EXTENDING THE EFFECTIVE-BANDWIDTH CONCEPT TO NETWORKS WITH PRIORITY CLASSES

by

Arthur W. Berger[1] and Ward Whitt[2]
AT&T Labs

March 25, 1998

IEEE Communications Magazine 36 (1998) 78–84

[1]Room 1K-211, Holmdel, NJ 07733-3030; awberger@att.com
[2]Room A117, 180 Park Avenue, Florham Park, NJ 07932-0971; wow@research.att.com

Abstract

ATM switches are now being designed to allow connections to be partitioned into priority classes, with packets being emitted for higher priority classes before packets are emitted for lower priority classes. Accordingly, allocation of network resources based on different priority levels is becoming a realistic possibility. Thus we need new methods to do connection admission control and capacity planning that take account of the priority structure. In this paper we show that the notion of effective bandwidths can be used for these purposes when appropriately extended. The key is to have admissibility of a set of connections determined by a linear constraint for each priority level, involving a performance criterion for each priority level. For this purpose, connections are assigned more than one effective bandwidth, one for its own priority level and one for each lower priority level. Candidate effective bandwidths for each priority level can be determined by using previous methods associated with the first-in first-out discipline, including the method based on large-buffer asymptotics. The proposed effective-bandwidth structure makes it possible to apply product-form stochastic loss network models to do dimensioning.

1. Introduction

Emerging high-speed communication networks, such as broadband ISDN networks that employ ATM technology, tend to be packet networks rather than circuit-switched networks because the packet structure allows for better resource sharing. In a packet network, sources do not require dedicated bandwidth (e.g. circuits) for the entire duration of a connection. Unfortunately, however, the enhanced flexibility of packet networks also makes it more difficult to effectively control the admission of connections seeking to enter an existing network and to plan the capacity of future networks when they are designed.

The problems of admission control and capacity planning in a packet network may be addressed by a concept known as the "effective bandwidth" or "equivalent bandwidth" of a connection. When employing this concept, an appropriate effective bandwidth is assigned to each connection and each connection is treated as if it required this effective bandwidth throughout the active period of the connection. The feasibility of admitting a given set of connections may then be determined by ensuring that the sum of the effective bandwidths is less than or equal to the total available bandwidth (i.e., the capacity). By using effective bandwidths in this manner, the problems of admission control and capacity planning are addressed in a fashion similar to that employed in circuit-switched networks.

Of course, the actual bandwidth (bit rate) needed by each variable-bit-rate connection is uncertain and fluctuates over time, as depicted in Figure 1. The actual required bandwidth fluctuates between some minimal level, perhaps 0, and a peak rate, which is typically determined by the speed of the access line. When there are many independent connections, the aggregate required bandwidth should usually be close to the sum of the average rates, by virtue of the law of large numbers. The aggregate required bandwidth should fluctuate around the overall average rate. Clearly, if the required bandwidth should only rarely exceed the capacity (the maximal output rate), the effective bandwidth for each connection should be at least its average rate. In summary, it is evident that the effective bandwidth of a connection should be some value between its average rate and its peak rate. Any particular value that is used is necessarily an approximation, but potentially a very useful approximation.

Given effective bandwidths, the problems of connection admission control and dimensioning in packet networks simplify to well understood techniques for multi-rate circuit-switched networks. A new connection is admitted if the sum of the effective bandwidths is less than the capacity. Let e_i

be the effective bandwidth of a connection of type i; let n_i be the number of connections of type i; let I be the number of connection types; and let c be the capacity of a link. (For ATM networks the link could be replaced by a virtual path (VP).) The set of connections determined by the vector $(n_1, \ldots, n_I)$ is said to be *admissible* if

$$\sum_{i=1}^{I} e_i n_i \leq c \; . \tag{1}$$

When the network contains multiple constrained resources, there is such a constraint for each resource. Then a set of connections is deemed admissible if inequality (1) holds for each resource.

A candidate new connection is admitted if the set of existing connections plus the new connection produces a feasible set of connections. Otherwise, the candidate new connection is rejected. To do dimensioning, we can specify arrival rates and average holding times for each connection type. Then, assuming a product-form stochastic loss network model, as in Ross [22], we can compute blocking probabilities for each connection type for any given capacity. These blocking probability calculations can be efficiently performed by numerically inverting the generating function of the normalization constant in these product-form models, as in Choudhury, Leung and Whitt [5], [6]. We then choose the capacity of each resource so that the blocking probabilities are suitably small. Moreover, we need not use a complete-sharing policy. We can improve performance by imposing upper-limit and guaranteed-minimum constraints on the connection classes. An upper limit of U_i on type i restricts the number of type-i connections that can be simultaneously present to be at most U_i. A guaranteed minimum of M_i for type i is a constraint on all other types, ensuring that there is always room for at least M_i type-i connections. With these constraints, the blocking probabilities can still be efficiently computed by numerical inversion. Moreover, the capacities and sharing parameters can be found by a search algorithm, as in Choudhury, Leung and Whitt [7].

