Multiple Time Scale Regulation and Worst Case Processes for ATM Network Control

Debasis Mitra and *John A. Morrison*
AT&T Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Recent research on congestion control in high speed ATM networks has sought to extract statistical multiplexing efficiencies while guaranteeing quality of service. To facilitate network control, traffic regulators are positioned at network edges, with Leaky Bucket regulation ubiquitous. We consider the composite of conventional Leaky Bucket regulators, one for each time scale. Over longer time scales the enforcement of traffic to the mean in the regulation is less stringent, in so far as the bucket is larger, while the mean rate is smaller. The regulated process is required to be compliant with all the Leaky Bucket regulators. In the approach taken here, traffic sources are assumed to be mutually independent but otherwise adversarial to the utmost extent permitted by regulators. To mitigate against an overly conservative design, small losses are allowed in accessing the shared unbuffered network resource. We rely on the Chernoff bound, which also gives an asymptotic, large deviations approximation, to estimate the loss probability. This paper resolves a question regarding the constitution of independent, regulated traffic processes which maximize the estimate of the loss probability. We show that the estimated loss probability is maximized by certain periodic on-off processes with random phases. Importantly, this loss probability depends on the composite regulators' parameters only through the ratio of the mean rate in the longest time scale to the peak rate. This is true even in a heterogeneous environment.

Multiple Time Scale Regulation and Worst Case Processes for ATM Network Control

Debasis Mitra and *John A. Morrison*
AT&T Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

Research in high speed communications has been dominated in recent years by issues concerning services based on the Asynchronous Transfer Mode (ATM). An important issue is congestion control. The statistical nature of a significant part of the traffic, its burstiness and variability, has posed challenges in designs with guaranteed quality of service. To provide any guarantee it is essential to regulate traffic at network edges, to check compliance with contracts between users and network service providers. In particular, the Leaky Bucket regulator [17, 15, 3] has become ubiquitous in current service offerings, and it appears very likely that it will be an important regulator in higher speed ATM-based services in the near future. One approach to network dimensioning is based on statistical source models, their regulation and the congestion effects at the internal switching nodes of the network. See for instance [15, 5] for various detailed studies. The reliance on statistical source models in this approach is a matter of concern, however. A contrasting approach is to base network and switch dimensioning on the worst case congestion effects at the network nodes, given that regulators exist at the entry points. That is, traffic sources are considered to be adversarial to the utmost extent permitted by regulators. Even with this worst case approach there are interesting choices to be made. One option is to proceed with deterministic bounds to ensure that there are no losses in the network nodes. See [14] for an example of work along these lines. Another option is to allow very small losses, say 10^{-9}, and to extract multiplexing gains from the statistical independence of traffic processes. We believe that the latter option leads to better use of network resources. See [6, 9, 16] for examples of work along these lines. This paper resolves a question regarding the constitution of independent, regulated traffic processes which maximize the estimate of the probability of loss in an unbuffered network resource. Prior work on this question has been reported in [4, 18, 19]. We show that the estimated loss probability is maximized by certain periodic on-off processes with random phases. Bounds and asymptotic estimates of the worst case probability of loss are also obtained.

A second separate theme of this paper is the behavior of the sources over multiple time scales and their appropriate regulation. The motivation for considering multiple time scales is rooted in the observation that, quite commonly, application processes, especially data, have traffic rates which are relatively small and have small range in long time scales and, in contrast, over shorter time scales the range of rates is much larger. (Arguably, this feature is tautologous to burstiness.) Here the rate over a time scale is obtained from the total volume of source data over time intervals which are representative of the time scale. It is advantageous to take this dependence of rates on time scales into account, rather than work with the worst case values over all time scales, which are dominated by those from the short time scale. By so doing, the network service provider can increase resource utilizations, while the service user is entitled to lower tariffs. The cost of such an arrangement is a more elaborate contract between the user and the network, and the concomitant policing of such a contract. However, we believe that a reasonable solution is a system of, say S, Leaky Bucket regulators, one for each time scale. The regulated process is required to comply with all S regulators and, additionally, a peak rate constraint.

