9-1-1991

Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks

Roch A. Guérin
University of Pennsylvania, guerin@acm.org

Hamid Ahmadi
IBM Research Division

Mahmoud Naghshineh
IBM Research Division

Recommended Citation

Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks

Abstract

High-speed fast-packet-switched network architectures are capable of supporting a wide range of connections with different bandwidth requirements and traffic characteristics. While this environment provides increased flexibility in supporting various services, its dynamic nature poses difficult traffic control problems when trying to achieve efficient use of network resources. One such a problem is the issue of bandwidth management and allocation. Because of the statistical multiplexing of all connections at the physical layer and the variations of connections bit rate, it is important to characterize, for a given Grade-Of-Service (GOS), both the effective bandwidth requirement of a single connection and the aggregate bandwidth usage of multiplexed connections. In this paper, we propose a computationally simple approximate expression for the "*equivalent capacity*" or bandwidth requirement of both individual and multiplexed connections, based on their statistical characteristics and the desired GOS. The purpose of such an expression is to provide a unified metric to represent the effective bandwidth used by connections and the corresponding effective load of network links. These link metrics can then be used for efficient bandwidth management, routing, and call control procedures aimed at optimizing network usage. While the methodology proposed in the paper can provide an exact approach to the computation of the equivalent capacity, the associated complexity makes it infeasible for real-time network traffic control applications, hence, an approximation is required. The validity of the approximation developed in the paper is verified by comparison to both exact computations and simulation results.

Keywords

Equivalent Bandwidth, Call admission, network

Disciplines

Digital Communications and Networking

Comments

Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks

Roch Guérin, *Senior Member, IEEE,* Hamid Ahmadi, and Mahmoud Naghshineh

Abstract—High-speed fast-packet-switched network architectures are capable of supporting a wide range of connections with different bandwidth requirements and traffic characteristics. While this environment provides increased flexibility in supporting various services, its dynamic nature poses difficult traffic control problems when trying to achieve efficient use of network resources. One such problem is the issue of bandwidth management and allocation. Because of the statistical multiplexing of all connections at the physical layer and the variations of connections bit rate, it is important to characterize, for a given grade of service (GOS), both the effective bandwidth requirement of a single connection and the aggregate bandwidth usage of multiplexed connections. In this paper, we propose a computationally simple approximate expression for the *"equivalent capacity"* or bandwidth requirement of both individual and multiplexed connections, based on their statistical characteristics and the desired GOS. The purpose of such an expression is to provide a unified metric to represent the effective bandwidth used by connections and the corresponding effective load on network links. These link metrics can then be used for efficient bandwidth management, routing, and call control procedures aimed at optimizing network usage. While the methodology proposed in this paper can provide an exact approach to the computation of the equivalent capacity, the associated complexity makes it infeasible for real-time network traffic control applications. Hence, an approximation is required. The validity of the approximation developed in this paper is verified by comparison to both exact computations and simulation results.

I. INTRODUCTION

IN high-speed packet-switched network architectures such as ATM [1] and PARIS [6], several classes of traffic streams with widely varying traffic characteristics are statistically multiplexed and share common switching and transmission resources. Because of the potentially dramatic differences in the statistical behavior of connections, the problems of bandwidth management and traffic control pose new challenges and create difficulties very different from the ones present in traditional packet- or circuit-switched networks. Because all connections are statistically multiplexed at the physical layer and the bit rate of connections varies, a challenging problem is to characterize, as a function of the desired grade of service

(GOS), the effective bandwidth requirement of both individual connections and the aggregate bandwidth usage of connections multiplexed on a given link. This problem is the main focus of this paper.

The basic objective of a bandwidth management and traffic control strategy is to allow for high utilization of network resources, while sustaining an acceptable grade of service for all connections. Several network traffic control functions, such as flow and congestion control and routing and call admission policies, depend on the characterization of the effective bandwidth of individual connections and the resulting load on network links. In particular, this information is needed by the network to decide if and how to accept incoming connection requests. Routes are typically selected and calls accepted so as to "optimize" some measure of resources utilization while providing adequate GOS to the carried traffic. This requires knowledge of both the current traffic conditions and the impact of adding a new connection.

This information is provided by accounting (on each link) for the amount of bandwidth currently allocated to accommodate existing connections, and by identifying how much additional bandwidth needs to be *reserved* on links over which a new connection is to be routed. Because of the statistical multiplexing of connections and shared buffering points in the network, both the accounting and reservation are based on some aggregate statistical measures matching the overall traffic demand rather than on physically dedicated bandwidth or buffer space per connection. In addition to the inherent complexity of such a matching, another major challenge is to provide these traffic control functions in *real-time*, upon the arrival of a connection request. The corresponding procedures must, therefore, be computationally simple enough so their overall complexity is consistent with real-time requirements.

The problem of bandwidth allocation in an ATM environment has been addressed in a number of very interesting papers [8], [9], [11], [13], [15], [21], [22], [26]. While the approaches taken in these papers may differ, they are all either restricted to a limited range of connection characteristics [11], [15], [22] or of limited flexibility [8], [9], [13], [21], [26] since they essentially rely on sets of curves, obtained by analysis or simulation, which are to be used as guidelines to determine the equivalent capacity of connections. The major drawback of the latter approach is, that in addition to requiring storage of the

precomputed curves at the bandwidth management control points, the static nature of the information may not accurately reflect the dynamic and changing nature of real-time network traffic conditions and connections characteristics.

In this paper, we propose a computationally simple approximation for the *equivalent capacity* or bandwidth requirement of a single or multiplexed connections on the basis of their statistical characteristics. When connections are statistically multiplexed, their aggregate statistical behavior differs from their individual statistical representation. One needs, therefore, to define new metrics to represent the effective bandwidth requirement of an individual connection as well as the total effective bandwidth requirement of connections multiplexed on each link. The purpose of the equivalent capacity expression is to provide a unified metric to represent the effective bandwidth of a connection as well as the effective aggregated load on network links at any given time. These link metrics can then be used for efficient bandwidth management, routing, and call control procedures.

The equivalent bandwidth is computed from the combination of two different approaches, one based on a fluid-flow model [3], [17], [20], [23] and the other on an approximation of the stationary bit rate distribution. These two approaches have been selected as they complement each other, capturing different aspects of the behavior of multiplexed connections, while remaining computationally simple. This aspect is of significance when real-time decisions have to be made on "if" and "how" to handle new incoming connections. Because of its simplicity, the approximation developed in this paper allows for the real-time tracking of bandwidth allocations on network links. It should be emphasized that the approach taken in the paper, rather than trying to apply exact but intractable models that do not necessarily capture all the impact of complex network interactions, has been to develop a simple and yet reasonably accurate metric to compare load levels on network links.

The goal is to capture the key connection parameters that influence bandwidth allocation. Therefore, the equivalent capacity computation focuses on the bandwidth requirement of the bit rate generated by sources[1], and not on the different interactions that take place within the network. As mentioned, such interactions are often too complex to even be accurately described. Furthermore, focusing directly on the bit rate requirement of a source is reasonable in the context of a high-speed network attempting to provide "transparent" service to sources. The methodology proposed in this paper can provide an exact computational procedure for the equivalent capacity, but the complexity of an exact computation is not practical for real-time network traffic control applications. Hence, an accurate and computationally simple approximation is developed. The accuracy of the final approximation is ex-

amined by comparing it to both exact computations and simulation results, and is found to be acceptable across the range of possible connection characteristics.

The organization of the paper is as follows. Section II presents the model used to characterize a connection. In Section III, we describe the equivalent capacity model and its derivation. Section IV presents numerical results illustrating the use of the equivalent capacity model, while Section V is devoted to discussions of some limiting behaviors and extensions of the equivalent capacity expression. Finally, Section VI summarizes the findings of the paper and outlines possible extensions.

II. TRAFFIC MODEL FOR A SINGLE CONNECTION

A single connection has a variable bit rate bounded by the maximum bit rate of its physical attachment. In order to characterize the effective bit rate or equivalent capacity of a connection, we need to select an appropriate model to specify its characteristics in terms of known parameters or metrics. In this paper, we adopt a two-state fluid-flow model that captures the basic behavior of the data source associated with a connection. The rationale for such a model is that a source is either in an "idle state," transmitting at zero bit rate, or in a "burst state" and transmitting at its peak rate. Such a source model has the advantage of being both simple and flexible as it can be used to either represent connections ranging from bursty to continuous bit streams (e.g., [9]) or approximate more complex sources (e.g., [19], [21]).

Based on this two-state fluid-flow model, we define idle and burst periods to be the times during which the source is idle or active, respectively. The peak rate of a connection R_{peak} and distributions of idle and burst periods completely[2] identify the traffic statistics of a connection. Assuming the parameters of a connection are stationary, its peak rate R_{peak} and utilization ρ, i.e., fraction of time the source is active, completely identify other quantities of interest such as mean m and variance σ^2 of the bit rate. For exponentially distributed burst and idle periods, the source is furthermore completely characterized by only three parameters, namely R_{peak}, ρ, and b, where b is the mean of the burst period. The mean burst period b gives some indications on how data is being generated by the source. Two sources, with identical mean and peak bit rates but different burst periods, have different impacts on the network [10], [13], [21], [26]. It is our belief that the connection metric vector (R_{peak}, ρ, b) represents the most significant aspects of a source behavior.