Over the last ten years considerable work has been done on effective bandwidths. A theoretical basis was developed in the context of large-deviation asymptotics, which we will briefly review in Section 4.2. Early work was done by Hui [16], Kelly [18], Gibbens and Hunt [14] and Guerin, Ahmadi and Naghshineh [15]. The theory was extended by Chang [2], Kesidis, Walrand and Chang [20], Elwalid and Mitra [12], Whitt [23] and others. Recent summaries can be found in Chang and Thomas [3], de Veciana, Kesidis and Walrand [10] and Kelly [19]. Unfortunately, however, the effective-bandwidth approach based completely on the large deviations theory is often not a very accurate approximation; see Choudhury, Lucantoni and Whitt [8]. Hence, various heuristic refinements have been proposed, many abandoning the linear structure (1). However, as indicated above, the linear structure in (1) can greatly assist engineering. Thus we keep (1) and allow the effective bandwidths e_i to be adjusted as needed. In particular, given a nonlinear admissible set associated with some other admission control procedure, we can obtain effective bandwidths by introducing a linear approximation to the nonlinear admissible set. To do so, we might exploit knowledge of the typical operating region. We could use a linear hyperplane at the boundary of the admissible set near the typical operating region. For example, consider the case of two classes. A nonlinear admissible set might look as depicted in Figure 2. We might know that the typical operating region is the shaded region in Figure 2. Then we might approximate the admissible set by a linear hyperplane, chosen to be tangent to the admissible set at a point near the typical operating region. This line implicitly defines effective bandwidths for the two classes. In particular, the approximate effective bandwidths are $e_i^* = c/n_i^*$ where n_i^* is the point on the n_i axis intersected by the tangent line.

The concept of effective bandwidths has been developed for buffers using the first-in first-out (FIFO) service discipline. However, now ATM switches are being designed to allow the connections to be partitioned into priority classes with packets being emitted from higher priority classes before lower priority classes. This priority structure is useful to meet the different requirements of the diverse traffic that will be carried at ATM networks. Typical implementations have from two to four priority classes. The highest priority class might be constant-bit-rate (CBR) traffic. The next priority class might be real-time (interactive) video traffic. Non-real-time variable-bit-rate (VBR) traffic could be a lower priority class, which might be further divided into two priorities, making a lowest priority class for best-effort or available-bit-rate (ABR) traffic.

It is natural, then, to consider how the concept of effective bandwidths should be modified to properly take account of priority classes. And that is the topic of this paper. Here we present a more informal treatment focusing on engineering insights. We have given a more formal treatment in Berger and Whitt [1], to which we refer for additional details.

2. Modifications of Effective Bandwidths for Priorities

In this section we present an informal engineering argument to show that, regardless of the method used for computing effective bandwidths, if delay priorities are implemented in the network node, then practical and efficient engineering rules should use a linear constraint per-priority (leading to a trapezoidal admissible set for two priorities) and wherein a given connection type is associated with *multiple* effective bandwidths.

Before doing so, we point out that a more formal mathematical development based on large-buffer asymptotics is presented in [1]. There we show that the admissible set resulting from full asymptotic analysis does *not* actually have the proposed structure, but that a reasonable approximation does. We also review the related literature in [1]. Here we simply point out that other researchers previously began to examine the impact of non-FIFO queueing on bandwidth allocation and admission control in high-speed networks. See de Veciana and Kesidis [9] for the generalized processor-sharing policy and Chang and Zajic [4], Elwalid and Mitra [11, 13], Kulkarni, Gun and Chimento [21] and Zhang [24] for priority disciplines. With priorities, Mitra [11] and Kulkarni et al. [21] focus on the case in which all classes are in a single queue sharing a common buffer, whereas we and the others consider the case in which each class has its own queue with its own buffer. The paper by Elwalid and Mitra [13] is closest to this paper and [1]. Of particular relevance to the present paper, they point out that the admissible set can often be approximated by one with linear boundaries, i.e., a trapezoid as in Figure 4 herein. Their analysis can be interpreted as providing additional support for our proposal.

To consider how effective bandwidths might be extended to accommodate priority classes, consider the simple case of a single link with two connection types, with type 1 having priority over type 2. Before introducing priorities, the admissible set is the set of pairs (n_1, n_2) such that $e_1 n_1 + e_2 n_2 \leq c$, where e_i is the effective bandwidth of class i and c is the capacity (output rate) as depicted in Figure 3.

The primary reason we should want to use priority service is that the lower priority class has a looser performance criterion than the higher priority class. Thus, when the higher priority class has "filled the link" according to its performance criterion, there should still be some bandwidth leftover for the lower priority class.

