Adv. Appl. Prob. **21**, 376–397 (1989)
Printed in N. Ireland

OPTIMAL TIME TO REPAIR A BROKEN SERVER

AWI FEDERGRUEN,* *Columbia University*
KUT C. SO,** *AT&T Bell Laboratories*

Abstract

We consider a single-server queueing system with Poisson arrivals and general service times. While the server is up, it is subject to breakdowns according to a Poisson process. When the server breaks down, we may either repair the server immediately or postpone the repair until some future point in time. The operating costs of the system include customer holding costs, repair costs and running costs. The objective is to find a corrective maintenance policy which minimizes the long-run average operating costs of the system. The problem is formulated as a semi-Markov decision process. Under some mild conditions on the repair time and service time distributions and the customer holding cost rate function, we prove that there exists an optimal stationary policy which is characterized by a single threshold parameter: a repair is initiated if and only if the number of customers in the system exceeds this threshold. We also show how the average cost under such policies may be computed and how an optimal policy may efficiently be determined.

QUEUES WITH BREAKDOWNS; SEMI-MARKOV DECISION PROCESS; CORRECTIVE MAINTENANCE POLICY; MONOTONE POLICIES

1. Introduction

This paper considers the problem of determining optimal repair policies for operating devices that are subject to breakdowns. This is a common problem in reliability models and a large amount of research on optimal operating policies for maintenance systems (see e.g. McCall (1965), Pierskalla and Voelker (1976), Sherif and Smith (1981)) exists in the literature. The major difference between our model and most existing repair models is that we view the operating devices as servers in a queueing system, providing service to arriving customers.

We consider a single-server queueing system with Poisson arrivals and general i.i.d. service times. While the server is up, it is subject to breakdowns according to a Poisson process. When the server breaks down, we may either repair the server immediately or postpone the repair until some future point in time. The operating costs of the system include customer holding costs, repair costs and running costs. The objective is to find a corrective maintenance policy which minimizes the

Received 8 December 1987; revision received 6 May 1988.

* Postal address: Graduate School of Business. 412 Uris Hall, Columbia University, New York, NY 10027, USA. Research partly supported by NSF Grant no. ECS-8604409, and partly carried out at the Hebrew University, Jerusalem.

** Postal address: AT&T, Bell Laboratories, P.O. Box 900, Princeton, NJ 08540, USA.

long-run average operating costs of the system. We refer to this problem as the wait/repair problem.

Basically, the advantage of considering the operating device as a server in a queueing system is that we can take into account the costs due to customers waiting in the system, which in turn implies that we take into account the workload of the system at the time it fails in considering whether to repair the system immediately or not. This formulation has the flexibility of assigning a higher operating cost when there are more customers in the system. This higher cost may be due to the higher workload in the system. It also may be due to greater customer dissatisfaction because of longer waiting times, or to lost productivity (profit) if the customers are internal to the organization.

One possible extension of our basic model is to include the option of turning off the server even when it is operational, at the expense of a fixed shutdown cost. This would permit saving on the operating costs. This extended problem can be viewed as an extension of the problem of optimally operating an $M/G/1$ queueing system with a removable server and can be solved by a slight modification of the analysis in this paper. Our results thus generalize those in Yadin and Naor (1963), Sobel (1969) and Bell (1971) for $M/G/1$ systems without breakdowns and repairs. See also Balanchandran (1973), Balanchandran and Tijms (1975) and Heyman (1977) for the analysis of special classes of policies for this model. We shall discuss this in greater detail in Section 6.

There is also similarity between queues with breakdowns and preemptive priority queues (see e.g. Gaver (1962), White and Christie (1958) and Federgruen and Green (1986), (1988)). For example, assume that there are two priority job classes in which the high-priority jobs can preempt any low-priority job. Customers correspond to the low-priority jobs. A breakdown corresponds to an arrival of a high-priority job which preempts the current low-priority job in service. The repair time corresponds to the service time of a high-priority job or to the busy period initiated by a high-priority job. Therefore, queues with breakdowns may also find applications for the corresponding priority queueing models.

The rest of the paper is organized as follows. In Section 2 we formulate the wait/repair problem as a semi-Markov decision process. In Section 3 we show under some mild conditions with respect to the repair and service time distributions as well as the holding cost function, that an optimal stationary policy exists. In Section 4 we prove that an optimal stationary policy exists which is characterized by a single threshold parameter: a repair is initiated if and only if the number of customers in the system exceeds this threshold. We refer to such policies as *monotone*. In Section 5 we demonstrate how the average cost of a monotone policy may be computed and how an optimal monotone policy may efficiently be determined. Finally, in Section 6, we give a few remarks on some special cases and extensions of our model.

Our proof that *optimal monotone* policies exist is, to our knowledge, based on a somewhat novel approach. In the literature on optimal control problems of stochastic systems, most structural results are obtained by verifying that value functions or solutions to certain optimality equations have specific properties such as monotonicity, convexity, K-convexity, etc. Our proof is based on a parametric variation of one of the model parameters. The existence of a monotone optimal policy is first verified when this parameter takes on sufficiently small values and is then inductively extended from interval to interval. We believe that this approach may prove to be fruitful in many other models as well.

2. Model formulation

We consider a single-server queue with Poisson arrivals at rate λ and general i.i.d. service times with service time distribution $F(.)$. While the server is up and working, it is subject to breakdowns according to a Poisson process with rate σ where $0<\sigma<\infty$, and possibly with a different rate σ_0 with $0\leqq\sigma_0<\infty$ when the server is up but idle. We assume that the arrival, service and breakdown mechanisms are independent of each other. When a service is interrupted, it needs to be restarted from scratch, i.e., when the service of a job is restarted, its service time is a new independent sample from the distribution $F(.)$. In the special case where service times are exponential, it is immaterial whether an interrupted service due to breakdowns is resumable or not. When the server breaks down, we can either initiate a repair or postpone the repair (wait) until some future point in time. Repair times are i.i.d. with general distribution $G(.)$ and mean v, $0<v<\infty$.

The cost structure includes a customer holding cost rate $H(i)\geqq 0$ where i represents the number of customers in the system, a fixed cost $c\geqq 0$ for each repair, a running cost rate r when the server is working, and a possibly different running cost rate r_0 when the server is idle. The running cost is 0 when the server is broken. In some cases, $r_0\leqq r$ because additional costs (resources) are needed to serve the customers, e.g., electricity cost to run the machine. However, we do not impose any restriction on the relation between r and r_0 in our model. We can also include a reward R for each service completion. The long-run average value of this reward component is, however, constant for any policy under which the queueing system is stable. We therefore assume $R=0$. Our objective is to minimize the average operating costs of the system among all possible policies.

The wait/repair problem is formulated as a semi-Markov decision process with state space $S=\{(i,j)\mid i=0,1,2,\cdots,\text{ and } j=0,1\}$. $(X(t),Y(t))\in S$ denotes the state of the system at time t where $X(t)$ represents the number of customers in the system at time t and $Y(t)$ the status of the server at time t, with value 0 or 1 indicating that the server is down or up, respectively. The decision epochs include (i) breakdowns or service completions when the server is up and working; (ii) arrivals of customers when the server is idle or down and no repair has been initiated yet; and (iii) repair completions.

For each decision epoch at which $Y(t)=1$, there is only one action available, namely, do nothing and let the server work. For those epochs with $Y(t)=0$, there are two available actions: *wait* (denoted by action 1) and *repair* (denoted by action 2). A pure stationary policy f is called *monotone* if for all $i<j$, $f(i,0)\leqq f(j,0)$. In particular, we call a monotone policy an *n-policy* where $n=\min\{i \mid f(i,0)=2\}$, i.e., an n-policy repairs the server when the number of customers in the system is equal to or exceeds n.

Let $Z(x,t)$ be the total cost incurred up to time t when starting in state x. The average cost vector under policy π, ϕ_π, is defined as $\phi_\pi(x)=\limsup_{t\to\infty} E_\pi\{Z(x,t)/t\}$ and the minimum average cost vector ϕ is given by $\phi(x)=\inf_\pi \phi_\pi(x)$. A policy π^* is *average optimal* if $\phi_{\pi^*}(x)=\phi(x)$ for all $x\in S$. For policies whose average cost is independent of the initial state x, we write ϕ_π instead of $\phi_\pi(x)$.

3. Existence of an optimal stationary policy

We first impose the following three conditions.

Condition 1. $H(i)$ is non-decreasing in i and bounded above by a polynomial. Specifically, $H(i)\leqq hi^m$ for some constant $h>0$ and integer $m\geqq 1$. Furthermore, $H(i)\to\infty$ as $i\to\infty$.