A process is compliant with a Leaky Bucket regulator with parameters (ρ, σ) if the content of the token buffer in the regulator, which is governed by the following algorithm, is always nonnegative.

> The token buffer is incremented at rate ρ up to a maximum of σ and decremented at the rate of the regulated process.

As mentioned above, we have a system of S Leaky Bucket regulators, (ρ_s, σ_s) $(s = 1, 2, \ldots, S)$ and the processes are required to comply with all S regulators. Let P denote the peak rate, i.e., the maximum permissible process rate. Also let

$$\rho_1 < \rho_2 < \cdots < \rho_S < P \; , \tag{1.1}$$

i.e., (ρ_1, σ_1) and (ρ_S, σ_S) are respectively associated with the longest and the shortest time scales. To make this correspondence between time scales and regulators meaningful, we have to allow greater latitude in the enforcement to the mean over longer time scales (see Figure 1), i.e.,

$$\sigma_1 > \sigma_2 > \cdots \sigma_S \; . \tag{1.2}$$

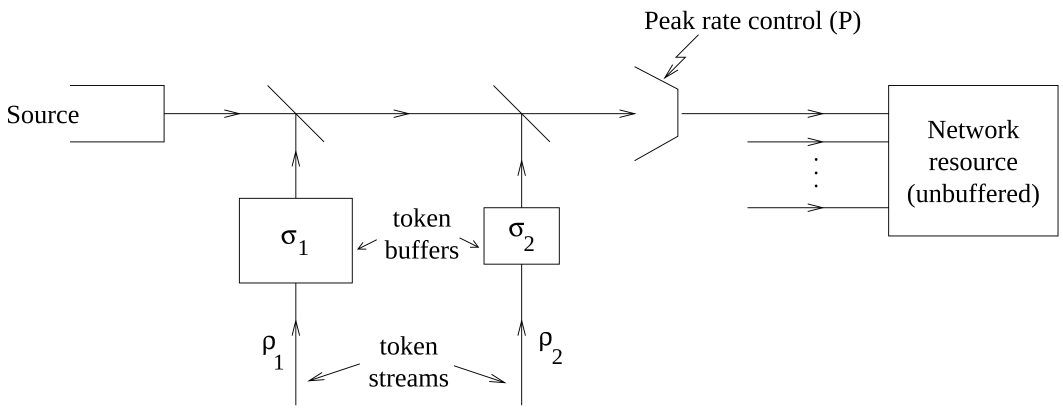

Figure 1: Schematic of system of regulators ($S = 2$).

Coupling between regulators arises when the token buffer of a regulator, say s, is empty. In this situation, the behavior of a nonempty token buffer of another regulator is determined in part by ρ_s. The analysis in this paper is based on fluid flow, which is an approximation to discrete packet/cell transmission. We should also point out that in contrast to some prior work, where the peak rate P plays only a cursory role, here it has a dominant role.

To keep the exposition simple we restrict our attention to the case of only two time scales, i.e., $S = 2$. We allow traffic heterogeneity by means of classes. Thus, regulators for class j are (ρ_{j1}, σ_{j1}), (ρ_{j2}, σ_{j2}) and the peak rate is P_j ($j = 1, 2, \ldots, J$). There are K_j processes of class j.

Our main result is on the maximum estimated loss probability at the network resource, where the maximum is taken over all regulated processes. We show that this quantity depends on the regulators' parameters only through $\{\rho_{j1}/P_j\}$ ($j = 1, 2, \ldots, J$). That is, the relevant regulator parameter for each class is the ratio of the mean rate in the longest time scale to the peak rate. The templates of the worst case regulated processes depend on all the parameters of the regulators.

Systems of dual Leaky Bucket regulators are commonly used in the ATM industry to simultaneously enforce a mean rate over longer time intervals and a peak rate over shorter time intervals [13]. An identical application is considered by Lakshman and Ott [10] in a multiplexing system.

