Probability in the Engineering and Informational Sciences, **20**, 2006, 543–570. Printed in the U.S.A.

OPTIMAL ADMISSION CONTROL IN QUEUES WITH WORKLOAD-DEPENDENT SERVICE RATES

RENÉ BEKKER*
Vrije Universiteit, Amsterdam
1081 HV Amsterdam, The Netherlands
E-mail: rbekker@few.vu.nl

SEM C. BORST
Department of Mathematics & Computer Science
Eindhoven University of Technology
5600 MB Eindhoven, The Netherlands
and
CWI
1090 GB Amsterdam, The Netherlands
and
Bell Laboratories
Lucent Technologies
Murray Hill, NJ 07974
E-mail: sem@cwi.nl

We consider a queuing system with a workload-dependent service rate. We specifically assume that the service rate is first increasing and then decreasing as a function of the amount of work. The latter qualitative behavior is quite common in practical situations, such as production systems. The admission of work into the system is controlled by a policy for accepting or rejecting jobs, depending on the state of the system. We seek an admission control policy that maximizes the long-run throughput. Under certain conditions, we show that a threshold policy is optimal, and we derive a criterion for determining the optimal threshold value.

*This work was done while the author was affiliated with the Department of Mathematics & Computer Science, Eindhoven University of Technology, The Netherlands, and with CWI, Amsterdam.

© 2006 Cambridge University Press 0269-9648/06 $16.00

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

1. INTRODUCTION

Queuing systems with state-dependent rates occur in many practical situations. A specific example concerns production systems where the productivity of the shop floor personnel depends on the level of work-in-process (workload). In particular, the productivity (i.e., the speed of the server) first increases when the workload is low until a certain optimum is attained and then decreases when the system reaches overload (caused by, e.g., stress factors); see, for instance, [3,10]. The latter qualitative behavior is quite characteristic of efficiency patterns observed in many practical scenarios.

In this article we consider an $M/G/1$ queue with a service rate that is first increasing and then decreasing as a function of the workload. We assume that the speed of the server is intrinsically determined by the workload and the system characteristics and thus cannot be directly controlled. However, there is an admission policy to control the amount of work present. Depending on the state of the system, arriving customers might be either accepted or rejected, or, equivalently, the facility might be either open or closed for potential customers. The aim of this article is to find the admission policy that maximizes the long-run throughput. In particular, under some assumption, we show that a threshold policy is optimal and give an intuitively appealing expression for the optimal threshold value.

The above-described $M/G/1$ queue with admission control may be modeled as a semi-Markov decision process (MDP). Most of the theory on MDPs concerns models with finite or countable state spaces. Because in the present queuing model both the admission policy and the service speed depend on the workload, we are dealing with an MDP with uncountable state space $[0,\infty)$; see, for instance, [11,12,14] for some general MDPs with infinite state spaces. To derive structural properties of the optimal policy, a commonly used approach in MDPs is the construction of value functions that possess certain concavity properties. Because the value functions in our model typically do not exhibit such behavior, we apply sample-path techniques to compare different policies.

An interesting study showing several similarities with our model is [7]. In [7], the author considers an $M/G/1$ queuing system with continuous-time arrival control and a fixed reward rate R when the server is busy and holding cost rate cx when the workload is x. Hence, such an $M/G/1$ queue can also be modeled as an MDP with the admission control depending on the system state and the state space is infinite $[0,\infty)$. Using sample-path arguments and general theory on continuous-time MDPs developed in [6], the author proves the average-cost optimality of threshold policies.

Another branch of single-server queues with uncountable state spaces concerns $M/G/1$ queues with service control. Specifically, the service speed can be continuously adapted based on the residual amount of work. In [5], the service speed equals r_1 when the workload is less than some fixed level K and equals r_2 when the workload exceeds K. Under some fairly general cost functions, the author determines the optimal switching level K. In [8,15], the server works at constant

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

speed, but can be switched on and off. The cost function includes holding cost and switching cost for turning the server on. The average-cost optimality of D policies is shown in [8,15]. In D policies, the server is turned off only when the system becomes empty (while the server was on) and the server is turned on only when the workload exceeds level D (and the server was off).

This article is organized as follows. We give a detailed model description and several representations of the throughput in Section 2. In Section 3 the optimality of threshold policies under Assumption 2.1 (see Section 2) is shown. A criterion for the optimal threshold value is derived in Section 4. In Section 5 we present several examples of (combinations of) service speed functions and service requirement distributions satisfying Assumption 2.1. We explicitly determine the optimal threshold value and corresponding throughput in Section 6 for some special cases. Some concluding remarks and suggestions for further research are given in Section 7.

2. MODEL DESCRIPTION

We consider an $M/G/1$ queue with a workload-dependent service rate. The customers (or jobs) arrive according to a Poisson process of rate λ. The service requirement of the nth customer is B_n, $n = 1,2,\ldots$, where the B_n are assumed to be independent and identically distributed (i.i.d.) copies of a random variable B with distribution $B(\cdot)$ and mean β. We also assume that the sequences of interarrival intervals and service requirements are independent.

The server works at a rate that depends on the amount of work in the system as described by some function $r(\cdot)$; that is, the service rate is $r(x)$ when the amount of work is x. As in [1,2,9], we assume that $r(0) = 0$ and that $r(\cdot)$ is strictly positive and left-continuous and has a right limit on $(0,\infty)$. In addition, we specifically focus on the case that $r(\cdot)$ is increasing on $(0, r_{\max}]$ and decreasing on $(r_{\max},\infty)$ for some $r_{\max} \geq 0$.

The admission of work into the system is governed by a control policy that prescribes whether arriving customers are accepted or rejected, depending on the state of the system. We assume that the service requirement of a customer only becomes known after the acceptance decision; see Section 7 for a further discussion. Thus, the admission control policy might equivalently be interpreted as a rule for closing or opening access to the system.

We seek an admission control policy that maximizes the long-run throughput. The long-run throughput under policy π is defined as

$$TH^{\pi} := \lim_{t\to\infty} \frac{B^{\pi}(0,t)}{t},$$

assuming the limit to exist. Here $B^{\pi}(0,t)$ denotes the amount of work completed during $[0,t]$ under policy π. A policy π^* is said to be (strictly) optimal if $TH^{\pi^*} \geq TH^{\pi}$ ($TH^{\pi^*} > TH^{\pi}$) for all policies $\pi \neq \pi^*$.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

For now, we restrict the attention to the class of stationary and deterministic policies that base their actions on the current amount of work in the system only. For a given policy π, we use $\pi(x) = 1$ to denote that it accepts a customer that arrives when the workload equals x and write $\pi(x) = 0$ otherwise. Later we will show that the found optimal policy is in fact optimal within a broader class that includes nonstationary and randomized policies as well.

Let V_t^π be the workload at time t and let W_n^π be the workload just before the nth arrival epoch. Denote by V^π and W^π the random variables with the corresponding steady-state distributions, if they exist, and let $v^\pi(\cdot)$ be the density of V^π.

We first consider the case $\lambda\beta < r_\infty$, with $r_\infty := \lim_{x\to\infty} r(x)$. In that case, the system remains stable under the greedy policy that always accepts customers. Thus, the throughput achieved under the latter policy equals $\lambda\beta$, which is optimal, since the maximum achievable long-run throughput is bounded by the offered traffic load.

In the remainder of the article we focus on the case $\lambda\beta > r_\infty$. (The boundary case $\lambda\beta = r_\infty$ is rather delicate, and a full analysis is beyond the scope of the present article.) In that case, the system is unstable under the greedy policy that always accepts customers. Henceforth, we restrict the attention to policies π such that $\pi(x) = 0$ for all $x > M$ for some large M, which ensures the existence of the steady-state workload distribution. Even though the policy that always accepts customers might continue to be optimal, the maximum achievable throughput can be approached arbitrarily close for sufficiently large M.

Since the steady-state workload distribution exists, the throughput TH^π under policy π as defined earlier might in fact be expressed in several alternative ways. Observing that $B^\pi(0,t) = \int_0^t r(V_u^\pi)\,du$, the throughput can be equivalently written as

$$TH^\pi = \lim_{t\to\infty} \frac{1}{t} \int_0^t r(V_u^\pi)\,du = \mathbb{E}[r(V^\pi)] = \int_0^\infty r(x) v^\pi(x)\,dx.$$

Invoking the further identity relation (with $A^\pi(0,t)$ denoting the amount of work accepted during $[0,t]$ under policy π)

$$B^\pi(0,t) = V_0^\pi + A^\pi(0,t) - V_t^\pi$$

and noting that $V_t^\pi/t \to 0$ as $t \to \infty$, we observe that the throughput can also be expressed as

$$TH^\pi = \beta \lim_{t\to\infty} \frac{N^\pi(0,t)}{t},$$

where $N^\pi(0,t)$ denotes the number of accepted customers during $[0,t]$ under policy π. Using the PASTA property, the above expression can be further rewritten as

$$TH^\pi = \lambda\beta \mathbb{P}(\pi(V^\pi) = 1) = \lambda\beta \left(\pi(0)\mathbb{P}(V^\pi = 0) + \int_0^\infty \pi(x) v^\pi(x)\,dx \right). \quad \textbf{(1)}$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Finally, we introduce some additional notation. Define

$$R(x) := \int_0^x \frac{1}{r(y)}\, dy, \qquad 0 < x < \infty,$$

representing the time required for the system to empty in the absence of any arrivals, starting from workload x. In order to avoid technicalities, we assume that $R(x) < \infty$ for all $x > 0$, as in [9]. Moreover, we assume that $\mathbb{E}[R(x + B + \delta) - R(x + B)] \to 0$, as $\delta \to 0$. The latter condition only rules out cases where the workload process is being absorbed in some positive workload level and is satisfied if, for instance, $r_\infty > 0$ or if B has finite support. Further, define

$$Z(x) := \int_0^\infty \int_x^{x+y} \frac{1}{r(z)}\, dz\, dB(y) = \int_x^\infty \frac{1}{r(z)} (1 - B(z - x))\, dz, \tag{2}$$

representing the expected time required for the system to return to workload x after a customer has been accepted, in the absence of any further arrivals. In the remainder of the article, we make the following assumption with regard to $Z(x)$.

ASSUMPTION 2.1: *There exists some* $z_{\min} \geq 0$ *such that* $Z(x)$ *is decreasing on* $[0, z_{\min}]$ *and increasing on* $[z_{\min}, \infty)$.

The above assumption is satisfied for a wide class of $M/G/1$-type models with workload-dependent service rates. We give several illustrative examples in Section 5.

