Queueing Systems 38, 125–148, 2001

Call Admission: A New Approach to Quality of Service

MARTIN I. REIMAN
Bell Laboratories, Murray Hill, NJ, USA

ADAM SHWARTZ*
Electrical Engineering Technion, Israel

Received 25 March 1999; Revised 10 December 2000

Abstract. Quality of service guarantees are an important and much discussed aspect of ATM network design. However, there is no standard definition of quality of service. Moreover, some often-used criteria seem quite crude. We consider call admission to a bufferless ATM multiplexer with on/off sources. A new criterion for a guarantee on average cell loss is proposed. This criterion represents the quality of service from the point of view of the user, and is thus more reliable. We calculate the optimal policy that minimizes blocking subject to the guarantee, when there is only one type of user.

The measure of cell-loss we propose is applicable to a wide range of models. It gives rise to a mathematical programming formulation, which we derive explicitly for our case.

Keywords: admission control, semi-Markov decision process, near complete decomposability, mathematical programming, ATM

1. Introduction

Quality of service guarantees are a key issue in the design of modern computer communications networks. In the context of ATM (Asynchronous Transfer Mode) networks, the two most obvious performance issues – from the users' point of view – are call blocking and cell loss. When statistical multiplexing is used to improve performance, these criteria necessarily take a probabilistic form, such as probability of blocking, or average cell loss.

The precise meaning of "average cell loss", however, has not been agreed upon. Reiman et al. [9] consider the problem of dynamic admission control of an ATM multiplexer. Two types of QoS (Quality of Service) guarantees on maximal cell loss are proposed. The conservative QoS guarantee involves a "probabilistic worst case", in that it measures the average cell loss in a given state of the system. The aggressive QoS guarantee is based on a steady state average; see section 5 for details. These are both common approaches to QoS guarantees in analytic studies. However, in real systems QoS guarantees are made to individual users, and from the users' point of view, both measures are too crude. A specific user may have a pattern of usage that results in

* Work performed while on sabbatical leave, visiting Bell Laboratories, Murray Hill and Rutgers.

his/her requiring service in a manner such that the "average" state of the system he/she sees is different from that of steady state. In addition, one would like to make adherence to the guarantees as strict as possible, in the sense that little averaging is done, so that the guarantee is as often met pathwise as possible. On the other hand, an overly conservative approach to these guarantees would obviate some of the benefits of statistical multiplexing. Nagarajan et al. [8] suggest that short term criteria may be more appropriate for modern applications: see [8, and references therein] for further discussion of this point. They analyze a different model – an $M/M/1/K$ queue – and consider criteria such as the probability that more than a given fraction of cells are lost over a fixed (deterministic) interval of time. The length of this interval represents the average time of a connection. We introduce a new type of QoS guarantee, which we consider in the context of the model of [9]. We say the "individual cell-loss guarantee" is met if the expected number of lost cells over a call's duration of service is smaller than a specified fraction of (the expected number of) cells that were submitted by all calls during that time. For a precise definition see (5.12) for the case where all calls are identical, or (8.8) when several call types are considered. The criteria in [8] are more accurate in that they measure the probability that more than a given fraction of cells are lost, whereas we compare expected values. On the other hand, our criterion is more precise in that it follows an individual call until its completion, rather than considering an average call duration.

The emphasis here is on what a call experiences during its sojourn time in the system, that is, while the call is being served. The averaging is only with respect to events that occur once the call is admitted: if a call is admitted in a given state of the system, then the average performance *given that initial state* should be within the guaranteed limits.

In this paper we investigate dynamic call admission to an ATM multiplexer. We show how to compute an optimal policy that maximizes call admissions (minimizes blocking), subject to a guarantee on individual cell loss. One immediate consequence of our criteria is that the dynamic policies we consider are necessarily stationary, since all calls entering at the same state should be treated equally. In addition, as in [7], we restrict our search to non randomized (deterministic) policies (although for constrained problems of this type, typically optimal policies require randomization).

The paper is organized as follows. Section 2 is based on [9], and describes a model of dynamic call admission, with on–off sources of several types. A simplification based on time scale separation (also based on [9]) is delineated in section 3. Section 4 introduces the basic blocking and cell loss measures. In sections 5–7 we consider a single type of user. Section 5 provids precise definitions of the performance criteria in this context. In section 6 we prove a basic structural result for the single type case: the new "individual" criterion leads to a call admission region that is larger than with the worst-case, but more cautious than the average criterion. We provide an efficient procedure for calculating this admission region in section 7, and compute optimal thresholds for several values of the parameters. The derivation of this procedure relies on an analysis of difference equations with boundary conditions. In section 8 we consider the general problem, with several different types of users. We show that the solution of our opti-

mization problem solves a mathematical program. This program is actually applicable to much more general models. It is typically composed of a linear program, coupled to a second set of linear constraints, but the coupling is via nonlinear equalities. Both the size of the problem and the nonlinear equalities make this program essentially impossible to solve for any realistically sized system. Some background for the mathematical program is provided in the appendix.

2. Admission to an ATM trunk

We model an ATM link as a fluid "pipe" with transmission capacity R, no buffer, and J types of calls (connections). Calls of type i, $i = 1, 2, \ldots, J$, arrive according to a Poisson process of rate λ_i. If a call of type i is admitted, it alternates (independently of everything else) between "on" and "off" states as a two state Markov process, with average "on" period of β_i^{-1} and average "off" period of α_i^{-1}. If a type i call were to stay forever, the fraction of time it spends in the "on" state would be $p_i = \alpha_i/(\alpha_i + \beta_i)$. Upon turning off, a type i call departs with probability q_i. When on, a type i call generates cells, modeled as a fluid, with rate ν_i. When cells are generated at a rate exceeding R, the excess cells are lost. This is precisely the model considered in [7,9], except that here we are considering an arbitrary number J of call types, and more general criteria.

As in [9], we consider two aspects of Quality of Service (QoS): cell level QoS and call level QoS. At the cell level, we want the cell loss ratio of type i to be smaller than f_i, typically of the order of $f_i \approx 10^{-9}$. The call level QoS is the call blocking probability. To allow for different relative importance of blocking different types of calls, we let each admitted type i call pay a reward w_i. We then consider the problem of maximizing the long run average reward rate, subject to the cell level QoS constraints.

3. Problem reduction using time scale decomposition

The optimization problem described above can be formulated as a Semi-Markov Decision Process (SMDP) with constraints; for further details, see [9]. Let $k_i(t)$ denote the number of type i calls in progress at time t, and let $n_i(t)$ denote the number of type i calls in the "on" state at time t, $i = 1, 2, \ldots, J$. Denote $\mathbf{k} = \{k_1, k_2, \ldots, k_J\}$, $\mathbf{n} = \{n_1, n_2, \ldots, n_J\}$, $\mathbf{k}(t) = \{k_1(t), k_2(t), \ldots, k_J(t)\}$ and $\mathbf{n}(t) = \{n_1(t), n_2(t), \ldots, n_J(t)\}$. Under a stationary admission control strategy $\{(\mathbf{k}(t), \mathbf{n}(t)),\ t \geqslant 0\}$ is a Markov process. As in [9], we use the notion of Nearly Completely Decomposable (NCD) Markov chains to reduce the dimension of this problem from $2J$ to J. We present a brief heuristic discussion; see [9] for more details and justification.

During a call's "lifetime" it goes through many on/off cycles. Intuitively, when making a call admission decision, the number of calls of each type in progress is important, but the number of calls of each type in the on state is not, because these quantities oscillate too rapidly. Thus the $\mathbf{n}$ component of the state becomes noise and can be ignored for admission control purposes. This part of the state *does* affect the cell loss rate,

so it must be "averaged" properly. The **n** process reaches equilibrium between changes in the **k** process. The equilibrium corresponds to fixed **k** and is given by the binomial distribution:

$$\psi(\mathbf{k}, \mathbf{n}) = \prod_{i=1}^{J} \binom{k_i}{n_i} p_i^{n_i} (1 - p_i)^{k_i - n_i}.$$

When the total cell arrival rate is

$$\lambda_a = \lambda_a(\mathbf{n}) = \sum_{i=1}^{J} n_i \nu_i,$$

the cell loss rate is $[\lambda_a(\mathbf{n}) - R]^+$. The decision of which cells to discard may, in general, depend on cell type, as well as on the detailed state $(\mathbf{k}, \mathbf{n})$. If we adopt the simple rule that cell loss is proportional to the number of cells submitted for transmission, then the average type i cell loss rate in state **k** is

$$b_i(\mathbf{k}) = \sum_{n_1=0}^{k_1} \cdots \sum_{n_J=0}^{k_J} \psi(\mathbf{k}, \mathbf{n}) \cdot \left[\lambda_a(\mathbf{n}) - R\right]^+ \frac{n_i \nu_i}{\lambda_a(\mathbf{n})}. \tag{3.1}$$

As indicated above, in the NCD limit the detailed on/off behavior of calls disappears when viewing call arrivals and departures. Calls of type i arrive (as before) in a Poisson process of rate λ_i, and have exponential holding times with rate $\mu_i = q_i \alpha_i \beta_i / (\alpha_i + \beta_i) = q_i \beta_i p_i$.