Hui [7] proposed the use of the Chernoff bound and the corresponding large deviations approximation to the loss probability at an unbuffered network resource. See also [8, 11]. In this paper we are considering only the analytically simpler case of unbuffered network resources. However, as shown in [6] there are ways to turn a buffered resource problem into an unbuffered resource problem. This paper extends earlier work [12] which dealt only with a single time scale.

2. REGULATED PROCESSES AND UNBUFFERED RESOURCES

$R_{ji}(t)$ are stationary random processes. At each t, $R_{ji}(t)$ are independent, nonnegative random variables with moment generating functions

$$M_j(s) = E[e^{sR_{ji}}] = \int_0^\infty e^{sx} dV_j(x) \tag{2.1}$$

where

$$V_j(x) = Pr(R_{ji} \le x) \ . \tag{2.2}$$

Interpret $R_{ji}(t)$ as the instantaneous load placed on an unbuffered resource at time t by a regulated process of class j ($j = 1, 2, \ldots, J$). In the framework of fluid models $R_{ji}(t)$ is a rate. In the first part of the analysis we consider only two constraints implied by regulation:

$$R_{ji}(t) \le P_j : \quad \text{peak rate constraint} \tag{2.3}$$

$$\int_0^{P_j} x dV_j(x) \le \rho_{j1} : \quad \text{mean rate constraint} \tag{2.4}$$

Each of the regulators imposes a mean rate constraint. However, only the one shown in (2.4) need be considered in this part of the analysis, since $\rho_{j1} = \min_s \rho_{js}$.

A simple example is

$$R_{ji}(t) = W_j(t + \theta_{ji}) \ , \tag{2.5}$$

where $W_j(t)$ are periodic "template" functions of period T_j and θ_{ji} are independent random variables uniformly distributed in the interval $[0, T_j)$. Then $V_j(x)$ is the fraction of the period T_j in which $W_j \le x$. The example may be extended by scaling time, with each process having its own scale parameter. Of course such scalings keep the distribution $V_j(x)$ invariant.

We are interested in the loss probability L,

$$L = Pr(R > C) \ , \tag{2.6}$$

where the total load $R = \sum_{j=1}^{J} \sum_{i=1}^{K_j} R_{ji}$ and C is the capacity of the resource. Chernoff's bound gives [1],

$$\log L \le -F(s^*) \ , \tag{2.7}$$

where

$$F(s) = sC - \sum_{j=1}^{J} K_j \log M_j(s) \tag{2.8}$$

and $F(s^*) = \sup_{s \geq 0} F(s)$.

If $K_j = \gamma_j C$, with $\gamma_j = O(1)$ and $C \to \infty$, then [2]

$$\log L = -F(s^*) \left[1 + O\left(\frac{\log C}{C} \right) \right] \ . \tag{2.9}$$

Hence the asymptotic, large deviations approximation is $L \sim \exp(-F(s^*))$.

Our goal in this paper is to characterize regulated processes which minimize $F(s^*)$, i.e., which maximize the asymptotic loss, or the bound on the loss. First, in Section 3 we will prove that, subject to the constraints (2.3) and (2.4), $F(s^*)$ is minimized when

$$V_j(0) = 1 - \frac{\rho_{j1}}{P_j}, \quad V_j(P_j) - V_j(P_j - 0) = \frac{\rho_{j1}}{P_j} \quad (j = 1, 2, \ldots, J) \ . \tag{2.10}$$

That is, the random variable takes values 0 and P_j with probabilities $(1 - \rho_{j1}/P_j)$ and (ρ_{j1}/P_j), respectively. The proof will also exhibit monotonic behavior of the asymptotic loss.

Secondly, in Section 4 we will demonstrate that there are corresponding on-off processes with periodic templates $\{W_j(t)\}$ and uniformly distributed random phases which satisfy (2.10) and are also compliant with the system of Leaky Bucket regulation. The token buffer sizes as well as the periods of the templates may be arbitrarily specified.

