Quality of Service of Packet Data Streams over ATM Networks

Jordi Domingo-Pascual and Judit Sánchez-Valdivia

Universitat Politècnica de Catalunya. Departament d´Arquitectura de Computadors. Campus Nord. Mòdul D6. 08071 BARCELONA (Catalunya) SPAIN. Phone: +34 3 4016981, Fax: +34 3 4017055, E-mail: jordid@ac.upc.es

Abstract

The Asynchronous Transfer Mode (ATM) has been defined by the ITU-T (formerly CCITT) as the technique to support B-ISDN [1]. The B-ISDN can be seen as the universal information transmission infrastructure which merges the concepts of Integration of Services and High Speed Transmission into a unique bearer service.

Basically, the QoS parameters considered in ATM Networks are the cell loss probability and the end-to-end delay and its variation (delay jitter). For packet oriented communications the cell loss probability is not really a measurement of the quality of service as the data unit transmitted is a packet, that is a series of ATM cells. Even if only one cell is missing the receiving entity will discard the whole packet, so the real QoS measurement is the packet loss probability. This packet loss probability is not easy to define as it depends on the cell loss probability, the size of the packet and the actual rate cells are transmitted depending whether spacing control is performed or not. When congestion occurs at the multiplexing buffers, the mechanism proposed here reduces the load discarding the remaining cells of a given packet once one cell has been lost. This alleviates the congestion and allows cells of other connections to be stored into the buffer. This reactive congestion control mechanism is to be applied at the ATM Level and must be implemented in hardware.

This paper presents a performance evaluation study measuring the packet losses of a reference connection as the background traffic mean load is increased, and we compare the results of the two situations: without discarding cells and using the proposed cell discarding algorithm. Results show that the proposed discarding solution has a good behaviour reducing the overall packet losses of the connections and performs as a good congestion control mechanism.

1. INTRODUCTION

ATM Networks provide connection oriented fast packet-switching transmission of data. The small and fixed size of the packets, named cells, allow simple fast switching, different transmission rates and variables ones, and short transmission delay [2]. On the other hand internal buffering at the switching nodes makes that packet switched networks have the drawback of allowing variable delay and a small probability of error. Connection

This work has been funded by the Spanish Ministry of Education (CICYT) under grant number TIC92-1289-PB.

oriented networks can provide the user with a quality of service, but ATM Networks do not implement error control within the link layer and, thus, lost cells are detected in a end-to-end basis, that is, error control and flow control are end-to-end functions. With nowadays transmission technology on fiber optics the main cause of transmission errors is due to the loss of cells in the switching node buffers. One approach is to dimension the buffers on the multiplexing and switching stages so that the cell loss probability is very small, but if the same solution wants to be applied to bursty traffic this implies long buffers. A trade-off must be made between the cell loss probability (CLP decreases with long buffers) and the delay introduced by the buffering.

One approach is to dimension multiplexing and switching buffers to cope with cell level contention (that due to the multiplexing of cell streams in an output buffer) only. In this way, end-to-end delays are short and cell delay variations are small also. With this approach the main problem is to control the burst level contention in order to maintain the desired CLR. On the other hand, this design allows heterogeneous traffic to share switching equipment and links if good traffic and congestion control procedures are applied.

2. QUALITY OF SERVICE FOR PACKET ORIENTED COMMUNICATIONS

The main quality of service parameters for ATM connections are the cell loss probability (CLP), the end-to-end delay (D), and the cell delay variation (CDV), also known as delay jitter. Each network management procedure must take into account these parameters; bandwidth allocation procedures, routing decisions, call acceptance control functions should work together to provide the desired values for these QoS parameters.

While data applications can accept moderate to long delays and variable delay for the cells of the same connection, real-time applications and more concretely interactive applications are more stringent and do not allow long delays nor variable ones. For interactive traffic and for Cell Delay Variation (CDV) sensitive traffic, small buffers provide short end-to-end delay and small CDV. In this paper we focus on the performance of ATM networks for packet oriented applications rather than cell stream oriented. Stream oriented transmission over ATM is supported by the ATM Adaptation Layer 1 (AAL1); this is suitable for circuit emulation, uncompressed voice (64 Kbps) and uncompressed video. Packet oriented applications comprise a wide range of video, image and voice applications as well as all data transmission ones, including LAN-to-LAN interconnection, LAN emulation, and generic data transmission (X.25 and TCP/IP based applications). These latter types of communication are known as ABR (“Available Bit Rate”) and all of them are supported by the AAL5 [3], [4]. This ABR traffic coexists with the other types of traffic and accepts low grade service in the sense of a small Cell Loss Probability (CLP). But for packet oriented communications the cell loss probability is not a real measure of the quality of service. As the receiving application, or the user, will discard a whole packet when only one cell has been lost, due to buffer overflow, or network congestion, a more appropriate measure of the QoS will be the packet loss probability (PLP). But as traffic and congestion functions work at cell level, when congestion appears, they must choose the cells to be dropped. To choose the cells to be discarded we will take advantage of the fact that when one cell is lost all the cells of this packet becomes useless; so, discarding all these cells alleviates congestion and allows cells of the other connections to be stored in the buffer.

