Chapter 5

# Part Two ATM Traffic Management and Control

A Comparative Performance Analysis of Call Admission Control Schemes in ATM Networks

Khaled Elsayed

Department of Electronics and Communications Engineering Faculty of Engineering, Cairo University, Giza, Egypt 12613 E-mail: khaled@ieee.org

Harry Perros

Department of Computer Science North Carolina State University, Raleigh, NC 27695, USA E-mail: hp@eos.ncsu.edu

- Key words: Call Admission Control, Traffic Management, ATM Networks, Quality of Service, Effective Bandwidth, Diffusion Approximation, PGPS, EDF, SP
- Abstract: Connection Admission Control (CAC) is one of the primary mechanisms for preventive congestion control and bandwidth allocation in ATM networks. A substantial number of CAC schemes have been proposed. In this paper, we review the salient features of some of these algorithms. We also provide a comparative study of the performance of CAC schemes devised to meet certain quality of service requirements expressed in terms of cell loss probability and maximum delay.

113

## 1. INTRODUCTION

In recent years, there has been a tremendous growth in the development and deployment of ATM networks. One area which is of significant importance to ATM networks is traffic management. Congestion control is one of the primary mechanisms for traffic management. The primary role of a network congestion control procedure is to protect the network and the user in order to achieve network performance objectives and optimize the usage of network resources. In ATM-based B-ISDN, congestion control should support a set of ATM quality-of-service classes sufficient for all foreseeable B-ISDN services.

Congestion control schemes can be classified into preventive control and reactive control. In preventive congestion control, one sets up schemes which prevent the occurrence of congestion. In reactive congestion control, one relies on feedback information for controlling the level of congestion. Both approaches have advantages and disadvantages. In ATM networks, a combination of these two approaches is currently used in order to provide effective congestion control. For instance, CBR and VBR services use preventive schemes and ABR service is based on a reactive scheme.

Preventive congestion control involves the following two procedures: connection admission control (CAC) and bandwidth enforcement. ATM is a connection-oriented service. Before a user starts transmitting over an ATM network, a connection has to be established. This is done at connection setup time. The main objective of this procedure is to establish a path between the sender and the receiver. This path may involve one or more ATM switches/routers. On each of these ATM switches, resources have to be allocated to the new connection.

The connection set-up procedure runs on a resource manager (which is typically a workstation attached to the switch). The resource manager controls the operations of the switch, accepts new connections, tears down old connections, and performs other management functions. If a new connection is accepted, bandwidth and/or buffer space in the switch is allocated for this connection. The allocated resources are released when the connection is terminated.

Call admission control deals with the question as to whether a switch can accept a new connection or not. Typically, the decision to accept or reject a new connection is based on the following two questions:

- 1. Does the new connection affect the quality-of-service of the connections that are currently being carried by the switch?
- 2. Can the switch provide the quality-of-service (QOS) requested by the new connection?

The answer to these questions is a function of the connections' traffic characteristics, the QOS requested, and the network state.

Call admission control schemes have been developed so that they satisfy a particular quality of service. In packet networks, the two major QOS attributes are packet loss and packet delay. A new connection may request from the network a certain bound on packet loss and packet delay. Moreover, these bounds can be deterministic or statistical. For deterministic QOS, a new connection would request a maximum end-to-end packet/cell delay or a maximum threshold on the value of packet/cell loss probability. On the other hand, for statistical QOS, a connection would request that its packets experience, for example, a mean end-to-end delay or a mean packet/cell loss probability.

Call admission control schemes may be classified into a) non-statistical allocation, or peak bandwidth allocation, and b) statistical allocation. Non-statistical allocation can be used to enforce deterministic bounds on the requested QOS of a connection. Statistical allocation can be used to enforce either deterministic or statistical QOS bounds. Below we examine the two types of call admission control. The advantage of peak bandwidth allocation is that it is easy to decide whether to accept a new connection or not. The disadvantage of peak allocation is that unless connections transmit at peak rates, the output port link may be grossly under-utilized.

In statistical allocation, bandwidth for a new connection is not allocated on per peak rate basis. Rather, the allocated bandwidth is less than the peak rate of the source. As a result, the sum of all peak rates may be greater than the capacity of the output link. Statistical allocation makes economic sense when dealing with bursty sources, but it is difficult to carry out effectively. This is because of difficulties in characterizing the arrival process of ATM cells and the lack of understanding as to how this arrival process is shaped deep in the ATM network.

Another difficulty in designing a connection admission control algorithm for statistical allocation is that decisions have to be done on the fly, and therefore they cannot be CPU intensive. Typically, the problem of deciding whether to accept a new connection or not may be formulated as a queueing problem. The connection admission control algorithm has to be applied to the buffer of each output port. If we isolate an output port and its buffer from the rest of the switch, we will obtain the queueing model shown in figure 1. This type of queueing structure is known as an ATM multiplexer. It represents a number of ATM sources feeding a finite capacity queue, which is served by a server (the output port). The service time is constant equal to the time it takes to transmit an ATM cell.



Figure 1. An ATM multiplexer

Now, let us consider the cell loss probability as the requested QOS, and let us assume that the cell loss probability of the existing connections is satisfied. The question that arises is whether the cell loss probability will still be maintained if the new connection is added. This can answered by solving this ATM multiplexer with the existing and new connections. However, the solution to this problem is very difficult and CPU intensive (see for example Elsayed and Perros [8] and Li [17]). It gets even more complicated if we assume complex arrival processes. In view of this, a variety of different bandwidth allocation algorithms have been proposed which are based on different approximations, or different types of schemes which do not require the solution of such a queueing problem.

In this paper, we will examine some of the connection admission control algorithms that have been proposed for statistical allocation. Before we proceed, however, we review briefly the various traffic models that have been proposed in the literature.

#### **1.1** Characterization of an arrival process

Prior to the advent of ATM networks, performance models of telecommunication systems were typically developed based on the assumption that arrival processes are Poisson distributed. That is, the time between successive arrivals is exponentially distributed. In some cases, such as in public PSTN switching, extensive data collection actually supported the Poisson assumption.

Over the last few years, we have gone through several paradigm shifts regarding our understanding as to how to model an ATM source. Following the first performance models which were based on the Poisson assumption, or the Bernoulli assumption, it became apparent that these traffic models did not capture the notion of burstiness that is present in traffic resulting from applications such as moving a data file and packetized encoded video. Thus, there was a major shift towards using arrival processes of the on/off type.

