ARTICLES

INDIVIDUAL EQUILIBRIUM AND LEARNING IN PROCESSOR SHARING SYSTEMS

EITAN ALTMAN

INRIA, Sophia-Antipolis Cedex, France

NAHUM SHIMKIN

Technion—Israel Institute of Technology, Haifa, Israel

(Received July 1996; revisions received September 1997, October 1997, November 1997; accepted December 1997)

We consider a processor-sharing service system, where the service rate to individual customers decreases as the load increases. Each arriving customer may observe the current load and should then choose whether to join the shared system. The alternative is a constant-cost option, modeled here for concreteness as a private server (e.g., a personal computer that serves as an alternative to a central mainframe computer). The customers wish to minimize their individual service times (or an increasing function thereof). However, the optimal choice for each customer depends on the decisions of subsequent ones, through their effect on the future load in the shared server. This decision problem is analyzed as a noncooperative dynamic game among the customers. We first show that any Nash equilibrium point consists of threshold decision rules and establish the existence and uniqueness of a *symmetric* equilibrium point. Computation of the equilibrium threshold is demonstrated for the case of Poisson arrivals, and some of its properties are delineated. We next consider a reasonable dynamic learning scheme, which converges to the symmetric Nash equilibrium point. In this model customers simply choose the better option based on available performance history. Convergence of this scheme is illustrated here via a simulation example and is established analytically in subsequent work.

The quality of service experienced by an individual customer in a shared service system often depends on the current load. Thus, in order to estimate the expected service quality, a user must consider not only the current load at her arrival instant, but also how this load might develop throughout her service period. We study in this paper the implications of this observation for self-optimizing users who consider joining the shared system. The model we consider is motivated in part by the following scenario.

Potential computer users, each requiring the use of a computer to execute a given job, arrive sequentially at a computer facility. Each user, upon arrival, may choose between the following two options: either connect to a central mainframe computer (MF), which is normally serving many users in parallel; or use a personal computer (PC). Each user is solely interested in minimizing her own service time (which coincides here with the sojourn time).

Service at the MF computer is performed according to the Processor Sharing discipline (e.g., Jaiswal 1982, Ramaswami 1984), where available computing power is equally divided among all users present. Consequently, the service rate to each user decreases as the load increases (although the *total* service rate may actually increase). An arriving user may observe the current load, namely the number of users already in MF. However, to evaluate her expected service time at MF she must take into account the possible load on this computer throughout her service time, which in turn is affected by the decisions of subsequent users. This leads us to consider the resulting decision problem in a game theoretic framework and to explore the Nash equilibrium solution for the resulting dynamic game.

Although we find it convenient to refer to this PC-MF application, the basic model is relevant to a wide range of application areas that involve shared service, ranging from transportation and recreation to computing and telecommunications. We mention here two relevant applications in the latter area:

(1) Consider a situation where users can communicate with each other either through a Local Area Network (LAN) or through the public network, e.g., by connecting to the telephone network via a modem, which is typically slower. However, the throughput available to each user on the LAN decreases as the total workload increases. This is especially the case in LANs where a single channel should be shared between all users, e.g., the FDDI (Fiber Distributed Data Interface). The LAN can thus be approximated by a processor sharing queue, whereas the public network can be viewed as assigning a private server to each session.

Subject classifications: Queues, optimization: processor sharing. Games, stochastic: dynamic equilibrium and learning.
Area of review: STOCHASTIC MODELS.

0030-364X/98/4606-0776 $05.00

(2) Consider a non-real-time application, such as data transfer, on an ATM (Asynchronous Transfer Mode) network (see ATM Forum 1996). ATM networks support both guaranteed services as well as best-effort services. Guaranteed services are CBR (Constant Bit Rate), in which a fixed amount of bandwidth is assigned to a session, and VBR (Variable Bit Rate), in which some average and peak bit-rates are assigned to a session. Best effort services are ABR (Available Bit Rate) and UBR (Unspecified Bit Rate); in both cases, some available bandwidth is shared among the connections that use these services. At a session level, ABR and UBR services can be approximated by a processor sharing queue, whereas CBR and VBR services can be approximated by a single server, dedicated for one session.

Our study focuses on two main issues. In the first part of the paper, we explore the properties of the Nash equilibrium solution. The main results here are the existence, uniqueness, and structural characterization of a *symmetric* equilibrium policy. The required analysis of the processor-sharing queue relies on stochastic coupling arguments rather than explicit calculations, which facilitates the consideration of both general interarrival times and of state-dependent total service rates.

The second issue concerns the descriptive power of the Nash equilibrium solution. The question arises as to what extent the Nash equilibrium actually describes the system operation under realistic conditions, and what mechanism might lead to this equilibrium. We consider a dynamic learning scenario, where users make simple decisions based on past performance statistics. Since the observed statistics depend on past decisions, this leads to a closed-loop adaptive decision problem. We discuss certain issues that are pertinent to the convergence of this scheme to that Nash equilibrium point and illustrate this convergence via a simulation experiment. Analysis and proof of convergence can be found in Altman and Shimkin (1997).

Dynamic control of queueing systems has been the subject of considerable research, and surveys can be found in Stidham (1985) and Walrand (1988). Social optimization of a processor-sharing queue has been considered in de Waal (1988). Individual optimality has been studied and compared to social optimality in various models, under the first-come-first-served (FCFS) service discipline; see Naor (1969), Yechiali (1972), Bartroli and Stidham (1992), and the above-mentioned surveys. Under a FCFS discipline the expected sojourn time of any single customer is completely determined by the queue length at her arrival, so the individually optimal decision policy is trivial in that case.

There has been some work on the Nash equilibrium concept for the study of individually optimal dynamic control of queueing systems; see Glazer and Hassin (1986) and Hassin and Haviv (1994). Game theoretical analysis has been applied to other queueing control problems—e.g., Kulkarni (1983), Lee and Cohen (1985), Bovopoulos and Lazar (1987), Shenker (1990), Hsiao and Lazar (1991), Altman (1992), Altman and Koole (1992), Altman and Hordijk (1995), and Shimkin and Shwartz (1993)—where the last five consider dynamic problems. Finally, results related to our work have been obtained in Assaf and Haviv (1990), Haviv (1991), and Xu and Shantikumar (1993).