Roughly, we can approximate the PLP as P times CLP, being P the number of cells per packet and being CLP small enough, as it is in ATM networks where the buffers are dimensioned to cope with cell multiplexing contention. As the packet length is an important parameter we have studied LAN-to-LAN interconnection traffic over ATM in an heterogeneous environment. LAN-to-LAN traffic is envisaged to be the main source of traffic, at least during the first stage of the deployment of ATM networks.

3. SELECTIVE DISCARDING MECHANISM

An ATM network with small buffers will improve its performance if burst level congestion is avoided. To achieve this target a selective cell discarding mechanism is proposed. It is a reactive congestion control function at cell level. This mechanism seems to be useful for bursty traffic and ABR traffic. In addition to the fact that the loss of one cell makes useless all the other cells of the packet, being they already transmitted or not yet, the buffering of some of these "useless" cells can cause the loss of other cells of the remaining connections that share the buffer; when these cells are discarded no connection is punished in particular, as these are useless cells. Once a cell is dropped at the output port a signal is sent to the corresponding input port so that all the remaining cells with that VCI are discarded until the last ATM cell arrives. The last cell of the packet is tagged as End-of-SDU type cell [4]. Then, the discard signal is reset. This last cell is transmitted so that the receiving entity recognises easily a truncated or lost packet.

The present approach is different from that presented in [5], though the concept is similar. There they propose a Partial Discard and Early Packet Discard procedures and evaluate them for TCP connections over ATM, modelling by simulation the behaviour of the TCP protocol when multiplexing a given number of TCP connections. The study we present in this paper is more general and does not take into account the effect of the transport protocol. We focus on the packet loss probability for packet oriented communications which includes all data communications using any transport protocol. Particularly, we study LAN-to-LAN traffic over ATM.

4. EVALUATION AND RESULTS

The model used to carry out the simulation is a simple one. There are two sources of traffic: a reference connection and a background one. Both share a multiplexing buffer. The capacity of the buffer is small compared with the packet length (that is, the burst size).

The reference connection traffic is modelled by a MMDP ON-OFF source. It has two states "Idle" and "Send". It changes from the Idle state to the Send state if a generated random number q is lower than the parameter "a" that represents the mean load of the reference connection traffic. Varying "*a*" we can simulate different given mean loads for the reference connection traffic. In the Send state the cells of a packet are transmitted at a constant rate. Each "K" cycles a cell is generated, and a complete packet has "P" cells, so the source remains P*K cycles in the Send state transmitting at a constant rate (1/K).

The background traffic simulates a traffic that is the result of a wide variety of different sources, and generates cells with a given mean load "*x*".

A simulation works as follows. The three input parameters are "P", "K" and "*a*", that is, the length of the reference connection packet, the rate of the reference connection (1/K) normalised to the output link capacity, and the mean load of the reference connection. For each value of the background mean load "*x*", the number of lost background cells and the number of lost reference connection packets are measured and are given as a percentage. The range of "*x*" varies from low load to congestion. Each simulation is done without and with the proposed discarding mechanism, and the results for both situations are compared.

The simulations have modelled several scenarios. First, we represent an Ethernet LAN-to-LAN interconnection over an ATM link at 155 Mbps. Each packet is about 32 ATM cells ("P" = 32), which are sent one each fifteen cell slot time ("K" = 15) (that corresponds to a 10 Mbps. rate), the next three graphics correspond to a reference connection mean load, "*a*", of 0.2, 0.5 and 0.7, respectively.

Study 1 ("S1") means that the simulation has been carried out without applying the discarding mechanism, and Study 2 ("S2") means that the mechanism has been applied. Each legend shows the percentage of background lost cells: (% bk S1) for the S1 case, and (% bk S2) for case S2. The percentage of reference connection lost packets are labelled as (% pk S1) and (% pk S2) respectively. The reference axis "load" means the total load of both sources of traffic, that is, x + a/K. This is the reason why the saturation point is different for each case, depending of the load.

