A New Approach for Allocating Buffers and Bandwidth to Heterogeneous, Regulated Traffic in an ATM Node

Anwar Elwalid & Debasis Mitra

AT&T Bell Laboratories
Murray Hill, NJ 07974

Robert H. Wentworth

AT&T Bell Laboratories
200 Schulz Drive
Red Bank, NJ 07701

Abstract

A new approach to determining the admissibility of variable bit rate (VBR) traffic in buffered digital networks is developed. In this approach all traffic presented to the network is assumed to have been subjected to *leaky-bucket* regulation, and extremal, periodic, on-off regulated traffic is considered; the analysis is based on fluid models. Each regulated traffic stream is allocated bandwidth and buffer resources which are independent of other traffic. Bandwidth and buffer allocations are traded off in a manner optimal for an adversarial situation involving minimal knowledge of other traffic. This leads to a single-resource statistical-multiplexing problem which is solved using techniques previously used for unbuffered traffic. VBR traffic is found to be divisible into two classes, one for which statistical multiplexing is effective and one for which statistical multiplexing is ineffective in the sense that accepting small losses provides no advantage over requiring lossless performance. The boundary of the set of admissible traffic sources is examined, and is found to be sufficiently linear that an *effective bandwidth* can be meaningfully assigned to each VBR source, so long as only statistically-multiplexable sources are considered, or only non-statistically-multiplexable sources are considered. If these two types of sources are intermixed, then nonlinear interactions occur and fewer sources can be admitted than a linear theory would predict. A qualitative characterization of the nonlinearities is presented. The complete analysis involves conservative approximations; however, admission decisions based on this work are expected to be less overly conservative than decisions based on alternative approaches.

1 INTRODUCTION

Research in high speed communications has been dominated in recent years by issues related to services based on the Asynchronous Transfer Mode (ATM). An important issue is congestion control. This issue is intertwined with the notion of the *capacity* of the network measured in users for given quality of service, and the administration of admission control in real-time, where the goal is to admit users upto capacity. The statistical nature of a significant part of the traffic, its burstiness and variability, and the stringency of the quality of service requirements combine to pose challenges of unprecedented difficulty. An essential prerequisite is the regulation of traffic at network edges, as shown in Figure 1.1. The Leaky Bucket regulator [20] has become ubiquitous in current service offerings. In our work we assume that all traffic is subject to regulation by such devices, which have three parameters: the token rate r, which bounds the long-term average rate of the regulated traffic; B_T, the token buffer size, which bounds the burst size;

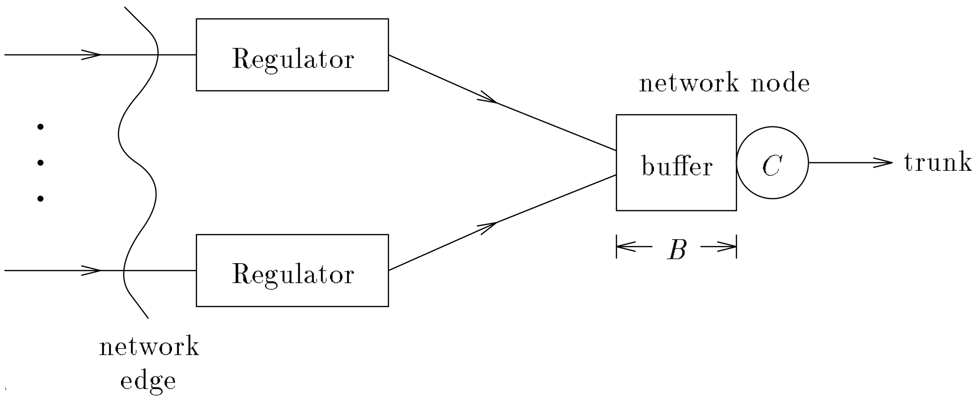

Figure 1.1: Regulated traffic offered to an ATM node.

P, which bounds the peak rate. Also shown in Figure 1.1 is the network node, which for our purposes is composed of two resources, a buffer of size B and a trunk of transmission bandwidth C.

One approach to network capacity calculations and admission control has been based on statistical source models. See for instance [1, 18], and also [5] for analysis of the combined effects of such models and regulators. The reliability of source models in this approach is a matter of concern, however. A contrasting approach is based on worst case congestion effects at the network nodes, given that regulators exist at the entry points. That is, traffic sources are adversarial to the 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 [16] for an example along these lines. Another option is to allow very small loss probabilities, say 10^{-9}, and to extract multiplexing gains from the statistical independence of traffic processes. This approach, which combines independent, worst case source behavior with a quality of service guarantee in the form of a small loss probability, is taken in this paper. It should counter the overly conservative usage implied by the deterministic bounds [24], and lead to more efficient use of network resources. This is little prior statistical analysis of burst-scale congestion in the multiplexing of periodic, on-off source traffic. Examples are [12] and [19]. In these papers the analysis is approximate and the calculations of the overflow probability, as pointed out in [12], are quite laborious. Also, no attempt is made to carry the analysis through to the stage where it can be used in real-time administration of admission control.

Before proceeding to details we give an overview of the paper in three parts.

(i) We examine the capacity of a single network node and its admission control based on a probabilistic quality of service guarantee, specifically that the cell loss probability does not exceed L, a small number. All traffic submitted to the network node is leaky-bucket-regulated and, moreover, is on-off, periodic and extremal in the sense that when the traffic source is on its rate is the maximum permitted, i.e. P, the amount of data generated in an on period, Q, is also maximum, where

$$Q = B_T \frac{P}{P - r} , \tag{1.1}$$

and the mean rate in each period is r, the maximum permitted. The randomness in the source model stems from the assumption of statistical independence of the traffic processes, specifically in their phases. A key aspect of our work is the heterogeneity of the regulators; in fact, we expect a proliferation quite early in the evolution of the networks. The analysis is based throughout on fluid models.

(ii) The approach taken in this paper, where one of the two key network resources is the buffer, is novel in its use of established techniques for analyzing *unbuffered* resources, such as the Chernoff large deviations approximation [9, 10, 11, 14]. The important bridge to the case where bandwidth is the other network resource is based on a technique that allows the two resources to be exchangeable. The first phase of the analysis concludes with a reduction from a two-resource to a one-resource problem, and in the second phase the aforementioned established techniques take over. Even in the second phase we have discovered non-linear phenomenon which is of utmost importance to multiplexing. When only a single class of statistically homogeneous traffic sources is present, we find that the statistical multiplexing gain exceeds unity only if the capacity of the single resource, say, the bandwidth, exceeds a critical value, C_c, which depends on the source characteristics. Therefore, for given bandwidth at the ATM node, C, in a heterogeneous environment there may exist source classes for which $C_c < C$ and other classes with $C_c \geq C$. We call the two classes "statistically-multiplexable" and "non-statistically-multiplexable" VBR sources, respectively, and abbreviate the terms to S-VBR and NS-VBR. While this dichotomy has long been appreciated at an intuitive level by researchers, this appears to be the first time that it has been analytically and quantifiably stated. We give a simple and explicit expression for C_c for a source class. The simultaneous presence of both types of sources makes the design and admission control problems harder.

(iii) This paper pays considerable attention to molding the analytic and qualitative results for the purposes of real-time administration of admission control. A key concept here is effective bandwidth, which has recently received considerable attention for both buffered [7, 8, 6, 21] and unbuffered [10, 11, 14] network resources. Underlying the concept is linearity of $\partial \mathcal{A}_L$, the boundary of the admissible set $\mathcal{A}_L$, which is the set of combinations of sources of various classes for which the quality of service requirement specified by L is satisfied. The linearity of $\partial \mathcal{A}_L$ is sometimes exact in an appropriate asymptotic limit, and sometimes approximate and empirical. We find in our numerical investigations that the boundary of the admissible set often deviates negligibly from linearity if all source classes are either S-VBR or NS-VBR. When both source types are present, the boundary is nonlinear, albeit approximately piece-wise linear, and with very different effective bandwidths in the different regions. Our analytic results identify the boundaries of the regions in $\partial \mathcal{A}_L$.

We proceed to a more detailed description of the traffic sources. As mentioned earlier, and also shown in Figure 1.1, all traffic offered to an ATM node is regulated at the network edge. The regulation is performed by leaky bucket devices with parameters (r, B_T, P); the ancillary burst-size parameter Q is obtained as indicated in (1.1). We assume that the departure process from such a regulator is extremal, on-off and periodic with indeterminate phases. We do not have conclusive proof that such processes are worst-case in the sense of maximizing the steady-state loss probability in the ATM node. However, this is suggested by the results of Doshi [4], Mitra and Morrison [15] and Worster [22] (see also [23]). Doshi has various related results in general frameworks. In [15] it is specifically shown that, for a single-resource problem, which is what we have in this paper after the initial phase of the analysis, the Chernoff estimate of the loss is maximized by extremal, on-off, periodic processes with independent, uniformly distributed random phases. This result holds for heterogeneous regulators.