The organization of the paper is as follows. The model is presented in Section 1. In Section 2 the Nash equilibrium is studied, assuming i.i.d. interarrival times, possibly state-dependent service rate, and exponential service requirements. It is shown that in any equilibrium point, the decision rule of each user is a threshold rule. Existence and uniqueness of a *symmetric* equilibrium point are then established, using certain monotonicity and continuity properties of the service time at MF. (The proofs of these results are established in the Appendix using stochastic-coupling arguments.) Section 3 concerns the actual calculation of the symmetric equilibrium. Formulas are derived for the case of Poisson arrivals and are illustrated by numerical examples. In Section 4 the (individually optimal) Nash solution is compared with the socially optimal one. Section 5 considers the proposed dynamic learning scheme. We close with some concluding remarks in Section 6.

1. THE MODEL

Consider a service system that consists of two service facilities, Q_{MF} and Q_{PC}. Customers (users) arrive at this system sequentially, with interarrival times that are independent and identically distributed and have finite mean λ^{-1}. Simultaneous arrivals are excluded. Each arriving customer observes the number of customers in Q_{MF} at her arrival instant, and should choose whether to join Q_{MF} or Q_{PC}.

We assume that Q_{MF} has a buffer size B, which may be finite or infinite. A customer that arrives when the buffer is full cannot be admitted and must turn to Q_{PC}.

The service at Q_{MF} is exponential with rate $\mu(x)$, where $x \geqslant 1$ is the number of customers in Q_{MF}. The service discipline is processor sharing, so the service intensity for each customer equals $\nu(x) \triangleq \mu(x)/x$. The alternative Q_{PC} offers a fixed expected service time θ^{-1}. In both queues, service commences immediately upon admission, so the sojourn time coincides with the service time.

We make the following assumptions on the service rate at Q_{MF}:

(i) $0 < \mu(x) \leqslant \mu_{\max}$ for every $x \geqslant 1$ (bounded service rate);
(ii) $\nu(x) = \mu(x)/x$ is strictly decreasing in x. Thus, the service rate applied to each customer decreases as the load increases.

Let $X(t)$ denote the number of customers at Q_{MF} at time t, and suppose that the system starts at $t = 0$ with initial state $X(0) = x_0$. Let T_k, $k \geqslant 0$, denote the arrival time of customer C_k, where $0 = T_0 < T_1 < T_2 < \dots$. Hence $t = 0$ is the arrival time of customer C_0.

Upon arrival, customer C_k observes the current queue length $X(T_k)$, and should then decide which queue to join. A randomized decision rule for customer C_k is therefore defined by $u_k = \{u_k(x), 0 \leqslant x < \mathsf{B}\}$, where $u_k(x) \in [0, 1]$ is the probability of joining Q_{MF} if a queue length $X(T_k) = x$ is observed. Let U denote the collection of such decision rules, and let $\pi = (u_0, u_1, \ldots) \in \Pi \triangleq U^\infty$ denote the vector of decision rules of all the customers, to which we refer as a *policy*.

Denote by W_k the service duration of customer C_k, and let $\bar{W}_k(x, \pi)$ be the expected value of W_k given that x customers are present at Q_{MF} upon arrival of C_k, and that all customers (including C_k) follow the policy π. Then

$$\bar{W}_k(x, \pi) = u_k(x)\bar{V}_k(x, \pi) + (1 - u_k(x))\theta^{-1},$$

where $\bar{V}_k(x, \pi)$ is the expected service duration of C_k at Q_{MF} under the same conditions. We observe that $\bar{V}_k$ depends on π through $\{u_l, l > k\}$, the decision rules of subsequent customers.

We assume that the customers are self-optimizing, so that each wishes to minimize her own service time. To this end, she should obviously evaluate her expected service time at the two queues, namely $\bar{V}_k(x, \pi)$ and θ^{-1}, and choose the lower one.

The dependence of $\bar{V}_k$ on the decisions of other users leads us to study this problem within a game theoretic framework.

2. NASH EQUILIBRIUM SOLUTION

We now consider the system as a noncooperative game in which each customer wishes to minimize her own expected service time. The main results of this section concern the characterization of the equilibrium points of this game.

For each policy $\pi = (u_0, u_1, \ldots)$, let π^{-k} denote the collection of all decision rules in π, excluding the u_k, and let $[\pi^{-k}|u'_k]$ be the policy which replaces u_k by u'_k.

Definition 1. *A decision rule u_k is an* optimal response *for C_k against a policy π if*

$$\bar{W}_k(x, [\pi^{-k}|u_k]) \leqslant \bar{W}_k(x, [\pi^{-k}|u'_k]) \quad \forall x \geqslant 0, u'_k \in U. \tag{1}$$

Definition 2. *A policy $\pi = (u_0, u_1, \ldots)$ is a Nash equilibrium policy if u_k is an optimal response for C_k against π, for every $k \geqslant 0$.*

Thus, in equilibrium no one can gain by a unilateral change of her decision rule. Since the queue length x is observed prior to decision, we require this to hold for every possible value of x.

In the sequel we shall take special interest in equilibrium policies that are *symmetric*, namely the decision rules of all customers are identical. Such policies are natural here, since the specifications of all customers are the same, and they all face the same decision problem. (In the terminology of Schelling 1960, symmetric equilibria are the natural candidates for the *focal equilibrium* of the game.)

We define next a special class of decision rules, namely threshold rules. For any $0 \leqslant q \leqslant 1$ and integer $L \geqslant 0$, the decision rule u is an $[L, q]$-threshold rule if

$$u(x) = \begin{cases} 1 & \text{if } x < L, \\ q & \text{if } x = L, \\ 0 & \text{if } x > L. \end{cases} \tag{2}$$

A customer who employs this rule joins Q_{MF} if the queue length x is smaller than L, while if $x = L$ she does so with probability q. Otherwise she joins Q_{PC}. An $[L, q]$ threshold rule will be denoted by $[L, q]$, or more compactly by $[g]$ with $g = L + q$. Note that $[L, 1]$ and $[L + 1, 0]$ are identical. When the buffer size is finite, any $[g]$ with $g > \mathsf{B}$ is equivalent to $[\mathsf{B}]$.

We now turn to the main results of this section.

Theorem 1. (i) *For any equilibrium policy $\pi^* = (u_0, u_1, \ldots)$, each decision rule u_k is a threshold rule.*

(ii) *A symmetric equilibrium policy $\pi^* = (u^*, u^*, \ldots)$ exists, is unique, and u^* is a threshold rule.*

The proof proceeds through some lemmas. The first two establish basic monotonicity and continuity properties of the service time in Q_{MF}.

Recall that $\bar{V}_k(x, \pi)$ denotes the expected service time at Q_{MF}.