Condition 2. The $(m+1)$th moment of the repair time distribution function $G(.)$ is finite; let $v^{(l)}$ be the lth moment of the repair time, $l=1,2,\cdots,m+1$. (We write v instead of $v^{(1)}$.)

Condition 3.

$$\lambda\{\tilde{F}(\sigma)^{-1}s+(\tilde{F}(\sigma)^{-1}-1)v\}<1 \quad \text{where} \quad \tilde{F}(\sigma)=\int_0^\infty e^{-\sigma t}\,dF(t).$$

For any $x,y\in S$, let $p_{xy}(a)$ denote the one-step transition probability from state x to state y if action a is taken. Note that $\hat{F}(.)$ defined by $\hat{F}(t)=1-e^{-\sigma t}(1-F(t))$ $(t>0)$ represents the c.d.f. of a *truncated service time*, defined as the time between the initiation of a service and the next decision epoch (either the completion of this service or a breakdown, whatever comes first). Let $s^{(l)}$ be the lth moment of $\hat{F}(.)$, $l\geqq 1$. Since the truncated service time is always smaller than or equal to an exponential breakdown time, all of its moments are finite. We write s instead of $s^{(1)}$. The one-step expected holding times $t(x;a)$ are clearly given by

$$t((i,1);1)=\begin{cases} s, & i>0 \\ (\lambda+\sigma_0)^{-1}, & i=0 \end{cases}$$

$$t((i,0);1)=\lambda^{-1}$$

$$t((i,0);2)=v.$$

Note that all expected holding times are finite and uniformly bounded by two positive numbers $t_{\min}$ and $t_{\max}$, i.e., $0 < t_{\min} \leqq t(x; a) \leqq t_{\max}$ for all $x \in S$ and $a = 1, 2$.

Let $c(x; a)$ denote the one-step expected costs in state $x \in S$ when action a is taken. Clearly,

(1) $c((i, 0); 1) = H(i)\lambda^{-1}$

$$c((i, 0); 2) = c + \sum_{k=0}^{\infty} \int_0^{\infty} \frac{e^{-\lambda t}(\lambda t)^k}{k!} \left\{ \sum_{l=0}^{k} H(i+l) \frac{t}{(k+1)} \right\} dG(t)$$

$$\leqq c + h \int_0^{\infty} \left\{ \sum_{k=0}^{\infty} \frac{e^{-\lambda t}(\lambda t)^k t}{k!} (i+k)^m \right\} dG(t)$$

$$= c + h \sum_{l=0}^{m} \binom{m}{l} i^{m-l} \int_0^{\infty} \sum_{k=0}^{\infty} \frac{e^{-\lambda t}(\lambda t)^k t k^l}{k!} dG(t) \tag{2}$$

$$= c + h \sum_{l=0}^{m} \binom{m}{l} i^{m-l} \sum_{p=0}^{l} S_{l,p} \int_0^{\infty} \sum_{k=0}^{\infty} \frac{e^{-\lambda t}(\lambda t)^k t k \cdots (k-p+1)}{k!} dG(t)$$

$$= c + h \sum_{l=0}^{m} \binom{m}{l} i^{m-l} \sum_{p=0}^{l} S_{l,p} \lambda^p \nu^{(p+1)} = O(i^m)$$

where $S_{l,p}$ is the (l, p)th Stirling number of the second kind, see e.g. Selby (1980). (The first identity follows by conditioning on the repair time t and k, the number of customers that arrive during this repair time, as well as by the observation that conditional upon k Poisson arrivals occurring in an interval of length t, the interarrival times have mean $t/(k+1)$, see e.g. Ross (1970). The first inequality follows from $H(.)$ being non-decreasing and the polynomial bound in Condition 1.) Similarly,

$$c((0, 1); 1) = \frac{H(0) + r_0}{(\lambda + \sigma_0)} \tag{3}$$

$$c((i, 1); 1) = O(i^m). \tag{4}$$

(Verification of (4) is analogous to that of (2), replacing G by $\hat{F}$, ν by s, and c by 0.)

Conditions 1 and 2 thus ensure that all one-step expected costs are finite and polynomially bounded. As shown below, Condition 3 guarantees that a stationary policy exists with finite long-run average cost, namely the *no-wait* or 0-policy under which the server is repaired as soon as broken.

First, define C_i as the expected total cost incurred under the no-wait policy when starting in state $(i+1, 1)$ until reaching state $(i, 1)$.

Lemma 3.1.

(a) $C_i = O(i^m)$, $i \geqq 1$.

(b) The long-run average cost of the no-wait policy is finite.

Proof. (a) Consider the queueing system with service interruptions that arises under the no-wait policy when the holding cost rate function is given by $H_i(.)$ with $H_i(k) = H(i+k)$ $(k \geqq 0)$, while in all states the operating cost rate is r and breakdowns occur with rate σ. A busy period is defined to start with the arrival of a customer to an 'empty' system and to terminate when the system empties out next. Note that C_i is the conditional expected cost in a *single busy period* of this queueing system given that the busy period starts with the server up. The cost process in such a busy period is equivalent to that in an ordinary $M/G/1$ queue, with a customer's *completion time* Γ, i.e., the time between the start of a customer's service and its termination, as the 'service time'.

One easily verifies, see e.g. Gaver (1962), that a completion time consists of the convolution of N independent truncated service times and $(N-1)$ independent repair times where N is a geometrically distributed random variable with parameter $\tilde{F}(\sigma)=\int_0^\infty e^{-\sigma t}\,dF(t)$ (i.e., with mean $[\tilde{F}(\sigma)]^{-1}$), i.e.,

$$\Gamma=\sum_{p=1}^{N} S_p^T+\sum_{p=1}^{N-1} R_p, \tag{5}$$

$\{S_1^T,\cdots,S_N^T\}$ denote independent truncated service times and $\{R_1,\cdots,R_{N-1}\}$ independent repair times. One easily verifies from (5) that the first $(m+1)$ moments of Γ are finite since the corresponding moments of $G(.)$ are. (See e.g. Lemmas 2 and 3 in Federgruen and Green (1986)). For example, $E(\Gamma)=[\tilde{F}(\sigma)]^{-1}s+([\tilde{F}(\sigma)]-1)v$, as is easily verified with Wald's lemma.

Condition 3 represents the stability condition for the queueing system that arises under the no-wait policy; the system is thus regenerative with terminations of busy periods as regeneration epochs. Let γ_i represent the long-run average cost under the no-wait policy with holding cost rate function $H_i(.)$ (and all other cost components as specified). Let T represent the length of a busy cycle (defined as the interval between two consecutive regeneration epochs) and L the number of service completions in a busy cycle. Since the first $(m+1)$ moments of the completion time Γ are finite, it follows as in ordinary $M/G/1$ systems that the first m moments of L are finite, so also Wolff (1984). We conclude that

$$hE(L+i)^m+r+c\{E(L)(\tilde{F}(\sigma)^{-1}-1)+\sigma_0/\lambda\}/E(T)\geqq\gamma_i\geqq C_i/E(T)-r. \tag{6}$$

The first inequality follows from $H_i(k)=H(i+k)\leqq h(i+k)^m$ and the fact that the average repair costs per unit time are given by $E(T)^{-1}E\{$expected total repair costs in a busy cycle$\}\leqq E(T)^{-1}c\{E(L)\sigma/\mu+\sigma_0/\lambda\}$ since the expected number of breakdowns in a busy cycle prior to the first arrival is bounded by σ_0/λ while the expected number of breakdowns thereafter equals, by Wald's lemma, $E(L)$ times the expected number of breakdowns that occur during a customer's service, and the latter equals $[\tilde{F}(\sigma)]^{-1}-1$. The second inequality follows from the observation that $C_i-rE(T)\leqq E\{$total holding and repair costs in a busy period which starts with the server up$\}\leqq E\{$total holding and repair costs in an arbitrary busy period$\}\leqq\gamma_iE(T)$. We conclude that $C_i=O(i^m)$.

(b) Immediate from the bound for γ_0 in (6).

For any integer $u>0$, let $S_u=\{(i,j);\,0\leqq i\leqq u$ and $j=0$ or $1\}$. Let $M(x,y)$ denote the minimal expected cost incurred when starting in state $x \in S$ until the first visit to state $y \in S$.

Lemma 3.2. Fix $u>0$. There exists a number $K>0$ such that $M(x, y) \leqq K \max\{i^{m+1}, 1\}$ for all $x=(i, j)$ ($j=0$ or 1) and all $y \in S_u$.