3. EXTREMAL DISTRIBUTIONS

To avoid trivialities assume that $\sum_{j=1}^{J} K_j \rho_{j1} < C$ and that

$$\lim_{s \to \infty} \sum_{j=1}^{J} \frac{K_j M_j'(s)}{M_j(s)} > C \ ,$$

(the prime denotes derivative) since otherwise there is no loss. Also note from (2.8) that

$$F'(s) = C - \sum_{j=1}^{J} \frac{K_j M_j'(s)}{M_j(s)} \ . \tag{3.1}$$

It is readily seen that $F'(0) > 0$, since $M_j'(0) \leq \rho_{j1}$, and that $F'(\infty) < 0$. Also, from (2.1) and Schwarz's inequality, $M_j''(s) M_j(s) \geq [M_j'(s)]^2$ for $s \geq 0$, and hence $F''(s) < 0$. In summary, $F(s)$ is a strictly concave function with a unique maximum at $s = s^*$, which is the positive root of the equation $F'(s) = 0$.

Given $V_j(x)$, which satisfies (2.4), we introduce $\widehat{V}_j(x)$, a matched distribution of independent random variables which take values 0 and P_j with probabilities $(1 - w_j)$ and w_j, respectively,

$$\widehat{V}_j(0) = 1 - w_j, \quad \widehat{V}_j(P_j) - \widehat{V}_j(P_j - 0) = w_j \quad (j = 1, 2, \ldots, J) \tag{3.2}$$

where

$$w_j = \frac{1}{P_j} \int_0^{P_j} x dV_j(x) \; . \tag{3.3}$$

Also, $\widehat{F}(s)$ is obtained from the moment generating function $\widehat{M}_j(s)$, and $\widehat{s}^*$ is where $\widehat{F}(s)$ is maximized.

Since

$$\widehat{M}_j(s) = 1 - w_j + w_j e^{sP_j} \; , \tag{3.4}$$

it follows that

$$\begin{aligned} \widehat{M}_j(s) - M_j(s) &= \frac{1}{P_j} \int_0^{P_j} \left[x(e^{sP_j} - 1) - P_j(e^{sx} - 1) \right] dV_j(x) \\ &\geq 0 \; , \end{aligned} \tag{3.5}$$

from the convexity of $(e^{sx} - 1)$. Equality in (3.5) holds only if $V_j(x) \equiv \widehat{V}_j(x)$. Hence, importantly,

$$F(s) - \widehat{F}(s) = \sum_{j=1}^{J} K_j \log \left[\frac{\widehat{M}_j(s)}{M_j(s)} \right] \geq 0 \tag{3.6}$$

for all $s \geq 0$. In particular, $F(\widehat{s}^*) \geq \widehat{F}(\widehat{s}^*)$. Hence,

$$F(s^*) \geq F(\widehat{s}^*) \geq \widehat{F}(\widehat{s}^*) \; . \tag{3.7}$$

Observe that $F(s^*) = \widehat{F}(\widehat{s}^*)$ if and only if $V_j(x) \equiv \widehat{V}_j(x)$ $(j = 1, 2, \ldots, J)$.

We next show that subject to the constraint $w_j \leq \rho_{j1}/P_j$, $\widehat{F}(\widehat{s}^*)$ is minimized at the extreme point $w_j = \rho_{j1}/P_j$, for $j = 1, 2, \ldots, J$. This follows from

$$\frac{\partial \widehat{F}(\widehat{s}^*)}{\partial w_j} = - \frac{K_j(e^{\widehat{s}^* P_j} - 1)}{(1 - w_j + w_j e^{\widehat{s}^* P_j})} < 0 \; , \tag{3.8}$$

since $\widehat{F}'(\widehat{s}^*) = 0$.