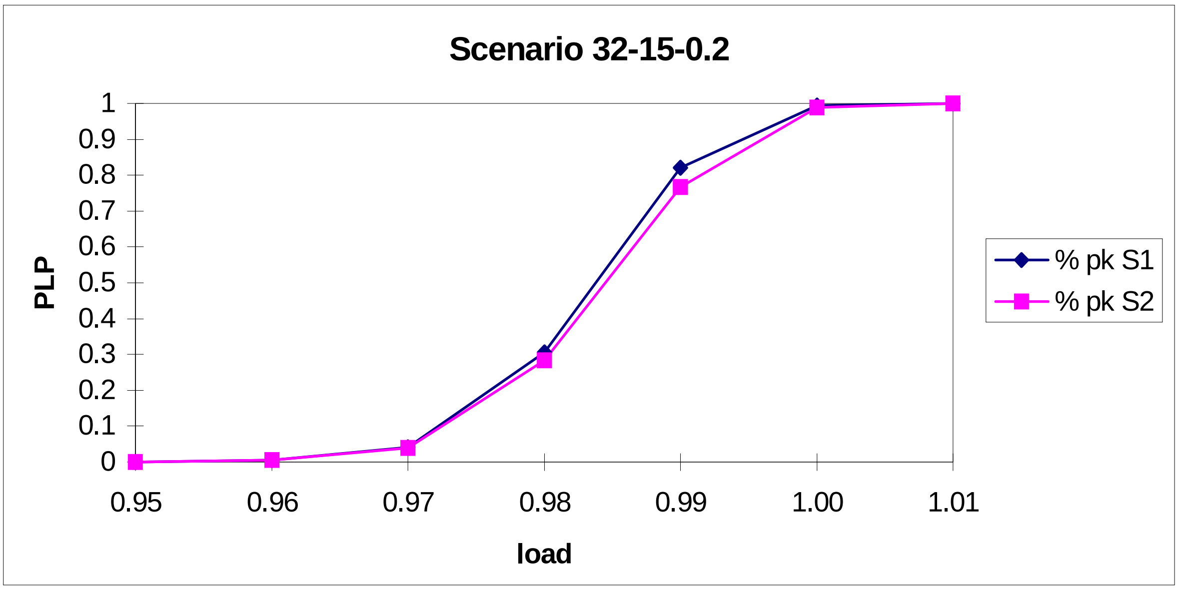

Figure 1. Packet losses for P = 32, K = 15, a = 0.2

Figures 1 to 6 show the packet losses for the reference connection and the cell losses for the background connection when congestion occurs at the multiplexing buffer.

