

Cooperating Leaky Bucket for Average Rate Enforcement of VBR Video Traffic in ATM Networks

Joseph S.M. Ho, Hüseyin Uzunalioglu and Ian F. Akyildiz

School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332
U.S.A.

Abstract

Because of the burstiness of video traffic, a very large sampling interval is needed to obtain acceptable cell loss probability. Previous studies reported very high cell loss probabilities for well-behaved sources when using the Leaky Bucket mechanism (LB) to police VBR video sources. In this paper, a mean rate policing mechanism called Cooperating Leaky Bucket (CLB) is introduced. This mechanism can attain better cell loss probability by sharing the leaky buckets among all sources. In addition, it guarantees that the cell loss probability will never be higher than that obtained by LB. Depending on the number of traffic sources available, the cell loss probability can be significantly reduced when compared to LB. At the same time, the combined traffic entering the network is enforced to below the total negotiated traffic rate. Moreover, the additional overhead due to the implementation of CLB is minimal.

1 Introduction

Preventive Congestion Control (PCC) method is employed by ATM networks to ensure the Quality of Service (QOS) such as cell loss probability, delay and delay jitter of the connections. Under PCC, each traffic source has to provide the network a *traffic descriptor* during connection setup. The traffic descriptor consists of three primitive parameters: peak rate, average rate, and peak rate duration. A *call admission control* procedure is then run to make sure that the new traffic source can be supported without degradation to the QOS of other connections. Once the connection is established, a *policing* function is used to ensure that the source adheres to the negotiated

parameters specified by the traffic descriptor.

Possibilities for the sources to violate the negotiated traffic parameters exist. This may be due to the inability of the sources to accurately estimate the traffic characteristics before the actual communications or it is because of attempts to obtain economical advantages. In any case, violations of the negotiated parameters can lead to degradation of network performance. Excessive cell losses may occur as a result of the unexpected traffic load.

A number of policing mechanisms have been proposed. Among others, the Leaky Bucket mechanism (LB) [11] is the most widely accepted. The LB mechanism can be viewed as a counter which increases by one when a cell arrives and decreases at a specified rate called the *leak rate*. If the counter is at its maximum allowable value, further cell arrival constitutes traffic violation and the cell is discarded. Two parameters of the LB mechanism, leak rate and counter limit, can be varied to tailor the policer for particular requirements. The leak rate is generally set to the policed traffic rate as LB can never allow traffic rate higher than the leak rate to enter the network. Adjusting the counter from one extreme to another represents a trade-off between the sensitivity to violation and the allowable burstiness of the arriving traffic.

While it is generally easy to enforce peak traffic rate, there are several problems to be addressed in average rate enforcement. The bursty nature of some ATM traffic types makes it very difficult to detect violations effectively. In the case of VBR video traffic, the peak rate may be many times higher than the average rate. A sudden arrival of a burst of cells can easily overflow the leaky bucket causing excessive cell losses. This happens even when the traffic source has been adhering to the negotiated mean rate in the long run. This problem can be alleviated by increasing the LB

counter limit so that the leaky bucket can accommodate larger burst of cell arrivals. However, the larger the counter limit, the longer it takes to detect a traffic violation. Damages may have already been done to the network by the time the violation is detected.

A simulation study reported in [10] demonstrates that the violation probability is extremely high when the LB mechanism is used for mean rate policing of VBR sources. The simulations employ real-life video data including the *Star Wars* sequence and video-phone and video conference data. It is shown that the violation probability can be as high as 0.5 when the LB counter is small. Even if the counter limit is set to 100,000, the violation probability is still at a very high level. The study concludes that the LB counter values needed to police the selected video sources close to their mean rate are not realistic.

Several policing mechanisms were introduced that attempt to improve the classical LB mechanism in different aspects. In [11], an overdimensioning factor $C \geq 1$ is added to the LB mechanism. Under this scheme, the policed traffic rate is equal to C times the negotiated mean traffic rate. Depending on the value of C selected, the cell loss probability can be reduced close to 0. However, this method also reduces the sensitivity of the LB mechanism to violating traffic. A traffic source can effectively transmit at the policed traffic rate, which is higher than the negotiated traffic rate, throughout the whole connection. The QOS of other connections can therefore be severely degraded.

