Throughput-Delay Tradeoff

with Impatient Arrivals

Rachel Levi and Adam Shwartz

Electrical Engineering
Technion—Israel Institute of Technology
Haifa 32000, Israel

April 14, 1995

Abstract

Consider a queue where the (single) server is free to choose between two service modes (fast and slow), while arrivals may decline to join the queue if the queue is too long. The server is interested in Throughput—Cost trade-off: using the fast service is expensive, while slow service is free. Can the server guarantee a long run, time average Cost and Throughput that approach a specific, desired region, regardless the behavior of the customers?

We characterize explicitly all convex approachable sets, as well as all non-approachable sets, as follows. There is a polygon so that:

Any convex set that intersects two sides of the polygon is approachable,

Any set whose closure fails to intersect these sides is not approachable.

The mathematical model that describes such systems is a Semi-Markov decision process with two independent decision makers, and vector payoff, where each coordinate is a long-run average. Under very mild conditions we obtain in [3] a general theory of approachability for such processes, which is applicable here.

1 Introduction

In modern communication networks, users not only compete for resources, but can also exert some control over the network (e.g. by deciding if and when to send large files, or submit heavy jobs.) Users also often have a choice between several networks, so that if the performance of one is not satisfactory, they can connect to another. Network management, on the other hand, may control the network to some extent in real time.

In this context, the classical building block of communication networks—the single server queue—gives rise to the following model. Consider a single server queue ($M/\bullet/1$ type) with a Poisson arrival process and two allowed service rates (fast and slow.) The customers are impatient: they may decline to join the queue if the server is busy or if the queue is too long, although they always join if the queue is empty. Fig. 1 illustrates the structure of the queuing system.

The server is interested in Throughput—Cost trade-off: using the fast service mode entails a fixed cost per customer, and an additional cost per unit of fast service time, while slow service is free. In a worst-case scenario, the question is: can the server guarantee a long run, time average Cost and Throughput that approach a specific, desired region, regardless the behavior of the customers? For example, can the server guarantee Throughput that exceeds a certain threshold while the Cost is kept below another threshold, no matter what the customers do? Let us call a set for which this can be achieved an "approachable set" (cf. Figure 2 below).

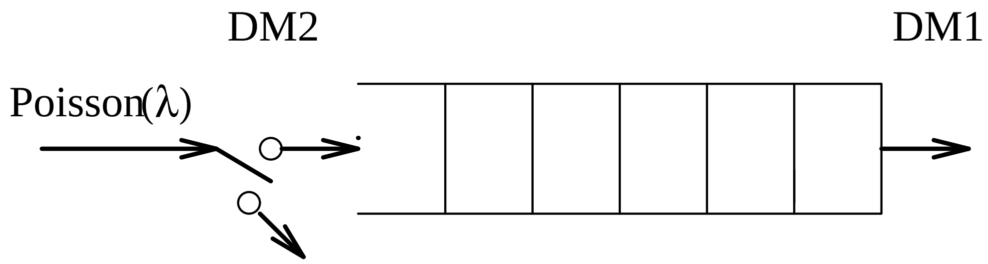

Figure 1: The Model

In this paper we restrict to exponential service times, although the general case leads to similar results: see [3]. We give an explicit characterization of all convex approachable sets, as well as a characterization of all non-approachable sets. This characterization is surprisingly simple. There is a polygon (with 4 vertices. In general there may be 5: see [3]) so that

Any convex set that intersects two specified sides of the polygon is approachable,

Any set whose closure fails to intersect the same sides is not approachable.

We provide closed-from (algebraic) expressions for the vertices (and hence a complete description of the polygons.) In addition, for each approachable set we provide an algorithm for computing a policy that guarantees that the long-time performance indeed approaches the desired set. In the case of exponential service we can compute a policy explicitly.

This model is a special case of a dynamic system with many decision makers and a vector of performance measures. Such systems appear in other engineering and communication systems such as many-user computers, and Data and Cellular communication networks. This approach to the analysis of performance regions of Markovian systems was proposed in Shimkin and Shwartz [2], where the discrete-time case was analyzed. Shimkin and Shwartz extended the idea of approachability due to Blackwell [1] from repeated games to stochastic games.