Proof. Since S_u is a finite subset of the state space, it suffices to prove that the inequality holds for a fixed state $y^0 \in S_u$. Let $M^*(x, y^0)$ be the expected total cost incurred when starting in state $x \in S$ until the first visit to state y^0 under the no-wait policy. Similarly, let $M^1(x)$ denote the expected total cost incurred under the no-wait policy when starting in state $x \in S$ up until the next visit to the set S_u. We show that a constant $K>0$ exists such that

$$(7) \qquad M^1((i, j)) \leqq K i^{m+1} \quad \text{for all} \quad i \geqq 1 \quad \text{and} \quad j=0, 1.$$

This implies in particular that $M^1(x)$ is *finite* for all $x \in S_u$.

Note that $M^*(x, y^0)$ may be decomposed as:

$$(8) \qquad M^*(x, y^0) = M^1(x) + M^2(x, y^0)$$

with $M^2(x, y^0)$ the expected total cost incurred under the no-wait policy between the first entry into the set S_u (starting in state x) and the first visit to state y^0. (If the set S_u is entered in state y^0, the cost on this trajectory is defined to be 0). Consider under the no-wait policy the Markov chain embedded on visits to the set S_u. All states in S_u belong to a single recurrent set of states on this (finite state) embedded Markov chain, since they all communicate with each other in the original semi-Markov process. Thus, let F be the largest expected first passage time from any state in S_u to y^0. Clearly,

$$(9) \qquad M^2(x, y^0) \leqq F \max_{z \in S_u} \{M^1(z)\} < \infty.$$

Thus, by proving (7), we have due to (8) and (9) that a constant $K>0$ exists such that $M^*((i, j), y^0) \leqq K i^{m+1}$ for all $i \geqq 0$ and $j=0, 1$.

We prove (7) first for all $i>u$ and $j=1$. Note that when starting the process in state $(i, 1)$, the set S_u is entered in state $(u, 1)$ after sequential visits to states $(i-1, 1)$, $(i-2, 1), \cdots, (u, 1)$. We conclude, in view of Lemma 3.1, that a constant $K>0$ exists such that

$$(10) \quad M^*((i, j), y^0) = \sum_{l=u}^{i-1} C_l \leqq K \sum_{l=u}^{i-1} l^m \leqq K \int_u^i x^m \, dx \leqq \frac{K}{(m+1)} i^{m+1}, \qquad i>u.$$

For $x=(i, 0) \in S$, we have

$$(11) \quad \begin{aligned} M^*((i, 0), y^0) &= c((i, 0); 2) + \sum_{k=(u+1-i)^+}^{\infty} \int_0^{\infty} \frac{e^{-\lambda t}(\lambda t)^k}{k!} M^*((i+k, 1), y^0) \, dG(t) \\ &\leqq c((i, 0); 2) + \frac{K}{(m+1)} \sum_{k=0}^{\infty} \int_0^{\infty} \frac{e^{-\lambda t}(\lambda t)^k}{k!} (i+k)^{m+1} \, dG(t). \end{aligned}$$

We conclude that $M^*((i, 0), y^0) = O(i^{m+1})$ since $c((i, 0); 2) = O(i^m)$ (see (2)) while the second term to the right of (11) is $O(i^{m+1})$—the verification of which is analogous to that of (2).

It remains to be shown that $M^*((i, 1), y^0) < \infty$ for $i \leqq u$; but this follows from $c((i, 0); 1) < \infty$ and a verification which is analogous to that of (11).

We are now ready to prove the existence of a solution to the average cost optimality equation as well as the existence of a stationary optimal policy.

Theorem 3.3.

(a) There exist a constant g^* and a non-negative function $\{v^*(x): x \in S\}$ with the property $v^*((i, j)) = O(i^{m+1})$ for all $i \geqq 0$, $j = 0, 1$, which satisfy the optimality equation

$$(12) \qquad v(x) = \min_{a=1,2} \left\{ c(x; a) - g t(x; a) + \sum_{y \in S} p_{xy}(a) v(y) \right\}, \qquad x \in S.$$

(b) Assume the constant $\tilde{g}$ and the non-negative function $\{\tilde{v}(x): x \in S\}$, with $\tilde{v}((i, j)) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 0, 1$, solve the optimality equation (12). Let π be a stationary policy which in each state $x \in S$ prescribes an action which achieves the minimum in (12). Then π is an optimal policy and $\tilde{g}$ is the minimum long-run average cost.

The proof of Theorem 3.3 is based on the following lemma. Let δ_{ij} denote the Kronecker delta, i.e., $\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

Lemma 3.4 Consider the discrete-time Markov decision model with the same state and action spaces as our original semi-Markov decision model but with one-step expected costs and one-step transition probabilities given by $\hat{c}(x; a) \equiv c(x; a) t_{\min} / t(x; a)$ and $\hat{p}_{xy}(a) \equiv \delta_{xy} + [t_{\min}/t(x; a)][p_{xy}(a) - \delta_{xy}]$, respectively. A constant g^* and a non-negative function $\{v^*(x): x \in S\}$ solve (12) if and only if $g^* t_{\min}$ and the function $\{v^*(x): x \in S\}$ solve the optimality equation for the discrete-time Markov decision model:

$$(13) \qquad v(x) = \min_{a=1,2} \left\{ \hat{c}(x; a) - g + \sum_{y \in S} \hat{p}_{xy}(a) v(y) \right\}, \qquad x \in S.$$

Proof. See Schweitzer (1971); note that $\hat{p}_{xy}(a) \geqq 0$ and $\sum_{y \in S} \hat{p}_{xy}(a) = 1$ for all $x, y \in S$.

Proof of Theorem 3.3.

(a) We make the following observations with respect to the discrete-time Markov decision model constructed in Lemma 3.4:

(i) Its state space S is countable and all action sets are finite.

(ii) All one-step expected costs are non-negative.

(iii) All states in S communicate under some (i.e. the no-wait) policy.

In addition, let $\hat{\gamma}^0$ be the long-run average cost of the 0-policy (i.e. the no-wait policy) in the discrete-time model. It is again elementary to verify that $\hat{\gamma}^0 = t_{\min}\gamma^0$ with γ^0 the long-run average cost of the no-wait policy in the original model. Thus $\hat{\gamma}^0 < \infty$, see Lemma 3.1(b). Let $u = \min\{i : H(i) > \hat{\gamma}^0, i \geqq 1\} < \infty$ in view of $\lim_{i\to\infty} H(i) = \infty$, see Condition 1. Finally let $\hat{M}(x, y)$ denote the minimum expected total costs incurred in the discrete-time model when starting in state x until the first visit to state y.

(iv) $\max_{y \in S_u} \hat{M}((i, j), y) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 1, 2$.

Observe that (iv) follows from Lemma 3.2 and the fact that for all $x \neq y \in S$ and $a = 1, 2$,

$$\frac{t_{\min}}{t_{\max}} c(x; a) \leqq \hat{c}(x; a) \leqq c(x; a)$$

and

$$\frac{t_{\min}}{t_{\max}} p_{xy}(a) \leqq \hat{p}_{xy}(a) \leqq p_{xy}(a), \qquad x \neq y.$$

Let $V_\alpha(x)$ denote the minimum expected total discounted cost over the infinite horizon when starting in state $x \in S$ and discounting at rate $\alpha > 0$. Let x_α denote the most favourable starting state, i.e. $V_\alpha(x_\alpha) \leqq V_\alpha(x)$ for all $x \in S$. It follows from observation (ii) and Proposition 8 on p. 253 of Bertsekas (1976) that

$$V_\alpha(x) = \min_{a=1,2} \left\{ \hat{c}(x; a) + (1-\alpha) \sum_{y \in S} \hat{p}_{xy}(a) V_\alpha(y) \right\}, \qquad x \in S.$$

Subtracting $V_\alpha(x_\alpha)$ from both sides, we obtain for all $x \in S$,

$$\begin{aligned} V_\alpha(x) - V_\alpha(x_\alpha) &= -\alpha V_\alpha(x_\alpha) \\ &\quad + \min_{a=1,2} \left\{ \hat{c}(x; a) + (1-\alpha) \sum_{y \in S} \hat{p}_{xy}(a)[V_\alpha(y) - V_\alpha(x_\alpha)] \right\}. \end{aligned} \tag{14}$$