4. The criteria

As in [9], we augment the state **k** with an additional variable, j, indicating an arrival of a j-type call, or a non-arrival state (in which case $j = 0$). In the interest of brevity, we retain the notation of **k** for this expanded state. For an explicit derivation of the transitions of the ensuing semi-Markov process, see [9]. Denote the transition probabilities from state **k** to state $\mathbf{k}'$ when action a is chosen by $p(\mathbf{k}, \mathbf{k}'; a)$. We are interested in minimizing the steady-state blocking probability, possibly weighted between different types, subject to constraints on cell loss. This is formalized as follows.

For each call that we accept in state **k** (which now includes the arrival type), we associate a positive reward $r(\mathbf{k}, a)$: here a is the action taken, where $a = 1$ indicates acceptance of the call, and $a = 0$ corresponds to blocking (rejection). For example, in [9], $r(\mathbf{k}, a) = w_j$ whenever $a = 1$ and $j \neq 0$. The (expected) steady state of the resulting cost is to be minimized.

In addition, we let g_i, $i = 1, \ldots, J$, be cost functions associated with cell-loss by a type-i call. In section 5, the following two alternatives are used:

$$g_i^0(\mathbf{k}, a) = b_i(\mathbf{k}) - f_i k_i p_i \nu_i, \tag{4.1}$$

$$g_i^1(\mathbf{k}, a) = \frac{g_i^0(\mathbf{k}, a)}{k_i} \tag{4.2}$$

with the interpretation that these measure the number of cells lost by type-i calls, as a fraction of cells submitted for transmission by this type. The cell-loss guarantee entails that the average of this measure, over the sojourn time of the arriving customer, be kept below the promised level of service.

5. Admission to an ATM link with identical users

In this section we restrict our attention to the one type version of the problem described above, so that $J = 1$. This leads to a one-dimensional problem for which we can prove structural results and easily perform numerical calculations.

The state of the system at time t is denoted by $k(t)$. For the reasons described in sections 8 and 6 we only consider stationary, deterministic, threshold policies: Admit a call if and only if the state immediately preceding its arrival is less than some threshold.

Let K denote the threshold (which we hold fixed and, in the interest of brevity, omit from the notation). Consider the one dimensional birth–death chain $k_u(m)$ "induced" by this threshold policy. This discrete time Markov chain is obtained by uniformizing the associated birth–death process. Let $\{p_{ij},\ 0 \leqslant i, j \leqslant K\}$ denote the one-step transition probabilities for this chain. Then the total event rate is $\Lambda = \lambda + K\mu$, and

$$\begin{aligned} p_{ii+1} &= \Lambda^{-1}\lambda, && 0 \leqslant i < K, \\ p_{ii-1} &= \Lambda^{-1} i\mu, && 1 \leqslant i \leqslant K, \\ p_{ii} &= \Lambda^{-1}(\Lambda - \lambda - i\mu), && 0 \leqslant i < K, \\ p_{KK} &= \Lambda^{-1}\lambda, \end{aligned}$$

and $p_{ij} = 0$ otherwise. Let $\pi_\ell^{(k_0)}(m) = P\{k_u(m) = \ell \mid k(0) = k_0\}$ for $0 \leqslant \ell, k_0 \leqslant K$ and $m \geqslant 0$. In addition, let π_ℓ, $0 \leqslant \ell \leqslant K$, denote the stationary distribution of this process (so that, by ergodicity, $\pi_\ell = \lim_{m\to\infty} \pi_\ell^{(k_0)}(m)$). This process is precisely an Erlang loss model, so that

$$\pi_\ell = \left[\sum_{m=0}^{K} \frac{(\lambda/\mu)^m}{m!}\right]^{-1} \frac{(\lambda/\mu)^\ell}{\ell!}, \quad 0 \leqslant \ell \leqslant K.$$

Note that, for a fixed ℓ, π_ℓ is monotone decreasing in K.

The constraint associated with the conservative view of QoS is defined in [9], as follows. Let

$$b(k) = \sum_{n=0}^{k} \binom{k}{n} p^n (1-p)^{k-n} [nv - R]^+ \tag{5.1}$$

denote the cell-loss rate with k calls in progress (averaged over the stationary distribution of the number of calls in the on state). Then the conservative QoS constraint is that

$$b(k) \leqslant fkpv, \quad 0 \leqslant k \leqslant K; \tag{5.2}$$

that is, the average cell loss per unit time, when in state k, is at most a fraction f of the average amount $kp\nu$ produced, per unit time, in that state. It will be convenient to rewrite this condition in terms of the function g_0

$$g_0 = b(k) - fkp\nu. \tag{5.3}$$

Let K_{C} denote the largest threshold allowed by the conservative QoS constraint, which is defined as

$$K_{\mathrm{C}} = \max\bigl\{k\colon\ g_0(k) \leqslant 0\bigr\}. \tag{5.4}$$

The aggressive QoS constraint [9] is based on steady-state behavior, and can be written as

$$\sum_{k=0}^{K} g_0(k)\pi_k \leqslant 0. \tag{5.5}$$

Let K_{A} denote the largest threshold allowed by the aggressive QoS constraint, defined as

$$K_{\mathrm{A}} = \max\left\{K\colon\ \sum_{k=0}^{K} g_0(k)\pi_k \leqslant 0\right\}. \tag{5.6}$$

The individual call-based QoS constraint is based on the cell loss experienced during the life of a "marked" call in the system. To determine this cell loss we need to define an appropriate absorbing process and associated "taboo" probabilities.

Let $\{X(m),\ m \geqslant 0\}$ be an absorbing Markov chain on the state space $\{0, 1, \ldots, K\}$ (with the state 0 as the absorbing state), obtained by modifying the process k_u so that it jumps to state 0 when the marked call departs (note that state 0 for this chain is different from state 0 for the original chain). Let

$${}_a\pi_{\ell}^{(k_0)}(m) = P\bigl\{X(m) = \ell \mid X(0) = k_0\bigr\}, \quad 0 \leqslant \ell, k_0 \leqslant K. \tag{5.7}$$

For $0 < \ell \leqslant K$, ${}_a\pi_{\ell}^{(k_0)}(m)$ is the probability that the "absorbing" system is in state ℓ at time m, with the marked call still in the system, given that the initial state is k_0 including the marked call. Conceptually, the marked call arrives at $m = 0$. However, due to the exponential service assumption, it suffices to assume that it is in the system at $m = 0$.

Let $\{p_{ij}^a,\ 0 \leqslant i, j \leqslant K\}$ denote the one-step transition probabilities for this absorbing chain. Then

$$\begin{cases} p_{ii+1}^a = \Lambda^{-1}\lambda, & 1 \leqslant i < K, \\ p_{ii-1}^a = \Lambda^{-1}(i-1)\mu, & 1 \leqslant i \leqslant K, \\ p_{i0}^a = \Lambda^{-1}\mu, & 1 \leqslant i \leqslant K, \\ p_{ii}^a = \Lambda^{-1}(\Lambda - \lambda - i\mu), & 1 \leqslant i < K, \\ p_{KK}^a = \Lambda^{-1}\lambda \quad \text{and} \quad p_{00}^a = 1. \end{cases} \tag{5.8}$$

Note that the "defect" (the probability of jumping to the absorbing state) is $\Lambda^{-1}\mu$ for all states $1 \leqslant i \leqslant K$. Let

$$ {}_{a}^{s}\pi_{\ell}^{(k_0)} = \sum_{m=0}^{\infty} {}_{a}\pi_{\ell}^{(k_0)}(m), \tag{5.9} $$

denote the sum of the marginal probabilities, and define

$$ {}_{a}^{n}\pi_{\ell}^{(k_0)} = {}_{a}^{s}\pi_{\ell}^{(k_0)} \left[\sum_{k=1}^{K} {}_{a}^{s}\pi_{k}^{(k_0)} \right]^{-1}, \quad 1 \leqslant k_0, \ell \leqslant K. \tag{5.10} $$

Note that ${}_{a}^{n}\pi_{\ell}^{(k_0)}$ is a "normalized" version of the ${}_{a}^{s}\pi_{\ell}^{(k_0)}$, in that

$$ \sum_{\ell=1}^{K} {}_{a}^{n}\pi_{\ell}^{(k_0)} = 1, \quad 1 \leqslant k_0 \leqslant K. \tag{5.11} $$

Define the individual call-based QoS constraint as

$$ \sum_{m=0}^{\infty} \sum_{k=0}^{K} g_0(k) {}_{a}\pi_{k}^{(k_0)}(m) \leqslant 0, \quad 0 < k_0 \leqslant K. \tag{5.12} $$

Note that

$$ \sum_{m=0}^{\infty} \sum_{k=0}^{K} b(k) {}_{a}\pi_{k}^{(k_0)}(m) \tag{5.13} $$

is the expected amount of "fluid" lost over the duration of the marked call, while

$$ p\nu \sum_{m=0}^{\infty} \sum_{k=0}^{K} k_a \pi_{k}^{(k_0)}(m) \tag{5.14} $$

is the expected amount of fluid produced over the duration of the marked call. Thus the QoS constraint requires that the average loss does not exceed a given fraction of the average produced, both over the duration of the marked call.