Lemma 1. *For every policy π and $k \geqslant 0$, $\bar{V}_k(x, \pi)$ is strictly increasing in x. In fact, for every $x \geqslant 0$, $\bar{V}_k(x + 1, \pi) - \bar{V}_k(x, \pi) > \delta_x$ for some positive δ_x, independent of π.*

A detailed proof, based on stochastic coupling arguments, is given in the Appendix. The idea is simple. Consider two cases, the first when a customer (say C_0) enters at queue length x, and the second when she enters at $x + 1$. Since the queue lengths change by at most one customer at a time, the queue length in the second case will be higher than in the first one, until such a time τ (possibly infinite) when they coincide. From then on the queue lengths will remain equal, since the decision policies for incoming customers are the same in both cases. It then follows that the service rate applied to C_0 throughout her stay is higher in the first case than in the second, and strictly so up until τ. This implies a lower expected service time. □

Recall that $[g]$ stands for a threshold rule, and let $[g]^\infty$ denote the stationary policy $\pi = ([g], [g], \ldots)$. Thus $\bar{V}_k(x, [g]^\infty)$ is the expected service time of customer C_k if she joins Q_{MF} at queue length x, while all subsequent customers are using the threshold rule $[g]$.

Lemma 2. *For every $k \geqslant 0$ and $x \geqslant 0$, $\bar{V}_k(x, [g]^\infty)$ is:*

(i) *strictly increasing in $g \in [0, \mathsf{B}]$, and*

(ii) *continuous in $g \in [0, \mathsf{B}]$.*

The proof, again using stochastic coupling arguments, is presented in the Appendix. □

Lemma 3. *Let π be an arbitrary policy, and let U_k^* be the set of decision rules for C_k that are optimal against π.*

(i) *Any $u_k^* \in U_k^*$ is a threshold rule, with finite threshold, and is given by:*

$$u_k^*(x) = \begin{cases} 1 & \text{if } \bar{V}_k(x, \pi) < \theta^{-1}, \\ q_x & \text{if } \bar{V}_k(x, \pi) = \theta^{-1}, \\ 0 & \text{if } \bar{V}_k(x, \pi) > \theta^{-1}, \end{cases} \tag{3}$$

where $0 \leq q_x \leq 1$ is arbitrary. (Recall that θ^{-1} is the expected service time in Q_{PC}.)

(ii) *Consequently, the set U_k^* is given as follows. Let L^* be the smallest nonnegative integer such that $\bar{V}_k(L^*, \pi) \geq \theta^{-1}$. If $\bar{V}_k(L^*, \pi) = \theta^{-1}$, then $U_k^* = \{[L^*, q]: 0 \leq q \leq 1\}$. Otherwise, U_k^* consists of the single threshold rule $[L^*, 0]$.*

Proof. Assume that C_k observes x customers in Q_{MF} at her arrival. She should now choose to join either Q_{MF}, where her expected service time would be $\bar{V}_k(x, \pi)$, or Q_{PC}, where her expected service time is θ^{-1}. Obviously, the optimal decision is to choose the lower one and is thus given by (3). By Lemma 1, $\bar{V}_k(x, \pi)$ is strictly increasing in x, so that (3) is indeed a threshold rule. Finiteness of the threshold should be checked for $\mathsf{B} = \infty$; it follows from the easily verified fact that $\bar{V}_k(x, \pi) \to \infty$ as $x \to \infty$. Thus (i) is established, and (ii) follows immediately from (i) after noting again that $\bar{V}_k(x, \pi)$ is increasing in x. □

Proof of Theorem 1. (i) Let $\pi^* = (u_0, u_1, \ldots)$ be an equilibrium policy. By Definition 2 each u_k must be optimal against π^*, and from Lemma 3 it follows that u_k is a threshold rule.

(ii) By (i), a symmetric equilibrium policy π^* must consist of identical *threshold* policies, i.e., $\pi^* = [g]^\infty$ for some $g \in [0, \mathsf{B}]$. It remains to establish existence and uniqueness of a threshold $g^* \in [0, \mathsf{B}]$ such that $[g^*]$ is optimal for C_0 (hence, by symmetry, for any customer C_k) against $[g^*]^\infty$.

Define the point-to-set mapping $G^*: [0, \mathsf{B}] \to 2^{[0,\mathsf{B}]}$, which associates with every threshold $g \in [0, \mathsf{B}]$ the set of "optimal thresholds" against $[g]^\infty$, namely

$$G^*(g) = \{g' \in [0, \mathsf{B}]: [g'] \text{ is optimal against } [g]^\infty\}. \tag{4}$$

Thus, it is required to prove that G^* possesses a unique fixed point, i.e., a unique $g^* \in [0, \mathsf{B}]$ such that $g^* \in G^*(g^*)$. In essence, the required existence and uniqueness follow, respectively, from continuity and monotonicity properties of G^*. In fact, existence may be deduced by applying the Kakutani fixed point theorem, which is commonly used to establish existence of Nash equilibria; see, e.g., Başar and Olsder (1995). However, it will be more instructive to construct explicitly the graph of G^*.

By Lemma 3(ii), the set $G^*(g)$ can be expressed as follows. Let $\ell(g)$ denote the minimal integer $L \geq 0$ for which $\bar{V}_0(L, [g]^\infty) > \theta^{-1}$ (note that $\ell(g)$ is finite since $\bar{V}_0(x, [g]^\infty) \to \infty$ as $x \to \infty$). If $\ell(g) \geq \mathsf{B}$ then $G^*(g) = \{\mathsf{B}\}$; otherwise,

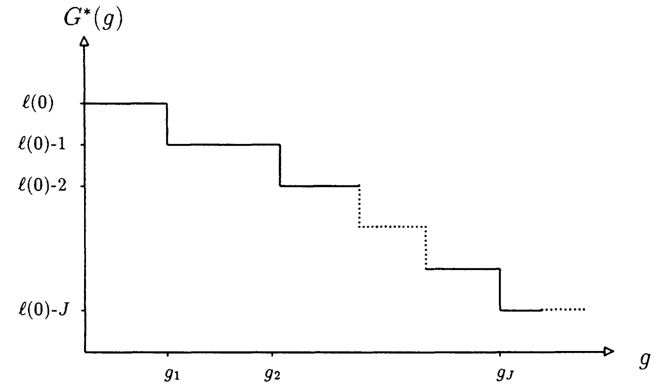

Figure 1. The optimal thresholds.

$$G^*(g) = \begin{cases} \{\ell(g) + q : 0 \leq q \leq 1\} & \text{if } \bar{V}_0(\ell(g), [g]^\infty) = \theta^{-1}, \\ \{\ell(g)\} & \text{if } \bar{V}_0(\ell(g), [g]^\infty) > \theta^{-1}. \end{cases} \tag{5}$$