To be more precise, suppose that each priority class has its own buffer, with the class-i buffer having capacity b_i. The performance criterion for class i might be that the long-run proportion of cells lost due to buffer overflow be less than p_i. The class-2 criterion might be weaker because b_2 is greater than b_1 or because p_2 is greater than p_1, or both.

Of course, if the high-priority class is CBR traffic or nearly CBR traffic, then its peak rate would be very close to its average rate, so that there would be negligible room for class 2 when class 1 reaches its performance limit. However, if class 1 traffic has considerable variability, then its peak rate might be much greater than its average rate, so that there might be considerable room in the link in terms of average rate when the class-1 priority limit is reached. For example, in ATM if the variable-bit-rate (VBR) real-time connections have filled the link according to their effective bandwidth, there would likely be room for some lower-priority available-bit-rate (ABR) connections. Indeed, the occupancy (in average sense) might well be 50% or less when the VBR traffic is at its upper limit. Even with CBR high-priority connections, there may be some spare bandwidths for lower-priority connections, because performance criteria on cell delay variation might limit the occupancy of CBR connections to, say, 90%.

Given that some class-2 connections can be admitted when class 1 is at its upper limit, we expect a vertical segment on the right of the admissible set. Instead of the triangular admissible set in Figure 3, we should anticipate the trapezoidal admissible set in Figure 4.

In order to have the trapezoidal admissible set in Figure 4, we need a second linear constraint. We now should have the pair of constraints:

$$\begin{aligned} e_1 n_1 &\leq c \\ e_1^2 n_1 &+ e_2 n_2 \leq c \ . \end{aligned} \tag{2}$$

The first constraint is the same constraint for class 1 above. The second constraint is the new linear constraint, which agrees with the old constraint for class 2 alone. The new parameter e_1^2 is determined by the height of the vertical segment on the right of the trapezoidal admissible set in Figure 4.

In constructing the trapezoidal admissible set in Figure 4, we have assumed that we know the class-2 limits when class-1 is at its lower and upper limits. The linearity in between can be regarded as the effective-bandwidth approximation.

It is useful to interpret the new parameter e_1^2 in (2). The parameter e_1^2 can be regarded as the effective bandwidth for a priority-1 connection that is subject to the priority-2 performance criterion. We say that e_1^2 is the effective bandwidth for a priority-1 connection *as seen by* priority 2. Given the sensible case in which the priority-1 criterion is tighter than the priority-2 criterion, we have

$$e_1^2 < e_1 \ . \tag{3}$$

In the construction of Figure 4 from Figure 3, we relied on the inequality (3). If instead we have $e_1 < e_1^2$, then the first constraint in (2) would be vacuous.

Figures 3 and 4 are also useful to graphically see the advantage of introducing priority classes. When there is a large vertical segment at the right in the trapezoidal admissible set, then the trapezoidal admissible set is much larger than the triangular admissible set. Constructing the admissible sets with and without priority classes can be very helpful to see the advantage of having priorities, where in the case of no priorities (FIFO service) the admission of connections of any type would be subject to the strictest (otherwise priority-1) performance criterion. In some cases priorities may provide a big gain, while in other cases they may only provide a modest gain. If priorities were used incorrectly (so that inequality (3) were reduced), then the admissible set with priorities would actually be strictly smaller than without priorities.

The notion of per-priority effective bandwidth generalizes to an arbitrary number of priority classes. For three priority classes, the admissible set is

$$\begin{array}{lllllll} e_1^1 n_1 & \leq & c & & & & \\ e_1^2 n_1 & + & e_2^2 n_2 & \leq & c & & \\ e_1^3 n_1 & + & e_2^3 n_2 & + & e_3^3 n_3 & \leq & c \; , \end{array} \tag{4}$$

where e_i^k is the effective bandwidth for a priority-class-i connection as seen by priority k, with $i \leq k$ in all cases. In (4) we have used e_i^i to denote e_i. Multiple connection types within a given priority class are treated just as with FIFO. Let i denote the priority level and let j denote the connection type, where $1 \leq j \leq J_i$ and $1 \leq i \leq I$. Let e_{ij}^k denote the effective bandwidth of a priority-i type-j connection as seen by priority k. We need e_{ij}^k only for $k \geq i$. With I priority levels, the admissible set is determined by the I constraints

$$\sum_{i=1}^{k} \sum_{j=1}^{J_i} e_{ij}^k n_{ij} \leq c \; , \quad k = 1, \ldots, I \; . \tag{5}$$

The sum over i in (5) could be extended to all i (up to I) provided that we set $e_{ij}^k = 0$ for $k < i$.