Another possible way of reducing the cell loss probability is to allow both the violating and non-violating cells to enter the network. However, the violating cells are marked as low priority which will be dropped first whenever congestion occurs [5]. This mechanism allows the population of cells to grow uncontrollably. The QOS of both marked and unmarked cells can be affected. In [3], this scheme is extended to allowing the unmarking of previously marked cells. This mechanism requires the addition of a second Cell Loss Priority (CLP) bit in the ATM cell header. Besides, information on network condition is needed to decide whether a cell should be unmarked. Another modification to the original LB mechanism is proposed in [7]. The basic idea is to increase the counter limit of the leaky bucket from its original value, N , to a higher value K . However, cells are discarded when the counter reaches N . However, this mechanism does little to reduce the cell loss probability of non-violating traffic sources. It is proposed in [9] to use multiple leaky buckets to police a single source. An example using two leaky buckets is given such that short term

and long term mean traffic rates are enforced by two separate leaky buckets. Each cell has to obtain a token from each leaky bucket in order to be admitted to the network. By choosing the parameters correctly, the maximum long term mean rate can be enforced to the negotiated value while the maximum burst size can be kept reasonably small. A convenient way of selecting the parameters for the leaky buckets is needed.

A policing mechanism called Credit Banking (CB) is introduced in [1]. Under CB, a source is allowed to borrow transmission credits from its own future. The problem with this mechanism is that initial conditions are hard to define. Also, there is no simple way of determining the length of various time intervals that can provide optimal result.

In this paper, we introduce a mechanism called Cooperating Leaky Bucket (CLB.) CLB is a variation of LB in which buckets from all the traffic sources are coordinated. With CLB, the cell loss probabilities for "good" (always adhere to the negotiated mean rate value) traffic sources can be greatly reduced without adding significant overhead. The cell loss probability for each source can be guaranteed to be at least as small as that obtained by the LB mechanism. At the same time, combined traffic entering the network can be enforced to strictly below the total policed mean traffic rate.

Simulation studies are conducted to demonstrate the performance of CLB compared to the LB mechanisms. Results are obtained using both non-violating and violating traffic sources under various parameters. It can be seen in Section 3 that CLB outperforms LB in all setups when there are more than one traffic sources. Cell loss probability can be reduced significantly depending on the number of traffic sources. Moreover, the overhead introduced by CLB as compared to LB is minimal.

The organization of this paper is as follows. In Section 2, we describe the CLB mechanism. Section 3 gives the simulation results using VBR video traffic sources. A conclusion is given in Section 4.

2 Cooperating Leaky Bucket

The LB mechanism ensures that the traffic rate entering the network can never exceed the leak rate of the bucket. Non-bursty traffic is a good candidate for mean rate policing using the LB mechanism. However, as the burstiness of the traffic increases, the LB counter limit has to be increased accordingly to accommodate the arrival of traffic bursts. This, in turn, increases the time required to detect traffic violations.

One possible way to deal with this problem is to perform virtual path policing [2]. Traffic entering the virtual path comes from the statistical multiplexing of virtual channels. As a result, this traffic stream is much less bursty and the LB mechanism should perform well in policing virtual path traffic. However, fairness is a major problem if the policer is located after the multiplexer. A violating traffic source can easily take away most of the available bandwidth. The QOS of other sources can be significantly degraded.

The CLB mechanism takes advantage of the fact that multiplexed traffic is much less bursty, but at the same time ensures the fairness in the allocation of bandwidth. The implementation of CLB is similar to that of LB. A leaky bucket is located at the entrance to the network. The leak rate is set to the negotiated mean traffic rate and the counter limit is set based on the burstiness of the traffic. The major difference that distinguishes CLB from LB is that once a bucket becomes empty, its leak rate will be distributed to other buckets according to the following.

Assume there are n traffic sources. Let \mathbf{m} be a $1 \times n$ dimensional vector such that its i^{th} element is the negotiated mean rate of the i^{th} traffic source. Also let \mathbf{e} be an $n \times 1$ dimensional vector such that its i^{th} element is 0 if the i^{th} leaky bucket is empty and 1 otherwise. If the leaky buckets from all sources are empty, then the leak rate for each leaky bucket is 0. Otherwise, the leak rate for the i^{th} leaky bucket, denoted by l_i , can be calculated as follows:

$$l_i = \mathbf{m} \cdot \mathbf{1} \cdot \frac{m_i}{\mathbf{m} \cdot \mathbf{e}} \cdot e_i \quad (1)$$

where $\mathbf{1}$ is an $n \times 1$ vector containing all 1's and m_i and e_i are the i^{th} elements of the \mathbf{m} and \mathbf{e} vectors, respectively. The leak interval of the i^{th} leaky bucket, denoted by t_i , is given by:

$$t_i = \frac{1}{l_i} \quad (2)$$

The leak rate for each source depends on the number of empty leaky buckets. This means that leak rates has to be updated dynamically in response to changes in the state of the leaky buckets. We will introduce later in this section a method that can be used to obtain the leak rates as given by equation (1) without constantly going through this bucket state checking and leak rate calculation process.