The model described above can be extended to "handle" sources with nonexponential burst and idle periods, through the use of simple approximation techniques. These techniques rely essentially on standard moment matching approximations and are discussed in Section V-A, where a number of illustrative examples are also presented.

---

[1] The types of "sources" handled by the method are, however, quite general. They include both individual users as well as more complex sources, such as the output of a multiplexer.

[2] Assuming i.i.d. burst and idle periods.

III. Equivalent Capacity and Link Metric

The equivalent capacity of a set of connections multiplexed on a link is defined as the amount of bandwidth required to achieve a desired GOS, e.g., buffer overflow probability, given the offered aggregate bit rate generated by the connections. It is a function of individual connection characteristics and available network resources such as buffers. Note that the GOS criterion could also be expressed in terms of delay rather than loss. However, if several different GOS (be it delay or loss) are to be supported by the network, it is necessary that either links be dimensioned to meet the most stringent GOS constraint or that additional mechanisms be provided to distinguish between connections. For example, different delay or jitter constraints can possibly be supported through dedicated priority and/or buffering schemes (see [7] for an example). In this paper, loss or buffer overflow probability was selected as the common measure of link loads. As mentioned above, other constraints, e.g., jitter, can typically be handled through other link-level mechanisms.

The equivalent capacity, therefore, determines the bandwidth allocation levels on network links. It provides a unified metric for link loads, which can then be used by network control functions such as routing and congestion control. Because of the "real-time" requirements of these functions, it is critical that the complexity of the equivalent capacity computation be kept as low as possible while still accounting for connections characteristics, existing network traffic, and desired GOS. Exact solutions are either intractable or, when available, incompatible with real-time requirements. As the goal is to provide a simple while still reasonably accurate metric to measure and compare link loads, approximations must therefore be developed.

The approach taken combines two different approximations. The first one relies on a fluid-flow model that builds on the source representation introduced in the previous section. The second approximation focuses on the distribution of the stationary bit rate on a link. As mentioned earlier, the two approximations were selected both for their computational simplicity and because they capture different aspects of connection behavior. The first approximation accurately estimates the equivalent capacity when the impact of individual connection characteristics is critical. The second approximation is representative of bandwidth requirements when the effect of statistical multiplexing is of significance. As both aspects are typically exclusive, the two approximations complement each other and can be combined to predict relatively accurately the equivalent capacity of connections.

A. Fluid-Flow Approximation

The first approximation for the equivalent capacity is based on a fluid-flow model [3], [17], [20], [23]. In such a model, the bit rate generated by a number of multiplexed connections is represented as a continuous flow of bits with intensity varying according to the state of an underlying continuous-time Markov chain. This Markov chain is obtained from the superposition of the sources associated with each connection. For example, in the case of multiplexing N identical two-state Markov sources (each capable of a peak bit rate of R_{peak}), the intensity of the aggregate bit rate is jR_{peak} when j sources are active. The aggregate bit rate is offered to a buffer which is emptied at a constant rate of c. We are interested in determining the smallest value $\hat{C}$ of c that, for a given buffer size x, ensures a buffer overflow probability (GOS) smaller than ϵ. The value $\hat{C}$ is the equivalent capacity of the multiplexed connections.

The determination of the equivalent capacity $\hat{C}$ requires that we first obtain an expression giving the distribution of the buffer contents as a function of the connections characteristics and the service rate. This expression must then be inverted to determine the value of the service rate, which ensures an overflow probability smaller than ϵ for the available buffer size x. This value is the equivalent capacity that should be allocated to the connections. The distribution of the buffer content can be derived using available methods, i.e., [3] for identical sources and [17], [20], [23] for more general input Markov chains.[3] Although these methods provide the necessary information to determine the equivalent capacity, the associated computational complexity is often not compatible with the real-time requirements mentioned earlier. This is because, even when the buffer content distribution can be explicitly derived, the resulting expression cannot be easily inverted to yield the equivalent capacity as a function of other known parameters. Iterative numerical procedures must then be used to determine the value of the equivalent capacity (see Section III-A-1 and A-2 for additional details, and [21] for an example). In order to overcome this problem, approximations must be used and are now developed in the rest of this section.

1) Single Source: We first address the case of a single two-state Markov source, where we wish to determine the capacity or bandwidth to allocate to the associated connection in isolation. Using the notation introduced earlier, a two-state Markov source is characterized by its peak rate R_{peak}, utilization ρ, and mean burst period b. The distribution of the buffer contents, when such a source is feeding a buffer served by a constant rate server, can be derived using standard techniques for either infinite or finite buffer systems.[4] From this distribution, it is then possible to determine the service rate, or equivalent capacity $\hat{c}$, needed to achieve a given GOS. The GOS, as discussed earlier, is the buffer overflow probability or, more specifically, the buffer overflow probability when the source is active and transmitting. The use of such a conditional overflow probability is required to ensure that sources with

---

[3]Note that the techniques available from [17], [20], [23] are not restricted to two-state sources and therefore would allow the use of more general source models, i.e., hyperexponential, phase-type, etc. (see [19], [21] for examples).

[4]In the case of infinite buffers, the probability that the queue length exceeds the buffer size x is an upper bound for the overflow probability.

very low utilizations are not overly penalized (the capacity allocated to sources with utilizations below ϵ would otherwise be zero), but has little impact for reasonable source utilizations, i.e., above 10%. Assuming a finite buffer of size x, the equation satisfied by the equivalent capacity for an overflow probability of ϵ (the desired GOS) is then easily found to be of the form:

$$\epsilon = \beta \times \exp\left(-\frac{x(c - \rho R_{peak})}{b(1 - \rho)(R_{peak} - c)c}\right)$$

where

$$\beta = \frac{(c - \rho R_{peak}) + \epsilon\rho(R_{peak} - c)}{(1 - \rho)c}. \tag{1}$$

Note that an infinite buffer system satisfies the same equation, only with a different value for β.

The equivalent capacity for a single source can then be obtained by solving for c in (1). It should be noted that, even for the simple single-source case, no explicit expression is available for the equivalent capacity, and (1) must be solved numerically. A natural simplification is, however, available, as the term β can be shown to be typically close to 1 (in fact, always smaller). Approximating β by 1 in (1) provides an explicit upper bound for $\hat{c}$ which, furthermore, is close to the exact value (extensive numerical experiments never produced a relative error above 10%). The equivalent capacity associated with a single connection in isolation is then taken to be:

The distribution $F(x)$ is completely determined from the values of the associated eigenvalues, eigenvectors, and corresponding coefficients. There are typically[7] no explicit expressions for these quantities, which must then be determined numerically.[8] The equivalent capacity corresponding to a set of multiplexed sources can then be obtained again using iterative numerical techniques, where at each iteration a new solution to (3) must be computed. Such a procedure, although exact, is unfortunately not compatible with a dynamic and real-time environment. Approximations must be developed. As mentioned earlier, the final expression is tested against exact computations and simulation results, with which it is found to be in reasonably good agreement.

A first natural step is to consider an asymptotic approximation. The system is assumed to have an "infinite" buffer with a probability of buffer overflow approximately equal to the probability that the buffer content exceeds the actual buffer size. For large buffers, this quantity can be further approximated by only considering the contribution of the term corresponding to the largest negative, or dominant eigenvalue in (3) (other terms become negligible as the buffer size increases.) Under such assumptions, the buffer overflow probability $G(x)$ is known [3], [17] to be of the form:

$$G(x) \simeq \beta e^{z_0 x} \tag{4}$$

$$\hat{c} \simeq \frac{\alpha b(1 - \rho)R_{peak} - x + \sqrt{[\alpha b(1 - \rho)R_{peak} - x]^2 + 4x\alpha b\rho(1 - \rho)R_{peak}}}{2\alpha b(1 - \rho)} \tag{2}$$

where $\alpha = \ln(1/\epsilon)$. Note, that in the case of a continuous bit stream connection, we have $\rho = 1$ and $b = \infty$, and taking limits in (2) yields the expected result $\hat{c} = R_{peak}$.

2) Multiple Sources: In the case of multiple superposed sources, the distribution of the content of a buffer, fed by the aggregate bit rate and served by a constant rate[5] server, can also be determined. The approach is more complex than for single sources, and essentially numerical in spite of the existence of powerful simplification techniques, i.e., the "decomposition methods" of [17], [20], [23]. In general, when the input bit rate is characterized by an N-state Markov chain, the distribution of the buffer contents is of the form:

$$F(x) = \sum_{i=1}^{N} a_i \Phi_i e^{z_i x} \tag{3}$$

where the z_i and Φ_i are, respectively, generalized eigenvalues and eigenvectors associated with the solution of the differential equation satisfied by the stationary probabilities of the system, and the a_i's are coefficients determined from boundary conditions[6] (see [3], [17], [20], [23] for details).

where z_O is the largest negative eigenvalue and β is a constant term independent of x.