The theoretical foundation for the models we investigate is a "worst case scenario," where a single decision maker (called DM1) competes against all other decision makers, who have their own goals, which may contradicts those of DM1. DM1 is interested in a vector of time averaged performances, and he defines a required performance region. Can he guarantee a time averaged performance that reaches that desired set? What policy should he employ?

The mathematical model, introduced here to describe such systems is a Semi-Markov decision process with two independent decision makers, and with a vector payoff, where each coordinate is of the long-run average type. Under very mild conditions we obtain in [3] a general theory of approachability. This theory includes necessary and sufficient conditions for a convex set to be approachable by DM1, sufficient conditions for any set to be approachable by DM1 as well as an "approaching" policy that achieves the desired performance when this is possible. This theory is the basis for the analysis of our queueing models, and includes the results of [1] and [2] as a special case.

2 The model

Consider a single queue with Poisson arrivals with rate λ, where arriving customers may decide, upon arrival, not to join the queue. However, they always join if the queue is empty, and if they join they stay until served. Service is non-preemptive and first come first served (although the latter is immaterial.) Service times are independent and exponential, and the server can choose between fast service with parameter μ_s and slow service with parameter μ_f. We make the assumptions (the first is relaxed in [3]):

- Service mode can be changed only when the queue is empty,
- The queue was initially empty, and $\lambda < \mu_s < \mu_f$.

We define a stage to be any instant when a customer arrives to the queue, or a service was completed. For each stage m, we define the corresponding state of the queue x_m to be the pair $x_m = (i_m, j_m)$, where $i_m \in \{0, 1, 2, \ldots\}$ denotes the number of customers in the system including the customer who is currently being served and $j_m \in \{empty, fast, slow\}$ denotes the service mode. (The system is at service mode *empty* when there are no customers in the system). We also denote by t_m the m-th transition time (i.e. from stage m to stage m+1), and by (a_m^1, a_m^2) the choices made by the server and the customer at stage m, respectively.

At each stage, we define the 'single stage throughput' and 'single stage cost' as follows: The single stage throughput Λ_m is the probability that the next decision instant will occur when a customer leaves the system, given the current state x_m, decisions (a_m^1, a_m^2) and transition time t_m. Equivalently, it is the probability a service will be completed at stage m+1 given (x_m, a_m^1, a_m^2, t_m). This probability does not dependent on the past behavior of either the server or the customers, nor on the initial state (see [3]). Thus we can write:

$$\Lambda_m = \Pr(\text{a service will be completed at stage m+1}). \tag{1}$$

The single stage cost C_m consists of a fixed cost, namely $c_1 \geq 0$, for each customer who left the system right after being served with the fast service, and $c_2 \geq 0$, for each unit time of using the fast service mode. Thus

$$C_m \equiv 1\{\text{a fast service was employed at stage m}\} \cdot (c_1 \Lambda_m + c_2 t_m) \tag{2}$$

where $1\{\}$ is the indicator function, i.e., $1\{A\} = 1$ iff the statement A is satisfied, and otherwise the value of the function is zero.

The performance criteria, i.e. Time-averaged-Throughput and Cost are:

$$\bar{\Lambda}_n = \frac{\sum_{m=0}^{n-1} \Lambda_m}{\sum_{m=0}^{n-1} t_m}, \quad \text{and} \quad \bar{C}_n = \frac{\sum_{m=0}^{n-1} C_m}{\sum_{m=0}^{n-1} t_m} \ . \tag{3}$$

Let $r_m = (C_m, \Lambda_m)$ be the vector payoff and $\bar{r}_m = (\bar{C}_m, \bar{\Lambda}_m)$ its time average.