3. Loss Versus Delay Performance Criteria

There are two different performance criteria that are commonly considered: cell loss probabilities and delay tail probabilities. With the FIFO discipline, these two criteria are closely related. It is common to use the tail probability of the queue-length distribution in an unlimited-buffer model to approximate the cell loss probability. However, assuming constant-size cells, the delay at any time is a constant multiple of the queue length. Thus, with the FIFO discipline and an unlimited-buffer approximation, any delay performance criterion translates into an equivalent cell loss probability requirement.

However, with priority classes the equivalence between delay and cell loss no longer holds. A lower priority class cell has to wait not only for all cells of its priority and all higher priorities that are currently in the system; it also has to wait for new higher-priority cells that arrive *after* the lower-priority cell arrives, but before it can receive service. Thus, the delay can be much greater than determined by the workload vector seen upon arrival.

Thus, to be specific and to avoid confusion, in the beginning of Section 2 we stipulated that each priority class had its own buffer and that the performance criterion for each class was based on the cell loss probability. However, our approach to effective bandwidth with priorities is quite general, so that it should accommodate variations in the model and performance criteria.

4. Determining the Effective Bandwidths

Our analysis so far holds independently of how the given effective bandwidths are calculated. In the context of FIFO service, various methods have been proposed for making such calculations. Any of these methods could be extended to incorporate the per-priority effective bandwidths proposed herein. In the present section we review three methods that have been used in the FIFO context and show how they can be adapted to priority service. All three methods, both in the FIFO context and in the generalization to priorities, exploit the assumed linearity in the effective bandwidth constraints, equations (1) and (5). (The reader can skip any of the following subsections without loss of continuity.)

4.1. Measurements at Boundary Points of the Admissible Set

Suppose that the FIFO-service method is based on determining the maximum number of admissible connections of a given type when no other connection types are present. In particular, to determine e^i_{ij}, consider only priority-i type-j connections for one fixed j. Find the upper limit $\bar{n}_{ij}$ for each connection type alone to obtain parameter specification.

$$e^i_{ij} = c/\bar{n}_{ij} \ , \tag{6}$$

which corresponds to the constraint

$$e^i_{ij} n_{ij} \le c \ . \tag{7}$$

(In using (6) we ignore integrality constraints, i.e., the requirement that the number of connections must be some integer. Assuming that the capacity c is relatively large, this effect should be minor.)

So far we have determined the effective bandwidths e^k_{ij} for $k = i$. Now we determine e^k_{ij} for $k > i$. First fix i and k with $k > i$. We consider a feasible number of priority-i type-j connections established on the link, say n^o_{ij}. This number might be the maximum number admissible given the priority-i criterion, $\bar{n}_{ij}$, or it might be a lower value that corresponds to a designed engineering point. Given n_{ij}^o, we then see how many priority-k type-ℓ connections can be admitted for any fixed ℓ, considering the priority-k performance criterion. Suppose that this number is $m_{k\ell}^o$. We then let

$$e_{ij}^k = (c - e_{k\ell}^k m_{k\ell}^o)/n_{ij}^o \ . \tag{8}$$

Equation (8) corresponds to the constraint

$$e_{ij}^k n_{ij} + e_{k\ell}^k n_{k\ell} \le c \ . \tag{9}$$

In (9), we first determine a value for n_{ij}, n_{ij}^o. Then, with that value n_{ij}^o in place, we determine the upper limit on $n_{k\ell}$, $m_{k\ell}^o$. Since the inequality (9) should be an equality at the upper limit (again ignoring integrality problems) and since $e_{k\ell}^k$ has previously been determined, we can solve for the single missing parameter e_{ij}^k, obtaining the equation (8).

In the case where n_{ij}^o is chosen to be the maximum number admissible, $\bar{n}_{ij}$, then $m_{k\ell}^o$ is a natural measure of the benefit from using per-priority effective bandwidths, since $m_{k\ell}^o$ would be zero with effective bandwidths based on FIFO service. Moreover, when n_{ij}^o equals $\bar{n}_{ij}$, equation (8) can be expressed as:

$$e_{ij}^k = e_{ij}^i \left(1 - \frac{e_{k\ell}^k m_{k\ell}^o}{c}\right) \ . \tag{10}$$

In (10) e_{ij}^k equals e_{ij}^i times a factor that is between zero and one. The larger the value of $m_{k\ell}^o$, the smaller is the value of e_{ij}^k relative to e_{ij}^i. Thus, another measure of the benefit of per-priority effective bandwidths is how much smaller e_{ij}^k is relative to e_{ij}^i. In cases when e_{ij}^k is close to e_{ij}^i, the complexity of using distinct effective-bandwidths probably outweighs the potential efficiency gains.