As mentioned in [17], z_O can usually be obtained from relatively simple expressions, while the determination of the constant β is typically more difficult. We shall illustrate this with the example of N identical two-state Markov sources, for which explicit expressions are actually available [3]. This example will furthermore provide us with some indications to further simplify (4) so as to obtain a simple, explicit approximate expression for the equivalent capacity of multiple connections. In the case of N identical two-state Markov sources, we have:

$$G(x) \simeq \left(\frac{N\rho R_{peak}}{C}\right)^N \left[\prod_{i=1}^{N-[C/R_{peak}]-1} \frac{z_i}{z_i + z_O}\right]$$
$$\cdot \exp\left(-\frac{N(C - N\rho R_{peak})}{b(1 - \rho)(NR_{peak} - C)C}x\right) \tag{5}$$

where the z_i are all negative eigenvalues and the largest one, z_O, is explicitly given in the exponential term, C is

---

[5] [20] and [23] actually provide techniques that are applicable to systems where both the service and the input rates are modulated by Markov chains.

[6] For example, the coefficients associated with eigenvalues with positive real part are equal to zero for stable infinite buffer systems.

[7] One notable exception is the case of identical two-state Markov sources [3].

[8] The previously mentioned decomposition techniques partition the problem into a number of lower dimensions ones, therefore reducing its overall complexity.

the capacity allocated to the N connections, and $[C/R_{peak}]$ is the largest integer smaller than C/R_{peak}.

As can be seen from (5), the expression for β is quite complex and, in fact, involves not only the largest eigenvalue but all the other negative eigenvalues as well. In addition, since the value of the allocated capacity appears as a variable determining the number of coefficients in the product term present in β, general, simple approximations for β seem difficult to derive (see [17] for similar remarks and suggestions on the use of simulations to determine β). In contrast, the expression for z_O is a relatively simple function of the unknown allocated capacity. A natural approach is then to approximate β, at least over a range of connection characteristics,[9] by a constant providing a reasonable estimate.

A possible choice, consistent with the single-source approximation, is to again approximate β by 1. Although the actual value of β is unfortunately not always close to 1, there exists a range of connection characteristics for which this assumption is reasonable. For example, in the case of (5), β is the product of two factors. The first one is the Nth power of a term smaller than 1, and the second is itself the product of about N terms all greater than 1. One can then expect, that at least for some range, these two terms will compensate for each other and yield a product of magnitude reasonably close to 1. It can, in fact, be verified numerically that this is approximately the case when either the number of connections is small or the actual total equivalent capacity is reasonably close to the overall mean rate. In other cases, approximating β by 1 can, however, result in a gross overestimate[10] of the equivalent capacity. We nevertheless decide to approximate β by 1 for the fluid-flow model, and develop a second approximation, based on the distribution of the stationary bit rate, for cases where the assumption $\beta \simeq 1$ clearly does not hold.

Once β has been set equal to 1, a simple explicit expression can then be obtained for the equivalent capacity as a function of individual source characteristics. Letting $\beta = 1$ in (5) and requiring a buffer overflow probability of ϵ gives:

$$\epsilon = e^{z_0 x} \Rightarrow z_O = -\frac{\alpha}{x} \tag{6}$$

where again $\alpha = \ln 1/\epsilon$.

In the case of two-state Markov sources (not necessarily identical), there exists a simple relation between the dominant eigenvalue z_O and the allocated capacity (see (6.3) of [17]). It is given by:

where $1/\lambda_i$ and $1/\mu_i$ are the mean idle and burst periods of source i, respectively, and $R_{peak}^{(i)}$ is the corresponding peak rate. Combining (6) and (7) to obtain the value of the equivalent capacity $\hat{C}_{(F)}$ given by the flow approximation for N multiplexed sources, we get:

$$\hat{C}_{(F)} = \sum_{i=1}^{N} \hat{c}_i \tag{8}$$

where the $\hat{c}_i$ are determined from (2) and we have used the relations:

$$\mu_i = \frac{1}{b_i}, \qquad \lambda_i = \frac{\rho_i}{b_i(1 - \rho_i)}$$

and

$$\lambda_i + \mu_i = \frac{1}{b_i(1 - \rho_i)}.$$

The major advantages of (8) are its computational simplicity and its explicit dependency on source parameters. However, although the linearity of (8) certainly simplifies the accounting of how bandwidth is allocated to connections, it also clearly identifies the limitations of the expression. The linearity explicitly indicates that the simplifying assumption $\beta \simeq 1$ amounts to ignoring the effect of statistical multiplexing. In particular, unless the equivalent capacities of individual connections are themselves close to their mean bit rates, their sum is typically an overestimate of their equivalent capacity. Another approximation is, therefore, needed to accurately determine the required bandwidth allocation for cases in which statistical multiplexing is significant. This is the purpose of the next section.

B. Stationary Approximation

The base for an approximation, when the effect of statistical multiplexing is the dominant factor, can be obtained by studying the impact of the assumption $\beta \simeq 1$. Turning again to the case of the N identical two-state Markov sources described in (5), we recall earlier observations on the behavior of β. We noted that β is significantly different from 1 when a number of connections with equivalent capacity much larger than their mean bit rates are multiplexed. This is essentially the case for connections with long burst periods and relatively low utilizations. This aspect is illustrated in Figs. 1 and 2, which plot the value of β for different mean burst periods and different numbers of multiplexed sources. Fig. 1 consid-

$$C = \frac{\sum_{i=1}^{N} [(\lambda_i + \mu_i + R_{peak}^{(i)} z_O) - \sqrt{(\lambda_i + \mu_i + R_{peak}^{(i)} z_O)^2 - 4\lambda_i R_{peak}^{(i)} z_O}]}{2z_O} \tag{7}$$

---

[9]See the next paragraph and Section III-B for additional information.

[10]In the case of (5), it can be seen that as C becomes substantially larger than the total mean bit rate $N\rho R_{peak}$, both the first and second terms in β decrease. They, therefore, clearly cannot compensate for each other.

ers the case of sources with a 10% utilization while, in Fig. 2, sources are 50% utilized.
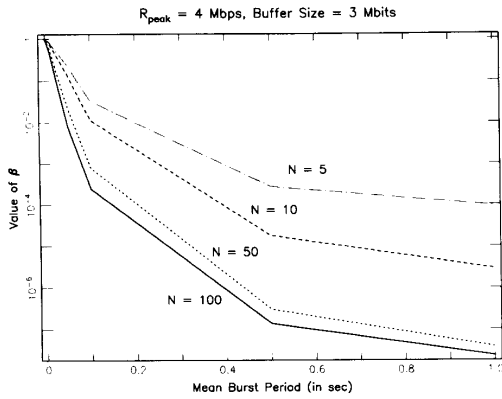
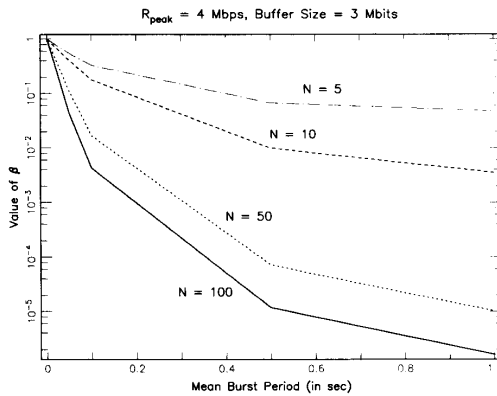Fig. 1. Value of β for sources with 10% load.

Fig. 2. Value of β for sources with 50% load.

The two figures clearly illustrate the impact of the different parameters, and in particular point to the importance of the mean burst period, and the number of sources on the value of β. In addition, in the case of long[11] burst periods, the relation between buffer overflows and packet losses may be inaccurate.[12] Similarly, the asymptotic (large buffer) approximation itself is likely to be inaccurate as well in the case of large burst periods. The significance of such inaccuracies must be tempered by the fact that even an exact model does not provide a correct measure of the loss probability seen by connections, as it cannot fully capture the impact of interactions within the network. However, this points to a limitation of the fluid-flow model and the assumption $\beta = 1$, in the case of many connections[13] with long burst periods. Fortunately, when a number of connections with relatively long burst periods are multiplexed, a reasonably accurate es-

---

[11]i.e., the amount of data generated in an average burst is not small compared to the buffer size.

[12]See [18] for some related discussions in an ATM environment.

[13]In the case of a single connection, (2) rapidly converges to the peak rate.

timate of the required bandwidth can be obtained from the stationary bit rate distribution.