To provide some intuition, suppose that the system operates according to the last-come first-served preemptive-resume (LCFS-PR) discipline, which does not affect the workload process in any way. With that view in mind, $Z(x)$ might be thought of as the expected service time of a customer that arrives when the workload equals x, and $z_{\min}$ represents the workload level at which arriving customers have the minimum expected service time. Thus, from the LCFS-PR perspective, the direct reward of accepting customers is first increasing (on $(0, z_{\min}]$) and then decreasing (on $(z_{\min}, \infty)$). However, the decision to either accept or reject also affects future rewards (service times). In Section 3 insights from the LCFS-PR discipline are applied to show that the optimal policy has a threshold structure when Assumption 2.1 is satisfied.

3. OPTIMALITY OF THRESHOLD POLICIES

In the first part of this section we only consider stationary deterministic policies. Since the actions of the admission control policy then only depend on the workload level x, we will also, for brevity, refer to the value of x as the state of the system. An excursion from state x is then the period that starts with the acceptance of a customer in state x and ends with the first subsequent return to state x. For conciseness, we will frequently write that a policy accepts/rejects in an interval $[v, w]$ when it accepts/rejects customers that arrive when the workload is in the interval $[v, w]$. In

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

the second part of this section we show that the found optimal policy is in fact optimal within a broader class that also includes nonstationary and randomized policies.

Define $N^{\pi}(x)$ and $T^{\pi}(x)$ as follows:

$N^{\pi}(x) \equiv$ expected number of accepted customers during an excursion from state x under policy π.

$T^{\pi}(x) \equiv$ expected duration of an excursion from state x under policy π.

It can be verified that $N^{\pi}(x)$ and $T^{\pi}(x)$ are continuous; see also the proof of Lemma 3.2.

Consider an arbitrary policy π that rejects in $[x, x+\delta]$. Let π' be a modified policy, which does the same as π except that it accepts in $[x, x+\delta]$. Let $G^{\pi}(y)$ be the expected number of excursions during a busy cycle that start from a workload level below y under policy π, which are not part of an excursion starting from a level $z \in [x, y]$, $y \geq x$.

LEMMA 3.1: *For some $\gamma \in (0,1)$, we have*

$$TH^{\pi'} = (1-\gamma)TH^{\pi} + \gamma \frac{\displaystyle\int_x^{x+\delta} \beta N^{\pi'}(y)\, dG^{\pi'}(y)}{\displaystyle\int_x^{x+\delta} T^{\pi'}(y)\, dG^{\pi'}(y)}.$$

PROOF: By [13, Thm. 1], the throughput under policy π can be equivalently expressed as $TH^{\pi} = \mathbb{E}R^{\pi}/\mathbb{E}T^{\pi}$, where R^{π} is the reward (i.e., amount of work served) during a busy cycle and T^{π} is the cycle length under policy π. Consider a busy cycle and take an arbitrary sample path of the workload process $\{V_t^{\pi'}, t \geq 0\}$ under policy π'. We construct a stochastic process $\hat{V}_t$ by deleting the excursions from level $y \in [x, x+\delta]$ and pasting together the remaining parts. First note that the residual interarrival time at a downcrossing of y is still exponential (see, e.g., [1]). Now, it may be readily checked that $\hat{V}_t$ and V_t^{π} have the same statistical properties. Thus, for the expected number of accepted customers during a busy cycle under policy π', $\mathbb{E}N^{\pi'}$, we have

$$\mathbb{E}N^{\pi'} = \mathbb{E}N^{\pi} + \int_x^{x+\delta} N^{\pi'}(y)\, dG^{\pi'}(y)$$

and, equivalently, for the expected duration of a busy cycle, we have

$$\mathbb{E}T^{\pi'} = \mathbb{E}T^{\pi} + \int_x^{x+\delta} T^{\pi'}(y)\, dG^{\pi'}(y).$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Using Wald's theorem, we derive

$$\frac{\mathbb{E}R^{\pi'}}{\mathbb{E}T^{\pi'}} = \frac{\beta\mathbb{E}N^{\pi'}}{\mathbb{E}T^{\pi'}} = \frac{\beta\left(\mathbb{E}N^{\pi} + \int_x^{x+\delta} N^{\pi'}(y)\, dG^{\pi'}(y)\right)}{\mathbb{E}T^{\pi} + \int_x^{x+\delta} T^{\pi'}(y)\, dG^{\pi'}(y)}$$

$$= (1-\gamma)\frac{\mathbb{E}R^{\pi}}{\mathbb{E}T^{\pi}} + \gamma\frac{\int_x^{x+\delta} \beta N^{\pi'}(y)\, dG^{\pi'}(y)}{\int_x^{x+\delta} T^{\pi'}(y)\, dG^{\pi'}(y)},$$

where

$$\gamma = \frac{\int_x^{x+\delta} T^{\pi'}(y) dG^{\pi'}(y)}{\mathbb{E}T^{\pi} + \int_x^{x+\delta} T^{\pi'}(y) dG^{\pi'}(y)}$$

represents the fraction of time spent on excursions starting between x and $x+\delta$. This completes the proof. ■

Let π^* denote an optimal policy, with corresponding throughput $TH^* = \mathbb{E}[r(V^{\pi^*})]$.

LEMMA 3.2 (Optimality properties):

(i) It is strictly optimal to reject in $[v,w] \Rightarrow \beta N^{\pi^*}(x)/T^{\pi^*}(x) < TH^*$, *for almost every* $x \in [v,w]$.

(ii) It is optimal to accept in $[v,w] \Rightarrow \beta N^{\pi^*}(x)/T^{\pi^*}(x) \geq TH^*$, $\forall x \in [v,w]$.

Note that the inequality in (i) may hold with equality for some $x \in [v,w]$.

PROOF: We first prove that

$$\frac{\int_x^{x+\delta} N^{\pi'}(y)\, dG^{\pi'}(y)}{\int_x^{x+\delta} T^{\pi'}(y)\, dG^{\pi'}(y)} \to \frac{N^{\pi}(x)}{T^{\pi}(x)} \quad \text{as } \delta \downarrow 0. \tag{3}$$

For some small $\delta > 0$ and $y \in [x, x+\delta]$, we have

$$T^{\pi'}(y) - T^{\pi}(x+\delta) \leq (R(x+\delta) - R(y))\left(1 + \lambda \max_{y \leq u \leq x+\delta} T^{\pi'}(u)\right) \to 0 \quad \text{as } \delta \downarrow 0,$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

where $R(x+\delta) - R(y)$ is the time required to go from $x+\delta$ to y in the absence of any arrivals. Similarly, as $\delta \downarrow 0$, $T^{\pi'}(y) - T^{\pi}(x+\delta)$ can be bounded from below by

$$-\mathbb{E}[R(x+B+\delta) - R(y+B)]\left(1 + \lambda \max_{y+B \le u \le x+B+\delta} \pi'(u) T^{\pi'}(u)\right) \to 0.$$

Applying similar arguments to $N^{\pi'}(y)$ then yields (3). (Another way to see that (3) holds is to observe that the density $dG^{\pi}(\cdot)$ is well defined.)

The remainder of the proof is by contradiction. For part (i), assume that the strictly optimal policy π^* rejects in $[v,w]$, but there is some interval $(u,u+\delta) \subseteq [v,w]$, with $\delta > 0$, such that $\beta N^{\pi^*}(x)/T^{\pi^*}(x) \ge TH^*$ for $x \in (u,u+\delta)$. Consider a modified policy π that accepts in $[u,u+\delta]$ and follows π^* otherwise. First using Lemma 3.1 and then letting $\delta \downarrow 0$ (and using (3)), it follows that $\mathbb{E}[r(V^{\pi})] \ge TH^*$, contradicting the strict optimality of π^*. For part (ii), assume that the optimal policy π^* accepts in $[v,w]$ but, for some $x \in [v,w]$, $\beta N^{\pi^*}(x)/T^{\pi^*}(x) < TH^*$. Using (3), it follows that there is some interval $U := (u-\delta, u+\delta) \subseteq [v,w]$ such that $\beta N^{\pi^*}(x)/T^{\pi^*}(x) < TH^*$ for every $x \in U$. Consider the modified policy π that rejects in U and follows π^* otherwise. Using Lemma 3.1 (with $\pi' \equiv \pi^*$), it is easily seen that $\mathbb{E}[r(V^{\pi})] > TH^*$, contradicting the optimality of π^*. ■

LEMMA 3.3: *It is optimal to accept in* $[0, z_{\min}]$.

PROOF: It is obvious that it is optimal to accept in an empty system. Now assume that it is not optimal to accept in $[0, z_{\min}]$. Then there is some policy π such that $\pi(x) = 1$ for $x \in [0, u^*]$, but $\pi(x) = 0$ for $x \in (u^*, u^*+\delta]$, with $u^* + \delta \le z_{\min}$ and $\delta > 0$, that is strictly optimal.

Take some arbitrary $0 < y < \delta$. In the proof, we compare $N^{\pi}(u^*)$ and $T^{\pi}(u^*)$ with $N^{\pi}(u^*+y)$ and $T^{\pi}(u^*+y)$. Using stochastic coupling, we show that $\beta N^{\pi}(u^*+y)/T^{\pi}(u^*+y)$ can be written as a combination of $\beta N^{\pi}(u^*)/T^{\pi}(u^*)$ and possibly contributions from some additional excursions. Since π is assumed to be optimal, both terms provide an average reward of at least TH^* by Lemma 3.2(ii). By Lemma 3.2(i), this contradicts the strict optimality of rejecting in $(u^*, u^*+\delta]$, because the coupling holds for any $y \in (0,\delta)$.

For the first part in the stochastic coupling (i.e., the part of the excursion from u^*+y related to $\beta N^{\pi}(u^*)/T^{\pi}(u^*)$), observe that it follows from Assumption 2.1 that $Z(u^*) \ge Z(u^*+y)$, implying that the direct reward of accepting customers at level u^*+y is at least as high as the direct reward of accepting at level u^*. For the second part, we use the fact that we only make additional excursions if they are advantageous.