117

The ATM Forum has defined a standard mechanism for specifying a connection's traffic [1]. A connection is specified by the tuple (PCR, CDVT, SCR, MBS) where PCR is the peak cell rate, CDVT is the cell delay variation tolerance, SCR is the sustainable cell rate, and MBS is maximum burst size. Using the peak rate and the cell delay variation, one can effectively police the peak rate. Also, using the maximum burst length, one can estimate a cell delay variation that can be used to police the sustainable rate. These parameters can be enforced using the GCRA algorithm of the ATM forum, which is equivalent to a dual leaky-bucket mechanism [1].

Most of the CAC schemes use the tuple of parameters (PCR, SCR, MBS) of the existing and new connections when making a decision on accepting or rejecting a connection. The parameter CDVT is a function of the user and network equipment and has little effect on traffic characterization of the connection. The tuple (PCR, SCR, MBS) can be used to specify a variety of traffic models. A model that introduces statistical variation into the model specified by (PCR, SCR, MBS) is the on/off source model. A popular instance of on/off sources is the Interrupted Poisson Process (IPP) or its discrete-time counterpart the Interrupted Bernoulli Process (IBP). In an IPP, there is an active period during which arrivals occur in a Poisson fashion, followed by an idle period during which no arrivals occur. These two periods are exponentially distributed, and they alternate continuously. An IBP is defined similarly, only the arrivals during the active period are Bernoulli distributed, and the two periods are geometrically distributed. Another way of describing a source is using the fluid approach. Here arrivals occur with a continuous rate during the active period. This defines an on/off fluid source or equivalently an Interrupted Fluid Process (IFP).

#### 1.2 Classification of Connection Admission Schemes

In this paper we consider two main categories of CAC schemes: a) schemes for bounding cell loss probability for connections, and b) schemes for bounding cell delay. A variety of different connection admission schemes have been proposed in the literature. Some of these schemes require an explicit traffic model and some only require traffic parameters such as the peak and average rate. In this paper we review some of these schemes. For presentation purposes, the schemes have been classified as follows:

• CAC schemes based on the cell loss probability. These include

- 1. Effective Bandwidth (Equivalent Capacity)
- 2. Diffusion Approximation
- 3. Upper Bounds of the cell loss probability

- CAC schemes based on cell delay. These are usually associated with certain scheduling methods. Our study includes
  - 1. Weighted Fair Queueing (WFQ) or Packet-by-Packet Generalized Processor Sharing (PGPS) scheduling
  - 2. Delay-Earliest Deadline First (EDF) scheduling
  - 3. Static Priority (SP) scheduling

This classification was based on the underlying principle that was used to develop a CAC scheme and its targeted QOS objective. The remainder of this paper is organized as follows. In section 2, we review the salient features of the four CAC schemes mentioned above that are based on the cell loss probability. Extensive numerical comparisons between three of these schemes are then given in subsections 2.5 to 2.11. In section 3, we review the CAC schemes mentioned above that are based on the cell delay. Numerical comparisons between three of these schemes are given in section 3.5. Other CAC schemes are described in section 4.

# 2. CAC SCHEMES FOR THE CELL LOSS PROBABILITY QOS

### 2.1 Effective Bandwidth/Equivalent Capacity

Let us consider a single source feeding a finite capacity queue. Then, the effective bandwidth of the source is the service rate of the queue that corresponds to a cell loss of e. The effective bandwidth for a single source can be derived as follows (see Guerin, Ahmadi, and Naghshineh [14]). Each source is assumed to be an IFP. Let R be its peak rate, r the fraction of time the source is active, and b the mean duration of the active period. An IFP source can be completely characterised by the vector (R, r, b). Let us now assume that the source feeds a finite capacity queue with constant service time. Let K be the capacity of the queue. The effective bandwidth e is given by:

$$e = \frac{a - K + \sqrt{(a - K)^2 + 4Kar}}{2a}R$$
(1)  
where  $a = \ln(1/e)b(1 - r)R$ .

In the case of N sources, and given that the buffer has a capacity K, the effective bandwidth is again the service rate e which ensure that the cell loss

119

for all sources is less than or equal to *e*. Guerin, Ahmadi, and Naghshineh [14] proposed the following approximation for multiple sources:

$$e = \min(\mathbf{r} + a'\mathbf{s}, \sum_{i=1}^{N} e_i)$$
(2)

where  $e_i$  is the effective bandwidth of the *i*th source calculated using expression (1), and  $\sum_{i=1}^{N} e_i$  is the sum of all the individual effective

bandwidths,  $\mathbf{r}$  is the total average bit rate, i.e.  $\mathbf{r} = \sum_{i=1}^{N} \mathbf{r}_{i}$ , where

 $\mathbf{r}_i = r_i R_i$  is the mean bit rate of the *i*th source,  $\mathbf{s}^2 = \sum_{i=1}^N \mathbf{s}_i^2$ , where  $\mathbf{s}_i^2$  is the variance of the bit rate of the *i*th source,  $\mathbf{s}_i^2 = \mathbf{r}_i (R_i - \mathbf{r}_i)$ , and  $a' = \sqrt{-2\ln(\mathbf{e}) - \ln 2\mathbf{p}}$ .

Some studies (see Choudhury, Lucantoni, and Whitt [4] and Elsayed and Perros [8]) have clearly indicated the inaccuracy of effective bandwidth methods in some situations. In particular, the effective bandwidth method fails when a bufferless system subject to an input traffic has a small probability that the traffic load exceeds the link capacity. In the effective bandwidth approach, this probability is assumed to be close to one (and is taken as one in the calculations). Rege [22] compares various approaches for effective bandwidth and proposes some modifications to enhance the accuracy of the scheme. Elwalid et al. [9] proposed a method in which they combined Chernoff bounds with the effective bandwidth. This method permits better accuracy than effective bandwidth for the case of a bufferless (or a small buffer for that matter) multiplexer that can achieve substantial statistical gain. However, in some other cases, the method does not improve the accuracy of the effective bandwidth.

Kulkarni, Gun, and Chimento [15] considered the effective bandwidth vector for two-priority on/off source. Chang and Thomas [3] introduced a calculus for evaluating source effective bandwidth at output of multiplexers and upon demultiplexing or routing. On-line evaluation of effective bandwidth have been proposed by De Veciana, Kesidis and Walrand [24]. Duffield et al. [7] proposed maximum entropy as a method for characterizing traffic sources and their effective bandwidth.

