the reader is referred to the references in [15] and [42, 43].

The queueing network studied in Seshadri and Pinedo [28] captures the effects of productivity loss through service rate reductions and downtime by allowing the work per unit time to depend on an efficiency factor. The queueing models most closely related to ours are those of Koyanagi and Kawai [17] and Yao, in Chapter 3 of [42]. Both of these works include multiple server states. An important assumption that is made in [17] is that at the beginning of maintenance, the customers in the system are rejected and so are the customers arriving during maintenance. We make no such assumptions. Yao [42] also includes multiple server deterioration states; however, they are modeled differently. The lifetime of the server depends on the number of customers served since the last (non-instantaneous) repair. If the technical state of the server, the number of service completions since last repair, reaches some state L , then repair is mandatory. More important than this modeling difference is the fact that Yao [42] assumes a single service rate. The fact that deteriorations may result in reduced service rates is an important characteristic for the models we consider. Conditions are shown in [42] for optimal policies that are monotone in the technical state of the server: for a fixed number in queue, q , there exists an optimal server-state threshold $s^*(q)$ such that it is optimal to start repair if and only if the state of the server is worse than $s^*(q)$. Similar results are shown here in Sections 3 and 4. Yao [42] does not address whether or not the optimal policy for his model is monotone in q .

Reliability issues have recently received increased attention in the controlled queueing research community as Flexible Manufacturing Systems (FMS) have become more prevalent. An important benefit of flexible servers is that they can alleviate congestion due to failures and loss of resources. Earlier works on reliability and FMS include [37], [39], and [21]. More recently, the models of Andradottir et al. [1] and Wu et al. [40, 41] incorporate reliability in the control of agile servers; however, they do not include preventive maintenance of the servers. They model servers that are either ‘up’ or ‘down,’ and they do not incorporate maintenance control decisions. The focus in the present work is instead on models of maintenance control.

The remainder of the paper is organized as follows. Section 2 contains formal definitions of the models. Results for the repair model are given in Section 3; the main results are presented in Theorem 3.2. Section 4 contains the results for the replacement model; the main results are Theorem 4.2 and Proposition 4.10. Additional examples and two heuristics are considered in Section 5. Section 5 also concludes the paper.

2 Model Formulations

Consider a single queue being served by a single machine (server). While the server state is s , the machine works at rate μ_s , $s \in \{0, 1, \dots, B\}$. It is assumed that server state zero is a repair state with $\mu_0 \equiv 0$, and $0 < \mu_1 \leq \mu_2 \leq \dots \leq \mu_B < \infty$. Customers arrive to the queue according to a Poisson process with

rate $\lambda > 0$. The service requirements for the customers are i.i.d. and exponentially distributed with mean 1. Note that this implies that the service time for a customer that receives service at rate μ throughout is exponentially distributed with rate μ .

So long as the machine is not being repaired, the machine in server state s deteriorates asynchronously (independent of the work it is doing) down to server state $(s - 1)$ in a random amount of time. Assume that deterioration times for such transitions out of server state s are i.i.d. and exponentially distributed with means $1/m_s > 0$. If repairs are not done preemptively, as soon as the server state reaches 0, repairs are initiated to bring the server state back up to B . The replacement model has instantaneous repair times while the repair model has random, not necessarily exponentially distributed, repair times. Rather, assume that they are i.i.d. with general distribution $G(\cdot)$ that has finite first and second moments. Denote the mean repair time by $1/m_0 > 0$.

The costs associated with operating the system are holding costs h per customer per unit time and fixed repair/replacement costs $K \geq 0$. Repairs can be initiated preemptively; that is, before the machine fails by reaching server state 0. It is assumed that both the queue length and the server state are perfectly observable to the decision maker at all times.

The problem is modeled as a semi-Markov decision process (SMDP). Let Π be the set of all non-anticipating policies. A policy $\pi \in \Pi$ prescribes when to repair the server, given the number of customers in the system (the queue length process) and the state of the server (the server state process). For a fixed policy $\pi \in \Pi$ and $t \geq 0$, let $Q^\pi(t)$ and $S^\pi(t)$ represent the queue length process and the server state process at time t , respectively. Denote the set of decision epochs (under π) by $\mathbb{D} = \{\sigma_n, n \geq 0\}$. The decision epochs during non-repair times form a Poisson process with rate $\Psi \equiv \lambda + \sum_{s=1}^B (\mu_s + m_s)$. This follows from the exponential distribution assumptions and *uniformization* (see [19]). Without loss of generality, assume that $\Psi = 1$. The (random) times \mathbb{D} then include server deteriorations, repair completions, and, for non-repair times, customer arrivals and departures, and *dummy* transitions due to uniformization. The state of the “embedded” process after the n th transition is given by $X_n^\pi = (Q^\pi(\sigma_n), S^\pi(\sigma_n))$, $n \geq 0$. That is, the state space $\mathbb{X} = \mathbb{Z}^+ \times \{0, 1, \dots, B\}$, where \mathbb{Z}^+ is the set of non-negative integers. A change in the server state process corresponds to a change for the embedded process. However, a change in the queue length process may not coincide with a change in the embedded process since the queue length may increase during a repair. From an optimization standpoint, this is of no consequence since the two processes coincide at decision epochs.

The sets of feasible actions for state $x = (q, s)$ are

$$\mathbb{A}_x = \begin{cases} \{R\} & \text{if } s = 0, \\ \{D, R\} & \text{else,} \end{cases}$$

where R indicates repair and D indicates do not repair. Let

$$k(s, a) = \begin{cases} 0 & \text{if } a = D, \\ K & \text{if } a = R. \end{cases} \quad (2.1)$$

Let $\{(X_n^\pi, A_n^\pi), n \geq 0\}$ be the sequence of states of the embedded process and the actions chosen. Suppose δ_t is the number of decision epochs up to time t ; it is assumed that $t = 0$ is a decision epoch. Define

$$g_t^\pi(x) = \mathbb{E}_x \left\{ \sum_{i=0}^{\delta_t-1} k(S^\pi(\sigma_i), A_i^\pi) + \int_0^t hQ^\pi(u)du \right\}, \quad (2.2)$$

$$v_{n,\theta}^\pi(x) = \mathbb{E}_x \left\{ \sum_{i=0}^{n-1} \left[e^{-\theta\sigma_i} k(S^\pi(\sigma_i), A_i^\pi) + \int_{\sigma_i}^{\sigma_{i+1}} e^{-\theta u} hQ^\pi(u)du \right] \right\}, \quad (2.3)$$

where $\theta > 0$ is the discount factor. The expression (2.2) defines the expected total cost incurred by time t when the initial state is x . Likewise, (2.3) defines the expected total discounted cost incurred by the n th decision epoch. The expressions in (2.4) and (2.5) below define the infinite horizon expected discounted cost and average cost under π , respectively.

$$v_\theta^\pi(x) = \lim_{n \rightarrow \infty} v_{n,\theta}^\pi(x), \quad (2.4)$$

$$g^\pi(x) = \limsup_{t \rightarrow \infty} \frac{g_t^\pi(x)}{t}. \quad (2.5)$$

Define the optimal values $v_\theta(x) = \inf_{\pi \in \Pi} v_\theta^\pi(x)$, $g(x) = \inf_{\pi \in \Pi} g^\pi(x)$, where any policy that achieves the infimum of the respective criteria is deemed optimal.

In either model considered below, we provide conditions under which there exists an optimal policy that is monotone in the server state s . This type of policy is often called a *switching curve* policy.

Definition 2.1 Suppose in each state $x \in \mathbb{X}$ there are but two actions, D and R . A (deterministic) stationary policy is a *switching curve policy* if it may be described by a curve in \mathbb{X} that separates \mathbb{X} into two connected

regions. In one region the policy calls for action D to be used, while in the other region action R is used. Furthermore, a switching curve policy is called a *monotone switching curve policy* if the curve dividing \mathbb{X} into two regions is monotone.

Although typically more general, for the purposes of this paper we say that a policy that is monotone in s has *switching curve structure that is monotone in s* .

3 The Repair Model

This section is devoted to the repair model. The main results are that there exists an optimal policy that is monotone in the server state s , but not necessarily in the number of customers in the system. That is, there exists an optimal policy that has switching curve structure that is monotone in s (see Figure 1).

For $\theta \geq 0$, let

$$C_\theta(q, a) = \begin{cases} hq/(\theta + 1) & \text{if } a = D, \\ K + h \int_0^\infty \int_0^u [e^{-\theta t} (q + \lambda t) dt] G(du) & \text{if } a = R, \end{cases} \quad (3.1)$$

$$p_\theta(j) = \int_0^\infty e^{-\theta t} \frac{e^{-\lambda t} (\lambda t)^j}{j!} G(dt). \quad (3.2)$$

The function $C_\theta(q, a)$ is the one-stage expected discounted cost incurred when the process is in state (q, s) and action a is chosen. Recall that $\Psi = 1$. In (3.1), when $a = R$, λt is the expected number of customer arrivals within t time units. The function $p_\theta(j)$ is the discounted probability that exactly j customers arrive during a repair. For the discounted cost case, with $\theta > 0$ and $(q, s) \in \mathbb{X}$, define the discounted cost optimality equations (DCOE):

$$\begin{aligned} f_\theta^D(q, s) &= C_\theta(q, D) + \frac{1}{\theta + 1} [\lambda v_\theta(q + 1, s) + \mu_s v_\theta((q - 1)^+, s) \\ &\quad + m_s v_\theta(q, s - 1) + (1 - \lambda - \mu_s - m_s) v_\theta(q, s)], \quad s \in \{1, 2, \dots, B\}, \end{aligned} \quad (3.3)$$

$$f_\theta^R(q) = C_\theta(q, R) + \sum_{j=0}^{\infty} p_\theta(j) v_\theta(q + j, B), \quad (3.4)$$

$$v_\theta(q, s) = \min \{ f_\theta^D(q, s), f_\theta^R(q) \}, \quad s \in \{1, 2, \dots, B\}, \quad (3.5)$$

$$v_\theta(q, 0) = f_\theta^R(q), \quad (3.6)$$

where y^+ represents the positive part of y . Similarly, define the average cost optimality inequalities (ACOI):

$$\begin{aligned} f^D(q, s) &= -g + C_0(q, D) + \lambda w(q + 1, s) + \mu_s w((q - 1)^+, s) \\ &\quad + m_s w(q, s - 1) + (1 - \lambda - \mu_s - m_s) w(q, s), \quad s \in \{1, 2, \dots, B\}, \end{aligned} \quad (3.7)$$

$$f^R(q) = -\frac{g}{m_0} + C_0(q, R) + \sum_{j=0}^{\infty} p_0(j) w(q + j, B), \quad (3.8)$$

$$w(q, s) \geq \min \{f^D(q, s), f^R(q)\}, \quad s \in \{1, 2, \dots, B\}, \quad (3.9)$$

$$w(q, 0) \geq f^R(q). \quad (3.10)$$

It is proved next that under a sufficient stability condition there exists w and g that satisfy (3.7) - (3.10). Any stationary policy that minimizes the right-hand side of (3.5) is θ -discounted cost optimal with optimal values v_θ . Furthermore, any stationary policy minimizing the right hand side of (3.9) is average cost optimal and $g(x) = g$ for every initial state x .

