Dynamic Bandwidth Allocation for ATM Switches [1]

Ivy Hsu and Jean Walrand [2]

Department of EECS, University of California, Berkeley, CA 94720

submitted to Journal of Applied Probability, 1994; revised Feb. 5, 1995

Abstract

We explore a dynamic approach to the problems of call admission and resource allocation for communication networks with connections that are differentiated by their quality of service requirements. In a dynamic approach, the amount of spare resources is estimated on-line based on feedbacks from the network's quality of service monitoring mechanism. The schemes we propose remove the dependence on accurate traffic models and thus simplify the tasks of supplying traffic statistics required of network users. In this paper we present two dynamic algorithms. The objective of these algorithms is to find the minimum bandwidth necessary to satisfy a cell loss probability constraint at an Asynchronous Transfer Mode (ATM) switch. We show that in both schemes the bandwidth chosen by the algorithm approaches the optimal value almost surely. Furthermore, in the second scheme, which determines the point closest to the optimal bandwidth from a finite number of choices, the expected learning time is finite.

1 Introduction

The Asynchronous Transfer Mode (ATM) enables broadband networks to support a wide range of communication services. A major advantage of ATM is its statistical multiplexing that allows efficient sharing of network resources (e.g. switch buffers and transmission bandwidth) among users. However, unlike datagram networks, an ATM network is expected to satisfy certain quality of service (QoS) requirements that are agreed upon between the network and the user during the call set-up phase of a connection. This requires a resource allocation scheme that takes into account the traffic statistics and the QoS requirements of each connection.

Prior studies propose the following *static* approach to the resource allocation problem [9, 12]. During the call set-up phase, the user supplies the network with two

[1]This research is supported in part by a grant of the Networking and Communications Research Program of the National Science Foundation.

[2]Email addresses: ihsu@eecs.berkeley.edu, wlr@eecs.berkeley.edu

sets of parameters, one indicating its QoS requirements such as cell loss probability and maximum acceptable delay, and the other its anticipated traffic statistics. If the network has enough spare resources to support the additional traffic without violating the QoS of existing connections, it accepts the call and allocates the needed resources. Otherwise it rejects the new call.

However, we see several problems with this static approach, due to its reliance on the user-supplied traffic parameters: (1) there is not yet a consensus on the choices of traffic parameters that can both effectively characterize all traffic for the purpose of resource allocation and be easy to enforce by the policing mechanism; (2) often the user application may not be able to accurately anticipate its traffic characteristics; (3) the statistics of an incoming traffic can change once it is multiplexed with other cell streams. Thus the same traffic parameters supplied by the user may not be equally accurate for all switches that the connection traverses; (4) the static approach does not provide any feedback for the network to verify if the QoS requirements are indeed satisfied.

One solution to these problems is to require the users to supply only a set of simple parameters, such as the peak and mean cell rates. For example, this can be achieved for traffic regulated by the Leaky Bucket mechanism. The network then makes a conservative call admission and resource allocation decision by assuming the worst-case traffic patterns conforming to these parameters. Despite its simplicity, this conservative policy alone would result in a significant under-utilization of the network resources. Thus we propose to couple it with a *dynamic* mechanism which estimates the true resource requirements based on the network's QoS measurements. This in turn provides the network an estimate of its spare resources at the next call acceptance request. Using direct QoS feedbacks, our algorithms contrast with other dynamic approaches based on model fitting or approximation of moment generating function. We choose this direct approach because it eliminates the complexity of traffic modeling and the inherent inaccuracy present in the indirect methods.

In this paper we propose two dynamic schemes. We place the emphasis on bandwidth allocation, although dynamic buffer allocation can be also incorporated. The QoS parameter of interest is the probability of cell loss due to buffer overflow in an ATM switch. The remainder of this paper is organized as follows: in section 1.1 we introduce the output-buffered switch architecture and show how the contended resources, i.e. output bandwidth and buffer, can be divided among multiple QoS classes. In section 1.2 we present the iterative framework in which dynamic bandwidth estimation can be used to improve utilization. Section 1.3 discusses the Markov modulated fluid model of the traffic used in our analysis. In section 2 we present a dynamic scheme based on an adaptive algorithm. Using a counter to emulate buffer occupancy, the switch can estimate the cell loss probability over a fixed period of time for a given bandwidth. If the fraction of cell loss in this period exceeds the target, the bandwidth over the next period is increased, and vice versa. We show that the

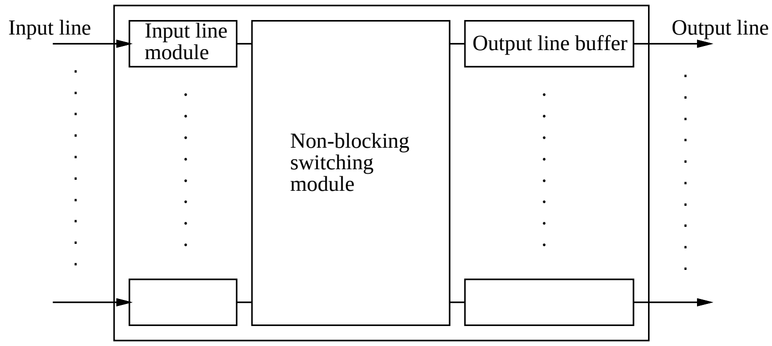

Figure 1: Model of a non-blocking, output-buffered switching node.

bandwidth chosen by this scheme converges to the optimal value almost surely. In section 3 another approach called parallel monitoring is proposed. Its objective is to select the optimal from a finite set of choices. With multiple counters as virtual buffers, the algorithm can monitor the performance of all candidate bandwidths in parallel. We show that this scheme converges to the optimal value within finite expected time. Section 4 offers some concluding remarks.