Following the proof on pp. 213–215 of Weber and Stidham (1987), one verifies that a sequence $\{\alpha_k\}_{k=1}^{\infty} \downarrow 0$ exists such that

$$\begin{aligned} v^*((i, j)) &\equiv \lim_{k\to\infty} \{V_{\alpha_k}((i, j)) - V_{\alpha_k}(x_{\alpha_k})\} = O(i^{m+1}) \quad \text{for all} \quad i \geqq 0 \quad \text{and} \quad j = 0, 1, \\ g^* &\equiv \lim_{k\to\infty} \alpha_k V_{\alpha_k}(x_{\alpha_k}). \end{aligned} \tag{15}$$

Taking limits on both sides of (14) we obtain for all $x \in S$:

$$\begin{aligned} v^*(x) &= -g^* + \lim_{a=1,2} \left\{ \hat{c}(x; a) + \lim_{k\to\infty} \sum_{y \in S} \hat{p}_{xy}(a)[V_{\alpha_k}(y) - V_{\alpha_k}(x_{\alpha_k})] \right\} \\ &= \min_{a=1,2} \left\{ \hat{c}(x; a) - g^* + \sum_{y \in S} \hat{p}_{xy}(a) v^*(y) \right\} \end{aligned}$$

where the last equality follows from Lebesgue's dominated convergence theorem,

(15) and the observation that $\sum_{i=0}^{\infty} \hat{p}_{x(i,j)}(a) i^{m+1} < \infty$ for $j = 0, 1$. (The latter may again be verified as with (2) using the finiteness of the first $(m + 1)$ moments of the repair time distribution.) Now (a) follows from Lemma 3.4.

(b) follows from the proof on p. 215 in Weber and Stidham (1987).

Remark 3.1. An alternative proof may be obtained by relatively minor adaptations of the analysis in Federgruen et al. (1979). This approach allows for a direct verification of the existence of a solution to the optimality equation (12) in the semi-Markov decision model (rather than via the equivalent discrete-time model).

In the next section we use the following corollary with respect to the n-policies ($n \geqq 1$). For any stationary policy f, let $c(x;f)$, $t(x;f)$ and $p_{xy}(f)$ denote respectively, the one-step expected cost in state x, the one-step expected holding time in state x and the one-step transition probability from state x to state y, under policy $f (x, y \in S)$.

Corollary 3.5. Let f be an n-policy ($n \geqq 0$),

(a) There exist a constant g_n and a non-negative function $\{v_n(x); x \in S\}$ with $v_n((i, j)) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 0, 1$, which satisfy the system of equations:

$$v(x) = c(x;f) - gt(x;f) + \sum_{y \in S} p_{xy}(f) v(y), \qquad x \in S. \tag{16}$$

(b) If $\{g, v(.)\}$ solve (16) with $v((i, j)) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 0, 1$, then g represents the long-run average cost under policy f; moreover, the function $v(.)$ is unique up to an additive constant.

(c) If f is optimal and $\{g, v(.)\}$ solve (16) with $v((i, j)) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 0, 1$, then $\{g, v(.)\}$ solve the optimality equation (12).

Proof. (a) Repeat the proof of Theorem 3.3, restricting the action set for $x = (i, 0)$ as follows: if $i < n$ only action $a = 1$ is allowed; if $i \geqq n$ only action $a = 2$ is allowed. Redefine $M_f(x, y)$ as the expected total cost incurred under policy f when starting in state x up until the first visit to state y ($x, y \in S$). A minor adaptation of Lemma 3.2 establishes for any fixed $y \in S$ that $M_f((i, j), y) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 0, 1$.

(b) The fact that g equals the long-run average cost under policy f follows again as in the proof of Theorem 3.3. Note that all states in S belong to a single ergodic set under policy f. Thus let $\pi(.)$ be the unique steady-state distribution of the (embedded) Markov chain under policy f. Assume $\{g, v^1(.)\}$ and $\{g, v^2(.)\}$ are two alternative solutions to (16) with $v^l((i, j)) = O(i^{m+1})$ for all $i \geqq 0$ and $j = 0, 1$ ($l = 1, 2$). One easily verifies, as for ordinary $M/G/1$ queueing systems that the steady-state probabilities $\pi((i, j))$ decline to 0 at an asymptotically exponential rate, i.e., there exists a number η, with $0 \leqq \eta < 1$, such that $\pi((i, j)) = O(\eta^i)$ for all $i \geqq 0$ and $j = 0, 1$; see e.g. Neuts (1981b). Thus $\sum_{y \in S} \pi(y) v^l(y) < \infty$ for $l = 1, 2$. Subtracting

Equation (16) with $v = v^2$ from Equation (16) with $v = v^1$ we obtain for all $x \in S$

$$v^1(x) - v^2(x) = \sum_{y \in S} p_{xy}(f)[v^1(y) - v^2(y)].$$

By a standard argument, one concludes that for all $x \in S$

$$v^1(x) - v^2(x) = \sum_{y \in S} \pi(y)v^1(y) - \sum_{y \in S} \pi(y)v^2(y)$$

where the right side is a constant, independent of x.

(c) Assume to the contrary that for some $x^0 \in S$,

$$c(x^0; a) - gt(x^0; a) + \sum_{y \in S} p_{x^0y}(a)v(y) < c(x^0; f) - gt(x^0; f) + \sum_{y \in S} p_{x^0y}(f)v(y),$$

for $a \neq f(x^0)$, the action prescribed by policy f. Let policy $\bar{f}$ be defined as the policy which prescribes action a in state x^0 but is otherwise identical to policy f. Thus, since $\{g, v(.)\}$ solve (16),

$$c(x; \bar{f}) - gt(x; \bar{f}) + \sum_{y \in S} p_{xy}(\bar{f})v(y) \leqq v(x), \tag{17}$$

with strict inequality for $x = x^0$.

As in part (b), one easily verifies that under policy $\bar{f}$, all states in S are positive recurrent with the unique steady-state distribution $\bar{\pi}$ satisfying $\bar{\pi}((i, j)) = O(\eta^i)$ for all $i \geqq 0$, $j = 0, 1$, and some η with $0 < \eta < 1$. Thus, multiply both sides of (17) by $\bar{\pi}(x)$ and add over $x \in S$ to conclude that the long-run average cost under policy f: $\sum_{x \in S} \bar{\pi}(x)c(x; \bar{f}) / \sum_{x \in S} \bar{\pi}(x)t(x; \bar{f}) < g$. This contradicts the optimality of policy f.

Thus, for any n-policy, let $\{g_n, V_n(.)\}$ be the unique solution to (16) with $V_n((0, 1)) = 0$ ($n \geqq 0$, see Corollary 3.5(b)). We refer to $V_n(.)$ as the relative cost function of the n-policy. The characterization of the structure of an optimal policy in the next section is achieved by varying the parameter r_0 over the entire real line. Under the n-policy, let $\pi_n(0, 1)$ denote the steady-state probability of the system being empty and the server up. We clearly have the following result.

Lemma 3.6. (a) $g_n = \pi_n(0, 1)r_0 + \bar{g}_n$ where $\bar{g}_n$ is independent of r_0.

(b) $\pi_n(0, 1)$ is decreasing in n.

Proof. Part (a) is immediate. Note that the first passage times from states (0, 1) and (0, 0) to themselves are strictly increasing in n. This verifies part (b).

4. Characterization of an optimal policy

In this section, we prove that an n-policy is optimal, provided that the holding cost rate function $H(.)$ is convex. We thus add the following Condition 4.

Condition 4. $H(.)$ is convex.

For any stationary policy f and any pair of states $x, y \in S$, let $T_f(x, y)$ denote the expected first passage time from state x to state y under policy f, and let $M_f(x, y)$ denote the expected total cost when starting in state x until reaching state y. For n-policies, we write $T_n(.)$ and $M_n(.)$ instead of $T_f(.)$ and $M_f(.)$.

We first prove a more limited result: there exists an optimal (stationary) policy which prescribes immediate repairs whenever the number of customers in the system is sufficiently large.

Lemma 4.1. Let $\{g^*, v^*\}$ be a solution of optimality equation (12) with $v^* \geqq 0$ and $v^*((i, j)) = O(i^{m+1})$ for all $i \geqq 1$, $j = 0, 1$. Let $i^* = \min \{i : H(i) > g^*\}$. For $i \geqq i^*$, only action 2 achieves the minimum in (12). In particular, there exists an optimal (stationary) policy f which prescribes immediate repairs (action 2) in all states $(i, 0)$ with $i \geqq i^*$.