We denote by $\Omega(t)$ the rate process which is the output of the regulator, see Figure 1.2. In the figure

$$T_{\text{on}} = \frac{B_T}{P - r} \ , \quad \text{and} \quad T_{\text{off}} = \frac{B_T}{r} \ . \tag{1.2}$$

We let the period be denoted by $T = T_{\text{on}} + T_{\text{off}} = Q/r$. Also, the source activity factor,

$$w_\Omega = \ \Pr(\Omega > 0) = \frac{r}{P} \ . \tag{1.3}$$

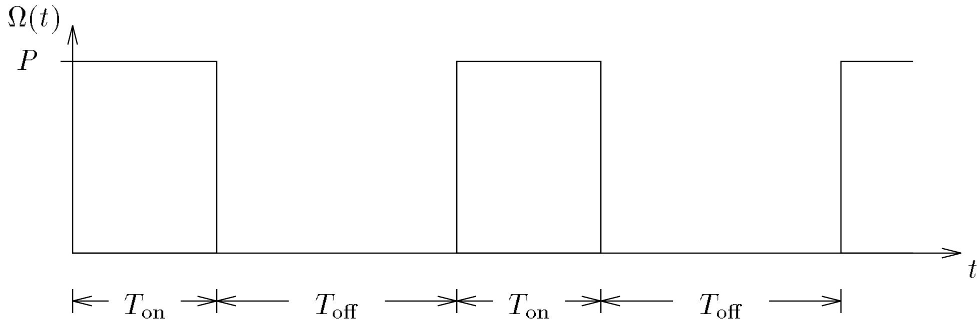

Figure 1.2: Periodic, on-off rate process departing from regulator.

The paper is organized as follows. Section 2, which follows, is on lossless multiplexing. Here the basic idea of exchangeability of buffer and bandwidth is introduced and an optimum trade-off is struck, which gives rise to the effective bandwidth for lossless multiplexing. This quantity, denoted by e_0, is the benchmark to which comparisons are made later in calculating the statistical multiplexing gain. Section 3 on statistical multiplexing exploits the non-coincidence of bursts due to the random phases of the regulated traffic. The loss probability is estimated by the Chernoff bound, and a refinement of it. The admissible set and the calculations leading to the effective bandwidth, e, which is based primarily on the work of Hui [10], are reviewed. Section 4 on qualitative theory gives, first, the analysis for a homogeneous class of traffic sources, which leads to C_c, the critical bandwidth for the statistical multiplexing gain given by e_0/e to exceed unity. The ideas are generalized to the heterogeneous environment, giving rise to results on the nonlinear composition of the boundary of the admissible set. Finally the sensitivity of the statistical multiplexing gain on the token rate r is studied and a corollary establishes a critical value of r_c. Section 5 reports on extensive numerical investigations on all the topics highlighted by the analysis. Finally Section 6 concludes with discussions on miscellaneous topics.

2 LOSSLESS MULTIPLEXING

In this section we consider the base case where the quality of service requirement is lossless performance. Each virtual circuit is allocated buffer b_0 and bandwidth e_0, and admission is denied if uncommitted resources at the node are inadequate to satisfy the allocations. A key feature is the heterogeneity of the virtual circuits and hence the goal is to characterize the set of admissible combinations of virtual circuits with various associated regulator parameters. We first consider a single source in a framework which we call its Virtual Buffer/Trunk System. Next, in Section 2.2, the multiplexing of sources is considered and the admissible set obtained.

2.1 Virtual Buffer/Trunk System

Consider a single source with periodic, on-off rate process $\Omega(t)$, see Figure 1.2, which supplies an infinite buffer with a trunk of bandwidth c, as shown in Figure 2.1. We assume that $c \geq r$, the stability condition. Let $v(t)$ and $u(t)$ respectively denote the *buffer content* and the *utilized bandwidth* at time t, where the latter is an on-off process which takes values c and 0 depending on whether the buffer is not empty or empty, respectively. See Figure 2.2.

The stability condition ensures that the buffer is emptied in every cycle of length T. Let D_{on} and D_{off} respectively denote the time in each cycle that the buffer and the trunk is utilized. Clearly, $D_{\text{on}} + D_{\text{off}} = T$,

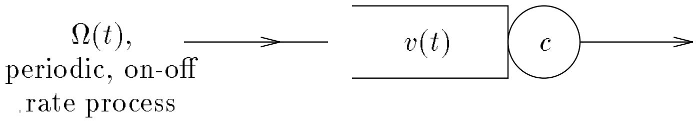

Figure 2.1: Virtual buffer/trunk system for a single source.

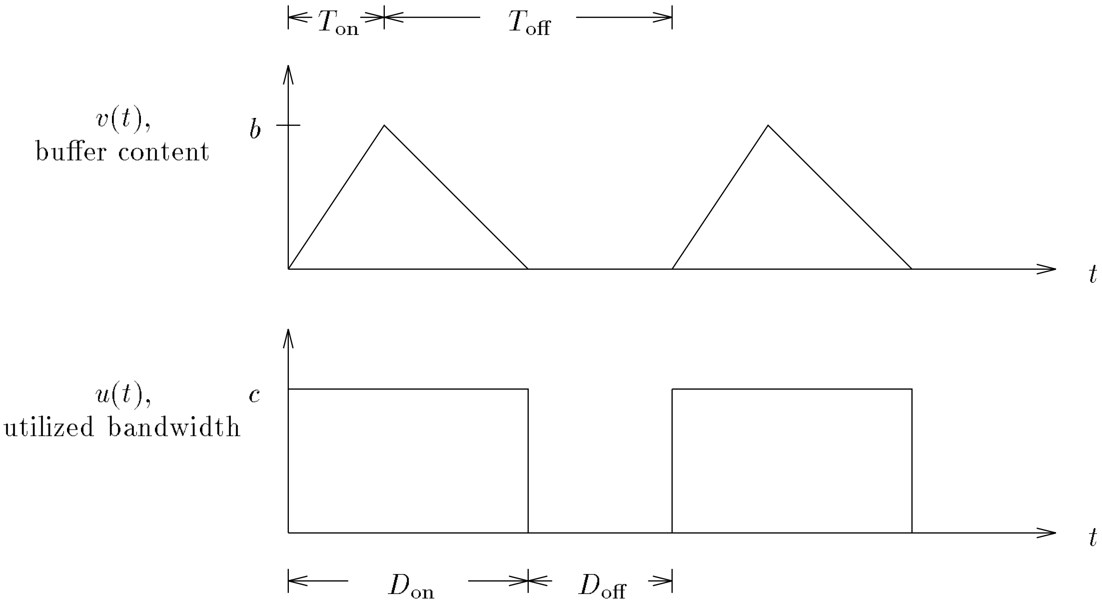

Figure 2.2: Processes of the virtual system for a single source.

and D_{on} exceeds T_{on} by the time taken by the buffer to go from its state of maximum occupancy to empty. Let b denote the maximum buffer content, $\sup_{t>0} v(t)$. Then

$$D_{\text{on}} = T_{\text{on}} + \frac{b}{c}, \quad D_{\text{off}} = T_{\text{off}} - \frac{b}{c} \ . \tag{2.1}$$

Hence w, fraction of time that the virtual buffer/trunk system is busy,

$$w = \frac{D_{\text{on}}}{D_{\text{on}} + D_{\text{off}}} = \frac{T_{\text{on}}}{T_{\text{on}} + T_{\text{off}}} + \frac{b}{c} \cdot \frac{1}{T_{\text{on}} + T_{\text{off}}} = w_{\Omega} + \frac{b}{c} \cdot \frac{1}{T} \ . \tag{2.2}$$

Importantly, $w \geq w_{\Omega}$. Also,

$$b = B_T - \frac{B_T(c-r)}{P-r} \ . \tag{2.3}$$

Hence $b \leq B_T$. Note that b decreases linearly with increasing c for $r \leq c \leq P$, and that the extreme points of (b, c) are (B_T, r) and $(0, P)$.

2.2 Multiplexing

We now use the behavior of the processes of the virtual buffer/trunk system for single sources to design admission control for heterogeneous sources sharing the nodal resources.

First, suppose there is a set of I virtual circuits such that the traffic source for circuit i $(1 \leq i \leq I)$, is an on-off process $\Omega_i(t)$. The aggregate traffic $\sum_{i=1}^{I} \Omega_i(t)$ will be associated with buffer content $V(t)$ and a utilized bandwidth $U(t)$, where the latter is an on-off process which takes values C and 0.