More specifically, the value $\hat{C}_{(S)}$ of the equivalent capacity can be selected to ensure that the aggregate stationary bit rate exceeds $\hat{C}_{(S)}$ only with a probability smaller than ϵ, the desired buffer overflow probability. This clearly ensures a buffer overflow probability below ϵ, but is often a substantial overestimate of the actual bandwidth required as it ignores the "smoothing" effect of the buffer, i.e., the buffer allows the input rate to exceed the output rate for "short" periods. In the case of connections with long burst periods, it can, however, be argued that because the aggregate bit rate remains at the same value for relatively long periods, overload situations (when present) are likely to last sufficiently long to result in buffer overflow. It is then reasonable to allocate enough bandwidth to make the probability of an overload condition equal to the desired buffer overflow probability. For a set of multiplexed connections, the value of the equivalent capacity which satisfies such a requirement on the aggregate stationary bit rate is, therefore, the smallest value $\hat{C}_{(S)}$ which satisfies:

$$Pr(B > \hat{C}_{(S)}) \leq \epsilon \qquad (9)$$

where B is the aggregate bit rate and ϵ the desired buffer overflow probability. This amounts to imposing that the frequency of "overload periods" be less than ϵ.

The distribution of B can be determined from the stationary distribution of the underlying Markov chain formed by the superposition of sources. For example, in the case of identical two-state Markov sources, the probability p_k that k out of N sources are active is given by a binomial distribution.

$$p_k = \binom{N}{k} \rho^k (1 - \rho)^{N-k}.$$

The value of $\hat{C}_{(S)}$ is then obtained by computing the smallest integer k' such that:

$$\sum_{k=k'+1}^{N} p_k \leq \epsilon.$$

The stationary approximation then gives $\hat{C}_{(S)} = k' R_{peak}$.

In the above example, the determination of k', and therefore $\hat{C}_{(S)}$, is possible, but in general (e.g., nonhomogeneous sources) the computation of $\hat{C}_{(S)}$ can be difficult. However, in most cases[14] where the effect of statistical multiplexing is of significance, the distribution of the stationary bit rate can be rather accurately approximated by a Gaussian distribution [11], [15], [22].

The assumption of a Gaussian distribution allows us to use standard approximations to estimate the tail of the bit rate distribution. In particular, we want to determine the value $\hat{C}_{(S)}$ so that the cumulative tail probability beyond $\hat{C}_{(S)}$ does not exceed ϵ. The value of $\hat{C}_{(S)}$ can then be ob-

---

[14]As we shall briefly discuss below, exceptions can be handled by some simple checks.

tained from approximations for the inverse of the Gaussian distribution. A good approximation is given by:

$$\hat{C}_{(S)} \simeq m + \alpha'\sigma, \quad \text{with } \alpha' = \sqrt{-2 \ln(\epsilon) - \ln(2\pi)}$$

$$(10)$$

where m is the mean aggregate bit rate $(m = \Sigma_{i=1}^{N} m_i)$, and σ is the standard deviation of the aggregate bit rate $(\sigma^2 = \Sigma_{i=1}^{N} \sigma_i^2)$. Note that (10) is also computationally simple and only depends on the GOS and the means and variances of the bit rates generated by individual connections. These are directly available from the source characteristics defined earlier.

C. Equivalent Capacity

It now remains to combine the flow [(2) and (8)] and stationary [(10)] approximations into a single expression giving the equivalent capacity of a set of connections. As both approximations overestimate the actual value of the equivalent capacity[15] and are inaccurate for different ranges of connections characteristics, the equivalent capacity $\hat{C}$ is taken to be the minimum of $\hat{C}_{(F)}$ and $\hat{C}_{(S)}$. That is:

$$\hat{C} = \min\left\{ m + \alpha'\sigma, \ \sum_{i=1}^{N} \hat{c}_i \right\} \quad (11)$$

where N is the number of multiplexed connections and the different quantities involved are defined in (2), (8), and (10).

It should be mentioned, that although (11) is the minimum of the values obtained from both the flow and stationary approximations, it still overestimates the actual bandwidth requirements. As we shall see in the next section, this overestimation is, however, typically reasonable. Furthermore, a somewhat conservative allocation is preferable to underestimating the bandwidth requirements of connections, which could result in network congestion or even failure.

As a minor additional comment on (11), it should be recalled that some care must be exercised to avoid situations where the Gaussian assumption used in the stationary approximation does not hold. This typically happens with small numbers of very bursty connections with high peak rates, low utilizations, and long burst periods. In such cases, the stationary approximation can yield a lower capacity than actually required. These types of connections can, however, be easily detected (i.e., check how many can be multiplexed on a network link) and should be treated as constant bit rate connections with intensity equal to their individual equivalent capacity as given by the flow approximation [(2)].

D. Link Metrics

Based on the expression given in (11), it is then possible to define *link metrics* that characterize the current capacity allocation on network links. This metric must

provide a simple and compact form to store the bandwidth allocation information, while allowing for real-time updates and computations of allocation levels. A three-dimensional vector representation meeting these requirements is provided in (12).

$$L_j = \left(m = \sum_{i=1}^{N} m_i, \ \ \sigma^2 = \sum_{i=1}^{N} \sigma_i^2, \ \ \hat{C}_{(F)} = \sum_{i=1}^{N} \hat{c}_i \right)$$

$$(12)$$

where L_j is the link metric vector on link j, and as previously defined, m and σ^2 are, respectively, the mean and variance of the aggregate bit rate corresponding to all the connections routed on link j, and $\hat{C}_{(F)}$ is the sum of their individual equivalent capacities as defined in (2) and (8).

It is easily seen that few computations are required to obtain the bandwidth allocation on link j from the information stored in the link metric vector described in (12). Specifically, (11) indicates that the required operations consist of 1 square root, 1 multiplication, 1 addition, and a comparison. Such computations can clearly be performed in real-time by modern processors, even when large numbers of computations need to be executed nearly simultaneously.

The second major aspect of (12) is the "incremental" property of the link metric vector. In particular, updates can be performed based solely on the combined information contained in the link metric vector and the characteristics of the connections being established or released. There is no need for specific information on the individual characteristics of the connections currently routed over the link. In particular, a link metric update is executed as follows: 1) Upon receipt of a connect or disconnect request, the origin node associated with the connection computes, from the connection metric vector $(R_{peak}^{(i)}, \rho_i, b_i)$, a request vector of the form $r_i = \pm(m_i, \sigma_i^2, \hat{c}_i)$ (a + indicates a connect request and a − a disconnect). 2) The request vector r_i is sent to all nodes with links over which the connection is or should be routed, and the corresponding link metric vectors are updated by performing a simple vector addition/subtraction.[16]

Note that if the amount of computation needs to be reduced further, it is possible to assign connections to predefined traffic classes for which equivalent capacity values have already been computed. The definition of traffic classes could be continuously refined based on the information gathered by traffic-monitoring devices at access points to the network.

IV. NUMERICAL RESULTS

In this section, we provide a number of numerical examples that illustrate the equivalent capacity concept, check the accuracy and investigate limitations of the proposed approximations, and probe other important aspects

---

[15]Statistical multiplexing is ignored and $\beta \leq 1$ for the flow approximation, and the buffer is ignored in the stationary approximation.

[16]Note that in the case of a connect request, it is also necessary to verify if enough capacity is available on the link to accommodate the new connection.

such as the effect of bandwidth allocation inside the network. Numerical results are obtained from simulations or exact analysis, when available. They consider the effect of different parameters such as source characteristics, number and types of sources, and interactions between different traffic types within the network.

The first set of examples compares exact and approximate values of the equivalent capacity for cases that are representative of the range of behaviors of the approximation. Exact equivalent capacity values for one or more identical sources are computed by iteratively solving (1) or (3) for the unknown capacity. Approximate values are obtained from (11). Both the single and multiple sources examples demonstrate the relative accuracy of the expression over a wide range of system parameters, and also show that the two approaches, the fluid-flow and the stationary approximations, complement each other well.

Fig. 3 shows, as a function of source utilization, the actual amount of capacity that needs to be allocated to a single 4 Mb/s source to achieve a buffer overflow probability of 10^{-5} with a 3 Mbits buffer and a mean burst period of 100 ms. Three curves are displayed in Fig. 3, and give the mean bit rate, the exact equivalent capacity, and the value obtained from the flow approximation (2). The results illustrate the fact that either peak or mean bit rate allocation can be a substantial over- or underallocation, respectively. They also show the good agreement, for a single source, between exact and approximate values.

Fig. 4 illustrates the behavior, again as a function of source utilization, of the equivalent capacity when a small number, 5, of high-speed sources are multiplexed. It shows, as expected, that for small numbers of sources, the stationary approximation results in a substantial overestimate of the required capacity. On the other hand, the flow approximation, although also an overestimate, tracks the required capacity reasonably well especially at high loads. Fig. 5 considers the multiplexing of a larger number, 50, of sources (the aggregate peak rate is kept constant). It demonstrates that, as the number of connections is increased, the stationary approximation performs better. For the chosen connection characteristics, it provides a better approximation of the equivalent capacity at low loads, while at high loads the flow approximation is again more accurate. Of importance here is the fact that the flow and stationary approximations complement each other over different ranges of connection characteristics.