To arrive at the sufficient stability condition, consider the following policy. For $\ell \in \{1, 2, \dots, B\}$ and $(q, s) \in \mathbb{X}$ we call the stationary policy ϕ_ℓ that uses the decision rule

$$\phi_\ell(q, s) = \begin{cases} D & \text{if } s \geq \ell, \\ R & \text{if } s < \ell \end{cases}$$

the *server-state threshold policy* with *threshold* ℓ . Note that ϕ_ℓ is a switching curve policy where the action is repair if and only if $s < \ell$ (independent of q). The next proposition states that the limiting queue length process is finite under policy ϕ_ℓ provided that the arrival rate is less than the average service capacity.

Proposition 3.1 *Let $\ell \in \{1, 2, \dots, B\}$. If*

$$\lambda < \frac{\sum_{s=\ell}^B \mu_s / m_s}{1/m_0 + \sum_{s=\ell}^B 1/m_s}, \quad (3.11)$$

then there exists an invariant probability measure φ such that $\lim_{t \rightarrow \infty} \mathbb{E}_x[Q^{\phi_\ell}(t)] = \mathbb{E}_\varphi[Q^{\phi_\ell}(0)] < \infty$, and $g^{\phi_\ell}(x) < \infty$, for all $x \in \mathbb{X}$. Furthermore, for the Markov chain induced by ϕ_ℓ , the class of states $\mathbb{X}_\ell = \{(q, s) | q \in \mathbb{Z}^+, s \in \{\ell - 1, \ell, \dots, B\}\}$ are ergodic (aperiodic and positive recurrent) and the states $\mathbb{X} \setminus \mathbb{X}_\ell$ are transient.

Proof: See Section 7.1 in the Appendix. ■

Theorem 3.2 *For either the discounted cost criterion or the average cost criterion suppose*

$$\lambda < \max_{l \in \{1, 2, \dots, B\}} \left\{ \frac{\sum_{s=l}^B \mu_s / m_s}{1/m_0 + \sum_{s=l}^B 1/m_s} \right\}. \quad (3.12)$$

The following then hold.

1. *There exists an optimal stationary policy that has switching curve structure that is monotone in s .*
2. *There may not exist an optimal switching curve policy that is a monotone switching curve policy (monotone in both q and s).*

The proof of Theorem 3.2 is divided into several pieces. The first states the intuitive result that it is better to start in a state with fewer customers and a higher service rate. This implies the first result of the theorem for the discounted cost case and implies the existence of a “minimal” state $(0, B)$ where the value functions (under each criterion) are minimized.

Proposition 3.3 *For $(q_1, s_1), (q_2, s_2) \in \mathbb{X}$, suppose $q_1 \leq q_2$ and $s_1 \geq s_2$. Then, $v_\theta(q_1, s_1) \leq v_\theta(q_2, s_2)$.*

Proof: The result is proved via a sample-path argument by considering two processes. Process 1, which starts with a better server state and fewer initial customers, mimics the actions of the policy governing Process 2 by repairing at the same times (the policy employed by Process 1 is potentially non-stationary). Furthermore, Process 2 only sees customer departures if Process 1 sees potential (there may not be a customer to serve) departures. The result is that both processes see the same fixed costs while Process 2 sees higher holding costs.

Suppose Process 1 starts in state (q_1^0, s_1^0) at time 0 (the superscript ‘0’ indicates time 0). Process 2 starts in state (q_2^0, s_2^0) where $q_1^0 \leq q_2^0$ and $s_1^0 \geq s_2^0$. Process 2 follows an arbitrary policy $\pi_2 \in \Pi$. Denote the policy under which Process 1 mimics Process 2 repairs by π_1 . We will now show how to construct Process 2 so that $\pi_1 \in \Pi$; recall it needs to be non-anticipating.

Assume that the two processes are defined on the same probability space. Both processes see the same customer arrivals and sequence of repair times. Recall that there are no decision epochs during repair times. For non-repair times, deteriorations and customer departures are constructed by thinning a Poisson process with rate $\Psi = 1$. We must construct the processes so that the probability of service and server state

transitions in Process 1 (2) are μ_{s_1} (μ_{s_2}) and m_{s_1} (m_{s_2}), respectively. The dummy transitions should then be $1 - \mu_{s_i} - m_{s_i}$ for Process i , $i = 1, 2$. Consider first the case when $s_2^0 > 0$. Suppose there is an event for the Poisson process at time t and that at $t-$ (just prior to the event) Process 1 is in state (q_1, s_1) and Process 2 is in state (q_2, s_2) . Suppose for now that $s_1 > s_2 > 0$. With probability $(1 - \mu_{s_1} - m_{s_1} - m_{s_2})$ there is a dummy transition for both processes; with probability m_{s_1} there is a deterioration in Process 1 and a dummy transition for Process 2; with probability m_{s_2} there is a deterioration in Process 2 and a dummy transition for Process 1; with probability μ_{s_1} there is a potential departure for Process 1. In this case, with probability μ_{s_2}/μ_{s_1} there is a potential departure for Process 2 as well, otherwise there is a dummy transition. Note that the probability of a dummy transition in Process 2 is $1 - \mu_{s_1} - m_{s_1} - m_{s_2} + m_{s_1} + \mu_{s_1} \left(1 - \frac{\mu_{s_2}}{\mu_{s_1}}\right) = 1 - \mu_{s_2} - m_{s_2}$. Similarly, for Process 1 there is a dummy transition with probability $1 - \mu_{s_1} - m_{s_1} - m_{s_2} + m_{s_2} = 1 - \mu_{s_1} - m_{s_1}$ as desired. For $s_1 = s_2 > 0$, a potential departure occurs in both processes with probability μ_{s_1} and a deterioration occurs in both Process 1 and Process 2 with probability m_{s_1} . In this case, since π_1 calls for repair only when π_2 repairs, when $s_1 = s_2$ the two processes are coupled in their server state processes; both processes see the same deteriorations and potential customer departures from that time forward. It is not necessary to consider cases where $s_1 < s_2$ because these states are never reached.

In either case, under the construction above, Process 1 always maintains a service rate that is at least as high as that of Process 2. Moreover, Process 2 never has a customer departure without a corresponding potential departure for Process 1. Thus, Process 1 incurs total holding costs that are less than or equal to those incurred by Process 2. The total repair costs are the same for both processes. Thus, $v_\theta^{\pi_1}(q_1^0, s_1^0) \leq v_\theta^{\pi_2}(q_2^0, s_2^0)$. Policy π_1 is a (randomized) non-anticipating policy and hence $\pi_1 \in \Pi$. Since π_2 is an arbitrary policy, it follows that $v_\theta(q_1^0, s_1^0) \leq v_\theta(q_2^0, s_2^0)$.

It remains to consider the case when $s_2^0 = 0$. In this case, suppose Process 1 also repairs in state (q_1^0, s_1^0) . Both processes incur fixed costs K at time zero. Under the construction above, after time zero Processes 1 and 2 remain coupled in the server state and the result follows as in the case when $s_1 = s_2$. ■

To complete the proof of Theorem 3.2 we need to show that the DCOE and ACOI have solutions, that the average cost solution can be obtained via limits from the discounted cost value functions, and that the structure of an optimal policy is the same in both cases. To this end, we provide conditions under which the following set of conditions hold (see [26]).

- *SENSMI*: There exist $\Delta > 0$ and $\epsilon > 0$ such that for every state and action, there is a probability of

at least ϵ that the transition time will be greater than Δ .

- *SENSM2*: There exists \mathfrak{B} such that $\tau(i, a) \leq \mathfrak{B}$ for every i and a , where $\tau(i, a)$ is the mean transition time out of state i when action a is chosen.
- *SENSM3*: $v_\theta(i) < \infty$ for every state i and $\theta > 0$.
- *SENSM4*: There exists $\theta_0 > 0$ and non-negative numbers M_i such that $w_\theta(i) \leq M_i$ for every state i and $0 < \theta < \theta_0$, where $w_\theta(i) = v_\theta(i) - v_\theta(\mathbf{0})$, for a distinguished state $\mathbf{0}$. For every state i , there exists an action $a(i)$ such that $\sum_j P_{ij}(a(i))M_j < \infty$, where P_{ij} is the probability of transitioning to state j from state i when action $a(i)$ is chosen.
- *SENSM5*: There exists $\theta_0 > 0$ and a non-negative number N such that $-N \leq w_\theta(i)$ for every i and $0 < \theta < \theta_0$.

Proposition 3.4 *Under (3.12) the following hold.*

1. For $\theta > 0$, $v_\theta(x)$ satisfies the DCOE (3.3) - (3.6) and any stationary policy that minimizes the right-hand side of (3.5) is θ -discounted cost optimal.
2. There exists a constant g and a limit point $w(q, s) = \lim_{k \rightarrow \infty} [v_{\theta_k}(q, s) - v_{\theta_k}(0, B)]$, where $\theta_k \downarrow 0$, such that (g, w) satisfy the ACOI (3.7) - (3.10).
3. Any stationary policy that achieves the minimum in the ACOI (3.9) with (g, w) defined as above is average cost optimal. Moreover, the optimal average cost is $g(x) = g$, for every initial state $x \in \mathbb{X}$.
4. The optimal average cost may be computed by $g = \lim_{\theta \downarrow 0} \theta v_\theta(x)$ for any $x \in \mathbb{X}$.

Proof: Because $\Psi < \infty$ and repair times have a positive mean, SENS1 holds and it follows that there are only a finite number of decision epochs in a finite time interval, i.e., the processes are regular. SENS2 is satisfied since the transition rates m_0 and Ψ are uniformly bounded above. SENS3-4 require the stability results. Let the distinguished state, $\mathbf{0}$, be $(0, B)$. Under (3.12), Proposition 3.1 implies that there exists $\ell \in \{1, 2, \dots, B\}$ such that the average cost for policy ϕ_ℓ is finite with a single ergodic class $\mathbb{X}_\ell \supset (0, B)$. Note that \mathbb{X}_ℓ can be reached by any transient state in one (repair) transition. SENS3-4 then follow from Lemma 2 of [26]. SENS5 holds since, by Proposition 3.3, the distinguished state $(0, B)$ has the lowest cost, implying that $w_\theta(q, s) \geq 0$.

SENSM1 and SENSM3 are the assumptions of Theorem 1 of [26], which implies Result 1. SENSM1-5 and linear holding costs imply that Proposition 4 of [26] holds. In turn, Results 2 - 4 then follow from Theorem 2 of [26]; Result 2 is contained in the proof of the theorem, which references [25]. ■

The following is immediate from Proposition 3.3 and Result 2 of Proposition 3.4.

Corollary 3.5 *For $(q_1, s_1), (q_2, s_2) \in \mathbb{X}$, suppose $q_1 \leq q_2$ and $s_1 \geq s_2$. Then, under (3.12), $w(q_1, s_1) \leq w(q_2, s_2)$.*

We are now ready to prove the first result of Theorem 3.2.