If we let K_{I} denote the largest threshold allowed by the individual call-based QoS constraint, we can write it as

$$ K_{\mathrm{I}} = \max \left\{ K \colon \sum_{k=0}^{K} g_0(k) {}_{a}^{n}\pi_{k}^{(k_0)} \leqslant 0, \ 1 \leqslant k_0 \leqslant K \right\}. \tag{5.15} $$

The definitions of K_{C}, K_{A} or K_{I} can be generalized in the following way. Let K_{C}^{ℓ} (K_{A}^{ℓ} or K_{I}^{ℓ}) be defined through (5.4) (respectively (5.6) or (5.15)), but with the function g_0 replaced by some g_ℓ. Of particular interest is the function

$$ g_1(k) = \begin{cases} 0 & \text{if } k = 0, \\ \dfrac{g_0(k)}{k} & \text{if } k \geqslant 1. \end{cases} \tag{5.16} $$

With this definition, we have $K_C^1 = K_C$. However, the average as well as the individual thresholds are, in general, different. Since k is the number of calls in progress, the individual criterion under g_1 has a natural interpretation. If we assume that lost cells are shared equally among all active users, then $b(k)/k$ is the amount of fluid lost by the marked call, whereas $fp\nu$ is the amount of fluid produced by that call. So, K_I^1 is the largest threshold under which the expected fraction of cells lost by the individual, marked call is not greater than f.

The class of functions g_ℓ we consider is described in the following assumption.

Assumption 1. The function g_ℓ satisfies the following: for some finite $K_C^\ell > 0$,

$$g_\ell(0) = 0, \tag{5.17}$$

$$g_\ell(k) \leqslant 0 \qquad \text{for all } 1 \leqslant k \leqslant K_C^\ell, \tag{5.18}$$

$$g_\ell(k) > 0 \text{ and increasing} \qquad \text{for } k > K_C^\ell. \tag{5.19}$$

Lemma 2. The functions g_0, defined in (5.3), and g_1, defined in (5.16), satisfy the conditions of assumption 1. The function g_0 is strictly convex for $k \geqslant K_C$.

Proof. Note that g_0 is defined only for integer values of its argument, so that the appropriate definition of convexity needs be applied. Consider the effect of an additional, distinguished call on the function b, defined in (5.1). This distinguished call is either active (with probability p), or not producing fluid (with probability $1 - p$). Define $n_0 = \max\{n\colon n\nu \leqslant R\}$, and

$$\mathbb{P}_k^{=} = \mathbb{P}(n_0 \text{ calls are active} \mid k \text{ calls in progress}) \tag{5.20}$$

$$\mathbb{P}_k^{>} = \mathbb{P}(\text{at least } n_0 + 1 \text{ calls are active} \mid k \text{ calls in progress}). \tag{5.21}$$

Then

$$b(k+1) - b(k) = (1-p)\cdot 0 + p(n_0\nu - R + \nu)\mathbb{P}_k^{=} + p\nu\mathbb{P}_k^{>} \tag{5.22}$$

$$= p(n_0\nu - R + \nu)\big(\mathbb{P}_k^{=} + \mathbb{P}_k^{>}\big) + p\big(\nu - (n_0\nu - R)\big)\mathbb{P}_k^{>}. \tag{5.23}$$

However,

$$\mathbb{P}_k^{=} + \mathbb{P}_k^{>} = \mathbb{P}(\text{at least } n_0 \text{ calls are active} \mid k \text{ calls in progress}). \tag{5.24}$$

Clearly then, both $\mathbb{P}_k^{>}$ and $(\mathbb{P}_k^{=} + \mathbb{P}_k^{>})$ are strictly increasing in k as soon as $k > n_0$, and since $(\nu - (n_0\nu - R)) \geqslant 0$ we conclude that $b(k+1) - b(k)$ is strictly increasing in k in this range.

Since g_0 is the difference of b and a linear function, it is strictly convex wherever b is. By definition of K_C, this is clearly the case for $k \geqslant K_C$. Now by definition, $g_0(0) = 0$ and $g_0(k) < 0$ for $k < K_C$. Since $g_0(K_C+1) > 0$ it is necessarily increasing for $k > K_C$.

By definition, $g_1(0) = 0$ and $g_1(k) \leqslant 0$ if and only if $g_0(k)$ is. Using the strict convexity of g_0 and $g_0(0) = 0$,

$$\frac{g_0(k+1)}{k+1} = \frac{1}{k}\left(\frac{k}{k+1}g_0(k+1) + \frac{1}{k+1}g_0(0)\right) \tag{5.25}$$

$$> \frac{1}{k}g_0\left(\frac{k}{k+1}(k+1) + \frac{1}{k+1}\cdot 0\right) \tag{5.26}$$

$$= \frac{1}{k}g_0(k). \tag{5.27}$$

Thus $g_1(k)$ is strictly increasing for $k > K_C$. □

It will be convenient to obtain a different representation of K_I, which holds for any g_ℓ that satisfies assumption 1. Let τ denote the time at which the marked call completes service. Define

$$V(K, k_0) = E_{k_0}^K \sum_{m=0}^{\tau-1} g_\ell\big(X(m)\big), \quad 1 \leqslant k_0 \leqslant K < \infty, \tag{5.28}$$

where $E_{k_0}^K$ denotes the expected value under threshold K with initial state k_0. For notational convenience we define $V(K, 0) = 0$, $0 \leqslant K < \infty$. In addition, since the results below apply for any ℓ, we drop it from the notation (and in particular use K instead of K^ℓ). Since $g_\ell(0) = 0$, we can write

$$V(K, k_0) = \sum_{m=0}^{\infty}\sum_{k=1}^{K} g_\ell(k)_a\pi_k^{(k_0)}(m) = \sum_{k=1}^{K} g_\ell(k)_a^s\pi_k^{(k_0)}. \tag{5.29}$$

By the definition of K_I,

$$V(K_I, k_0) \leqslant 0, \quad 0 < k_0 \leqslant K_I, \tag{5.30}$$

and

$$V(K_I + 1, k_0) > 0 \quad \text{for some } k_0 \in \{1, \ldots, K_I + 1\}. \tag{5.31}$$

Let

$$V^n(K, k_0) = \sum_{k=1}^{K} g_\ell(k)_a^n\pi_k^{(k_0)}, \quad 1 \leqslant k_0 \leqslant K.$$

Then $V(K, k_0) \leqslant 0$ if and only if $V^n(K, k_0) \leqslant 0$.

6. Identical users: structure of the optimal policies

Our optimization problem is to minimize the call blocking probability, subject to a constraint on the individual cell loss. We call a policy feasible if, under the policy, the

relevant constraints are satisfied, and optimal if it is feasible and, in addition, provides the lowest possible blocking probability among all feasible policies. The constraint is to hold for every accepted call, and consequently we must use stationary policies only. In [9], the authors restrict attention to threshold policies: these policies are characterized by a single number K_T: a call is admitted in state k if and only if $k < K_T$. In the same vein, below we discuss optimality in the class of threshold policies. Our goal here is to show how the three thresholds K_{C}, K_{A}, and K_{I} relate. In section 7 we provide an efficient numerical procedure for calculating K_{I}.

All the results of this section apply for any function g_ℓ satisfying assumption 1, and we shall denote, e.g., by K_{A} the threshold corresponding to the function under consideration. But first, a few comments on these policies.

1. The optimal policy may well be randomized; this is true of both the individual and the aggressive criteria. The interpretation of randomization when the individual criterion is used goes as follows. If a call is admitted, then its performance guarantee should be met: that is, the averaging is done under the condition that the call is accepted. However, we average with respect to later arrivals to that state. Thus, the randomization can be used to make the performance guarantee a tight bound, thereby improving the call blocking part.

2. The average ("aggressive") bound K_{A} may be infinite: this would be the case if, for example, the arrival rate is very low. This is not the case for the individual criterion, as demonstrated by lemma 3 below.