Consider assigning each virtual circuit a virtual buffer/trunk. Circuit i would be allocated bandwidth c_i ($r \le c_i \le P$), and also buffer b_i, the peak buffer occupancy given in (2.3). Then, if

$$\sum_{i=1}^{I} c_i \le C \ , \tag{2.4}$$

it follows that

$$\sum_{i=1}^{I} u_i(t) \le C \tag{2.5}$$

and

$$\sum_{i=1}^{I} v_i(t) \ge V(t) \ . \tag{2.6}$$

Thus, the sum of the virtual buffer usages bounds the aggregate buffer usage. We shall conservatively use the sum of the virtual buffer usages, $\sum_{i=1}^{I} v_i(t)$, as an estimate of the aggregate buffer usage $V(t)$. Equation (2.6) implies that if our estimate does not exceed the available buffer B, then neither will the actual buffer usage.

In the case where all sources are identical except for phase, one finds

$$\sum_{i=1}^{I} b_i = \sup_{\text{phases}, t} \left\{ \sum_{i=1}^{I} v_i(t) \right\} = \sup_{\text{phases}, t} \{V(t)\} \ . \tag{2.7}$$

Thus, in this case our estimate accurately predicts the maximum buffer usage, and hence the border between lossless and lossy transmission. The estimate will be less precise when dissimilar sources are involved, or when small losses are allowed.

Our main concern now is the selection of the parameters c_i, which are free at this point. If all the sources are known and are heterogeneous, then determining the values c_i which produce the lowest estimated peak buffer usage without exceeding the total bandwidth constraint (2.4) involves a complicated optimization. A simplification results if the c_i for a particular source must be chosen independent of any knowledge of the character of other sources. Then a little thought makes it clear that optimal resource management requires that the per-circuit allocations (b_i, c_i) be such that both network resources will always be exhausted at the same time for any aggregation of circuits using the same allocation rule. This occurs when

$$\frac{b_i}{B} = \frac{c_i}{C} \ , \tag{2.8}$$

i.e., when the allocations are in proportion to their respective nodal capacities. Although we shall sometimes consider cases where all sources are taken to be known, we shall continue to abide by (2.8) because of the simplification it provides. Using (2.8) in lieu of doing a calculation invoking simultaneous knowledge of all circuits leads to admission decisions which are exact if all sources are identical, but conservative if sources are heterogeneous.

Observe that for given B and C, (2.3) and (2.8) determine a unique allocation which we denote by (b_0, e_0). The solution is depicted in Figure 2.3. The quantity e_0, which is of obvious importance throughout this paper, is referred to as the "effective bandwidth for lossless performance". The justification for the term stems from the fact that if $e_{0,i}$ is the effective bandwidth of the i^{th} virtual circuit, then the set of circuits $\{1, 2, \ldots, I\}$ is admissible if

$$\sum_{i=1}^{I} e_{0,i} \le C \ . \tag{2.9}$$

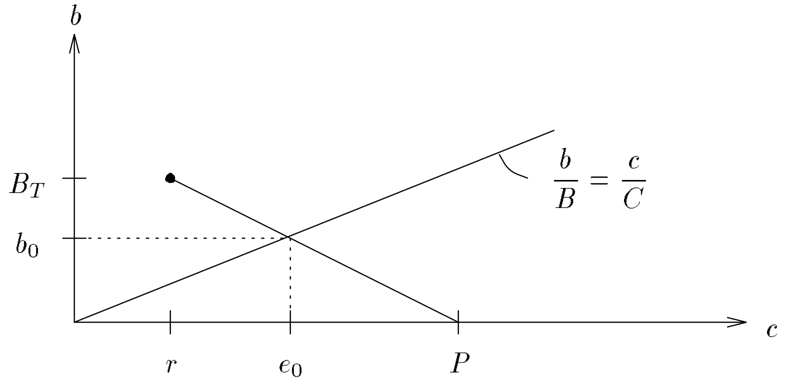

Figure 2.3: Determination of buffer and bandwidth allocation (b_0, e_0).

The admissible set may equivalently be given by $\sum_{i=1}^{I} b_{0,i} \leq B$, where $b_{0,i}$ is the "effective buffer requirement" for virtual circuit i; however, we will use the former description.

Exceptions to the above procedure arise if there is no intersection of the straight lines in Figure 2.3. An intersection at (b_0, e_0) exists if and only if

$$\frac{C}{B} > \frac{r}{B_T} \ . \tag{2.10}$$

Sources which violate (2.10) are "bandwidth-limited" and in such cases a natural choice for the allocated bandwidth is

$$e_0 = r \ . \tag{2.11}$$

There are the following options available in the corresponding buffer allocations: (i) $b_0 = rB/C$, or (ii) $b_0 = B_T$. The former allocation, which in effect inflates B_T to rB/C and is therefore more liberal, has the advantage that the administration of admission control and resource allocation is routine, i.e., as if (2.10) holds. This option is recommended if the margin of violation of (2.10) is small, which is more typical. On the other hand, if the margin is large and the second option is followed then it is necessary to go to the extreme length of recalculating the residual B and C available to the non-bandwidth-limited sources whenever such a bandwidth-limited virtual circuit is admitted to the network.

We end this section by giving formulas for e_0 and w, which will be useful later:

$$e_0 = \begin{cases} \dfrac{P}{1 + \dfrac{B/C}{B_T}(P - r)} & \text{if } \ r \leq \dfrac{B_T}{B/C} \\[2ex] r & \text{if } \ \dfrac{B_T}{B/C} \leq r < P \end{cases} \tag{2.12}$$

$$w = \frac{w_\Omega P}{e_0} = \frac{r}{e_0} \tag{2.13}$$

Note that in the above expressions e_0 and w depend on the nodal buffer B and bandwidth C only through B/C. This important fact should not be surprising since the ratio of buffer and bandwidth allocated to individual sources is equal to the ratio of the buffer and bandwidth at the node.

The formula for e_0 perhaps is most easily remembered when expressed in terms of different variables:

$$e_0 = \begin{cases} \dfrac{P}{1 + \dfrac{T_{\text{buf}}}{T_{\text{on}}}} & \text{if } T_{\text{buf}} \le T_{\text{off}} \\ r & \text{if } T_{\text{off}} \le T_{\text{buf}} \end{cases} \tag{2.14}$$

where

$$T_{\text{buf}} = \frac{B}{C} \tag{2.15}$$

is the maximum delay time of the buffer.

3 STATISTICAL MULTIPLEXING

In contrast to the framework of the preceding section, here we espouse the statistical viewpoint and allow small loss probabilities. We establish a framework for extracting statistical multiplexing gains from (i) the assumption of statistical independence of traffic sources, and (ii) the exploitation of the unutilized time-varying portions of the allocated resources to virtual circuits. This typically leads to better utilization of network resources.

It will be helpful to assume that the virtual circuits are grouped by classes. Let J denote the number of classes, where each class is associated with a particular set of parameters for the regulator, (r, B_T, P), and K_j denotes the number of virtual circuits of class j.

3.1 Reduction to a Single Resource Loss Model

We begin with the characterization of the i^{th} virtual circuit of class j by two stationary random processes $u_{ji}(t)$ and $v_{ji}(t)$, which respectively denote the utilized bandwidth and buffer content for the circuit's virtual buffer/trunk system. As shown in Figure 2.2, these two processes are synchronized in their on-off behavior. To make their correspondence closer for ease of analysis, we shall conservatively bound the process $v_{ji}(t)$ by an on-off process, which takes the peak value for the on period. Thus we shall view $u_{ji}(t)$ and $v_{ji}(t)$ as two synchronized on-off processes, which respectively take values $e_{0,j}$ and $b_{0,j}$ while on and 0 while off. The parameters $e_{0,j}$ and $b_{0,j}$ are calculated by the procedure in the preceding section. With this formalism in place we need only to deal with either $u_{ji}(t)$ or $v_{ji}(t)$, and we choose the former.

On account of the assumed statistical independence of traffic sources, the processes $u_{ji}(t)$ $(i = 1, 2, \ldots, K_j)$ of the same source class have identical templates and differ only in their phase, i.e.,

$$u_{ji}(t) = u_j(t + \theta_{ji}) \ , \tag{3.1}$$

where $u_j(t)$ is a deterministic, periodic on-off function with period T_j, which takes values $e_{0,j}$ and 0 for fractions w_j and $(1 - w_j)$, respectively; the phases θ_{ji} are independent random variables uniformly distributed in the interval T_j.

Our performance measure is the loss probability P_{loss},

$$P_{\text{loss}} = \Pr(U > C) \ , \tag{3.2}$$

where the total instantaneous load $U = \sum_{j=1}^{J} \sum_{i=1}^{K_j} u_{ji}$. Hence P_{loss} is the fraction of time that the aggregated demand for bandwidth from all sources exceeds the nodal bandwidth. From the earlier discussions, time periods during which losses occur at the network node due to a full buffer are contained in periods during which $U(t) > C$. We shall take for our quality of service requirement,

$$P_{\text{loss}} \leq L \ , \tag{3.3}$$

where L is a small number, such as 10^{-9}.