First consider the expected duration of an excursion from level u^* under policy π and the expected number of accepted customers during such an excursion (i.e., $N^{\pi}(u^*)$ and $T^{\pi}(u^*)$). Let the first jump, initiating an excursion, occur at time 0 and

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

observe that the workload level right after the first jump equals $u^* + B$ (i.e., $V_{0^+}^{\pi} = u^* + B$). Note that the workload process attains local minima just before arrival instants at which customers are going to be accepted. Using terminology of random walks, define a stopping time $\tau_s^{\pi} := \inf\{t \geq 0 : V_t^{\pi} \leq u^*\}$, an equivalent notion measured in the number of arrivals $\tau^{\pi} := \inf\{k \geq 0 : W_k^{\pi} \leq u^*\}$, and a sequence of descending ladder epochs $\tau^{\pi}(1) < \cdots < \tau^{\pi}(N) < \tau^{\pi}$ with corresponding descending ladder heights $u^* + B > W_{\tau^{\pi}(1)}^{\pi} > \cdots > W_{\tau^{\pi}(N)}^{\pi} > u^*$ as follows: $\tau^{\pi}(1) := \inf\{0 \leq k \leq \tau^{\pi} : \pi(W_k^{\pi}) = 1\}$, and for $n = 2, \ldots, N$ (if $\tau^{\pi}(1) < \tau^{\pi}$),

$$\tau^{\pi}(n+1) := \inf\{\tau^{\pi}(n) < k < \tau^{\pi} : W_k^{\pi} < W_{\tau^{\pi}(n)}^{\pi}, \pi(W_k^{\pi}) = 1\}.$$

Note that $W_{\tau^{\pi}(N)}^{\pi} > u^* + \delta$, since $\pi(x) = 0$ for $x \in [u^*, u^* + \delta]$. A typical sample path in case $N = 2$ is depicted in Figure 1a. Using the above, we write

$$N^{\pi}(u^*) = 1 + \sum_{n=1}^{N} N^{\pi}(W_{\tau^{\pi}(n)}^{\pi}), \tag{4}$$

$$T^{\pi}(u^*) = Z(u^*) + \sum_{n=1}^{N} T^{\pi}(W_{\tau^{\pi}(n)}^{\pi}). \tag{5}$$

Now consider $N^{\pi}(u^* + y)$ and $T^{\pi}(u^* + y)$. In this case, at time 0 the workload jumps to $u^* + y + B$ (i.e., $V_{0^+}^{\pi} = u^* + y + B$). As defined earlier, we have a stopping time $\tilde{\tau}_s^{\pi}$, a discrete-time equivalent $\tilde{\tau}^{\pi}$, and a sequence of descending ladder epochs $0 < \tilde{\tau}^{\pi}(1) < \cdots < \tilde{\tau}^{\pi}(M) < \tilde{\tau}^{\pi}$ with corresponding descending ladder heights $u^* + y + B > W_{\tilde{\tau}^{\pi}(1)}^{\pi} > \cdots > W_{\tilde{\tau}^{\pi}(M)}^{\pi} > u^* + y$ (see Figure 1b for a typical realization). Observe that the residual interarrival time at a downcrossing of $u^* + B$ is still exponential. Hence, using stochastic coupling and the fact that $W_{\tau^{\pi}(N)}^{\pi} > u^* + \delta$, the descending ladder epochs can be divided into two sets: (i) $\tilde{\tau}^{\pi}(1), \ldots, \tilde{\tau}^{\pi}(M-N)$ with $u^* + y + B > W_{\tilde{\tau}^{\pi}(1)}^{\pi} > \cdots > W_{\tilde{\tau}^{\pi}(M-N)}^{\pi} > u^* + B$ and (ii) $\tilde{\tau}^{\pi}(M-N+1), \ldots, \tilde{\tau}^{\pi}(M)$ such that $W_{\tilde{\tau}^{\pi}(n+M-N)}^{\pi} =^d W_{\tau^{\pi}(n)}^{\pi}$ for $n = 1, \ldots, N$. This coupling is illustrated in Figure 1 (with $N = 2$ and $M = 3$). In this figure, the sample paths in the range of the solid arrow (i.e., between $[0, s]$ and $[s', \tilde{\tau}_s^{\pi}]$ respectively) are identical. Using the above arguments, we have

$$N^{\pi}(u^* + y) = 1 + \sum_{n=1}^{M-N} N^{\pi}(W_{\tilde{\tau}^{\pi}(n)}^{\pi}) + \sum_{n=1}^{N} N^{\pi}(W_{\tau^{\pi}(n)}^{\pi}), \tag{6}$$

$$T^{\pi}(u^* + y) = Z(u^* + y) + \sum_{n=1}^{M-N} T^{\pi}(W_{\tilde{\tau}^{\pi}(n)}^{\pi}) + \sum_{n=1}^{N} T^{\pi}(W_{\tau^{\pi}(n)}^{\pi}). \tag{7}$$

Since $W_{\tilde{\tau}^{\pi}(n)}^{\pi}$, $n = 1, \ldots, M-N$, are the workloads just before an arriving customer is accepted and π is the supposed optimal policy, Lemma 3.2 yields

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

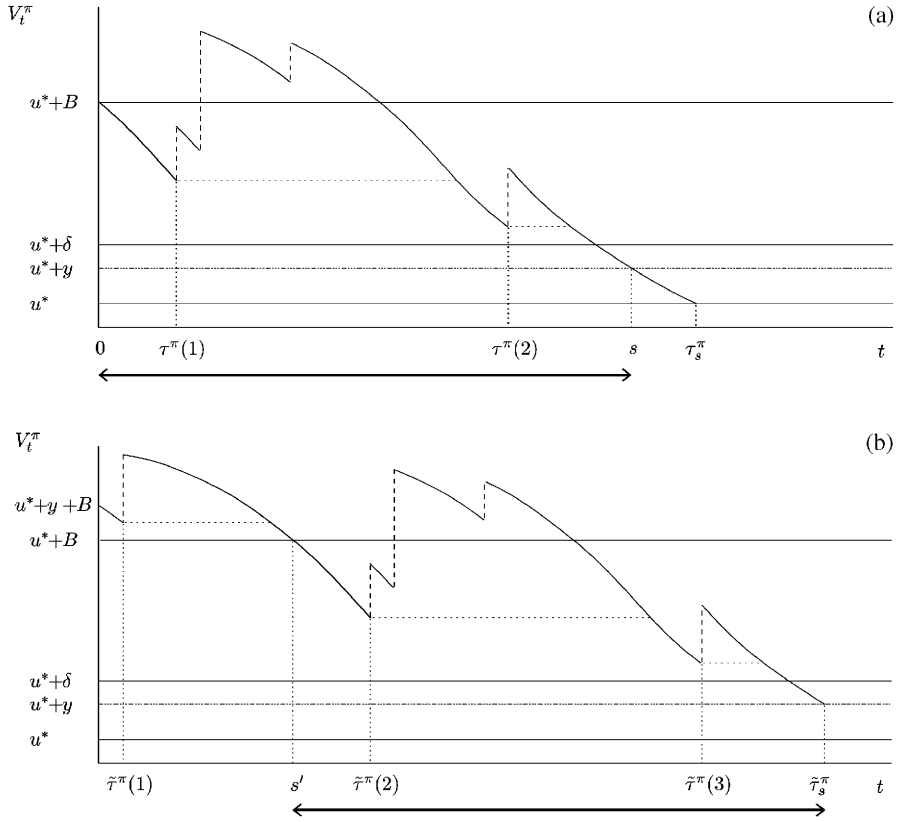


FIGURE 1. The sample paths of two excursions of V_t^π: one excursion from state u^* and one excursion from state $u^* + y$. In this example, $N = 2$ and $M = 3$.

$$\frac{\beta \sum_{n=1}^{M-N} N^\pi(W^\pi_{\tilde{\tau}^\pi(n)})}{\sum_{n=1}^{M-N} T^\pi(W^\pi_{\tilde{\tau}^\pi(n)})} \geq \beta \min_{n=1,\ldots,M-N} \frac{N^\pi(W^\pi_{\tilde{\tau}^\pi(n)})}{T^\pi(W^\pi_{\tilde{\tau}^\pi(n)})} \geq TH^*. \tag{8}$$

Moreover, using (4) and (5) in addition to Assumption 2.1, we obtain

$$\frac{\beta\left(1 + \sum_{n=1}^{N} N^\pi(W^\pi_{\tau^\pi(n)})\right)}{Z(u^* + y) + \sum_{n=1}^{N} T^\pi(W^\pi_{\tau^\pi(n)})} \geq \frac{\beta N^\pi(u^*)}{T^\pi(u^*)} \geq TH^*, \tag{9}$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

where the second inequality relies on the fact that it is optimal to accept at level u^*. Combining (6)–(9) yields

$$\frac{\beta N^{\pi'}(u^*+y)}{T^{\pi'}(u^*+y)}$$

$$\geq \min\left\{\frac{\beta\sum_{n=1}^{M-N} N^{\pi}(W^{\pi}_{\tilde{\tau}^{\pi}(n)})}{\sum_{n=1}^{M-N} T^{\pi}(W^{\pi}_{\tilde{\tau}^{\pi}(n)})}, \frac{\beta\left(1+\sum_{n=1}^{N} N^{\pi}(W^{\pi}_{\tau^{\pi}(n)})\right)}{Z(u^*+y)+\sum_{n=1}^{N} T^{\pi}(W^{\pi}_{\tau^{\pi}(n)})}\right\}$$

$$\geq TH^*.$$

By Lemma 3.2 it cannot be strictly optimal to reject at level u^*+y, $0<y<\delta$. ■

THEOREM 3.1: *There exists a threshold policy that is optimal among the class of stationary deterministic policies.*

PROOF: It follows from Lemma 3.3 that it is optimal to accept when the workload is in $[0, z_{\min}]$. Suppose that a threshold policy is not optimal; that is, there exists some policy π that is strictly better than any threshold policy. Let $n^{\pi} := \int_0^{\infty} \max(\pi(x^+)-\pi(x),0)\,dx$ be the number of "gaps" of policy π (i.e., the number of times $\pi(\cdot)$ switches from zero to one). Let π be an optimal policy, which is strictly better than any threshold policy, with the least number of gaps (i.e., $\pi = \arg\min_{\pi\in\Pi^*} n^{\pi}$, with Π^* the class of optimal policies). This implies that there is some $u^* > z_{\min}$ and $\delta_2 > \delta_1 > 0$ such that $\pi(x)=0$ on $(u^*, u^*+\delta_1)$ and $\pi(x)=1$ on $(u^*+\delta_1, u^*+\delta_2)$. We note that gaps consisting of singular points can be removed.

Take some arbitrary $0<y<\delta_1$. In the proof, we consider $N^{\pi}(u^*+y)$ and $T^{\pi}(u^*+y)$. Using the fact that it is optimal to accept in $(u^*+\delta_1, u^*+\delta_2)$, we show that $\beta N^{\pi}(u^*+y)/T^{\pi}(u^*+y) \geq TH^*$ (contradicting the fact that π contains the least number of gaps among policies in Π^*). This follows from the fact that the direct reward of accepting at level u^*+y exceeds the reward of accepting at any level $x > u^*+y$. Moreover, additional excursions are only made when they are advantageous.