1.1 Switch Model

Figure 1 illustrates the architecture of an ATM switching node we will consider. It uses a non-blocking switching module that operates at a multiple of the input and output line rate. Thus if cells from several input lines are simultaneously designated for the same output line, they will be routed to that output in one cell time. This requires a buffer to be placed at each output to store the occasional burst of incoming cells. Thus the resources to be shared among users are the output buffer space and the output bandwidth.

Assume that there are a fixed number I of traffic classes that are differentiated by their QoS requirements. For the output line buffer, we adopt the architecture proposed in [9] (see figure 2). The total buffer space is divided into I first-in-first-out buffers (FIFOs), with each FIFO serving one traffic class. Note that the sizes of the FIFOs are not necessarily the same. The output bandwidth is shared among the classes through a *framing* scheme as follows. Time is divided into slots, with each slot equal to one cell transmission time. A fixed number T_{total} of time slots constitute a *frame*. Define Θ as the total bandwidth of an output line. Then class i traffic is guaranteed $\Theta \cdot T_i/T_{\text{total}}$ transmission bandwidth when it is given priority in T_i slots in each frame. [9] suggests that the slots assigned to each class should be evenly distributed in a frame to reduce delay variance. To simplify the analysis, we assume that the size of each FIFO is fixed; only the bandwidth is allocated dynamically. Note that given the buffer size and the bandwidth allocation for a traffic class, an

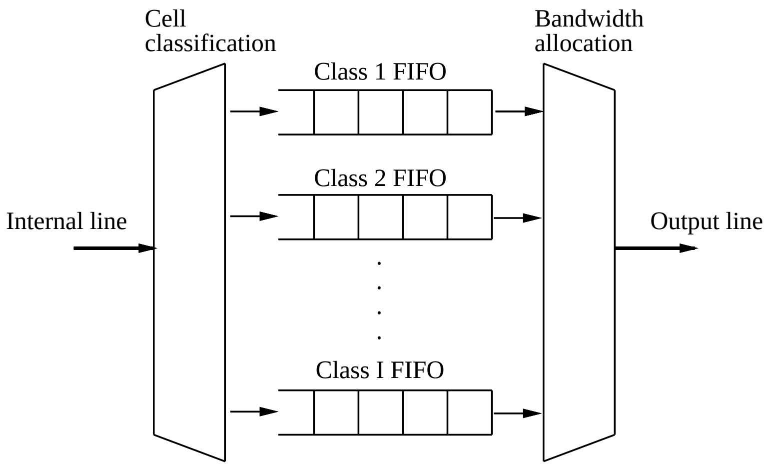

Figure 2: Architecture of the output line buffer.

upper bound on the queueing delay can be guaranteed. This provides a way to constrain delays for delay-sensitive connections. In this paper we focus on the cell loss probability requirement.

1.2 Bandwidth Estimation

The problem of determining the worst-case traffic behavior has been treated in [2, 10, 5, 6]. We briefly state the result of [2] here. First of all, the fraction of cell loss in a bufferless system is used as a conservative estimate of the true cell loss probability. It can be shown that, given the mean and peak rates of the sources that are multiplexed, the worst cell loss probability is achieved when each source is on-off, namely that it alternates between its peak rate and silence, with a probability of being at the peat rate equal to mean/peak.

For simplicity, we assume that sources of each class are homogeneous. Furthermore, we assume that the FIFO of class i is drained at a fluid rate $\Theta_i \equiv \Theta \cdot T_i/T_{\mathrm{total}}$, and ignore the effect introduced by the framing implementation. Thus by the worst case result above, one can obtain a conservative call acceptance boundary $\underline{\mathbf{M}}$ such that for each point on the boundary, $\underline{M} = (M_1, \ldots, M_I) \in \underline{\mathbf{M}}$, there exists a correponding set of bandwidth partition $(\Theta_1, \ldots, \Theta_I)$, $\sum_{i=1}^{I} \Theta_i = \Theta$ at which the cell loss probability constraints $(\alpha_1, \ldots, \alpha_I)$ are satisfied for M_i sources of class i.

This conservative call acceptance boundary can then be refined in an iterative process. Initially, calls can be accepted until a point on the boundary is reached. At this time, a bandwidth estimation algorithm is used to determined the actual bandwidth requirement of each class. Denote the estimates by $(\Theta_1', \ldots, \Theta_I')$. By the conservative nature of the initial estimates, $\Theta_1' + \cdots + \Theta_I' \leq \Theta$. Suppose the next call request is of class i. Overlooking the increase in multiplexing gain, the capacity requirement of this new call can be conservatively estimated as Θ_i'/M_i. Thus the call can be accepted if $\Theta_1' + \cdots + \Theta_I' + \Theta_i'/M_i \leq \Theta$. The bandwidth estimation process is then repeated.

We should emphasize the *virtual* nature of our proposed estimation algorithms. By using counters to serve as virtual buffers, the performance of different bandwidth values can be measured without impacting the service quality perceived by the users. The actual bandwidth allocation is not modified over the course of the estimation process, and only changed at call set-up (and similarly disconnection).

1.3 Traffic Model

In the remainder of this paper we present the dynamic algorithms for determining the minimum capacity for a given QoS class i with M_i traffic streams, and will supress the subscript i in subsequent notations. The operation of the proposed algorithms does not require knowledge of the traffic characteristics. For the analysis we model the superposed incoming traffic to the FIFO as a Markov modulated fluid. In a Markov modulated fluid model, the source is controlled by a finite-state, irreducible, continuous-time Markov chain with state space $\Lambda = \{1, 2, \ldots, N\}$. The source generates cells at rate λ_n when the Markov chain is at state $n, n \in \Lambda$. Without loss of generality, assume $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N$. Denote B as the size of the FIFO. When a FIFO is full, all incoming cells are dropped. Denote α as the target cell loss probability.