The dynamic model is a Semi-Markov Decision Process with two independent Decision Makers. For simplicity we denote the server as DM1 (Decision Maker 1) and the customers as DM2 (Decision Maker 2). Define $a_n^1 \in \{fast, slow\}$ the service mode chosen by DM1 at stage n,
$a_n^2 \in \{0, 1\}$ denotes DM2's action at stage n,
where $a_n^2 = 0$ means that he left the queue and $a_n^2 = 1$ means he joined it.
$\mu_{j_n} \in \{\mu_f, \mu_s\}$ is the service parameter at stage n.
The transition probabilities are:

$$p(i_{n+1} = l, j_{n+1} = k \mid i_n, j_n, a_n^1, a_n^2)$$

$$= \begin{cases} 1 & i_n = 0, j_n = empty, l = 1, k = a_n^1 \\ \mu_{j_n}/(\lambda + \mu_{j_n}) & i_n = 1, j_n, l = 0, k = empty \\ (1 - a_n^2)\lambda/(\lambda + \mu_{j_n}) & i_n = 1, j_n, l = 1, k = j_n \\ a_n^2\lambda/(\lambda + \mu_{j_n}) & i_n = 1, j_n, l = 2, k = j_n \\ \mu_{j_n}/(\lambda + \mu_{j_n}) & i_n > 1, j_n, l = i_n - 1, k = j_n \\ (1 - a_n^2)\lambda/(\lambda + \mu_{j_n}) & i_n > 1, j_n, l = i_n, k = j_n \\ a_n^2\lambda/(\lambda + \mu_{j_n}) & i_n > 1, j_n, l = i_n + 1, k = j_n \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

The transition times have the following distribution function (see [3]):

$$F_{x_n, x'}(t \mid a_n^1, a_n^2)$$

$$= \begin{cases} 1 - \exp(-\lambda t) & x_n = (0, empty), x' = (1, y) & y \in \{fast, slow\} \\ 1 - \exp(-(\lambda - \mu_f)t) & x_n = (l, fast), x' = (l \pm \{0, 1\}, fast) & l = 1, 2, \ldots \\ 1 - \exp(-(\lambda - \mu_s)t) & x_n = (l, slow), x' = (l \pm \{0, 1\}, slow) & l = 1, 2, \ldots \end{cases} \tag{5}$$

By (1)–(5), the Time-averaged-Throughput and Cost satisfy

$$\Lambda_n = \begin{cases} 0 & x_n = (0, empty) \\ \mu_f/(\lambda + \mu_f) & x_n = (l, fast), l = 1, 2, \ldots \\ \mu_s/(\lambda + \mu_s) & x_n = (l, slow), l = 1, 2, \ldots \end{cases} \tag{6}$$

$$C_n = \begin{cases} 0 & x_n = (0, empty) \\ (c_1\Lambda_n + c_2 t_n) & x_n = (l, fast), l = 1, 2, \ldots \\ 0 & x_n = (l, slow), l = 1, 2, \ldots \end{cases} \tag{7}$$

We let state 0 to be the state $0 := (0, empty)$, which will be proved below to satisfy a recurrence condition, and we let N to be the first passage stage to state 0, i.e. $N = \inf\{n > 0 : x_n = 0\}$.

To state the main result of this paper, we use the terms approachability and excludability: see definition 3.2 below. The results refer to Figure 2: explicit expressions for J, K, L and M are given in section 3.

Theorem 2.1. *(i) Let $B \subseteq \mathbb{R}^2$ be any closed convex set. Then B is approachable by the server* iff *it intersects both JM and KL. Otherwise is excludable by the arrivals.*
(ii) Any closed set $B' \subseteq \mathbb{R}^2$, which does not intersects either JM or KL is excludable by the arrivals.