We next argue that the graph of $G^*(g)$ has the "staircase" form depicted in Figure 1. Recall from Lemmas 1 and 2 that $\bar{V}_0(x, [g]^\infty)$ is strictly increasing in x and g and is continuous in g. This implies that, as g increases from 0 to B, the map $g \to \ell(g)$ is nonincreasing, piecewise constant, left continuous, and its downward jumps are of exactly one unit. Let $J \geq 0$ be the number of jumps of $\ell(\cdot)$, and, provided $J \geq 1$, let $0 < g_1 < \cdots < g_J$ be the jump points of $\ell(\cdot)$. By the above-mentioned properties of $\bar{V}_0(x, [g]^\infty)$ it also follows that $\bar{V}_0(\ell(g_j), [g_j]^\infty) = \theta^{-1}$ at each jump point g_j. Defining for notational convenience $g_0 = 0$ and $g_{J+1} = \infty$, $\ell(\cdot)$ is given by: $\ell(g) = \ell(0) - j$ for $g_j < g \leq g_{j+1}$, $j = 0, \ldots, J$. It then follows from (5) that (see Figure 1)

$$G^*(g) = \begin{cases} \{\ell(0)\} & \text{if } g < g_1, \\ [\ell(0) - j,\ \ell(0) - j + 1] & \text{if } g = g_j,\ 1 \leq j \leq J, \\ \{\ell(0) - j\} & \text{if } g_j < g < g_{j+1},\ 1 \leq j \leq J. \end{cases} \tag{6}$$

Obviously, the graph of G^* is intersected exactly once by the line with unit slope, say at point (g^*, g^*). It follows that g^* is the unique fixed point of $g \to G^*(g)$, and therefore $\pi^* \triangleq [g^*]^\infty$ is the unique symmetric equilibrium point. □

Remark. It is interesting to note that the equilibrium threshold g^* can be either an integer or a noninteger (corresponding to a deterministic or a randomized equilibrium policy), with neither case being generic. This observation, which is clearly illustrated in Figure 5, may be understood with the help of Figure 1, where the line of unit slope may intersect the graph of $G^*(g)$ in its horizontal (integer) or vertical (noninteger) part.

We close this section by pointing out some interesting generalizations.

1. For concreteness we used for the performance criterion the expected service time; however, our sample path proofs actually show that the waiting times in equilibrium are optimized in the sense of stochastic dominance, implying in particular that the cost could be defined as the expected value of any increasing function of the waiting time.

2. We can allow for a certain fraction (realized via random splitting) of the arrivals to be *uncontrolled*, in the sense they must join Q_{MF} whenever its buffer is not full. We shall find these useful in Section 5.

3. Lemma 1, and consequently part (i) of Theorem 1, hold even if the interarrival times are not identically distributed, and the expected service times in Q_{PC} are different for different customers.

3. COMPUTATION OF THE EQUILIBRIUM THRESHOLD

Given the existence and uniqueness of the symmetric equilibrium, we consider in this section the computation of the corresponding equilibrium threshold. In the following subsection we characterize this threshold in terms of the expected waiting times in Q_{PC} under threshold policies. The computation of the equilibrium threshold, and in particular of the random part q^*, requires in general the numeric solution of an implicit equation. We then specialize our discussion to the case of Poisson arrivals and constant total service rate at Q_{MF}. Closed-form solutions are derived for this case, and the results are illustrated through a numerical example.

Our computations involve the expected service times $\bar{V}_k(x, [L, q]^\infty)$ in Q_{MF}. By symmetry these do not depend on the customer index k, which will be omitted.

3.1. Characterization

Lemma 4. *The equilibrium threshold* $[L^*, q^*]$ *may be determined by the following procedure.*

(a) *If* $\bar{V}(\mathsf{B} - 1, [\mathsf{B}]^\infty) < \theta^{-1}$ *then* $L^* = \mathsf{B}$, $q^* = 0$. *Otherwise,*

(b) $L^* = \min\{L \geq 0 : \bar{V}(L, [L, 1]^\infty) > \theta^{-1}\}$.

(b1) *If* $\bar{V}(L^*, [L^*, 0]^\infty) \geq \theta^{-1}$, *then the equilibrium threshold is* $[L^*, 0]$.

(b2) *If* $\bar{V}(L^*, [L^*, 0]^\infty) < \theta^{-1}$, *then the equilibrium threshold is* $[L^*, q^*]$, *where* $0 < q^* < 1$ *is the unique solution of*

$$\bar{V}(L^*, [L^*, q^*]^\infty) = \theta^{-1}. \tag{7}$$

Proof. For (L^*, q^*) as defined above, it is required to show that $\pi^* := [L^*, q^*]^\infty$ is an equilibrium policy, namely that the threshold rule $[L^*, q^*]$ is an optimal response against π^*. Case (a) is straightforward and follows from Lemma 3(ii). As for (b), note first that L^* is well defined even for $\mathsf{B} = \infty$, since $\lim_{L\to\infty} \bar{V}(L, [L, 1]^\infty) = \infty$. In case (b1), Lemma 3 implies that $[L^*, 0]$ is an optimal response against $[L^*, 0]^\infty$ if (i) $\bar{V}(L^* - 1, [L^*, 0]^\infty) \leq \theta^{-1}$ and (ii) $\bar{V}(L^*, [L^*, 0]^\infty) \geq \theta^{-1}$. But (i) follows by the definition of L^* (note that $[L^*, 0]$ and $[L^* - 1, 1]$ are the same), and (ii) is the defining condition for case (b1). Concerning case (b2), note first that existence and uniqueness of the solution q^* to Equation (7) follow from the continuity and monotonicity of $\bar{V}$ in q. Now, Lemma 3 implies that $[L^*, q^*]$ is an optimal response against $[L^*, q^*]^\infty$ if (7) is satisfied. □

3.2. Some Explicit Calculations

We now specialize our discussion to the case of Poisson arrivals and constant total service rate at Q_{MF} ($\mu(x) \equiv \mu$). For this case we determine the service times $\bar{V}(L, [L, q]^\infty)$ and obtain the solution q^* of Equation (7).