Suppose that at time 0 an arriving customer with service requirement B is accepted when the workload equals u^*+y (i.e., $V^{\pi}_{0^+} = u^*+y+B$). As in the proof of Lemma 3.3 (see also Figure 1a, with $\delta_1 \equiv \delta$), we can define "stopping times" τ^{π}_s and τ^{π} and a sequence of descending ladder epochs $\tau^{\pi}(1) < \cdots < \tau^{\pi}(N) < \tau^{\pi}$ with corresponding descending ladder heights $u^*+y+B > W^{\pi}_{\tau^{\pi}(1)} > \cdots > W^{\pi}_{\tau^{\pi}(N)} > u^*+y$. Note that $W^{\pi}_{\tau^{\pi}(N)} > u^*+\delta_1$ (if $N>0$), since $\pi(x)=0$ for $x\in[u^*, u^*+\delta_1]$. Applying this construction yields

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

$$N^{\pi}(u^* + y) = 1 + \sum_{n=1}^{N} N^{\pi}(W^{\pi}_{\tau^{\pi}(n)}),$$

$$T^{\pi}(u^* + y) = Z(u^* + y) + \sum_{n=1}^{N} T^{\pi}(W^{\pi}_{\tau^{\pi}(n)}).$$

By Lemma 3.2, $\beta \sum_{n=1}^{N} N^{\pi}(W^{\pi}_{\tau^{\pi}(n)}) / \sum_{n=1}^{N} T^{\pi}(W^{\pi}_{\tau^{\pi}(n)}) \geq TH^*$ since π is assumed to be an optimal policy. Moreover, using a similar ladder height construction, it can be easily checked (in general) that

$$\frac{N^{\pi}(x)}{T^{\pi}(x)} \leq \max_{v \geq x} \frac{1}{Z(v)}. \tag{10}$$

Hence, invoking Assumption 2.1 yields

$$\frac{\beta}{Z(u^* + y)} \geq \frac{\beta}{Z(u^* + \delta_1)} \geq TH^*.$$

Combining the above, we obtain $\beta N^{\pi}(u^* + y)/T^{\pi}(u^* + y) \geq TH^*$ for any $y \in (0, \delta_1)$. By Lemma 3.2, this contradicts the fact that policy π has the minimum number of gaps among the class of optimal policies Π^*. ■

The ladder height construction in the proof of Theorem 3.1 allows us to generalize (10):

PROPOSITION 3.1: *For the throughput during an excursion from level x, we have the following bounds:*

$$\min_{v \geq x : \pi(v) = 1} \frac{1}{Z(v)} \leq \frac{N^{\pi}(x)}{T^{\pi}(x)} \leq \max_{v \geq x : \pi(v) = 1} \frac{1}{Z(v)}.$$

These bounds are especially natural from the perspective of the LCFS-PR discipline. In that view, the proposition simply states that the throughput during an excursion from level x is at least the minimum (and at most the maximum) of one over the mean service time of accepting at any level above x if policy π is applied.

Remark 3.1: The proof of Theorem 3.1 crucially depends on the fact that $Z(\cdot)$ has only one local minimum (i.e., Assumption 2.1). Suppose for the moment that $Z(\cdot)$ has L local minima. Thus, $Z(\cdot)$ is decreasing on $[z^k_{\max}, z^k_{\min})$ and increasing on $[z^k_{\min}, z^{k+1}_{\max})$, $k = 1, \ldots, L$, where $z^1_{\max} = 0$ and $z^{L+1}_{\max} = \infty$. Similar to the proof of Lemma 3.3, we deduce that if $\pi(x) = 1$ for some $x \in [z^k_{\max}, z^k_{\min})$, then $\pi(y) = 1$ for all $y \in [x, z^k_{\min})$ (note that $\pi(0) = 1$ and accepting is thus optimal in $[0, z^1_{\min})$). Also, it follows from the proof of Theorem 3.1 that if $\pi(x) = 1$ for some $x \in [z^L_{\max}, \infty)$, then $\pi(y) = 1$ for all $y \in [z^L_{\max}, x)$. However, the intervals $[z^k_{\min}, z^{k+1}_{\max})$, $k = 1, \ldots, L - 1$, are not covered by the proof. In particular, the trade-off between direct and future rewards remains undecided there.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Theorem 3.1 shows that the threshold policy is optimal among the class of stationary and deterministic policies. To prove that a (stationary and deterministic) threshold policy is also optimal within the broader class of policies considered in [13], we use insights from this section to construct an appropriate (value) function satisfying [13, Thm. 2]. The class of policies in [13] consists of all measurable decision rules and includes nonstationary and nondeterministic policies.

THEOREM 3.2: *There exists a threshold policy that is optimal within the class of policies considered in [13].*

PROOF: Let π be a threshold policy with threshold value x^* that is optimal within the class of stationary and deterministic policies. Now define $n^\pi(x)$ and $t^\pi(x)$ as follows:

$n^\pi(x) \equiv$ expected amount of work served in a period starting with workload level x until the end of the busy cycle under policy π.

$t^\pi(x) \equiv$ expected length of a period starting with workload level x until the end of the busy cycle under policy π.

Similar to [7], let

$$\tilde{f}(x) := n^\pi(x) - TH^\pi t^\pi(x).$$

Consider $\mathbb{E}[\tilde{f}(x+B)]$ and divide the busy cycle in two parts; first we have an excursion from state x followed by the remaining part of the cycle starting with a downcrossing of level x. Hence (see also [7, Lemma 6.3]),

$$\mathbb{E}[\tilde{f}(x+B)] = \beta(N^\pi(x) - 1) - TH^\pi T^\pi(x) + \tilde{f}(x), \tag{11}$$

where the $N^\pi(x) - 1$ stems from the fact that the arrival in state x is not counted in $\mathbb{E}[\tilde{f}(x+B)]$.

Define, for $x \geq 0$,

$$f(x) := \begin{cases} \beta + \mathbb{E}[\tilde{f}(x+B)] & \text{for } 0 \leq x \leq x^* \\ \tilde{f}(x) & \text{for } x > x^*, \end{cases} \tag{12}$$

where x is the state of the system just before a decision epoch. By conditioning on the first arrival, we also obtain the following relationship between $\tilde{f}(\cdot)$ and $f(\cdot)$:

$$\tilde{f}(x) = \int_0^\infty f(R^{-1}((R(x) - y)^+))\lambda e^{-\lambda y}\, dy - \frac{TH^\pi}{\lambda}, \tag{13}$$

with $R^{-1}(\cdot)$ the inverse function of $R(\cdot)$; see for example, [1,2,9] for details. Because π is assumed to be an optimal stationary deterministic policy, Lemma 3.2 yields that $\beta N^\pi(x) - TH^\pi T^\pi(x)$ is positive for $x \in [0, x^*)$ and nonpositive for $x \in [x^*, \infty)$. Using the above in addition to (11) and (12), we obtain

$$f(x) = \max\{\beta N^\pi(x) - TH^\pi T^\pi(x), 0\} + \tilde{f}(x). \tag{14}$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Combining (11) with (13), we can rewrite (14) as

$$f(x) = \max\left\{\beta + \int_0^\infty \int_0^\infty f(R^{-1}((R(x+z)-y)^+))\lambda e^{-\lambda y}\,dy\,dB(z), \right.$$
$$\left. \int_0^\infty f(R^{-1}((R(x)-y)^+))\lambda e^{-\lambda y}\,dy\right\} - \frac{TH^\pi}{\lambda}.$$

Thus, the function $f(\cdot)$ satisfies the optimality equation for the average-cost criterion (i.e. Eq. (3) in [13]). The theorem now follows directly from [13, Thm. 2]. ■

4. CRITERION FOR THE OPTIMAL THRESHOLD

In Section 3 we showed that if Assumption 2.1 is satisfied, a threshold policy is optimal. The derivation of that result also suggested the following criterion for the optimal threshold:

$$TH^{\pi_{\bar{x}}} = \frac{\beta}{Z(\bar{x})}, \tag{15}$$

where $\pi_{\bar{x}}$ denotes a threshold policy with parameter $\bar{x}$. The above criterion is intuitively appealing when we consider marginal arguments. Informally speaking, the optimal threshold will be chosen such that the throughput just equals the expected reward of customers accepted in state $\bar{x}$ (which has reward $\beta/Z(\bar{x})$).

Moreover, the above criterion allows us to deduce some properties of the optimal threshold value. Using a similar construction as in (some of) the proofs of Section 3, it can be shown that $TH^{\pi_{\bar{x}}}$ is increasing as a function of λ. (To see this, we note that a higher λ yields additional arrivals that are only accepted if the resulting excursions are advantageous.) Because $\beta/Z(\bar{x})$ is independent of λ, we can directly conclude from (15) that the optimal threshold value is decreasing in λ. It can also easily be checked that the optimal threshold approaches $z_{\min}$ as $\lambda \to \infty$. This behavior of the optimal threshold reveals the typical trade-off between direct and future rewards; the upper bound for the throughput is attained by accepting customers in state $z_{\min}$, but the optimal policy anticipates decreasing arrival rates by starting to accept customers at increasing workload levels to compensate for the increased probability of reaching an empty system (where the server is idle).

In the remainder of this section we use another method to derive a criterion for the optimal threshold value and give some properties of $TH^{\pi_{\bar{x}}}$ as a function of $\bar{x}$. Moreover, when $Z(\cdot)$ does not satisfy Assumption 2.1, we show that a criterion similar to (15) holds for the optimal threshold value, which provides the optimal policy within the class of threshold policies. (Note that a threshold strategy might then not be optimal among the class of stationary and deterministic policies.) However, we start with the general form of the throughput under a threshold strategy with some fixed threshold $\bar{x}$.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Observe that for fixed $\bar{x}$, the workload under policy $\pi_{\bar{x}}$ has the same dynamics as an $M/G/1$ queue with a general service rate and impatience of customers depending on the amount of work found upon arrival. Under policy $\pi_{\bar{x}}$ the model is in fact a special case of the finite-buffer queue in [1], with

$$v^{\pi_{\bar{x}}}(x) = \begin{cases} \mathbb{P}(V^{\pi_{\bar{x}}} = 0)K^*(x,0), & 0 < x \le \bar{x} \\ \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left[K(x,0) + \int_0^{\bar{x}} K(x,y)K^*(y,0)\,dy\right], & x > \bar{x}, \end{cases}$$

where $\mathbb{P}(V^{\pi_{\bar{x}}} = 0)$ follows from normalization:

$$\mathbb{P}(V^{\pi_{\bar{x}}} = 0) = \left[1 + \int_0^{\bar{x}} K^*(x,0)\,dx + \int_{\bar{x}}^{\infty} K(x,0)\,dx + \int_{\bar{x}}^{\infty}\int_0^{\bar{x}} K(x,y)K^*(y,0)\,dy\,dx\right]^{-1}. \tag{16}$$

Here the (iterated) kernels are defined as in [1,9]; that is, for $0 \le y < x < \infty$, $K(x,y) := \lambda(1 - B(x-y))/r(x)$,

$$K_{n+1}(x,y) := \int_y^x K(x,z)K_n(z,y)\,dz, \tag{17}$$

and $K^*(x,0) := \sum_{n=1}^{\infty} K_n(x,0)$. Using the representation in (1) for the throughput, we obtain

$$TH^{\pi_{\bar{x}}} = \lambda\beta\mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left(1 + \int_0^{\bar{x}} K^*(x,0)\,dx\right). \tag{18}$$

Note that $Z(\bar{x})$ and $TH^{\pi_{\bar{x}}}$ are continuous and differentiable functions of $\bar{x}$. In order to determine the optimal threshold, it is useful to consider the derivative of $TH^{\pi_{\bar{x}}}$ with respect to $\bar{x}$.