#### 2.2 Diffusion Approximation

Gelenbe, Mang and Onvural [12] proposed a scheme that uses statistical bandwidth obtained from a closed-form expression based on the diffusion approximation models. Specifically, a diffusion process with absorbing boundaries and jumps was used to analyze approximately a discrete-time ATM multiplexer with N IFP sources. Two models are used: one for a finite buffer (FB) ATM multiplexer and the other for an infinite buffer (IB) ATM multiplexer. In the IB model, the cell loss probability is estimated by the overflow probability, which is the overall probability of exceeding the actual buffer capacity (K) in a system with an unlimited buffer size. The cell loss probability calculated from these two models is:

$$L_{FB} = \frac{1}{\sqrt{2p}} e^{\frac{2K}{a}(r-C)} e^{-\frac{(r-C)^2}{2s^2}}$$
(3)

$$L_{IB} = \frac{s}{r\sqrt{2p}} e^{\frac{2K}{a}(r-C)} e^{-\frac{(r-C)^2}{2s^2}}$$
(4)

For *N* IFP sources with parameters ( $R_i$ ,  $r_i$ ,  $b_i$ ), we have  $\mathbf{s}^2 = \sum_{i=1}^{N} \mathbf{s}_i^2$  is

the total variance where  $\mathbf{s}_i^2 = \mathbf{r}_i(R_i - \mathbf{r}_i)$  and  $\mathbf{r}_i = r_i R_i$ ,  $\mathbf{r} = \sum_{i=1}^N \mathbf{r}_i$ ,

 $\boldsymbol{a} = \sum_{i=1}^{N} \boldsymbol{r}_{i} C V_{i}^{2} \text{ is the instantaneous variance of the arrival process where}$  $C V_{i}^{2} = \frac{1 - (1 - \boldsymbol{b}_{i} T_{i})^{2}}{(\boldsymbol{b} T_{i} + \boldsymbol{\sigma}_{i} T_{i})^{2}} \text{ and } T_{i} = \frac{1}{R_{i}}, \quad \frac{1}{\boldsymbol{b}_{i}} = b_{i} \text{ the mean on period, and}$ 

 $\frac{1}{g_i}$  is the mean off period of the *i*th source. Finally,  $C^{-1}$  is the time required to

transmit one cell.

Let us define the statistical bandwidth as the bandwidth that needs to be allocated for the multiplexed connections in order to keep the cell loss probability below e (the required cell loss probability). We get two expressions (one for the FB and the other for the IB model respectively) for the statistical bandwidth:

$$C_{FB} = \boldsymbol{r} - \boldsymbol{d} + \sqrt{\boldsymbol{d}^2 - 2\boldsymbol{s}^2 \boldsymbol{w}_1}$$
(5)

$$C_{IB} = \boldsymbol{r} - \boldsymbol{d} + \sqrt{\boldsymbol{d}^2 - 2\boldsymbol{s}^2 \boldsymbol{w}_2}$$
(6)

where  $\boldsymbol{d} = \frac{2K}{\boldsymbol{a}}\boldsymbol{s}^2$ ,  $\boldsymbol{w}_1 = \ln(\boldsymbol{e} \sqrt{2\boldsymbol{p}})$ , and  $\boldsymbol{w}_2 = \ln(\boldsymbol{e} \boldsymbol{r} \sqrt{2\boldsymbol{p}}) - \ln(\boldsymbol{s})$ . It is possible to take:

$$C_{df} = \max(C_{FB}, C_{IB})$$

as the (worst-case) estimate of the statistical bandwidth. The procedure to admit or reject a new connection is then summarized as follows:

- 1) At any time keep a record of the quantities  $\mathbf{r} = \sum \mathbf{r}_i, \mathbf{s}^2 = \sum \mathbf{s}_i^2, \mathbf{a} = \sum \mathbf{r}_i C V_i^2$  of the existing connection 2) When a new connection arrives, update  $\mathbf{r} = \sum \mathbf{r}_i, \mathbf{s}^2 = \sum \mathbf{s}_i^2, \mathbf{a} = \sum \mathbf{r}_i C V_i^2$  to include the new connection
- 3) If the resulting  $C_{df} < C$ , then admit the new connection.
- 4) Else reject the new connection and update  $\mathbf{r} = \sum \mathbf{r}_i, \mathbf{s}^2 = \sum \mathbf{s}_i^2, \mathbf{a} = \sum \mathbf{r}_i C V_i^2$  to exclude the rejected connection effect.

### 2.3 Upper Bounds of the Cell Loss Probability

Several connection admission schemes have been proposed which are based on an upper bound for the cell loss probability. Saito [23] proposed an upper bound based on the average number of cells that arrive during a fixed interval (*ANA*), and the maximum number of cells that arrive in the same fixed interval (*MNA*). The fixed interval was taken to be equal to D=2, where D is the maximum admissible delay in a buffer. Using these parameters, the following upper bound was derived. Let us consider a link serving Nconnections, and let  $p_i(j), i = 1, 2, \dots, N$ , and  $j = 1, 2, \dots$ , be the probability that j cells belonging to the *i*th connection arrive during the period D=2. Then, the cell loss probability *CLP* can be bounded by:

$$CLP \le B(p_1, p_2, \dots, p_N; D/2) = \frac{\sum_{k=0}^{\infty} [k - D/2]^+ p_1 * p_2 * \dots p_N(k)}{\sum_{k=0}^{\infty} k p_1 * p_2 * \dots p_N(k)}$$

where \* is the convolution operation. Let  $q_i(j)$  be the following functions:

$$\boldsymbol{q}_{i}(j) = \begin{cases} ANA_{i} / MNA_{i}, & j = MNA_{i}, \\ 1 - ANA_{i} / MNA_{i}, & j = 0, \\ 0, & \text{otherwise.} \end{cases}$$

Then it can be shown that

$$CLP \leq B(p_1, p_2, \dots, p_N; D/2)$$
  
$$\leq B(q_1, q_2, \dots, q_N; D/2)$$
  
$$= \frac{\sum_{k=0}^{\infty} [k - D/2]^* q_1 * q_2 * \dots q_N(k)}{\sum_{k=0}^{\infty} k q_1 * q_2 * \dots q_N(k)}$$

A new connection is admitted if the resulting  $B(p_1, p_2, \dots, p_N; D/2)$  is less than the allowable cell loss probability. Saito proposes a scheme for calculating  $q_1 * q_2 * \cdots q_N(k)$  efficiently. He also obtained a different upper bound based on the average and the variance of the number of cells that arrive during D/2.