Proof of Statement 1 of Theorem 3.2. Consider the average cost criterion. From Corollary 3.5, $w(q, s) \geq w(q, s + 1)$, $s \in \{1, 2, \dots, B\}$. From Proposition 3.4, the ACOI are satisfied and any policy that satisfies the ACOI is average cost optimal. Let $q \in \mathbb{Z}^+$ and $s \in \{1, 2, \dots, B - 1\}$. Assume that repair is optimal in state $(q, s + 1)$. Suppose that repair is not optimal in state (q, s) . Then by (3.9), $w(q, s) < f^R(q)$ and $w(q, s + 1) \geq f^R(q)$. This implies $w(q, s + 1) \geq f^R(q) > w(q, s)$, a contradiction to Corollary 3.5. Therefore, if repair is optimal in state $(q, s + 1)$, repair is also optimal in state (q, s) . The proof is similar for the discounted cost criterion. ■

It seems intuitive that if it is optimal to repair in state (q, s) , then it would also be optimal to repair in state $(q + 1, s)$. Such intuition is in line with the work of Federgruen and So [12] where when there are more customers in the system the decision-maker is more likely to start repair immediately as opposed to postponing. Similarly, in [13], when there are more customers in the system the decision-maker is more likely to choose faster, more expensive repairs; the optimal policies are monotone in the number of customers in the system. It was originally conjectured by the authors that this would continue to hold in the current models. However, surprisingly, this is not the case.

For all of the repair model numerical examples in this paper it is assumed that the repair times are exponentially distributed. All of the numerical examples are for the average cost criterion. For the calculations, the queue length is truncated. Since any feasible stationary policy is unichain, the results of [23], Chapter 8, apply under queue length truncation and the computations can be done using either policy iteration or value iteration. The next (counter-)example completes the proof of Theorem 3.2.

Proof of Statement 2 of Theorem 3.2.

Example 3.6 *Suppose that the parameter settings are: $K = 0$; $h = 1$; $\lambda = 1$; $B = 4$; $\mu_1 = 1/2$, $\mu_2 = 1$, $\mu_3 = 3/2$, $\mu_4 = 2$; $m_s = 1/5$, $s \in \{0, 1, \dots, 4\}$. Figure 1 displays the optimal repair policy.*




Figure 1: Example of an optimal non-monotone switching curve policy. The optimal repair policy for Example 3.6.

The optimal repair policy in Figure 1 is not monotone. This is evident by an initial decrease in the switching curve between $q = 0$ and $q = 1$. We note that there is not a tie; for example, it is not also optimal to choose action 0 in state $(0, 2)$. For $s = 2$, when there are zero customers in the system, there is actually more incentive to repair than when there are one or two customers in the system. As the number of customers increases, then so again does the incentive to repair. It should also be noted that in every example we considered that if an example does not have an optimal monotone switching curve, the non-monotonicity occurred at $q = 0$. In essence, when there are zero customers in the system, server capacity is not being used, so there is less penalty for downtime during repairs as compared to when customers are present.

4 The Replacement Model

In this section we consider the replacement model where the fixed costs are positive (and may vary as a function of the server state). Thus, replace K with $K(s)$ in (2.1), where $K(s)$ is a strictly positive function on $\{0, 1, \dots, B\}$. Since all non-repair times are exponentially distributed, all decision epochs occur according to a Poisson process with uniform rate $\Psi = 1$. The model then is an MDP, not a more general SMDP. Due to the memoryless property of the exponential distribution, it should be clear that do not replace is optimal when $s = B$; otherwise, under replace a fixed cost $K(B) > 0$ would be incurred for an instantaneous transition back to the same state.




Figure 2: Example of an optimal replacement policy. The optimal replacement policy for Example 4.1.

The DCOE for the replacement model (with server state dependent costs) are:

$$f_\theta(q, s) = \frac{1}{\theta + 1} [hq + \lambda v_\theta(q + 1, s) + \mu_s v_\theta((q - 1)^+, s) + m_s v_\theta(q, s - 1) + (1 - \lambda - \mu_s - m_s) v_\theta(q, s)], \quad s \in \{1, 2, \dots, B\}, \quad (4.1)$$

$$v_\theta(q, B) = f_\theta(q, B), \quad (4.2)$$

$$v_\theta(q, s) = \min \{f_\theta(q, s), K(s) + f_\theta(q, B)\}, \quad s \in \{1, 2, \dots, B - 1\}, \quad (4.3)$$

$$v_\theta(q, 0) = K(0) + f_\theta(q, B). \quad (4.4)$$

For the average cost criterion, the replacement model ACOI are:

$$f(q, s) = -g + hq + \lambda w(q + 1, s) + \mu_s w((q - 1)^+, s) + m_s w(q, s - 1) + (1 - \lambda - \mu_s - m_s) w(q, s), \quad s \in \{1, 2, \dots, B\}. \quad (4.5)$$

$$w(q, B) \geq f(q, B), \quad (4.6)$$

$$w(q, s) \geq \min \{f(q, s), K(s) + f(q, B)\}, \quad s \in \{1, 2, \dots, B - 1\}, \quad (4.7)$$

$$w(q, 0) \geq K(0) + f(q, B). \quad (4.8)$$

Consider the following illustrative example.

Example 4.1 For the replacement model, suppose the parameter settings are: $K(s) = K = 20/\psi$; $h = 1$; $\lambda = .4$; $B = 4$; $\mu_1 = 1/4$, $\mu_2 = 1/2$, $\mu_3 = 3/4$, $\mu_4 = 1$; $m_s = 1/2$, $s \in \{1, 2, \dots, 4\}$. Figure 2 displays the optimal replacement policy.

Several observations from Figure 2 are worth noting. First, the optimal policy is a monotone switching curve. Second, when $q = 0$ or $s = B$ do not replace is optimal. Third, as q gets large, it is optimal to replace in states where the service rate is less than μ_B ; the switching curve levels off at $B - 1$. With the exception of the monotonicity of the switching curve, each of these are stated in general in the following theorem, the main theorem of the section. Again the proof is divided into several parts.

Theorem 4.2 Suppose $K(s) = K$ for all s . Under the discounted cost criterion or under

$$\lambda < \max_{l \in \{1, 2, \dots, B\}} \left\{ \frac{\sum_{s=l}^B \mu_s / m_s}{\sum_{s=l}^B 1 / m_s} \right\} \quad (4.9)$$

and the average cost criterion, the following hold.

1. There exists an optimal policy with switching curve structure that is monotone in s .
2. Do not replace is optimal for $s > 0$ if either $q = 0$ or $\mu_s = \mu_B$.

Suppose now that $K(s)$ is allowed to vary with s . In the discounted case there exists θ_0 such that for all $0 < \theta < \theta_0$, or in the average cost case under (4.9), the following holds.

3. There exists a (finite) $q_0 \in \mathbb{Z}^+$ such that for $q \geq q_0$ replace is optimal in (q, s) if $\mu_s < \mu_B$ and do not replace is optimal if $\mu_s = \mu_B$; that is, for large q , it is optimal to replace as soon as the service rate deteriorates to a rate strictly less than μ_B , but no sooner.

Statement 2 of Theorem 4.2 implies that for $q = 0$ there is little incentive to replace. Since replacements are instantaneous, it seems reasonable to wait until customers are in the system. Given the observation from Theorem 3.2 that optimal switching curves in repair models may not be monotone at $q = 0$, the question is raised: are the optimal switching curves for replacement problems monotone? Statement 3 of Theorem 4.2 provides a partial answer to this question as q approaches infinity. Furthermore, all of the examples we ran exhibited an optimal monotone switching curve. To date we do not have a proof that this holds in general.

We first show that the optimality equations are satisfied. We use the optimality equations to prove that a switching curve policy that is monotone in s is optimal and that do not replace is optimal when the queue is either empty or serving at the highest rate. We then prove that the switching curve is monotone in q when q becomes large. This will complete the proof of Theorem 4.2. To end the section, monotonicity in s is extended under a set of sufficient conditions to cases where replacement costs $K(s)$ vary with s . Numerical examples illustrate that the optimal policies may not be monotone in s when the sufficient conditions are not met. The first result is analogous to Proposition 3.3 for repair models.

Proposition 4.3 For $(q_1, s_1), (q_2, s_2) \in \mathbb{X}$, suppose $q_1 \leq q_2$ and $s_1 \geq s_2$. Then,

$$v_\theta(q_1, s_1) \leq v_\theta(q_2, s_1), \quad (4.10)$$

$$v_\theta(q_1, s_1) \leq v_\theta(q_2, s_2) + \max_{s \in \{0, 1, \dots, B-1\}} K(s) - \min_{s \in \{0, 1, \dots, B-1\}} K(s). \quad (4.11)$$

Proof: The proof is via precisely the same sample-path argument as that for Proposition 3.3; let all repair times be zero. Expression (4.10) follows directly since Processes 1 and 2 are always coupled in the server state. For expression (4.11), note that once Processes 1 and 2 couple in the server state, they remain coupled. The two processes couple in the server state when they both replace, simultaneously, for the first time. The holding costs are always lower for Process 1; however, the first replacement cost may be higher for Process 1 than Process 2 if $K(s)$ is not non-increasing in s . (Assume that Process 1 does not actually replace and pay a positive fixed cost if $s = B$.) The difference between the cost of the first replacement for Process 1 and that of Process 2 is bounded above by the difference $\max_{s \in \{0, 1, \dots, \min\{s_1, B-1\}\}} K(s) - \min_{s \in \{0, 1, \dots, \min\{s_2, B-1\}\}} K(s)$. ■

Proposition 4.4 *For the replacement model, the following hold.*

1. For $\theta > 0$, $v_\theta(x)$ satisfies the DCOE (4.1) - (4.4) and any stationary policy that minimizes the right-hand side of (4.3) is θ -discounted cost optimal.

Moreover, under (4.9)

2. There exists a constant g and a limit point $w(q, s) = \lim_{k \rightarrow \infty} [v_{\theta_k}(q, s) - v_{\theta_k}(0, B)]$, where $\theta_k \downarrow 0$, such that (g, w) satisfy the ACOI (4.5) - (4.8).
3. Any stationary policy π^* that achieves the minimum in the ACOI (4.7) with (g, w) defined as above is average cost optimal with average cost g . Moreover, the average cost optimality equalities (ACOE) (the ACOI with “=” replacing “ \geq ”) hold at any state that is positive recurrent under the Markov chain induced by π^* .
4. The optimal average cost may be computed by $g = \lim_{\theta \downarrow 0} \theta v_\theta(x)$ for any $x \in \mathbb{X}$.

Proof: For the discounted cost case, since the model is an MDP, Result 1 follows from Propositions 1.5 and 1.7 of Chapter 3 of [2] and does not require a stability condition. For the average cost case, [27] requires that SENI-3 hold. The conditions are analogous to the semi-Markov conditions SENSM3-5, respectively.