Lemma 3. Let g satisfy assumption 1. For the individual criterion, there exists a number $\overline{K}_{\mathrm{I}} < \infty$ so that, under any feasible policy, a call arriving to state $k > \overline{K}_{\mathrm{I}}$ is always blocked.

Proof. Note that the definitions (5.15)–(5.29) make sense under any stationary policy; however, they are based on our uniformized process. But when no threshold is imposed, there is no upper bound on the maximal rate of events. We shall therefore need to work with the original process $k(t)$, and not with the uniformized process.

So, if $g = g_0$, let $V(K, k_0)$ denote the difference between the average amount of fluid lost, and the allowed loss by the average number of calls in the system, both over the duration of the "marked" call. This value may be calculated through either the original, continuous-time process as in (8.2) (although we do not provide the explicit formulas here) or through the uniformized process. For general g, $V(K, k_0)$ would be defined as a sum if the computation is done for the uniformized process, or the corresponding integral in continuous time. $V(\infty, k_0)$ is defined only through the continuous-time process, and we need to show that $V(K, k_0)$ is well defined (including the case $K = \infty$), and that $V(\infty, k_0) > 0$ for all k_0 large enough. That is, unless a threshold is imposed, the individual cell-loss guarantee for a call which is accepted when the system is very busy would be violated. We shall in fact establish this under *any stationary policy*.

A stationary policy is specified through a sequence $\eta = \{\eta_1, \eta_2, \ldots\}$, where η_k is the probability of accepting a call when there are k calls in the system. So fix a stationary policy, and an arbitrary state k_1 so that $K_{\mathrm{C}} < k_1 < k_0$. Let τ denote the time that the marked call leaves (in the continuous time process!), and τ_{k_1} the first time the process enters state k_1. Denote $\underline{g} = -\inf_k g(k)$. Then $\underline{g}$ is positive and finite, and under any stationary policy,

$$V(K, k_0) \geqslant g(k_1) E^K_{k_0}(\tau \wedge \tau_{k_1}) - \underline{g} E^K_{k_0}\big((\tau - \tau_{k_1})\mathbb{1}(\tau > \tau_{k_1})\big) \tag{6.1}$$

where we used the fact that $g(k) \geqslant g(k_1)$ for all $k \geqslant k_1$. Equation (6.1) holds for all K including the case $K = \infty$. Since the second term is bounded by $\underline{g} E^K_{k_0}\tau$ and the first term is positive, we conclude that $V(K, k_0)$ is well defined and bounded below. Now set $K = \infty$. Since τ_{k_1} increases monotonically to $+\infty$ as $k_0 \to \infty$ while τ is not changed, we have

$$\underline{g} E^K_{k_0}\big((\tau - \tau_{k_1})\mathbb{1}(\tau > \tau_{k_1})\big) \to 0 \quad \text{as } k_0 \to \infty \tag{6.2}$$

by the monotone convergence theorem, and, in addition, $E^K_{k_0}(\tau \wedge \tau_{k_1}) \to E^K_{k_0}\tau$. The result follows since $g(k_0)$ is positive for $k_0 > K_{\mathrm{C}}$. □

Theorem 4. If g satisfies assumption 1 then $K_{\mathrm{C}} \leqslant K_{\mathrm{I}} \leqslant K_{\mathrm{A}}$.

Proof. We prove the two inequalities separately. The first one is trivial: by the definition of K_{C},

$$g(k) \leqslant 0, \quad 0 \leqslant k \leqslant K_{\mathrm{C}},$$

so that

$$\sum_{k=0}^{K} g(k) {}^n_a\pi^{(k_0)}_k \leqslant 0, \quad 1 \leqslant k_0 \leqslant K \leqslant K_{\mathrm{C}}.$$

It follows immediately from the definition of K_{I} that $K_{\mathrm{C}} \leqslant K_{\mathrm{I}}$.

The proof of the second inequality requires more effort. Our proof is by contradiction. We assume that $K_{\mathrm{A}} + 1$ is feasible for the individual call based QoS, and show a contradiction.

With $K = K_{\mathrm{A}} + 1$ feasible we have

$$V(K, k_0) \leqslant 0, \quad 1 \leqslant k_0 \leqslant K.$$

This implies that

$$\sum_{k=1}^{K} g(k) {}^n_a\pi^{(k_0)}_k \leqslant 0, \quad 1 \leqslant k_0 \leqslant K. \tag{6.3}$$

Because $X(m) = k_u(m)$ for $0 \leqslant m < \tau$ we can write

$${}^n_a\pi^{(k_0)}_k = \frac{E^K_{k_0}\big[\sum_{m=0}^{\tau-1} \mathbb{1}(k_u(m) = k)\big]}{E^K_{k_0}[\tau]}.$$

The probability that the marked call leaves in the next step does not depend on the state, so $E^K_{k_0}[\tau]$ does not depend on k_0. Thus

$$ {}^{n}_{a}\pi_k^{(k_0)} = \frac{E^K_{k_0}\big[\sum_{m=0}^{\tau-1} \mathbb{1}(k_u(m)=k)\big]}{E^K[\tau]}. $$

Consider the following discrete time Markov chain on the state space $\{1,\ldots,K = K_{\mathrm{A}}+1\}$. The dynamics of the process are the same as $\{X(m),\ m \geqslant 1\}$ until τ. At time τ, instead of being absorbed in state 0 the process jumps to state $k \in \{1,\ldots,K\}$ with probability γ_k, and one call is chosen as "marked". (We assume that $\gamma_k \geqslant 0$, $1 \leqslant k \leqslant K$ and $\sum_{k=1}^K \gamma_k = 1$.) Let $\boldsymbol{\gamma} = (\gamma_1,\ldots,\gamma_K)$, and let $\psi_k(\boldsymbol{\gamma})$ denote the invariant (steady-state) probability of this process being in state k. A simple regenerative argument enables us to write

$$ \psi_k(\boldsymbol{\gamma}) = \frac{\sum_{k_0=1}^K \gamma_{k_0} E^K_{k_0}\big[\sum_{m=0}^{\tau-1} \mathbb{1}(k_u(m)=k)\big]}{E^K[\tau]} = \sum_{k_0=1}^K \gamma_{k_0}\, {}^{n}_{a}\pi_k^{(k_0)}, \quad 1 \leqslant k \leqslant K. \tag{6.4} $$

There is another way to calculate $\psi_k(\boldsymbol{\gamma})$, $1 \leqslant k \leqslant K$, namely through the one step transition probabilities for the Markov chain, which we denote by $\{r_{ij},\ 1 \leqslant i, j \leqslant K\}$. Based on the description of the process given above we can write

$$ \begin{aligned}
&r_{ii+1} = \Lambda^{-1}[\lambda + \mu\gamma_{i+1}] = p_{ii+1} + \Lambda^{-1}\mu\gamma_{i+1}, && 1 \leqslant i \leqslant K-1,\\
&r_{ii-1} = \Lambda^{-1}\big[(i-1)\mu + \mu\gamma_{i-1}\big] = p_{ii-1} + \Lambda^{-1}\mu(\gamma_{i-1}-1), && 1 < i \leqslant K,\\
&r_{ii} = \Lambda^{-1}\big[(\Lambda - \lambda - i\mu) + \mu\gamma_i\big] = p_{ii} + \Lambda^{-1}\mu\gamma_i, && 1 \leqslant i \leqslant K-1,\\
&r_{KK} = \Lambda^{-1}\big[(\Lambda - K\mu) + \mu\gamma_K\big] = p_{KK} + \Lambda^{-1}\mu\gamma_K,\\
&r_{ij} = \Lambda^{-1}\mu\gamma_j, && \text{otherwise.}
\end{aligned} $$

This is a finite state, ergodic, aperiodic Markov chain, so $(\psi_1(\boldsymbol{\gamma}),\ldots,\psi_K(\boldsymbol{\gamma}))$ is the unique solution to

$$ \sum_{i=1}^K \psi_i(\boldsymbol{\gamma}) r_{ij} = \psi_j(\boldsymbol{\gamma}), \quad 1 \leqslant j \leqslant K, \tag{6.5} $$

with

$$ \sum_{i=1}^K \psi_i(\boldsymbol{\gamma}) = 1. \tag{6.6} $$

Let $\boldsymbol{\pi}^{(0)} = (\pi_1^{(0)},\ldots,\pi_K^{(0)})$ be the probability vector obtained from $\boldsymbol{\pi} = (\pi_0, \pi_1, \ldots, \pi_K)$ by renormalizing:

$$ \pi_i^{(0)} = \frac{\pi_i}{1-\pi_0}, \quad 1 \leqslant i \leqslant K. $$

Set

$$\begin{aligned}
\hat{\gamma}_1 &= \pi_1^{(0)} + \pi_2^{(0)},\\
\hat{\gamma}_j &= \pi_{j+1}^{(0)}, \quad 1 < j \leqslant K-1,\\
\hat{\gamma}_K &= 0.
\end{aligned}$$