3.2 Chernoff Bound, Admissible Set, Effective Bandwidth

Our primary means for estimating P_{loss} is the Chernoff bound, which is reviewed here, and a refined large deviations approximation based on it, which is described in Sec. 3.3. The instantaneous loads u_{ji} are independent, nonnegative random variables with moment generating functions

$$M_j(s) = E[\exp(su_{ji})] = \int_0^\infty e^{sx} dW_j(x) \ , \tag{3.4}$$

where

$$W_j(x) = \Pr(u_{ji} \leq x) \ . \tag{3.5}$$

Chernoff's bound gives [2],

$$\log P_{\text{loss}} \leq -F_{\boldsymbol{K}}(s^*) \ , \tag{3.6}$$

where

$$F_{\boldsymbol{K}}(s) = sC - \sum_{j=1}^{J} K_j \log M_j(s) \tag{3.7}$$

and $F_{\boldsymbol{K}}(s^*) = \sup_{s \geq 0} F_{\boldsymbol{K}}(s)$.

If $C \to \infty$ and $K_j/C = O(1)$, then [3], [17]

$$\log P_{\text{loss}} = -F_{\boldsymbol{K}}(s^*) \left[1 + O\left(\frac{\log C}{C} \right) \right] \ . \tag{3.8}$$

Hence the asymptotic large deviations approximation is $P_{\text{loss}} \sim \exp(-F_{\boldsymbol{K}}(s^*))$.

To avoid trivialities assume the stability condition,

$$\sum_{j=1}^{J} K_j E(u_{ji}) < C \ , \quad \text{and} \quad \lim_{s \to \infty} \sum_{j=1}^{J} K_j \frac{M_j'(s)}{M_j(s)} > C \ , \tag{3.9}$$

(the prime denotes derivative) since without the latter condition there is no loss. Also note that

$$F_{\boldsymbol{K}}'(s) = C - \sum_{j=1}^{J} K_j \frac{M_j'(s)}{M_j(s)} \ . \tag{3.10}$$

It is easy to verify that $F_{\boldsymbol{K}}(s)$ is a strictly concave function with a unique maximum at $s = s^*$, which is the positive root of the equation $F_{\boldsymbol{K}}'(s) = 0$.

Of particular interest here is the case of binomially distributed u_{ji}, where $w_j = \Pr(u_{ji} = e_{0,j})$ and $1 - w_j = \Pr(u_{ji} = 0)$. Then

$$F_{\boldsymbol{K}}(s) = sC - \sum_{j=1}^{J} K_j \log\{1 - w_j + w_j \exp(se_{0,j})\} \ , \tag{3.11}$$

and s^* is obtained by solving the equation

$$\sum_{j=1}^{J} \frac{K_j w_j e_{0,j} \exp(s e_{0,j})}{1 - w_j + w_j \exp(s e_{0,j})} = C \ . \tag{3.12}$$

In the single-class case, i.e. $J = 1$, the resulting expressions are illuminating: for $a = (C/e_0)/K$,

$$s^* = \frac{1}{e_0} \log \left[\frac{a}{1-a} \cdot \frac{1-w}{w} \right] , \tag{3.13}$$

$$F_K(s^*) = K \left[a \log \left(\frac{a}{w} \right) + (1-a) \log \left(\frac{1-a}{1-w} \right) \right] . \tag{3.14}$$

The above expression may be used to obtain $K_{\max}$, which is the value of K for which

$$F_K(s^*) = \log(1/L) . \tag{3.15}$$

In light of (3.3) and (3.6), the quality of service requirement on P_{loss} is satisfied for all $K \leq K_{\max}$. Similarly, in the case of multiple classes, we are interested in the *admissible set*

$$\mathcal{A}_L(B, C) = \{ \boldsymbol{K} : F_{\boldsymbol{K}}(s^*) \geq \log(1/L) \} \tag{3.16}$$

Hui [10] has shown that the complement of the admissible set is convex; hence the points on the boundary, which we denote by $\partial\mathcal{A}_L$, also satisfy (3.15) with K replaced by J-tuples $(K_1, \ldots, K_J)$.

For a wide variety of conditions the boundary $\partial\mathcal{A}_L$ is closely approximated by an appropriately chosen linear hyperplane. (There are notable exceptions, which the theory in the next section will systematically treat; however, even in these cases the boundary is approximately piece-wise linear and hence the following comments apply with obvious modification.) For instance, following Hui [10] and Kelly [11], we may pick a point $\tilde{\boldsymbol{K}}$ on $\partial\mathcal{A}_L$ and obtain the tangent hyperplane to $\partial\mathcal{A}_L$ which touches $\partial\mathcal{A}_L$ at $\tilde{\boldsymbol{K}}$. Let the region constrained by this hyperplane by denoted by $\tilde{\mathcal{A}}_{L,\tilde{K}}$. Clearly, from the aforementioned concavity of $\mathcal{A}_L$, $\tilde{\mathcal{A}}_{L,\tilde{K}}$ is a conservative bound to $\mathcal{A}_L$. Moreover, denoting

$$\tilde{\mathcal{A}}_{L,\tilde{K}} = \left\{ \boldsymbol{K} : \sum_{j=1}^{J} K_j e_j \leq C \right\} , \tag{3.17}$$

the parameters e_j which define $\tilde{\mathcal{A}}_{L,\tilde{K}}$ are readily obtained [10], [11]:

$$e_j = \frac{\log M_j(\tilde{s}^*)}{\tilde{s}^* + (\log L)/C} \quad (j = 1, 2, \ldots, J) , \tag{3.18}$$

where $\tilde{s}^*$ is where $F_{\tilde{\boldsymbol{K}}}(s)$ is maximized.

The quantity e_j may justifiably be called the effective bandwidth of class j traffic sources. The statistical multiplexing gain g compares this quantity to $e_{0,j}$, which is the effective bandwidth for lossless performance,

$$g_j = e_{0,j}/e_j \quad (j = 1, 2, \ldots, J) . \tag{3.19}$$

In our numerical investigations we have found the end results to be fairly insensitive to reasonable choices of $\tilde{\boldsymbol{K}}$. One selection which works well is $\tilde{K}_j = K_{\max,j}/2$, where $K_{\max,j}$ is the maximum number of admissible sources when only class j sources are admitted.

An optimistic (upper) bound of $\mathcal{A}_L$, which follows from the aforementioned concavity of the admissible region is

$$\left\{ K : \sum_{j=1}^{J} K_j e_j^u = C \right\} , \tag{3.20}$$

where $e_j^u = C/K_{\max,j}$. These optimistic estimates of the effective bandwidths are considerably easier to calculate, since they are obtained by considering each source class in isolation, and in many cases are quite close to the conservative estimates in (3.18). Such is the case when $\partial\mathcal{A}_L$ is close to being linear. Examples in Section 5 will demonstrate that this is quite common.

3.3 Refined Large Deviations Approximations

In our numerical investigations we use a refinement to the large deviations approximation (3.8) and bound (3.6) [3, 17]. For the asymptotic scaling in which $K_j/C = O(1)$ and $C \to \infty$, the refined approximation is

$$P_{\text{loss}} = \frac{\exp\{-F_{\boldsymbol{K}}(s^*)\}}{s^*\sigma(s^*)\sqrt{2\pi}}[1 + o(1)] , \tag{3.21}$$

where s^* is obtained by solving $F'_{\boldsymbol{K}}(s) = 0$, i.e., exactly as before, and $\sigma^2(s) = \frac{\partial^2}{\partial s^2} \log E[e^{sU}]$, where U is the total instantaneous load, see (3.2). More specifically,

$$\sigma^2(s) = \sum_{j=1}^{J} K_j \left[\frac{M_j''(s)}{M_j(s)} - \left\{ \frac{M_j'(s)}{M_j(s)} \right\}^2 \right] . \tag{3.22}$$

In the particular case of binomially distributed random variables u_{ji}, the quantities in the above expression have appeared before, except for $M_j''(s)$ which is given by $w_j e_{0,j}^2 \exp(se_{0,j})$.

The dominant exponential (in C) term is common to both the base and refined large deviations approximations. Hence, the fundamental qualitative properties are not affected. We will find it more convenient to conduct our qualitative investigation in terms of the simpler base approximation. In quantitative terms a rough rule of thumb is that the refinement adds about 10% to the calculated nodal capacity.

4 QUALITATIVE THEORY

In this section we develop a qualitative theory based on the contents of the preceding section. We examine more carefully the existence of solutions $\boldsymbol{K}$ to the equation given earlier for determining the boundary $\partial\mathcal{A}_L$ of the admissible set, namely,

$$F_{\boldsymbol{K}}(s^*) = \log(1/L) . \tag{4.1}$$

The topic addressed here is whether the admissible set $\mathcal{A}_0$ for lossless performance is a proper subset of $\mathcal{A}_L$, i.e., whether statistical multiplexing gains exist. We also determine conditions under which the sets $\mathcal{A}_0$ and $\mathcal{A}_L$ coincide either completely or partially. When only a single source class exists these results illuminate conditions under which the statistical multiplexing gain, $g = e_0/e$ either exceeds or equals unity. We show that there exist critical values C_c, such that if $C \le C_c$ then $g = 1$ and if $C > C_c$ then $g > 1$. Similar critical values r_c exist for the token rate. Simple expressions for C_c and r_c are obtained. For given nodal bandwidth C we thus obtain statistically-multiplexable and non-statistically-multiplexable VBR source classes. In the context of multiple source classes we obtain results here on the constitution of the boundary of the admissible set. The results here help to explain the observations from numerical investigations, which are reported in the next section, on the almost-linearity of the boundary of the admissible in the presence of only S-VBR sources, and the qualitatively different boundary, with piece-wise linear segments, which is obtained when both S-VBR and NS-VBR sources are present.