Consider the policy ϕ_ℓ . The fluid limit analysis in Section 7.1 in the Appendix also applies to the replacement model under ϕ_ℓ by removing the states $\{0, \dots, \ell - 1\}$. The same analysis implies stability for the analogous replacement process $\{X(t), t \geq 0\}$ on state space $\mathcal{X} = \mathbb{Z}^+ \times \{\ell, \dots, B\}$. This in turn implies that, under (4.9), ϕ_ℓ induces an irreducible ergodic Markov chain on $\mathbb{X}_\ell = \mathbb{Z}^+ \times \{\ell - 1, \ell, \dots, B\}$.

Furthermore, the long-run average queue length of such a process is finite. Hence, the long-run average holding costs are finite. Since the fixed costs are bounded, the long-run average replacement costs are also finite, and $g^{\phi_\ell}(x) < \infty$ for all $x \in \mathbb{X}_\ell$. It follows from [27] C.2.2(iv) and C.1.4(v) that the expected time and cost to reach the distinguished state $(0, B)$ for the first time starting in any positive recurrent state \mathbb{X}_ℓ is finite. In turn, since replacement costs are finite, the expected time and cost to reach $(0, B)$ starting from a transient state, with $s \in \{1, 2, \dots, \ell - 2\}$, is also finite. Therefore, under (4.9), ϕ_ℓ satisfies the assumption of Proposition 7.5.3 of [27], and SEN1-2 hold.

By Proposition 4.3, $-\max_{i \in \{0, 1, \dots, B-1\}} K(i) \leq v_\theta(q, s) - v_\theta(0, B)$, for all $(q, s) \in \mathbb{X}$. Since $\max_{i \in \{0, 1, \dots, B-1\}} K(i) > 0$, SEN3 holds for the distinguished state $(0, B)$. Applying Theorems 7.2.3 and 7.4.3 of [27] then yields the desired Results 2-4. ■

The following corollary is immediate from Proposition 4.3 and Result 2 of Proposition 4.4.

Corollary 4.5 *For $(q_1, s_1), (q_2, s_2) \in \mathbb{X}$ suppose $q_1 \leq q_2$ and $s_1 \geq s_2$. Then, under (4.9),*

$$w(q_1, s_1) \leq w(q_2, s_2) + \max_{s \in \{0, 1, \dots, B-1\}} K(s) - \min_{s \in \{0, 1, \dots, B-1\}} K(s). \quad (4.12)$$

This leads to the proof of the first statement of Theorem 4.2.

Proof of Statement 1 of Theorem 4.2. Consider the average cost criterion. Since $K(s) = K$, Corollary 4.5 implies $w(q, s) \geq w(q, s + 1)$, $s \in \{1, 2, \dots, B - 1\}$. From Proposition 4.4, the ACOI are satisfied and any policy that satisfies the ACOI is average cost optimal. From (4.7),

$$w(q, s) \geq \min \{f(q, s), K + f(q, B)\}, \quad s \in \{1, 2, \dots, B - 1\}.$$

Let $q \in \mathbb{Z}^+$ and $s \in \{1, 2, \dots, B - 2\}$. Assume that replace is optimal in state $(q, s + 1)$. Suppose that replace is not optimal in state (q, s) . Then $w(q, s) < K + f(q, B)$ and $w(q, s + 1) \geq K + f(q, B)$. This implies $w(q, s + 1) \geq K + f(q, B) > w(q, s)$; a contradiction. Therefore, if replace is optimal in state $(q, s + 1)$, replace is also optimal in state (q, s) . The proof is similar for the discounted cost criterion. ■

Proof of Statement 2 of Theorem 4.2. Suppose $s \in \{1, 2, \dots, B - 1\}$ and $q = 0$. From (4.1) we have

$$\begin{aligned} f_\theta(0, s) &= (\theta + 1)^{-1} [\lambda v_\theta(1, s) + \mu_s v_\theta(0, s) + m_s v_\theta(0, s - 1) + (1 - \lambda - \mu_s - m_s) v_\theta(0, s)] \\ &= (\theta + 1)^{-1} [\lambda v_\theta(1, s) + m_s v_\theta(0, s - 1) + (1 - \lambda - m_s) v_\theta(0, s)]. \end{aligned}$$

From (4.3) and (4.4) we get

$$\begin{aligned}
f_\theta(0, s) &\leq (\theta + 1)^{-1} [\lambda[K + f_\theta(1, B)] + m_s[K + f_\theta(0, B)] + (1 - \lambda - m_s)[K + f_\theta(0, B)]] \\
&= (\theta + 1)^{-1} [K + \lambda f_\theta(1, B) + \mu_B f_\theta(0, B) + (1 - \lambda - \mu_B) f_\theta(0, B)] \\
&= (\theta + 1)^{-1} [K + \lambda f_\theta(1, B) + \mu_B f_\theta(0, B) + m_B f_\theta(0, B - 1) \\
&\quad + (1 - \lambda - \mu_B - m_B) f_\theta(0, B) + m_B (f_\theta(0, B) - f_\theta(0, B - 1))] \\
&= (\theta + 1)^{-1} K + f_\theta(0, B) + (\theta + 1)^{-1} m_B (f_\theta(0, B) - f_\theta(0, B - 1)) \\
&\leq K + f_\theta(0, B),
\end{aligned}$$

where the last inequality follows since Proposition 4.3 implies that $f_\theta(0, B) - f_\theta(0, B - 1) \leq 0$. From (4.3) and Result 1 of Proposition 4.4, it follows that do not replace is optimal in state $(0, s)$.

Consider now $q \in \{1, 2, \dots\}$ and $s \in \{1, 2, \dots, B - 1\}$ such that $\mu_s = \mu_B$ (as discussed, the result holds for $s = B$). From the DCOE we have

$$\begin{aligned}
f_\theta(q, s) &= (\theta + 1)^{-1} [hq + \lambda v_\theta(q + 1, s) + \mu_B v_\theta(q - 1, s) + m_s v_\theta(q, s - 1) \\
&\quad + (1 - \lambda - \mu_B - m_s) v_\theta(q, s)] \\
&\leq (\theta + 1)^{-1} [hq + \lambda [f_\theta(q + 1, B) + K] + \mu_B [f_\theta(q - 1, B) + K] + m_s [f_\theta(q, B) + K] \\
&\quad + (1 - \lambda - \mu_B - m_s) [f_\theta(q, B) + K]] \\
&= (\theta + 1)^{-1} [hq + K + \lambda f_\theta(q + 1, B) + \mu_B f_\theta(q - 1, B) + m_B f_\theta(q, B - 1) \\
&\quad + (1 - \lambda - \mu_B - m_B) f_\theta(q, B) + m_B (f_\theta(q, B) - f_\theta(q, B - 1))] \\
&= (\theta + 1)^{-1} K + f_\theta(q, B) + (\theta + 1)^{-1} m_B (f_\theta(q, B) - f_\theta(q, B - 1)) \\
&\leq K + f_\theta(q, B),
\end{aligned}$$

where the last inequality follows since Proposition 4.3 implies that $f_\theta(q, B) - f_\theta(q, B - 1) \leq 0$. From (4.3) and Result 1 of Proposition 4.4, it follows that do not replace is optimal in state $(0, s)$. The proof for the average cost criterion is similar. ■

Lemma 4.6 For $(q, s) \in \mathbb{X}$,

$$v_\theta(q + 1, s) - v_\theta(q, s) \geq h \sum_{n=1}^q \left(\frac{1}{\theta + 1} \right)^n,$$

$$w(q + 1, s) - w(q, s) \geq hq.$$

Proof: Consider the same sample-path argument as in the proof of Proposition 4.3. Suppose that Process 1 starts in state (q, s) and Process 2 starts in state $(q + 1, s)$. Recall that Process 1 mimics Process 2 by replacing at the same times so that both processes have the same replacement costs and both processes see the same (potential) departures. The only differences in costs are due to the holding costs. The queue length under Process 1 is never larger than the queue length under Process 2. Furthermore, note that the difference in queue lengths between the two processes is exactly one before the first time that the system under Process 1 empties. Hence, the difference in queue lengths is exactly one for at least q decision epochs. The result follows. ■

Proof of Statement 3 of Theorem 4.2. Consider the discounted cost criterion. The case $\mu_s = \mu_B$ is proved by Statement 2 of Theorem 4.2. Consider a state (q, s) where $q > 0$, $s \in \{1, 2, \dots, B - 1\}$, and $\mu_s < \mu_B$. Suppose that do not replace is optimal in state (q, s) . Then, by (4.2) and (4.3)

$$K(s) + v_\theta(q, B) - v_\theta(q, s) \geq 0, \quad (4.13)$$

and by (4.1)

$$\begin{aligned} v_\theta(q, s) &= (\theta + 1)^{-1}(hq + \lambda v_\theta(q + 1, s) + \mu_s v_\theta(q - 1, s) + (\mu_B - \mu_s)v_\theta(q, s) \\ &\quad + m_s v_\theta(q, s - 1) + m_B v_\theta(q, s) + (1 - \lambda - \mu_B - m_s - m_B)v_\theta(q, s)). \end{aligned}$$

Then,

$$\begin{aligned} v_\theta(q, B) - v_\theta(q, s) &= (\theta + 1)^{-1}(\lambda(v_\theta(q + 1, B) - v_\theta(q + 1, s)) + \mu_s(v_\theta(q - 1, B) - v_\theta(q - 1, s)) \\ &\quad + (\mu_B - \mu_s)(v_\theta(q - 1, B) - v_\theta(q, s)) + m_s(v_\theta(q, B) - v_\theta(q, s - 1)) \\ &\quad + m_B(v_\theta(q, B - 1) - v_\theta(q, s)) \\ &\quad + (1 - \lambda - \mu_B - m_s - m_B)(v_\theta(q, B) - v_\theta(q, s))) \\ &\leq (\theta + 1)^{-1} \left((\mu_B - \mu_s)(v_\theta(q - 1, s) - v_\theta(q, s)) + \max_{s \in \{0, 1, \dots, B-1\}} K(s) \right), \end{aligned}$$

by Proposition 4.3. It follows from Lemma 4.6 that

$$K(s) + v_\theta(q, B) - v_\theta(q, s) \leq K(s) - h(\mu_B - \mu_s) \sum_{n=1}^{q-1} \left(\frac{1}{\theta + 1} \right)^{n+1} + \max_{s \in \{0, 1, \dots, B-1\}} K(s)/(\theta + 1).$$

Thus,

$$\lim_{q \rightarrow \infty} (K(s) + v_\theta(q, B) - v_\theta(q, s)) \leq K(s) - h(\mu_B - \mu_s)/\theta + \max_{s \in \{0, 1, \dots, B-1\}} K(s)/(\theta + 1),$$




Figure 3: Example of an optimal replacement policy that is non-monotone in the server state s . The optimal replacement policy for Example 4.7.

which, since $\mu_s < \mu_B$, goes to $-\infty$ as $\theta \downarrow 0$. Therefore, when q is large we have a contradiction to (4.13) for θ close to 0, and replace is optimal in state (q, s) . The proof for the average cost criterion is similar. ■

We remark that this result does not hold for the repair model. Indeed, in Example 3.6 (see Figure 1) the switching curve levels off as q becomes large, but do not replace is optimal in server state 3 while $\mu_3 < \mu_B$.