We now verify that $\boldsymbol{\pi}^{(0)}$ solves (6.5), with $\boldsymbol{\gamma} = \hat{\boldsymbol{\gamma}}$, so that $\psi_i(\hat{\boldsymbol{\gamma}}) = \pi_i^{(0)}$, $1 \leqslant i \leqslant K$. We have

$$\begin{aligned}
\sum_{i=1}^{K} \pi_i^{(0)} r_{ij} &= (1-\pi_0)^{-1} \sum_{i=1}^{K} \pi_i r_{ij}\\
&= (1-\pi_0)^{-1} \sum_{i=1}^{K} \pi_i \big(p_{ij} + \Lambda^{-1}\mu\hat{\gamma}_j - \Lambda^{-1}\mu \mathbb{1}_{\{j=i-1\}}\mathbb{1}_{\{j\leqslant K-1\}}\big)\\
&= (1-\pi_0)^{-1}\left[\sum_{i=1}^{K} \pi_i p_{ij} + (1-\pi_0)\Lambda^{-1}\mu\hat{\gamma}_j - \Lambda^{-1}\mu\pi_{j+1}\mathbb{1}_{\{j\leqslant K-1\}}\right]\\
&= (1-\pi_0)^{-1}\big[\pi_j - \pi_0 p_{0j} + (1-\pi_0)\Lambda^{-1}\mu\hat{\gamma}_j - \Lambda^{-1}\mu\pi_{j+1}\mathbb{1}_{\{j\leqslant K-1\}}\big]\\
&= \begin{cases}
(1-\pi_0)^{-1}\big[\pi_1 - \pi_0\Lambda^{-1}\lambda + \Lambda^{-1}\mu(\pi_1+\pi_2) - \Lambda^{-1}\mu\pi_2\big], & j = 1,\\
(1-\pi_0)^{-1}\big[\pi_j + \Lambda^{-1}\mu\pi_{j+1} - \Lambda^{-1}\mu\pi_{j+1}\big], & 1 < j \leqslant K-1,\\
(1-\pi_0)^{-1}\pi_K, & j = K.
\end{cases}
\end{aligned}$$

For $1 < j \leqslant K-1$ and for $j = K$ it is immediate that the expression is equal to $\pi_j^{(0)}$. This same conclusion follows for $j = 1$ upon noting that $\pi_0 = \pi_1\mu/\lambda$.

We are now finally ready to display the contradiction. As a consequence of $\boldsymbol{\pi}^{(0)}$ satisfying (6.5) with $\boldsymbol{\gamma} = \hat{\boldsymbol{\gamma}}$, we can write, using (6.4),

$$\pi_k^{(0)} = \sum_{k_0=1}^{K} \hat{\gamma}_{k_0} {}_a^n \pi_k^{(k_0)}, \quad 1 \leqslant k \leqslant K. \tag{6.7}$$

This enables us to write (using $g(0) = 0$)

$$\begin{aligned}
\sum_{k=0}^{K} g(k)\pi_k &= (1-\pi_0)\sum_{k=1}^{K} g(k) \sum_{k_0=1}^{K} \hat{\gamma}_{k_0} {}_a^n \pi_k^{(k_0)}\\
&= (1-\pi_0)\sum_{k_0=1}^{K} \hat{\gamma}_{k_0} \sum_{k=1}^{K} g(k)_a^n \pi_k^{(k_0)}.
\end{aligned}$$

By (6.3) we can conclude that

$$\sum_{k=0}^{K} g(k)\pi_k \leqslant 0,$$

which contradicts the definition of K_{A}. Thus $K_{\mathrm{I}} \leqslant K_{\mathrm{A}}$. □

By definition, in order to compute K_{I} we need to evaluate $\sum_{k=0}^{K} g(k)^n_a \pi_k^{(k_0)}$ for $1 \leqslant k_0 \leqslant K$. It actually suffices to evaluate this for $k_0 = K$, as we now show.

Lemma 5. Assume g satisfies assumption 1. If $K > K_{\mathrm{C}}$ and $V(K, K) \leqslant 0$ then $V(K, k_0) \leqslant 0$ for $1 \leqslant k_0 \leqslant K$.

Proof. We assume that $K > K_{\mathrm{C}}$ throughout the proof. Let

$$W(K, k_0) = E_{k_0}^K \sum_{m=0}^{(\tau_{K_{\mathrm{C}}} \wedge \tau)-1} g\big(X(m)\big),$$

where $\tau_{K_{\mathrm{C}}} = \min\{m\colon\ X(m) = K_{\mathrm{C}}\}$. We can write

$$V(K, k_0) = W(K, k_0) + P_{k_0}^K(\tau_{K_{\mathrm{C}}} < \tau) V(K, K_{\mathrm{C}}). \tag{6.8}$$

Consider first the case $k_0 > K_{\mathrm{C}}$. By assumption, $g(k)$ is positive and increasing when restricted to $k > K_{\mathrm{C}}$. Thus,

$$W(K, k_0) \leqslant W(K, k_0 + 1), \quad K_{\mathrm{C}} < k_0 < K, \tag{6.9}$$

and $W(K, K) > 0$. In addition,

$$P_{k_0+1}^K(\tau_{K_{\mathrm{C}}} < \tau) \leqslant P_{k_0}^K(\tau_{K_{\mathrm{C}}} < \tau), \quad K_{\mathrm{C}} < k_0 < K. \tag{6.10}$$

Using (6.8) with $k_0 = K$ yields

$$V(K, K_{\mathrm{C}}) = \frac{V(K, K) - W(K, K)}{P_K^K(\tau_{K_{\mathrm{C}}} < \tau)}. \tag{6.11}$$

Thus (since we have $V(K, K) \leqslant 0$ by assumption)

$$V(K, K_{\mathrm{C}}) \leqslant 0. \tag{6.12}$$

We can write

$$\begin{aligned} &V(K, k_0) - V(K, k_0 + 1) \\ &\quad = W(K, k_0) - W(K, k_0 + 1) + \big[P_{k_0}^K(\tau_{K_{\mathrm{C}}} < \tau) - P_{k_0+1}^K(\tau_{K_{\mathrm{C}}} < \tau)\big] V(K, K_{\mathrm{C}}). \end{aligned} \tag{6.13}$$

Combining (6.9), (6.10), and (6.12) with (6.13) yields

$$V(K, k_0) \leqslant V(K, k_0 + 1), \quad K_{\mathrm{C}} < k_0 < K. \tag{6.14}$$

By (6.14), the assumption that $V(K, K) \leqslant 0$, and (6.11),

$$V(K, k_0) \leqslant 0, \quad K_{\mathrm{C}} \leqslant k_0 \leqslant K.$$

For $k < K_{\mathrm{C}}$, $g(k) \leqslant 0$. Thus

$$W(K, k_0) \leqslant 0, \quad 1 \leqslant k_0 < K_{\mathrm{C}}. \tag{6.15}$$

Combining (6.8), (6.12), and (6.15) yields

$$V(K, k_0) \leqslant 0, \quad 1 \leqslant k_0 < K_{\mathrm{C}}. \qquad \square$$

7. Identical users: computation of optimal policies

By theorem 4, $K_{\mathrm{I}} \geqslant K_{\mathrm{C}}$. By lemma 5, in order to calculate K_{I}, it suffices to compute $V(K, K)$ for $K \geqslant K_{\mathrm{C}}$. Although it is possible to solve for $V(K, K)$ using (5.29) (with suitable truncation of the infinite sum in (5.9)), there is a substantially more efficient technique, using difference equations. We first derive a second order difference equation for $V(K, k)$, and then provide a solution that is simple to evaluate numerically.