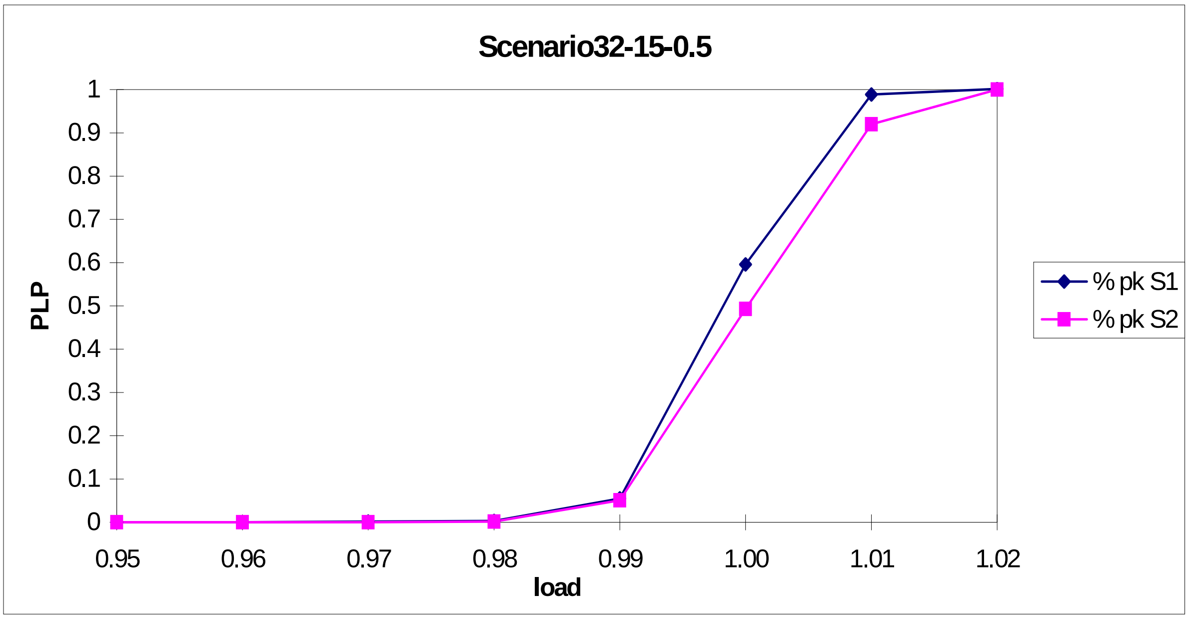

Figure 2. Packet losses for P = 32, K = 15, a = 0.5

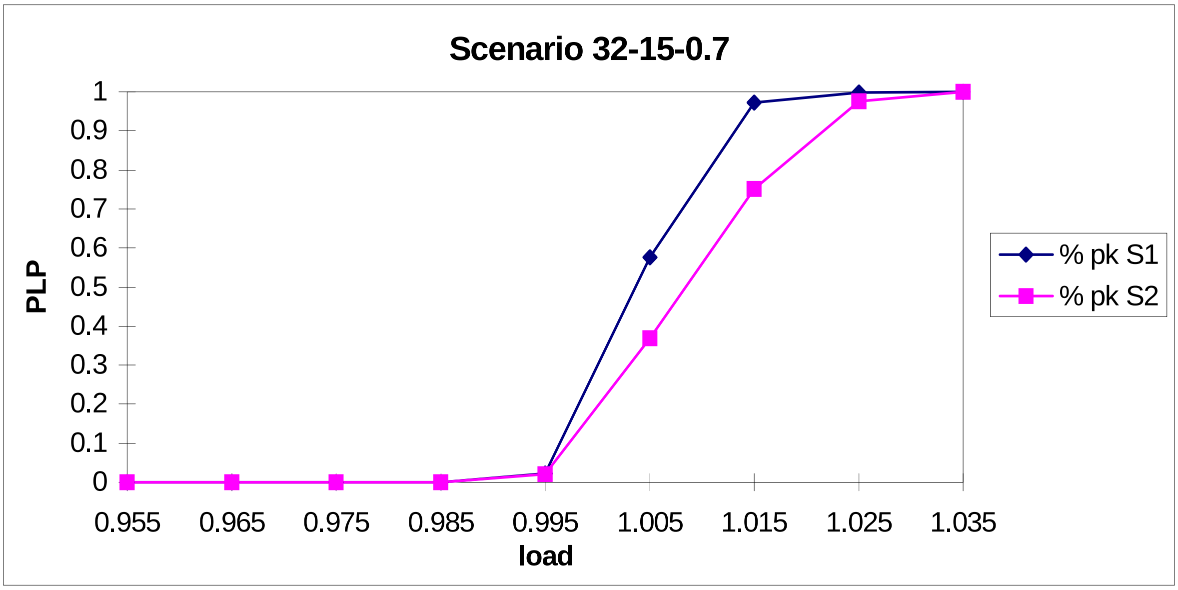

Figure 3.Packet losses for P = 32, K = 15, a = 0.7

The following figures show, with more detail, how the number of lost background cells is reduced when the discarding mechanism is applied.