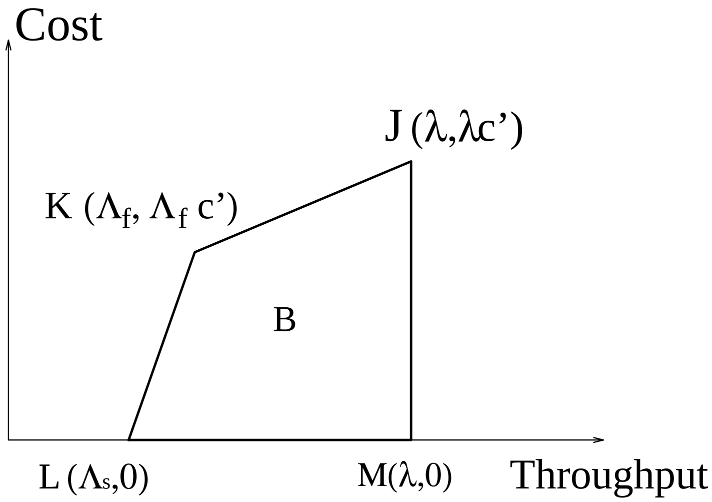

Figure 2: Performance Space

3 Approachability—definition and result

A policy for a decision maker can be any way of choosing an action based on the history of the system (i.e. its previous states and transition times) and the past actions (which includes a randomized decisions as well). We denote by π_i a policy for DMi (i=1,2), and by Π_i the class of all policies for DMi. A stationary policy is a policy that depends on the current state only. We denote by f (g) stationary policies for DM1 (DM2 respectively), and by $\Pi_1(S)$ $(\Pi_2(S))$ the class of all stationary policies.

For a given initial state $x_0 = x$, each pair of policies (π_1, π_2) induces a probability measure $P^x_{\pi_1,\pi_2}$. The corresponding expectation operator is denoted by $E^x_{\pi_1,\pi_2}$. A policy $\pi_1 \in \Pi_1$ is said to be stable if for some positive constants T_2 and R_2,

$$E^0_{\pi_1,\pi_2}\left(\sum_{m=0}^{N-1} t_m\right)^2 \le T_2 \qquad \forall \pi_2 \in \Pi_2\ , \tag{8}$$

$$E^0_{\pi_1,\pi_2}\left(\sum_{m=0}^{N-1} |r_m|\right)^2 \le R_2 \qquad \forall \pi_2 \in \Pi_2\ . \tag{9}$$

A set $\Pi'_1 \subseteq \Pi_1$ is uniformly stable if (8) and (9) are satisfied for every $\pi_1 \in \Pi'_1$ with the same constants T_2 and R_2. Stability for DM2's policies is defined symmetrically.

Definition 3.1. *Let $\{y_n, n \ge 0\}$ be a sequence of random variables over a measurable space $(\Omega, \Im)$ and let $\{P_\lambda, \lambda \in \Lambda\}$ be a collection of probability measures on $(\Omega, \Im)$. We say that $y_n \stackrel{n\to\infty}{\longrightarrow} 0$ P_λ-a.s. at a* uniform rate *over Λ if, for all $\varepsilon > 0$:*

$$\lim_{K\to\infty} \sup_{\lambda\in\Lambda} P_\lambda\left(\sup_{n\ge K} |y_n| > \varepsilon\right) = 0\ . \tag{10}$$

Let $d(r, B)$ denote the Euclidean point to set distance and $B^c_\delta = \{\beta \in \mathbb{R}^q : d(\beta, B) \ge \delta\}$.

Definition 3.2. *Let the initial state x be fixed. A set $B \subseteq \mathbb{R}^q$ is approachable (from x by DM1) if there exists a B-approaching policy $\pi^*_1 \in \Pi_1$ such that $d(\bar{r}_n, B) \stackrel{n\to\infty}{\longrightarrow} 0$ for every $\pi_2 \in \Pi_2$, at a uniform rate over Π_2. B is excludable (from x by DM2) if there* *exists a B-excluding policy* $\pi_2^* \in \Pi_2$ *such that for some* $\delta > 0$ $d(\bar{r}_n, B_\delta^c) \xrightarrow{n\to\infty} 0$ *for every* $\pi_1^* \in \Pi_1$, *at a uniform rate over* Π_1.

Remarks:

1. Approachability and Excludability are obviously the same for a set and for its closure, thus we can restrict ourselves to closed sets.

2. Obviously, a given set can't be both approachable by DM1 and excludable by DM2. However, a set can be neither approachable nor excludable.

3. Suppose a closed set B is approachable by DM2. Then by definition for all $\delta > 0$, B_δ^c is excludable by DM2. The same holds for DM1, thus a sufficient (necessary) condition for approachability is also a sufficient (necessary) condition for excludability.

4. The requirement of the uniform rate of convergence is of great importance. From this requirement it follows, that given any approachable set and a related desired accuracy, DM1 is able to satisfy it, with finitely many stages.