Fix $[L, q]$, and define $\bar{V}(x) := \bar{V}(x, [L, q]^\infty)$. Then $\bar{V}(x)$, $0 \leq x \leq L$, is the solution of the following set of $L + 1$ linear equations:

$$\bar{V}(x) = \frac{1}{\alpha} + \frac{\mu x}{\alpha(x + 1)} \bar{V}(x - 1) + \frac{\lambda}{\alpha} \bar{V}(x + 1), \quad 0 \leq x \leq L - 2, \tag{8}$$

$$\bar{V}(L - 1) = \frac{1}{\alpha} + \frac{\mu(L - 1)}{\alpha L} \bar{V}(L - 2) + \frac{q\lambda}{\alpha} \bar{V}(L) + \frac{(1 - q)\lambda}{\alpha} \bar{V}(L - 1), \tag{9}$$

$$\bar{V}(L) = \frac{1}{\mu} + \frac{L}{L + 1} \bar{V}(L - 1), \tag{10}$$

where $\alpha = \mu + \lambda$. These equations follow from the memoryless property of the system, which implies that $V(x)$ equals the expected remaining service time of any customer present at queue length $x + 1$. Thus, $\bar{V}(x)$ equals the expected time till the next transition (α^{-1} in the first equations), plus the expected remaining service time after that transition. These equations can obviously be solved numerically for each given $[L, q]$; however, in order to obtain the optimal threshold in closed form we derive a more explicit solution. By (8), $\bar{V}(x)$ can be expressed as

$$\bar{V}(x) = a(x)\bar{V}(0) + b(x), \quad 0 \leq x < L, \tag{11}$$

where the coefficients $a(x)$ and $b(x)$ are obtained recursively by substituting (11) into (8), which yields

$$a(x + 1) = \frac{1}{\lambda}\left[(\mu + \lambda)a(x) - \mu \frac{x}{x + 1} a(x - 1)\right], \quad x \geq 1, \tag{12}$$

$$b(x + 1) = \frac{1}{\lambda}\left[(\mu + \lambda)b(x) - \mu \frac{x}{x + 1} b(x - 1) - 1\right], \quad x \geq 1, \tag{13}$$

with initial values $a(0) = 1$, $a(1) = (\mu + \lambda)/\lambda$, $b(0) = 0$, $b(1) = \lambda^{-1}$. Note that these coefficients *do not depend on* L *and* q. Next we obtain $\bar{V}(0, [L, q])$. By substituting $\bar{V}(L)$ from (10) into (9) and then substituting $\bar{V}(L - 2)$ and $\bar{V}(L - 1)$ from (11), we obtain after some algebraic manipulation

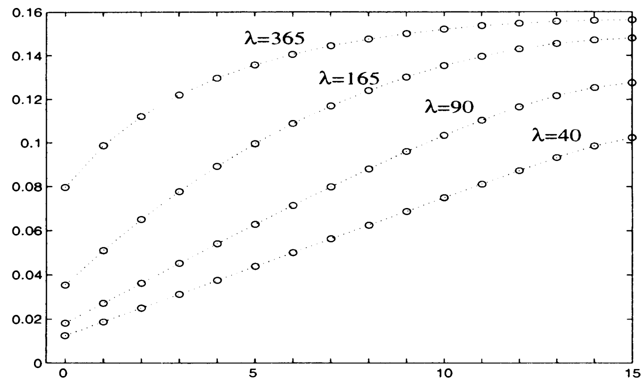

Figure 2. $\bar{V}(x[15, 1])$ as a function of x for different values of λ.

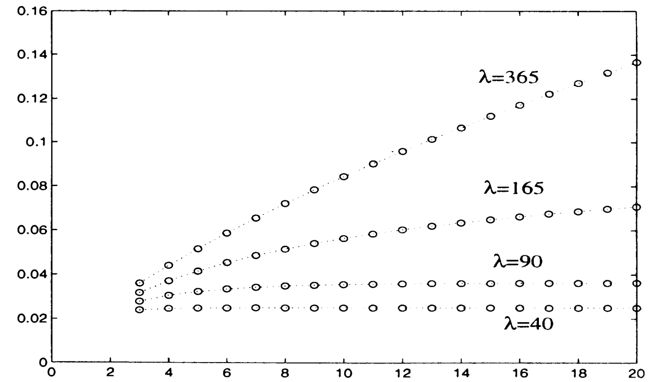

Figure 3. $\bar{V}(3, [L, 1])$ as a function of L for different values of λ.

$$\bar{V}(0) = \frac{\mu^{-1} - \frac{1}{L+1} b(L-1) - q^{-1}(b(L) - b(L-1))}{\frac{1}{L+1} a(L-1) + q^{-1}(a(L) - a(L-1))}. \tag{14}$$

$\bar{V}(x)$ can now be obtained from (11)–(13).

We can now calculate the equilibrium threshold. We first compute L^* by Lemma 4. If we are in case (a) or (b1) of Lemma 4 then $q^* = 0$. Otherwise, in case (b2), we compute $0 < q^* < 1$ as the unique solution of (7). Using (10), (11), and (14) to express $\bar{V}(L^*, [L^*, q]^\infty)$ as a function of q, rearranging and canceling terms we obtain a linear equation for q^*, whose solution is

$$q^* = C^{-1}\Big[\frac{a(L^*) - a(L^*-1)}{a(L^*-1)} \cdot \Big(\theta^{-1} - \mu^{-1} - \frac{L^*}{L^*+1} b(L^*-1)\Big) + \frac{L^*}{L^*+1}(b(L^*) - b(L^*-1))\Big], \tag{15}$$

where $C \triangleq \mu^{-1} - \theta^{-1}/(L^* + 1)$.

3.3. A Numeric Example

We illustrate these results and some properties of the relevant quantities through a numerical example. We consider the parameters $\mathsf{B} = \infty$, $\theta = 10$ and $\mu = 100$.

In Figure 2 we depict $\bar{V}(x, [15, 1])$ as a function of x for different values of λ. It can be seen that $\bar{V}(x, [15, 1])$ is bounded by 0.16 for any λ and x and that it increases with both x and λ. Indeed, for $\lambda \to \infty$ we have $\bar{V}(x, [15, 1]) \to 0.16$ since in that case there will always be 16 customers in the queue, and therefore the expected service time of each one of them is $[\mu/16]^{-1} = 0.16$.

In Figures 3 and 4 we depict $\bar{V}(3, [L, 1])$ and $\bar{V}(L, [L, 1])$ as a function of L for different values of λ. As expected, both are monotone increasing in L and in λ.

Finally, Figure 5 presents the equilibrium threshold $[L^*, q^*]$ as a function of λ. In that figure, a real number of the form 10.34 means $L^* = 10$ and $q^* = 0.34$. L^* can be obtained from Figure 4 as the first integer for which $\bar{V}(L, [L, 1])$ exceeds $\theta^{-1} = 0.1$ (see Lemma 4). Note that L^* is a decreasing function of λ and approaches $\mu/\theta = 10$. Indeed, for $\lambda \to \infty$ and threshold $[L, 0]$, the queue length approaches a constant L, and therefore the expected service time increases toward $(\mu/L)^{-1}$. Comparing with θ^{-1}, the individually optimal threshold decreases to $L^* = \mu\theta^{-1}$. Figure 5 also compares the equilibrium threshold to the socially optimal one. This will be discussed in the following section.