2 Adaptive Algorithm

In our first approach, let θ_0 be the initial bandwidth assignment of the given traffic class. Fix T as the length of the sampling period. As mentioned above, a counter is used to represent the buffer occupancy. Over the sampling period $[(n-1)T, nT)$, $n \in \mathbf{N}$, the counter is decremented at a fixed rate θ_{n-1}, and incremented by each arrival. The number of arrivals C_n and the number of lost cells L_n over this period are measured using another set of counters. At time nT, a new bandwidth is determined by the following algorithm:

$$\theta_n = \theta_{n-1} + \frac{c}{n}(L_n - \alpha C_n), \tag{1}$$

where c is a positive constant.

Theorem 1 *If there exists a minimum bandwidth $\theta_* \in [0, \Theta]$ that achieves the cell loss probability α, then*

$$\theta_n \to \theta_* \text{ a.s. as } n \to \infty.$$

Theorem 1 is based on a general result on adaptive algorithm by Benveniste et al. [1, 7] A simplified version of their theorem is given in Appendix A. We outline the proof of Theorem 1 below.

Proof

Let $R(t) \in \Lambda$ and $B(t) \in [0, B]$ represent the state of the source and the buffer occupancy at time $t \geq 0$, respectively. Define the state vector of the n^{th} sampling period $[(n-1)T, nT)$ as $X_n = (R_n, B_n, C_n, L_n)$, where $R_n = R(nT)$, $B_n = B(nT)$, and C_n and L_n are defined as before. From Eq. 1, the updating function in Eq. 24 is $H(\theta, x) = l_x - \alpha c_x$, where $x = (r_x, b_x, c_x, l_x)$. Note that it only depends on the measurable quantities.

We first verify assumptions (A.1) to (A.5) in Appendix A. For any function $f : \mathbf{R}^d \times \mathbf{R}^k \to \mathbf{R}$, we will use the following shorthand notations:

$$\pi_\theta f_\theta(x) = \int f(\theta, y) \pi_\theta(x, dy), \tag{2}$$

$$\Gamma_\theta f_\theta = \int f(\theta, y) \Gamma_\theta(dy). \tag{3}$$

Assumptions **(A.1)** to **(A.3)** are straightforward by the choices of $\gamma_n = \frac{c}{n}$ and X_n. For a fixed θ, $\{X_n\}$ is a Markov chain with a transition probability matrix Π_θ. Define Γ_θ as its invariant probability.

For **(A.4)(i)**, define $h(\theta) \equiv \Gamma_\theta H_\theta$. Without loss of generality, assume $\theta < \theta'$. Consider two systems with arrivals driven by the same Markov chain, and with the same initial conditions. For the system with service rate θ define the following:

$C_\theta(t)$ = number of arrivals in $[0, t]$,
$L_\theta(t)$ = number of lost cells in $[0, t]$,
$Y_\theta(t)$ = amount of time the buffer is nonempty in $[0, t]$,
$B_\theta(t)$ = buffer occupancy at time t,

and similarly for the system with rate θ'. Note that

$$C_\theta(t) = L_\theta(t) + \theta \cdot Y_\theta(t) + B_\theta(t), \quad \text{for all } t. \tag{4}$$

Furthermore,

$$\lim_{t \to \infty} \frac{C_\theta(t)}{t} \quad = \quad \frac{E[C_\theta]}{T}, \tag{5}$$

$$\lim_{t\to\infty} \frac{L_\theta(t)}{t} = \frac{E[L_\theta]}{T}, \tag{6}$$
$$\lim_{t\to\infty} \frac{Y_\theta(t)}{t} = \frac{E[Y_\theta]}{T}, \tag{7}$$

where $E[C_\theta]$, $E[L_\theta]$, and $E[Y_\theta]$ are the expected values of C, L, and Y in one sampling period when the system is in steady state. Since $0 \le B_\theta(t) \le B$, $\lim_{t\to\infty} B_\theta(t)/t = 0$. Therefore

$$E[C_\theta] = E[L_\theta] + \theta \cdot E[Y_\theta]. \tag{8}$$

The same also holds for the system with rate θ'. Thus

$$E[L_\theta] - E[L_{\theta'}] = \theta' \cdot E[Y_{\theta'}] - \theta \cdot E[Y_\theta]. \tag{9}$$

Consider the three possible cases: (1) $0 \le \theta < \theta'$. Since the two systems have coupled arrivals, $B_\theta(t) \ge B_{\theta'}(t)$, for all t. Therefore $L_\theta(t) \ge L_{\theta'}(t)$ and $Y_\theta(t) \ge Y_{\theta'}(t)$, for all t.

$$\begin{aligned}
\theta' \cdot E[Y_{\theta'}] - \theta \cdot E[Y_\theta] &= (\theta' - \theta)E[Y_{\theta'}] + \theta(E[Y_{\theta'}] - E[Y_\theta]) \\
&\le (\theta' - \theta)E[Y_{\theta'}] \\
&\le (\theta' - \theta)T.
\end{aligned} \tag{10}$$

(2) $\theta < \theta' \le 0$. Since $E[Y_\theta] = E[Y_{\theta'}] = T$, $\theta' \cdot E[Y_{\theta'}] - \theta \cdot E[Y_\theta] = (\theta' - \theta)T$. (3) $\theta < 0 < \theta'$. $\theta' \cdot E[Y_{\theta'}] - \theta \cdot E[Y_\theta] = \theta' \cdot E[Y_{\theta'}] - \theta T \le (\theta' - \theta)T$. Hence $|h(\theta) - h(\theta')| = |E[L_\theta] - \alpha \cdot E[C_\theta] - E[L_{\theta'}] + \alpha \cdot E[C_{\theta'}]| \le |\theta' - \theta| \cdot T$.