We take stock in the following

Proposition: (i) For any independent random variables R_{ji} with distributions $V_j(x)$, which satisfy the peak and mean rate constraints in (2.3) and (2.4), there exist corresponding distributions $\widehat{V}_j(x)$ of random variables which take values 0 and P_j only and have the same mean, for which the asymptotic approximation to the loss probability is greater, i.e.,

$$F(s^*) \geq \widehat{F}(\widehat{s}^*) ,$$

with equality if and only if $V_j \equiv \widehat{V}_j$ $(j = 1, 2, \ldots, J)$.

(ii) The asymptotic approximation to the loss probability is further maximized by making the mean rates as large as possible, i.e., by making $w_j = \rho_{j1}/P_j$ $(j = 1, 2, \ldots, J)$ in (3.2). ■

4. EXTREMAL PROCESSES

It remains to show that there are independent random processes $R_{ji}(t)$ which are compliant with the system of Leaky Bucket regulation and satisfy (2.10). These processes will then have been shown to maximize the asymptotic estimate of the loss probability. The constructions below are of the periodic, on-off template functions $W_j(t)$ which with the random phases θ_{ji} give $R_{ji}(t)$, as shown in (2.5). In these constructions the periods T_j, and the parameters (ρ_{j1}, σ_{j1}) and (ρ_{j2}, σ_{j2}) of the Leaky Bucket regulators, are prescribed, finite and positive. We assume that $P_j > \rho_{j2} > \rho_{j1}$ and $\sigma_{j1} > \sigma_{j2}$. We note that (2.10) is satisfied if $W_j(t) = P_j$ for a fraction ρ_{j1}/P_j of the period T_j, and $W_j(t) = 0$ for the remaining fraction. We assume that the small and large token buffers are both full at $t = 0$; we will ensure that they will both be full again at $t = T_j$.

To simplify the notation, we now omit the subscripts j. If $W(t) = P$, then the small and large token buffers empty at rates $P - \rho_2$ and $P - \rho_1$, respectively. If $W(t) = 0$, then the small and large token buffers fill at rates ρ_2 and ρ_1, respectively. Since $P > \rho_2 > \rho_1$, the large token buffer empties faster, and fills more slowly, than the small one. We assume that the process turns on at $t = 0$, so that $W(0) = P$. The process turns off when either token buffer empties. We adopt the rule that the process does not turn on again until the token buffer which emptied is full. If both token buffers empty simultaneously, then the process does not turn on again until the large token buffer is full. The small token buffer is full sooner, of course, and remains so until the large token buffer is. The process also turns off when it has been on for a fraction ρ_1/P of the period T.

We first consider the case $(P - \rho_2)\sigma_1 \leq (P - \rho_1)\sigma_2$. (This is possible, even though we have assumed that $\sigma_1 > \sigma_2$, since $P - \rho_2 < P - \rho_1$. If $\sigma_1 \leq \sigma_2$, then the inequality holds). If $\rho_1(P - \rho_1)T < P\sigma_1$, then the process turns off at $t = \rho_1 T/P$, and neither token buffer is empty. The small token buffer is full at $t = \rho_1 T/\rho_2$, and remains full until $t = T$, when