Let

$$R(f,g) := \lim_{n\to\infty} \frac{E^0_{f,g}\left(\sum_{m=0}^{n-1} r_m\right)}{E^0_{f,g}\left(\sum_{m=0}^{n-1} t_m\right)} \tag{11}$$

where r_m is the vector payoff, and we use the notation:
$R(\Pi_1(S), g) := \{R(f,g) : f \in \Pi_1(S)\}$
$\bar{R}(\Pi_1(S), g) := conv R(\Pi_1(S), g)$,
for the set of all possible ratios and its closed convex hull respectively. In a similar way we define $R(f, \Pi_2(S))$ and $\bar{R}(f, \Pi_2(S))$.

Theorem 3.3. *Assume the initial state is* 0 *and let* t *denote a generic transition time. Assume that the set of all policies are uniformly stable, and that there exist constants* a *and* b *such that for any pair of policies,* $a \leq E^0_{\pi_1,\pi_2}(t) \leq b$ *and the vector payoff is bounded. Let* $B \subseteq \mathbb{R}^q$ *be any closed convex set. Then (i)* B *is approachable by DM1 iff* $B \bigcap \bar{R}(\Pi_1(S), g) \neq \emptyset$ *for any* $g \in \Pi_2(S)$. *(ii) If* B *is not approachable by DM1 then it is excludable by DM2, with a stationary policy.*

Proof: see [3].

Remark: For conditions under which the set of all policies are uniformly stable see [3] Lemma 2.1. Theorem 3.3 holds under weaker assumptions, see [3, Chapter 3].

4 Proofs

In this section we prove Theorem 2.1. We use the following notation.
$fast^\infty/slow^\infty$—The stationary deterministic policies of DM1.
g_1—The stationary deterministic policy when all the customers join the queue.
g_2—The stationary deterministic policy when the customers do not join the queue unless they must (i.e. unless the queue is empty).

Lemma 4.1. *Assume $\lambda < \mu_s < \mu_f$. Then Theorem 3.3 holds.*

Proof: The proof follows the following steps: we prove by coupling that $E^0_{\pi_1,\pi_2}(N^2)$ is uniformly bounded for any pair of policies. In the current queuing system the payoff vector is bounded and the transition times satisfy $a \leq E^0_{\pi_1,\pi_2}(t) \leq b$. Thus the set of all policies is uniformly stable and Theorem 3.3 holds. (see proof in [3, chapter 4]). Q.E.D.

To obtain an explicit expression we define the following constants (see Figure 2):

$$c' := (c_1 + c_2/\mu_f), \tag{12}$$

$$\Lambda_s := \frac{\lambda\mu_s}{\lambda+\mu_s}\ , \qquad \Lambda_f := \frac{\lambda\mu_f}{\lambda+\mu_f}. \tag{13}$$

Note that $\Lambda_s < \Lambda_f < \lambda$ and $c' > 0$. Let JM be the line segment between $J = (\lambda, \lambda c')$ and $M = (\lambda, 0)$, and let KL be the line segment between $K = (\Lambda_f, \Lambda_f c')$ and $L = (\Lambda_s, 0)$. Similarly we define JK and LM.