(A.4)(ii) First note that $\nu_\theta(x) = \sum_{n\ge 0} \Pi_\theta^n (H_\theta - h(\theta))(x)$ satisfies (A.4)(ii) if we can show that this series converges. We use a coupling argument. Consider two systems δ and β, both with the same rate θ. System δ starts with initial state $X_0^\delta = x$, and system β with initial state $X_0^\beta \sim \Gamma_\theta$.

Let $N_r = \min\{n : R_n^\delta = R_n^\beta\}$. Define a third system γ, also with service rate θ as follows: for $n \le N_r$, $X_n^\gamma = X_n^\delta$. For time $t > N_r T$, system γ is then driven by the same Markov source as system β. Since systems γ and β have the same input and output rates at $t \ge N_r T$, the remaining of the state vector will couple if $B^\gamma(t) = B^\beta(t)$ for some $t \ge N_r T$. This happens when the buffers of the two systems are either both full or both empty. Define $N_b \equiv \min\{n : B_{n+N_r}^\beta = B_{n+N_r}^\gamma\}$. Since

$$\begin{aligned}
|\sum_{n\ge 0} \Pi_\theta^n (H_\theta - h(\theta))(x)| &\le \sum_{n\ge 0} |E[L_n^\gamma] - \alpha E[C_n^\gamma] - E[L_n^\beta] + \alpha E[C_n^\beta]| \\
&\le K \cdot (E[N_r + N_b \mid X_0^\delta = x] + 1),
\end{aligned} \tag{11}$$

for some constant K, it is enough to prove that $E[N_r + N_b \mid X_0^\delta = x] < \infty$.

Since the source Markov chain is irreducible and finite-state, there exists a positive constant $M < \infty$ such that $P(R_M^\delta = r \mid R_0^\delta = r) > 0$, for all $r \in \Lambda$.

This implies that $P(N_r < M, R_M^\delta = r \mid R_0^\delta = r) > 0$, for all $r \in \Lambda$. Define $\varepsilon \equiv \min_{r\in\Lambda} P(N_r < M, R_M^\delta = r \mid R_0^\delta = r)$. Then

$$\begin{aligned}
& P(N_r \geq n \mid R_0^\delta = r_x) \\
\leq\ & P(N_r \geq M\lfloor \frac{n}{M} \rfloor \mid R_0^\delta = r_x) \\
=\ & \sum_{r\in\Lambda} P(N_r \geq M\lfloor \frac{n}{M} \rfloor, N_r \geq M(\lfloor \frac{n}{M} \rfloor - 1), R^\delta_{M(\lfloor \frac{n}{M} \rfloor -1)} = r \mid R_0^\delta = r_x) \\
=\ & \sum_{r\in\Lambda} P(N_r \geq M \mid N_r \geq 0, R_0^\delta = r) P(N_r \geq M(\lfloor \frac{n}{M} \rfloor - 1), R^\delta_{M(\lfloor \frac{n}{M} \rfloor -1)} = r \mid R_0^\delta = r_x) \\
\leq\ & (1-\varepsilon) P(N_r \geq M(\lfloor \frac{n}{M} \rfloor - 1) \mid R_0^\delta = r_x) \\
\leq\ & (1-\varepsilon)^{\lfloor \frac{n}{M} \rfloor}. \qquad (12)
\end{aligned}$$

Therefore

$$\begin{aligned}
E[N_r \mid X_0^\delta = x] &= \sum_{n=0}^{\infty} P(N_r \geq n \mid R_0^\delta = r_x) \\
&\leq \sum_{k=0}^{\infty} \sum_{j=0}^{M-1} (1-\varepsilon)^{\lfloor \frac{kM+j}{M} \rfloor} \\
&= M\frac{1}{\varepsilon} < \infty, \text{ for all } x. \qquad (13)
\end{aligned}$$

$E[N_b \mid X_0^\delta = x] < \infty$ for all x is straightforward.

(A.4)(iii) By showing the convergence of the series in (A.4)(ii), we prove that there exist constants $C_1 < \infty, q_1 = 0$, such that $|\nu_\theta(x)| \leq C_1(1+|x|^{q_1})$.

For the second part of (A.4)(iii), recall that $\Pi_\theta \nu_\theta(x) = \sum_{n\geq 1} \Pi_\theta^n (H_\theta - h(\theta))(x)$. Again using a similar coupling argument, we can show that the inequality is satisfied by $\lambda = 1, q_2 = 0$ and $C_2 = 2\cdot T \cdot E[\text{coupling time}] < \infty$.

(A.5) For any compact subset Q of $\mathbf{R}$, given that $\theta_n \in Q$, $|L_{n+1}| \leq \max(0, (\lambda_N - \inf Q)T)$. Therefore

$$\begin{aligned}
E_{x,a}\{I(\theta_k \in Q, k \leq n)(1+|X_{n+1}|^q)\} &\leq E_{x,a}\{I(\theta_n \in Q)(1+|X_{n+1}|^q)\} \\
&\leq 1 + (N + B + \lambda_N T + \max(0, (\lambda_N - \inf Q)T))^q \\
&\equiv \mu_q(Q) < \infty, \qquad (14)
\end{aligned}$$

for Q is bounded if it is compact in $\mathbf{R}$. Since $|x|^q \geq 0$ for any $q > 0$, the condition is true for all n, x, a.