Proof. Let f be a stationary policy which, in each state $x \in S$, prescribes an action which achieves the minimum in (12). It follows from Theorem 3.3(b) that f is optimal. Assume to the contrary that policy f prescribes the wait action (action 1) for some state $(i', 0)$ with $i' \geqq i^*$. Then there exists an integer $i \geqq i^*$ such that action 1 is prescribed in state $(i, 0)$ and action 2 is prescribed in state $(i + 1, 0)$. (Otherwise, no repairs would be initiated when the number of customers in the system exceed i^* so that $\lim_{t \to \infty} \Pr \{(X(t), Y(t)) = (i, 0) \text{ with } i \geqq N\} = 1$ for all $N \geqq 1$, i.e., the policy would have infinite average cost.) Thus,

$$
\begin{aligned}
v^*((i, 0)) &= H(i)\lambda^{-1} - g^*\lambda^{-1} + v^*((i + 1, 0)) \\
&= H(i)\lambda^{-1} - g^*\lambda^{-1} + c((i + 1, 0); 2) - g^*\nu + \sum_{k=0}^{\infty} p_k v^*((i + k + 1, 1)) \qquad (18) \\
&\leqq c((i, 0); 2) - g^*\nu + \sum_{k=0}^{\infty} p_k v^*((i + k, 1))
\end{aligned}
$$

where $p_k = \int_0^\infty e^{-\lambda t}(\lambda t)^k / k! \, dG(t)$ is the probability that k customers arrive during a repair time ($k \geqq 0$). Note that the number of customers in the system cannot drop from $(i + k + 1)$ to any number less than $(i + k)$, without the system passing through state $(i + k, 1)$. Thus,

$$
\begin{aligned}
0 &\geqq [H(i) - g^*]\lambda^{-1} + [c((i + 1, 0); 2) - c((i, 0); 2)] \\
&\quad + \sum_{k=0}^{\infty} p_k[v^*((i + k + 1, 1)) - v^*((i + k, 1))] \\
&\geqq [H(i) - g^*]\lambda^{-1} + [c((i + 1, 0); 2) - c((i, 0); 2)] \\
&\quad + \sum_{k=0}^{\infty} p_k[H(i + k + 1) - g^*]T_f((i + k + 1, 1), (i + k, 1)).
\end{aligned}
$$

Each of the terms to the right of this inequality is however, non-negative while the first one is strictly positive, in view of $H(.)$ being non-decreasing, the definition of i^* and (2). This leads to a contradiction.

Theorem 4.2 (main result). There exists an increasing sequence $\{R_n\}_{n=0}^{\infty}$ with $R_0=-\infty$, $R_1>-\infty$, $R_n<\infty$ and $R_n\to\infty$ as $n\to\infty$ such that the n-policy is optimal for $r_0\in[R_n, R_{n+1}]$.

Proof. The proof is by induction. We first establish that a number $R_1>-\infty$ exists such that the 0-policy is optimal for all $r_0\leqq R_1$.

In view of Theorem 3.3 and Corollary 3.5, it suffices to show that V_0 satisfies optimality equation (12) for all r_0 sufficiently small. Fix an initial value r_0^0 for r_0, let $i^*(r_0)\equiv\min\{i:H(i)>g_0(r_0)\}$ and g_0^0 the long-run average cost under the 0-policy with $r_0=r_0^0$. For all $i\geqq i^0\equiv i^*(r_0^0)$, one verifies as in the proof of Lemma 4.1 that (12) holds for $v=V_0$, $x=(i,0)$ with $i\geqq i^0$ and all $r_0\leqq r_0^0$. (In (18), replace v^* by V_0 and g^* by g_0^0; note that $i^*(r_0)\leqq i^*(r_0^0)$ for all $r_0\leqq r_0^0$ since g_0 is increasing in r_0.)

It thus suffices to verify that for all r_0 sufficiently small and $i<i^0$, $\{V_0, g_0\}$ satisfy the optimality equation (12) as well, or in view of Corollary 3.5 that

$$
\begin{aligned}
&c((i,0);2)-g_0\nu+\sum_{k=0}^{\infty}p_kV_0(i+k,1)\\
(19)\qquad &\leqq H(i)\lambda^{-1}-g_0\lambda^{-1}+c((i+1,0);2)-g_0\nu+\sum_{k=0}^{\infty}p_kV_0((i+1+k,1)).
\end{aligned}
$$

(19) is equivalent to

$$
\begin{aligned}
g_0\lambda^{-1}\leqq{}& H(i)\lambda^{-1}+c((i+1,0);2)-c((i,0);2)\\
&+\sum_{k=0}^{\infty}p_k\{V_0((i+1+k,1))-V_0((i+k,1))\},
\end{aligned}
$$

and hence to

$$
\begin{aligned}
&g_0\left\{\lambda^{-1}+\sum_{k=0}^{\infty}p_kT_0((i+1+k,1),(i+k,1))\right\}\\
&\qquad\leqq H(i)\lambda^{-1}+c((i+1,0);2)-c((i,0);2)+\sum_{k=0}^{\infty}p_kM_0((i+1+k,1),(i+k,1)).
\end{aligned}
$$

The latter inequality holds for all $i\leqq i^0$ and r_0 sufficiently small since $g_0\to-\infty$ as $r_0\to-\infty$ and the right-hand side is independent of r_0.

Now assume that numbers $\{R_1,\cdots,R_{n+1}\}$ exist such that the l-policy is optimal for all $r_0\in[R_l, R_{l+1}]$ $(l=0,1,\ldots,n)$ with $R_{n+1}=\sup\{w:w\geqq R_n$ and the n-policy is optimal for all $R_n\leqq r_0\leqq w\}$. We now show that the $(n+1)$-policy is optimal for $r_0=R_{n+1}$. This establishes the existence of a number $R_{n+2}\geqq R_{n+1}$ such that the $(n+1)$-policy is optimal for all $r_0\in[R_{n+1}, R_{n+2}]$. (Choose $R_{n+2}=\sup\{w:w\geqq R_{n+1}$ and the $(n+1)$-policy is optimal for all $R_{n+1}\leqq r_0\leqq w\}$.)

Consider some $\varepsilon>0$ such that the n-policy is not optimal for $r_0=R_{n+1}+\varepsilon$. Let $V_n^*(.)$ be the relative cost function and g_n^* the average cost of the n-policy with $r_0=R_{n+1}+\varepsilon$. In view of Theorem 3.3, $V_n^*(.)$ fails to satisfy the optimality equation for $r_0=R_{n+1}+\varepsilon$. Thus, some non-negative integer i exists such that one of the following two cases holds:

Case (i). For some i with $0 \leqq i \leqq n-1$:

$$c((i,0);2) - g_n^* v + \sum_{k=0}^{\infty} p_k V_n^*((i+k,1)) < V_n^*((i,0)) = \frac{H(i)}{\lambda} - \frac{g_n^*}{\lambda} + V_n^*((i+1,0))$$

$$= \frac{H(i)}{\lambda} - \frac{g_n^*}{\lambda} + \frac{H(i+1)}{\lambda} - \frac{g_n^*}{\lambda} + V_n^*((i+2,0))$$

$$\vdots$$

$$= \frac{H(i) + \cdots + H(n-1)}{\lambda} - \frac{(n-i)g_n^*}{\lambda} + c((n,0);2) - g_n^* v + \sum_{k=0}^{\infty} p_k V_n^*((n+k,1)).$$

Rearranging terms on both sides gives

$$\frac{H(i) + \cdots + H(n-1)}{\lambda} + c((n,0);2) - c((i,0);2)$$
$$+ \sum_{k=0}^{\infty} p_k \{V_n^*((n+k,1)) - V_n^*((i+k,1))\} - \frac{(n-i)g_n^*}{\lambda} > 0,$$

or equivalently, $\psi_i(R_{n+1} + \varepsilon) > 0$ where

$$\begin{aligned} \psi_i(r_0) = {} & \lambda^{-1}\{H(i) + \cdots + H(n-1)\} + \{c((n,0);2) - c((i,0);2)\} \\ & + \sum_{k=0}^{\infty} p_k M_n((n+k,1),(i+k,1)) \\ & - g_n \left\{ (n-i)\lambda^{-1} + \sum_{k=0}^{\infty} p_k T_n((n+k,1),(i+k,1)) \right\}. \end{aligned} \tag{20}$$

Notice that all quantities on the right side of (20) are independent of r_0 except for g_n. Since g_n is increasing in r_0, $\psi_i(r_0)$ is decreasing in r_0. By the optimality of the n-policy at $r_0 = R_{n+1}$ and Corollary 3.5(c), we have $0 < \psi_i(R_{n+1} + \varepsilon) < \psi_i(R_{n+1}) \leqq 0$, a contradiction. Thus, Case (i) cannot occur and Case (ii) below must arise.