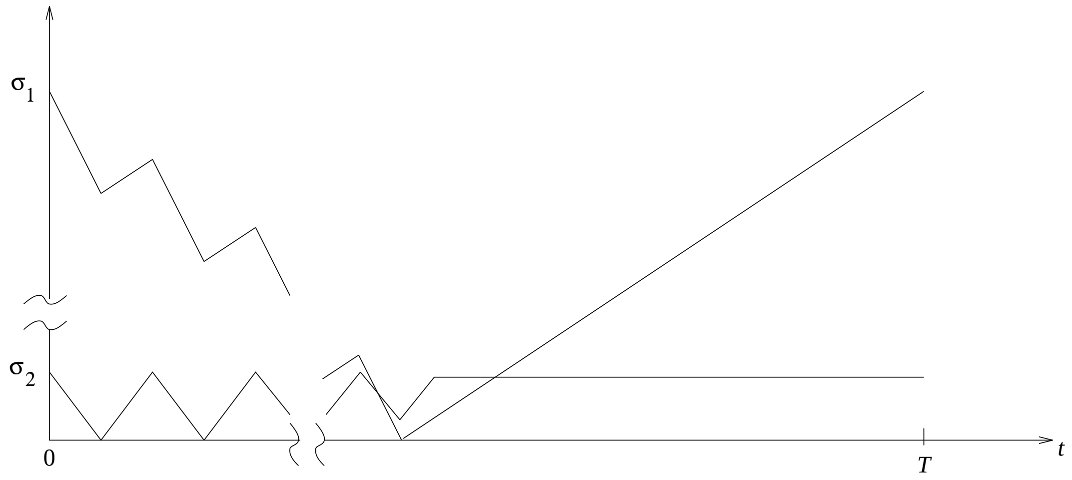

Figure 2: Contents of token buffers for $T = \frac{P}{(P-\rho_1)}\left(\frac{\sigma_1}{\rho_1} + \frac{N\sigma_2}{\rho_2}\right)$.

the large token buffer is full. If $\rho_1(P-\rho_1)T = P\sigma_1$, then the large token buffer is empty at $t = \rho_1 T/P$, and it is full at $t = T$. If $\ell P\sigma_1 < \rho_1(P-\rho_1)T \le (\ell+1)P\sigma_1$, where ℓ is a positive integer, then the pattern for $\rho_1(P-\rho_1)T = P\sigma_1$ is repeated ℓ times, followed by the pattern for $0 < \rho_1(P-\rho_1)T \le P\sigma_1$. We remark that for $(P-\rho_2)\sigma_1 \le (P-\rho_1)\sigma_2$ the large token buffer is dominant, and corresponds to single time scale regulation [12] on the longer time scale.

We now consider the case

$$\frac{(N-1)P(\rho_2-\rho_1)\sigma_2}{\rho_2(P-\rho_2)} < \sigma_1 - \frac{(P-\rho_1)\sigma_2}{(P-\rho_2)} \le \frac{NP(\rho_2-\rho_1)\sigma_2}{\rho_2(P-\rho_2)} , \tag{4.1}$$

where N is a positive integer. If $nP\sigma_2 < \rho_1(P-\rho_2)T \le (n+1)P\sigma_2$, where $0 \le n \le N-1$, then the small token buffer empties and fills n times, and then partially empties and fills. The process turns on at $t = mP\sigma_2/\rho_2(P-\rho_2)$, $m = 0, \ldots, n$, and off at $t = (mP+\rho_2)\sigma_2/\rho_2(P-\rho_2)$, $m = 0, \ldots, n-1$, $(n \ge 1)$, and finally off at $t = \rho_1 T/P + n\sigma_2/\rho_2$, when it has been on for a fraction ρ_1/P of the period. The small token buffer is full at $t = \rho_1 T/\rho_2$, and remains full until $t = T$, when the large token buffer is full. If

$$\frac{NP\sigma_2}{\rho_1(P-\rho_2)} < T \le \frac{P}{(P-\rho_1)}\left(\frac{\sigma_1}{\rho_1} + \frac{N\sigma_2}{\rho_2}\right) , \tag{4.2}$$

then the process turns on and off as above, with $n = N$.

If equality holds in (4.2), then the large token buffer is empty at $t = \rho_1 T/P + N\sigma_2/\rho_2$, and it is full at $t = T$. The contents of the large and small token buffers are depicted as a function of t in Figure 2. If

$$\frac{\ell P}{(P-\rho_1)}\left(\frac{\sigma_1}{\rho_1} + \frac{N\sigma_2}{\rho_2}\right) < T \le \frac{(\ell+1)P}{(P-\rho_1)}\left(\frac{\sigma_1}{\rho_1} + \frac{N\sigma_2}{\rho_2}\right) , \tag{4.3}$$

where ℓ is a positive integer, then the pattern for $T = \frac{P}{(P-\rho_1)}\left(\frac{\sigma_1}{\rho_1} + \frac{N\sigma_2}{\rho_2}\right)$ is repeated ℓ times, followed by the pattern for $0 < T < \frac{P}{(P-\rho_1)}\left(\frac{\sigma_1}{\rho_1} + \frac{N\sigma_2}{\rho_2}\right)$. This completes the construction.

