XIII International Switching Symposium

Stockholm, Sweden

May 27–June 1, 1990

Under the patronage of
Carl XVI Gustaf

Proceedings
Thursday afternoon, May 31, 1990

Volume 5 of 6

THE POLICING FUNCTION IN ATM NETWORKS

Erwin P. Rathgeb, Thomas H. Theimer
University of Stuttgart

ABSTRACT

ATM networks have been proposed by CCITT as the solution for future Integrated Broadband Communication Networks. They provide a high flexibility with regard to varying bandwidth requirements for different services as well as the momentary bitrate within a connection (bitrate on demand). For the network operator this results in a need to control the individual connections in order to ensure an acceptable quality of service for all existing connections by means of a so called policing or source monitoring function. In this paper, some basic aspects of this function including the dimensioning of its parameters are addressed using the "Leaky Bucket" mechanism as an example.

1. INTRODUCTION

Communication networks based on the Asynchronous Transfer Mode (ATM) have been proposed by CCITT as the solution for the future Broadband Integrated Services Digital Network (BISDN) [1]. In an ATM network any information is packetized and transferred in small, fixed size blocks called cells using the virtual connection (VC) concept. The statistical multiplexing of cells belonging to different virtual connections will provide the user with the possibility to have a bitrate varying in a wide range during a connection according to his needs (bitrate on demand). Flexible User/Network Interfaces (UNIs) in the range of 150 and 600 Mbit/s will allow connections ranging from low to very high bitrates to support services like voice, interactive video and high speed data.

Relevant service attributes characterizing the connection, like e.g. maximum and mean cell rate, have to be negotiated with the network at call setup. The admission control function of the network has to decide, whether the new virtual connection can be accepted or not taking into account the quality of service requirements of the new connection itself and also those of all other connections influenced by the new connection.

The broadband interfaces together with the fact, that a virtual connection can in principle exceed the negotiated throughput parameters up to the maximum capacity of the UNI requires a new network function called source monitoring or policing function. This function has to protect the network against congestion due to violation of the negotiated parameters resulting in a degradation of the quality of service for all connections sharing the same network resources in a statistical manner.

To protect the network, action has to be taken by the policing function after detecting a violation of the parameters. The most obvious action that can be taken is to discard all cells exceeding the range defined by the throughput parameters, which will be assumed in the following sections to show the effects of different dimensioning alternatives. When discarding cells, an ideal loss curve can be defined for the policing function taking into account two requirements, namely

- A well behaving source, i.e. a source not exceeding the negotiated parameters, shall experience no cell loss due to the policing function.
- All cells exceeding the parameters shall be discarded.

This ideal loss curve will be used in section 4.2 to assess the different dimensioning alternatives.

Another proposal is to mark the violating cells for preferred deletion in case of congestion within the network. This implies, that some sort of priority mechanism has to be implemented within the network. In this case also a special charging mechanism (volume based) should be applied in the network to prevent users from intentionally exceeding the parameters.

In order to protect all network resources, the policing function should be placed directly at the traffic source. On the other hand, this function has to remain under the control of the network, and therefore it can only be located as close as possible to the UNI. Furthermore, the source monitoring function has to be performed on a connection individual basis, which implies that it must be simple and especially efficient to implement in hardware.

To meet the rather conflicting requirements, several policing mechanisms have been proposed. The most prominent of them, the "Leaky Bucket" mechanism, will be used in the following to discuss the problem of dimensioning the policing function such that the best compromise between all requirements is obtained.

2. THE LEAKY BUCKET MECHANISM

The basic idea of the Leaky Bucket mechanism [2,4,7,8] is to control the difference between the negotiated mean cell rate and the actual cell rate of a source. This is achieved by a counter which is incremented each time a cell is generated by the source and decremented periodically with the mean cell rate. Thus, the state of the counter reflects the short-term bandwidth requirements of the source.

In order to avoid congestion inside the network, the cell rate of a source must be limited at the UNI by the policing function. If the negotiated throughput parameters are significantly violated, the counter value begins to increase and after some time it will reach a given limit. This indicates to the policing function that the source has exceeded the admissible parameter range, and the mechanism starts to take suitable actions (e.g. discard or mark cells) on all generated cells until the counter has fallen below its limit again.