4. SOCIAL, INDIVIDUAL, AND NAÏVE OPTIMALITY

Depending on the nature and goals of the decision makers in the system, different solution concepts might be appropriate under different circumstances. Here we briefly compare the Nash equilibrium solution with the socially optimal one and also touch upon individual optimality under simplified (naïve) assumptions.

Consider the social optimization problem, where the goal is to minimize the expected average sojourn time per customer. Here each customer's decision is evaluated not only for its effect on her own performance, but also for its effect on the others (the externality cost). Observing that this social cost does not depend on the service discipline in Q_{MF} (PS, FCFS, etc.), the problem becomes a standard

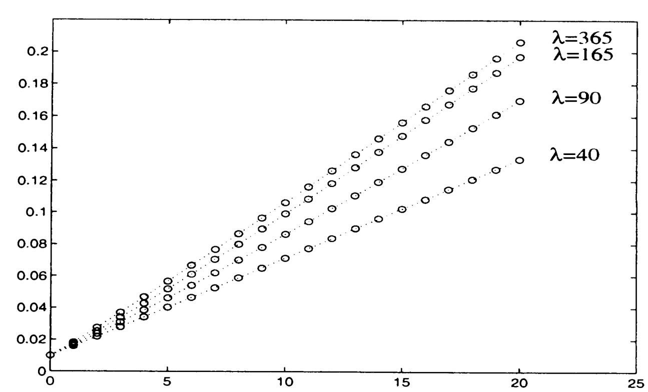

Figure 4. $\bar{V}(L, [L, 1])$ as a function of L for different values of λ.

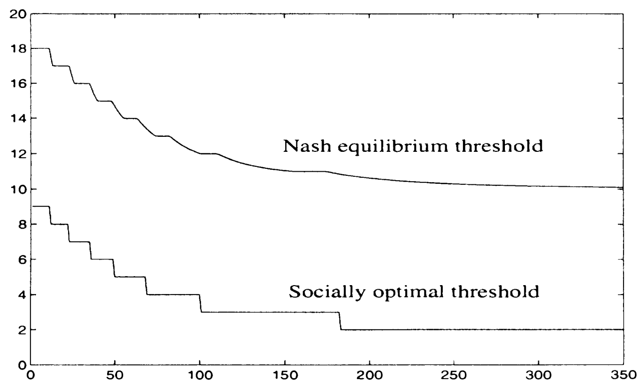

Figure 5. $g^* = L^* + q^*$ as a function of λ.

one in queueing control and is known to possess an optimal nonrandomized threshold policy (see, e.g., Stidham 1985). Let L_{so} be the value of this optimal threshold.

We assert that L_{so} is not larger than the Nash equilibrium threshold g^*. This is a direct consequence of the presence of externality costs in Q_{MF} and their absence in Q_{PC}. Indeed, if we increase the threshold beyond g^*, then any customer entering Q_{MF} (rather than Q_{PC}) at queue size larger than g^* increases both her own sojourn time, by definition of g^*, and the sojourn time of others in Q_{MF} by increasing the load there. Therefore such a threshold is socially worse than g^*.

Let us calculate L_{so} for the case of Poisson arrivals and constant service rate μ. Suppose that the threshold $[L, 0]$ is used. Denoting $\rho = \lambda/\mu$, and using standard results for finite-buffer queues, the expected sojourn time in the system is given by

$$E[W_L] = \frac{\rho^L}{1-\rho^{L+1}}[(1-\rho)\theta^{-1} - (L+1)\mu^{-1}] + \frac{\mu^{-1}}{1-\rho}. \tag{16}$$

L_{so} is obtained as the integer that minimizes $E[W_L]$.

As discussed above, it is seen from Figure 5 that the socially optimal threshold is indeed lower than the Nash equilibrium threshold. The threshold L_{so} decreases as λ (and thus ρ) grows, and it reaches asymptotically $L_{so} = 1$; our calculations show that $L_{so} = 1$ for $\lambda \geqslant 815$. We already saw that as λ grows, the Nash equilibrium threshold tends to $[10, 0]$, i.e., 10 times higher than the socially optimal threshold. The expected sojourn time in Q_{MF} is thus 10 times higher when the individual (equilibrium) criterion is used. It is obvious that individual users will have a strong incentive to deviate from the socially optimal policy in this case.

Finally, it should be of interest to consider the case of naïve individual optimality. Here we have self-optimizing customers, who make the simplifying assumption that the load perceived at their arrival will not change throughout their service periods. While false in general, this assumption may often be adopted in practice. Thus, a customer who sees x customers in Q_{PC} assumes that her expected sojourn time there will be $\mu(x+1)/(x+1)$ and compares this time to θ^{-1}. (Observe that for state-independent μ, this coincides with the case of a simple FCFS queue.) The resulting threshold is independent of the arrival rate λ, which indicates its deficiency for the processor sharing problem. In our example we obtain $L = \mu/\theta = 10$. It is not hard to verify that the naïve individual threshold will be lower than the individually optimal equilibrium threshold (as long as $\mu(x)$ is nondecreasing in x and there are no uncontrolled arrivals to Q_{PC}). This follows, since for a customer who enters just below the threshold, the assumption that the queue length will not change is the worst possible one.

5. LEARNING AND EQUILIBRIUM

The Nash equilibrium solution is defined from a normative viewpoint. It sets the "rational" choice for a sophisticated decision maker, who has global information about the system and can reason about the choices of others.

This section examines the relevance of the Nash equilibrium solution derived above from a different, descriptive viewpoint. We demonstrate that it naturally emerges in a dynamic learning scenario, which is a reasonable one for the system at hand. In the scenario considered here, customers base their decisions on statistical data that are accumulated by the server. No prior information is assumed regarding service in or arrivals to Q_{MF}, nor do customers employ game theoretic considerations to arrive at their decisions.

Other dynamic schemes with partially rational behavior that give rise to the Nash equilibrium have been extensively studied within the game theoretic and engineering literature, mostly for the case where the underlying game is static. See, e.g., Li and Başar (1987), Hsiao and Lazar (1991), Lakshmivarahan (1981), and Fudenberg and Levine (1998).

We consider the system of Section 1 and assume that the server monitors the average sojourn times of customers in Q_{MF}, depending on the queue length at their arrival instants. Newly arriving customers have access to these accumulated data, in addition to the queue length at their arrival, and may use it to assess their performance at Q_{MF} before deciding which queue to join.