According to equation (1), the total leak rate is distributed among the non-empty buckets in proportion to the magnitude of the negotiated mean rate. When there is no empty bucket, CLB works exactly the same as LB. However, when there exists one or




Figure 1: LB mechanism.




Figure 2: CLB mechanism.

more empty buckets, CLB distributes the unused leak rate to other buckets. In the case of LB, the leak rate is wasted if the bucket becomes empty before the arrival of the next burst of cells. This essentially reduces the long term rate of the traffic entering the network to below the leak rate. Since it is unlikely for most of the sources to have long bursts of cell arrivals at the same time, many leaky buckets become empty well before the end of each frame. CLB takes advantage of this statistical characteristic of the traffic sources by allocating the unused leak rate to the needed sources. The CLB mechanism makes the best use of the available leak rates to maximize the performance. This mechanism essentially polices the much less bursty combined traffic of all the input sources at the total mean rate. Unlike virtual path policing, CLB guarantees the minimum bandwidth allocated to each source in the form of a minimum leak rate. At the same time, the combined traffic rate entering the network can never exceed the total negotiated traffic rate.

Figure 1 shows the cell arrivals during a selected frame interval of two traffic sources. Assume that frames from both sources start at the same time. Also, cells within each frame are clustered at the beginning of the frame interval. It is shown that source 1 has a below average number of cell arrivals, its leaky bucket is empty well before the end of the frame. However, source 2 has an above average number of cell arrivals; its leaky bucket remains non-empty at the end of the frame. This carry-over of bucket content to the next

frame will decrease the number of cells that can be accepted during the next frame interval. Hence, the cell loss probability of subsequent frames will be increased. Figure 2 shows the two frames as in Figure 1. In this case, CLB is used for mean rate policing. As in the previous case, the leaky bucket of source 1 is empty well before the end of the frame. According to the CLB mechanism, the unused leak rate of source 1 is allocated to source 2 once source 1 becomes empty. With the increased leak rate, the leaky bucket of source 2 is able to return to empty by the end of the frame. The cell loss probability of subsequent frames is not affected. This example demonstrates how CLB can reduce the cell loss probability by sharing the leak rate. It may seem, at this point, that CLB will not be able to detect violations effectively because it allows larger cell bursts to enter the network. Simulation results given in section 3 demonstrate that CLB can effectively lower the cell loss probability of "good" sources while at the same time penalizing violating sources with high cell loss rate. This will be discussed in more detail in section 3.

For homogeneous traffic sources, the leak rate, as specified by equation (1), can be achieved by implementing the leaky buckets as a polling system [12, 4]. Figure 3 shows the polling system with n leaky buckets. The service rate of the server is deterministic and it is set to n times the policed mean rate. The service discipline is round robin with one-limited service [12]. According to this discipline, after serving queue k , the server will select queue $(k+1)$ modulo n for service. If the bucket selected is empty, the server will immediately switch to the next bucket until a non-empty bucket is selected. This polling model guarantees that the leak rate of each bucket is at least the mean service rate when it is not empty. When a leaky bucket is empty, its allocated leak rate is evenly distributed to the other non-empty buckets. This method eliminates the leak rate calculations required every time when a queue becomes empty or non-empty.

3 Experiments

Simulation studies are conducted to compare the performance of the CLB mechanism to the LB mechanism. A number of assumptions is made throughout these simulation studies. They are described in the following subsection.

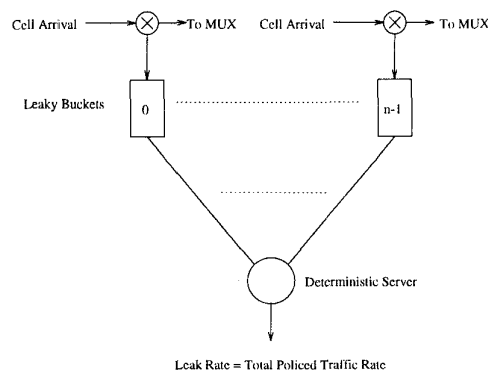


Figure 3: Polling system implementation of the leaky buckets.