After investigating the impact of utilization for sources with a given mean burst period, we now fix the source utilization and let the mean burst period vary. This case is considered in Fig. 6, where in order to better illustrate the impact of long burst periods on the error made by the flow approximation,[17] a large number, 500, of connections is assumed. All sources have a fixed utilization of 0.5. As expected, the flow approximation grossly over-

[17]Recall that the flow approximation ignores the effect of statistical multiplexing.
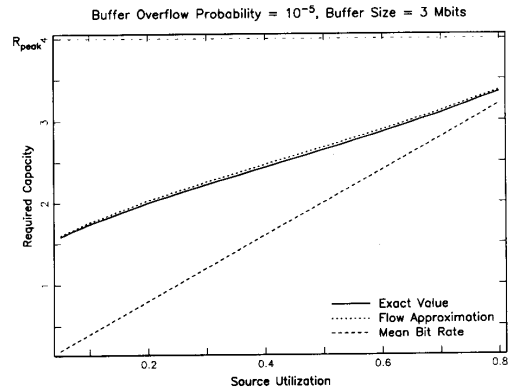

Fig. 3. Equivalent capacity for a single source with 4 Mb/s peak rate and 100 ms mean burst period.
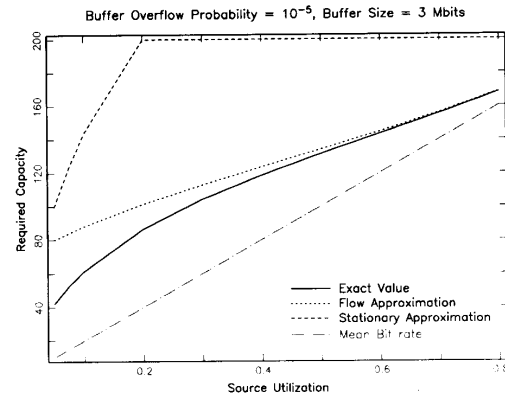

Fig. 4. Equivalent capacity for five sources, each with 40 Mb/s peak rate and 10 ms mean burst period.
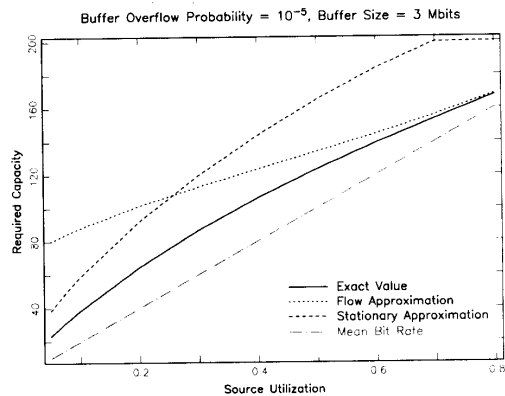

Fig. 5. Equivalent capacity for 50 sources, each with 4 Mb/s peak rate and 100 ms mean burst period.

estimates the equivalent capacity as the mean burst period is increased. However, for long burst periods the stationary approximation now provides a reasonable estimate, while the flow approximation remains more accurate for
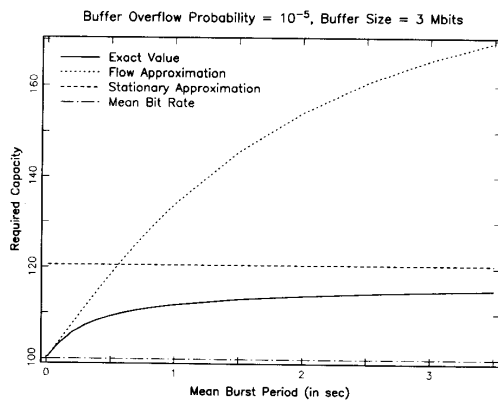
Buffer Overflow Probability = 10^{-5}, Buffer Size = 3 Mbits

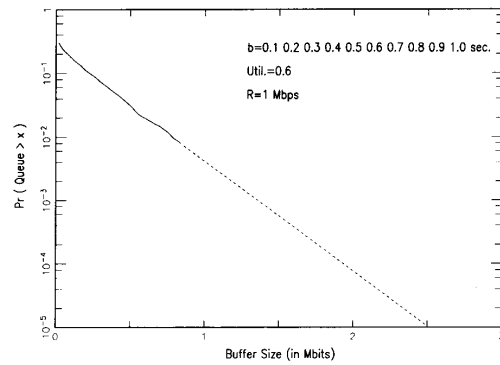Fig. 6. Equivalent capacity for 500 sources, each with 400 kb/s peak rate and 50% utilization.

Fig. 7. Buffer overflow probability for a mix of sources with different mean burst periods.

Fig. 8. Buffer overflow probability for a mix of sources with different peak rates.

Fig. 9. Example of insufficient bandwidth allocation when Gaussian assumption does not hold.

small burst periods. Here again, the flow and stationary approximations complement each other and provide a reasonable estimate of the equivalent capacity as connection characteristics vary.

The next set of examples briefly investigates the case of nonidentical sources. The results presented here are obtained by means of simulations, although the techniques of [17], [20], [23] could have been used to again solve numerically for the unknown capacity. A number of different cases are examined, and displayed in Figs. 7–9. The curves presented in these figures give the probability that the buffer content exceeds a given level, where the tail of the distribution has been obtained using a simple extension of the simulation results. (The transition from the simulation curves to the assumed exponential falloff is indicated by the change from plain to dotted lines.) The tail probabilities indicate whether the allocated capacity, obtained from the approximation, is sufficient or not to achieve the desired buffer overflow probability.

The impact of different mean burst periods is first considered in Fig. 7, where 10 sources, all with a peak rate of 1 Mb/s and a utilization of 0.6 but different mean burst periods ranging from 0.1 to 1 s, are multiplexed onto a single link. The individual equivalent capacities were computed using (2) for a 3 Mbits buffer and a desired buffer overflow probability of 10^{-5}. The corresponding individual bandwidth allocations range from 635 to 823 kb/s, for increasing mean burst periods. The total aggregate equivalent capacity was obtained from (11), and found to be the sum of the individual capacities. The figure shows that the aggregate equivalent capacity, although slightly conservative,[18] provides a reasonable estimate of how much capacity to allocate to meet the GOS requirement (the desired buffer overflow probability is achieved with a buffer size of about 2.5 Mb/s.)

The next example, shown in Fig. 8, considers a different mix of sources. The multiplexed sources consist of

[18]This is again expected, as allocating the sum of individual equivalent capacities ignores the effect of statistical multiplexing.
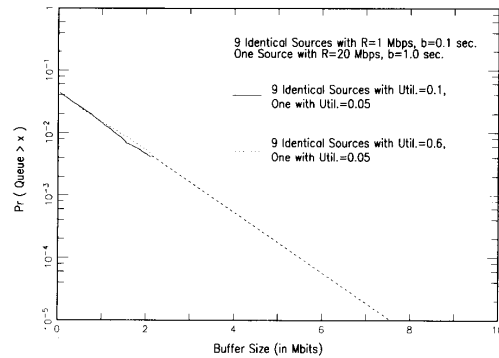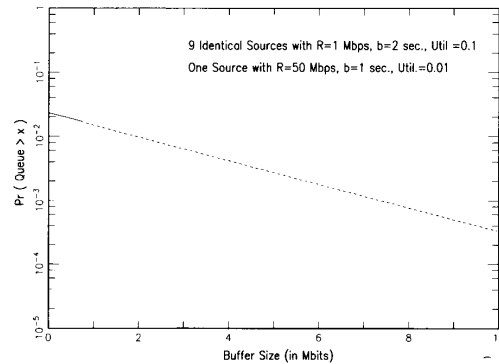
nine identical sources with a peak rate of 1 Mb/s and a mean burst period of 0.1 s, and a single source with a higher peak rate of 20 Mb/s and a 1 s mean burst period. This example examines the impact of very different peak rates in computing the required capacity allocation. The 20 Mb/s source has a 5% utilization, while two different utilizations (10% and 60%) are assumed for the 1 Mb/s

sources. Because of the large mean burst size (20 Mbits) of the high-speed source, a 10 Mbits buffer size is selected. The desired buffer overflow probability is again 10^{-5}, and the equivalent capacity of the high-speed source is found to be equal to 19.13 Mb/s, while the 1 Mb/s sources have equivalent capacities of 110 kb/s and 611 kb/s for utilizations of 0.1 and 0.6, respectively. The total equivalent capacity is once more the sum of individual equivalent capacities, and is equal to 20.12 and 24.63 Mb/s for the two different utilizations of the 1 Mb/s sources. As before, the figure indicates that the total equivalent capacity is a reasonable, although still conservative, estimate for the required capacity allocation.