Case (ii). For some i with $i \geqq n$:

$$c((i,0);2) - g_n^* v + \sum_{k=0}^{\infty} p_k V_n^*((i+k,1)) > H(i)\lambda^{-1} - g_n^*\lambda^{-1} + V_n^*((i+1,0))$$
$$= H(i)\lambda^{-1} - g_n^*\lambda^{-1} + c((i+1,0);2) - g_n^* v + \sum_{k=0}^{\infty} p_k V_n^*((i+1+k,1)),$$

or equivalently $\psi_i(R_{n+1} + \varepsilon) > 0$ where

$$\begin{aligned} \psi_i(r_0) = {} & \{c((i,0);2) - c((i+1,0);2)\} - H(i)\lambda^{-1} \\ & - \sum_{k=0}^{\infty} p_k M_n((i+1+k,1),(i+k,1)) \\ & + g_n \left\{ \lambda^{-1} + \sum_{k=0}^{\infty} p_k T_n((i+1+k,1),(i+k,1)) \right\}. \end{aligned} \tag{21}$$

It follows from (4) and the convexity of $H(.)$ that $c((i, 0); 2)$ is convex in i so that the first three terms to the right of (21) are non-increasing. Observe that $T_n((i+1+k, 1), (i+k, 1))$ represents the expected length of a busy period for an $M/G/1$ queue with arrival rate λ and mean service time $\tilde{F}(\sigma)^{-1}s + (\tilde{F}(\sigma)^{-1} - 1)v$. The last term to the right of (21) is thus independent of i and we conclude that $\psi_i(r_0)$ is non-increasing in i. It follows that $\psi_n(R_{n+1} + \varepsilon) \geqq \psi_i(R_{n+1} + \varepsilon) > 0$. It also follows from (21) that $\psi_n(r_0)$ is continuous in r_0. By the definition of R_{n+1} and the above analysis, there exists a sequence $\{\varepsilon_k\} \downarrow 0$ with $\psi_n(R_{n+1} + \varepsilon_k) > 0$. We conclude that $\psi_n(R_{n+1}) = 0$ since the n-policy is optimal for $r_0 = R_{n+1}$. (Use Corollary 3.5(c).)

Fix now $r_0 = R_{n+1}$. It follows from Corollary 3.5(c) that $\{g_n, V_n(.)\}$ solve the optimality equation (12) for this value of r_0. Since $\psi_n(R_{n+1}) = 0$, it follows from Theorem 3.3(b) that the $(n+1)$-policy is optimal for $r_0 = R_{n+1}$.

To show that $R_n < \infty$ for all $n \geqq 1$, note that $g_n = \pi_n(0, 1)r_0 + \bar{g}_n > \pi_{n+1}(0, 1)r_0 + \bar{g}_{n+1} = g_{n+1}$ for r_0 sufficiently large in view of parts (a) and (b) of Lemma 3.6 ($\bar{g}_n$ is independent of r_0, $n \geqq 1$). It remains to be shown that the above constructed sequence $\{R_n\} \uparrow \infty$. Assume to the contrary that $\lim_{n\to\infty} R_n = R^* < \infty$. As in Lemma 4.1, let $i^*(r_0) = \min\{i : H(i) > g^*(r_0)\lambda^{-1}\}$ and choose $n^* > i^*(R^*)$. Let $\{g^*, v^*\}$ be a solution of optimality equation (12) for $r_0 = R_{n^*}$ with $v^* \geqq 0$ and $v^*((i, j)) = O(i^{m+1})$. It follows from Lemma 4.1 that *only* action 2 achieves the minimum in (12) for some state $(i, 0)$ with $n^* > i \geqq i^*(R^*) \geqq i^*(R_{n^*})$. This contradicts the optimality of the n^*-policy for $r_0 = R_{n^*}$.

5. Computing optimal n-policies

In this section we describe an efficient algorithm for the determination of an optimal n-policy. We first need the following two lemmas.

Lemma 5.1. The optimal average cost g^* is piecewise linear, strictly increasing and concave in r_0.

Proof. From Theorem 4.2, it follows immediately that g^* is piecewise linear and strictly increasing in r_0. Since g_n is linear (and thus concave) in r_0 and $g^* = \inf_n \{g_n\}$, g^* is concave in r_0.

Lemma 5.2. For any fixed r_0, the function g_n is unimodal as a function of n.

Proof. We need to show that for any fixed r_0 and n,

(i) if $g_n \leqq g_{n+1}$, then $g_n \leqq g_{n+k}$ for all $k \geqq 2$ and

(ii) if $g_n \leqq g_{n-1}$, then $g_n \leqq g_{n-k}$ for all $k \geqq 2$.

To prove (i), recall from Lemma 3.6 that $g_n = \pi_n(0, 1)r_0 + \bar{g}_n$ with $\bar{g}_n$ independent of r_0 and $\pi_n(0, 1)$ decreasing in n. Thus, if for some fixed $r_0 = \bar{r}$ and n, $g_n \leqq g_{n+1}$, then $g_n < g_{n+1}$ for all $r_0 < \bar{r}$. Thus, to prove (i), assume to the contrary that $g_n > g_{n+k}$ and thus $g_{n+1} > g_{n+k}$. It follows again from the above observations that $g_{n+1} > g_{n+k}$ for all $r_0 \geqq \bar{r}$. Thus, the $(n+1)$-policy is not optimal for *any* value of r_0, contradicting Theorem 4.2.

The proof of (ii) is similar.

In view of Lemma 5.2, an optimal n-policy may be obtained by a simple bisection procedure. Let $N_1 < N_2$ be integers such that $g_{N_1+1} < g_{N_1}$ and $g_{N_2-1} \leqq g_{N_2}$.

Bisection procedure.

Step 0: Initialize $n_1 := N_1$; $n_2 := N_2$.

Step 1. Let $n := \lceil (n_1 + n_2)/2 \rceil$. If $g_n < g_{n+1}$, set $n_2 := n$; if $g_n > g_{n+1}$, set $n_1 := n$; if $g_n = g_{n+1}$ and $g_{n-1} \leqq g_n$, set $n_2 := n$; if $g_n = g_{n+1}$ and $g_{n-1} > g_n$, set $n_1 := n$. If $n_2 - n_1 > 1$, repeat Step 1.

The bisection procedure clearly requires no more than $O(\log_2 (N_2/N_1))$ evaluations of $\{g_n : n \geqq 1\}$. We now describe an efficient recursive procedure for the evaluation of the steady-state distribution $\{\pi_n(.)\}$ of the number of customers in system under the n-policy. This recursive scheme is analogous to that of standard $M/G/1$ queues, see e.g. Tijms (1986). With the help of the latter, the average cost g_n of the n-policy is easily computed as well, see below.

Let T be the length of a busy cycle, and for any $i \geqq 1$, let T_i be the total amount of time during a busy cycle with i customers present. Our derivation of a recursive scheme is based on the observation that a busy cycle may be divided into a random number of disjoint, so-called *service completion intervals* separated by the service completion epochs and that $E(T_i)$ may thus be calculated as the sum of the contributions to $E(T_i)$ of these intervals. Thus, define the quantities

A_{ik} = the expected amount of time that k customers are present during a service completion interval that is started with i customers present ($k \geqq i$).

We also need

B_{jk} = the expected amount of time that k customers are present during an interval which starts with j customers present and the initiation of a repair, and which terminates with the first service completion ($k \geqq j \geqq n$).

Similarly to the standard $M/G/1$ queue we obtain

$$(22)\quad E(T_k) = \frac{\lambda}{\sigma_0 + \lambda} A_{1k} + \frac{\sigma_0}{\sigma_0 + \lambda} \{1_{\{k<n\}} \lambda^{-1} + 1_{\{k \geqq n\}} B_{nk}\} + \sum_{j=1}^{k} \lambda E(T_j) A_{jk}, \qquad k \geqq 1,$$

see e.g. Tijms (1986). (The first two terms represent the contribution to $E(T_k)$ which occurs during the *first* interval of the busy cycle prior to the *first* service completion epoch; the first term represents the contribution when the first customer in the busy cycle arrives prior to a failure and the second term corresponds with the alternative case.)