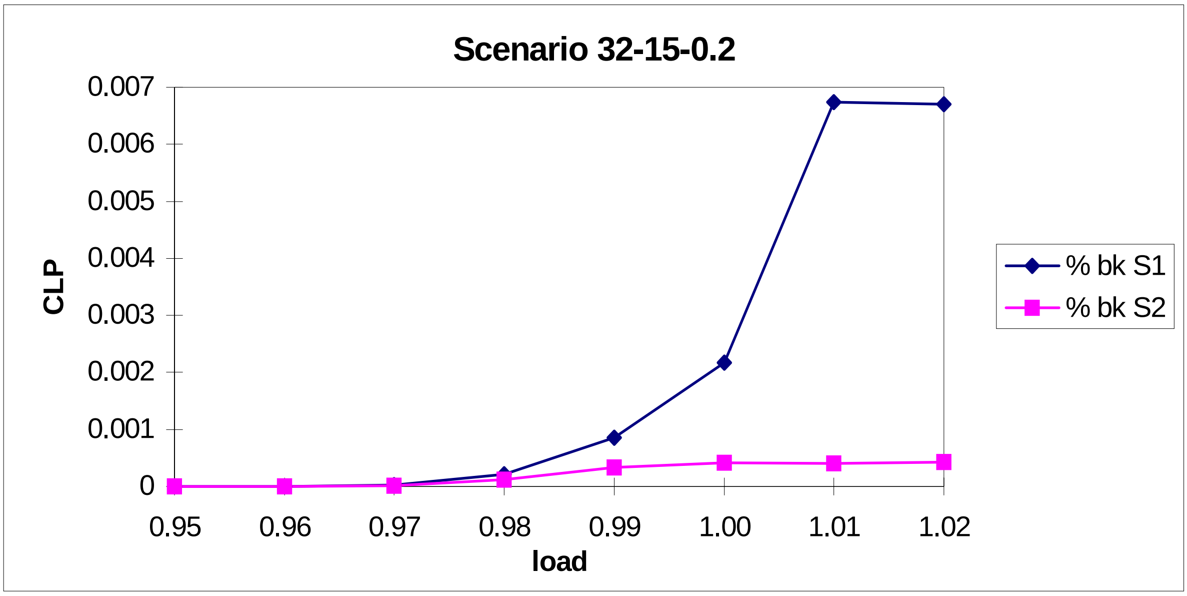

Figure 4. Background lost cells for P = 32, K = 15 , a = 0.2

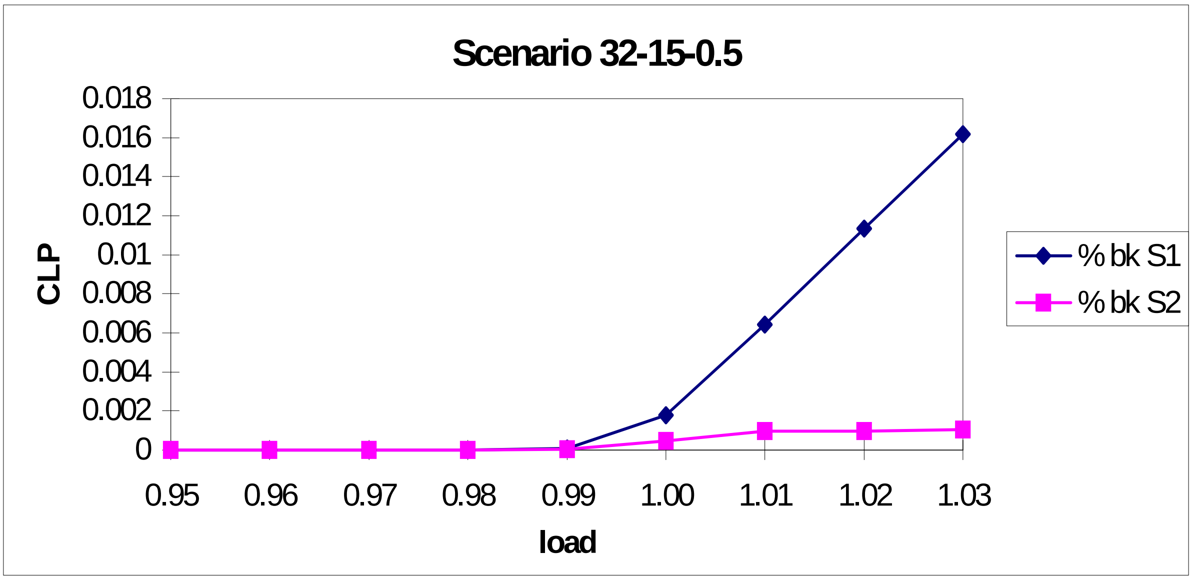

Figure 5. Background lost cells for P = 32, K = 15 , a = 0.5

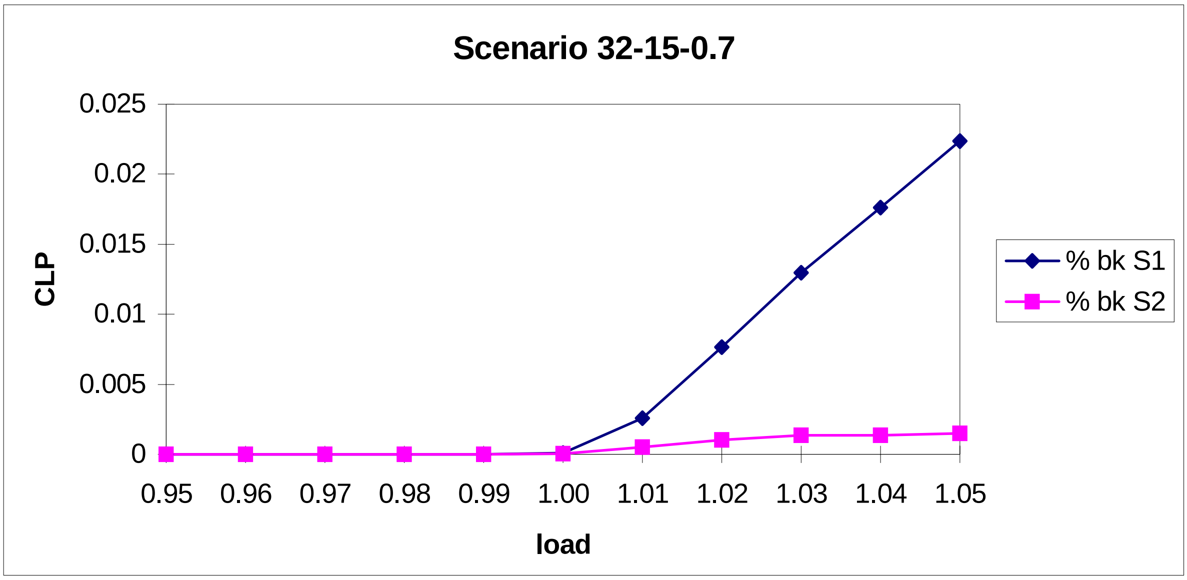

Figure 6. Background lost cells for P = 32, K = 15 , a = 0.7

The figures show that when the packet is short the mechanism can reduce the number of lost packets significantly.

The behaviour of the background lost cells is similar in the sense that when the mechanism is applied the number of lost cells is reduced dramatically.

These improvements are more remarkable as the reference connection load is higher.