4.1 Switching Curves when $K(s)$ is a Function of s

The main results for monotonicity in s so far, Statement 1 for Theorems 3.2 and 4.2, depend on the replacement cost being constant for all s . In this section we consider the model with $K(s)$ variable in s . We begin with the observation that when $K(s)$ varies with s , there may not exist an optimal policy that is monotone in s .

Example 4.7 For the replacement model, suppose the parameter settings are the same as in Example 4.1 except the replacement costs are $K(0) = K(1) = K(2) = 60/\psi$; $K(3) = K(4) = 20/\psi$. Figure 3 displays the optimal replacement policy.

The optimal switching curve is not monotone in s . For $q = 3$ there is an interesting pattern. Replace is optimal when $s = 3$ where it is cheaper than $s = 2$ or $s = 1$. At $s = 2$ do not replace is optimal, but at $s = 1$ replace is optimal because the service rate is too low. We note that there are no ties for the optimal actions.

To study monotone optimal policies, we find it useful to consider the value functions inductively. For all $(q, s) \in \mathbb{X}$, let $v_{0,\theta}(q, s) = K(s)$. For $n \in \{1, 2, \dots\}$, define the discrete-time finite horizon optimality

equations (FHOE):

$$\begin{aligned}
f_{n,\theta}(q, s) &= (\theta + 1)^{-1} [\hbar q + \lambda v_{n-1,\theta}(q + 1, s) + \mu_s v_{n-1,\theta}((q - 1)^+, s) \\
&\quad + m_s v_{n-1,\theta}(q, s - 1) + (1 - \lambda - \mu_s - m_s) v_{n-1,\theta}(q, s)], \quad s \in \{1, 2, \dots, B\}, \\
v_{n,\theta}(q, B) &= f_{n,\theta}(q, B), \\
v_{n,\theta}(q, s) &= \min \{f_{n,\theta}(q, s), K(s) + f_{n,\theta}(q, B)\}, \quad s \in \{1, 2, \dots, B - 1\}, \\
v_{n,\theta}(q, 0) &= K(0) + f_{n,\theta}(q, B).
\end{aligned}$$

The proofs of the following two lemmas are given in Section 7.2 in the Appendix.

Lemma 4.8 For all $q \in \mathbb{Z}^+$,

$$v_{n,\theta}(q, s) \leq v_{n,\theta}(q + 1, s), \quad s \in \{0, 1, \dots, B\}, \quad n \geq 0, \quad (4.14)$$

$$f_{n,\theta}(q, s) \leq f_{n,\theta}(q + 1, s), \quad s \in \{1, 2, \dots, B\}, \quad n \geq 1. \quad (4.15)$$

Lemma 4.9 Assume $K(s)$ is convex on $\{0, 1, \dots, B - 1\}$; that is, $K(s - 1) - K(s - 2) \leq K(s) - K(s - 1)$ for $s \in \{2, 3, \dots, B - 1\}$. Suppose one of the following conditions holds:

1. $K(s) \leq K(s - 1)$ and $m_{s-1} \geq m_s + \theta$, $s \in \{2, \dots, B - 1\}$;
2. $K(s) \geq K(s - 1)$ and $m_{s-1} \leq m_s + \theta$, $s \in \{2, \dots, B - 1\}$.

Then, the following hold for all $q \in \mathbb{Z}^+$.

$$v_{n,\theta}(q, s) - v_{n,\theta}(q, s - 1) \leq K(s) - K(s - 1), \quad s \in \{1, 2, \dots, B - 1\}, \quad n \geq 0, \quad (4.16)$$

$$f_{n,\theta}(q, s) - f_{n,\theta}(q, s - 1) \leq K(s) - K(s - 1), \quad s \in \{2, 3, \dots, B - 1\}, \quad n \geq 1. \quad (4.17)$$

Proposition 4.10 Assume $K(s)$ is convex on $\{0, 1, \dots, B - 1\}$. Suppose

1. for the discounted cost criterion, one of the conditions of Lemma 4.9 holds, or
2. for the average cost criterion, (4.9) holds and any of the following conditions hold:

$$(a) \quad m_{s-1} = m_s, \quad s \in \{2, 3, \dots, B - 1\};$$

$$(b) \quad K(s) \leq K(s - 1) \text{ and } m_{s-1} \geq m_s, \quad s \in \{2, 3, \dots, B - 1\};$$

$$(c) \quad K(s) \geq K(s - 1) \text{ and } m_{s-1} \leq m_s, \quad s \in \{2, 3, \dots, B - 1\}.$$

Then, there exists an optimal policy with switching curve structure that is monotone in s .

Proof: First consider the discounted cost criterion. We would like to show that if do not replace is optimal in state $(q, s - 1)$ then do not replace is also optimal in state (q, s) ; or, by (4.3),

$$f_\theta(q, s - 1) \leq K(s - 1) + f_\theta(q, B) \Rightarrow f_\theta(q, s) \leq K(s) + f_\theta(q, B), \quad s \in \{2, 3, \dots, B - 1\}. \quad (4.18)$$

Note that we only need to consider $s < B$ since do not replace is optimal when $s = B$. A sufficient condition guaranteeing that (4.18) holds is

$$f_\theta(q, s) - f_\theta(q, s - 1) \leq K(s) - K(s - 1), \quad s \in \{2, 3, \dots, B - 1\}. \quad (4.19)$$

Since $v_{0,\theta}(q, s) = K(s)$ is a bounded function, Propositions 1.5 and 1.7 of Chapter 3 of [2] imply that $\lim_{n \rightarrow \infty} v_{n,\theta}(q, s) = v_\theta(q, s)$ and $\lim_{n \rightarrow \infty} f_{n,\theta}(q, s) = f_\theta(q, s)$. Hence, by Lemma 4.9, (4.19) holds under either of the hypotheses of Lemma 4.9. Therefore, the result holds for the discounted cost criterion.

For the average cost criterion, it is sufficient to show that

$$w(q, s) - w(q, s - 1) \leq K(s) - K(s - 1), \quad s \in \{2, 3, \dots, B - 1\}. \quad (4.20)$$

To see that this is sufficient, assume that replace is optimal in state (q, s) , $s \in \{2, 3, \dots, B - 1\}$, and suppose that replace is not optimal in state $(q, s - 1)$. By Proposition 4.4, the ACOI are satisfied. From (4.7), $w(q, s) \geq K(s) + f(q, B)$ and $w(q, s - 1) < K(s - 1) + f(q, B)$. This implies $w(q, s) - w(q, s - 1) > K(s) - K(s - 1)$, which would contradict (4.20). Therefore, (4.20) is sufficient to guarantee that if replace is optimal in state (q, s) then replace is also optimal in state $(q, s - 1)$.

We now prove that (4.20) holds under either of the Conditions (a), (b), or (c). The proof is a sample-path argument that requires only a slight modification to that for Proposition 3.3. Consider two processes. Process 1 starts in state (q, s) , $s \in \{2, 3, \dots, B - 1\}$. Process 2 starts in $(q, s - 1)$ and uses the optimal policy. Process 1 follows a potentially suboptimal policy that mimics Process 2 replacements. In order to show that (4.20) holds, it is sufficient to show that the difference of total (non-discounted) costs for Process 1 minus the total costs for Process 2 is bounded above by $K(s) - K(s - 1)$.

The only difference in the construction here as compared to Proposition 3.3 is the way that server deteriorations are constructed. If the server state for Process i is s_i , then a potential deterioration occurs for both processes at rate $\max(m_{s_1}, m_{s_2})$. For Process i , given a potential deterioration there is an actual deterioration with probability $m_{s_i} / \max(m_{s_1}, m_{s_2})$ and a dummy transition with probability $(\max(m_{s_1}, m_{s_2}) - m_{s_i}) / \max(m_{s_1}, m_{s_2})$. Under this construction, $s_1 \geq s_2$ until Process 2 replaces for the first time, after which the two processes couple in the server state and $s_1 = s_2$. Again both processes

see the same arrivals and there are never more customers in the system under Process 1 than under Process 2; Process 1 has lower holding costs. The only difference in replacement costs is due to the first replacement. Since deteriorations are exponentially distributed, the processes will eventually replace, with probability one. Suppose that the server state under Process i just before the first replacement is s'_i ; $s'_1 \leq s$ and $s'_2 \leq s - 1$. To show (4.20), it is sufficient to show that


$$K(s'_1) - K(s'_2) \leq K(s) - K(s - 1), \quad s \in \{2, 3, \dots, B - 1\}. \quad (4.21)$$

Under Condition (a): the two processes deteriorate in their server states at precisely the same times, and $s'_1 = s'_2 + 1$. Hence, $K(s'_1) - K(s'_2) = K(s'_2 + 1) - K(s'_2)$, and (4.21) follows from convexity. Under Condition (b): since $m_{s-1} \geq m_s$, Process 1 does not deteriorate unless Process 2 also deteriorates, and $s'_1 \geq s'_2 + 1$. Either $s'_1 = s'_2 + 1$ and $K(s'_1) - K(s'_2) = K(s'_2 + 1) - K(s'_2)$, or $s'_1 > s'_2 + 1 \geq 1$ and, since $K(s) \leq K(s - 1)$ for $s \in \{2, 3, \dots, B - 1\}$, $K(s'_1) - K(s'_2) \leq K(s'_2 + 1) - K(s'_2)$. In either case, (4.21) then follows from convexity. Under Condition (c): since $m_{s-1} \leq m_s$, Process 2 does not deteriorate unless Process 1 also deteriorates. So, $0 \leq s'_1 - s'_2 \leq 1$. In the case $s'_1 = s'_2 + 1$, (4.21) follows from convexity. Otherwise, $s'_1 = s'_2$, $K(s'_1) - K(s'_2) = 0$, and (4.21) follows since $K(s) - K(s - 1) \geq 0$ for $s \in \{2, 3, \dots, B - 1\}$. ■


5 Additional Numerical Examples and Conclusions

This section includes additional examples under the average cost criterion and introduces some simple heuristics. The previous results stand to reduce computation of optimal policies, yet the non-monotonicity (in either q or s) might leave them somewhat difficult to implement. Consider again Example 3.6. The optimal repair policy is displayed again in Figure 4(a). Recall that a server-state threshold policy is characterized by a threshold ℓ ; the chosen action is repair (replace) if and only if $s < \ell$ (regardless of the number of customers in the system). The optimal average cost for Example 3.6 is 14.7024. The best server-state threshold policy has $\ell = 3$; see Figure 4(b). The threshold policy has an average cost of 15.0895, only 2.63% more than the optimal cost.


One might conjecture from this example that in fact one could use the optimal server-state threshold policy in place of the optimal policy without much loss in terms of cost. Moreover, if we are willing to do the work of computing it, a 2-level heuristic policy (defined formally below) may further improve the heuristic in terms of cost. Two questions that we would like to consider are:



(a) The optimal repair policy.