Although the mechanism clearly limits the used bandwidth of any source, it is unavoidable that occasionally some cells of a well behaving source are also dropped, introducing additional cell losses. It must be guaranteed by an appropriate dimensioning of the counter limit, that this probability is in the same order of magnitude as the cell loss probability for other network resources like multiplexers and switching nodes, if a unique quality of service for all services has to be provided by the network.

Since the policing function is performed on individual VCs, a service dependent loss probability could be introduced very easily at this point. This would allow to exploit the gap between the cell loss acceptable for a specific service and the unavoidable cell loss introduced by the rest of the network when dimensioning the policing function.

The evaluation of the cell loss probability for the policing function requires models for the source traffic as well as for the mechanism itself. With respect to the cell loss, the G/D/1-s delay loss system is an exact model for the Leaky Bucket mechanism. This model consists of a single server with deterministic service times, a finite capacity queue with *s* waiting places and a general arrival process. The service time of the model is equal to the mean interarrival time of cells, and the number of customers in the system (including server and queue) directly represents the state of the counter. It should be noted that the G/D/1-s queue is only a traffic model for the Leaky Bucket Mechanism. Of course, no cell is actually queued, and the resulting cell stream entering the network is not the output process of the queueing system. However, this model has stimulated the name "Leaky Bucket", because it is similar to a bucket which is "filled" by the source traffic and emptied with a constant rate through a hole in the bottom.

In addition to controlling the mean cell rate, the mechanism of the Leaky Bucket can also be used to police the peak cell rate of a source, which is equivalent to checking the minimum interarrival time of cells. This is simply realized by setting the counter limit to one and by decrementing the counter with the minimum interarrival time. Thus, it is guaranteed at least for the UNI that the distance between two cells of the same connection always exceeds a given minimum. If the customer premises network (CPN) already introduces a variable cell delay, the minimum distance may be violated at the UNI even by a well behaving source. In this case, the Leaky Bucket Mechanism can still be used, but it may be necessary to increase the counter limit to compensate the delay jitter of the CPN.

3. DIMENSIONING ASPECTS

The dimensioning of the source policing function requires realistic models for the traffic sources reflecting their main characteristics. However, most of the services which will be supported by a BISDN are still unknown today, making it impossible to forecast the particular traffic characteristics of future networks. Even the models and parameters for existing services like video telephony cannot be fixed yet, because the development of new video codecs is still in progress. Therefore, current studies are usually based on simple source models involving a high degree of abstraction. Although these models may not be very realistic, they will be used in the following to discuss the basic problems arising from the dimensioning of the Leaky Bucket Mechanism.

A traffic model which is frequently used for performance studies of ATM networks is the Bernoulli Process. Assuming a slotted operation of the UNI, this model is well suited for ATM networks due to its discrete time nature. In each timeslot, there is a fixed probability that a cell arrives at the UNI which is originating from the source under consideration. Consequently, the interarrival time between two cells has a shifted geometric distribution, and the Geo/D/1-s queue [5] is an exact model for the calculation of the cell loss probability.

If the Leaky Bucket Mechanism is dimensioned as described in section 2, the service time of the Geo/D/1-s queue is equal to the mean cell interarrival time, resulting in an offered traffic load A = 100%. However, Figure 1 shows that in this case a very high counter limit would be required to obtain an acceptable cell loss probability. Although the realization of a high counter limit is not very difficult, the dynamic behaviour of the mechanism becomes very poor, because it takes too long to detect a parameter violation.

Another problem with high counter limits is, that long bursts are admitted into the network, which could in principle be sent at the maximum UNI cell rate even if the negotiated maximum cell rate is much lower. Therefore, it might also be necessary to police the maximum cell rate, e.g. by cascading two Leaky Buckets with different maximum counter values and different decrementation intervals (see section 2), to avoid short term buffer congestion within the network [4].