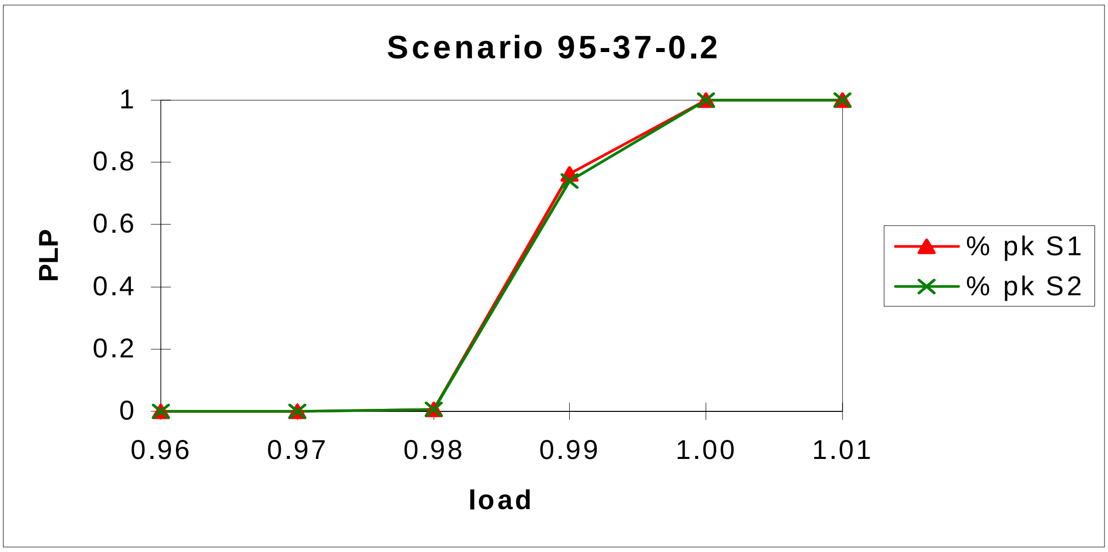

Figure 7. Packet losses for P = 95, K = 37 , a = 0.2

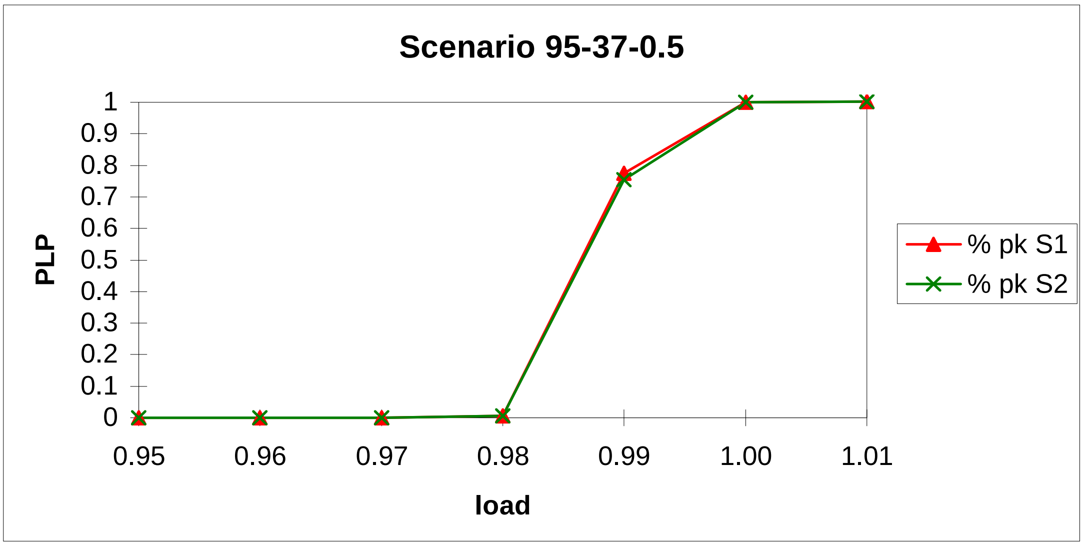

Figure 8. Packet losses for P = 95, K = 37 , a = 0.5

Another LAN-to-LAN connection over ATM that has been simulated is a Token Ring connection at 4 Mbps, that is, P = 95, K = 9, with different reference connection load figures, a = 0.2, 0.5, and 0.7. The results obtained are depicted in figures 7, 8 and 9.

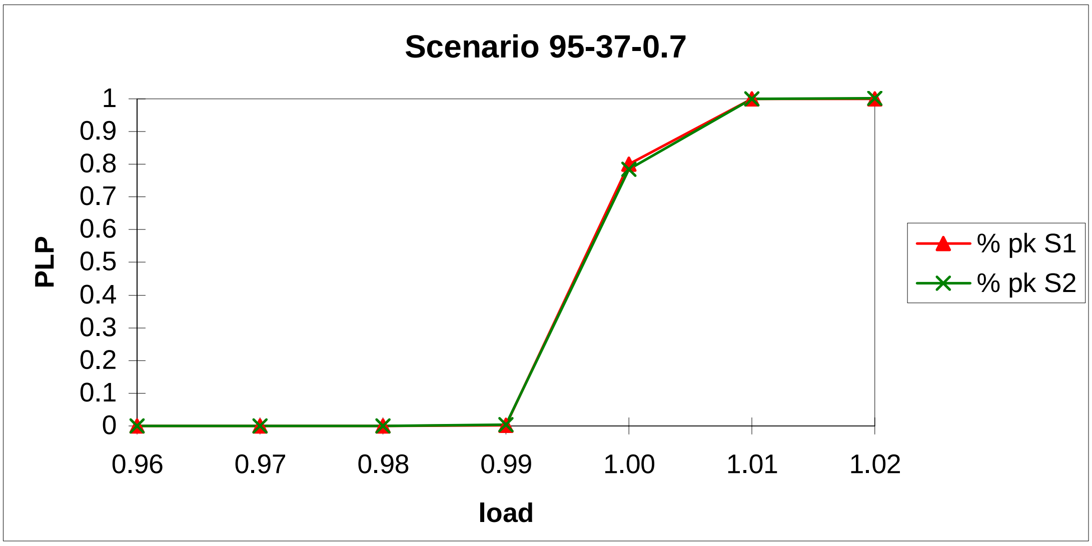

Figure 9. Packet losses for P = 95, K = 37 , a = 0.7

The values P = 380, K = 9, represent also a Token Ring, but this time at 16 Mbps. Also different loads have been considered. Figures 10, 11, 12 and 13 show the results obtained.