Proof of Theorem 2.1: The proof consists of the following three steps:
Step 1: Suppose that: (i) $JM = R(\Pi_1(S), g_1)$ and $KL = R(\Pi_1(S), g_2)$.
(ii) $JK = R(fast^\infty, \Pi_2(S))$ and $LM = R(slow^\infty, \Pi_2(S))$.
Let $B \subseteq \mathbb{R}^2$ be any closed convex set. Then B is approachable by DM1 *iff* it intersects both JM and KL, otherwise it is excludable by DM2. To see that, suppose the performance set B does not intersect JM and/or KL. Then DM2 has at least one stable policy, g_1 or g_2 respectively, such that $B \bigcap \bar{R}(\Pi_1(S), g_{1(2)}) = \emptyset$, and thus by Theorem 3.3 B is excludable by DM2. Now, suppose the performance set B does intersect JK and ML. B is a convex set and thus it intersects any line segment which connects any point in JM with any point in KL. From (ii) it then follows that for any $g \in \Pi_2(S)$, $B \bigcap \bar{R}(\Pi_1(S), g_1) \neq \emptyset$. Thus by Theorem 3.3, B is approachable by DM1. Thus, if we prove (i) and (ii) part (a) of Theorem 2.1 follows.
Step 2: Suppose (i) in step 1 holds. Then any closed set $B' \subseteq \mathbb{R}^2$, which does not intersect either JM or KL is excludable by DM2. To see that suppose $B' \bigcap JM \neq \emptyset$. By (i) $JM = \bar{R}(\Pi_1(S), g_1)$. Then by Theorem 3.3, JM is approachable by DM2. Now, since JM is a compact set and B' is closed set we have $d(B', \bar{R}(\Pi_1(S), g_1) = d(B', JM) > 0$. That means B' is excludable by DM2 (or in other words there exists a $\delta' > 0$, such that $B'^c_{\delta'}$ is approachable by DM2). By following the same arguments we see that B' is excludable by DM2 if $B' \bigcap KL \neq \emptyset$. Thus, if we prove (i) part (b) of the Theorem follows.
Step 3: Here we prove (i) and (ii). Then by step 1 and step 2 the Theorem follows. To this end, note that [see [4] Th. 7.5]:

$$R(f,g) = \lim_{n\to\infty} \frac{E^0_{f,g}\left(\sum_{m=0}^{n-1} r_m\right)}{E^0_{f,g}\left(\sum_{m=0}^{n-1} t_m\right)} = \frac{E^0_{f,g}\left(\sum_{m=0}^{N-1} r_m\right)}{E^0_{f,g}\left(\sum_{m=0}^{N-1} t_m\right)}. \tag{14}$$

Denote $\bar{\Lambda}_N := \frac{E^0_{f,g}(\sum_{m=0}^{N-1} \Lambda_m)}{E^0_{f,g}(\sum_{m=0}^{N-1} t_m)}$ and $\bar{C}_N := \frac{E^0_{f,g}(\sum_{m=0}^{N-1} C_m)}{E^0_{f,g}(\sum_{m=0}^{N-1} t_m)}$.

Let f be any stationary policy for DM1. The policy f is well defined by the probability to choose the fast-service mode when the queue is empty. We denote this probability by α_0, i.e. $\alpha_0 = \Pr_f(a_n^1 = fast | x_n = (0, empty))$. Then by a smoothing argument, (6), (7) and (14) we have:

$$\bar{\Lambda}_N = \frac{\frac{\alpha_0 \mu_f}{\mu_f + \lambda} E^0_{fast^\infty, g}(N-1) + \frac{(1-\alpha_0)\mu_s}{\mu_s + \lambda} E^0_{slow^\infty, g}(N-1)}{\frac{1}{\lambda} + \frac{\alpha_0}{\mu_f + \lambda} E^0_{fast^\infty, g}(N-1) + \frac{(1-\alpha_0)}{\mu_s + \lambda} E^0_{slow^\infty, g}(N-1)} \tag{15}$$

$$\bar{C}_N = \frac{\frac{c' \alpha_0 \mu_f}{\mu_f + \lambda} E^0_{fast^\infty, g}(N-1)}{\frac{1}{\lambda} + \frac{\alpha_0}{\mu_f + \lambda} E^0_{fast^\infty, g}(N-1) + \frac{(1-\alpha_0)}{\mu_s + \lambda} E^0_{slow^\infty, g}(N-1)} \tag{16}$$

Proof of (i): Let $g = g_1$, and let f be any stationary policy for DM1. We proceed by calculating $E^0_{fast^\infty, g_1}(N)$ and $E^0_{slow^\infty, g_1}(N)$. We note that for both cases, N is twice the number of customers being served at the busy period of the system. Then by well known result for M/M/1 queue [[5] Chapter 5 Section 9]:

$$E^0_{fast^\infty, g_1}(N) = \frac{2\mu_f}{\mu_f - \lambda}\ , \quad \text{and} \quad E^0_{slow^\infty, g_1}(N) = \frac{2\mu_s}{\mu_s - \lambda}.$$

Substituting the above in (15) and (16) we obtain:

$$\bar{\Lambda}_N = \lambda\ , \quad \text{and} \quad \bar{C}_N = \frac{\lambda \alpha_0 c' \mu_f (\mu_s - \lambda)}{\mu_s(\mu_f - \lambda) + \lambda \alpha_0 (\mu_s - \mu_f)}. \tag{17}$$

$\bar{C}_N$ in (16) is a continuous and increasing function of α_0 and thus

$$0 = \bar{C}_N(\alpha_0 = 0) \le \bar{C}_N(\alpha_0) \le \bar{C}_N(\alpha_0 = 1) = \lambda c',$$

which together with (14) implies $JM = \bar{R}(\Pi_1(S), g_1)$.