¿From equations (6) and (8), we obtain all the effective-bandwidth parameters e_{ij}^k with $i \le k$. We have obtained these parameters by exploiting the linearity of the constraint set (5). Given this linearity, it suffices to consider only priority-i type-j connections when we determine the effective-bandwidth parameters e_{ij}^i via (6). Similarly, for $i < k$, it suffices to consider only priority-i type-j connections and priority-k type-ℓ connections for any ℓ when we determine the effective-bandwidth parameters e_{ij}^k via (8). A significant point is that we need consider only two connection types in this calculation. To determine e_{ij}^k, we consider priority-i type-j connections and priority-k type-ℓ connections for some (any) ℓ.

Since the linear admissible set (5) is only an approximation, we might not actually want to fit the parameters by considering connections at their upper and lower limits. Instead, we might want to exploit knowledge of the typical operating region and determine a linear approximation to a more accurate admissible set by constructing a linear hyperplane tangent to the boundary for each priority class, as indicated at the end of Section 1. This observation applies to the determination of both e_{ij}^i and e_{ij}^k for $k > i$. For example, the more accurate admissible set might be determined by simulation, perhaps using source traces, or by system measurements.

The analysis so far indicates essential properties of the approximating linear admissible set. First, we should have one linear constraint (hyperplane) for each priority class, as given in (5). Moreover, assuming that higher-priority performance criteria are always tighter than lower-priority performance criteria, we should have the effective-bandwidth parameters (coefficients in the linear inequalities) ordered by

$$e_{ij}^k > e_{ij}^{k+1} \tag{11}$$

for all priority classes i and k and connection types j with $i \leq k$, extending inequality (3). In other words, the linear admissible set should have the form (5), which means I equations with $e_{ij}^k = 0$ for $k < i$, and the coefficients should be ordered as in (11).

4.2. Large-Buffer Asymptotics

Another way to obtain effective-bandwidth parameters, or to obtain a first cut on them, is to exploit large-buffer asymptotics, which involves the mathematical theory of large deviations. Specifically, we consider the limiting exponential decay rate of the steady-state buffer-content distribution in an unlimited-capacity buffer model. This mathematical framework has been very useful because with the FIFO discipline it provides a setting in which the linear admissible set in (1) is correct. More precisely, the admissible set is asymptotically correct as the buffer size increases (and the associated tail probability decreases) in the performance criterion. Moreover, it enables us to obtain relatively simple formulas for the effective-bandwidth parameters, assuming quite realistic stochastic models for the packet streams generated by the active connections.

Unfortunately, however, when priority classes are introduced, the large-buffer asymptotics no longer produces a linear admissible set (see [1]). However, numerical experience indicates that it is often reasonable to approximate the nonlinear admissible set obtained from the large-buffer asymptotics with priorities by a linear admissible set. Moreover, there is a nice physical interpretation for the linear approximation. As before, we use a different performance criterion for each priority class. When considering priority class i, the initial constraint would be on the workload of only priority class i. The approximation is to use an upper bound, considering instead the total workload for all connections of the first i priority classes. The approximation clearly has the desirable property of being conservative.

At first glance, this upper bound may seem far too crude, but since each successive lower priority class should have a substantially looser performance criterion, this modified criterion becomes intuitively plausible. More important, numerical experience indicates that it usually is an excellent approximation to the exact admissible set obtained from the large-buffer asymptotics. When there is substantial error in the (final) admissible set produced with this approach, it is usually due to the large-buffer asymptotics, not this approximation [1].

Thus, when we consider priority class i, we use a performance criterion based on the steady-state workload from all connections from the first i priorities. This implies that the problem for each priority class reduces to the previously considered FIFO problem. We obtain the effective bandwidths e_{ij}^k for $i \leq k$ by considering the FIFO problem involving the first k priority classes and the class-k performance criterion.

It thus remains to summarize the effective bandwidth formulas for the FIFO discipline. With the FIFO discipline the notion of effective bandwidth is based on the steady-state buffer-content in an unlimited-capacity buffer. We are given a performance criterion

$$P(B \geq b) \leq p \tag{12}$$

based on parameters b and p. Of course, the steady-state buffer content B depends on the connections present. We assume that the steady-state buffer-content distribution has an exponential tail, i.e.,

$$P(B \geq b) \approx e^{-\eta b} \ , \tag{13}$$

which is asymptotically correct as $b \to \infty$, where the decay rate η depends on the active sources. Then the effective bandwidth of a source i turns out to be

$$e_i = \psi_{A_i}(\eta^*)/\eta^* \ , \quad \text{where } \eta^* = -(\log p)/b \tag{14}$$

and $\psi_A(\theta)$ is the asymptotic-decay-rate function (also known as cumulant generating function).

$$\psi_{A_i}(\theta) = \lim_{t\to\infty} t^{-1} \log E e^{\theta A_i(t)} \ , \tag{15}$$

where $A_i \equiv \{A_i(t) : t \geq 0\}$ is the cell arrival process for source i, i.e., $A_i(t)$ is the input during the interval $[0, t]$.