Fig. 9 investigates a rather extreme case that illustrates an earlier comment on the limitations of the stationary approximation when the Gaussian assumption does not hold. We consider a situation where a number of relatively low-speed sources are multiplexed with a high-speed and very bursty source. We assume that nine 1 Mb/s sources are multiplexed with a single 50 Mb/s source. The 1 Mb/s sources have a utilization of 0.1 and a mean burst period of 2 s, while the 50 Mb/s source has a utilization of only 0.01 and a mean burst period of 1 s (an average burst is 50 Mbits long, but occurs very rarely!). The equivalent capacities of the 1 Mb/s sources are all equal to 0.6 Mb/s, while the 50 Mb/s source has an equivalent capacity close to its peak rate and equal to 49.13 Mb/s (a 10 Mbits buffer size is assumed.) The total equivalent capacity computed from (11) is this time given by the stationary approximation and equal to 24.66 Mb/s (the sum of the individual equivalent capacities is 54.53 Mb/s.) However, because of the large difference in connection characteristics, in particular peak rate, the Gaussian assumption used in the stationary approximation does not hold. The use of the stationary approximation can, therefore, result in insufficient capacity allocation. This is illustrated in Fig. 9, which shows, that for the assumed buffer size of 10 Mbits, the probability that the buffer overflow is about 10^{-3} instead of the desired 10^{-5}. As mentioned earlier, this problem can be overcome by isolating high-speed[19] connections and directly allocating to them their individual equivalent capacity.

The last example attempts to study some aspects of the interactions between connections inside the network. In particular, we investigate the potential impact of high-speed bursty connections on the bandwidth requirements of lower-speed ones in the following scenario. We take five groups of two connections each. In each group, one connection has characteristics R_{peak} = 25 Mb/s, ρ = 0.7, and b = 2 s, with an equivalent capacity of 24.58 Mb/s (a 10 Mbits buffer and a GOS of 10^{-5} are again assumed). The second connection is a lower-speed one with characteristics R_{peak} = 1 Mb/s, ρ = 0.8, and b = 1 s, and an equivalent capacity of 0.832 Mb/s. Each group is allocated a total equivalent capacity of 25.412 Mb/s, the sum
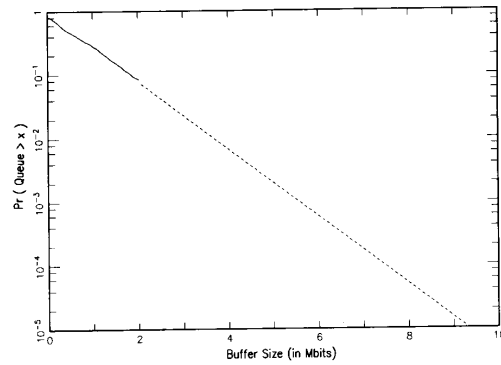

Fig. 10. Effect of interaction between high-speed and low-speed sources within the network.

of the two individual equivalent capacities. It is further assumed that this total equivalent capacity is actually the total link speed. The five low-speed connections coming over different links are then multiplexed onto a common link, while the high-speed connections are routed differently.

The purpose of this example is to study potential changes in the bandwidth requirements of the low-speed connections, as a result of their interactions with high-speed ones. In particular, we want to investigate the significance of the "gating" effect of high-speed bursts,[20] which can modify the effective peak rate of a low-speed connection and, therefore, its bandwidth requirements. This gating effect is caused by high-speed bursts, which when present force the queueing of data from the low-speed connection. The accumulated backlog is then transmitted at the end of the high-speed burst at a speed higher than the original peak rate of 1 Mb/s. This higher "effective" peak rate within the network can in turn require a higher bandwidth allocation. In order to estimate how much additional bandwidth may be needed, we first allocate the original equivalent capacity value and observe the resulting buffer content distribution. That is, we allocate a total of 4.16 Mb/s, the sum of the individual equivalent capacities, to the five low-speed connections routed on the same link, and obtain the corresponding buffer content distribution.

The result is shown in Fig. 10, which indicates that, despite the impact of the high-speed bursts, the original equivalent capacity remains sufficient to achieve the desired GOS. Two factors contribute to this result. The first one is the fact that the total equivalent capacity is originally an overestimate of the required bandwidth allocation. This built-in margin can, therefore, compensate for the increase in bandwidth requirement due to the gating effect. The second factor is that the impact of the gating effect itself is not too severe. Because of the interleaving effect of multiplexing, not all the data received from a low-speed connection is delayed during a high-speed

---

[19]Peak rate and mean burst size are such that only a few connections could be multiplexed on network links.

[20]They are, on the average, 50 Mb/s long.

burst. These conclusions hold across a wide range of conditions, as indicated by a number of other studies that further probed possible interactions inside the network. Results were obtained for a variety of connection mixes and characteristics, and consistently indicated that the equivalent capacity is sufficient to satisfy the GOS requirement. This provides further confidence in the general applicability of the bandwidth allocation techniques described in this paper.

V. EXTENSIONS AND ASYMPTOTIC BEHAVIORS

This section is dedicated to extending the equivalent capacity approximation developed in Section III to nonexponential sources, and to the study of some potentially interesting asymptotic behaviors. The extension to nonexponential sources is based on standard moment-matching techniques. It provides a simple approximation that allows us to apply the techniques developed for exponential sources to also estimate the bandwidth requirements of nonexponential sources. The study of asymptotic behaviors, e.g., as the load goes to 0, peak rate goes to infinity, etc., provides insight on the impact of different source characteristics and can help in classifying sources according to their requirements.

A. Extensions to Nonexponential Sources

A natural extension of the single-source fluid-flow model studied in Section III-A is to consider sources with generally distributed idle or burst periods. The study of such models is somewhat more complex, but some results are available. Sources with burst and idle periods distributed according to an underlying Markov chain can be analyzed using the results of [17], [20], [23].[21] The case of a single source with generally distributed burst periods is treated in [14], where the distributions of both the buffer content and the busy period are provided. Sources with general distributions for both idle and burst periods are more difficult to analyze, and the behavior of such systems can be shown to be described by a relation similar to that of a $GI/G/1$ queue. Because of these inherent difficulties in solving more general source models and our requirements for real-time solutions, we again resort to an approximation.

The approximation used to handle more general sources relies on moment-matching techniques that relate complex sources to simpler (e.g., exponential) "equivalent" ones. More specifically, we use a two-moment approximation that maps the first and second moments of the burst/idle period into the first moment of an equivalent, exponentially distributed burst/idle period. The use of such moment-matching approximations is quite standard (e.g., see [2], [12], [24], [25] for additional discussions and examples) and, as we shall see, provides relatively

---

[21]As mentioned earlier, the results provided in these papers are more general as they also allow the source bit rate to vary according to the state of the Markov chain.

accurate and computationally simple results. Based on this approximation, the equivalent capacity of the source is then computed from (2) using equivalent mean and idle periods.

The relation between the equivalent mean burst and idle periods, b' and I', and the original source characteristics must satisfy some basic constraints. First, the source utilization must be kept constant (load conservation) and second, the relation must be an identity for exponential sources. For simplicity, we assume that the relations giving the equivalent mean burst and idle periods are of the form $b' = \delta_b b$ and $I' = \delta_I I$, where b and I are the original mean burst and idle periods, and δ_b and δ_I are both functions of the first and second moments of the burst and idle periods. As the load ρ is directly related to the mean burst and idle periods, $\rho = b/(b + I)$, requiring that ρ be kept constant forces $\delta_b = \delta_I = \delta$. It now remains to select δ as a function of the first and second moments of the burst and idle periods, under the constraint that $\delta = 1$ for exponentially distributed burst and idle periods. A number of functions satisfy this constraint and, after a number of numerical studies, the following function was selected:

$$b' = \delta b, \quad \text{and } I' = \delta I, \quad \text{with } \delta = \frac{\overline{b^2}}{2b^2} \times \frac{\overline{I^2}}{2I^2} \quad (13)$$

where $\overline{b^2}$ and $\overline{I^2}$ are the second moments of the burst and idle periods, respectively. Note that, since ρ is constant and the mean idle period is not explicitly present in (2), the effect of the idle period distribution is only through δ.

Figs. 11 and 12 are used to illustrate the approximation proposed in (13). Fig. 11 covers the case of sources with exponential idle periods and deterministic, exponential, and hyperexponential burst periods, while Fig. 12 considers the symmetric case of sources with exponential burst periods and deterministic, exponential, and hyperexponential idle periods. All sources are taken to have the same peak rate ($R_{\text{peak}} = 4$ Mb/s), the same utilization ($\rho = 0.1$), and the same mean burst and idle periods ($b = 0.925$ ms, $I = 8.325$ ms). The buffer size is taken to be 1 Mbits, and the allocated capacity is determined to ensure a buffer overflow probability of 10^{-5}. In the case of a hyperexponentially distributed burst period, the mean duration of a burst state is either 0.25 or 2.5 ms with probabilities 0.7 and 0.3, respectively. This corresponds to burst sizes of 1 and 10 kbits, respectively. Similarly, a hyperexponentially distributed idle period has a mean duration of 1 ms with probability 0.8 and of 37.625 ms with probability 0.2.