(b) The best server-state threshold policy, within 2.63% of optimal.



(c) The best 2-level policy, within 1.13% of optimal.

Figure 4: Optimal repair policy and heuristics for Example 3.6.

- Can the non-monotonicity in q discussed in Example 3.6 be ignored?
- In the cases that it can be ignored, what kind of “smoothing” should be done?

If we smooth any “kinks” in the optimal policy depicted in Figure 4(a), we arrive at the server-state threshold policy in Figure 4(b). If we only smooth out all but one kink in Figure 4(a), we have the 2-level heuristic policy depicted in Figure 4(c) defined by two levels ℓ_1 and ℓ_2 and a queue threshold T : if $q < T$, repair (replace) if and only if $s < \ell_1$, otherwise, if $q \geq T$, repair (replace) if and only if $s < \ell_2$. Note that it is not assumed that $\ell_1 \leq \ell_2$. As one may expect, the best 2-level heuristic for Example 3.6 has $\ell_1 = 2$, $\ell_2 = 3$, and $T = 11$. The average cost is 14.8688, 1.13% more than the optimal cost.

If the observations from this example consistently hold true, the decision-maker would be able to completely ignore the non-monotonicity in q with minimal effect on the cost. The next example shows that this in fact is not always the case. Perhaps more importantly it shows that restricting attention to 2-level heuristics based on visual inspection can cause significant increases in cost.




Figure 5: Optimal repair policy for Example 5.1.

Example 5.1 For the repair model, suppose the parameter settings are: $K = 0$; $h = 1$; $\lambda = .3$; $B = 4$; $\mu_1 = 1/2$, $\mu_2 = 1/2$, $\mu_3 = 3/4$, $\mu_4 = 1$; $m_0 = 1/5$, $m_s = 1/10$, $s \in \{1, 2, \dots, 4\}$. Figure 5 displays the optimal repair policy.

The optimal average cost is 1.1612. The best server-state threshold policy has $\ell = 3$ and an average cost of 1.2200, 5.07% more than the optimal cost. As for 2-level heuristics, suppose we fix $\ell_1 = 1$ and $\ell_2 = 3$. It turns out that the best 2-level heuristic with these two levels has $T = 5$ and average cost of 1.3245, 14.06% above the optimal cost. This is more than the best server-state threshold policy. In fact, the 2-level heuristic that is best overall has $\ell_1 = \ell_2 = 3$, i.e., it is the best threshold policy with average cost 5.07% more than optimal.

Recall the replacement model of Example 4.1, Figure 2. The optimal policy for this example has an average cost of 1.6290. The best server-state threshold policy has $\ell = 3$ and an average cost of 1.8724, 15.01% more than optimal. The best 2-level heuristic has $\ell_1 = 1$, $\ell_2 = 3$, and $T = 2$ with an average cost of 1.6581, only 1.79% above optimal. The server-state threshold policy does not perform well. The 2-level heuristic performs much better.

One might consider the 2-level heuristics to be viable alternatives to the optimal policies. However, there are some difficulties. Example 5.1 shows that choosing the 2-level heuristic by visual inspection can have dire consequences. This coupled with the fact that computing the optimal 2-level heuristic may be more difficult than computing the optimal policy make it less than desirable.

In summary, there are three observations resulting from these numerical examples. First, it is important to keep track of the workload information. Without it, important state-action pairs may be ignored. Second, when optimal switching curves are non-monotone and decrease in q , such as in Examples 3.6 and 5.1, the effects of the decreases on the optimal average cost may be significant depending on how closely the heuristics can smooth the optimal policies. Finally, in some situations 2-level heuristics perform significantly better than server-state threshold policies.

In conclusion, we have considered two models for machine maintenance: repair and replacement. For both models we have shown in some generality that the optimal maintenance policies have switching curve structure. Indeed, for the repair model monotonicity in the server state is proven for generally distributed repair times. For the replacement model monotonicity in s holds for constant replacement costs and continues to hold, under some sufficient conditions, when replacement costs vary with s .

While the optimal policies are monotone in s , optimal repair decisions are not, in general, monotone in q . As congestion in the system grows so may the incentive to perform maintenance. On the other hand, there may be more incentive to repair when the system is empty compared to when there are a few customers present. In between is where the policy remains unclear although our numerics seem to confirm the conjecture that the policy is monotone when $q > 0$. For the replacement model we can only guarantee monotonicity in q for “large” q . Of course there is no way *a priori* to know when q is large enough.

The results developed in this paper allow a decision-maker to store only the switching curve, and not the space. As networks of single-server queues are developed, this stands to simplify the lookup table for optimal policies and hopefully makes their implementation easier. Finally, we would like to point out that although the server state is assumed to be discrete, we believe that similar results could be obtained from a continuous server state model.

6 Acknowledgments

We would like to thank the anonymous referees for helpful comments on the paper. This work is partially supported by the National Science Foundation, NSF-DMI-0132811.

References

- [1] S. Andradottir, H. Ayhan, and D. Down. Compensating for failures with flexible servers. *Operations Research*, to appear. 3
- [2] D. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, Belmont, MA, first edition, 1995. 15, 21
- [3] E. Boukas and Z. Liu. Production and maintenance control for manufacturing systems. *IEEE Transactions on Automatic Control*, 46:1455–1460, 2001. 2
- [4] H. Chen and D. Yao. *Fundamentals of Queueing Networks*. Springer-Verlag, New York, 2001. 31
- [5] J. Dai. On positive harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability*, 5:49–77, 1995. 28, 30, 31, 32

- [6] J. Dai and S. Meyn. Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Transactions on Automatic Control*, 40:1889–1904, 1995. 28, 30, 32
- [7] T. Das and S. Sarkar. Optimal preventive maintenance in a production inventory system. *IIE Transactions on Quality and Reliability Engineering*, 31:537–551, 1999. 2
- [8] M. Davis. Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *J. R. Statist. Soc. B*, 46(5):353–388, 1984. 28
- [9] R. Dekker. Applications of maintenance optimization models: a review and analysis. *Reliability Engineering and System Safety*, 51(3):229–240, March 1996. 2
- [10] C. Derman. On sequential decisions with Markov chains. *Management Science*, 9:16–24, 1962. 1
- [11] C. Derman. *Finite State Markovian Decision Processes*. Academic Press, New York, 1970. 1
- [12] A. Federgruen and K. So. Optimal time to repair a broken server. *Advances in Applied Probability*, 21:376–397, 1989. 2, 11
- [13] A. Federgruen and K. So. Optimal maintenance policies for single-server queueing-systems subject to breakdowns. *Operations Research*, 38(2):330–343, March-April 1990. 2, 11
- [14] A. Federgruen and K. So. Optimality of threshold policies in single-server queueing-systems with server vacations. *Advances in Applied Probability*, 23(2):388–405, June 1991. 2
- [15] S. Iravani and I. Duenyas. Integrated maintenance and production control of a deteriorating production system. *IIE Transactions*, 34:423–435, 2002. 2
- [16] D. Kaufman. *Dynamic Control of Production Systems with Varying Service Capacity*. PhD thesis, University of Michigan, Ann Arbor, Michigan, 2005. 32
- [17] J. Koyanagi and H. Kawai. An optimal maintenance policy for a deteriorating server of an M/G/1 queueing system. In *Stochastic Modeling in Innovative Manufacturing*, pages 215–224. Springer-Verlag, Berlin, 1997. 3
- [18] H. Lee and M. Rosenblatt. Simultaneous determination of production cycle and inspection schedules in a production system. *Management Science*, 33:1125–1136, 1987. 2
- [19] S. Lippman. Applying a new device in the optimization of exponential queueing system. *Operations Research*, 23(4):687–710, 1975. 4
- [20] S. Meyn and D. Down. Stability of generalized Jackson networks. *The Annals of Applied Probability*, 4(1):124–148, 1994. 32
- [21] K. Miriyala and N. Viswanadham. Reliability analysis of fms. *International Journal of Flexible Manufacturing Systems*, 2:145–162, 1989. 3
- [22] W. Pierskalla and J. Voelker. A survey of maintenance models: the control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly*, 23:353–388, 1976. 2
- [23] M. Puterman. *Markov decision processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, Inc., New York, 1994. 11
- [24] S. Ross. *Applied Probability Models with Optimization Applications*. Holden-Day Series Management Science. Holden-Day, Inc., San Francisco, 1970. 1, 2

- [25] L. Sennott. Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Operations Research*, 37(4):626–633, 1989. 11
- [26] L. Sennott. Average cost semi-Markov decision processes and the control of queueing systems. *Probability in the Engineering and Informational Sciences*, 3:247–272, 1989. 9, 10, 11
- [27] L. Sennott. *Stochastic Dynamic Programming and the Control of Queueing System*. Wiley Series in Probability and Statistics. John Wiley and Sons, Inc., New York, 1999. 15, 16
- [28] S. Seshadri and M. Pinedo. Optimal allocation of resources in a job shop environment. *IIE Transactions*, 1999. 3
- [29] Y. Sherif and M. Smith. Optimal maintenance models for systems subject to failure – a review. *Naval Research Logistics Quarterly*, 28:47–54, 1981. 2
- [30] T. Sloan. A periodic review production and maintenance model with random demand, deteriorating equipment, and binomial yield. *Journal of the Operational Research Society*, 55:647–656, 2004. 2
- [31] T. Sloan and J. Shanthikumar. Combined production and maintenance scheduling for a multiple-product, single machine production system. *Production and Operations Management*, 9:379–399, 2000. 2
- [32] T. Sloan and J. Shanthikumar. Using in-line equipment condition and yield information for maintenance scheduling and dispatching in semiconductor wafer fabs. *IIE Transaction*, 34:191–209, 2002. 2
- [33] K. So. Optimality of control limit policies in replacement models. *Naval Research Logistics*, 39:685–697, 1992. 2
- [34] M. Srinivasan and H. Lee. Production/inventory system with preventative maintenance. *IIE Transactions*, 28:879–890, 1996. 2
- [35] F. Van der Duyn Schouten and S. Vanneste. Maintenance optimization of a production system with buffer capacity. *European Journal of Operational Research*, 82:323–338, 1995. 2
- [36] G. van Dijkhuizen and A. van Harten. Two-stage generalized age maintenance of a queue-like production system. *European Journal of Operational Research*, 108:363–378, 1998. 1, 2
- [37] B. Vinod and T. John. On optimal capacities for repair facilities in flexible manufacturing systems. In *Flexible Manufacturing Systems: Methods and Studies*, pages 341–349. Elsevier, Amsterdam, 1986. 3
- [38] H. Wang. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3):469–89, June 2002. 2
- [39] M. Widmer and P. Solot. Do not forget the breakdowns and the maintenance operations in fms design problems. *International Journal of Production Research*, 28:421–430, 1990. 3
- [40] C. Wu, D. Down, and M. Lewis. Heuristics for allocation of reconfigurable resources in a serial line with reliability considerations. *preprint*, 2007. 3
- [41] C. Wu, M. Lewis, and M. Veatch. Dynamic allocation of reconfigurable resources in a two-stage tandem queueing system with reliability considerations. *IEEE Transactions on Automatic Control*, 51(2):309–314, 2005. 3

- [42] X. Yao. Optimal preventative maintenance policies for unreliable queueing and production systems. *Tech. Rep. PhD 2003-4, ISR, University of Maryland, College Park, MD*, 2003. 2, 3
- [43] X. Yao, X. Xie, M. Fu, and S. Marcus. Optimal joint preventative maintenance and production policies. *Naval Research Logistics*, 2005. 2

7 Appendix

7.1 Stability

In this section, fluid limit analysis and results of Dai [5] and Dai and Meyn [6] are used to prove Proposition 3.1. The result is first proved for $\ell = 1$. Recall ϕ_1 : do not repair until the server state process $S(t)$ deteriorates out of server state 1 (when repair is forced), independent of the number of customers in the system. It is assumed that all processes in this section operate under this policy; the dependence of quantities on ϕ_1 is suppressed.