LEMMA 4.1: *For the derivative of $TH^{\pi_{\bar{x}}}$, we have*

$$\frac{d}{d\bar{x}} TH^{\pi_{\bar{x}}} = \lambda\beta\mathbb{P}(V^{\pi_{\bar{x}}} = 0)K^*(\bar{x},0)\left[1 - \frac{TH^{\pi_{\bar{x}}}Z(\bar{x})}{\beta}\right].$$

PROOF: The proof is deferred to the Appendix. ■

Before we further discuss the optimal threshold criterion, we first derive some properties of $TH^{\pi_{\bar{x}}}$ as a function of $\bar{x}$. As in Lemma 3.1, consider a policy π that does not accept in $[a,b]$ and a modified policy π', which does the same as π except that $\pi'(x) = 1$ for $x \in [a,b]$. Then the throughput under policy π' can be written as a convex combination of the throughput under policy π and the throughput due to excursions starting from levels in $[a,b]$ (see Lemma 3.1). This relation is particularly useful in studying the relationship between $TH^{\pi_{\bar{x}}}$ and $Z(\cdot)$.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

LEMMA 4.2: *Suppose that (i) $dZ(x)/dx \leq 0$, for $x \in [a,b]$, and (ii) $TH^{\pi_a} \leq \beta/Z(a)$. Then*

$$TH^{\pi_x} \leq \frac{\beta}{Z(x)} \quad \textit{for all } x \in [a,b]. \tag{19}$$

If either (i) (for some $x \in [a,b]$) or (ii) holds with strict inequality, then (19) holds with strict inequality. Moreover, if the (strict) inequalities in (i) and (ii) are reversed, then the (strict) inequality in (19) is reversed.

PROOF: Fix an arbitrary $x \in (a,b]$. Lemma 3.1 yields that, for $\gamma \in (0,1)$,

$$TH^{\pi_x} = (1-\gamma)TH^{\pi_a} + \gamma \frac{\displaystyle\int_a^x \beta N^{\pi_x}(y)\, dG^{\pi_x}(y)}{\displaystyle\int_a^x T^{\pi_x}(y)\, dG^{\pi_x}(y)}.$$

From (i) and Proposition 3.1, we obtain $\beta N^{\pi_x}(y)/T^{\pi_x}(y) \leq \beta/Z(x)$ for every $y \in [a,x]$. Invoking (ii), it trivially follows that

$$TH^{\pi_x} \leq (1-\gamma)\frac{\beta}{Z(a)} + \gamma \frac{\beta}{Z(x)} \leq \frac{\beta}{Z(x)}, \tag{20}$$

where the last step is due to (i). Now if (i) holds with strict inequality for some $x \in [a,b]$, then the second inequality of (20) is strict, whereas the first one is strict if (ii) holds with strict inequality. The proof for the reversed signs is similar (use the lower bound in Proposition 3.1). ■

We now derive a criterion for the optimal threshold. Let π^*_{th} denote the optimal threshold strategy. Define the set $A := \{x \geq 0 : TH^{\pi_x} = \beta/Z(x)\}$. Note that, in general, A is a collection of N disjoint closed intervals A_i, $i = 1,\ldots,N$, where each interval can be a singleton. However, if A_i is not a singleton, then it follows directly from Lemma 4.2 that $Z(\cdot)$ is constant on A_i.

PROPOSITION 4.1: *If A is the empty set, then the greedy policy is optimal and $TH^{\pi^*_{th}} = r_\infty$. If A is nonempty,*

$$TH^{\pi^*_{th}} = \max\left\{r_\infty, \max_{x \in A} \beta/Z(x)\right\},$$

where the greedy policy is optimal when $TH^{\pi^*_{th}} = r_\infty$ and the optimal (finite) threshold is given by any $\bar{x} \in \arg\max_{x \in A} \beta/Z(x)$ otherwise.

PROOF: For the threshold at zero we have

$$TH^{\pi_0} = \frac{\beta}{Z(0)+\lambda} < \frac{\beta}{Z(0)}. \tag{21}$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

If A is the empty set, then we have from the continuity of $Z(\cdot)$ and $TH^{\pi_{\bar{x}}}$ that $TH^{\pi_x} < \beta/Z(x)$ for all x. Applying Lemma 4.1, we obtain that $d(TH^{\pi_x})/dx > 0$ for all x and the greedy policy is thus optimal.

If $N > 0$, then it follows from Lemma 4.1 that A contains all points satisfying $dTH^{\pi_x}/dx = 0$. Hence, A contains at least all extreme points. From (21) and Lemma 4.1 it follows that zero is a local minimum. Moreover, $TH^{\pi_x} \to r_\infty$ as $x \to \infty$. Because $TH^{\pi_{\bar{x}}}$ is continuous, finding the global maximum of $TH^{\pi_{\bar{x}}}$ reduces to finding the maximum of $\beta/Z(x)$, $x \in A$, and comparing it with r_∞. ■

Using Lemma 4.2, some additional properties of $TH^{\pi_{\bar{x}}}$ as a function of $\bar{x}$ can be derived. For instance, it can be shown that if $Z(\cdot)$ has m local maxima, then $N \leq 2m - 1$. In particular, if Assumption 2.1 is satisfied, then $N \leq 1$. This case is of special interest because a threshold policy is then optimal. Moreover, if in that case $N = 1$, then $A \cap [0, z_{\min})$ is empty and each value in A (possibly a singleton) is a global maximum of $TH^{\pi_{\bar{x}}}$. These arguments are summarized in the following corollary.

COROLLARY 4.1: *Suppose that Assumption 2.1 is satisfied. If A is the empty set, then the greedy policy is optimal. Otherwise, $N = 1$ and $x \in A$ is an optimal threshold with corresponding throughput*

$$TH^* = \beta/Z(x).$$

Finally, if $r(x)$ is constant for $x \geq L$ (and Assumption 2.1 holds), there is an easy way to determine directly whether the greedy policy is optimal. From Lemma 4.2 we then deduce the following:

COROLLARY 4.2: *Suppose Assumption 2.1 is satisfied and $r(x) = r_\infty$ for all $x \geq L$ for some $L > 0$. Then the greedy policy is optimal if and only if*

$$TH^{\pi_L} \leq \frac{\beta}{Z(L)} = r_\infty.$$

5. ASSUMPTION ON $Z(x)$

Although Assumption 2.1 is quite natural, it involves the service-rate function as well as the distribution of the service requirement. In this section we give some examples satisfying this assumption, assuming that $r(\cdot)$ is increasing on $(0, r_{\max}]$ and decreasing on $(r_{\max}, \infty)$ for some $r_{\max} \geq 0$ (as described in Section 2). We consider both cases with general service requirement distributions and cases with a wide class of service-rate functions. In addition, we provide a natural example that does not have the desired properties. This case reveals the strong dependence on both the service-rate function and the service requirement distribution.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

To show that Assumption 2.1 is satisfied, we frequently use the derivative of $Z(\cdot)$. Interchanging derivative and sum in addition to some rewriting yields

$$\frac{d}{dx} Z(x) = \int_0^\infty \frac{d}{dx} \left(R(x+b) - R(x)\right) dB(b)$$

$$= \int_0^\infty \frac{1}{r(x+b)} dB(b) - \frac{1}{r(x)} = \int_0^\infty \frac{r(x) - r(x+b)}{r(x+b)r(x)} dB(b).$$

For Assumption 2.1 to be satisfied, it remains to be shown that

$$r(x) - \mathbb{E}_B[r(x+B)] \leq 0, \qquad x \in [0, z_{\min}],$$

$$r(x) - \mathbb{E}_B[r(x+B)] \geq 0, \qquad x \in [z_{\min}, \infty).$$

Example 5.1: Suppose that $r_{\max} = 0$ (i.e., $r(\cdot)$ is decreasing on the positive half-line). By definition, $r(y) \geq r(x)$ for $y > x$, and it is readily seen that Assumption 2.1 is satisfied. Also, $z_{\min} = 0$ in this case.

Example 5.2: Suppose that $B(x) = I(x \geq \beta)$ (i.e., the service requirement is deterministic β). Observe that $\mathbb{E}_B[r(x+B)]$ is just the shifted $r(\cdot)$ function. Thus, $r(x) \leq \mathbb{E}_B[r(x+B)]$ if $x \in [0, (r_{\max} - \beta)^+]$ and $r(x) \geq \mathbb{E}_B[r(x+B)]$ if $x \in [r_{\max}, \infty)$. Moreover, $r(x)$ is increasing on $[(r_{\max} - \beta)^+, r_{\max}]$, whereas $\mathbb{E}_B[r(x+B)]$ is decreasing on the same interval. This directly yields the required property.