Now, let $g = g_2$. We calculate $E^0_{fast^\infty, g_1}(N)$ and $E^0_{slow^\infty, g_1}(N)$ directly. Suppose $f = fast^\infty$. Then the pair of policies $\{fast^\infty, g_2\}$ induce a Markov chain with only two states $\{(0, empty), (1, fast)\}$ and the transition matrix

$$\begin{pmatrix} 0 & 1 \\ \frac{\mu_f}{\lambda + \mu_f} & \frac{\lambda}{\lambda + \mu_f} \end{pmatrix}.$$

This Markov chain is irreducible and aperiodic and hence the required expectations are the inverse of the stationary probabilities [[6], Chapter 6, 2.19]. Thus
$E^0_{fast^\infty, g_1}(N) = 2 + \frac{\lambda}{\mu_f}$ and using similar steps $E^0_{fast^\infty, g_1}(N) = 2 + \frac{\lambda}{\mu_f}$. Substituting in (15) and (16) we obtain

$$\bar{\Lambda}_N = \frac{\lambda \mu_s \mu_f}{\lambda \alpha_0 \mu_s + \lambda(1 - \alpha_0)\mu_f + \mu_f \mu_s} \tag{18}$$

$$\bar{C}_N = (\bar{\Lambda}_N - \Lambda_s) \frac{\mu_f c'(\lambda + \mu_s)}{\lambda(\mu_f - \mu_s)} \tag{19}$$

Note that $(\bar{C}_N, \bar{\Lambda}_N)$ in (18)–(19) satisfies the equation of the linear line through the points K and L. Moreover $\bar{C}_N$ is a continuous and increasing function of α_0 and thus

$$0 = \bar{C}_N(\alpha_0 = 0) \le \bar{C}_N(\alpha_0) \le \bar{C}_N(\alpha_0 = 1) = \lambda \Lambda_f c'$$

which, together with (13) implies $KL = \bar{R}(\Pi_1(S), g_2)$.

Proof of (ii): Suppose $f = fast^\infty$, and DM2 applies any stationary policy g. Then by (15) and (16):

$$\bar{C}_N = c'\bar{\Lambda}_N \tag{20}$$

$$\bar{\Lambda}_N = \frac{\frac{\mu_f}{\mu_f+\lambda} E^0_{fast^\infty,g}(N-1)}{\frac{1}{\lambda} + \frac{1}{\mu_f+\lambda} E^0_{fast^\infty,g}(N-1)} \tag{21}$$

Note that $(\bar{C}_N, \bar{\Lambda}_N)$ in (20)–(21) satisfies the equation of the linear line through the points J and K. $\bar{\Lambda}_N$ is a continuous and increasing function of α_0. By coupling arguments it can be shown that $E^0_{fast^\infty,g}(N)$ achieves its maximal value whenever $g = g_2$, and achieves its minimal value whenever $g = g_1$. Thus

$$\Lambda_f \leq \bar{\Lambda}_N(g = g_2) \leq \bar{\Lambda}_N(g) \leq \bar{\Lambda}_N(g = g_1) = \lambda.$$

Thus, by using (13), we have $JK = \bar{R}(fast^\infty, \Pi_2(S))$.

Following the same arguments for $f = slow^\infty$ we have $LM = \bar{R}(slow^\infty, \Pi_2(S))$, hence (ii) is satisfied. Q.E.D.