The main disadvantage of this method is the absence of the burst size in the calculation and thus a worst case behaviour is assumed for the source. This method works well in the case when the actual source behaviour is close to the worst case behaviour assumed in the above calculation.

For other upper bounds on the cell loss probability see Rasmussen et al. [21], Castelli, Cavallero, and Tonietti [2], Doshi [6] and the closely related work by Elwalid, Mitra, and Wentworth [10].

# 2.4 Comparative Performance Analysis of the lossoriented CAC Schemes

In this section, we provide a numerical comparison among the following CAC schemes: a) the method proposed by Guerin, Ahmadi, and Naghshineh [14] for calculating the effective bandwidth (hereafter referred to as the

"equivalent capacity method"), b) the diffusion approximation method proposed by Gelenbe, Mang and Onvural [12] for calculating the statistical bandwidth (hereafter referred to as the "diffusion approximation method), and c) Saito's upper bound of the cell loss probability [23] (hereafter referred to as the CLP upper bound). These schemes were selected since they use the same set/subset of traffic descriptors. Namely, the peak bit rate, mean bit rate, and mean burst length of a call (R; r; b). (Note that the CLP upper bound scheme only utilizes the mean and peak bit rate information.) Before presenting the results, we define some necessary terms.

We will consider an ATM multiplexer consisting of a finite capacity queue of size K. This queue is served by a server (the outgoing link) of capacity C. The connections handled by this are classified into M classes, namely classes 1 through M. In this work, for illustration purposes, we limit M to 2. All the connections in the same class *i* have the same traffic descriptor ( $R_i$ ,  $r_i$ ,  $b_i$ ), where  $R_i$  is the connection's peak rate,  $r_i$  is the connection's average bit rate, and  $b_i$  is the connection's mean burst length.

Admission Region: This is the set of all values of  $(n_1; n_2)$  for which the cell loss probability is less than a small value e, where  $n_i$  is the number of allocated class *i* connections, i = 1; 2. In other words, this is the set of all combinations of the connections from the 2 classes for which the required cell loss probability e is achievable. In the numerical results given below, we obtain the outermost boundary of the region. All points enclosed between the boundary and the axes represent combinations of connections from each class which lie within the admission region.

**Statistical Gain**: Let  $N_{\min_i}$  be the number of class *i* connections admitted using peak rate allocation. So,  $N_{\min_i} = \lfloor C / R_i \rfloor$ . Likewise, define  $N_{\max_i}$  to be the number of class *i* connections that can be admitted using mean rate allocation. So,  $N_{\max_i} = \lfloor C / \mathbf{r}_i \rfloor$ . The statistical gain for a particular traffic class is defined as the maximum number,  $N_i$ , of connections admitted by a CAC scheme divided by the number of connections that can be accepted using peak rate allocation  $(N_{\min_i})$ , i.e.  $N_i / N_{\min_i}$  when a single class of calls is exclusively using the multiplexer. In order for a CAC scheme to be effective it should be able to provide some statistical gain when possible, i.e. achieve  $N_i / N_{\min_i} > 1$ .

Each of the three CAC schemes was implemented separately. The performance of these schemes relative to each other was compared for various regions of input traffic parameters, buffer size, and required cell loss probability. Also, operating regions for which a particular scheme provides statistical gain over peak rate allocation were identified. We fixed the link

speed at 150 Mbps and choose two classes of traffic with parameters given in table 1.

Table 1. Traffic parameters for the two classes

	R (Mbps)	r (Mbps)	b (Cells)
Class 1	10	1	340
Class 2	2	0.1	2600

#### 2.5 Case 1: Relatively Small Buffer Size

We consider the admission control of two classes assuming a relatively small buffer. The system parameters were chosen as follows. We set the required cell loss probability **e** is equal to 10<sup>-6</sup> and the buffer size K equal to 618 cells (32 Kbytes). The minimum,  $N_{\min_i}$ , and maximum number,  $N_{\max_i}$ , of connections for class 1 and 2 are respectively:  $(N_{\min_1}, N_{\max_1}) = (15; 150)$  and  $(N_{\min_2}, N_{\max_2}) = (75; 1500)$ .

The admission regions obtained for the three CAC methods are shown in figure 2.

The diffusion approximation provides the largest admission region for this example. For this method, the statistical gain for classes 1 and 2 is respectively 7.3 and 14.16. For the equivalent capacity method the gain is 6.13 and 11.37 respectively. For the equivalent capacity method, we note that the admission region is approximately bounded by the intersection of two regions bounded by two almost-linear boundaries: one is obtained by the Gaussian approximation and the other by the effective bandwidth calculation (the intersection near the (25,410) point). The CLP upper bound scheme provides a conservative admission regions yielding a statistical gain for classes 1 and 2 of 2.86 and 11.55 respectively. We note that for the case when the majority of connections belong to class 2, the CLP method is superior to equivalent capacity. However, this scheme is in general conservative with respect to the other schemes. It is obvious that for class 2 which has a much smaller mean to peak ratio the achieved gain for any of the methods is much higher than class 1 although it has a much longer on period. In general the larger the ratio of link capacity to the mean rate of connections, the larger the achieved statistical gain.



Figure 2. Admission regions for the CAC schemes, K=618 cells,  $e = 10^{-6}$ 



*Figure 3*. Admission regions for the CAC schemes, K=1236,  $e=10^{-6}$ 

# 2.6 Case 2: Relatively Large Buffer Size

The buffer size K was doubled to 1236 cells (64 Kbytes). The obtained admission regions for the three schemes are shown in figure 3.

Since the buffer size is increased to 1236, the admission region of all schemes increases. The diffusion approximation provides the largest admission region. When a single class share the multiplexer, the statistical gain that the diffusion approximation yields for classes 1 and 2 are respectively 8.4 and 15.75. For the equivalent capacity method the gain is 8

and 13.19 respectively. In this case, only the effective bandwidth calculation affects the admission region of the equivalent capacity.

For the CLP upper bound scheme, we observe that the maximum number of admitted connections from each class does not increase appropriately when doubling the buffer size. The achieved gains are 2.86 and 11.95 for class 1 and 2 respectively. The maximum number for class 1 remains at 43 while the maximum for class 2 increases slightly from 866 to 896. The reason for this is that class 2 has a lower peak rate and average rate than class 1. We note that in order for this scheme to yield a statistical gain, we need to have traffic sources with small peak and average rate relative to the link capacity.