Example 5.3: Suppose that $B(x) = 1 - e^{-\mu x}$, meaning that the service requirement is exponentially distributed. Observe that $r(x) - \mathbb{E}_B[r(x+B)] \geq 0$ for $x \geq r_{\max}$. Now take some arbitrary x and y, with $0 < x < y \leq r_{\max}$. Conditioning on the service requirement in case a customer arrives at level x and using the memoryless property of the exponential distribution, we have

$$\mathbb{E}_B[r(x+B)] = \int_x^y r(z) \mu e^{-\mu(z-x)} dz + e^{-\mu(y-x)} \mathbb{E}_B[r(y+B)]$$

$$\geq \int_x^y r(x) \mu e^{-\mu(z-x)} dz + e^{-\mu(y-x)} \mathbb{E}_B[r(y+B)]$$

and, thus,

$$r(x) - \mathbb{E}_B[r(x+B)] \leq e^{-\mu(y-x)} r(x) - e^{-\mu(y-x)} \mathbb{E}_B[r(y+B)].$$

Note that if $r(y) - \mathbb{E}_B[r(y+B)] \leq 0$, then $r(x) - \mathbb{E}_B[r(x+B)] \leq 0$ (since $r(y) \geq r(x)$). Similarly, if $r(x) - \mathbb{E}_B[r(x+B)] \geq 0$, then $r(y) - \mathbb{E}_B[r(y+B)] \geq 0$. This directly gives the desired property, where $z_{\min} = \arg\inf\{h : r(h) \geq \mathbb{E}_B[r(h+B)]\}$.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Example 5.4: Suppose that $r(\cdot)$ is defined as follows:

$$r(x) = \begin{cases} r_1, & 0 < x \le a \\ \text{increasing and concave}, & a < x \le r_{\max} \\ \text{decreasing}, & x > r_{\max}. \end{cases}$$

In addition, assume that $r_\infty \ge r_1$. From the properties of $r(\cdot)$, it is obvious that $r(x) - \mathbb{E}_B[r(x+B)] \le 0$ as $x \in (0, a]$ and $r(x) - \mathbb{E}_B[r(x+B)] \ge 0$ as $x \ge r_{\max}$. Hence, $a \le z_{\min} \le r_{\max}$. Now take arbitrary x, y, with $a \le x < y \le r_{\max}$. First consider the following:

$$\begin{aligned} &\mathbb{E}_B[r(y+B) - r(x+B)] \\ &\quad = \int_0^{r_{\max}-x} (r(y+b) - r(x+b))\, dB(b) \\ &\qquad + \int_{r_{\max}-x}^{\infty} (r(y+b) - r(x+b))\, dB(b) \\ &\quad \le \int_0^{r_{\max}-x} (r(y) - r(x))\, dB(b) + 0 = (r(y) - r(x))B(r_{\max} - x), \end{aligned}$$

where we used that $r(\cdot)$ is concave on $[a, r_{\max}]$ and decreasing on $[r_{\max}, \infty)$ in the second step. Using the above, we obtain

$$\begin{aligned} &r(x) - \mathbb{E}_B[r(x+B)] \\ &\quad = r(y) - \mathbb{E}_B[r(y+B)] + r(x) - r(y) + \mathbb{E}_B[r(y+B)] - \mathbb{E}_B[r(x+B)] \\ &\quad \le r(y) - \mathbb{E}_B[r(y+B)] + (r(x) - r(y))(1 - B(r_{\max} - x)). \end{aligned}$$

As in Example 5.3, note that if $r(y) - \mathbb{E}_B[r(y+B)] \le 0$, then $r(x) - \mathbb{E}_B[r(x+B)] \le 0$ (since $r(y) \ge r(x)$). Similarly, if $r(x) - \mathbb{E}_B[r(x+B)] \ge 0$, then $r(y) - \mathbb{E}_B[r(y+B)] \ge 0$. Hence, Assumption 2.1 is satisfied, with $z_{\min} = \arg\inf\{h : r(h) \ge \mathbb{E}_B[r(h+B)]\}$.

Finally, note that Example 5.1 is just a special case (take $a = r_{\max} = 0$). However, we believe that Example 5.1 is a natural special case, which admits an easy verification of Assumption 2.1.

Example 5.5: Here we provide an example for which Assumption 2.1 is not satisfied. For simplicity, we choose specific values for some model parameters. A slightly

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

more general model could be constructed by leaving some parameters unspecified while leaving the structure unaltered.

Consider the following service rate function:

$$r(x) = \begin{cases} r_1, & 0 < x \le a \\ (x-a)c + r_1, & a < x \le r_{\max} \\ (\hat{h}-x)c + r_1, & r_{\max} < x \le b \\ r_2 < r_1, & x > b, \end{cases}$$

with $\hat{h} = 2r_{\max} - a$, implying that $r(\hat{h}) = r_1$. Also, suppose that $B = a/3$ with probability $\frac{1}{2}$ and $B = \hat{h} - a/3$ with probability $\frac{1}{2}$, and take $c > 3(r_1 - r_2)/a$. After some calculations, we derive that $dZ(x)/dx$ is strictly positive on $(0, a/3)$ and $((2a/3) + (r_1 - r_2)/c, (4a/3) - (r_1 - r_2)/c)$ and strictly negative on $(a/3, (2a/3) + (r_1 - r_2)/c)$ and $((4a/3) - (r_1 - r_2)/c, \infty)$. Clearly, $Z(\cdot)$ has two local minima and Assumption 2.1 is not satisfied in this case.

6. SOME EXAMPLES

In general, (18) is suitable for a numerical calculation of the optimal threshold. Also, the characteristics of $TH^{\pi_{\bar{x}}}$ described in Section 4 suggest another numerical calculation of this optimal value, for instance, using a bisection method. In this section we give some examples in which we obtain an analytically more tractable expression for the optimal throughput with the corresponding optimal threshold value. In Section 6.1 we consider a two-level service rate: The service rate at time t is r_1 when $V_t^{\pi} \le a$ and $r_2 < r_1$ when $V_t^{\pi} \ge a$ (see, for instance, [5]). In Section 6.2, we generalize the service rate to an arbitrary step function, but we restrict ourselves to exponential service requirements there.

In any case, if the greedy policy is not optimal, the optimal threshold value must satisfy (15); see Proposition 4.1. Define

$$z(\bar{x}) := [\mathbb{P}(V^{\pi_{\bar{x}}} = 0)^{-1} - W(\bar{x})\lambda Z(\bar{x})], \tag{22}$$

where $W(x) := 1 + \int_0^x K^*(y,0)\,dy$ represents a nonnormalized workload distribution. Using (18) and some straightforward manipulations, we can rewrite (15) as

$$\frac{\mathbb{P}(V^{\pi_{\bar{x}}} = 0)}{Z(\bar{x})}\, z(\bar{x}) = 0.$$

Note that both $\mathbb{P}(V^{\pi_{\bar{x}}} = 0) > 0$ and $Z(\bar{x}) > 0$ and finite. Thus, finding the extremes of $TH^{\pi_{\bar{x}}}$ reduces to solving $z(\bar{x}) = 0$.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

6.1. Two-Level Service Rate

Suppose that the service rate is specified as

$$r(x) = \begin{cases} r_1 & \text{for } 0 < x \le a \\ r_2 & \text{for } x > a, \end{cases}$$

where $0 < r_2 < r_1$. Define $\rho_i := \lambda\beta/r_i$, $i = 1,2$. Because the service-rate function is decreasing, we obtain from Example 5.1 that Assumption 2.1 is satisfied and a threshold policy is thus optimal. To determine the optimal threshold $\bar{x}$, we derive from Corollary 4.2 that we only need to consider $\bar{x} \le a$.

Fix some $\bar{x} \in [0,a]$. Using results of [1,9], the stationary workload distribution can be easily reduced to a more tractable expression. Let

$$H(x) := \beta^{-1} \int_0^x (1 - B(y))\, dy \tag{23}$$

be the stationary residual service requirement distribution with density $h(\cdot)$. For $x \le a$, $K(x,y) = \rho_1 h(x-y)$ and it is well known that $K^*(x,y) = \sum_{n=1}^{\infty} \rho_1^n h_n(x-y)$, where $h_n(\cdot)$ is the density of the n-fold convolution $H_n(\cdot)$ (see, e.g., [1,9]).

Now we determine the three elements on the right-hand side of (22) separately, after which we combine them to determine $z(\bar{x})$. First consider $\lambda Z(\bar{x})$. Using the definitions of $Z(\cdot)$ and $H(\cdot)$, respectively (2) and (23), yields

$$\begin{aligned} \lambda Z(\bar{x}) &= \lambda \int_{\bar{x}}^{a} \frac{1}{r_1} (1 - B(x - \bar{x}))\, dx + \lambda \int_a^{\infty} \frac{1}{r_2} (1 - B(x - \bar{x}))\, dx \\ &= \rho_2 + (\rho_1 - \rho_2) H(a - \bar{x}). \end{aligned} \tag{24}$$

Second, consider the nonnormalized workload distribution $W(\cdot)$. Interchanging integral and sum in addition to the above results, we immediately obtain for each $x \in [0, \bar{x}]$,

$$W(x) = 1 + \int_{0^+}^{x} \sum_{n=1}^{\infty} \rho_1^n h_n(y)\, dy = \sum_{n=0}^{\infty} \rho_1^n H_n(x). \tag{25}$$

Remark 6.1: Note that $W(\cdot)/W(a)$ is the steady-state workload distribution in a finite dam with speed r_1 and buffer size a. In the case $\rho_1 < 1$, it is an easy exercise to see that the Laplace–Stieltjes transform of $W(\cdot)$ provides the well-known Pollaczek–Khinchine formula. If $\rho_1 \ge 1$, $\int_0^a W(x)\, dx$ is still finite and a steady-state workload distribution exists (see, e.g., [9]). However, Cohen [4,5] describes a more elegant way to determine $W(\cdot)$ in that case.