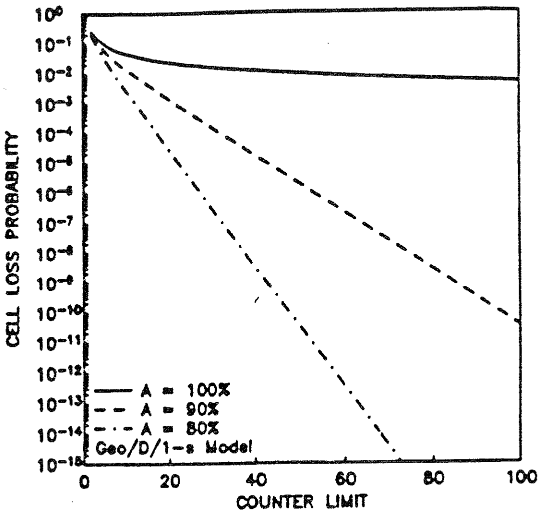

Figure 1 Loss Probability of the Leaky Bucket Mechanism vs Counter Limit

A solution to these problems could be to decrement the Leaky Bucket counter faster than it is given by the mean interarrival time. This means that the counter limit required for a given cell loss probability decreases rapidly, because the offered traffic load of the corresponding queueing model is now below 100%. Figure 1 indicates that a counter limit of about 50 is sufficient to obtain a cell loss probability of 10^{-10}, if the counter is decremented in intervals corresponding to 80% of the mean interarrival time. Thus, the dynamic behaviour of the system improves significantly, but the source may exceed the negotiated mean cell rate up to 25%.

The above discussion shows that there is always a tradeoff between the loss probability of a well behaving source, the dynamic behaviour of the mechanism and the amount of bandwidth by which the negotiated mean cell rate may be exceeded without detection. If the Leaky Bucket is dimensioned according to the mean value of the interarrival time, the bandwidth is limited correctly, but the reaction time of the system is not acceptable. On the other hand, if the mechanism is overdimensioned by increasing the rate at which the counter is decremented, the dynamic behaviour can be improved, but it must be accepted that the source may exceed its specified mean cell rate. A reasonable compromise between these alternatives has to be found, taking into account the network operators needs and especially the quality of service requirements for specific services.

4. INFLUENCE OF THE SOURCE CHARACTERISTICS

In the previous section, a Bernoulli source has been used to describe the dimensioning approach for the Leaky Bucket mechanism. In the following, we will investigate the sensitivity of the results with respect to the source characteristics. We will use a Burst-Silence source as described below to vary the coefficient of variation for the interarrival times and the burstiness of the individual sources. The model is in principle known from previous papers (e.g. [3]), where the parameters have been chosen to model a packetized voice connection.

4.1 THE BURST-SILENCE SOURCE MODEL

The Burst-Silence source can be characterized as shown in Figure 2. If the source is active on the connection level, it alternates between bursts of activity and silence phases on the cell level. During the bursts, cells are emitted with constant inter cell distance *T*. A burst consists of a number of *X* cells, where *X* is a geometrically distributed random variable with mean *EX*. There is at least one cell in every burst, and the silence phases are assumed to be negative-exponentially distributed with mean *ES*.

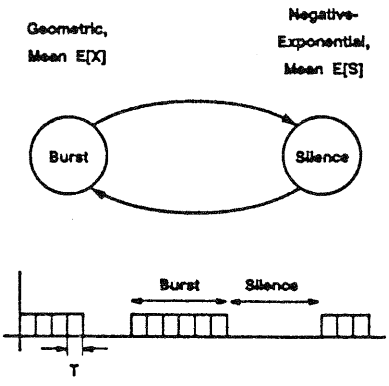

Figure 2 Burst — Silence Source Model

The time between two cell arrivals from one source can be characterized by the mean *EA*, the coefficient of variation *c* and the burstiness *B*.

$$EA = T + \frac{ES}{EX}, \quad c = \frac{ES}{EA} \times \sqrt{1 - \left(\frac{EX-1}{EX}\right)^2}, \quad B = \frac{EA}{T} \qquad (1)$$

The coefficient of variation indicates the variability of the cell interarrival times, and the burstiness gives the ratio of the maximum cell rate and the mean cell rate of the source. With this model, it is possible to keep two of the above characteristics constant while varying the third.

To analyze the arising GI/D/1-s system discrete time analysis methods have been used. The only adaptation that must be made for the analysis is to discretize the negative-exponential distribution of the silence phase. The discretization error, however, is small enough for the analysis to yield accurate results.

The steady state analysis has been carried out using an efficient direct approach involving the solution of the linear equation system characterizing the amount of work in the system at consecutive cell arrival instants. These equations can in principle be found in [6], but the iteration algorithm described there is better suited for an analysis of the transient behaviour of the system than for the evaluation of the system in the steady state.