3.1 Source Model

We assume that video traffic is produced by encoding subsequent video frames at a rate of 30 frames per second. The amount of information (number of bits) in a frame is generated according to the first order Continuous State Autoregressive Markov Model (AR) [8]. This model provides a rather accurate approximation of the bit rate of a single video source. The activity levels captured by the AR model are consistent with those encountered in the case of the picture-phone [8, 6]. With AR, the number of bits per pixel for frame n , denoted by $\lambda(n)$, is given as follows:

$$\lambda(n) = a\lambda(n-1) + bw(n) \quad (3)$$

where a and b are constants and $w(n)$ is a white Gaussian noise component with mean η and variance 1. Here we use the following values: $a=0.8781$, $b=0.1108$ and $\eta=0.572$. Assuming $\lambda(0)$ to be 0, bit rate of the subsequent frames can be generated recursively by equation (3). The expected value, $E(\lambda)$, of $\lambda(n)$ is 0.52 bits per pixel.

We use a channel speed of 155 Mbps. The number of pixels per frame is assumed to be 250,000. These assumptions plus the facts that there are 8 bits per byte and 44 bytes payload per cell give the following equation for the number of cells for frame n (denoted by $L(n)$):

$$L(n) = \frac{250000 \times \lambda(n)}{8 \times 44} = 710.23\lambda(n) \quad (4)$$

The expected value of $L(n)$ is $710.23 \times E(\lambda) = 369.32$ cells/frame. The mean traffic rate is $369.32/30 = 11079.6$ cells/second.

Two different modes of cell arrival patterns are considered: clustered (CF) and evenly distributed (EF).

In the CF mode, cells within a given frame are clustered at the beginning of the frame. After the arrival of the last cell in the frame, the source remains idle until the beginning of the next frame. In the EF mode, cells are evenly distributed within a frame. It is obvious that CF produces a much bursty traffic compared to EF. The LB mechanism should therefore perform better under EF. However, the additional delay introduced by EF may not be tolerated by time critical traffic such as video conferencing.

Frames from different sources can be either synchronized (SF) or unsynchronized (UF). In the case of SF, frames from different sources all start at the same time. On the other hand, when UF is used, frame start times are assumed to be randomly distributed.

3.2 Simulation Results

The first set of simulations is conducted to demonstrate the reduction in cell loss probability when CLB is used in place of LB. The traffic arrivals are generated as described in the previous subsection and the number of incoming sources is varied from 1 to 64. Unless specified otherwise, the LB and CLB counter limits are set to 450 for all the simulations to be performed. The following four cases are considered:

Case 1: Cells are clustered at the beginning of the frame; frames from different sources all start at the same time.

Case 2: Cells are evenly distributed in a frame; frames from different sources all start at the same time.

Case 3: Cells are clustered at the beginning of the frame; frame start times are randomly distributed.

Case 4: Cells are evenly distributed in a frame; frame start times are randomly distributed.

Case 1 is expected to generate the most bursty traffic. The other three cases generate varying degree of burstiness as will be shown by the simulation results.

Assume the cell loss probabilities obtained by using CLB to be p_{clb} . Also assume p_{lb} to be the cell loss probability when the same source (which has the same negotiated mean rate) is policed using LB. We define the normalized cell loss rate of CLB, denoted by R , as following:

$$R = \frac{p_{clb}}{p_{lb}} \quad (5)$$

Figure 4 gives the normalized cell loss rate for cases 1 to 4 as the number of sources is changed from 1 to 64.

In all cases, the normalized cell loss rates drop as the number of sources increases when CLB is applied. As expected, the cell loss probabilities for both mechanisms are the same and the normalized cell loss rate is




Figure 4: Normalized cell loss rate.

1 when there is only one source. No leak rate sharing can be done in this situation. As the number of sources increases, the normalized cell loss rate drops progressively. However, the drop is more significant when the number of sources is small. As can be seen in the graphs, the normalized cell loss rates are reduced by a much larger percentage when the number of sources is increased from 1 to 2 as compared to the increase from 32 to 64. This allows the network to capture most of the improvement due to CLB without requiring a large number of sources. In case 2, the normalized cell loss rate is about 0.15 when the number of sources is 64. The normalized cell loss rate is around 0.7 even when there are only two traffic sources. This set of simulations demonstrates the ability of CLB to take advantage of the statistical characteristics of bursty sources. The cell loss probabilities are drastically reduced when compared to LB. However the mean rate of each source entering the network is enforced to the policed value.