Recall that

$$V(K, k) = E_k^K \sum_{m=0}^{\tau-1} g\big(X(m)\big).$$

With the understanding that $p^a_{K\,K+1} = 0$ and $V(K, 0) = 0$ we can write

$$V(K, k) = g(k) + p^a_{kk+1} V(K, k+1) + p^a_{kk} V(K, k) + p^a_{kk-1} V(K, k-1). \tag{7.1}$$

Using equation (5.8), (7.1) becomes

$$\begin{aligned} V(K, k) = g(k) &+ \Lambda^{-1}\lambda V(K, k+1) + \Lambda^{-1}(\Lambda - \lambda - k\mu) V(K, k) \\ &+ \Lambda^{-1}(k-1)\mu V(K, k-1), \quad 1 \leqslant k < K \end{aligned} \tag{7.2}$$

and

$$V(K, K) = g(K) + \Lambda^{-1}\lambda V(K, K) + \Lambda^{-1}(K-1)\mu V(K, K-1). \tag{7.3}$$

These can be rewritten as

$$(\lambda + k\mu) V(K, k) = \Lambda g(k) + \lambda V(K, k+1) + (k-1)\mu V(K, k-1), \quad 1 \leqslant k \leqslant K \tag{7.4}$$

or

$$V(K, k+1) - \frac{\lambda + k\mu}{\lambda} V(K, k) + \frac{(k-1)\mu}{\lambda} V(K, k-1) = -\frac{\Lambda}{\lambda} g(k) \tag{7.5}$$

and

$$V(K, 0) = 0, \tag{7.6}$$

$$V(K, K) - \frac{(K-1)}{K} V(K, K-1) = \frac{\Lambda}{K\mu} g(K). \tag{7.7}$$

This is a second order linear recursion with boundary conditions. To obtain a solution, consider a general second-order linear recursion with varying coefficients, of the form

$$x_{t+1} + a_t^0 x_t + a_t^1 x_{t-1} = b_t^0 u_t + b_t^1 u_{t-1}, \quad 1 \leqslant t \leqslant T - 1. \tag{7.8}$$

Using the vector notation $\mathbf{x}_t = \binom{x_t}{x_{t-1}}$, $\mathbf{u}_t = \binom{u_t}{u_{t-1}}$, write (7.8) in matrix form as

$$\mathbf{x}_{t+1} = A_t \cdot \mathbf{x}_t + B_t \cdot \mathbf{u}_t, \quad 1 \leqslant t \leqslant T-1, \tag{7.9}$$

$$A_t = \begin{pmatrix} -a_t^0 & -a_t^1 \\ 1 & 0 \end{pmatrix}, \qquad B_t = \begin{pmatrix} b_t^0 & b_t^1 \\ 0 & 0 \end{pmatrix}. \tag{7.10}$$

Define the "fundamental matrix" by $\Phi(s,t) = A_{t-1} \cdot A_{t-2} \cdots A_s$. Any solution to (7.9) must satisfy

$$\mathbf{x}_T = \Phi(1,T) \cdot \mathbf{x}_1 + \sum_{t=2}^{T-1} \Phi(t,T) B_{t-1} \mathbf{u}_{t-1} + B_{T-1} \mathbf{u}_{T-1}. \tag{7.11}$$

Returning to our specific equation, we use (7.11) with $T = K$. We have

$$A_k = \begin{pmatrix} \dfrac{\lambda + k\mu}{\lambda} & -\dfrac{(k-1)\mu}{\lambda} \\ 1 & 0 \end{pmatrix}, \tag{7.12}$$

$$B_k \mathbf{u}_k = \begin{pmatrix} -\Lambda/\lambda \\ 0 \end{pmatrix} g(k). \tag{7.13}$$

The solution of (7.5)–(7.7) then satisfies

$$\begin{pmatrix} V(K,K) \\ V(K,K-1) \end{pmatrix} = \Phi(1,K) \begin{pmatrix} V(K,1) \\ 0 \end{pmatrix} + \sum_{k=2}^{K-1} \Phi(k,K) \begin{pmatrix} -\Lambda/\lambda \\ 0 \end{pmatrix} g(k-1) - \begin{pmatrix} -\Lambda/\lambda \\ 0 \end{pmatrix} g(K-1). \tag{7.14}$$

Equations (7.7), (7.14) now give us a system of three linear equations in the three unknowns $V(K,1)$, $V(K,K-1)$ and $V(K,K)$. The computation of $\Phi(k,K)$ is straightforward as it depends only on known values of the parameters. Note that this generic calculation is independent of the form of $g(k)$.

Lemma 6. There exists a unique solution to (7.5)–(7.7), which agrees with the unique solution of the system of equations (7.7) and (7.14).

Proof. By construction, any solution $\{V(K,k),\ k = 0,1,\ldots,K\}$ of (7.5)–(7.7) determines a solution $\{V(K,1),\ V(K,K-1),\ V(K,K)\}$ of (7.7) and (7.14). Conversely, a solution of (7.7) and (7.14) determines a solution of (7.5)–(7.7). Thus we can consider either, as convenient. To simplify the notation, abbreviate (7.7) and (7.14) as

$$v(K) - ev(K-1) = g', \tag{7.15}$$

$$\begin{pmatrix} v(K) \\ v(K-1) \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} v(1) \\ 0 \end{pmatrix} + G \tag{7.16}$$

for appropriate constants a, b, c, d, e, g' and matrix G. Using (7.15) to eliminate $v(K-1)$ from (7.16): denoting by G_1, G_2 the elements of the vector G,

$$v(K) = av(1) + G_1, \tag{7.17}$$

$$v(K-1) = cv(1) + G_2 \tag{7.18}$$

$$= e^{-1}\big(v(K) - g'\big) \tag{7.19}$$

and so

$$v(K) - g' = ecv(1) + eG_2. \tag{7.20}$$

Therefore, there exists a unique solution $v(1)$, $v(K)$ to the resulting equation if and only if $ae \neq c$. This is therefore equivalent to existence and uniqueness of solution to (7.15)–(7.16) but with $g' = 0$ and $G = 0$ (the zero matrix). So, translating back to the original equations, there exists a unique solution if and only if this holds when $g(k) = 0$ for all k.

In this case, however, $V(K, k) = 0$ for all k is a solution to the difference equation. Assume that it is not unique, at let $v^0(k)$ be the postulated nonzero solution. Then $v^0(1) \neq 0$, for otherwise the difference equation will dictate $v^0(k) = 0$ for all k. So, to complete the proof it suffices to show that $(ae - c)v^0(1) \neq 0$. However, note that when $g' = 0$ and $G = 0$,

$$av^0(1) = v^0(K), \tag{7.21}$$

$$cv^0(1) = v^0(K-1), \tag{7.22}$$

so that $(ae - c)v^0(1) = 0$ implies

$$v^0(K) \cdot \frac{K-1}{K} = v^0(K-1). \tag{7.23}$$

However, by (7.15),

$$v^0(K) = \frac{K-1}{K} v^0(K-1) \tag{7.24}$$

which can hold simultaneously only if $v^0(K) = v^0(K-1) = 0$. But in this case, the recursion equation implies that $v^0(k) = 0$ for all k, contradicting our assumption on v^0. □

The computational procedure is therefore the following. Starting with $\Phi(K, K) = I$ (the identity matrix), we calculate (backwards) $\Phi(k, K) = A_k \Phi(k+1, K)$, and add the term $\Phi(k, K)\binom{-\Lambda/\lambda}{0} g(k)$ to the sum in (7.14). After K steps, using the notation of (7.15)–(7.16), we have computed all the constants of those equations. From (7.17)

$$v(1) = \frac{g' - G_1 + eG_2}{a - ce}, \tag{7.25}$$

so we can add the final term of (7.14), and then solve explicitly for $v(K) = V(K, K)$.

The value $V(K, K)$ is then compared to the QoS guarantee, and the calculation repeats with a new K until the threshold is identified.

We provide numerical examples that correspond to parameters considered in [7,9]. In all cases we take $R = 45$ and $f = 10^{-9}$. We first consider $\mu = 0.1$, $\nu = 6.0$ and $p = 0.025$. For $\lambda = 1.125$ (which corresponds, according to [7], to the largest value of λ

for which the call blocking probability does not exceed 0.01; this is called $\lambda_{\max}$ in [9]), we obtain $K_C = 14$, $K_I = 16$ and $K_A = 20$. If we reduce the call arrival rate to $\lambda = 1.05$, then $K_A = \infty$ and $K_I = 16$ (and $K_C = 14$ because it does not depend on λ). The second case we consider has $\mu = 1.0$, $\nu = 1.5$, and $p = 0.1$. Taking $\lambda = 94.0$ (corresponding to $\lambda_{\max}$ for these parameters), we obtain $K_C = 100$, $K_I = 105$ and $K_A = 111$. Taking $\lambda = 85.0$ yields $K_A = \infty$ and $K_I = 107$. All of these numbers were calculated in a fraction of a second on a MIPs 4400/4010 processor running at 150 MHz.