Note that the boundary of $JKLM$ is defined by the extreme policies $g_1, g_2, fast^\infty$ and $slow^\infty$. The line through the points J and K satisfies the equation $\bar{C} = c'\bar{\Lambda}$, which is obviously the case when the server always uses the fast service mode. Similarly, when the server always apply the slow service mode we have $\bar{C} = 0$. Also note that since the queue is stable (or the embedded Markov chain is positive recurrent for any pair of stationary policies) we have $\bar{\Lambda} = \lambda$ when all the customers join the queue (i.e. $g = g_1$). Remark: It can be shown by (13), (15) and (16) that for any fixed g, a fixed stationary policy of the customers, $\bar{C}_N$ is an affine function of $\bar{\Lambda}_N$ for any policy f of the server.

The approaching policy

The approaching policy is as follows: whenever at 0, let ρ be the current time averaged payoff. For any $\rho \notin B$, we define $C(\rho)$ to be a closest point in B to ρ and $H(\rho)$ to be the line through $C(\rho)$ that is perpendicular to $\rho - C(\rho)$. Then for any $\rho \notin B$:

(i) If $H(\rho)$ separates JK from ρ—use $fast^\infty$.

(ii) If $H(\rho)$ separates LM from ρ—use $slow^\infty$.

Otherwise—use $fast^\infty$ or $slow^\infty$. For $\rho \in B$ use any policy.

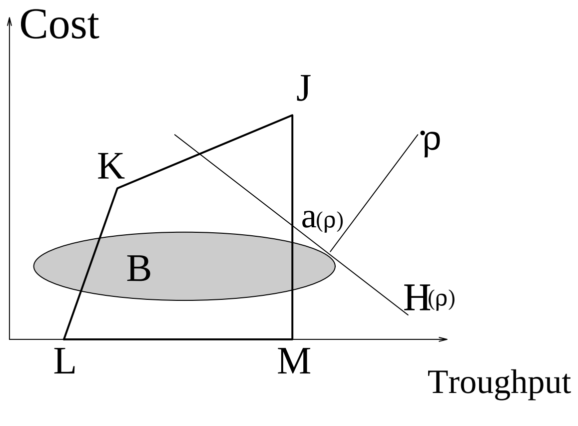

Figure 3: approaching policy

A performance set which is particularly interesting is a set with threshold values (Λ_0, C_0), so that the server (DM1), can guarantee a throughput of at least Λ_0 while the cost does not exceed C_0. We denote this set by B_0, $B_0 := \left\{(\bar{\Lambda}, \bar{C}) : \bar{\Lambda} \geq \Lambda_0, \bar{C} \leq C_0\right\}$. Then B_0 is approachable by the server *iff* it intersects JK. In other words the line JK is the set of all the optimal "thresholds sets," that the server can guarantee, in the sense that he cannot have better threshold sets with the two criteria simultaneously. JK represents the trade-off between the service cost and the throughput.

5 Extensions

Similar results were obtained for the case where the server can change the service mode at any stage. Under mild assumptions, we analyzed the approachable regions where the service times are arbitrary random variables: see [3]. These queuing systems are a special case of Semi-Markov decision process with two decision makers, for which we derive in [3] sufficient conditions for a set to be approachable, and necessary and sufficient conditions for a convex set to be approachable. The results in [1] and [2] are a special case of the results in [3].

Acknowledgments. Research of the second author was supported in part by the fund for promotion of research at the Technion.

References

[1] D. Blackwell, "An analog for the minmax theorem for vector payoffs", *Pacific J. Math.*, 6 pp. 1-9, 1956.

[2] N. Shimkin and A. Shwartz, "Guaranteed Performance Region in Markovian Systems with Competing Decision Makers", *IEEE Trans. Auto. Cont.*, Vol. 38, No. 1, Jan. 1993.

[3] R. Levi and A. Shwartz, *A Theory of Approachability and Throughput-Cost Trade-off in a Queue with Impatient Customers*, EE Pub. 936, Technion, Haifa August 1994.

[4] S. M. Ross, *Applied Probability models with Optimization Application*, Holden-Day, San Francisco, 1970.

[5] L. Kleinrock, *Queuing systems*, Wiley-Interscience, New-York, 1975.

[6] E. Cinlar, *Introduction to stochastic processes*, Prentice-Hall, NJ., 1975.