4.1 Single Source Class

From (3.14),

$$\frac{\partial}{\partial K} F_K(s^*) = \log\left(\frac{1-a}{1-w}\right) \tag{4.2}$$

where $a = C/(Ke_0)$. From the stability condition (3.9), $w < a$. Hence $\partial F_K(s^*)/\partial K < 0$. Also, for $K = C/e_0 + 0$, $\partial F_k(s^*)/\partial K = -\infty$, and, importantly,

$$F_K(s^*) = \frac{C}{e_0}\log(1/w)\ . \tag{4.3}$$

From the expression for s^* in (3.13),

$$s^* = \frac{1}{e_0}\log\left[\frac{1}{K - C/e_0}\cdot\frac{C}{e_0}\cdot\frac{1-w}{w}\right]\ , \tag{4.4}$$

it follows that $s^* \to \infty$ as $K \to C/e_0 + 0$, and $s^* = 0$ for $K = \frac{1}{w}\frac{C}{e_0}$. The following summarizes the salient features of $F_K(s^*)$. See also Figure 4.1.

Proposition 4.1 With C, w, e_0 held fixed, and K in the interval $[\frac{C}{e_0}, \frac{1}{w}\frac{C}{e_0}]$

(i) $F_K(s^*)$ is monotonic, strictly decreasing with increasing K.

(ii) For $K = C/e_0 + 0$, $s^* = \infty$, $\frac{\partial}{\partial K}F_K(s^*) = -\infty$ and $F_K(s^*) = \frac{C}{e_0}\log(1/w)$.

In particular, $\frac{C}{e_0}\log(1/w)$ is the maximum value of $F_K(s^*)$ and it is achieved at $K = C/e_0$. ∎

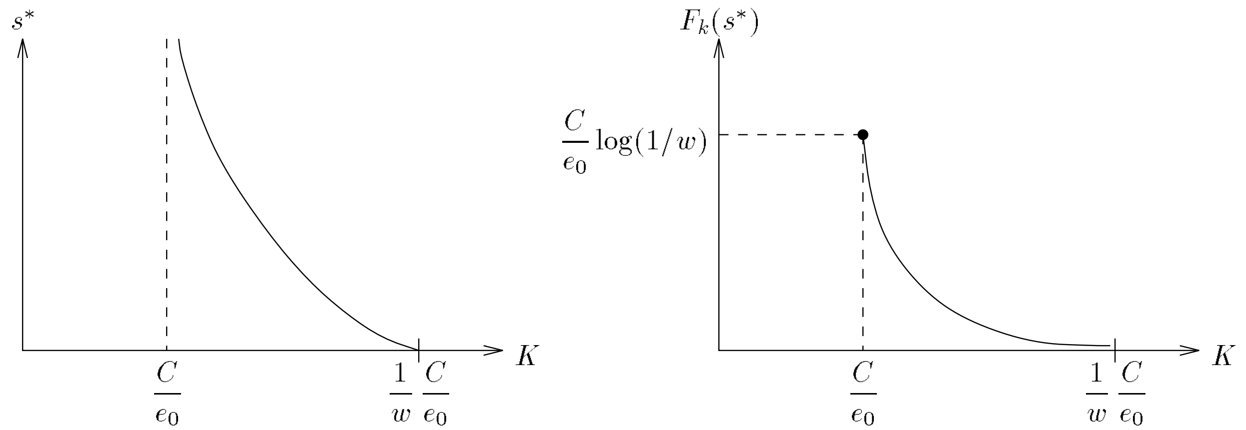

Figure 4.1: Sketches of s^* and $F_K(s^*)$ for model with single source class.

Now consider the calculation of $K_{\max}$, the largest value of K for which $P_{\text{loss}} \le L$, which is given by the solution K of $F_K(s^*) = \log(1/L)$. Observe the following dichotomy: if $\log(1/L) < \frac{C}{e_0}\log(1/w)$, the maximum value of $F_K(s^*)$, then $K_{\max} > C/e_0$. Hence the statistical multiplexing gain $g = e_0/e > 1$, since $e = C/K_{\max}$. If, on the other hand, $\log(1/L) \geq \frac{C}{e_0}\log(1/w)$ then no solution to (4.5) exists; however, we know independently that for $K = C/e_0$, $P_{\text{loss}} = 0$ and hence $K_{\max} = C/e_0$. In this case we obtain $g = 1$. The above facts are now summarized.

Definition: The critical bandwidth C_c is the smallest quantity such that the statistical multiplexing gain exceeds unity for all $C > C_c$.

Tacit to the above definition is the assumption that B varies with C in such a manner that B/C is held fixed.

Proposition 4.2

$$C_c = \frac{e_0 \log(1/L)}{\log(1/w)} \tag{4.5}$$

■

Recall from (2.12) and (2.13) that an implication of our resource allocation procedure is that e_0 and w depend on the nodal resource capacities B and C only through B/C. Hence this is also true for C_c.

4.2 Multiple Source Classes

Conditions in (3.9) give

$$\sum_{j=1}^{J} K_j w_j e_{0,j} < C < \sum_{j=1}^{J} K_j e_{0,j} \ . \tag{4.6}$$

Also, for $C \geq \sum K_j e_{0,j}$ obviously $P_{\text{loss}} = 0$. Our main result is

Proposition 4.3

(i) $F_{\boldsymbol{K}}(s^*)$ has a maximum value m_F, where

$$m_F = \max_{1 \leq j \leq J} \frac{C}{e_{0,j}} \log(1/w_j) \tag{4.7}$$

(ii) Say that the above maximum is reached at an unique j^*. Then $F_{\boldsymbol{K}}(s^*)$ reaches its maximum value at the following unique corner point $\boldsymbol{K}$:

$$\begin{aligned} K_j &= C/e_{0,j} \quad \text{if } j = j^* \\ &= 0 \qquad\quad \text{otherwise.} \end{aligned} \tag{4.8}$$

(iii)

$$\text{If} \qquad \log(1/L) > m_F \ \text{ then } \ \mathcal{A}_L = \mathcal{A}_0 \ . \tag{4.9}$$

$$\text{If} \qquad \log(1/L) < \min_j \frac{C}{e_{0,j}} \log(1/w) \ \text{ then } \ \mathcal{A}_0 \subset \mathcal{A}_L \ . \tag{4.10}$$

Otherwise $\mathcal{A}_0 \subseteq \mathcal{A}_L$, i.e., the boundaries of $\mathcal{A}_0$ and $\mathcal{A}_L$ coincide in part and elsewhere $\mathcal{A}_L$ subsumes $\mathcal{A}_0$. Specifically, $\boldsymbol{K}$ on the boundary of $\mathcal{A}_0$ such that $\sum K_j \log(1/w_j) \leq \log(1/L)$ is also on the boundary of $\mathcal{A}_L$, and when the inequality is false $\boldsymbol{K}$ is contained in the proper interior of $\mathcal{A}_L$.

Proof: Observe from (3.12) that as $s^* \to \infty$, $\sum_j K_j e_{0,j} \to C$. Hence, for $\sum_j K_j e_{0,j} = C$, from (3.11),

$$\begin{aligned} F_{\boldsymbol{K}}(s^*) &= \sum_{j=1}^{J} K_j \log(1/w_j) + s^* \left(C - \sum_{j=1}^{J} K_j e_{0,j} \right) \\ &= \sum_{j=1}^{J} K_j \log(1/w_j) \end{aligned} \tag{4.11}$$

Next,

$$\frac{\partial}{\partial K_i} F_{\boldsymbol{K}}(s^*) = -\log(1 - w_i + w_i e^{s^* e_{0,i}}) < 0 \ . \tag{4.12}$$

Similarly, by taking the derivative of both sides of (3.12) with K_i, we find that

$$\frac{\partial s^*}{\partial K_i} < 0 \ . \tag{4.13}$$

At this point we observe that various features of the analysis for the single class carry over naturally to the case of multiple classes, namely, on $\partial \mathcal{A}_0 = \{\boldsymbol{K} : \sum K_j e_{0,j} = C\}$, $s^* = \infty$ and $F_{\boldsymbol{K}}(s^*) = \sum K_j \log(1/w_j)$, and as $\boldsymbol{K}$ increases in any direction away from $\partial \mathcal{A}_0$, s^* and $F_{\boldsymbol{K}}(s^*)$ decrease.