The second set of simulations is conducted to demonstrate the fairness of the CLB mechanism. The CF and SF modes are used in these simulations. Two values for the number of sources, 4 and 8, are used. In both cases, the bit rate of one of the sources is multiplied by a violation factor.

Figures 5 and 6 present the results of this set of simulations. Here we evaluate a policing mechanism by measuring how close the traffic rate entering the network from the policed mean rate. The effectiveness ratio (ER) is defined as:

$$ER = \frac{\lambda_n - \lambda_p}{\lambda_p} \quad (6)$$




Figure 5: Effectiveness ratio for 4 sources




Figure 6: Effectiveness ratio for 8 sources

where λ_n is the traffic rate entering the network and λ_p is the policed mean rate. An ideal policing mechanism has an ER of 0. An ER value above or below 0 represents insufficient or excessive cell drops, respectively. Figures 5 and 6 show the graphs of ER versus violation factor for the two selected source counts. The ER for "good" traffic is negative and that for violating traffic is positive when CLB is used. The ER remains negative for both traffic types when LB is used. For "good" traffic, the LB mechanism generates 30% and 70% (for 4 and 8 sources, respectively) more cell losses. For violating traffic, LB performs better (closer to 0) when the violation probability is high. Inherently, LB is a good peak rate policer. It is therefore expected that LB can enforce the traffic rate to very close to the policed value when the violation factor is high.

As described in section 1, a method for reducing the

unnecessary cell drops for the "good" traffic sources is to multiply the policed traffic rate by an overdimensioning constant $C \geq 1$. Under the assumptions described earlier, simulations show that C values of 1.1 and 1.2 are required for LB to achieve similar "good" traffic cell loss probabilities as compared to CLB when the number of sources is 4 and 8, respectively. It can be seen in the graphs that the performance of LB is close to that of CLB when this overdimensioning constant is added. However, LB mechanism with a C value higher than 1 can allow excessive number of cells into the network. A violating source can effectively transmit at the policed mean rate, for the duration of the connection without being caught. This increases the cell loss probabilities of other sources at the multiplexer and hence degrades the QOS of other connections. Besides, it is not generally easy to select a C value that is good for all traffic conditions. In order for LB to match the performance of CLB, a new C value is needed for each set of source parameters (such as number of sources, traffic rate, etc.) CLB, on the other hand, will never allow the combined traffic entering the network to be higher than the total negotiated mean rate. The additional bandwidth allocated to each source, "good" or otherwise, is the unnecessary bandwidth from another source during its 'off' period. Even though more bandwidth is allocated to each traffic source as compared to LB, this does not constitute any degradation to the QOS of other connections when all traffic streams share the same virtual path.

The third set of simulations is designed to look into the effect of changing the counter limit on the cell loss probability for both the CLB and the LB mechanisms. As before, the number of traffic sources is set to 8 and the CF and SF modes are used. Traffic is generated according to AR as before. There is no violating traffic source in this case. Figure 7 gives the simulation results as the counter limit is increased from 50 to 2050. It can be seen that the normalized cell loss rate is close to 1 when the counter limit is small. The cell loss probabilities for both mechanisms are very close to each other when the counter limit is smaller than the average number of cells per frame. This is because the buckets are too small to accommodate the cell arrivals during most of the frames. The buckets stay full most of the time and several cells have to be discarded regardless of whether CLB or LB is used. As the counter limit increases, the normalized cell loss rate drops significant. The normalized cell loss rate constantly stays below 0.5 when the counter




Figure 7: Normalized cell loss rate as a function of counter limit.

limit is larger than 500. This demonstrates the ability of CLB to take advantage of the bursty nature of the traffic when the counter limit is sufficiently large. Consequently, more buckets can return to empty by the beginning of the next frame. The number of cells that can be accepted in the following frame is thus increased and hence reduces the cell loss probability.