8. Admission to an ATM link: the mathematical program

In this section we outline the derivation of a mathematical program, whose solution provides an optimal control for a general admission problem. Recall the notation of sections 2–4. With some abuse of notation (but hopefully no ensuing confusion) we denote by $\mathbf{k}_n$ and a_n also the state (respectively, action) just before the nth transition occurred. Let $N(t)$ be the number of events (call arrivals and departures) up to time t. The average reward corresponding to call blocking, under the control policy σ with initial state $\mathbf{k}$ can be written as

$$V_0(\mathbf{k}, \sigma) = \liminf_{T\to\infty} \frac{1}{T} \mathbb{E}^{\sigma}_{\mathbf{k}} \sum_{n=1}^{N(T)} r(\mathbf{k}_n, a_n). \tag{8.1}$$

Representing the constraints on cell loss is more involved, for two reasons. First, unlike call blocking, cell loss depends on the duration of events, since our functions g_i typically represent cell-loss rates. Secondly, our point of view is that of an individual call: we would like to provide a guarantee in terms of performance during the time a marked call is in the system. We shall in fact provide a separate constraint for each call type and for each possible state at which such a call may be accepted. Let

$$V_{\text{tot}}(\mathbf{k}, \sigma) = \mathbb{E}^{\sigma}_{\mathbf{k}} \int_0^{\tau} g_i(\mathbf{k}_t, a_t)\, \mathrm{d}t, \tag{8.2}$$

where τ is the time the marked call leaves. This quantity should fall below the guaranteed level for each state $\mathbf{k}$ in which new type i calls are accepted. Since the same guarantee is offered to all calls accepted in a given state, we must restrict our attention to stationary policies. Note, however, that we cannot use the standard uniformization approach of the theory of Semi-Markov Decision processes (SMDP), since our constraints may lead to randomized policies, under which the uniformization technique does not apply, as shown by Beutler and Ross [3].

To provide a mathematical program that combines the above objective and constraints we use the specific structure of our problem and then apply the general results of the appendix. In order to formulate the constraint pertaining to a type i call accepted in state $\mathbf{k}_s$, we construct a modified process with transition probabilities $p_{s,i}(\mathbf{k}, \mathbf{k}'; a)$ (to avoid cumbersome notation, we use s, i rather than $\mathbf{k}_s$, i as subscript or superscript). Because of the exponential holding time assumption, whenever an event occurs, the

probability that a specific call of type i (the "marked call") leaves the system equals the service rate for that type, divided by the total event rate in that state. Under the modified transition structure, in place of this event, the process jumps back to the initial state $\mathbf{k}_s$, a new call of the same type is chosen as "marked", and the process continues. Thus, we create a regeneration point for a process that would otherwise terminate when the marked call leaves.

Let $\mathbb{E}$ denote expectation with respect to the modified transition probabilities. Although $\mathbb{E}$ depends on the initial state $\mathbf{k}_s, i$ we omit this dependence from the notation: assume $\mathbf{k}_s$ and i are fixed. Let τ_ℓ denote the ℓth time that a marked call leaves the modified system. Let the random variable N_ℓ denote the total number of events until time τ_ℓ. Since we are restricting to stationary policies σ, we have

$$V_{\text{tot}}(\sigma) = \frac{1}{\ell}\mathbb{E}\int_0^{\tau_\ell} g_i(\mathbf{k}_t, a_t)\,\mathrm{d}t \tag{8.3}$$

$$= \frac{1}{\ell}\mathbb{E}\sum_{n=1}^{N_\ell} g_i(\mathbf{k}_n, a_n)\tau(\mathbf{k}_n, a_n), \tag{8.4}$$

where $\tau(\mathbf{k}, a)$ is the expected time for transition out of state $\mathbf{k}$ under action a. Since σ is a stationary policy and by definition of τ_ℓ, we have $\mathbb{E}\tau_\ell = \ell \cdot \mathbb{E}\tau_1$. Clearly, if the marked call is of type i then $\mathbb{E}\tau_1 = 1/\mu_i$. So

$$V_{\text{tot}}(\sigma) = \mathbb{E}\tau_1 \frac{\mathbb{E}\sum_{n=1}^{N_\ell} g_i(\mathbf{k}_n, a_n)\tau(\mathbf{k}_n, a_n)}{\mathbb{E}\sum_{n=1}^{N_\ell} \tau(\mathbf{k}_n, a_n)} \tag{8.5}$$

$$= \lim_{\ell\to\infty} \frac{(1/\mu_i)\mathbb{E}\sum_{n=1}^{\ell} g_i(\mathbf{k}_n, a_n)\tau(\mathbf{k}_n, a_n)}{\mathbb{E}\sum_{n=1}^{\ell} \tau(\mathbf{k}_n, a_n)} \tag{8.6}$$

which is exactly a ratio-average criterion [5].

We now apply the results of [5] to derive a mathematical program for the optimal policies for our call admission problem. In the appendix we describe this derivation by treating a discrete-time problem. Our semi-Markov problem is treated in the same way, by combining the results of [5] with those of the appendix. Our policy σ will be determined from the solution $z^0_{\mathbf{k}a}$ of the mathematical program as

$$\sigma(a \mid \mathbf{k}) = \frac{z^0_{\mathbf{k}a}}{\sum_{a'} z^0_{\mathbf{k}a'}} \tag{8.7}$$

if $\sum_{a'} z^0_{\mathbf{k}a'} \neq 0$, and is set arbitrarily otherwise so that $\sum_a \sigma(a \mid \mathbf{k}) = 1$. As in the appendix, the mathematical program is derived from the linear program of [5], with the additional condition that the policies ensuing from each of the constituting variables must agree. See also [9] for details of a related linear program. Let B_i denote the cell level quality of service guarantee (bound) to calls of type i; we say that the cell loss guarantee is met if

$$V_{\text{tot}}(\sigma) \leqslant B_i \quad \text{for all } \mathbf{k}_s \text{ and } i. \tag{8.8}$$

As in the appendix, the mathematical program is stated in terms of the variables $\{z^0_{\mathbf{k}a}, \mathbf{k}, a\}$ that arise from the average criterion $V_0(\mathbf{k}, \sigma)$, as well as the variables $\{z^{s,i}_{\mathbf{k}a}, \mathbf{k}, a\}$, which arise from the $V_{\text{tot}}(\mathbf{k}, \sigma)$. There is one set of such variables for each type i, as well as for each initial state $\mathbf{k}_s$. A priori, the number of states (as well as initial states) is infinite. However, applying the proof of lemma 3 below (section 6) to each type while rejecting all other types gives an upper bound, so that we can limit each type to a finite number of users. This reduces the problem to a finite state one. The mathematical program can now be stated as follows.

$$\max \sum_{\mathbf{k},a} r(\mathbf{k}, a) z^0_{\mathbf{k}a} \tag{8.9}$$

subject to

$$\sum_a z^0_{\mathbf{k}'a} - \sum_{\mathbf{k},a} p\big(\mathbf{k}, \mathbf{k}'; a\big) z^0_{\mathbf{k}a} = 0 \quad \text{for all } \mathbf{k}' \tag{8.10}$$

$$\sum_{\mathbf{k},a} \tau(\mathbf{k}, a) z^0_{\mathbf{k}a} = 1, \tag{8.11}$$

$$z^0_{\mathbf{k}a} \geqslant 0 \quad \text{for all } \mathbf{k}, a \tag{8.12}$$

and, with $\mathbf{k}_s$ denoting initial state, subject to

$$\sum_a z^{s,i}_{\mathbf{k}'a} - \sum_{\mathbf{k},a} p_{s,i}\big(\mathbf{k}, \mathbf{k}'; a\big) z^{s,i}_{\mathbf{k}a} = 0 \qquad \text{for all } i, \mathbf{k}_s, \mathbf{k}', \tag{8.13}$$

$$\sum_{\mathbf{k},a} \tau(\mathbf{k}, a) z^{s,i}_{\mathbf{k}a} = 1 \qquad \text{for all } i, \mathbf{k}_s, \tag{8.14}$$

$$z^{s,i}_{\mathbf{k}a} \geqslant 0 \qquad \text{for all } \mathbf{k}, a, i, \mathbf{k}_s, \tag{8.15}$$

$$\sum_{\mathbf{k},a} \frac{1}{\mu_i} g_i(\mathbf{k}, a) \tau(\mathbf{k}, a) z^{s,i}_{\mathbf{k}a} \leqslant B_i \qquad \text{for all } i, \mathbf{k}_s, \tag{8.16}$$

$$\frac{z^{s,i}_{\mathbf{k}a}}{\sum_{a'} z^{s,i}_{\mathbf{k}a'}} = \frac{z^0_{\mathbf{k}a}}{\sum_{a'} z^0_{\mathbf{k}a'}} \quad \text{for all } \mathbf{k}, a, i, \mathbf{k}_s. \tag{8.17}$$

This mathematical program may be theoretically quite satisfying, in that any optimal solution may be obtained by solving this program (see the appendix). However, from an implementation point of view, the task is quite formidable: most significant are the nonlinear constraints (8.17), for which no theory exists. In addition, the size of this program is quite large: the number of variables z^α_β is roughly twice (two actions) the number of types J times the square of the number of possible states. In addition, the number of both linear and nonlinear constraints is of the same order as the number of variables. Nonetheless, we believe that the equivalence between the non standard SMDP problem and the mathematical program is of interest, and further research into the consequences of this equivalence is warranted.