For practical purposes, it is important that $\psi_A(\theta)$ can be calculated for many stochastic processes. For example, suppose that source i is an on-off two-state Markov modulated Poisson process

(MMPP), having arrival rate λ_1 in the on state, mean on time r_1^{-1} and mean off time r_2^{-1} and where each arrival adds one unit of work (corresponding to a contant-size ATM cell). Then the asymptotic-decay-rate function is

$$\psi_{A_i}(\theta) = \frac{\sqrt{\alpha^2 + 4\lambda_1 r_2 (e^\theta - 1)} - \alpha}{2\theta} \ , \tag{16}$$

where $\alpha = r_1 + r_2 - \lambda_1(e^\theta - 1)$. If source i were a Poisson process with rate λ and where again each arrival adds one unit of work, then $\psi_{A_i}(\theta)$ is simply $\lambda(e^\theta - 1)$.

The extension of the above formulation to incorporate priority service is simple: For all connections subject to the priority-k constraint, use the priority-k performance criterion to determine η^* in (14). Thus, we let the effective bandwidth e_{ij}^k of a priority-i type-j connection as seen by priority k, with $i \le k$, be

$$e_{ij}^k = \psi_{A_{ij}}(\eta_k^*)/\eta_k^* \ , \quad \text{where } \eta_k^* = -(\log p_k)/b_k \tag{17}$$

and

$$\psi_{A_{ij}}(\theta) = \lim_{t\to\infty} t^{-1} \log E e^{\theta A_{ij}(t)} \ , \tag{18}$$

with η_k^* representing the priority-k performance constraint and $A_{ij} \equiv \{A_{ij}(t) : t \ge 0\}$ being the input process for a priority-i type-j connection.

4.3. Using a Standardized Traffic Descriptor

Consider Variable-Bit-Rate (VBR) ATM connections for which the Sustainable-Cell-Rate (SCR) traffic descriptor is specified [17]. The SCR constitutes an upper bound on the mean rate of the connection. Suppose that in the FIFO context the effective bandwidth for these connections is chosen to be some factor times the connection's SCR. Thus, for this subsection let e_n represent the effective bandwidth of the n^{th} connection established on the link, and let SCR_n denote the SCR for this connection. Then

$$e_n = \alpha \cdot SCR_n \ , \tag{19}$$

where the factor α is determined from historical measurements of realized connections. A conservative value for α might be picked initially, and then subsequently reduced as long as the performance commitment for the connections continues to be met. Given that there are N connections established on the link, equation (1) would now have the form:

$$\sum_{n=1}^{N} e_n = \sum_{n=1}^{N} \alpha \cdot SCR_n \le c \ . \tag{20}$$

Note that in this case connections are not grouped in "types."

To extend this method to account for priorities, multiple factors α are determined. Again for this subsection, let e_{in}^k represent the effective bandwidth of the n^{th} connection established at priority i, as seen by priority k, where:

$$e_{in}^k = \alpha_i^k \cdot SCR_n \ , \tag{21}$$

for chosen factors α_i^k, where $\alpha_i^{k+1} > \alpha_i^k$. Likewise, given N_i connections are established on the link at priority i, equation (5) has the form:

$$\sum_{i=1}^{k} \sum_{n=1}^{N_i} e_{in}^k = \sum_{i=1}^{k} \sum_{n=1}^{N_i} \alpha_i^k \cdot SCR_n \le c, \quad k = 1, \ldots, I \ . \tag{22}$$

5. Numerical Examples

In this section we present four examples to illustrate the benefits from using the per-priority effective bandwidths, including a case where the benefits are only marginal. A different issue is the accuracy of the effective-bandwidth approximations, of whatever type, as compared with the "exact" calculation of the admissible set. We do not pursue that issue here, but a detailed discussion is given in [1].

For the examples, we use effective bandwidths based on large-buffer asymptotics, as discussed in Section 4.2. These large-buffer asymptotics illustrate a range of possible results, and the calculations can be independently checked (whereas examples based on the other two methods in Section 4 would depend on unstated measurement studies for the key parameters $m_{k\ell}^o$ and α_i^k). For simplicity, we consider two priorities and one type of connection in each priority. For the first example we start with the simple case in which the connections are the same for each priority. (This could represent the case in which some users are given better service for a higher price.) Suppose that the connections are on-off two-state Markov modulated Poisson processes (MMPPs), where each arrival offers one unit of work (corresponding to one ATM cell), as in Section 4.2. Suppose that the mean rate is 0.01, the fraction of time on is 0.1 and the mean burst size is 20. Let the performance parameters be: $b_1 = 500$, $b_2 = 5,000$, and $p_1 = p_2 = 10^{-6}$. Lastly, let the link bandwidth be 1, which is is 100 times the mean rate. For these parameters, the effective bandwidths are $e_1 = 0.0174$, and $e_1^2 = e_2 = 0.0105$. Note that since the connection type is the same for both priorities, e_1^2 equals e_2. Also note that e_1 is larger than e_1^2, and that the priority-2 performance-criterion parameters are qualitatively looser than priority-1's.