*Figure 4*. Varying the Buffer Size,  $\varepsilon = 10^{-6}$ 

# 2.7 Effect of the Buffer Size

Assuming that only class 1 or class 2 connections are transported, we obtain the maximum number of admitted connections as a function of the buffer size. The buffer size is increased from a value of  $b_i$  /10 to 100  $b_i$ , where  $b_i$  is the mean burst length of class *i*, while the required cell loss probability is *e* fixed at 10<sup>-6</sup>. The results are shown in figure 4. The figure indicates that the diffusion approximation scheme and the equivalent capacity scheme asymptotically admit the same number of connections as the buffer size approaches infinity.

We observe that for small buffer sizes, the equivalent capacity method admits a fixed number of connections obtained through the Gaussian approximation (bufferless approximation). Furthermore, for class 1, the number of connections admitted by equivalent capacity is smaller than those admitted by the CLP upper bound for small buffer sizes.

The CLP upper bound scheme is less sensitive to the increase in buffer size. For this scheme, a temporary drop occurs to the maximum number of connections that can be admitted as the buffer increases. This is due to the effect of dividing ANA by MNA where MNA, a function of the buffer size and peak rate, must be an integer. So, by increasing the buffer size we get different values of ANA/MNA. We note also that increasing the buffer size beyond a specific value does not cause any increase in the number of admitted connections.

# 2.8 Effect of the Required Cell Loss Probability

Assuming that only class 1 or class 2 connections are transported, we obtain the maximum number of admitted connections as a function of the required cell loss probability. We fix the buffer size at 1236 cells and increase the cell loss probability from  $10^{-9}$  to  $10^{-3}$ . The results are shown in figure 5.

From this figure, we observe that for class 1 the diffusion approximation and the equivalent capacity scheme exhibit low sensitivity to the cell loss probability. In this particular example, the buffer size is large enough so that the two schemes admit a large number of connections even for a very small value of the required cell loss probability. For the diffusion approximation, the increase in the cell loss probability caused the maximum number of connections for class 1 to only increase from 118 to 138, not even reaching the maximum number of admittable connections, 150. The equivalent capacity scheme is more sensitive to the required cell loss probability than the diffusion approximation scheme. The maximum number of connections that can be admitted increased from 105 to 136 exhibiting higher sensitivity. This sensitivity is of course a function of buffer size as well. In general both methods become more sensitive when buffer sizes are small.

The CLP upper bound method is the most sensitive to the cell loss probability. In this example, the increase in the maximum number of connections is from 25 to 100 for class 1 and from 740 to 1320 for class 2. Since the sensitivity of the CLP upper bound method to buffer size is small, it seems, that the required cell loss probability affects the admission region and the achievable statistical gain.

Therefore, for the diffusion approximation and the equivalent capacity methods, if the buffer size is large their sensitivity to CLP is small whereas the CLP upper bound scheme is usually quite sensitive to the cell loss probability.



Figure 5. Varying the required CLP e, K = 1236



*Figure 6.* Varying the Activity Ratio r,  $e = 10^{-6}$ 

### 2.9 Effect of the Activity Ratio

In this section, we study the sensitivity of the three CAC schemes to changes in the activity ratio  $r_i = \mathbf{r}_i / R_i$ . Assuming that only class 1 or class 2 connections are transported, we obtain the maximum number of admitted connections as a function of  $r_i$ , as  $r_i$  increases from 0.05 to 0.5. We fix the buffer size at 1236 cells and the required cell loss probability at 10<sup>-6</sup>. The results are shown in figure 6.

We observe a strong dependence of all methods on the activity ratio. For class 1, when the activity ratio is 0.05, the two methods provide the maximum possible admitted number of connections (i.e. 150). The admitted number of connections drops sharply to 28, 25, and 15 respectively for the equivalent capacity, diffusion approximation, and CLP upper bound methods. The same behaviour is also observed in the case of class 2. The sensitivity to the activity ratio is greatest for the diffusion approximation and it is larger for the class with the smaller peak rates. We note that the

equivalent capacity methods admits more connections than the diffusion method when the activity ratio exceeds 0.25 for both class 1 and class 2.

# 2.10 Effect of the Ratio of the Buffer Size to the Mean Burst Length

We have already observed that the diffusion approximation scheme and the equivalent capacity scheme behave similarly when the buffer size is large. In this section, we study the effect of the ratio of the buffer size to the mean burst length of a connection, while keeping all other parameters fixed. We consider a multiplexer with either class 1 or 2 connections. The peak and average rates are given in table 1 while the mean burst size b was varied. For each value of *b*, the buffer size *K* was varied so that the ratio K/b varied from 0.1 to 100.

The results for the equivalent capacity, the diffusion approximation, and the CLP upper bound schemes are shown in figure 7. We note that for the equivalent capacity and the diffusion approximation methods, as long as the ratio K/b is kept constant, the maximum number of admitted connections is almost the same regardless of the value of the mean burst length b. This observation can be used in order to approximate the solution of a multiplexer with a large buffer size by that of a multiplexer with a smaller buffer. The mean burst length of the source must be scaled down accordingly in order to keep the ratio K/b constant.

The CLP upper bound scheme does not behave similarly, since it does not use any information about the burst length of the connection. This is reflected in figures 7(e) and 7(f). In this case, for each given value of b and K/b we get a new value of K. Since b is not taken into account in the calculation, the number of connections does not scale as in the other two schemes. As has already been observed this scheme's sensitivity to buffer size is poor.





# 3. CAC SCHEMES BASED ON THE CELL DELAY

For real-time applications, the network must be able to provide timely delivery of packets. For many applications, packets must be delivered within a bounded delay and/or bounded delay jitter. In this case, the CAC process has to ensure that the network will meet the required end-to-end delay and/or delay jitter for a new connection. Also, the CAC must insure that admitting the new connection would not affect those connections already in progress.

For delay bounded connections, we have two major categories of QOS: deterministic and statistical. For deterministic QOS, a connection requests that all its packets reach their destination within some finite delay D. Such a connection will be called a guaranteed service connection. For statistical QOS, a connection requests, for example, that the probability that the delay of a packet is smaller than a given bound D must be greater than a given value  $\Delta$ . Such a connection will be called a predictive service connection. In this paper, we concentrate on CAC schemes for guaranteed service connections.