The system starts at time 0 without any prior data or operating statistics regarding Q_{MF}. For every $t \geqslant 0$, let $N_t(x)$ denote the number of customers who had joined Q_{MF} at queue length x and already left it by time t. Let $\hat{V}_t(x)$ denote the (empirical) average service time of these $N_t(x)$ customers. A customer who arrives at time t may inspect the current vector $\hat{V}_t$.

Consider then a customer who arrives at time t, observes X_t customers in Q_{MF}, and has to choose between Q_{MF} and Q_{PC}. A natural decision rule for a customer who seeks to minimize her service time is:

$$\text{join } Q_{MF} \text{ if } \hat{V}_t(X_t) < \theta^{-1}, \text{ and join } Q_{PC} \text{ otherwise.} \tag{17}$$

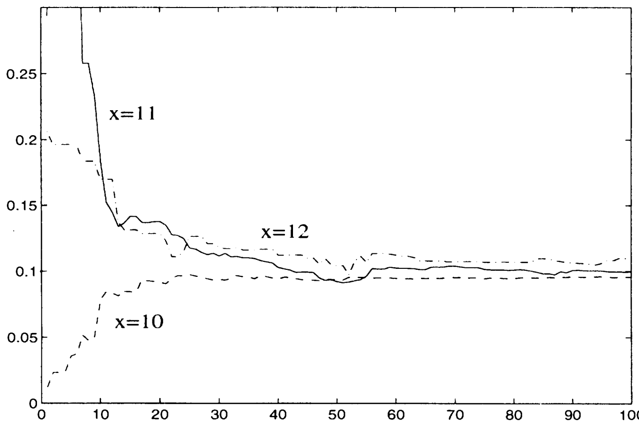

Figure 6. $\hat{V}_t(x)$ as a function of time.

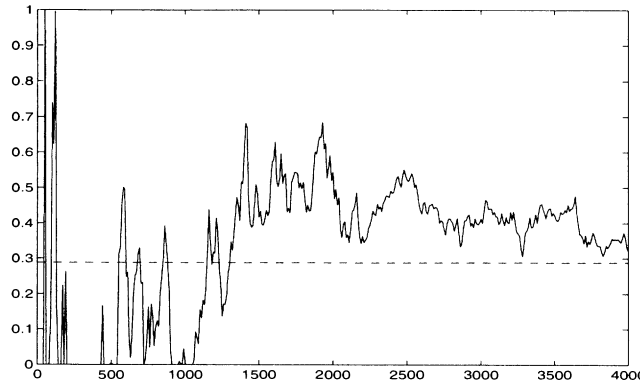

Figure 7. $u_t(11)$, the entry probability at queue size $x = 11$ as a function of time.

Under some mild modifications, it can be established that this learning decision rule, and consequently the relevant system performance measures, converges to the Nash equilibrium solution. An essential requirement for convergence is the existence of some fraction of *uncontrolled arrivals*, which always decide to enter Q_{MF}, and thus make sure that learning will continue at all relevant queue sizes. For further details and convergence analysis, which relies on the theory of the stochastic approximations algorithm, the readers are referred to Altman and Shimkin (1997). Further analysis of some related schemes may be found in Buche and Kushner (1998).

We illustrate the learning behavior using a simulated example. This simulation was performed with a constant service rate $\mu = 100$ and buffer size $B = 20$ at Q_{MF}, $\theta^{-1} = 0.1$, and Poisson arrivals with rate $\lambda = 120$, of which $\lambda_c = 100$ are controlled arrivals and $\lambda_u = 20$ are uncontrolled. The Nash equilibrium threshold for these parameters equals $g^* = 11.288$, as calculated numerically using the obvious extension of Equations (8) to include uncontrolled arrivals. Furthermore, to create complete correspondence with *randomized* Nash policies we incorporated a small uncertainty interval in the decision rule (17), meaning that when the difference $\hat{V}_t(X_t) - \theta^{-1}$ is very small ($\epsilon = 0.001$ was taken here), controlled customers take a randomized decision with the probability of entering Q_{MF} decreasing proportionally to that difference. For the above choice of parameters, randomized decisions will be taken whenever $\hat{V}$ is in the range $(\theta^{-1} - \epsilon, \theta^{-1} + \epsilon) = (0.099, 0.101)$.

The equilibrium values $\bar{V}(x, [g^*]^\infty)$ for queue lengths $x = 10$, 11 and 12 are (0.9750, 0.1000, 0.1030) respectively. The equality $\bar{V}(11, [g^*]) = \theta^{-1}$ is, of course, expected because of the equilibrium threshold value.

The system was simulated over a time interval of 4000 time units (where each time unit corresponds to 120 expected external arrivals). The behavior of the average waiting time statistics $\hat{V}_t(x)$ during the initial 100 time units is depicted in Figure 6, for entries at queue lengths $x = 10$, 11, and 12. We can see that fairly large average waiting times were obtained initially. This was caused mainly by the choice of null initial conditions for $\hat{V}_t$, which encouraged customers to initially join Q_{PC} at all queue sizes. However, as soon as these unfavorable statistics were observed, controlled arrivals temporarily stopped even at relatively low queue sizes, which caused low waiting times for those who did enter, and balanced the average statistics.

Beyond that initial period, reasonable convergence may be observed. The variable $\hat{V}_t(10)$ was the last to leave the band $\theta^{-1} \pm \epsilon$, around time 130, and beyond that time all arrivals at queue sizes other than $x = 11$ behaved exactly according to the Nash policy. As to the latter ($x = 11$), we can see in Figure 7 the value of the randomized decision at queue length $X_t = 11$ over the full simulation period. The dashed line denotes the equilibrium value $q^* = 0.288$. Good agreement is seen here also.

6. CONCLUDING REMARKS

The fundamental issue that was considered in this paper is the effect of expected load buildup on individual user decisions, and consequently on system performance, in shared service facility. Assuming symmetric users, we have shown the existence of a unique Nash equilibrium point and how this equilibrium might emerge as a result of a simple learning scenario.

We conclude by pointing to some issues that deserve further investigation. An important extension of the model would be to the case of multiple customer classes (for example, corresponding to different cost parameters). In this case, customers of different classes are expected to employ different decision rules, and the question of uniqueness of the equilibrium policy becomes multidimensional and harder to resolve. See Ben-Shahar et al. (1998) for some results regarding this problem.