The admissible set for three cases is given in Figure 5. The smallest admissible set (dashed line) labeled "eff.-bdwth. FIFO service" assumes priority service has *not* been implemented, the service discipline is FIFO, and the stricter performance criterion applies to all connections. In this case, e_2 would equal e_1, which is 0.0174, and the admissible set is given by (1) with $I = 2$. The middle admissible set (dotted line) labeled "priority-insensitive eff.-bdwth." assumes priority service *has* been implemented, and the looser performance criterion applies to priority 2, but just one effective bandwidth, e_1, is used for the priority-1 connections. Again the admissible set is given by (1) with $I = 2$. The largest admissible set (solid line) labeled "per-priority eff.-bdwth." uses two effective bandwidths, e_1 and e_1^2, for the priority-1 connections, and the admissible set is given by (2).

The main point of Example 1 is that at higher occupancies of priority 1, the admission of priority-2 connections is needlessly limited if the effective bandwidths are not adjusted for priorities. For example, when n_1 is 50, the priority-insensitive effective bandwidths limit n_2 to 12, whereas, n_2 would be 45 with per-priority effective bandwidths. If no priorities were used, then n_2 is 7. Note that half of the potential gain (measured in terms of area of admissible sets) from implementing priorities is not realized if the effective bandwidths are priority-insensitive.

Example 2 is the same as Example 1 except the priority-2 connections are more bursty: The mean burst size is changed from 20 to 100. Then e_2 changes from 0.0105 to 0.0128. The resulting admissible sets are shown in Figure 6. As in Example 1, at higher occupancies of priority 1, the admission of priority-2 connections is needlessly limited if the effective bandwidths are not adjusted for priorities. Also, in the present example we see more gain from the implementation of priorities, than in Example 1. For instance, in Example 2 when n_1 is small, say zero, the looser criterion used for priority-2 allows 5.2 times more connections to be admitted as compared with FIFO service. In Example 1, this factor was "only" 1.7.

The occupancy on the link due to the priority-1 connections influences the potential gain from the per-priority effective bandwidths. In Examples 1 and 2, when the number of priority-1 connections admitted is the maximum possible, and no priority-2 connections are present, the occupancy is 57%. Example 3 considers the case where this maximum priority-1 occupancy is lower, and Example 4 considers the case where it is higher.

Example 3 is the same as 2 except the priority-1 performance criterion is tighter: b_1 is reduced from 500 to 200, which is still ten times greater than the mean burst size. The resulting admissible sets are given in Figure 7. In Example 3 the maximum number of priority-1 connections admissible is 23, and thus the maximum priority-1 occupancy is only 23%. Here we see a very strong advantage of using per-priority effective bandwidths. For instance, when n_1 is 20, the priority-insensitive effective bandwidths restrict n_2 to 11, whereas, n_2 would be 61, five and half times greater, with per-priority effective bandwidths.

In Example 4, we consider the case where the higher-priority queue contains the superposition of constant bit rate ATM connections. We model this superposition as a Poisson process, where each arrival offers one unit of work. (If the ATM connections have not been jittered, then the Poisson assumption is conservative.) Let the priority-2 connections be the same as in Example 1: each connection is a two-state on-off MMPP with mean rate 0.01, fraction of time on 0.1, and mean burst size 20. Let the performance parameters be $b_1 = 100$, $b_2 = 5,000$, $p_1 = 10^{-9}$, and $p_2 = 10^{-6}$. The admissible sets are given in Figure 8. Here the maximum priority-1 occupancy is 90%, which is higher than in the previous examples, and the gain from the per-priority effective bandwidths is rather small. Although, at the larger priority-1 occupancies, we still see some gain. At 80% priority-1 occupancy, the priority-insensitive effective bandwidths restrict n_2 to 10, whereas, n_2 would be 18 with per-priority effective bandwidths. Overall, the additional complexity may outweigh the benefit in this example. As indicated before, we could then elect to set e_1^k equal to e_1^1 in each priority constraint of a multiple priority system, such as (4).

6. Conclusions

Our main conclusion is that to realize the gains from implementing service priorities at network nodes, the connection admission control and dimensioning policies using effective bandwidths should be revised. A given connection should be associated with multiple effective bandwidths: one corresponding to the priority level of the given connection and (potentially) one for each of the lower-level priorities.

It should be noted that for some service types, distinct effective bandwidths for all lower priorities may yield only modest efficiency gains, in which case to reduce complexity a given priority-i type-j connection would have the same value for the effective bandwidth e_{ij}^k for different priority levels k.