Finally, the first term of (22) (i.e., the inverse of the normalizing constant $\mathbb{P}(V^{\pi_{\bar{x}}} = 0)$), is the most complicated one. Using the expression for the steady-state workload density in addition to the above results, we derive for $\bar{x} < x \le a$,

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

$$V^{\pi_{\bar{x}}}(x) = V^{\pi_{\bar{x}}}(\bar{x}) + \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left[\int_{\bar{x}}^{x} \rho_1 h(y)\,dy + \int_{\bar{x}}^{x}\int_{0^+}^{\bar{x}} \rho_1 h(y-u) \sum_{n=1}^{\infty} \rho_1^n h_n(u)\,du\,dy\right]$$

$$= \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left[W(\bar{x}) + \rho_1(H(x) - H(\bar{x})) + \rho_1 \int_{0^+}^{\bar{x}} (H(x-u) - H(\bar{x}-u))\,dW(u)\right]$$

$$= \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left[W(\bar{x}) + \rho_1 \int_{0^-}^{\bar{x}} (H(x-u) - H(\bar{x}-u))\,dW(u)\right],$$

where we used $W(0) = 1$ in the final step. Note that for $x > a$, $K(x,y) = \rho_2 h(x-y)$. Using similar arguments, we obtain for $x > a$,

$$V^{\pi_{\bar{x}}}(x) = V^{\pi_{\bar{x}}}(a) + \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left[\int_{a}^{x} \rho_2 h(y)\,dy + \int_{a}^{x}\int_{0^+}^{\bar{x}} \rho_2 h(y-u) \sum_{n=1}^{\infty} \rho_1^n h_n(u)\,du\,dy\right]$$

$$= \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left[W(\bar{x}) + \rho_1 \int_{0^-}^{\bar{x}} (H(a-u) - H(\bar{x}-u))\,dW(u) + \rho_2 \int_{0^-}^{\bar{x}} (H(x-u) - H(a-u))\,dW(u)\right].$$

By (25), we have $\rho_1 \int_0^{\bar{x}} H(\bar{x}-u)\,dW(u) = W(\bar{x}) - 1$. Letting $x \to \infty$ and some rewriting then yields

$$\mathbb{P}(V^{\pi_{\bar{x}}} = 0)^{-1} = \rho_2 W(\bar{x}) + (\rho_1 - \rho_2) \int_{0^-}^{\bar{x}} H(a-u)\,dW(u) + 1. \tag{26}$$

It is now easy to get $z(\bar{x})$. Substituting (24)–(26) into (22) gives

$$z(\bar{x}) = 1 + (\rho_2 - \rho_1)\left[W(\bar{x})H(a-\bar{x}) - \int_{0^-}^{\bar{x}} H(a-u)\,dW(u)\right]. \tag{27}$$

Summarizing, Corollary 4.2 implies that the greedy policy is optimal if and only if

$$\rho_2 - (\rho_2 - \rho_1)W(a) \geq 0.$$

Otherwise, $TH^* = \rho/(\rho_2 + (\rho_1 - \rho_2)H(a - x^*))$, with $\rho := \lambda\beta$ and x^* is a solution to $z(x^*) = 0$.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

In general, the convolution in (27) can only be determined numerically. However, if the service requirement follows a phase-type distribution, explicit expressions can be obtained. For instance, if $B(x) = 1 - e^{-\mu x}$ (see also Section 6.2), then after quite lengthy but standard calculations, it follows that for $\rho_1 \neq 1$,

$$z(\bar{x}) = 1 + \frac{\rho_2 - \rho_1}{\rho_1 - 1} e^{-\mu a}(e^{\mu \bar{x}} - e^{\mu \rho_1 \bar{x}}). \tag{28}$$

In the case $\rho_1 = 1$, we obtain

$$z(\bar{x}) = 1 - (\rho_2 - 1)\mu \bar{x} e^{-\mu(a - \bar{x})}.$$

6.2. Exponential Service Requirements

Suppose that the service requirements are exponentially distributed with mean $1/\mu$ (i.e., $1 - B(x) = e^{-\mu x}$). Then, for fixed $\bar{x}$, the steady-state workload density is given in [1, Cor. 7.1]

$$v^{\pi_{\bar{x}}}(x) = \frac{\lambda \mathbb{P}(V^{\pi_{\bar{x}}} = 0)}{r(x)} \exp\{-\mu x + \lambda R(x \wedge \bar{x})\}, \tag{29}$$

where $\mathbb{P}(V^{\pi_{\bar{x}}} = 0)$ follows from normalization. In this subsection, we also assume that the service rate is a step function. More specifically, let $r(x) = r_i$ for $x \in [a_{i-1}, a_i)$, $i = 1, \ldots, N$ (where $a_0 = 0$), and let $r(x) = r_{N+1} < r_N$ for $x \geq a_N$. Denote $\rho_i = \lambda/(\mu r_i)$, $i = 1, \ldots, N+1$, and assume for simplicity that $\rho_i \neq 1$.

Example 5.3 shows that Assumption 2.1 is satisfied and a threshold policy is thus optimal. By Corollary 4.2, either the greedy policy is optimal or the optimal threshold x^* is less than a_N. Let $\bar{x} \in [a_n, a_{n+1})$ for some $n \leq N - 1$. Next we consider each of the three elements of $z(\bar{x})$ separately, after which we combine them into an expression for x^* satisfying (15). However, for later use, we first define the following three constants. (In the sequel, we follow the convention that empty sums are equal to zero.):

$$\gamma_n = \exp\left\{\sum_{k=1}^{n}\left(\frac{\lambda}{r_k} - \frac{\lambda}{r_{k+1}}\right)a_k\right\},$$

$$C_n = \frac{1}{1 - \rho_1} + \sum_{k=1}^{n}\left(\frac{\rho_k}{\rho_k - 1} - \frac{\rho_{k+1}}{\rho_{k+1} - 1}\right)\gamma_{k-1} e^{-\mu(1-\rho_k)a_k},$$

$$D_n = \sum_{k=n+1}^{N}(\rho_{k+1} - \rho_k)e^{-\mu a_k}.$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

First, consider $\lambda Z(\bar{x})$. Using (2) and rewriting the integral, we obtain

$$\begin{aligned}\lambda Z(\bar{x}) &= \int_{\bar{x}}^{a_{n+1}} \frac{\lambda}{r_{n+1}}\, e^{-\mu(x-\bar{x})}\, dx + \sum_{k=n+1}^{N-1} \int_{a_k}^{a_{k+1}} \frac{\lambda}{r_{k+1}}\, e^{-\mu(x-\bar{x})}\, dx \\ &\quad + \int_{a_N}^{\infty} \frac{\lambda}{r_{N+1}}\, e^{-\mu(x-\bar{x})}\, dx \\ &= \rho_{n+1}(1 - e^{-\mu(a_{n+1}-\bar{x})}) + \sum_{k=n+1}^{N-1} \rho_{k+1}(e^{-\mu(a_k-\bar{x})} - e^{-\mu(a_{k+1}-\bar{x})}) \\ &\quad + \rho_{N+1}\, e^{-\mu(a_N-\bar{x})} \\ &= \rho_{n+1} + D_n e^{\mu\bar{x}}. \end{aligned} \tag{30}$$

Second, consider $W(\cdot)$ (i.e., the workload distribution "without normalization"). It is easily checked that the time to empty the system starting from a_i, $i = 1, \ldots, n$, in the absence of any arrivals (i.e., $R(a_i)$) equals $a_i/r_i + \sum_{k=1}^{i-1}(1/r_k - 1/r_{k+1})a_k$. Hence, for $x \in [a_i, a_{i+1})$, we can deduce that

$$\exp\{\lambda R(x)\} = \exp\left\{\frac{\lambda(x-a_i)}{r_{i+1}} + \lambda R(a_i)\right\} = \gamma_i \exp\left\{\frac{\lambda x}{r_{i+1}}\right\}. \tag{31}$$

Now, for $i = 1, \ldots, n$, using (29) and (31), after some standard algebra we obtain that

$$\begin{aligned} V^{\pi_{\bar{x}}}(a_i)/\mathbb{P}(V^{\pi_{\bar{x}}} = 0) &= 1 + \sum_{k=0}^{i-1} \int_{a_k}^{a_{k+1}} \frac{\lambda}{r_{k+1}}\, e^{-\mu x + \lambda R(x)}\, dx \\ &= 1 + \sum_{k=0}^{i-1} \frac{\rho_{k+1}}{\rho_{k+1} - 1}\, \gamma_k (e^{-\mu(1-\rho_{k+1})a_{k+1}} - e^{-\mu(1-\rho_{k+1})a_k}) \\ &= \frac{1}{1-\rho_1} + \sum_{k=0}^{i-1} \frac{\rho_{k+1}}{\rho_{k+1} - 1}\, \gamma_k\, e^{-\mu(1-\rho_{k+1})a_{k+1}} - \sum_{k=1}^{i-1} \frac{\rho_{k+1}}{\rho_{k+1} - 1}\, \gamma_k\, e^{-\mu(1-\rho_{k+1})a_k} \\ &= \frac{1}{1-\rho_1} + \sum_{k=1}^{i} \frac{\rho_k}{\rho_k - 1}\, \gamma_{k-1}\, e^{-\mu(1-\rho_k)a_k} - \sum_{k=1}^{i-1} \frac{\rho_{k+1}}{\rho_{k+1} - 1}\, \gamma_{k-1}\, e^{-\mu(1-\rho_k)a_k} \\ &= C_{i-1} + \frac{\rho_i}{\rho_i - 1}\, \gamma_{i-1}\, e^{-\mu(1-\rho_i)a_i}, \end{aligned}$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

where we used $\gamma_k e^{\mu\rho_{k+1}a_k} = \gamma_{k-1} e^{\mu\rho_k a_k}$ in the fourth equality. Thus, combining (29) and (31) with the above, we obtain, after similar manipulations,

$$V^{\pi_{\bar{x}}}(\bar{x}) = V^{\pi_{\bar{x}}}(a_n) + \mathbb{P}(V^{\pi_{\bar{x}}} = 0) \int_{a_n}^{\bar{x}} \frac{\lambda}{r_{n+1}} e^{-\mu x + \lambda R(x)}\, dx \tag{32}$$

$$= V^{\pi_{\bar{x}}}(a_n) + \mathbb{P}(V^{\pi_{\bar{x}}} = 0) \frac{\rho_{n+1}}{\rho_{n+1} - 1} \gamma_n (e^{-\mu(1-\rho_{n+1})\bar{x}} - e^{-\mu(1-\rho_{n+1})a_n})$$

$$= \mathbb{P}(V^{\pi_{\bar{x}}} = 0) \left[C_n + \frac{\rho_{n+1}}{\rho_{n+1} - 1} \gamma_n e^{-\mu(1-\rho_{n+1})\bar{x}} \right], \tag{33}$$

which completes the calculation of $W(\cdot)$ (since $W(\bar{x})\mathbb{P}(V^{\pi_{\bar{x}}} = 0) = V^{\pi_{\bar{x}}}(\bar{x})$).