Now consider the behavior of $F_{\boldsymbol{K}}(s^*) = \sum K_j \log(1/w_j)$ on $\partial \mathcal{A}_0$. Since this function is linear in $\boldsymbol{K}$, its maximum on $\partial \mathcal{A}_0$ is reached at its boundary, typically a corner point. At the j^{th} corner point, where $K_i = C/e_{0,j}$ if $i = j$ and $K_i = 0$ otherwise, $F_{\boldsymbol{K}}(s^*) = (C/e_{0,j}) \log(1/w_j)$. Hence we arrive at the results in (i) and (ii).

Next consider the multiple class version of the problem considered previously for the single class, namely, the existence of solutions $\boldsymbol{K}$ in (4.1). If $\log(1/L) > m_F$ then clearly no solution exists. In this case, as with a single class, the requirement $P_{\text{loss}} \le L$ can only be met by having $P_{\text{loss}} = 0$, i.e., by restricting $\boldsymbol{K}$ to $\mathcal{A}_0$. Hence, in this case $\mathcal{A}_L = \mathcal{A}_0$. On the other and, if $F_{\boldsymbol{K}}(s^*) > \log(1/L)$ for all $\boldsymbol{K}$ on $\partial \mathcal{A}_0$ then, from the aforementioned range and monotonicity of $F_{\boldsymbol{K}}(s^*)$, solutions $\boldsymbol{K}$ to (4.1) exist in the complement of $\mathcal{A}_0$. Hence $\mathcal{A}_0 \subset \mathcal{A}_L$. Finally, consider the mixed case where $F_{\boldsymbol{K}}(s^*) \le \log(1/L)$ for some $\boldsymbol{K}$ on $\partial \mathcal{A}_0$ and the inequality is reversed for other $\boldsymbol{K}$ on $\partial \mathcal{A}_0$. In this case any $\boldsymbol{K}$ on $\partial \mathcal{A}_0$ for which the inequality holds is such that (4.1) does not have a solution and hence such a $\boldsymbol{K}$ is also an element of $\partial \mathcal{A}_L$. Hence, any $\boldsymbol{K}$ on $\partial \mathcal{A}_0$ for which $\sum K_j \log(1/w_j) \le \log(1/L)$ is also an element of $\partial \mathcal{A}_L$. Moreover, all other $\boldsymbol{K}$ on $\partial \mathcal{A}_0$ are contained in the proper interior of $\mathcal{A}_L$.

This concludes the proof of Proposition 4.3. ■

To illustrate the result for a two-class model we need to consider three cases:

(i) $\quad \log(1/L) > m_F, \quad \text{i.e.,} \quad C < \min(C_{c,1},\ C_{c,2})$ (4.14)

(ii) $\quad \log(1/L) < \min_j \left\{ \frac{C}{e_{0,j}} \log(1/w_j) \right\}, \quad \text{i.e.,} \quad C > \max(C_{c,1},\ C_{c,2})$ (4.15)

(iii) $\quad C_{c,1} < C < C_{c,2}$, (4.16)

where $C_{c,j}$ is the critical bandwidth for class j:

$$C_{c,j} = e_{0,j} \log(1/L) / \log(1/w_j), \qquad j = 1, 2 \ . \tag{4.17}$$

According to Proposition 4.3, for case (i) we have that $\mathcal{A}_L = \mathcal{A}_0$. This situation is sketched in Figure 4.2, (i); note that the admissible set is a simplex. The performance implication is that there is no statistical

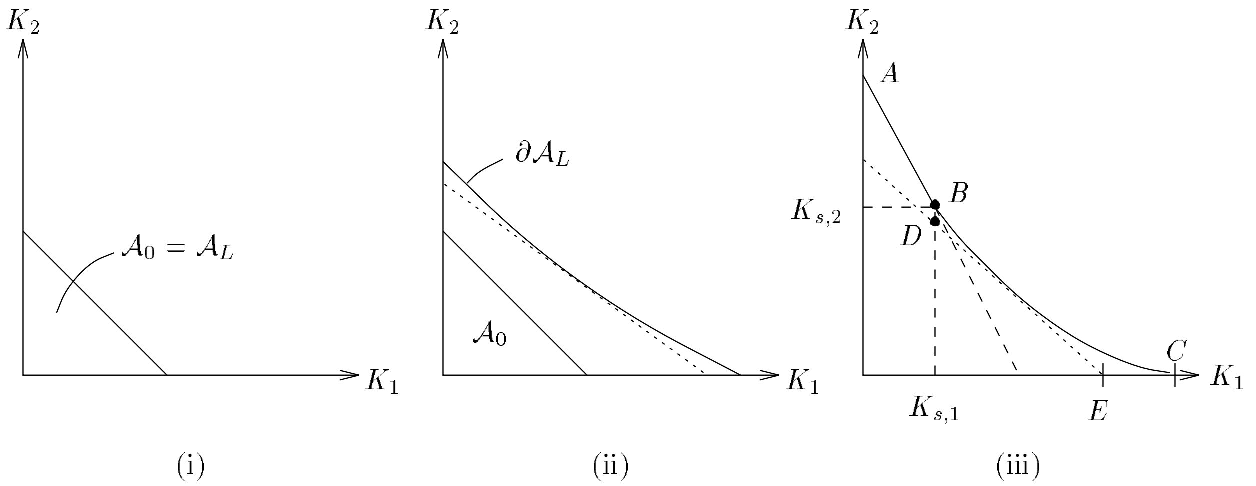

Figure 4.2: Illustration of various possibilities for admissible sets and statistical multiplexing gains for two traffic source classes.

multiplexing gain in this case. For this reason both source types are NS-VBR, i.e., "bad" in the sense that the sources are too "large" with respect to the node bandwidth. For case (ii), the result states that $\mathcal{A}_0$ is a proper subset of $\mathcal{A}_L$, as sketched in Figure 4.2, (ii). In this case the statistical multiplexing gain exists and both source types are S-VBR. The final case (iii) is sketched in Figure 4.2, (iii). The boundary of the admissible set is ABC, in which the segment AB is linear and coincides with the boundary of $\mathcal{A}_0$, namely, $\{\boldsymbol{K} : \sum K_j e_{0,j} = C\}$, while the segment BC is convex. The breakpoint B has coordinates $\boldsymbol{K}_s$ which the result gives:

$$K_{s,1}e_{0,1} + K_{s,2}e_{0,2} = C \tag{4.18,i}$$

$$K_{s,1}\log(1/w_1) + K_{s,2}\log(1/w_2) = \log(1/L) \tag{4.18,ii}$$

Note that for $K_1 > K_{s,1}$ the boundary segment BC lies entirely above the linear extension of AB, with the difference attributable to the statistical multiplexing gain. Of course no such gain exists for $K_1 < K_{s,1}$. In this case class 1 is S-VBR and class 2 is NS-VBR for the given node bandwidth (and, implicitly, node buffer). The breakpoint $\boldsymbol{K}_s$ is clearly important for operations and design because, first, it demarcates between two different sets of effective bandwidths, and secondly, there are obvious performance benefits in operating to the right of the breakpoint since the effective bandwidth of class 1 sources is smaller.

In the sketch of Figure 4.2, (iii) we have also shown the linear segment DE which is tangent to BC at an intermediate point. As our numerical results indicate, in many cases the departure from linearity of boundary segments represented by BC is slight, in which case DE is close to BC throughout.

4.3 Critical Value of the Token Rate r

An important observation reported in the next section is the sensitivity of the statistical multiplexing gain, $g = e_0/e$, on the token rate r. Here we present a complementary qualitative result which establishes that with all parameters except r held fixed, there exists a critical value of r, denoted by r_c, which separates token rates for which the statistical multiplexing gain exists, i.e. $g > 1$, from those for which $g = 1$. For simplicity we only consider the case of a single class.