We have indicated how the per-priority effective bandwidths can be determined in Section 4. First, these effective bandwidths can be estimated by measuring the admissible set, considering two connection types at a time. Second, these effective bandwidths may be computed by a minor modification of the now-familiar large-buffer asymptotics associated with the FIFO discipline, as reviewed in Section 4.2. For that purpose, we focus on the steady-state workload of the first i

priority classes, for each i. Third, effective bandwidths may be scaled versions of the sustainable cell rate (SCR) parameter. The general approach also allows effective bandwidths to be obtained in other ways.

Constructing the new admissible set with priorities shows the advantage of priorities when lower-priority classes have substantially looser performance criteria, because we can see that the admissible set is much larger than without priorities.

References

[1] A. W. Berger and W. Whitt, "Effective bandwidths with priorities and loss criteria," AT&T Labs, 1997, submitted for publication.

[2] C. S. Chang, "Stability, queue length and delay of deterministic and stochastic queueing networks," *IEEE Trans. Automat. Control*, vol. 39, pp. 913–931, 1994.

[3] C. S. Chang and J. A. Thomas, "Effective bandwidths in high-speed digital networks," *IEEE J. Sel. Areas Commun.*, vol. 13, pp. 1091–1100, 1995.

[4] C. S. Chang and T. Zajic, "Effective bandwidths of departure processes from queues with time varying capacities," *Proc. IEEE Infocom '95*, pp. 1101–1109, 1995.

[5] G. L. Choudhury, K. K. Leung and W. Whitt, "An algorithm to compute blocking probabilities in multi-rate multi-class multi-resource loss models," *Adv. Appl. Prob.*, vol. 27, pp. 1104–1143, 1995.

[6] G. L. Choudhury, K. K. Leung and W. Whitt, "An inversion algorithm to compute blocking probabilities in loss networks with state-dependent rates," *IEEE/ACM Trans. Networking*, vol. 3, pp. 585–601, 1995.

[7] G. L. Choudhury, K. K. Leung and W. Whitt, "Efficiently providing multiple grades of service with protection against overloads in shared resources," *AT&T Technical Journal*, vol. 74, pp. 50–63, 1995.

[8] G. L. Choudhury, D. M. Lucantoni and W. Whitt, "Squeezing the most out of ATM," *IEEE Trans. Commun.*, vol. 44, pp. 203–217, 1996.

[9] G. de Veciana and G. Kesidis, "Bandwidth allocation for multiple qualities of service using generalized processor sharing," *IEEE Trans. on Information Theory*, vol. 42, pp. 268–272, 1996.

[10] G. de Veciana, G. Kesidis and J. Walrand, "Resource management in wide-area ATM networks using effective bandwidths," *IEEE J. Sel. Areas Commun.* vol. 13, pp. 1081–1090, 1995.

[11] A. I. Elwalid and D. Mitra, "Fluid models for the analysis and design of statistical multiplexing with loss priorities on multiple classes of bursty traffic, *Proc. IEEE Infocom '92*, pp. 415–425, 1992.

[12] A. I. Elwalid and D. Mitra, "Effective bandwidths of general Markovian traffic sources and admission control of high speed networks," *IEEE/ACM Trans. Networking*, vol. 1, pp. 329–343, 1993.

[13] A. I. Elwalid and D. Mitra, "Analysis, approximations and admission control of a multi-service multiplexing system with priorities," *Proc. IEEE Infocom '95*, pp. 463–472, 1995.

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

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

[16] J. Y. Hui, "Resource allocation for broadband networks," *IEEE J. Sel. Areas Commun.*, vol. SAC-6, pp. 1598–1608, 1988.

[17] Inter. Telecommunications Union (ITU), "Traffic control and congestion control in B-ISDN," ITU-T Recommendation I.371, Geneva, May, 1996.

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

[19] F. P. Kelly, "Notes on effective bandwidths," in *Stochastic Networks*, Clarendon Press, Oxford, pp. 141–168, 1996.

[20] G. Kesidis, J. Walrand and C. S. Chang, "Effective bandwidths for multiclass Markov fluids and other ATM sources," *IEEE/ACM Trans. on Networking*, vol. 1, pp. 424-428, 1993.

[21] V. G. Kulkarni, L. Gun and P. F. Chimento, "Effective bandwidth vectors for multiclass traffic multiplexed in a partitioned buffer," *IEEE J. Sel. Areas Commun.*, vol. 13, pp. 1039–1047, 1995.

[22] K. W. Ross, *Multiservice Loss Models for Broadband Telecommunications Networks*, Springer-Verlag, London, 1995.

[23] W. Whitt, "Tail probabilities with statistical multiplexing and effective bandwidths in multi-class queues," *Telecommunication Systems*, vol. 2, pp. 71–107, 1993.

[24] J. Zhang, "Performance study of Markov modulated fluid flow models with priority traffic," *Proc. IEEE Infocom '93*, pp. 10–17, 1993.