Each figure plots the probability that the buffer contents exceed a given level for two different scenarios. In scenario 1, the allocated capacity is taken to be directly given by (2), irrespective of the actual distribution of the burst/idle period. In scenario 1, all sources are allocated a capacity of 414.2 kb/s, which is obtained under the assumption of exponential distributions. In scenario 2, the allocated capacity is again given by (2), but the burst/idle period distribution is taken into account through the use
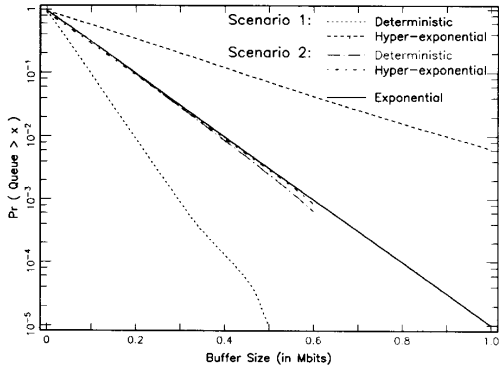
Fig. 11. Effect of moment matching approximation for sources with nonexponential burst periods. Scenario 1 does not use moment matching, while Scenario 2 does.
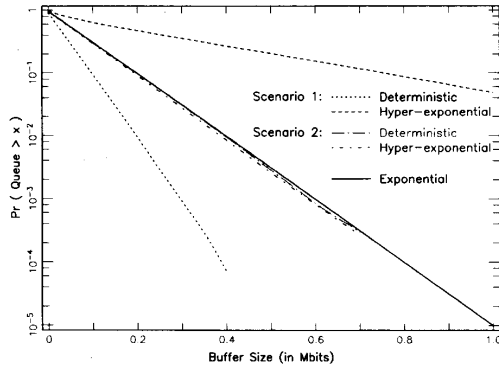

Fig. 12. Effect of moment matching approximation for sources with nonexponential idle periods. Scenario 1 does not use moment matching, while Scenario 2 does.

of equivalent mean burst or idle periods as given by (13). In scenario 2, sources with deterministic burst or idle periods are allocated a smaller capacity of 407 kb/s, while sources with hyperexponential burst or idle periods have larger capacity allocations of 433.2 and 464.5 kb/s, respectively. The capacity allocation of exponential sources is left unchanged at 414.2 kb/s.

As can be seen from the figures, the exponential model over- and underallocates capacity, respectively, when the burst or idle periods have deterministic or hyperexponential distributions. On the other hand, the simple approximation based on (13) yields relatively accurate estimates of the capacity needed to meet the desired GOS. There are, however, limitations to the accuracy of the proposed approximation, in particular for sources for which both the burst and idle periods are not exponential. A simple example is the case of deterministic burst and idle periods. Such a source is periodic and, provided the buffer can hold the data received in a burst, the required capacity allocation is simply the mean bit rate. However, for deterministic burst and idle periods, δ, as given by (13), is always equal to 0.25. This scaling, although lowering the capacity allocation, can be insufficient and yield values

higher than the mean bit rate. In spite of these limitations, numerical experiments have found the approximation to be generally accurate, in particular for sources where either the burst or idle periods remain exponential. This should prove useful to handle many practical, nonexponential sources.

B. Asymptotic Behavior

The study of the asymptotic behavior of (2) helps identify the relative impact of different source characteristics, and to understand some of the limitations of the approximation. Interesting behaviors take place for limit values of the utilization, burst duration, and peak rate.

The case $\rho \rightarrow 1$ was mentioned earlier and is of little interest, as it can be easily shown that a 100%, utilized source requires the expected peak rate allocation. The case $\rho \rightarrow 0$ is of more interest as it exhibits two possible limits, depending on the sign of the quantity $(\alpha R_{\text{peak}} b) - x$.

$$\lim_{\rho \rightarrow 0} \hat{c} = \begin{bmatrix} 0, & \text{if } \alpha R_{\text{peak}} b \leq x \\ R_{\text{peak}} \left(1 - \dfrac{x}{\alpha R_{\text{peak}} b}\right), & \text{if } \alpha R_{\text{peak}} b \geq x. \end{bmatrix}$$

(14)

Intuitively, (14) states that as $\rho \rightarrow 0$, the required capacity allocation also goes to 0 only if the buffer is "large enough" compared to the mean burst size. How large is large enough is a function of the desired GOS specified in α, i.e., the buffer should be able to hold α bursts of average size. When the buffer size is not sufficient, the required capacity has a nonzero limit since, although bursts are less and less frequent, the service rate must still handle large bursts whenever they arrive (recall that (2) approximates the buffer overflow probability when the source is active). In fact, this required capacity can be made arbitrarily close to the source peak rate, irrespective of the utilization. This indicates that standard bandwidth allocation may be rather inefficient for very bursty and very low utilization sources. A more efficient use of allocated bandwidth can be achieved if the information contained in a burst can be delayed, therefore reducing the source effective peak rate. In cases where the information generated by the source is delay sensitive, a possible alternative is, as suggested in [4] and [5], to allocate bandwidth only when a burst is present. This may, however, not always be feasible or even practical.

Another set of possible asymptotic behaviors involve limiting values for the mean burst period, i.e., $b \rightarrow 0$ or $b \rightarrow \infty$. In such cases, (2) behaves again as intuitively expected, and the capacity requirements are easily found to be the mean bit rate ρR_{peak} and the peak bit rate R_{peak} when $b \rightarrow 0$ and $b \rightarrow \infty$, respectively. Note that the limit for $b \rightarrow \infty$ is independent of the load, which is consistent with the expression in (14) for large b.

A more interesting behavior can be observed when we simultaneously let the mean burst duration go to 0 and the peak rate to ∞, while keeping the mean bit rate $m = \rho R_{\text{peak}}$ and the mean burst size $B = R_{\text{peak}} b$ constant. The interest

of such a scenario is that it reflects the impact of increased input link speed on the bandwidth requirements of a source.

As in the case $\rho \to 0$, two different behaviors are possible depending again on the relation between the buffer size x and the mean burst size B. Cases where the buffer is "small" compared to the mean burst size ($\alpha B > x$) are not further studied as the assumption of infinite peak rate then forces the allocated capacity to also be infinite. Instead, we concentrate on systems with reasonably sized buffers satisfying $\alpha B \le x$. For such systems, a standard asymptotic expansion of (2) gives:

$$\hat{c}_\infty = m \frac{x}{x - \alpha B}. \tag{15}$$

This indicates that the required capacity allocation $\hat{c}_\infty$ is proportional to the mean bit rate, with a scaling factor which is a function of the buffer size, mean burst size, and GOS.

It is interesting to notice that the expression given in (15) can also be obtained from an approximation for the required capacity allocation of a related $M/M/1$ queue. It is well known that the probability distribution of the waiting time in an $M/M/1$ queue (e.g., see p. 203 of [16]) is of the form:

$$P(w \le y) = 1 - \tilde{\rho} e^{-\mu(1 - \tilde{\rho})y} \tag{16}$$

where $1/\mu$ is the mean service rate, and $\tilde{\rho}$ is the load. These two quantities can be expressed as functions of the mean bit rate m, mean burst size B, and allocated capacity c:

$$\frac{1}{\mu} = \frac{B}{c}, \quad \text{and } \tilde{\rho} = \frac{m}{c}.$$

Using (16) with the above relations and the fact that the waiting time y in an $M/M/1$ queue is related to the buffer content x by the relation $y = x/c$, the probability that the buffer content exceeds a given level x is found to be:

$$P(Q > x) = F(x) = \frac{m}{c} e^{-(x(c - m)/cB)}. \tag{17}$$

Making the simplifying assumption $m/c = 1$, similar to that used in (1) ($\beta = 1$), one obtains the following expression for the allocated capacity necessary to achieve a buffer overflow probability below ϵ:

$$\hat{c}_{M/M/1} = m \frac{x}{x - \alpha B}$$

which is identical to the limiting expression for the fluid-flow model given in (15).

The difference between the equivalent capacity, computed using either the asymptotic approximation (15) or the fluid-flow model with the actual source peak rate (2), is illustrated in Fig. 13. The source used in this example has a peak rate of 4 Mb/s and a utilization of 10%; a 3 Mbits buffer size was assumed with a desired buffer overflow probability of 10^{-5}. Fig. 13 consists of three curves, giving the different values of the equivalent capacity as functions of the mean burst period. The three curves are obtained from (2), (15), and by solving (17)
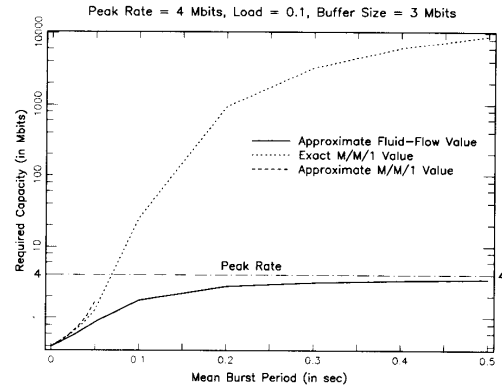

Fig. 13. Equivalent capacities for asymptotic and fluid-flow models.

for $F(x) = \epsilon$. It should be noted that, for small burst periods, the asymptotic approximation of (15) is close to the actual solution of (17). However, because of the large buffer assumption $\alpha B \le x$, it cannot be extended to long mean burst periods.