For the remaining conditions in Theorem 3, we first show that, with probability 1, $\theta_n \in S \equiv [-2\sqrt{Bc\alpha\lambda_N}, \lambda_N(1+cT)]$, for all n. Suppose $\theta_{n-1} \geq \lambda_N$. Since the service rate is higher than the maximum arrival rate, no cell will be lost in the next sampling period. Thus $L_n = 0$ and $\theta_n = \theta_{n-1} + \gamma_n(L_n - \alpha C_n) \leq \theta_{n-1}$. Therefore, $\theta_n > \theta_{n-1}$ only if $\theta_{n-1} < \lambda_N$. For all n

$$\begin{aligned}
\theta_n &= \theta_{n-1} + \gamma_n(L_n - \alpha C_n) \\
&\leq \theta_{n-1} + \gamma_n L_n \\
&\leq \lambda_N + c\lambda_N T = \lambda_N(1 + cT).
\end{aligned} \tag{15}$$

For the lower bound, fix $\epsilon > 0$. Suppose $\theta_{n-1} > -\epsilon$ and $\theta_n \leq -\epsilon$. Then if the service rate continues to decrease, it will take at most $B/\epsilon T$ sampling periods to fill up the buffer, after which the service rate can no longer decrease. Using $\theta_m \geq \theta_{m-1} - \gamma_m \alpha C_m \geq \theta_{m-1} - c\alpha\lambda_N T$ and the assumption that $\theta_{n-1} > -\epsilon$, we get

$$\theta_{n+B/\epsilon T} \geq -(\epsilon + \frac{Bc\alpha\lambda_N}{\epsilon}). \tag{16}$$

Choosing ϵ to maximize the right hand side, $\theta_n \geq -2\sqrt{Bc\alpha\lambda_N}$ for all n.

Let θ_* be such that $h(\theta_*) = E[L_{\theta_*}] - \alpha E[C_{\theta_*}] = 0$. Thus θ_* is the optimal service rate at which the fraction of cell loss equals the target. On the other hand, θ_* is an equilibrium point of the differential equation 30. Using Lyapunov function $V(\theta) = (\theta - \theta_*)^2$, one can show that it is a point of asymptotic stability with domain of attraction $= \mathbf{R} \supset S$.

3 Parallel Monitoring Algorithm

The main drawback of the adaptive algorithm is that it is slow. Since the cell loss probability is a very small value, the sampling period must be long. However, it is crucial to quickly find a good estimate of the optimal bandwidth, since the duration of the learning process must be within the timescale of the interarrival times of call requests. Our second algorithm, called *parallel monitoring*, addresses this issue.

We first observe that in the implementation of the framing scheme, the bandwidth allocated to a traffic class can only be changed at a fixed increment. Thus the objective is to determine the smallest bandwidth that satisfies the loss requirement within a finite number of choices. Let $\theta_1 < \theta_2 < \cdots < \theta_K$ be the choices of bandwidths, which should lie between the mean arrival rate of the aggregate source and the bandwidth determined by the worst-case assumption. Now that the number of candidates is finite, the fastest learning process would result if the performance of each possible bandwidth can be monitored simultaneously. The solution is to use K counters serving as virtual buffers to emulate the buffer occupancy process (Figure 3). Each counter is incremented at cell arrivals and decremented at one of the possible rates. Additional counters are used to store the numbers of arrivals and lost cells.

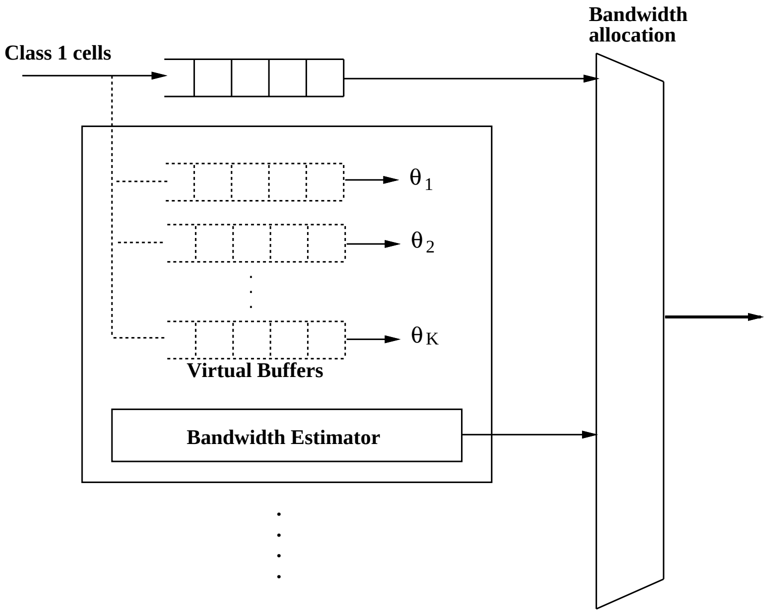

Figure 3: Implementation of the parallel monitoring algorithm.

Define $\bar{\lambda}$ as the expected cell arrival rate. Then the target cell loss rate is $\alpha \cdot \bar{\lambda}$. Define $E[Z^i]$ = the expected loss rate of virtual buffer i at steady state $-\alpha \cdot \bar{\lambda}$. Then $E[Z^1] \geq E[Z^2] \geq \cdots E[Z^K]$. Thus the optimal bandwidth $\theta_* \equiv \min\{\theta_i : E[Z^i] < 0\}$.

Define $C(t)$ = number of arrivals in $[0, t)$, and $L^i(t)$ = number of lost cells in virtual buffer i in $[0, t)$. Let $Z^i(t) = \frac{1}{t}[L^i(t) - \alpha C(t)]$. Let $\theta(t)$ be the optimal bandwidth at time t. That is,

$$\theta(t) = \min\{\theta_i : Z^i(t) < 0\}. \tag{17}$$

We have a stronger result than almost sure convergence of $\theta(t)$ to the optimal value θ_*. Define the *learning time* τ_c as the last time some value other than θ_* is identified, $\tau_c \equiv \sup\{t : \theta(t) \neq \theta_*\}$. τ_c is thus the maximum time it takes to obtain the correct estimate.