For the first term on the right-hand side of (22) (i.e., $\mathbb{P}(V^{\pi_{\bar{x}}} = 0)^{-1}$), we use similar arguments as for the previous term. We first consider $V^{\pi_{\bar{x}}}(x)$ with $x > a_{n+1}$ and let $i = \arg\max\{a_i : a_i \le x\}$ be the largest a_i smaller than x. Using (29) and applying (31) to determine $\lambda R(\bar{x})$, we obtain after algebra similar to that above, that

$$V^{\pi_{\bar{x}}}(x) = V^{\pi_{\bar{x}}}(\bar{x}) + \gamma_n e^{\mu\rho_{n+1}\bar{x}} \mathbb{P}(V^{\pi_{\bar{x}}} = 0)$$
$$\times \left[\int_{\bar{x}}^{a_{n+1}} \frac{\lambda}{r_{n+1}} e^{-\mu y}\, dy + \sum_{k=n+1}^{i-1} \int_{a_k}^{a_{k+1}} \frac{\lambda}{r_{k+1}} e^{-\mu y}\, dy + \int_{a_i}^{x} \frac{\lambda}{r_{i+1}} e^{-\mu y}\, dy \right]$$
$$= \mathbb{P}(V^{\pi_{\bar{x}}} = 0) \left[C_n + \frac{\rho_{n+1}^2}{\rho_{n+1} - 1} \gamma_n e^{-\mu(1-\rho_{n+1})\bar{x}} + \gamma_n e^{\mu\rho_{n+1}\bar{x}} \right.$$
$$\left. \times \sum_{k=n+1}^{i} (\rho_{k+1} - \rho_k) e^{-\mu a_k} - \gamma_n \rho_{i+1} e^{-\mu x + \mu\rho_{n+1}\bar{x}} \right].$$

Thus, letting $x \to \infty$, we obtain the normalizing constant:

$$\mathbb{P}(V^{\pi_{\bar{x}}} = 0)^{-1} = C_n + \frac{\rho_{n+1}^2}{\rho_{n+1} - 1} \gamma_n e^{-\mu(1-\rho_{n+1})\bar{x}} + \gamma_n D_n e^{\mu\rho_{n+1}\bar{x}}. \tag{34}$$

The function $z(\bar{x})$ can now easily be rewritten into a more appealing expression. In particular, substituting (30), (34), and $W(\bar{x})$ resulting from (33) into (22) and some reordering of terms yields

$$z(\bar{x}) = (1 - \rho_{n+1}) C_n + D_n \gamma_n e^{\mu\rho_{n+1}\bar{x}}$$
$$- D_n e^{\mu\bar{x}} \left(C_n + \frac{\rho_{n+1}}{\rho_{n+1} - 1} \gamma_n e^{-\mu(1-\rho_{n+1})\bar{x}} \right)$$
$$= (1 - \rho_{n+1}) C_n - C_n D_n e^{\mu\bar{x}} - \frac{D_n \gamma_n}{\rho_{n+1} - 1} e^{\mu\rho_{n+1}\bar{x}}.$$

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

Solving $z(\bar{x}) = 0$ is thus remarkably simple in this case, since the variable $\bar{x}$ only appears in two of the exponents. Summarizing, we conclude that the optimal policy is of the threshold type where the optimal threshold value is given by the solution of $z(\bar{x}) = 0$. Moreover, $TH^* = \beta/Z(\bar{x})$, where $\lambda Z(\bar{x})$ is given in (30).

Remark 6.2: It is easily checked that, in case $N = 1$, the formula for $z(\bar{x})$ indeed reduces to (28).

7. CONCLUDING REMARKS AND FURTHER RESEARCH

In the present article we considered the problem of optimal admission control in a system with a workload-dependent service rate. We assumed that the service requirement only becomes known right after the decision to accept or reject jobs. Our objective was to find a policy that maximizes the long-run throughput. Under some assumptions (in particular, Assumption 2.1), we showed that a threshold policy for accepting jobs is optimal and derived a criterion for the optimal threshold value.

We note that our main assumption (i.e., Assumption 2.1) involves sufficient conditions for optimality of threshold policies. An interesting subject for further research is to examine the structure of the optimal policy when Assumption 2.1 is not satisfied.

Moreover, there are various interesting model variations. For instance, the analysis is significantly changed if information about the service requirement is available. In that case, the decision will not only depend on the workload level but also on the size of the job, yielding a two-dimensional state space. A characterization of the optimal policy in that model might be a subject of further study. We note that a threshold policy will not be optimal in general. However, in some special cases, such as for deterministic service requirements or decreasing service rate functions, the optimal policy continues to be of the threshold type.

Other model variations are systems where jobs can be partly accepted (or rejected). The simplest version concerns a model where an infinite amount of work becomes available at Poisson instants and the policy prescribes the amount of work to accept. In some sense, this model is related to the case $\lambda \to \infty$ in the model of the present article, which might be interpreted as an infinite supply of jobs and the policy prescribes the time to accept a new job. More interesting are systems where the supply of work is bounded by the service requirements of arriving jobs and the decision is the amount of work to accept. In that case, the state space is two dimensional and the action space is continuous. The structure of the optimal policy in the latter model is also left for future investigation.

Acknowledgments

The authors are grateful to Professor J.W.M. Bertrand for posing questions that motivated the present work and to Dr. Phil Whiting for stimulating discussions.

The work of the first author was financially supported by a research grant from Philips Research. Part of the research was conducted in the framework of Euro-NGI.

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

References

1. Bekker, R. (2005). Finite-buffer queues with workload-dependent service and arrival rates. *Queueing Systems* 50: 231–253.
2. Bekker, R., Borst, S.C., Boxma, O.J., & Kella, O. (2004). Queues with workload-dependent arrival and service rates. *Queueing Systems* 46: 537–556.
3. Bertrand, J.W.M. & van Ooijen, H.P.G. (2002). Workload based order release and productivity: A missing link. *Production Planning & Control* 13: 665–678.
4. Cohen, J.W. (1976). *On regenerative processes in queueing theory*. Lecture Notes in Economics and Mathematical Systems Vol. 121. Berlin: Springer-Verlag.
5. Cohen, J.W. (1976). On the optimal switching level for an $M/G/1$ queueing system. *Stochastic Processes and Their Applications* 4: 297–316.
6. Doshi, B.T. (1974). Continuous time control of Markov processes on an arbitrary state space. Ph.D. thesis, Cornell University, Ithaca, NY.
7. Doshi, B.T. (1977). Continuous time control of the arrival process in an $M/G/1$ queue. *Stochastic Processes and Their Applications* 5: 265–284.
8. Feinberg, E.A. & Kella, O. (2002). Optimality of D-policies for an $M/G/1$ queue with a removable server. *Queueing Systems* 42: 355–376.
9. Harrison, J.M. & Resnick, S.I. (1976). The stationary distribution and first exit probabilities of a storage process with general release rule. *Mathematics of Operations Research* 1: 347–358.
10. Van Ooijen, H.P.G. & Bertrand, J.W.M. (2003). The effects of a simple arrival rate control policy on throughput and work-in-process in production systems with workload dependent processing rates. *International Journal of Production Economics* 85: 61–68.
11. Ritt, R.K. & Sennott, L.I. (1992). Optimal stationary policies in general state space Markov decision chains with finite action sets. *Mathematics of Operations Research* 17: 901–909.
12. Ross, S.M. (1968). Arbitrary state Markov decision processes. *Annals of Mathematical Statistics* 39: 2118–2122.
13. Ross, S.M. (1970). Average cost semi-Markov decision processes. *Journal of Applied Probability* 7: 649–656.
14. Schäl, M. (1993). Average optimality in dynamic programming with general state space. *Mathematics of Operations Research* 18: 163–172.
15. Tijms, H.C. (1976). Optimal control of the workload in an $M/G/1$ queueing system with removable server. *Mathematische Operationsforschung und Statistik* 7: 933–944.

APPENDIX
Proof of Lemma 4.1

LEMMA 4.1: *For the derivative of* $TH^{\pi_{\bar{x}}}$*, we have*

$$\frac{d}{d\bar{x}} TH^{\pi_{\bar{x}}} = \lambda\beta\mathbb{P}(V^{\pi_{\bar{x}}} = 0)K^*(\bar{x},0)[1 - TH^{\pi_{\bar{x}}}Z(\bar{x})/\beta].$$

PROOF: We first consider $\mathbb{P}(V^{\pi_{\bar{x}}} = 0)$. Observe that the double integration in (16) can be equivalently expressed as follows: $\int_{x=\bar{x}}^{\infty}\int_{y=0}^{\bar{x}} = \int_{y=0}^{\bar{x}}\int_{x=y}^{\infty} - \int_{x=0}^{\bar{x}}\int_{y=0}^{x}$. Using the definition of K^*, interchanging integral and sum, and applying (17), we can write

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press

$$\int_0^{\bar{x}} K(\bar{x},y)K^*(y,0)\,dy = \sum_{n=1}^{\infty} \int_0^{\bar{x}} K(\bar{x},y)K_n(y,0)\,dy$$

$$= \sum_{n=0}^{\infty} K_{n+1}(\bar{x},0) - K_1(\bar{x},0)$$

$$= K^*(\bar{x},0) - K(\bar{x},0). \tag{A.1}$$

Taking the derivative of $\mathbb{P}(V^{\pi_{\bar{x}}} = 0)$ with respect to $\bar{x}$, we obtain from (16) and the reordering of integration that

$$\frac{d}{d\bar{x}}\mathbb{P}(V^{\pi_{\bar{x}}} = 0) = -\mathbb{P}(V^{\pi_{\bar{x}}} = 0)^2 \left(K^*(\bar{x},0) - K(\bar{x},0) + \int_{\bar{x}}^{\infty} K(x,\bar{x})K^*(\bar{x},0)\,dx - \int_0^{\bar{x}} K(\bar{x},y)K^*(y,0)\,dy\right)$$

$$= -\mathbb{P}(V^{\pi_{\bar{x}}} = 0)^2 K^*(\bar{x},0) \int_{\bar{x}}^{\infty} K(x,\bar{x})\,dx,$$

where we used (A.1) in the second step. Now invoking (18) and taking the derivative of $TH^{\pi_{\bar{x}}}$ with respect to $\bar{x}$ yields

$$\frac{d}{d\bar{x}} TH^{\pi_{\bar{x}}} = \lambda\beta\mathbb{P}(V^{\pi_{\bar{x}}} = 0)K^*(\bar{x},0) + \lambda\beta \frac{d}{d\bar{x}} \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\left(1 + \int_0^{\bar{x}} K^*(x,0)\,dx\right)$$

$$= \lambda\beta\mathbb{P}(V^{\pi_{\bar{x}}} = 0)K^*(\bar{x},0) \times \left[1 - \mathbb{P}(V^{\pi_{\bar{x}}} = 0)\int_{\bar{x}}^{\infty} K(x,\bar{x})\,dx\left(1 + \int_0^{\bar{x}} K^*(x,0)\,dx\right)\right]$$

$$= \lambda\beta\mathbb{P}(V^{\pi_{\bar{x}}} = 0)K^*(\bar{x},0)[1 - TH^{\pi_{\bar{x}}}Z(\bar{x})/\beta],$$

where the final step follows from (18) and the fact that $\lambda Z(\bar{x}) = \int_{\bar{x}}^{\infty} K(x,\bar{x})\,dx$. This completes the proof. ■

https://doi.org/10.1017/S0269964806060335 Published online by Cambridge University Press