Acknowledgements

Research of A.S. was performed in part while he was visiting the Mathematical Sciences Research Center, Bell Laboratories, Murray Hill, and later also the Graduate School of Business, Rutgers University. Their support and hospitality are greatly appreciated. The research of A.S. was also supported in part by the fund for promotion of research at the Technion, in part by the consortium for broadband communication administered by the Chief Scientist of the Israeli Ministry of Commerce and Industry, and in part by the Israel Science Foundation, administered by the Israel Academy of Sciences and Humanities.

Appendix: Markov decision problems and linear programs

In this appendix we collect some results from the theory of MDPs in order to make the presentation more self contained. The results herein are either known, or immediate extensions of known results, but some of the results have not been published. This appendix serves both to illustrate as well as substantiate the derivation, in section 8, of the mathematical program.

Consider first the following optimization problem involving a finite state, finite action (discrete time) Markov decision process. We are given transition probabilities for the controlled process: $p(i, j; a)$ is the probability that the next state is j given that the current state is i and action a is chosen. We are also given immediate costs $c(i, a)$ incurred when action a is taken in state i. Let σ be a stationary randomized policy: $\sigma = \{\sigma(a \mid i)\}_{i,a}$, where $\sigma(a \mid i)$ is the probability of using action a when in state i. Let x_t denote the state at the (discrete) time t and a_t the action at t. Denote the average reward under policy σ, with initial state i by

$$V_{\text{av}}(i, \sigma) = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}_i^\sigma \sum_{t=1}^{T} c(x_t, a_t). \tag{A.1}$$

(The limit exists since the state space is finite, and the policy is stationary.) Assume that under every stationary policy, the resulting Markov chain contains at most one ergodic class (and possibly transient states): such a Markov decision process is called unichain. Then, under any stationary policy and for any initial state i, the limits below exist, are independent of the initial state i and

$$V_{\text{av}}(\sigma) = \lim_{T\to\infty} \mathbb{E}_i^\sigma \frac{1}{T} \sum_{t=1}^{T} c(x_t,\ a_t) \tag{A.2}$$

$$= \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{y,a} \mathbb{P}_i^\sigma(x_t = y,\ a_t = a) c(y, a) \tag{A.3}$$

$$= \sum_{y,a} \left(\lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{P}_i^{\sigma}(x_t = y,\ a_t = a) \right) c(y,a) \tag{A.4}$$

$$= \sum_{y,a} z_{ya}(\sigma) c(y,a), \tag{A.5}$$

where the last equality defines z. The state-action occupation measure z (the $\{z_{ya}\}$ are also known as the state-action frequencies) depends on the policy σ but, due to the ergodicity assumption, does not depend on the initial state. Given z we can recover $\sigma(a \mid y)$ from

$$\sigma(a \mid y) = \frac{z_{ya}}{\sum_{a'} z_{ya'}} \tag{A.6}$$

provided $\sum_a z_{ya} > 0$. If $\sum_a z_{ya} = 0$, then $\sigma(a \mid y)$ is immaterial and can be selected arbitrarily, subject to $\sum_a \sigma(a \mid y) = 1$. Conversely, if we let π^{σ} denote the invariant distribution of the process under policy σ we can compute z by

$$z_{ya} = \pi_y^{\sigma} \cdot \sigma(a \mid y). \tag{A.7}$$

We are interested in problems where additional constraints are imposed, of the form

$$V_m(\sigma) = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}_i^{\sigma} \sum_{t=1}^{T} c_m(x_t, a_t) \leqslant C_m, \quad m = 1, 2, \ldots, M, \tag{A.8}$$

for some immediate cost functions c_m and bounds C_m. Due to (A.2)–(A.5), the $V_m(\sigma)$ can also be written as linear functions of $z(\sigma)$. Consequently, we can write a linear program that corresponds to the Markov decision problem.

Theorem 7 [2,4,6]. Under the unichain assumption, a stationary policy σ is feasible (optimal among stationary policies) for the Markov decision problem

$$\text{Max } V_{\text{av}}(i, \sigma) \quad \text{subject to the constraints (A.8)} \tag{A.9}$$

if and only if $\{z_{ya}\}$ is feasible (optimal, respectively) for the associated linear program:

$$\max \sum_{ya} z_{ya} c(y,a) \tag{A.10}$$

such that

$$z_{ya} \geqslant 0, \qquad \text{all } y, a, \tag{A.11}$$

$$\sum_{ya} z_{ya} = 1, \tag{A.12}$$

$$\sum_{ya} z_{ya} c_m(y,a) \leqslant C_m, \quad m = 1, 2, \ldots, M, \tag{A.13}$$

$$\sum_{a} z_{sa} - \sum_{ya} z_{ya} p(s \mid y, a) = 0, \qquad \text{all } s. \tag{A.14}$$

The occupation measure $\{z_{ya}\}$ and the policy σ are related through (A.6), (A.7).

In section 8 we consider some auxiliary processes, that start at a fixed initial state $\mathbf{k}_s$, and are "restarted" when a marked call leaves. More generally, let each $V_m(\sigma)$ be defined through a different expectation operator $\mathbb{E}^{m,\sigma}$, corresponding to different transition probabilities p^m. Consider now the optimization problem (A.9) where maximization is with respect to stationary policies. Applying theorem 7 we see that for each m there corresponds a $z^m(\sigma)$, and σ is feasible if and only if $z^m(\sigma)$ satisfies the appropriate linear constraints. In addition, any feasible $\{z^m(\sigma),\ m = 1, \ldots, M\}$ must be related to σ via (A.6), (A.7). Denoting the original transition probabilities p by p^0 we have the following result.

Theorem 8. Under the unichain assumption for each m, a stationary policy σ is feasible (optimal among stationary policies) for the Markov decision problem

$$\text{Max } V_{av}(i, \sigma) \tag{A.15}$$

$$\text{subject to } V_m(\sigma) \leqslant C_m, \quad m = 1, 2, \ldots, M, \tag{A.16}$$

if and only if $\{z^m_{ya},\ m = 0, \ldots, M\}$ is feasible (optimal, respectively) for the associated linear program:

$$\max \sum_{ya} z^0_{ya} c(y, a) \tag{A.17}$$

such that

$$z^m_{ya} \geqslant 0, \qquad \text{all } y, a, \text{ and } 0 \leqslant m \leqslant M, \tag{A.18}$$

$$\sum_{ya} z^m_{ya} = 1, \qquad \text{all } 0 \leqslant m \leqslant M, \tag{A.19}$$

$$\sum_{ya} z^m_{ya} c_m(y, a) \leqslant C_m, \qquad 1 \leqslant m \leqslant M, \tag{A.20}$$

$$\sum_{a} z^m_{sa} - \sum_{ya} z^m_{ya} p^m(s \mid y, a) = 0 \qquad \text{all } s \text{ and } 0 \leqslant m \leqslant M, \tag{A.21}$$

$$\frac{z^m_{ya}}{\sum_{a'} z^m_{ya'}} = \frac{z^0_{ya}}{\sum_{a'} z^0_{ya'}} \qquad \text{all } y, a, \text{ and } 1 \leqslant m \leqslant M. \tag{A.22}$$

References

[1] E. Altman and V.A. Gaitsgory, Stability and singular perturbations in constrained Markov decision problems, IEEE Trans. Automat. Control 38 (1993) 971–975.

[2] E. Altman and A. Shwartz, Adaptive control of constrained Markov chains: Criteria and policies, Ann. Oper. Res. 28 (1991) 101–134.

[3] F.J. Beutler and K.W. Ross, Uniformization for Markov decision processes under stationary policies, J. Appl. Probab. 24 (1987) 644–656.

[4] C. Derman, *Finite State Markovian Decision Processes* (Academic Press, New York, 1970).
[5] E.A. Feinberg, Constrained semi-Markov decision processes with average rewards, ZOR – Math. Methods Oper. Res. 39 (1994) 257–288.
[6] A. Hordijk and L.C.M. Kallenberg, Linear programming and Markov decision chains, Managm. Sci. 25 (1979) 352–362.
[7] D. Mitra, M.I. Reiman and J. Wang, Robust admission control for heterogeneous ATM systems with both cell and call QoS requirements, in: *Proc. of ITC 15*, Washington DC, 1997.
[8] R. Nagarajan, J. Kurose and D. Towsley, Finite horizon statistical quality-of-service measures for high-speed networks, J. High Speed Networks 3 (1994) 351–373.
[9] M.I. Reiman, J. Wang and D. Mitra, Dynamic call admission control of an ATM multiplexer with on/off sources, in: *Proc. of 34th Conf. on Decision and Control*, New Orleans, LA, 1995, pp. 1382–1388.