Recall from Section 4.1 that if $(C/e_0)\log(1/w) \le \log(1/L)$ then the maximum number of sources that is admissible, $K_{\max} = C/e_0$ and hence $e = e_0$. On the other hand, if $(C/e_0)\log(1/w) > \log(1/L)$ then

$K_{\max} > C/e_0$, and $e_0 > e$. Denote

$$f(r) \triangleq \frac{C}{e_0} \log(1/w) , \tag{4.19}$$

so that $g > 1$ if and only if $f(r) > \log(1/L)$. In (4.19) let r be varied, while the other parameters, B, C, P, and B_T, are held fixed; note that e_0 and w depend on r as given in (2.12) and (2.13). As sketched in Figure 4.3, $f(r)$ is a monotonic, strictly decreasing function of r for r in $(0, B_T C/B)$. This is because both e_0 and w increase with r. Moreover, $f(r) \to \infty$ as $r \to 0+$ and $f(B_T C/B) = 0$. For $r > B_T C/B$, as

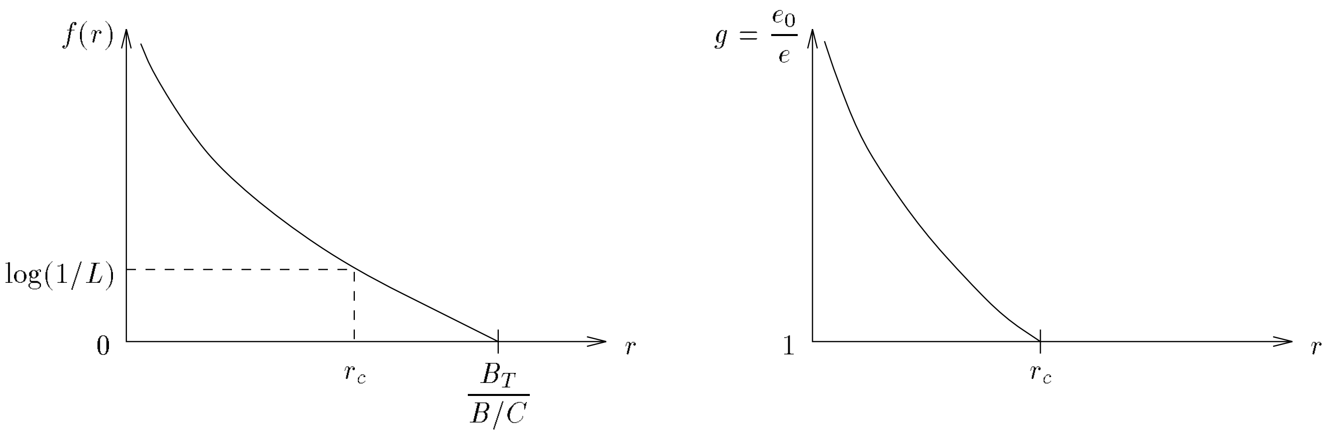

Figure 4.3: The token rate r is varied while all other parameters are held fixed in a single source class.

discussed earlier in Section 2, $e_0 = r$ and $w = 1$, and hence $f(r) = 0$. Consequently, there exists an unique solution r_c to

$$f(r) = \log(1/L) , \tag{4.20}$$

with the property that $f(r)$ is greater or less than $\log(1/L)$ depending on whether r is less or greater than r_c. Hence we have

Proposition 4.4 Let B, C and the source parameters P and B_T be fixed. The statistical multiplexing gain g exceeds 1 if and only if $r < r_c$. The critical token rate r_c is the unique solution r to the equation

$$\frac{C}{e_0} \log(1/w) = \log(1/L) , \tag{4.21}$$

where e_0 and w depend on r as given in (2.12) and (2.13). ■

It is not hard to show the monotonicity of g with respect to r for $r < r_c$, which is also depicted in Figure 4.3.

4.4 Critical Values of the Source Activity Factor w_Ω

It is also of interest to examine the case where the mean rate r and period T are held fixed, and the source activity factor w_Ω is varied. Varying w_Ω is equivalent to varying the peak rate P, since $P = r/w_\Omega$. Some algebraic manipulation of (4.21) reveals that the critical token rate r_c may be written as

$$r_c = \frac{C}{\log\left(\frac{1}{L}\right)} \left(w_\Omega + \frac{T_{\text{buf}}}{T}\right) \log\left[\left(w_\Omega + \frac{T_{\text{buf}}}{T}\right)^{-1}\right] \tag{4.22}$$

for

$$w_\Omega + \frac{T_{\text{buf}}}{T} < 1 \ . \tag{4.23}$$

When (4.23) is violated, $e_0 = r$ and $r_c = 0$.

The character of (4.22) is as follows. At $w_\Omega = 0$, r_c has the value $r_{c,0} = C/[e \log(1/L)]$. As w_Ω is varied, r_c achieves a peak value $r_{c,\max}$

$$r_{c,\max} = \begin{cases} \dfrac{C T_{\text{buf}}}{T \log(1/L)} \log\left(\dfrac{T}{T_{\text{buf}}}\right) & \text{if } \dfrac{T_{\text{buf}}}{T} < e^{-1} \\ r_{c,0} & \text{if } \dfrac{T_{\text{buf}}}{T} >= e^{-1} \ . \end{cases} \tag{4.24}$$

If $r_{c,\max} < r$ then the multiplexing gain g will be 1 for all values of w_Ω. If $r < r_{c,0}$ then there will be a critical value $w_{\Omega,c2}$ such that $g > 1$ for $w_\Omega < w_{\Omega,c2}$ and $g = 1$ for $w_\Omega < w_{\Omega,c2}$. If $T_{\text{buf}}/T < e_{\text{nat}}^{-1}$, then there is a third possibility, namely, that $r_{c,0} < r < r_{c,\max}$. In this case, there is an additional critical value, $w_\Omega, c1$, such that if $w_\Omega < w_{\Omega,c1}$, then $g = 1$.

This analysis provides a possible perspective on the ways in which NS-VBR traffic can arise: statistical multiplexing can be impossible because (i) the mean rate r is too large; or (ii) the traffic is too smooth (w_Ω is large); or (iii) there is little buffering and the peak rate P is large (w_Ω is small).

5 NUMERICAL STUDY

The purpose of this section is to numerically illustrate the mathematical results derived in prior sections and, also, to obtain qualitative results and design guidelines. One of the highlights of the results, which has considerable practical implications, is the linearity of the boundary of the admissible region and the specification of the conditions for it to hold. This result allows us to justifiably define the effective bandwidth for each traffic class independently of the others. The numerical examples in this section address the following issues: (i) the tightness of the Chernoff bound, (ii) the behavior of the effective bandwidth, e, with respect to nodal resources (bandwidth and buffer), (iii) the sensitivity of e and the multiplexing gain with respect to the mean rate, peak rate and token buffer size, (iv) the admission region when the source classes are all S-VBR, (v) the admission region when both S-VBR and NS-VBR sources are present, and (vi) the accuracy of our analytical approach to resource allocation and comparison with simulation.

We begin by considering results for a single class of sources. In Figure 5.1 we examine the effective bandwidth, $e = C/K_{\max}$ ($K_{\max}$ is integer) vs. C with B/C held fixed. Our objectives in this figure are to test the accuracy of the Chernoff bound, demonstrate the existence of C_c, and compare the values obtained numerically (for integer K) and from (4.5). Observe that for $C < 11.1$, e is equal to e_0, the effective bandwidth in lossless multiplexing; for $C > 11.1$ statistical multiplexing gain is in effect and e begins to drop. C_c obtained from (4.5) is 10.65. (The small discrepancy is due to the integrality of $K_{\max}$ in the figures.) Observe that the Chernoff bound on the loss probability results in a bound on e which is quite tight for all values of C and gets tighter as C increases; also, e approaches r for large values of C, as expected. In Figure 5.2 we consider a different class and obtain similar qualitative results; in particular, the values of C_c in the figure and that obtained from (4.5) are 13.1 and 12.6 respectively.

In contrast to Figures 5.1 and 5.2, in Figure 5.3 we consider two classes with very different values of C_c. In the figure C_c is 22.1 and 108.5 for classes 1 and 2, respectively, while (4.5) gives 20.9 and 100.7. This suggests that if the two classes are multiplexed over a link having bandwidth 45, for instance, then class 1 will exhibit statistical multiplexing gain while class 2 will not. The exact value of C_c for the multiplexed classes is examined further below.

Figure 5.4 considers the sensitivity of e with respect to r, the mean rate. The figure shows e and the multiplexing gain (e_0/e) as functions of r, for various values of B_T, with all other parameters held fixed, and strikingly exhibits the sensitivity of the multiplexing gain e_0/e with respect to r. For instance, when r increases from 0.1 to 0.2, the gain decreases from 2.7 to 1.54.

We next examine the multiplexing of multiple classes of sources. Figure 5.5 considers two classes, which were individually considered in Figures 5.1 and 5.2. The boundary of the admissible region is plotted for different value of loss probability L. For each L we plot the boundary obtained using Chernoff bound and its linear conservative approximation obtained as the tangent at $K_j = K_{\max,j}/2$, $j = 1, 2$. There are two important observations to make: (1) Statistical multiplexing, even with very small loss probability, provides considerable gain in network capacity (number of connections). For instance, more than 3-fold increase in the maximum number of sources of each class is obtained as L is increased from 0 to 10^{-7}, (2) The boundary of the admissible region is very close to its linear conservative approximation for all values of L; it is also very close to the optimistic linear approximation defined in (3.20). The linearity of the admissible region is due to the high degree of sharing allowed among all sources and can be best explained by examining Figure 5.5 in conjunction with Figures 5.1 and 5.2. We first note that the value of C_c for each class ($C_{c,1} = 11.1$ and $C_{c,2} = 13.1$) is less than the link capacity $C = 45$, hence the two classes are S-VBR. Second, and more interestingly, e_1, and e_2, obtained from the linearized boundary, correspond to the values read off from Figures 5.1 and 5.2 respectively. The linearity of the admissible region is further tested in Figures 5.6 for different values of C with B/C and L held fixed.