The learning framework suggested here seems quite general and applicable to other similar models. In the present context, an important extension would be to the case of user-based learning. We have assumed that keeping record of the performance statistics is handled by a central entity (the server), which monitors all customers, and makes this information available to all. In certain situations it might be more appropriate to consider learning by (a finite number of) users who repeatedly use the same service facility, and each one learns out of its own personal experience. Convergence of this distributed learning scheme is currently under investigation. More complicated distributed learning scenarios that incorporate partial information sharing between users (see, e.g., Kushner and Yin 1987) might be similarly considered.

ACKNOWLEDGMENT

This research was partially supported by a joint grant from the Israeli Ministry of Science and the French Ministry of Research and Technology.

APPENDIX

The appendix can be found at the *Operations Research* Home Page: http://opim.wharton.upenn.edu/~harker/opsresearch.html in the Online Collection.

REFERENCES

ATM Forum Technical Committee. 1996. *Traffic Management Specification*. Version 4.0, April.

Altman, E. 1994. Flow Control Using the Theory of Zero-Sum Markov Games. *IEEE Trans. Automat. Contr.* **39,** 4, 814–818.

Altman E. and A. Hordijk. 1995. Zero-Sum Markov Games and Worst-Case Optimal Control of Queueing Systems. Invited Paper. *QUESTA* **21** (special issue on optimization of queueing systems), S. Stidham (Ed.), 415–447.

Altman E. and G. Koole. 1993. Stochastic Scheduling Games with Markov Decision Arrival Processes. *J. Comput. Math. Appl.* (third special issue on "Differential Games"), 141–148.

Altman, E. and N. Shimkin. 1997. Learning and Equilibrium in Processor Sharing Systems. Report No. EE-1113, Electrical Engineering, Technion.

Assaf, D. and M. Haviv. 1990. Reneging from Processor Sharing Systems and Random Queues. *Math. O. R.,* **15,** 129–138.

Başar, T. and G. J. Olsder. 1995. *Dynamic Noncooperative Game Theory*. Second Edition. Academic Press, London.

Benveniste, A., M. Metivier, and P. Priouret. 1990. *Adaptive Algorithms and Stochastic Approximations*. Springer Verlag, Berlin.

Bovopoulos, A. D. and A. A. Lazar. 1987. Decentralized Algorithms for Optimal Flow Control. *Proc. 25th Allerton Conf. on Communications, Control and Comput.* University of Illinois at Urbana-Champaign, 979–987.

Bartroli, M. and S. Stidham. 1992. Admission to a General Stochastic Congestion System: Comparison of Individually and Socially Optimal Policies. *Queueing and Related Models*. U. N. Bhat and I. V. Basawa (Eds.). Oxford University Press, New York, 301–324.

Ben-Shahar, I., A. Orda, and N. Shimkin. 1998. Dynamic Service Sharing with Heterogeneous Preferences, preprint, submitted to *Queueing Systems*, August 1998.

Buche, R. and H. J. Kushner. 1998. Stochastic Approximation and User Adaptation in a Competitive Resource Sharing System. Technical Report No. LCDS-98-9, Brown University, April 1998.

de Waal, P. 1988. Stochastic Control of a Processor Sharing Queue with Partial State Information. *Proc. of the 27th IEEE Conf. Decision and Control*, Austin, TX, 1135–1136.

Fudenberg, D. and D. K. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Boston.

Glazer, A. and R. Hassin. 1986. Stable Priority Purchasing in Queues. *O. R. Lett.* **6,** 285–288.

Hassin, R. and M. Haviv. 1994. Equilibrium Strategies and the Value of Information in a Two Line Queueing System with Threshold Jockeying. *Commun. Statis.—Stochastic Models* **10,** 415–435.

Haviv, M. 1991. Stable Strategies for Processor Sharing Systems. *Euro. J. Opns. Res.* **52,** 103–106.

Hsiao, M. T. and A. A. Lazar. 1991. Optimal Decentralized Flow Control of Markovian Queueing Networks with Multiple Controllers. *Performance Eval.* **13,** 181–204.

Jaiswal, N. K. 1982. Performance Evaluation Studies for Time-Sharing Computer Systems. *Performance Eval.* **2,** 223–236.

Kulkarni, V. G. 1983. A Game Theoretic Model for Two Types of Customers Competing for Service. *O. R. Lett.* **2,** 119–122.

Kushner, H. J. and G. Yin. 1987. Asymptotic Properties of Distributed and Communicating Stochastic Approximations Algorithms. *SIAM J. Control Optimiz.* **25,** 1266–1290.

Lakshmivarahan, S. 1981. *Learning Algorithms Theory and Applications*. Springer, NY.

Lee, H. L. and M. A. Cohen. 1985. Multi-Agent Customer Allocation in a Stochastic Service System. *Mgmt. Sci.* **31,** 752–763.

Li, S. and T. Başar. 1987. Distributed Algorithms for the Computation of Noncooperative Equilibria. *Automatica* **23,** 523–533.

Lippman, S. A. and S. Stidham. 1977. Individual Versus Social Optimization in Exponential Congestion Systems. *Opns. Res.* **25,** 233–247.

Naor, P. 1969. On the Regulation of Queue Size By Levying Tolls. *Econometrica* **37,** 15–24.

Ramaswami, V. 1984. The Sojourn Time in the GI/M/1 Queue with Processor Sharing. *J. Appl. Prob.* **21,** 437–442.

Schelling, T. C. 1960. *The Strategy of Conflict.* Harvard University Press, Cambridge, MA.

Shenker, S. 1990. Efficient Network Allocation with Selfish Users. *Performance '90—Proc. 14th IFIP WG 7.3 International Symposium on Comput. Performance Modeling*. Edinburgh, 279–285.

Shimkin, N. and A. Shwartz. 1993. Guaranteed Performance Region in Markovian Systems with Competing Decision Makers. *IEEE Trans. Automat. Contr.* **90,** 84–95.

Stidham, S. 1985. Optimal Control of Admission to a Queueing System. *IEEE Trans. Automat. Contr.* **30,** 705–713.

Walrand, J. 1988. *An Introduction to Queueing Networks*. Prentice-Hall, NJ.

Xu, S. H. and J. G. Shantikumar. 1993. Optimal Expulsion Control—A Dual Approach to Admission Control of an Ordered Entry. *Oper. Res.* **41,** 1137–1152.

Yechiali, U. 1972. Customers' Optimal Joining Rules for the GI/M/s Queue. *Mgmt. Sci.* **18,** 434–443.

LINKED CITATIONS

References

The Regulation of Queue Size by Levying Tolls
P. Naor
Econometrica, Vol. 37, No. 1. (Jan., 1969), pp. 15-24.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28196901%2937%3A1%3C15%3ATROQSB%3E2.0.CO%3B2-3