For the analysis, it is constructive to augment the state space by a third process, $R(t)$, the residual repair process, defined as the time remaining until the current repair in progress at time t is completed. Define the Markovian state of the system at time t under policy ϕ_1 as

$$X(t) = (Q(t), S(t), R(t)).$$

When $S(t) > 0$, and a repair is not in progress, $R(t) = 0$. As soon as the server-state process enters state 0, there is a jump in $R(t)$. Assume that $R(t)$ is right continuous. Note that after a (random) jump to a positive value, the trajectory of $R(t)$ is decreasing at rate 1 until it reaches 0 again. In between positive jumps, the trajectory of $R(t)$ is deterministic. That is, it is “piecewise-deterministic.” Similar to the processes considered in [5], $\{X(t), t \geq 0\}$ is a piecewise-deterministic Markov process, with state space $\mathcal{X} = \{\mathbb{Z}^+ \times \{0\} \times [0, \infty)\} \cup \{\mathbb{Z}^+ \times \{1, \dots, B\} \times \{0\}\}$. It follows from [8] p. 362 that $\{X(t), t \geq 0\}$ is a strong Markov process.

For a state $x = (q, s, r) \in \mathcal{X}$, define the norm of x to be $|x| \equiv q + s + r$. Consider the scaled process

$$\bar{Q}^x(t) \equiv \frac{1}{|x|} Q^x(|x|t),$$

where the superscript now denotes the dependence on the initial state $x = X(0)$. Similar law-of-large-numbers type scalings for other processes will also be denoted with the bar symbol. Letting $|x| \rightarrow \infty$, any limit point $\bar{Q}(t)$ is called a *fluid limit* of the queue length process. We will show that every fluid limit is a solution to a set of equations known as the *fluid model*. The fluid model is said to be *stable* if there exists

a fixed time t_0 such that $\bar{Q}(t) = 0$ for all $t \geq t_0$. That is, the fluid model is stable if the queue eventually drains and once drained stays empty.

Let the customer inter-arrival times be $\xi(n)$, $n = 1, 2, \dots$, and the service requirements be $\eta(n)$, $n = 1, 2, \dots$. Recall that it is assumed that ξ and η are sequences of i.i.d. exponential random variables with means $1/\lambda$ and 1 , respectively. Let the deterioration times out of server state s be $\gamma_s(n)$, $n = 1, 2, \dots$, i.i.d. random variables with mean $1/m_s$; $s \in \{0, 1, \dots, B\}$. The inter-arrival times, service requirements, and deterioration times are all mutually independent. The times γ_0 are actually repair times, not deterioration times. Define the cumulative processes

$$E(t) = \max\{n \geq 0 : \xi(1) + \xi(2) + \dots + \xi(n) \leq t\}, \quad t \geq 0,$$

$$D(t) = \max\{n \geq 0 : \eta(1) + \eta(2) + \dots + \eta(n) \leq t\}, \quad t \geq 0,$$

$$Z_s(t) = \max\{n \geq 0 : \gamma_s(1) + \gamma_s(2) + \dots + \gamma_s(n) \leq t\}, \quad t \geq 0, \quad s \in \{1, 2, \dots, B\},$$

$$Z_0(t) = \max\{n \geq 0 : R(0) + \gamma_0(1) + \gamma_0(2) + \dots + \gamma_0(n-1) \leq t\}, \quad t \geq 0.$$

Let $Y_s^x(t)$ be the cumulative amount of time the server is in state s in $[0, t)$ given the initial state x . Then $Z_s(Y_s^x(t))$, $s \in \{1, 2, \dots, B\}$, is the number of transitions from s to $s-1$ completed in $[0, t)$. $Z_0(Y_0^x(t))$ is similar. Let $I_s^x(t)$ be the cumulative amount of time that there are no customers in the system while the server state is s , and let $T_s^x(t) = Y_s^x(t) - I_s^x(t)$. $T_s^x(t)$ is the cumulative amount of time that there are customers in the system while the server state is s , by time t . Let $W^x(t)$ be the cumulative work done by time t , so that $W^x(t) = \mu_0 T_0^x(t) + \mu_1 T_1^x(t) + \dots + \mu_B T_B^x(t)$; the fluid limit results do not require $\mu_0 = 0$. The definitions above imply the following system of equations:

$$Q^x(t) = Q^x(0) + E^x(t) - D^x(W^x(t)), \quad (7.1)$$

$$Q^x(t) \geq 0, \quad (7.2)$$

$$\sum_{s=0}^B Y_s^x(t) = t, \quad (7.3)$$

$$T_s^x(t) + I_s^x(t) = Y_s^x(t), \quad s = 0, 1, \dots, B, \quad (7.4)$$

$$W^x(t) = \sum_{s=0}^B \mu_s T_s^x(t), \quad (7.5)$$

$$Y_s^x(t), I_s^x(t), T_s^x(t), \text{ and } W^x(t) \text{ are non-decreasing and start from 0,} \quad (7.6)$$

$$\int_0^\infty Q^x(t) d\left(\sum_{s=0}^B I_s^x(t)\right) = 0, \quad (7.7)$$

$$|[Z_s(Y_s^x(u)) - Z_s(Y_s^x(t))] - [Z_l(Y_l^x(u)) - Z_l(Y_l^x(t))]| \leq 1, \quad t \leq u, \quad s, l \in \{0, 1, \dots, B\}. \quad (7.8)$$

Note that (7.7) is the non-idling constraint and guarantees that the idle time increases only if the total queue length is zero. The constraints (7.8) result from the cyclic nature of the deteriorations and repairs under the policy that only repairs when the machine fails. They follow since over any interval of time, the number of transitions out of server state s differs from the number of transitions out of any other server state l by at most 1 (recall the definition of ϕ_1). These constraints are similar to constraints for polling models presented in [6].

The next proposition, a variant of Theorem 4.1 in [5], presents the fluid model and establishes convergence of the scaled processes. This convergence is *uniform on compact sets (u.o.c.)*.

Proposition 7.1 *The following holds with probability one. For any sequence of initial states $\{x_j\} \in \mathcal{X}$ with $|x_j| \rightarrow \infty$ and $\bar{R}^{x_j}(0) \rightarrow 0$, there exists a subsequence $\{x_i\}$, $\{i\} \subseteq \{j\}$, with $|x_i| \rightarrow \infty$ such that*

$$(\bar{Q}^{x_i}(0), \bar{S}^{x_i}(0), \bar{R}^{x_i}(0)) \rightarrow (\bar{Q}(0), 0, 0), \quad (7.9)$$

$$(\bar{Q}^{x_i}(t), \bar{T}^{x_i}(t), \bar{Y}^{x_i}(t)) \rightarrow (\bar{Q}(t), \bar{T}(t), \bar{Y}(t)) \text{ u.o.c.}, \quad (7.10)$$

where $(\bar{Q}(t), \bar{T}(t), \bar{Y}(t))$ satisfies the following set of equations:

$$\bar{Q}(t) = \bar{Q}(0) + \lambda t - \bar{W}(t), \quad (7.11)$$

$$\bar{Q}(t) \geq 0, \quad (7.12)$$

$$\sum_{s=0}^B \bar{Y}_s(t) = t, \quad (7.13)$$

$$\bar{T}_s(t) + \bar{I}_s(t) = \bar{Y}_s(t), \quad s = 0, 1, \dots, B, \quad (7.14)$$

$$\bar{W}(t) = \sum_{s=0}^B \mu_s \bar{T}_s(t), \quad (7.15)$$

$$\bar{Y}_s(t), \bar{I}_s(t), \bar{T}_s(t), \text{ and } \bar{W}^x(t) \text{ are non-decreasing and start from 0,} \quad (7.16)$$

$$\int_0^\infty \bar{Q}(t) d\left(\sum_{s=0}^B \bar{I}_s(t)\right) = 0, \quad (7.17)$$

$$m_s [\bar{Y}_s(u) - \bar{Y}_s(t)] - m_l [\bar{Y}_l(u) - \bar{Y}_l(t)] = 0, \quad t \leq u, \quad s, l \in \{0, 1, \dots, B\}. \quad (7.18)$$

Proof: Notice that $\frac{1}{|x_j|}Q^{x_j}(0) \leq 1$, $\frac{1}{|x_j|}S^{x_j}(0) \leq 1$, $\frac{1}{|x_j|}R^{x_j}(0) \leq 1$, so there exists a subsequence along which each of these converge. Also note, $0 \leq S^{x_j}(0) \leq B$, so $\lim S^{x_j}/|x_j| \rightarrow 0$ as expected. Therefore, there exists a subsequence $|x_i| \rightarrow \infty$ such that (7.9) holds. For any $0 \leq t_1 \leq t_2$, and each s , we have $0 \leq \bar{Y}_s^x(t_2) - \bar{Y}_s^x(t_1) \leq t_2 - t_1$. That is, $\{\bar{Y}_s^x(t), |x| \geq 1\}$ is uniformly Lipschitz, and hence equicontinuous. Setting $t_1 = 0$ and $t_2 = t$ yields uniform bounds $0 \leq \bar{Y}_s^x(t) \leq t$. Therefore, by the Arzelá-Ascoli theorem, any subsequence of $\bar{Y}_s^x(t)$ has a u.o.c. convergent subsequence. The families $\{\bar{T}_s^x(t), |x| \geq 1\}$ and $\{\bar{I}_s^x(t), |x| \geq 1\}$ can be shown to have a u.o.c. convergent subsequence in the same manner. Similar to the proof of Lemma 4.2 in [5] we have

$$\bar{E}(t) = \lambda t, \quad (7.19)$$

$$\bar{D}(t) = 1t, \quad (7.20)$$

$$\bar{Z}_s(t) = m_s t, \quad s \in \{0, 1, \dots, B\}. \quad (7.21)$$

Equation (7.11) follows from (7.19), (7.20), and (7.1), where the random time change is valid by Theorem 5.3 of [4]. From (7.8), $[Z_s(Y_s^x(u)) - Z_s(Y_s^x(t))]$ differs from $[Z_l(Y_l^x(u)) - Z_l(Y_l^x(t))]$ by at most one. In the fluid limit, this difference is negligible. Therefore, (7.18) follows from (7.8) and (7.21); again, the random time change is valid by Theorem 5.3 of [4]. Equations (7.12)–(7.16) are a consequence of (7.2)–(7.6) and (7.17) follows from (7.7) and Lemma 4.4 of [5]. ■

The remainder of this section is dedicated to completing the proof of Proposition 3.1.