4.2 RESULTS

Figure 3 clearly shows the drastic influence of the interarrival coefficient of variation and of the burstiness on the loss probability. The Bernoulli source (B = 20, c = 0.975) yields far too optimistic results for most cases. The results for different counter decrementation intervals resulting in an offered traffic load of 70%, 53%, 40% and 37%, respectively, are shown in Figure 4. In this case the parameters for a packetized voice source (T = 16 ms, EX = 22, ES = 650 ms, EA = 45.54 ms [3]) have been used. It is obvious, that for a required cell loss probability of 10^{-10} and a maximum counter limit of about 50, the decrementation interval must be chosen to be 18 ms (A = 40%), which is close to the minimum cell interarrival time. Even for an admissible cell loss probability of 10^{-4}, which could be acceptable for voice traffic, and a counter limit close to 200 the operating point of the mechanism would be only 70% (decrementation interval: 32 ms).

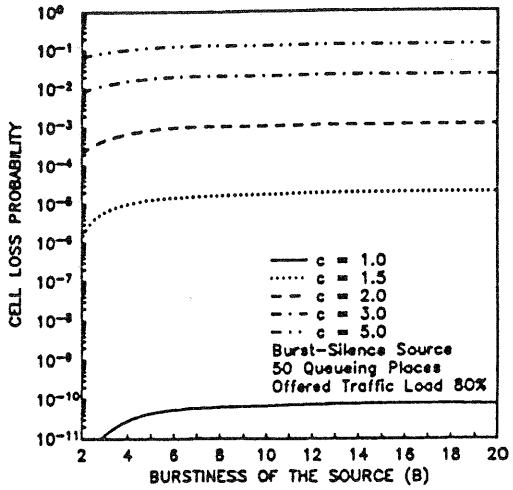

Figure 3 Influence of the Source Characteristics on the Cell Loss Probability

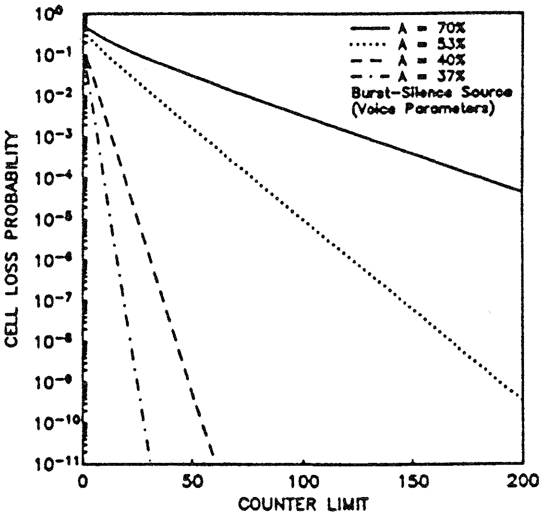

Figure 4 Loss Probability of the Leaky Bucket Mechanism for a Burst-Silence Source

The effect of the different operating points is shown in Figure 5 for a counter limit of 50. The ideal loss curve in this figure is characterized by zero cell loss for mean cell rates below the specified value (i.e. relative mean cell rates below one), and by a cell loss probability given by a carried load of 1.0 in the range above, which means that in this range all violating cells are discarded.

With the high operating point, the overload detection is relatively good, which is indicated by the approximation to the ideal loss curve in the range where the mean cell rate is higher than the negotiated cell rate. However, the actual loss probability for a well behaving source (relative mean cell rate below 1.0) is too high to be accepted.

The dotted curve gives the loss probability for a source emitting cells in fixed intervals. Such a source could exceed the negotiated mean cell rate by a factor of about 1.4 without suffering any cell loss, which has to be taken into account when designing the charging function.

This effect is even more drastic for the 40% operating point, where a deterministic source could exceed the negotiated mean cell rate up to a factor of 2.5 without cell loss. The dimensioning with poor overload detection fulfills the requirements for the loss

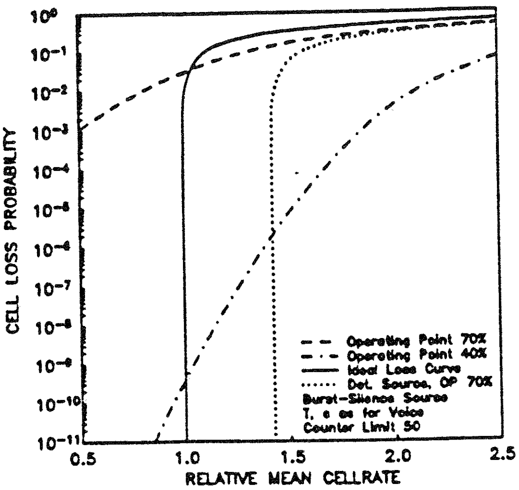

Figure 5 Influence of the Operating Point on the Loss Probability of the Leaky Bucket Mechanism

of a well behaving source, but it is very close to policing the peak cell rate.