ACKNOWLEDGEMENT: One of the authors (DM) gratefully acknowledges the benefit of many discussions with his colleagues Anwar Elwalid and Robert Wentworth.

References

[1] P. Billingsley, *Probability and Measure*, 2nd ed. (Wiley, New York) 1986.

[2] N. R. Chaganty and J. Sethuraman, "Strong large deviation and local limit theorems," Ann. Prob. **21**, No. 3, 1993, pp. 1671–1690.

[3] R. L. Cruz, "A calculus for network delay, Part I: Network elements in isolation," IEEE Trans. Inform. Theory **37**, 1991, pp. 114–131.

[4] B. T. Doshi, "Deterministic rule based traffic descriptors for broadband ISDN: worst case behavior and connection acceptance control," Proc. ITC-14, Eds. J. Labetoulle and J. W. Roberts, Elsevier, 1994, pp. 591–600.

[5] A. I. Elwalid and D. Mitra, "Analysis and design of rate-based congestion control of high speed networks, I: stochastic fluid models, access regulation," Queueing Systems **9**, 1991, pp. 29–63.

[6] A. I. Elwalid, D. Mitra and R. H. Wentworth, "A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM mode," submitted to IEEE JSAC special issue on "Advances in the Fundamentals of Networking," to appear 1995.

[7] J. Y. Hui, "Resource allocation for broadband networks," IEEE J. Select. Areas Commun. **6**, 1988, pp. 1598–1608.

[8] F. P. Kelly, "Effective bandwidths at multi-class queues," Queueing Systems **9**, 1991, pp. 5–16.

[9] K. Kvols and S. Blaabjerg, "Bounds and approximations for the periodic on/off queue with applications to ATM traffic control," IEEE INFOCOM, 1992, pp. 487–494.

[10] T. V. Lakshman and T. J. Ott, "Multi-rate traffic shapers for ATM," preprint, 1994.

[11] D. Mitra and J. A. Morrison, "Erlang capacity and uniform approximations for shared unbuffered resources," IEEE/ACM Trans. Networking **2**, No. 6, 1994, pp. 558–570.

[12] D. Mitra and J. A. Morrison, "Independent regulated processes to a shared unbuffered resource which maximize the loss probability," preprint, 1994.

[13] G. Niestegge, "The leaky bucket policing method in the ATM network," Int. J. Digital and Analog Commun. Syst. **2**, 1990, pp. 187–197.

[14] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks — the single-node case," IEEE/ACM Trans. Networking **1**, No. 3, June 1993, pp. 344–357.

[15] J. W. Roberts, ed., "Performance evaluation and design of multiservice networks," Final Report of the COST 224 project, Commission of the European Communities, 1992.

[16] J. W. Roberts, B. Bensaou and Y. Canetti, "A traffic control framework for high speed data transmission," Proc. IFIP Workshop on "Modelling and Performance Evaluation of ATM Technology," Eds. H. Perros, G. Pujolle, Y. Takahashi, Elsevier, 1993, pp. 243–262.

[17] J. S. Turner, "New directions in communications (or which way to the information age?)," IEEE Commun. Magazine, Oct. 1986.

[18] T. Worster, "Modelling deterministic queues: the Leaky Bucket as an arrival process," Proc. ITC-14, Eds. J. Labetoulle and J. W. Roberts, Elsevier, 1994, pp. 581–590.

[19] N. Yamanaka, Y. Sato and K. I. Sato, "Performance limitations of leaky bucket algorithm for usage parameter control of bandwidth allocation methods," IEICE Trans. Commun. **E75-B**, No. 2, 1992, pp. 82–86.