Theorem 2 $E[\tau_c] < \infty$.

Before proceeding with the proof, we first state a general lemma for functions of Markov chain. Let $\{Y_n\}$ be an irreducible finite-state Markov chain with state space Σ and transition probability $\pi(y, y')$. Let W_n be a bounded random function of (Y_{n-1}, Y_n) with conditional distribution function $F_{y,y'}(w) = P(W_n \leq w \mid Y_{n-1} = y, Y_n = y'), y, y \in \Sigma$. Let $Z_n = \frac{1}{n}\sum_{k=1}^{n} W_k$. Define $E[W]$ as the expected value of W_n at steady state.

Lemma 1 *If* $\tau = \sup\{n : Z_n \le a\}$ *for some* $a < E[W]$, *then* $E[\tau] < \infty$.

The proof of Lemma 1 is presented in Appendix B.

Proof of Theorem 2

Define

$$\tau_i = \begin{cases} \sup\{t : Z^i(t) > 0\}, & \text{if } \theta_i \ge \theta_* \\ \sup\{t : Z^i(t) < 0\}, & \text{if } \theta_i < \theta_* \end{cases} \tag{18}$$

Since $\tau_c \le \max_i \tau_i$ *a.s*, it is enough to prove that $E[\tau_i] < \infty, i \in \{1, 2, \ldots, K\}$. We discuss the proof for $\theta_i < \theta_*$ below. A similar argument can be applied to $\theta_i \ge \theta_*$. For simplicity of notation we drop the subscript i in the rest of the proof.

First fix a constant $T > 0$. Define the state vector $X_n \equiv (R(nT), B(nT))$, the state of the source Markov chain and the buffer occupancy at time nT, respectively. Let L_n = losses in $[(n-1)T, nT)$ and C_n = arrivals in $[(n-1)T, nT)$. $Z_n \equiv \frac{1}{n}\sum_{k=1}^{n}(L_k - \alpha C_k)$. By ergodicity, $\lim_{n\to\infty} Z_n = A$ *a.s.*, for some constant $A > 0$.

Choose a constant $\Delta > 0$ such that $\frac{B}{\Delta} \in \mathbf{N}$. Define $B^\Delta = \{0, \Delta, 2\Delta, \ldots, \frac{B}{\Delta}\Delta\}$. For any $0 \le b \le B$, call $\lfloor b \rfloor_\Delta$ the quantized value of b to the set B^Δ: $\lfloor b \rfloor_\Delta \equiv \max\{x \in B^\Delta : x \le b\}$. Consider a modified Markov chain with the same arrival process as the original, and state vector $\tilde{X}_n = (R_n, \tilde{B}_n)$, where $\tilde{B}_n = \lfloor \tilde{B}_{n-1} + \int_{(n-1)T}^{nT}(R(s) - \theta)1(0 < \tilde{B}(s) < \mathbf{B})ds \rfloor_\Delta$. Note that the difference in losses between the two systems over one period T is bounded by the difference in their buffer occupancies at the beginning of the period: $L_n - \tilde{L}_n \le B_{n-1} - \tilde{B}_{n-1}$.

Furthermore, note that $\tilde{B}(t) \le B(t), \forall t$. Thus $\tilde{B}(t) = 0$ whenever $B(t) = 0$ and $B(t) = B$ whenever $\tilde{B}(t) = B$. Define regenerative points $D_i \equiv \inf\{t > D_{i-1} : R(t) = N, \tilde{B}(t) = B$, and there exists $D_{i-1} < s < t$ such that $R(s) \ne 1\}$. Define U_n as the length of the regenerative cycle that includes the point nT. Thus

$$L_n - \tilde{L}_n \le B_{n-1} - \tilde{B}_{n-1} \le \frac{1}{T}(U_{n-1} + T)\Delta. \tag{19}$$

Therefore,

$$\begin{aligned} \tilde{Z}_n &\equiv \frac{1}{n}\sum_{k=1}^{n}(\tilde{L}_k - \alpha C_k) \\ &\ge \frac{1}{n}\left\{ \sum_{k=1}^{\lceil D_1/T \rceil}(\tilde{L}_k - \alpha C_k) + \sum_{k=\lceil D_1/T\rceil+1}^{n}(L_k - \alpha C_k) \right. \qquad (20) \\ &- \left. \frac{\Delta}{T}\sum_{k=\lceil D_1/T\rceil+1}^{n}(U_{k-1} + T) \right\}, \qquad (21) \end{aligned}$$

and

$$\lim_{n\to\infty} \tilde{Z}_n \geq 0 + A - \frac{\Delta}{T} \cdot \left(\frac{E[U^2]}{E[U]} + 2T \right). \tag{22}$$

One can show that the first two moments of the regenerative cycle are finite. Thus there exists a $\Delta > 0$ small enough for which $\lim_{n\to\infty} \tilde{Z}_n \geq \tilde{A} \equiv A - \Delta[\frac{E[U^2]}{TE[U]} + 2] > 0$ a.s. Choose a constant a such that $0 < a < \tilde{A}$. Define $M = \max\{n : Z_n \leq a\}$ and $\tilde{M} = \max\{n : \tilde{Z}_n \leq a\}$. Since $\tilde{Z}_n \leq Z_n$ a.s., $\tilde{M} \geq M$ a.s. Now $\{\tilde{X}_n\}$ is a finite-state Markov chain, $(\tilde{L}_n - \alpha C_n)$ is a bounded random function of $(\tilde{X}_{n-1}, \tilde{X}_n)$, and $\lim_{n\to\infty} \tilde{Z}_n > a$ a.s. Applying Lemma 1, we obtain $E[\tilde{M}] < \infty$ and therefore $E[M] < \infty$.