The last set of simulations is used to demonstrate the performance of CLB with heterogeneous traffic sources. Traffic is generated as described in subsection 3.1. However, frame arrival rates of half of the sources are increased from 30 to 60 frames per second. We use the terms *slow* and *fast* to identify the original and the modified traffic sources, respectively. The leak rate of each leaky bucket is set to the mean arrival rate of the policed source. The total number of sources is changed from 2 to 64 and the normalized cell loss rate are recorded for both traffic types. There is no violating source in these simulations. Figure 8 gives the simulation results. It can be seen that the normalized cell loss rates for both traffic types are very close to each other. The normalized cell loss rates decrease as the number of traffic sources increases. Similar percentage gain in cell loss probability can be obtained regardless of the negotiated mean traffic rate of the sources. This demonstrated the ability of CLB to obtain better cell loss probability in heterogeneous environment.

4 Conclusions

In this paper, we introduced a policing mechanism called Cooperating Leaky Bucket (CLB.) This mechanism takes advantage of the bursty characteristic of




Figure 8: Normalized cell loss rate with equal number of fast and slow sources

VBR video traffic sources to improve the performance of the classical Leaky Bucket mechanism. CLB can reduce the cell loss probability when there exists a minimum of two traffic sources. The improvement is more significant if more traffic sources are present. The CLB mechanism can decrease the cell loss probability of "good" traffic sources, and hence increases the traffic rate entering the network. But at the same time, the CLB mechanism enforces the total traffic rate entering the network to below the total negotiated mean rate. CLB also penalizes violating traffic with high cell drop rate. The cell loss probabilities for all traffic are guaranteed to be as small as that produced by the LB mechanism.

We have to note, however, that CLB may not be appropriate for policing individual virtual circuit (VC) traffic. Congestion may be created at the switches where traffic streams from different CLB's merge. Currently, we are working on the extension of CLB mechanism for use in VC policing. This may be achieved by incrementing the bucket size by the ratio of current leak rate to the policed leak rate when a cell arrives. This ratio is always equal to 1 for classical LB. The result presented in the paper is not intended as a final solution of the VBR video traffic policing problem. We will investigate in combining CLB with other policing mechanisms to achieve even better cell loss probability under various traffic parameters. The dimensioning problem for CLB will also be addressed. A mechanism is needed to select the appropriate bucket size and leak rate for a specific traffic source.

References

- [1] I.F. Akyildiz, Y. Gong, "Average Rate Enforcement for Real-Time Traffic in ATM Networks," *Technical Report, Georgia Institute of Technology*, March 1992; Submitted for Publication.
- [2] A. Atkinson, "A Traffic Control Scheme for Virtual Paths in an Asynchronous Transfer Mode Network," *IEEE GLOBECOM 94*, pp.1768-1773
- [3] V. Bemmell, M. Ilyas, "A Unified Congestion Control Strategy in ATM Networks," *IEEE GLOBECOM 93*, pp. 1600-1604.
- [4] A. Demers, S. Keshav, S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," *Journal of Internetworking Research and Experience*, pp. 2-26, Oct. 1990.
- [5] A. Gravey, et-al., "Tagged Versus Strict Rate Enforcement in ATM Networks," *IEEE GLOBECOM 91*, pp. 218-288.
- [6] B.G. Haskell, "Buffer and Channel Sharing by Several Interframe Picturephone Coders," *BSTJ*, Vol. 51, No. 1, pp. 261-289, Jan. 1972.
- [7] B. Lague, C. Rosenberg, F. Guillemin, "A Generalization of Some Policing Mechanisms," *IEEE INFOCOM 92*, Vol. 2, pp. 767-775, May 1992.
- [8] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, J.D. Robbins, "Performance Models of Statistical Multiplexing in Packet Video Communications," *IEEE Transactions on Communications*, Vol. 36, No. 7, pp. 834-844, July 1988.
- [9] A. Ortega, M. Vetterli, "Multiple Leaky Buckets for Increased Statistical Multiplexing of ATM Video," *Proc. of Packet Video Workshop '94*, Sept. 1994.
- [10] E.P. Rathgeb, "Policing of Realistic VBR Video Traffic in an ATM Network," *International Journal of Digital and Analog Communication Systems*, Vol. 6, pp. 213-226, 1993.
- [11] E.P. Rathgeb, "Modeling and Performance comparison of Policing Mechanisms for ATM Networks," *IEEE JSAC*, Vol. SAC-9:325-334, April 1991.
- [12] H. Takagi, *Analysis of Polling Systems*, The MIT Press, 1986.