By the theory of regenerative processes, we have $\pi_n(k) = E(T_k)/E(T)$, $k \geqq 1$ and

$\pi_n(0) = [\lambda E(T)]^{-1}$. We conclude that

$$\pi_n(k) = \lambda\pi_n(0)\left[\frac{\lambda}{\sigma_0+\lambda}A_{1k} + \frac{\sigma_0}{\sigma_0+\lambda}\{1_{\{k<n\}}\lambda^{-1} + 1_{\{k\geqq n\}}B_k\}\right] + \sum_{j=1}^{k}\lambda\pi_n(j)A_{jk}, \qquad k \geqq 1. \tag{23}$$

We now derive expressions for the quantities $\{A_{jk}\}$. Let $\tilde{G}(.)$ denote the cumulative distribution function of a customer's completion time Γ, see (5). For $j \geqq n$, the expression for A_{jk} and its derivation are standard: assuming that at epoch 0 a new service starts while j customers are present, define $\chi_t = 1$ if there are k customers in the system and this service completion interval is still in progress at time t; otherwise, define $\chi_t = 0$. Thus, $A_{jk} = E(\int_0^\infty \chi_t\,dt)$, $j \geqq n$. Since $E(\chi_t) = \Pr(\chi_t = 1) = (1-\tilde{G}(t))e^{-\lambda t}(\lambda t)^{k-j}/(k-j)!$ using the fact that the number of arrivals in $[0, t]$ is Poisson distributed with mean λt, we find

$$A_{jk} = \int_0^\infty (1-\tilde{G}(t))\frac{e^{-\lambda t}(\lambda t)^{k-j}}{(k-j)!}\,dt, \qquad k \geqq j, \qquad j \geqq n. \tag{24}$$

The integral in (24) may be computed in closed form when $\tilde{G}(.)$ is of phase type which is the case when the repair time and service time distributions are of phase type, see Tijms (1986) and Neuts (1981b) and the Appendix.

Now, consider the case where $j < n$. For a service completion interval that starts with j customers present, let

A'_{jk} = the expected amount of time that k customers are present during the *first* part of the service process which precedes the *first* failure.

(If no failure occurs during the service process, $A'_{jk} = A_{jk}$.) One verifies as above (see (24)) that

$$\begin{aligned} A'_{jk} &= \int_0^\infty (1-\hat{F}(t))\frac{e^{-\lambda t}(\lambda t)^{k-j}}{(k-j)!}\,dt \\ &= \int_0^\infty (1-F(t))\frac{e^{-(\lambda+\sigma)t}(\lambda t)^{k-j}}{(k-j)!}\,dt, \qquad k \geqq j \quad \text{and} \quad j < n. \end{aligned}$$

To characterize $(A_{jk} - A'_{jk})$, condition upon the event that t time units after the beginning of the service completion interval a breakdown occurs while the service is still in process and that l customers are present. Clearly,

$$\begin{aligned} A_{jk} &= A'_{jk} + \sum_{l=j}^{k}\left\{\int_0^\infty \sigma e^{-\sigma t}(1-F(t))\frac{e^{-\lambda t}(\lambda t)^{l-j}}{(l-j)!}\,dt\right\}\{1_{\{k<n\}}\lambda^{-1} + 1_{\{k\geqq n\}}B_{\max\{l,n\},k}\} \\ &= A'_{jk} + \sigma\sum_{l=j}^{k}A'_{jl}\{1_{\{k<n\}}\lambda^{-1} + 1_{\{k\geqq n\}}B_{\max(l,n),k}\}, \qquad (k \geqq j). \end{aligned}$$

Finally, as in (24),

$$B_{jk} = \int_0^\infty (1 - \hat{G}(t)) \frac{e^{-\lambda t}(\lambda t)^{k-j}}{(k-j)!} dt, \qquad (k \geqq j \geqq n) \tag{25}$$

where $\hat{G}(.)$ is the cumulative distribution function of the convolution of Γ with an additional repair time. Once again, $\hat{G}(.)$ is of phase type, provided $G(.)$ and $F(.)$ are and the integrals in (25) may be computed in closed form.

The steady-state distribution $\{\pi_n(.)\}$ is then determined via the recursive scheme (23): guess an initial value for $\pi_n(0)$, evaluate $\pi_n(j)$ for consecutive values of j until the series $\sum_j \pi_n(j)$ appears to have converged and rescale the steady-state probabilities by the obtained estimate for $\sum_j \pi_n(j)$. The long-run average cost of the n-policy is now easily determined:

$$\begin{aligned} g_n = &\sum_{j=0}^{\infty} \pi_n(j)H(j) + r\lambda\bar{F}(\sigma)^{-1}s + \frac{r_0}{\lambda+\sigma_0}\lambda\pi_n(0) \\ &\times \left\{1_{\{n>0\}} + 1_{\{n=0\}}\left[1 - \frac{\sigma_0}{\lambda+\sigma_0}\int_0^\infty e^{-\lambda t}\, dG(t)\right]^{-1}\right\} \\ &+ c\lambda[\bar{F}(\sigma)^{-1} - 1] + c\lambda\pi_n(0) \\ &\times \left\{1_{\{n>0\}}\frac{\sigma_0}{(\lambda+\sigma_0)} + 1_{\{n=0\}}\left[\frac{\lambda}{\sigma_0} + 1 - \int_0^\infty e^{-\lambda t}\, dG(t)\right]^{-1}\right\}. \end{aligned} \tag{26}$$

The first term represents the average holding costs, the next two terms the average running costs and the last two terms the average repair costs. To verify the repair costs, attribute any repair that interrupts a service to the customer whose service is interrupted, and a repair that starts while the system is empty to the customer who initiates the next busy period. The average number of repairs of the first type is thus given by $\lambda[\bar{F}(\sigma)^{-1} - 1]$ since on average λ customers arrive per unit of time and each experiences $[\bar{F}(\sigma)^{-1} - 1]$ breakdowns on average. (Apply Wald's lemma.) This verifies the fourth term in (26). The last term denotes the average number of repairs of the second type. The average arrival rate to an empty system is given by $\lambda\pi_n(0)$. Let b denote the average number of breakdowns that occur in a busy cycle prior to the arrival of the first customer. If $n > 0$, $b = \sigma_0/(\lambda + \sigma_0)$; if $n = 0$ (the no-wait policy), b is obtained as the solution of the linear equation

$$b = \frac{\sigma_0}{(\lambda+\sigma_0)}\left[1 + b\int_0^\infty e^{-\lambda t}\, dG(t)\right].$$

The running costs can be verified similarly.

6. Remarks

We first give some remarks on the case where the repair time is negligible, i.e. $v = 0$. Such a case might correspond to replacing the broken server (machine) by an available new server (machine). In this case, $t_{\min} = 0$. However, all the results in Section 3 (and following) may be obtained by minor adaptations of the analysis in Section 3. In particular, we do not consider repair completion epochs as decision epochs. Instead, we define the one-step transition probability $\tilde{p}_{(i,0)x}(2) \equiv p_{(i,1)x}(1)$ and the one-step expected cost $\tilde{c}[(i, 0); 2] \equiv c + c[(i, 1); 1]$. Furthermore, one easily verifies that convexity of the customer holding cost rate $H(i)$ is not required to prove the optimality of monotone policies in Theorem 4.2.

When the service time distribution is exponential and the customer holding cost is linear, an alternative efficient algorithm based on difference equations may be found to compute an optimal monotone policy (see So (1985)).

As pointed out in Section 1, one extension to the wait/repair problem is to include the option of turning off the server even when it is operational by paying a shutdown cost. This would permit saving the running cost when the number of customers in the system is sufficiently small (perhaps zero) to justify this action. However, we require that when the server is down, either because of a breakdown or being turned off, we will pay the same cost and require the same operation to bring the server back to work again. This model can apply in the situation where the server *can* be switched to do other types of work (corresponding to a shutdown) or it *has* to be switched to process some higher priority jobs when these jobs arrive (corresponding to a breakdown), and a repair corresponds to bringing the server back to serve the lower priority jobs. We can use a similar approach to show that there exists an optimal (stationary) policy which either never turns the server off or turns the server off only when there is no customer in the system, and repairs the server if and only if the number of customers in the system exceeds a certain level for this problem.

To see this, consider any stationary policy f which turns the server off when there are n ($\geqq 1$) customers in the system. First observe that under policy f, there will be at least n customers in the system once the number of customers in the system exceeds n, i.e., the states (i, j) for all $i < n$, $j = 0, 1$ are transient states. Now consider another (stationary) policy f^* with $f^*(i, j) = f(i + n, j)$ for all $i \geqq 0$, $j = 0, 1$. Starting from any state (i_0, j_0) and state $(i_0 + n, j_0)$ under policies f^* and f, respectively, the two systems are identical except that the system under policy f^* has n customers less than that under policy f. Since the customer holding rate function is non-decreasing, the average cost under policy f^* will be less than that under policy f. Therefore, we only need to consider two classes of policies. The first class of policies corresponds to those which never turn off the server and the problem then reduces to the wait/repair problem. The second class of policies corresponds to those which turn off the server only when there is no customer in the system. For this class of policies, the average cost will be independent of r_0. However, we can use a similar parametric analysis on the shutdown cost to show that there exists an optimal policy which only turns the server off when there is no customer in the system and repairs the server if and only if the number of customers in the system exceeds a certain level. As pointed out in the introduction, these results generalize those obtained for $M/G/1$ systems without breakdowns and repairs.