Proof of Proposition 3.1. First consider the case $\ell = 1$. We want to show that the fluid model for ϕ_1 is stable. The non-idling constraint (7.17) implies that when $\bar{Q}(t) > 0$, $\bar{I}_s = 0$ for all s . Hence, when $\bar{Q}(t) > 0$, from (7.14), $\bar{Y}_s(t) = \bar{T}_s(t)$, $s \in \{0, 1, \dots, B\}$. Because $\bar{T}_s(t)$, $\bar{I}_s(t)$, and $\bar{Y}_s(t)$ are Lipschitz continuous, they are absolutely continuous and differentiable almost everywhere. Taking the derivative with respect to t , $\dot{\bar{Y}}_s(t) = \dot{\bar{T}}_s(t)$ when $\bar{Q}(t) > 0$. From (7.18) it follows that $m_s \dot{\bar{Y}}_s(t) = m_l \dot{\bar{Y}}_l(t)$, which implies $m_s \dot{\bar{T}}_s(t) = m_l \dot{\bar{T}}_l(t)$, $s, l \in \{0, 1, \dots, B\}$. Substituting into (7.13) gives $\frac{m_s}{m_0} \dot{\bar{T}}_s(t) + \frac{m_s}{m_1} \dot{\bar{T}}_s(t) + \dots + \frac{m_s}{m_B} \dot{\bar{T}}_s(t) = 1$, or $\dot{\bar{T}}_s(t) = (1/m_s)(1/m_0 + 1/m_1 + \dots + 1/m_B)^{-1}$, $s \in \{0, 1, \dots, B\}$. So, from (7.11), for $\bar{Q}(t) > 0$,

$$\dot{\bar{Q}}(t) = \lambda - \sum_{s=0}^B \frac{\mu_s/m_s}{1/m_0 + 1/m_1 + \dots + 1/m_B}.$$

The fluid limit of the queue length process drains when $\dot{Q}(t) < 0$. The queue will drain when (3.11) holds for $\ell = 1$. Furthermore, under (3.11) with $\ell = 1$, Lemma 5.2 of [5] implies that once the fluid limit is drained, it will remain drained. Therefore, the fluid model is stable.

It follows from Theorem 4.2 of [5] that $\{X(t), t \geq 0\}$ is positive Harris recurrent and a stationary distribution exists if the sublevel set $\{x : |x| \leq \kappa\}$ is a *petite* set for any $\kappa > 0$. The proof of this fact follows in much the same manner as Lemma 3.7 of [20] and is omitted for brevity; the interested reader is referred to [16], Lemma 3.5 for complete details.

Since $\{X(t), t \geq 0\}$ is positive recurrent, so is the semi-Markov process (SMP) induced by the policy ϕ_1 . Moreover, since the transition rates are bounded below, the embedded (discrete-time) Markov chain is also positive recurrent. Irreducibility and aperiodicity of this Markov chain are trivial.

It remains to show that $g^{\phi_1} < \infty$. Using the same analysis for the residual service times in [6] for the repair process $R(t)$, we get from Theorem 4.1 of [6] that $\limsup_{t \rightarrow \infty} \frac{\mathbb{E}^{\phi_1} \int_0^t Q(u) du}{t} < \infty$. Since we have linear holding costs, the long-run average holding cost rate is finite. The long-run average number of repairs is less than the long-run average number of renewals for a renewal process with average inter-arrival time $1/m_0$. This together with the assumption that $K < \infty$ imply $g^{\phi_1} < \infty$ as desired.

The result for $\ell > 1$ follows by simply relabeling states 1 and 0 in the analysis above as ℓ and $\ell - 1$, respectively. The states with $s \in \{0, 1, \dots, \ell - 2\}$, for which repair is chosen, are transient since they are visited at most once. Since the amount of time spent in one of these transient states is finite with probability 1, the long-run average cost for starting in a transient state is the same as it is for starting in any recurrent state. ■

7.2 Replacement Model Properties

Proof of Lemma 4.8. The proof is by induction on n . For $n = 0$, $v_{0,\theta}(q, s) - v_{0,\theta}(q + 1, s) = 0$; (4.14) holds. Assume that (4.14) holds for n . Consider $n + 1$. For $s \in \{1, 2, \dots, B\}$,

$$\begin{aligned}
& f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q + 1, s) \\
&= (\theta + 1)^{-1} \left[-h + \lambda(v_{n-1,\theta}(q + 1, s) - v_{n-1,\theta}(q + 2, s)) \right. \\
&\quad + \mu_s(v_{n-1,\theta}((q - 1)^+, s) - v_{n-1,\theta}(q, s)) + m_s(v_{n-1,\theta}(q, s - 1) - v_{n-1,\theta}(q + 1, s - 1)) \\
&\quad \left. + (1 - \lambda - \mu_s - m_s)(v_{n-1,\theta}(q, s) - v_{n-1,\theta}(q + 1, s)) \right] \\
&\leq 0,
\end{aligned}$$

where the inequality follows by the induction assumption. Hence, (4.15) holds for $n + 1$ given that (4.14) holds for n . This implies, for $s = 0$,

$$v_{n+1,\theta}(q, 0) - v_{n+1,\theta}(q + 1, 0) = f_{n+1,\theta}(q, B) - f_{n+1,\theta}(q + 1, B) \leq 0.$$

For $s \in \{1, 2, \dots, B\}$,

$$\begin{aligned} & v_{n+1,\theta}(q, s) - v_{n+1,\theta}(q + 1, s) \\ &= \min \{f_{n+1,\theta}(q, s), K(s) + f_{n+1,\theta}(q, B)\} - \min \{f_{n+1,\theta}(q + 1, s), K(s) + f_{n+1,\theta}(q + 1, B)\} \\ &\leq \max \{f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q + 1, s), f_{n+1,\theta}(q, B) - f_{n+1,\theta}(q + 1, B)\} \\ &\leq 0, \end{aligned}$$

since (4.15) holds for $n + 1$. Hence, (4.14) holds for $n + 1$. Therefore, by induction, (4.14) and (4.15) hold for all $n \geq 1$. \blacksquare

Proof of Lemma 4.9. The proof is by induction on n . For $n = 0$, $v_{0,\theta}(q, s) - v_{0,\theta}(q, s - 1) = K(s) - K(s - 1)$; (4.16) holds. Assume that (4.16) holds for n . Consider $n + 1$. For $s \in \{2, 3, \dots, B - 1\}$,

$$\begin{aligned} & f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q, s - 1) \\ &= (\theta + 1)^{-1} [\lambda v_{n,\theta}(q + 1, s) + \mu_s v_{n,\theta}((q - 1)^+, s) + m_s v_{n,\theta}(q, s - 1) \\ &\quad + (1 - \lambda - \mu_s - m_s) v_{n,\theta}(q, s) - \lambda v_{n,\theta}(q + 1, s - 1) - \mu_{s-1} v_{n,\theta}((q - 1)^+, s - 1) \\ &\quad - m_{s-1} v_{n,\theta}(q, s - 2) - (1 - \lambda - \mu_{s-1} - m_{s-1}) v_{n,\theta}(q, s - 1)] \\ &\leq (\theta + 1)^{-1} [\lambda (v_{n,\theta}(q + 1, s - 1) + [K(s) - K(s - 1)]) \\ &\quad + \mu_s (v_{n,\theta}((q - 1)^+, s - 1) + [K(s) - K(s - 1)]) + m_s v_{n,\theta}(q, s - 1) \\ &\quad + (1 - \lambda - \mu_s - m_s) (v_{n,\theta}(q, s - 1) + [K(s) - K(s - 1)]) \\ &\quad - \lambda v_{n,\theta}(q + 1, s - 1) - \mu_{s-1} v_{n,\theta}((q - 1)^+, s - 1) - m_{s-1} v_{n,\theta}(q, s - 2) \\ &\quad - (1 - \lambda - \mu_{s-1} - m_{s-1}) v_{n,\theta}(q, s - 1)], \end{aligned}$$

where the inequality follows from the induction assumption. After again applying the induction assumption, this expression reduces to

$$\begin{aligned} f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q, s - 1) &\leq (\theta + 1)^{-1} [(\mu_s - \mu_{s-1}) (v_{n,\theta}((q - 1)^+, s - 1) - v_{n,\theta}(q, s - 1)) \\ &\quad + m_{s-1} [K(s - 1) - K(s - 2)] + (1 - m_s) [K(s) - K(s - 1)]]. \end{aligned}$$

Since $K(s)$ is convex, $K(s-1) - K(s-2) \leq K(s) - K(s-1)$, and

$$\begin{aligned} f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q, s-1) &\leq (\theta+1)^{-1}(\mu_s - \mu_{s-1})(v_{n,\theta}((q-1)^+, s-1) - v_{n,\theta}(q, s-1)) \\ &\quad + (\theta+1)^{-1}(1 - (m_s - m_{s-1}))[K(s) - K(s-1)]. \end{aligned} \quad (7.22)$$

The first term on the right hand side of (7.22) is non-positive by (4.14) and $\mu_s \geq \mu_{s-1}$. Hence,

$$f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q, s-1) \leq (\theta+1)^{-1}(1 - (m_s - m_{s-1}))[K(s) - K(s-1)]. \quad (7.23)$$

When $m_{s-1} \geq (\leq) m_s + \theta$, $(\theta+1)^{-1}(1 - (m_s - m_{s-1})) \geq (\leq) 1$. So, under either Condition 1 or 2, (4.17) holds for $n+1$ given that (4.16) holds for n .

To complete the proof, we need to show that (4.16) holds for $n+1$. For $s=1$,

$$\begin{aligned} v_{n+1,\theta}(q, 1) - v_{n+1,\theta}(q, 0) &= \min\{f_{n+1,\theta}(q, 1), K(1) + f_{n+1,\theta}(q, B)\} - K(0) - f_{n+1,\theta}(q, B) \\ &\leq K(1) - K(0). \end{aligned}$$

For $s \in \{2, 3, \dots, B-1\}$,

$$\begin{aligned} &v_{n+1,\theta}(q, s) - v_{n+1,\theta}(q, s-1) \\ &= \min\{f_{n+1,\theta}(q, s), K(s) + f_{n+1,\theta}(q, B)\} - \min\{f_{n+1,\theta}(q, s-1), K(s-1) + f_{n+1,\theta}(q, B)\} \\ &\leq \max\{f_{n+1,\theta}(q, s) - f_{n+1,\theta}(q, s-1), K(s) - K(s-1)\} \\ &\leq K(s) - K(s-1), \end{aligned}$$

where the last inequality follows since (4.17) holds for $n+1$. Hence, (4.16) holds for $n+1$. Therefore, by induction, (4.16) and (4.17) hold for all $n \geq 1$. ■