CAC schemes for the cell delay are closely associated with the packet scheduling mechanism implemented in the network switches. The

scheduling mechanism determines to a large extent the packet queueing delay at each switch. A lot of work has been done in the area of calculating packet delays for various scheduling disciplines such as First-In-First-Out [10, 11], Static Priority [10, 11], Weighted Fair Queueing [20], and Earliest Deadline First [19]. When comparing scheduling disciplines it is necessary to evaluate the following aspects:

- Admission/schedulability region: how many connections from each class can be admitted without violating their requested delay bounds?
- Isolation and fairness among connections
- Ease of implementation and complexity of the calculation needed to perform the admissibility/schedulability test

In our model, we assume that connections are constrained by a leakybucket like traffic filter and each connection *i* has the traffic descriptor ( $R_i$ ,  $r_i$ ,  $b_i$ ) where  $R_i$  is the peak rate,  $r_i$  the average rate, and  $b_i$  is the maximum burst size. With this traffic model, it is possible to calculate the worst-case end-to-end delay for many scheduling disciplines.

#### 3.1 Packet-by-Packet Generalized Processor Sharing

Weighted fair queueing (WFQ) and packet-by-packet generalized processor sharing (PGPS) are approximations of the Generalized Processor Sharing (GPS) discipline. WFQ and PGPS are identical so we will only consider WFQ. In GPS, packets are served as if they are in separate logical queues, the server visits each nonempty queue in turn and serves an infinitesimally small amount of data from each queue, so that, in any finite time interval, it can visit each logical queue at least once, independent of the number of queues. The scheduler in WFQ works as follows: compute the time that a packet would finish its service if the packet is served by a GPS server; then serve packets in order of their finishing times. The calculation of the packet finishing times under (weighted) GPS is illustrated in Keshav [15].

To determine the worst-case end-to-end packet delay, consider a connection constrained by  $(b_i, \mathbf{r}_i)$  passing through *L* schedulers, where the *l*th scheduler has a link rate  $C_1$ . Let  $g_{i,l}$  be the service rate assigned to that connection at the *l*th scheduler. Let  $g_i = \min_l g_{i,l}$ , where  $g_i \ge \mathbf{r}_i$  for stability of the queues. Let  $P_{\max_i}$  be the largest packet from connection *i*, and assume that  $P_{\max}$  is the largest size of packet allowed in the network. Then, the end-to-end network delay  $d_i$  for a packet from connection *i* satisfies [44, 55]:

$$d_{i} \leq \frac{b_{i}}{g_{i}} + \sum_{l=1}^{L-1} \frac{P_{\max_{i}}}{g_{i,l}} + \sum_{l=1}^{L} \frac{P_{\max_{i}}}{C_{l}}$$
(8)

independently of the behaviour of other connections.

It is very important to note that, when the link speed is very large compared to  $P_{\text{max}}$ , the above bound of  $d_i$  simplifies to  $\frac{b_i}{g_i}$ , i.e. packetization is very important for providing small end-to-end delay

is very important for providing small end-to-end delay.

A CAC scheme based on WFQ scheduling works as follows. When a connection is setup, the connection parameters  $(b_i, \mathbf{r}_i, d_i)$  are signaled to the network. The network calculates the required  $g_i$  to satisfy the delay constraint using equation (8). If  $g_i \ge \mathbf{r}_i$  and the sum of  $g_i$  plus the reserved bandwidth of the existing connections is smaller than  $C_1$  and the sum  $\mathbf{r}_i$  of plus the overall average rate of the connections is smaller than  $C_1$  at all intermediate switches, the connection is admitted; otherwise it is rejected.

#### 3.2 Delay Earliest-Deadline-First Scheduling

In earliest-deadline-first (EDF) schedulers, each packet is assigned a deadline and the scheduler serves packets in order of their deadline. Delay-EDF is an extension of EDF that describes how a scheduler assigns deadlines to packets. At connection setup time, the connection declares a peak rate and a desired delay bound for worst-case delay. The scheduler performs a schedulability test to ensure that every connection meets its delay bound even when they are transmitting at peak rate.

A delay-EDF scheduler needs to sort packets in order of their deadline, which is also done by WFQ. The scheduler also needs to store finishing times as in WFQ. The main advantage of delay-EDF over WFQ is that its delay bound is independent of the allocated bandwidth to the connection at the expense of peak bandwidth allocation (this, however, can be relaxed for connections constrained by a leaky-bucket). EDF has been proven to be an optimal scheduling discipline in the sense that if a set of connections is schedulable under any scheduling discipline then the set is also EDFschedulable in the single node case.

Consider leaky-bucket constrained connections with traffic descriptor ( $b_i$ ,  $\mathbf{r}_i$ ) and a delay bound  $d_i$  at scheduler l. Assume that two connections i and j

are ordered such that  $d_i^l < d_j^l$  if i < j. Then as long as  $\sum_{i=1}^{N} \mathbf{r}_i < C_l$ , we have the following schedulability condition at scheduler l (due to Libeherr, Werge, and Ferrari [19])

$$d_{j}^{l} \geq \frac{b_{j} + \sum_{i=1}^{j-1} (b_{i} - \mathbf{r}_{i} d_{i}^{l}) + \max_{k > j} P_{\max_{k}}}{C_{l} - \sum_{i=1}^{j-1} \mathbf{r}_{i}}$$
(9)

The schedulability test of delay-EDF schedulers is complex since the check for condition (9) is computationally expensive.

Liebherr and Werge [18] simplified the implementation of EDF scheduling by discretizing the range of packet deadline values. The search time for the next packet to schedule is brought to O(1). Firiou, Kurose, and Towsley [11], suggested an efficient algorithm for schedulability testing given that connections are constrained by  $(R, \mathbf{r}, b)$ . The complexity is O(N), where N is the number of admitted connection at the time of invocation of the schedulability test.

A possible CAC scheme based on EDF scheduling is the following:

- 1. A set-up message for connection *i* is sent along the connection's selected path. The set-up message contains connection's i traffic descriptor ( $R_i$ ,  $r_i$ ,  $b_i$ ) and its end-to-end delay bound  $d_i$ . A variable  $d_i$  is initialized to zero and included in the set-up message.
- 2. At each intermediate scheduler l, a minimum value for the maximum delay  $d_i^{\iota}$  that can be assured for connection i is calculated. The variable  $d_i$  is incremented by  $d_i^{\iota}$ . At the same time, CAC checks if  $\sum \mathbf{r} < C_l$  for all connections passing through the link including the new connection.
- 3. At the destination node, CAC checks if  $d_i \leq d_i$ . If yes, the connection is accepted.
- 4. On the reverse path, a local delay bound  $d'_i$  is calculated. This is the local deadline of connection *i* at link *l*.