5. CONCLUSION

Summarizing the discussion in the previous sections, the main aspects of the policing function are

- the remaining loss probability for a source which is in accordance with the negotiated parameters,
- the accuracy for the detection of parameter violations and
- the dynamic behaviour of the function.

These aspects and the necessary tradeoffs are not at all specific for the Leaky Bucket mechanism used as an example in this paper. They are common to all mechanisms proposed so far, including the different variants of window mechanisms.

The optimum compromise for both the network and the user not only depends on the characteristics and the quality of service requirements of the individual services, but also involves other functions like admission control and charging. Our studies have shown that it is not possible to police the mean cell rate of sources with highly bursty characteristics, if a very low cell loss probability has to be guaranteed for all sources complying with the negotiated parameters. Consequently, for some of these services it will only be possible to control the maximum cell rate resulting in a need for peak bandwidth allocation. To avoid a peak bandwidth allocation for services with a high burstiness and a relatively long duration of burst and silence periods (e.g. interactive file transfer), dynamic bandwidth allocation schemes (burst switching) will be useful.

For connections with high bitrate requirements, where a peak rate bandwidth allocation has to be applied anyway, it may be sufficient to police the maximum cell rate in conjunction with a volume based charging scheme that makes it unattractive to exceed the negotiated parameters. However, in this case it is important to have a very fast dynamic reaction, because a single misbehaving source can cause severe congestion.

Slow sources, on the other hand, that can be efficiently multiplexed and emit only few cells per second, can be allowed to send long bursts at their maximum rate without harming the network. These sources can be policed according to a cell rate close to their mean cell rate, especially if they are not very loss sensitive and service dependent loss probabilities are acceptable.

A basic assumption in the above discussions is, that the relevant characteristics are known exactly at call setup. This may be true for the maximum cell rate, but is more unlikely for e.g. mean cell rates and the coefficient of variation of the cell interarrival times, especially when looking at interactive retrieval services. Another assumption is, that the traffic characteristics experienced by the policing function are exactly the same as generated by the terminal. Taking into account e.g. the jitter introduced by complex customer premises networks (NT2) with multiple access protocols this assumption may not be realistic in all cases. Inaccuracies and uncertainties of this kind make it particularly difficult to dimension the policing function to low loss probabilities for sources complying with the parameters, because in this range the loss probabilities are very sensitive to parameter variations.

ACKNOWLEDGEMENTS

This work has been supported by Philips Kommunikations Industrie AG, Nuernberg, FRG. The authors would like to thank H. Kroener for his stimulating ideas concerning the efficient solution of the discrete time GI/D/1-s queue and Dr. R. Lehnert for his helpful comments.

REFERENCES

[1] *CCITT, Recommendation I.121; "Broadband Aspects of ISDN", Blue Book, Vol. III.7, Geneva 1989.*

[2] *CCITT, Study Group XVIII, Temporary Document 74; "Likely Solution Package", San Diego, February 1989.*

[3] **Heffes H., Lucantoni D.M.;** *"A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Performance", IEEE Journal on Selected Areas in Communications, Vol. SAC-4, No. 6, September 1986, pp. 856-868.*

[4] **Kowalk W., Lehnert R.;** *"The 'Policing Function' to Control User Access in ATM Networks - Definition and Implementation", Proceedings of the ISSLS'88, Boston, September 1988.*

[5] **Louvion J.R., Boyer P., Gravey A.;** *"A Discrete-Time Single Server Queue with Bernoulli Arrivals and Constant Service Time", Proceedings of the 12th International Teletraffic Congress, Torino, June 1988.*

[6] **Tran-Gia P., Ahmadi H.;** *"Analysis of a Discrete-Time $G^{[X]}$/D/1-s Queueing System with Applications in Packet-Switching Systems", Proceedings of the INFOCOM'88, New Orleans, 1988.*

[7] **Turner J.S.;** *"New Directions in Communications (or Which Way in the Information Age ?)", Proceedings of the Zurich Seminar on Digital Communications, Zurich, March 1986, pp. 25-32.*

[8] **Turner J.S.;** *"On Integrated Networks for Diverse Applications", Tutorial Notes of the INFOCOM'87, San Francisco, March 1987.*