For $n > (M + 1 + \frac{\alpha\lambda_N}{\varepsilon})$, consider $Z(t)$ for $nT \leq t < (n+1)T$. Since $L(t) \geq L(nT)$ and $C(t) - C(nT) \leq \lambda_N T$, we have the following relationship:

$$\begin{aligned} Z(t) &= \frac{nT}{t} Z(nT) + \frac{1}{t}[L(t) - L(nT) - \alpha(C(t) - C(nT))] \\ &\geq \frac{nT}{t} Z(nT) - \frac{1}{t}\alpha(C(t) - C(nT)) \\ &\geq \frac{nT}{t} Z(nT) - \frac{1}{t}\alpha\lambda_N T \\ &> \frac{nT}{(n+1)T}\varepsilon - \frac{1}{nT}\alpha\lambda_N T \\ &= \frac{n}{n+1}\varepsilon - \frac{\alpha\lambda_N}{n} > 0, \end{aligned} \tag{23}$$

where the third inequality is due to the fact that $nT \leq t < (n+1)T$ and $Z_n > \varepsilon$ for $n > M$, and the fourth inequality results from $n > 1 + \frac{\alpha\lambda_N}{\varepsilon}$. Thus $Z(t) > 0$ for all $t > T \cdot (M \vee (1 + \frac{\alpha\lambda_N}{\varepsilon}))$. Therefore $\tau \leq T \cdot (M + 1 + \frac{\alpha\lambda_N}{\varepsilon})$ and $E[\tau] < \infty$.

4 Conclusions

For broadband networks that are intended to support multiple traffic classes, dynamic resource allocation provides an attractive solution to the problems suffered by the conventional static approach. It eliminates the reliance on a detailed traffic model to guarantee the quality of service provided to the user, and thus simplifies the tasks on the user's end. This benefit can become crucial as new varieties of multimedia applications pose increasing challenges to the existing modeling techniques. A dynamic algorithm can also complement a conservative call acceptance policy in achieving a high network resourse utilization.

In this paper we have presented two dynamic bandwidth allocation schemes and shown their convergence property. Both require only simple operations that can be implemented in hardware. For the parallel monitoring algorithm, we showed that its learning time has a finite expected value. Furthermore, since the algorithm simultaneously tracks the performance of multiple allocations, one would expect it to be one of the fastest dynamic schemes. One common concern about dynamic schemes is their speed of convergence. In our application, it is crucial for the algorithms to arrive at a good estimate of the actual bandwidth within the timescale of the call interarrival and set-up times. Although it can be seen from the proof of Theorem 2 that the probability of not choosing the optimal bandwidth decreases inversely exponentially with time, the rate of decrease depends on the source statistics. Thus we turn to simulations to study the applicability of the dynamic algorithms. Preliminary simulation results have shown that the parallel monitoring algorithm is effective in identifying the optimal choice within very short time. The simulation results and heuristic techniques for speeding up the estimation process will be reported in a subsequent paper.

A Almost Sure Convergence of Adaptive Algorithm

Consider the general form of adaptive algorithm,

$$\theta_n = \theta_{n-1} + \gamma_n H(\theta_{n-1}, X_n), \tag{24}$$

with state vector $X_n \in \mathbf{R}^k$, control vector $\theta_n \in \mathbf{R}^d$, and updating function $H : \mathbf{R}^d \times \mathbf{R}^k \to \mathbf{R}^d$. $(\gamma_n)_{n \in \mathbf{N}}$ is a sequence of gains. We first state the assumptions:

(A.1) $(\gamma_n)_{n \in \mathbf{N}}$ is a decreasing sequence of positive real numbers such that $\sum_n \gamma_n = +\infty$.

(A.2) There exists a family $\Pi_\theta : \theta \in \mathbf{R}^d$ of transition probabilities $\Pi_\theta(x, \mathcal{A})$ on $\mathbf{R}^k$ such that for any Borel subset $\mathcal{A}$ of $\mathbf{R}^k$,

$$P(X_{n+1} \in \mathcal{A} \mid \mathcal{F}_n) = \Pi_{\theta_n}(X_n, \mathcal{A}). \tag{25}$$

(A.3) For any compact subset Q of $\mathbf{R}^d$, there exist constants C, q (depending on Q) such that for all $\theta \in Q$,

$$|H(\theta, x)| \leq C(1 + |x|^q). \tag{26}$$

(A.4) There exists a function h on $\mathbf{R}^d$, and for each $\theta \in \mathbf{R}^d$ a function $\nu_\theta(.)$ on $\mathbf{R}^k$ such that

(i) h is Lipschitz on $\mathbf{R}^d$;

(ii) $(I - \Pi_\theta)\nu_\theta = H_\theta - h(\theta)$ for all $\theta \in \mathbf{R}^d$;

(iii) for all compact subsets Q of $\mathbf{R}^d$, there exist constants $C_1, C_2, q_1, q_2, \lambda \in [\frac{1}{2}, 1]$, such that for all $\theta, \theta' \in Q$,

$$|\nu_\theta(x)| \leq C_1(1 + |x|^{q_1}), \tag{27}$$

$$|\Pi_\theta \nu_\theta(x) - \Pi_{\theta'} \nu_{\theta'}(x)| \leq C_2 |\theta - \theta'|^\lambda (1 + |x|^{q_2}). \tag{28}$$

(A.5) For any compact subset Q of $\mathbf{R}^d$ and any $q > 0$, there exists $\mu_q(Q) < \infty$ such that for all n, $x \in \mathbf{R}^k$, $a \in \mathbf{R}^d$

$$E_{x,a} I(\theta_k \in Q, k \leq n)(1 + |X_{n+1}|^q) \leq \mu_q(Q)(1 + |x|^q). \tag{29}$$

For $t \geq t_0$, let $\bar{\theta}(t; t_0, a_0)$ denote the solution of the differential equation

$$\bar{\theta}'(t) = h(\bar{\theta}(t)), \quad t \geq t_0, \quad \bar{\theta}(t_0) = a_0. \tag{30}$$

Theorem 3 *[1] Suppose that* **(A.1)** *to* **(A.5)** *hold. Let θ_* be a point of asymptotic stability of the differential equation 30. If θ_n is bounded a.s. and θ_n visits a.s. a compact subset of the domain of attraction of θ_* infinitely often, then θ_n converges to θ_* a.s.*

B Proof of Lemma 1

We first state a theorem of large deviation for finite-state Markov chains. The theorem is a simple extension of Theorem 3.1.2 in [4].