In Figure 5.7 the effect of varying the peak rate is examined. The boundary of the admissible region is plotted for different values of P_1 and P_2 (the peak rates of class 1 and class 2 respectively) while all other parameters are held fixed. Note that increasing the peak rate reduces the admissible region, but does not alter the linearity of its boundary.

In Figures 5.8, we consider three classes of S-VBR sources and plot the admissible region in the space of K_1 and K_2 for different fixed values of K_3. Again, the linearity of the boundary persists.

Next, we consider the multiplexing of two source classes, where class 1 is S-VBR and class 2 is NS-VBR, i.e., $C_{c,1} < C < C_{c,2}$. As discussed in Section 4.2 such mixture of VBR classes results in a breakpoint in the boundary of the admissible region. Figure 5.9 demonstrates this result. Class 1 and class 2 are the same classes considered individually is Figure 5.3. The breakpoint in the figure is at $K_1 = 10$ and $K_2 = 8$ which closely matches that given by equation (4.18) where $K_1 = 10.517$ and $K_2 = 7.78$.

The analysis leading to the preceding numerical results involved a number of (generally conservative) approximations which are reviewed in Section 6. The accuracy of these approximations may be assessed by comparing the number of sources predicted by the analysis with that predicted by simulation. We have done some preliminary work in this regard. We consider a node with a single class of sources, where the source parameters, the link buffer size and the link bandwidth are as given in Figure 5.2. Table 5.1 gives

Loss Probability	2.2×10^{-8}	1.7×10^{-7}	1.3×10^{-6}	6.0×10^{-6}	2.8×10^{-5}
K (Simulation)	150	160	170	180	190
K (Analysis)	110	119	129	137	147

Table 5.1: Fluid-model simulation vs. analysis for $r = 0.15$ Mbps $P = 1.5$ Mbps and $Q = 250$ cells.

the number of admissible sources found from analysis and from a fluid-model simulation for different loss probabilities. In the simulation, each periodic on-off source generates fluid at a constant rate during the on period, and the fluid is drained from the buffer at a constant rate. The table is obtained by determining the loss probability from the simulation for each value of K, the number of admissible sources, and then determining the corresponding number of sources for that loss probability from analysis. Note that the number of admissible sources obtained from the analysis is a lower bound on the number found from the simulation.

6 Discussion

Our analysis has been based on a number of approximations. Most of the approximations used have been conservative. Each of these approximations essentially replaces the number of sources that can be admitted for a given loss probability with an upper bound on the admissible number of sources. Conservative approximations include (A) use of the Chernoff approximation, which we know to be quite accurate; (B) the use of segregated bandwidth allocations for each circuit so that inequality (2.6) is taken to be an equality; (C) the optimization of each circuit buffer and bandwidth allocation independent of the nature of the traffic sources for other circuits, leading to (2.8); (D) the replacement of the virtual buffer usage $v_{ji}(t)$ by an on-off process, which in conjunction with approximations (B) and (C) reduces the problem to one involving a single resource; (E) the use of worst-case leaky-bucket output rather than more typical traffic patterns, albeit it is difficult to know what usage patterns will be typical in ATM networks; and optionally (F) the approximation of the concave admissible set boundary by a tangent plane, leading to (3.18). Optimistic approximations include (G) the use of the fluid-model as a proxy for a system with discrete traffic; and optionally (H) the approximation of the concave admissible set boundary by a plane determined by single-class effective bandwidths, leading to (3.20).

The fluid-model simulation results in Section 5 are suggestive with regard to the magnitude the error introduced by approximations (A), (B), and (D). In the examples considered, the analytic results admit 22–27 percent fewer sources than the fluid simulation would allow. Of course, much more work would need to be done before any broad generalizations could be made regarding the accuracy of these approximations. Similarly, work would be needed to evaluate the accuracy of (C); this could be studied via simulations of multiple source classes.

Regarding (G), the use of the fluid model in lieu of a discrete model is optimistic because it ignores the possible need to buffer cells during some portion of the cell interarrival time of each source. If the number of cells involved in a burst is large and the buffering required in the fluid model is substantial, then the fluid-model predictions are expected to be very close to those produced by a discrete model. In cases where the cell interarrival times for individual sources are large and the number of sources in a burst is small, the predictions of the fluid model may be significantly optimistic. While optimism is dangerous in an admission control criterion, (i) this optimism may be counterbalanced by the conservatism of other approximations, and (ii) we anticipate that relatively simple methods of accounting for errors due to the fluid approximation will be forthcoming.

The numerical results of Section 5 suggest that approximations (F) and (H), linear approximations to the boundary of the admissible set, should be fairly accurate as long as S-VBR and NS-VBR traffic is not intermixed. The potential for using the optimistic linear approximation, (H), is tremendously important for practical on-line admission control, insofar as it allows one to incrementally tally the total effective bandwidth used in a buffer/trunk system as sources are added and dropped, without any need to take into account future traffic or recalculate the effective bandwidths of previously added traffic. How to attain this sort of efficiency in cases where S-VBR and NS-VBR traffic are intermixed will be the subject of future work.

Throughout this work we have focused on a single buffer/trunk system. It is important to examine the effect of buffering and multiplexing on the characteristics of the source traffic as it traverses the network. While the total traffic multiplexed into a node gets smoothed out due to buffering, and hence its effective bandwidth decreases, some sources experience increase and others experience decrease in their effective bandwidth. The traffic burstiness of an individual source increases if the traffic passes through the multiplexing node when its buffer is nonempty and few other sources are on. In this case, the peak rate of the output traffic increases, and can reach the link bandwidth if the traffic arrives to nonempty buffer while all other sources are off. This increase in burstiness will be of concern in subsequent downstream nodes if the source follows a downstream path which is shared by only few other sources. However, when the source peak rate is small compared to the link rate, i.e., the number of sources multiplexed is large, the probability that only few sources are on is small. In [13] simulation studies using source models with exponentially-distributed on and off periods confirm that the the change in the characteristics of the source traffic due to buffering and multiplexing in a network node is negligible provided that the source peak rate is a small fraction of the link bandwidth, and, thus, network analysis based on nodal decomposition is justified.

7 Acknowledgement

We gratefully acknowledge the benefit of discussions and support of our colleague Flavio Bonomi who saw early the need for a new approach to resource allocation and admission control to ATM networks.

References

[1] D. Anick, D. Mitra and M. M. Sondhi, "Stochastic theory of a data handling system with multiple sources", *Bell Syst. Tech. J.*, **61**, pp. 1871–1894, 1982.

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

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

[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, pp. 591-600, 1994.

[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**, pp. 29–63, 1991.

[6] A. I. Elwalid and D. Mitra, "Effective bandwidth of general Markovian traffic sources and admission control of high speed networks", *IEEE/ACM Trans. Networking*, **1**, No. 3, pp. 329–343, 1993.

[7] R. Guerin, H. Ahmadi and M. Nagshineh, "Equivalent capacity and its application to bandwidth allocation in high-speed networks", *IEEE J. Select. Areas Commun.*, **9**, pp. 968–981, 1991.

[8] R. J. Gibbens and P. J. Hunt, "Effective bandwidths for the multi-type UAS channel", *Queueing Systems*, **9**, pp. 17–28, 1991.

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

[10] J. Y. Hui, *Switching and Traffic Theory for Integrated Broadband Networks*, (Kluwer, Boston), 1990.

[11] F. P. Kelly, "Effective bandwidths at multi-type queues", *Queueing Systems*, **9**, pp. 5–15, 1991.

[12] 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.

[13] W. C. Lau and S. Q. Li, "Traffic analysis in large-scale high-speed integrated networks: validation of nodal decomposition approach," *IEEE INFOCOM*, 1993, pp. 1320–1329.

[14] D. Mitra and J. A. Morrison, "Erlang capacity and uniform approximations for shared unbuffered resources", to appear in *IEEE/ACM Trans. Networking.* A condensed version of this paper is in *Proc. ITC-14*, pp. 875-885, 1994.

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

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

[17] V. V. Petrov, "On the probabilities of large deviations for sums of independent random variables", *Theory of Prob. and its Applications*, Vol. X, No. 2, pp. 287–298, 1965.

[18] J. W. Roberts, ed., *Performance Evaluation and Design of Multiservice Networks*, Final Report of the COST 224 project, Commission of the European Communities, 1992.

[19] 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 and Y. Takahashi, Elsevier, 1993, pp. 243-262.

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

[21] W. Whitt, “Tail probabilities with statistical multiplexing and effective bandwidths for multi-class queues”, *Telecommun. Syst.*, Vol. 2, 1993, pp. 71-107.

[22] 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.

[23] 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.*, Vol. E75-B, No. 2, 1992, pp. 82-86.

[24] D. Yates, J. Kurose, D. Towsley and M. Hluchyj, “On per-session end-to-end delay and call admission problem for real-time applications with QOS requirements”, SIGCOMM ‘93, 1993.