### **3.3** Static Priority Scheduling

A static-priority (SP) scheduler assigns each connection to a fixed priority level p, where  $1 \le p \le P$ , where P is the number of priority levels. All connections in priority level p will have the same delay bound  $d_p$ , with  $d_p < d_q$  for p < q, i.e. the priority of a connection is high if its delay bound is low. The SP scheduler always selects the first arriving packet packet from the highest priority backlogged queue. It is fairly easy to implement an SP scheduler since it consists of a fixed number of FIFO queues, one for each priority level. For leaky-bucket constrained connections, Cruz [5] has derived necessary and sufficient schedulability conditions for SP schedulers to satisfy a given delay bound.

Consider connections that are  $(\mathbf{r}, b)$  constrained, where  $\mathbf{r}$  is the average rate and b is the maximum burst size. Let  $P_{\max_p}$  be the largest packet size for connections belonging to priority level p. Assuming only one connection in each priority level and that the minimum packet size is zero. Let P be the number of sessions,  $(\mathbf{r}_p, b_p, d_p)$  be the traffic descriptor and delay bound for connection p, where  $1 \le p \le P$ , then the set of connections is schedulable at link l if

$$d_{p} \geq \frac{\sum_{q=1}^{p} b_{q} + \max_{r>p} P_{\max_{r}}}{C_{l} - \sum_{q=1}^{p-1} r_{q}}$$
(10)

for  $1 \le p \le P$ . A CAC scheme for SP schedulers can be devised in a similar manner to the EDF-based CAC as shown in section 3.2.

# 3.4 Comparative Performance Analysis of the CAC Schemes for Connections with a Maximum Delay Bound

In this section, we carry out a comparative study of the WFQ, EDF, and SP scheduling disciplines for connections with guaranteed service for delay. We consider a source/destination pair interconnected by a path of L hops. The nodes are homogeneous and have a fixed link capacity C. We also assume that all connections are identical, and that they all traverse the same set of nodes and links from source to destination. We evaluate the maximum number of connections that can be admitted in the network without violating the delay bound, for various traffic characteristics and different values of the delay bound. Certainly, the above does not describe a realistic network configuration, since in a real network there will be more than one source destination pair, connections would traverse different paths, and connections would have varying traffic descriptors and delay bounds. This fictitious configuration, however, is useful in illustrating the basic behavior of the scheduling disciplines and their performance relative to each other.

In all experiments we set the link capacity *C* to 155 Mbps and all packets are constrained to be of fixed length of 53 bytes. We assume a connection of characterised by the descriptor ( $\mathbf{r}$ , b), where  $\mathbf{r}$  is the average rate and b is maximum burst size. We fix  $\mathbf{r}$  at 3550 bps such that the condition  $N\mathbf{r} < C$ , where *N* is the maximum number of connections admitted by a specific scheduling discipline, is not violated in all the cases considered. This is deliberately chosen such that the number of connections is constrained by the deadline schedulability conditions of a particular scheduling discipline and not the stability condition of the system. The number of hops traversed by the connections is varied from 1 to 10, the burst size *b* takes values from the set {0.1, 1, 8}Kbytes, and the delay bound *d* takes values from set {10, 50, 100} msec.

Figure 8 shows the maximum number of connections as a function of the number of hops, burst size, and delay bound for WFQ, EDF, and SP scheduling disciplines. Since we have only one class of connections, SP becomes equivalent to the FIFO discipline.

For the three disciplines, there is a clear strong dependence of the number of admitted connections on the burst size and the number of hops traversed. As the burst size increases, the number of admitted connections decreases. Likewise, as the number of hops increases, we observe a decrease on the number of admitted connections. The burst size has the strongest influence on WFQ. Also, in the case of WFQ, the increased number of hops does not affect the number of connections when burst size is very large compared to the maximum packet size (in the reported results, b=8 kbytes, and  $P_{max} = 53$  bytes). In general for small burst sizes (i.e. b=0.1 Kbytes), the number of admitted connections decreases rapidly as the number of hops increases. The influence of the number of hops on the number of admitted connections decreases (i.e. when the sources become burstier).

In figure 9, we compare the performance of the three scheduling disciplines, assuming an end-to-end delay of d = 50 msec and b = 0.1 and 8 Kbytes. For the examples considered, both WFQ and EDF always provide the same number of connections for L = 1. It is obvious that WFQ performs consistently better than EDF in a multi-hop network. Does this violate the fact that EDF is optimal? No, the reason behind this is that in WFQ, the reserved bandwidth is calculated using a methodology that takes into account the network as a whole. In EDF, however, local delays are added up in each node without taking into account how the connection's traffic is distributed among the multi-hop path. This shows that there is a need to modify the schedulability conditions of EDF (and SP) schedulers to take the distribution of traffic among the network nodes. It is also clear that by increasing the number of hops, EDF and SP become identical. Also, for large burst size, there is no difference between EDF and SP.



Figure 8. Admission region under WFQ, EDF, and SP

Figure 10 shows the dependence on burst size for the case where L = 1 (note that WFQ and EDF provide identical results for L=1). The burst size is varied from 50 bytes (about one ATM cell) to 50 Kbytes (about 1000 ATM cells). For large bursts, the scheduling disciplines provide identical results and the number of admitted connections decreases dramatically. Since most real-time applications have some knowledge about the required end-to-end delay but they can not tell in advance their burst size, we suggest that the network should be able to negotiate a burst size with the application requesting guaranteed service. For predictive service, we suggest adjusting

the burst size dynamically at run-time to allow for better network performance.



*Figure 9* Comparison of WFQ, EDF, and SP, with  $D^* = 50$  mses



Figure 10. Effect of burst size on admission region

### 4. CONCLUSIONS

In this paper, we have provided an evaluation of CAC schemes for cell loss and delay sensitive services in an ATM network. For cell loss sensitive services, we evaluated the equivalent capacity, the diffusion approximation and the CLP upper bound methods. Below, we summarize the findings reported in this paper:

**Performance**: In most of the scenarios considered in this paper, the diffusion approximation has outperformed the other methods in terms of providing a larger admission region than the other two schemes. The CLP upper bound scheme is usually a pessimistic scheme but still provides some

statistical gain over peak rate allocation. The CLP upper bound scheme as the activity ratio decreases and approaches the same performance levels attained by the other two schemes.