Define

$$\Lambda_{y,y'}(\lambda) = \log \int_w e^{\lambda w} dF_{y,y'}(w). \tag{31}$$

Associated with every $\lambda \in \mathbf{R}$ a nonnegative matrix Π_λ, whose elements are

$$\pi_\lambda(y, y') = \pi(y, y') e^{\Lambda_{y,y'}(\lambda)}, \quad y, y' \in \Sigma. \tag{32}$$

Let $\rho(\Pi_\lambda)$ denote the Perron-Frobenius eigenvalue of the matrix Π_λ. For every $z \in R$, define

$$I(z) = \sup_{\lambda \in \mathbf{R}} [\lambda z - \log \rho(\Pi_\lambda)]. \tag{33}$$

Then the empirical mean Z_n satisfies the large deviation principle with rate function $I(\cdot)$. Namely,

Theorem 4 *For any set $\Gamma \subseteq \mathbf{R}$, and any initial state $\sigma \in \Sigma$,*

$$\begin{aligned} -\inf_{z \in \Gamma^\circ} I(z) &\leq \liminf_{n \to \infty} \frac{1}{n} \log P_\sigma^\pi(Z_n \in \Gamma) \\ &\leq \limsup_{n \to \infty} \frac{1}{n} \log P_\sigma^\pi(Z_n \in \Gamma) \leq -\inf_{z \in \bar{\Gamma}} I(z). \end{aligned} \tag{34}$$

Proof of Lemma 1

Let $I \equiv inf_{z \in (-\infty, a]} I(z) > 0$. The theorem implies that, for any initial state σ and any $0 < \varepsilon < I$, there exists a constant n_ε, such that

$$\sup_{m \geq n} \frac{1}{m} \log P_\sigma^\pi(Z_m \leq a) \leq -I + \varepsilon, \quad \text{for all } n \geq n_\varepsilon. \tag{35}$$

Therefore,

$$P_\sigma^\pi(Z_m \leq a) \leq e^{-m(I-\varepsilon)}, \quad \text{for all } m \geq n_\varepsilon. \tag{36}$$

In addition, $\{\tau > n\} = \bigcup_{k=1}^\infty \{Z_{n+k} \leq a\}$. Thus we have

$$\begin{aligned} E[\tau] &= \sum_{n=0}^\infty P_\sigma^\pi(\tau > n) \\ &\leq \sum_{n=0}^\infty \sum_{k=1}^\infty P_\sigma^\pi(Z_{n+k} \leq a) \\ &= \sum_{k=1}^\infty k \cdot P_\sigma^\pi(Z_k \leq a) \\ &\leq \sum_{k=1}^{n_\varepsilon - 1} k \cdot P_\sigma^\pi(Z_k \leq a) + \sum_{k=n_\varepsilon}^\infty k \cdot e^{-k(I-\varepsilon)} < \infty. \end{aligned} \tag{37}$$

References

[1] A. Benveniste, M. Metivier, and P. Priouret (1987)
Adaptive Algorithms and Stochastic Approximations, Springer-Verlag.

[2] M. Bonatti and A. A. Gaivoronski (1994?)
"Worst Case Analysis of ATM Sources with Application to Access Engineering of Broadband Multiservice Networks," *preprint.*

[3] J. Bucklew (1990)
Large Deviation Techniques in Decision, Simulation, and Estimation, Wiley.

[4] A. Dembo and O. Zeitouni (1993)
Large Deviations Techniques and Applications, Jones and Bartlett.

[5] B. T. Doshi (1993)
"Deterministic Rule Based Traffic Descriptors for Broadband ISDN: Worst Case Behavior and Connection Acceptance Control," *IEEE Globecom'93*, pp. 1759-1764.

[6] I. Hsu and J. Walrand (1994)
"Admission Control for Multi-Class ATM Traffic," submitted to *Computer Networks and ISDN Systems Journal.*

[7] M. Metivier and P. Priouret (1984)
"Applications of a Kushner and Clark Lemma to General Classes of Stochastic Algorithms," *IEEE Transactions on Information Theory*, March 1984, vol. IT-30, no. 2, pp. 140-151.

[8] S. P. Meyn and R. L. Tweedie (1993)
Markov Chains and Stochastic Stability, Springer-Verlag.

[9] Y. Takagi and T. Takahashi (1991)
"Priority Assignment Control of ATM Line Buffers with Multiple QOS Classes," *IEEE Journal on Selected Areas in Communications*, Sept. 1991, vol. 9, no. 7, pp. 1078-1092.

[10] E. A. Varvarigos (1994)
"Universal Upper Bounds on the Packet Loss Probability of Statistical Multiplexing in High-Speed Networks," *32nd Allerton Conference on Communication, Control, and Computing.*

[11] M. Vidyasagar (1978)
Nonlinear Systems Analysis, Prentice Hall.

[12] S. Yazid and H. T. Mouftah (1992)
"Congestion Control Methods for BISDN," *IEEE Communications Magazine*, July 1992, vol. 30, no. 7, pp. 42-47.