As expected, the assumption of an infinite input rate results in larger equivalent capacities, in particular for relatively large burst periods. This points to the importance of the source peak rate in determining bandwidth allocation.

VI. Conclusion

In this paper, we have presented a general approach to the problem of bandwidth allocation in high-speed fast-packet-switched networks. The approach relies on computationally simple approximations that estimate the bandwidth requirements, or equivalent capacity, of both individual and multiplexed connections. Based on the resulting simple expression for the equivalent capacity of connections, we defined link metrics that allow for real-time updates and checking of bandwidth allocations on network links while only requiring minimum storage of information about existing connections. This provides the basis for efficient and practical implementations of advanced network control functions, such as admission control and dynamic routing.

A number of possible extensions and topics for future research can also be identified. An important one is the investigation of better approximations for the parameter β of Section III. Although, as mentioned in [17], this problem seems to be difficult, approximations taking the impact of the mean burst period into account are desirable and could further improve the accuracy of the final equivalent capacity expression. A possible approach is through the use of "curve-fitting" techniques, based for example, on studies similar to those shown in Figs. 1 and 2. Another promising area for extensions is that of approximations for more general sources that improve upon the simple expression used in (13). This can greatly improve the estimation of the bandwidth requirements of a number of complex real-life sources.

ACNKNOWLEDGMENT

REFERENCES

[1] Special Issue on Asynchronous Transfer Mode, *Int. J. Digit. Analog Cabled Syst.*, vol. 1, no. 4, 1988.

[2] S. Albin, "Approximating a point process by a renewal process II: Superposition arrival process to queues," *Operat. Res.*, vol. 32, no. 2, pp. 1133–1162, Sept./Oct. 1984.

[3] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of a data-handling system with multiple sources," *Bell Syst. Tech. J.*, vol. 61, no. 8, pp. 1871–1894, Oct. 1982.

[4] P. Boyer, "A congestion control for the ATM," presented at the 7th ITC Sem., Morristown, NJ, Oct. 1990.

[5] P. Boyer, J.-R. Louvion, and D. Tranchier, "Intelligent multiplexing in ATM based networks," presented at the IEEE MULTIMEDIA'90, Bordeaux, France, Nov. 1990.

[6] I. Cidon and I. Gopal, "PARIS: An approach to integrated high-speed private networks," *Int. J. Digit. Analog Cabled Syst.*, vol. 1, no. 2, pp. 77–86, Apr./June 1988.

[7] I. Cidon, I. Gopal, G. Grover, and M. Sidi, "Real-time packet switching: A performance analysis," *IEEE J. Select. Areas Commun.*, vol. 6, no. 9, pp. 1576–1586, Dec. 1988.

[8] M. Decina and T. Toniatti, "On bandwidth allocation to bursty virtual connections in ATM networks," in *Proc. ICC'90*, 1990, paper 318.6, pp. 844–851.

[9] M. Decina, T. Toniatti, P. Vaccari, and L. Verri, "Bandwidth assignment and virtual call blocking in ATM networks," in *Proc. INFOCOM'90*, 1990, pp. 881–888.

[10] L. Dittman and S. Jacobsen, "Statistical multiplexing of identicals bursty sources in an ATM network," in *Proc. GLOBECOM'88*, 1988, paper 39.6, pp. 1293–1297.

[11] S. Erfani and M. Malek, "Dynamic access capacity management in a multiservice packet-mode environment," in *Proceedings of the 1989 Network Management and Control Workshop*, A. Kershenbaum, M. Malek, and M. Wall, Eds. New York: Plenum, 1990, pp. 277–285.

[12] A. A. Fredericks, "Congestion in blocking systems—A simple approximation technique," *Bell Syst. Tech. J.*, vol. 59, no. 6, pp. 805–827, July/Aug. 1980.

[13] G. Gallassi, G. Rigolio, and L. Fratta, "ATM: Bandwidth assignment and bandwidth enforcement policies," in *Proc. GLOBECOM'89*, 1989, paper 49.6, pp. 1788–1793.

[14] D. P. Gaver, Jr. and R. G. Miller, Jr., "Limiting distributions for some storage problems," in *Studies in Applied Probability and Management Science*, K. J. Arrow, S. Karlin, and H. Scarf, Eds. Stanford, CA: Stanford University Press, 1962, pp. 110–126.

[15] P. Joos and W. Verbiest, "A statistical bandwidth allocation and usage monitoring algorithm for ATM networks," in *Proc. ICC'89*, 1989, paper 13.5, pp. 415–422.

[16] L. Kleinrock, *Queueing Systems, Vol. 1: Theory.* New York: Wiley, 1975.

[17] L. Kosten, "Liquid models for a type of information storage problems," *Delft Prog. Rep.: Math. Eng., Math. and Inform. Eng.*, vol. 11, pp. 71–86, 1986.

[18] H. Kröner, T. H. Theimer, and U. Briem, "Queueing models for ATM systems—A comparison," presented at the 7th ITC Sem., Morristown, NJ, Oct. 1990.

[19] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J. D. Robbins, "Performance models of statistical multiplexing in packet video communications," *IEEE Trans. Commun.*, vol. 36, no. 7, pp. 834–844, July 1988.

[20] D. Mitra, "Stochastic theory of a fluid model of producers and consumers coupled by a buffer," *Advanc. Appl. Prob.*, vol. 20, pp. 646–676, Sept. 1988.

[21] J. A. S. Monteiro, M. Gerla, and L. Fratta, "Statistical multiplexing in ATM networks," in *Proc. 4th Int. Conf. Data Commun. Perform.*, Barcelona, Spain, 1990, paper 49.6, pp. 148–162.

[22] F. C. Schoute, "Simple decision rules for acceptance of mixed traffic streams," in *Proc. 12th Int. Teletraffic Cong.*, 1988, pp. 771–777.

[23] T. E. Stern, "Analysis of separable Markov-modulated queueing processes," Tech. Rep. CU/CTR/TR 113-8841, Columbia Univ., New York, NY, 1988.

[24] W. Whitt, "Approximating a point process by a renewal process I: Two basic methods," *Operat. Res.*, vol. 30, no. 1, pp. 125–147, Jan./Feb. 1982.

[25] ——, "Heavy traffic approximations for service systems with blocking," *AT&T Bell Lab. Tech. J.*, vol. 63, no. 5, pp. 689–708, May/June 1984.

[26] G. M. Woodruff and R. Kositpaiboon, "Multimedia traffic management principles for guaranteed ATM network performance," *IEEE J. Select. Areas Commun.*, vol. 8, no. 3, pp. 437–446, Apr. 1990.

Roch Guérin (S'84–M'86–SM'91) was born in Paris, France, on June 26, 1959. He received the "Diplôme d'Ingénieur" from the École Nationale Supérieure des Télécommunications, Paris, France, in 1983, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1984 and 1986, respectively.

Since August 1986, he has been with the IBM Research Division, T. J. Watson Research Center, Yorktown Heights, New York, where he now is a member of the Broadband Networking Department, where he is working on the architecture, implementation, and field trial testing of high-speed networks. In addition to these activities, his current research interests are in the area of performance analysis and modeling of communications systems, in particular congestion control, bandwidth management, dynamic routing, and their interactions in high-speed networks. He is a Technical Editor for the *IEEE Communications Magazine*.

Dr. Guérin is a member of Sigma Xi, the IEEE Communications and Information Theory Societies.
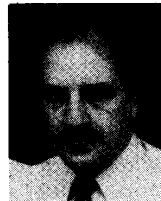
Hamid Ahmadi received the B.S., M.S., and Ph.D. degrees in electrical engineering from Columbia University, New York, NY, in 1976, 1978, and 1983, respectively.

From 1980 to 1984, he was with the Switching and Signaling Systems Department at Bell Laboratories, Holmdel, NJ, where he was involved in the area of performance analysis and protocol studies of signaling networks. He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, in 1984, where he has been a member of the Telecommunication Systems Department. He spent a two-year assignment at the IBM Research Division in Zurich, Switzerland, during 1986–1987. He is currently Manager of the Network Architecture and Analysis Group responsible for research into high-speed network traffic control architectures, performance analysis of communication systems and wireless mobile communication networks. He is a Technical Editor of the IEEE TRANSACTIONS ON COMMUNICATIONS, and is currently an Adjunct Professor in the Graduate Center at the Polytechnic University, where he teaches graduate courses in computer communication networks.

Mahmoud Naghshineh received the Vordiplom in electrical engineering from TH Aachen, Germany, in 1985, the B.S. degree in computer engineering and the M.S. degree in electrical engineering from the Polytechnic University, Brooklyn, NY, in 1988 and 1991, respectively.

He joined IBM in 1988. From 1988 to 1989, he worked on the architecture and performance evaluation of large systems. Since 1990, he has been working in the Network Analysis Group at IBM's T. J. Watson Research Center, Hawthorne, NY. His work involves modeling and performance analysis of communication networks.