The assumption that the server's up times are exponential (while the repair time and service time distributions are arbitrary) appears to be essential: for different up-time distributions, it does not appear to be possible to compute (exactly) even such aggregate performance measures as the expected number of customers in the system, even if the simplest of all policies, the no-wait policy is adopted; see Gaver (1962) and Federgruen and Green (1986), (1988).

The proof technique used in this paper can be used to produce a similar type of structural result in many other models as well. In Federgruen and So (1989) we apply the technique to obtain structural results for other control problems in queueing systems subject to breakdowns and in priority queueing systems. We now sketch the application of the technique to the following control problem in $M/M/1$ queues with two discrete service rates.

Consider an $M/M/1$ queue with batch arrivals with rate λ and i.i.d. batch sizes with mean x. Two service rates (parameters μ_1 and μ_2 with $\mu_1 < \mu_2$) are available. The cost structure includes a non-decreasing and unbounded customer holding cost rate and two different service cost rates for the two available service rates. The problem is to choose the service rate to minimize the average operating costs of the system.

The problem can be formulated as a Markov decision process. The decision epochs include service completions and customer arrivals. To use our proof technique, we introduce an idle running cost rate r_0 when the system is empty. Under the conditions that $\lambda x < \mu_2$ and that the customer holding cost rate is bounded above by a polynomial, we can use similar arguments to prove the analogous results in Lemma 4.1 and Theorem 4.2 in this problem. In particular, we can prove that an optimal monotone policy exists when $r_0 = 0$, i.e., a policy which uses the *faster* service rate μ_2 if and only if the number of customers in the system exceeds a certain level.

As shown, our proof technique appears to be very useful in proving structural results for control problems under the average cost criterion. It appears that the technique should be readily usable to obtain similar structural results under infinite- or finite-horizon discounted cost criteria. Note that our proof technique is primarily based on the existence of solutions to an optimality equation with a policy prescribing actions that achieve the minima in this equation being optimal. In most models where these results can be established, they can usually be established under the discounted cost criterion when they can be verified under the average cost criterion. (As in this paper, the proof for the latter is often based on that of the former.) We are currently investigating these conjectures.

Appendix

In this appendix we demonstrate that the distributions $\tilde{G}(.)$ and $\hat{G}(.)$ which are used in (24) and (25) are of phase-type, provided the service time and repair time distributions $F(.)$ and $G(.)$ are and that their parameters are easily obtained from those of $F(.)$ and $G(.)$.

A phase-type distribution represents the time until absorption in a transient finite state continuous-time Markov chain. The class of phase-type distributions is dense in the class of continuous distributions. Let

$$Q^s = \begin{vmatrix} U & U^0 \\ 0 & 0 \end{vmatrix}$$

be the infinitesimal generator of the Markov chain underlying the service-time distribution $F(.)$ where the $R^s \times R^s$ matrix U satisfies $U_{ii} < 0$ $(1 \leqq i \leqq R^s)$ and $U_{ik} \geqq 0$ $(i \neq k)$ and U^0 is a column vector with

$$U_i^0 = -\sum_{k=1}^{R^s} U_{ik}, \qquad i = 1, \cdots, R^s.$$

Let α denote the probability row-vector of the initial state of the chain $(\alpha \in R^1 \times R^s)$. Similarly, let

$$Q^r = \begin{vmatrix} W & W^0 \\ 0 & 0 \end{vmatrix}$$

be the infinitesimal generator of the Markov chain underlying the repair time distribution $G(.)$ where W is an $R^r \times R^r$ matrix and let β denote the probability row-vector $(\beta \in R^1 \times R^r)$.

One easily verifies that $\tilde{G}(.)$ is a phase-type distribution with an underlying chain with $(R^s + R^r + 1)$ states, $[\alpha, 0]$ as the initial state probability vector and $(R^s + R^r + 1) \times (R^s + R^r + 1)$ infinitesimal matrix generator

$$\begin{vmatrix} U - \sigma I & \sigma 1 \beta & U^0 \\ W^0 \alpha & W & 0 \\ 0 & 0 & 0 \end{vmatrix}.$$

Similarly $\hat{G}(.)$ is a phase-type distribution with the same infinitesimal matrix generator but initial state probability vector $[0, \beta]$.

References

[1] BALACHANDRAN, K. R. (1973) Control policies for a single server system. *Management Sci.* **19,** 1013–1018.

[2] BALACHANDRAN, K. R. AND TIJMS, H. C. (1975) On the D-policy for the $M/D/1$ queue. *Management Sci.* **21,** 1073–1076.

[3] BELL, C. (1971) Characterization and computation of optimal policies for operating an $M/G/1$ queueing system with removable server. *Operat. Res.* **19,** 208–218.

[4] BERTSEKAS, D. P. (1976) *Dynamic Programming and Stochastic Control.* Academic Press, New York.

[5] FEDERGRUEN, A. AND GREEN, L. (1986) Queueing systems with service interruptions. *Operat. Res.* **34,** 752–768.

[6] FEDERGRUEN, A. AND GREEN, L. (1988) Queueing systems with service interruptions II. *Naval Res. Logist. Quart.* **35,** 345–358.

[7] FEDERGRUEN, A., HORDIJK, A., AND TIJMS, H. C. (1979) Denumberable state semi-Markov decision processes with unbounded costs, average cost criterion. *Stoch. Proc. Appl.* **9,** 222–235.

[8] FEDERGUEN, A. AND SO, K. C. (1989) Optimal maintenance policies for single-server queueing systems subject to breakdowns. *Operat. Res.* To appear.

[9] GAVER, D. P. Jr. (1962) A waiting line with interrupted service, including priorities. *J. R. Statist. Soc.* **B24,** 73–90.

[10] HEYMAN, D. (1977) The T-policy of the $M/G/1$ queue. *Management Sci.* **23,** 775–778.

[11] MCCALL, J. J. (1965) Maintenance policies for stochastically failing equipment: a survey. *Management Sci.* **11,** 493–624.

[12] NEUTS, M. F. (1981a) *Matrix Geometric Solutions in Stochastic Models.* Johns Hopkins University Press, Baltimore, Maryland.

[13] NEUTS, M. F. (1981b) Stationary waiting time distributions in the $GI/PH/1$ queue. *J. Appl. Prob.* **18,** 901–912.

[14] PIERSKALLA, W. P. AND VOELKER, J. A. (1976) A survey of maintenance models: the control and surveillance of deteriorating systems. *Naval Res. Logist. Quart.* **23,** 353–388.

[15] ROSS, S. M. (1970) *Applied Probability Models with Optimization Applications.* Holden-Day, San Francisco.

[16] SCHWEITZER, P. J. (1971) Iterative solution of the functional equations for undiscounted Markov renewal programming. *J. Math. Anal. Appl.* **34,** 495–501.

[17] SELBY, S. M. (1980) *Standard Mathematical Tables.* CRC Press, Cleveland.

[18] SHERIF, Y. S. AND SMITH, M. L. (1981) Optimal maintenance models for systems subject to failure—a review. *Naval Res. Logist. Quart.* **28,** 47–54.

[19] SO, K. C. (1985) Optimal Maintenance Policies for Single-Server Queueing Systems Subject to Breakdowns. Unpublished Ph.D. dissertation. Stanford University.

[20] SOBEL, M. F. (1969) Optimal average cost policy for a queue with start-up and shut-down costs. *Operat. Res.* **17,** 145–162.

[21] TIJMS, H. C. (1986) *Stochastic Modelling and Analysis: A Computational Approach.* Wiley, Chichester.

[22] WEBER, R. R. AND STIDHAM, Jr., S. (1987) Optimal control of service rates in networks of queues. *Adv. Appl. Prob.* **19,** 202–218.

[23] WHITE, H. AND CHRISTIE, L. S. (1958) Queueing with preemptive priorities or with breakdowns. *Operat. Res.* **6,** 79–95.

[24] WOLFF, R. W. (1984) Conditions for finite ladder height and delay moments. *Operat. Res.* **32,** 909–916.

[25] YADIN, M. AND NAOR, P. (1963) Queueing systems with a removable service station. *Operat. Res. Quart.* **14,** 393–405.