**Complexity of CAC decision making**: The CLP upper bound method is the most complex to implement because of the expensive and computationally intensive convolution operation (note that a simplification for the evaluation exists). The diffusion approximation method is computationally more expensive than the equivalent capacity method since  $C_{FB}$  and/or  $C_{IB}$  needs to be computed each time a new connection is negotiated with the system. The calculation of  $C_{FB}$  and  $C_{IB}$  involves the computation of a square root. If the Gaussian approximation (see equation (2)) is used to complement the equivalent capacity basic calculation, then the equivalent capacity and diffusion approximation schemes are approximately equivalent from the complexity point of view.

**Suitability for implementation in a real-life network**: The CLP upper bound method is more suitable for implementation in a real-life scenario since it requires only two traffic parameters which are peak rate and average rate. In scenarios where traffic is shaped, it would be possible to get estimates for the burst size and the other schemes (equivalent capacity and diffusion approximation) will provide better performance than CLP.

**Sensitivity**: All schemes have exhibited some dependency on the variations of the system parameters. The equivalent capacity and diffusion approximation methods show significant dependency on buffer size and the buffer size to mean burst size ratio. The CLP upper bound scheme does take the mean burst size into account and is therefore not sensitive to uncertainties in this parameter. All schemes have exhibited large sensitivity to the activity ratio. Therefore any change in the average/peak rate values will affect the accuracy of the decision made by the CAC scheme applied. This calls for the need to do apply dynamic bandwidth allocation and calculation when applying these schemes to assure the accuracy of the CAC decision in real-life networks.

In the second part of the paper, we considered delay sensitive services. Specifically, we discussed CAC schemes associated with three scheduling disciplines: WFQ, EDF, and SP scheduling. We believe WFQ to be the best available candidate for deployment in real-life networks due to its superiority when it comes to the issues of complexity, implementation and performance. However, WFQ is the most sensitive scheme to burst size uncertainties. EDF can be competitive to WFQ in the case of networks with small number of hops.

### 5. **REFERENCES**

- [1] ATM Forum User Network Interface (UNI) 4.0 Specification, 1996.
- [2] P. Castelli, E. Cavallero, and A. Tonietti, Policing And Call Admission Problems in ATM Networks, in: A. Jensen and V.B. Iversen (Eds.), Teletraffic and datatraffic in a period of change, (North-Holland, 847-852, 1991.
- [3] C.-S. Chang and J. A. Thomas. Effective Bandwidth in High Speed Networks. IEEE Journal on Selected Areas in Communications, 13:1091--1100, 1995.
- [4] G. L. Choudhury, D. M. Lucantoni, and W. Whitt, On the Effective Bandwidths for Admission Control in ATM Networks. In Proceedings of 14th International Teletraffic Congress (ITC), pages 411--420, 1994.
- [5] R. Cruz, A calculus for network delay: part I: Network elements in isolation, IEEE Trans. Inform. Theory, vol. 37, 114-131, 1991.
- [6] B. T. Doshi. Deterministic Rule Based Traffic Descriptors for Broadband ISDN: Worst Case Behavior and Connection Acceptance Control. In Proceedings of 14th International Teletraffic Congress (ITC), 591-600, 1994.
- [7] N. G. Duffield, J. T. Lewis, N. O'Connel, R. Russell, and F. Toomey. Entropy of ATM Traffic Streams: A Tool for Estimating QoS Parameters. IEEE Journal on Selected Areas in Communications, 13:981-990, 1995.
- [8] K. Elsayed and H. G. Perros. Analysis of an ATM Statistical Multiplexer with Heterogeneous Markovian On/Off Sources and Applications to Call Admission Control. Journal of High Speed Networks, vol. 6 no. 2, pp. 123-139, 1997.
- [9] A. Elwalid and D. Mitra, Effective bandwidth of general Markovian traffic sources and admission control of high speed networks, IEEE/ACM Trans. Networking 1, 329-343, 1993.
- [10] A. Elwalid, D. Heyman, T. V. Lakshman, D. Mitra, and A. Weiss. Fundamental Bounds and Approximations for ATM Multiplexers with Applications to Video Teleconferencing. IEEE Journal on Selected Areas in Communications, 13:1004-1016, 1995.
- [11] V. Firoiu, J. Kurose, and D. Towsley, Efficient admission control for EDF Schedulers, IEEE Infocom'97, 1997.
- [12] E. Gelenbe, X. Mang, and R. Onvural, Diffusion based statistical call admission control in ATM, Performance Evaluation, vol. 27, 411-36, 1996.

- [13] S. J. Golestani, Congestion-free communication in broadband packet networks, IEEE Trans. Comm. 39, 1802-1812, 1991.
- [14] R. Gu'erin, H. Ahmadi, M. Naghshineh, Equivalent capacity and its application to bandwidth allocation in high-speed networks, IEEE Journal on Selected Areas in Communications, vol. 9, 968-981, 1991.
- [15] S. Keshav, An Engineering Approach to Computer Networks, Chapter 9, Addison-Wesley, 1996.
- [16] V. Kulkarani, L. Chimento, Effective bandwidth vector for two-priority ATM traffic, INFOCOM '94, 1056-1064, 1994.
- [17] S.-Q. Li. A General Solution Technique for Discrete Queueing Analysis of Multimedia Traffic on ATM. IEEE Transactions on Communications, 39:1115--1132, 1991.
- [18] J. Liebherr and D. E. Werge, Design and analysis of a high-performance packet multiplexer for multiservice networks, with delay guarantees, Technical report CS-94-03, University of Virginia, 1994.
- [19] J. Liebherr, D. E. Werge, and D. Ferrari, Exact admission control for networks with a bounded delay service, IEEE Trans. on Networking, vol. 4, 885-901, 1996.
- [20] A. K. Parekh and R. G. Gallager, A generalized processor sharing approach to flow control in integrated services networks: The multiple node case, IEEE Trans. on Networking, vol. 2, 137-150, 1994.
- [21] C. Rasmussen, J.H. Kvols, and S.B. Jacobsen, Sourceindependent call acceptance procedures in ATM networks, IEEE JSAC 9, 351-358, 1991.
- [22] K. M. Rege. Equivalent Bandwidth and Related Admission Criteria for ATM Systems-A Performance Study. International Journal of Communications Systems, 7:181--197, 1994.
- [23] H. Saito, Call admission control in an ATM network using upper bound of cell loss probability, IEEE Trans. Comm. vol. 40, 1512-1521, 1992.
- [24] G. De Veciana, G. Kesidis, and J. Walrand. Resource Management in Wide-Area ATM Networks Using Effective Bandwidth. IEEE Journal on Selected Areas in Communications, 13:1081--1090, 1995.