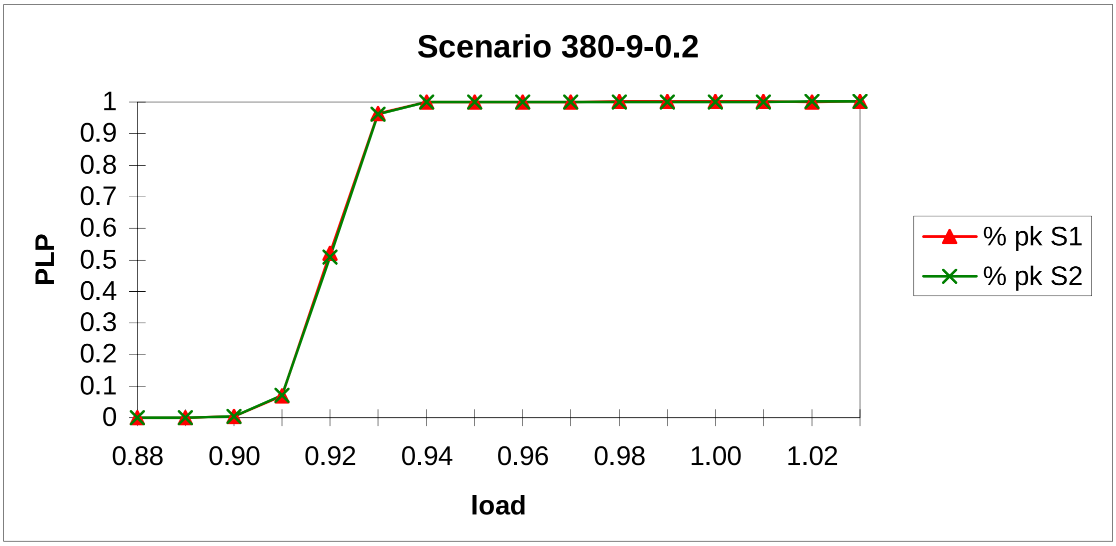

Figure 10. Packet losses for P = 380, K = 9, a = 0.2

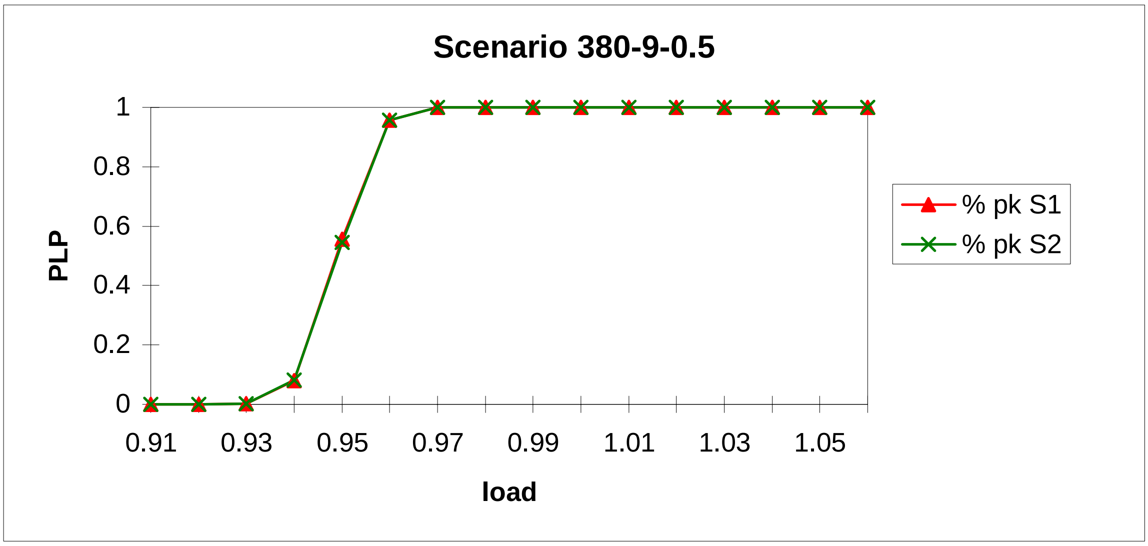

Figure11. Packet losses for P = 380, K = 9, a = 0.5

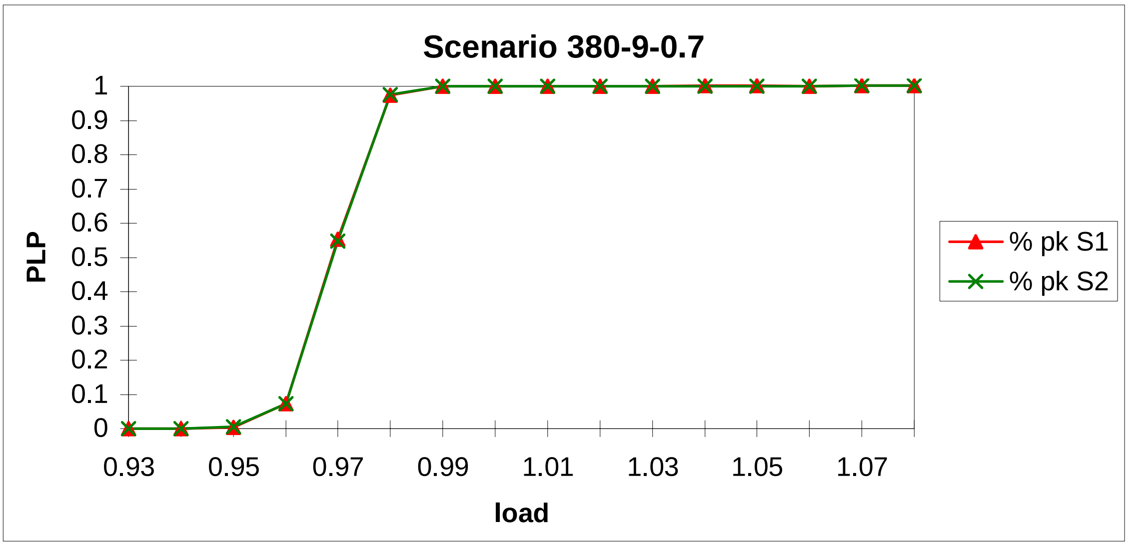

Figure 12. Packet losses for P = 380, K = 9, a = 0.7

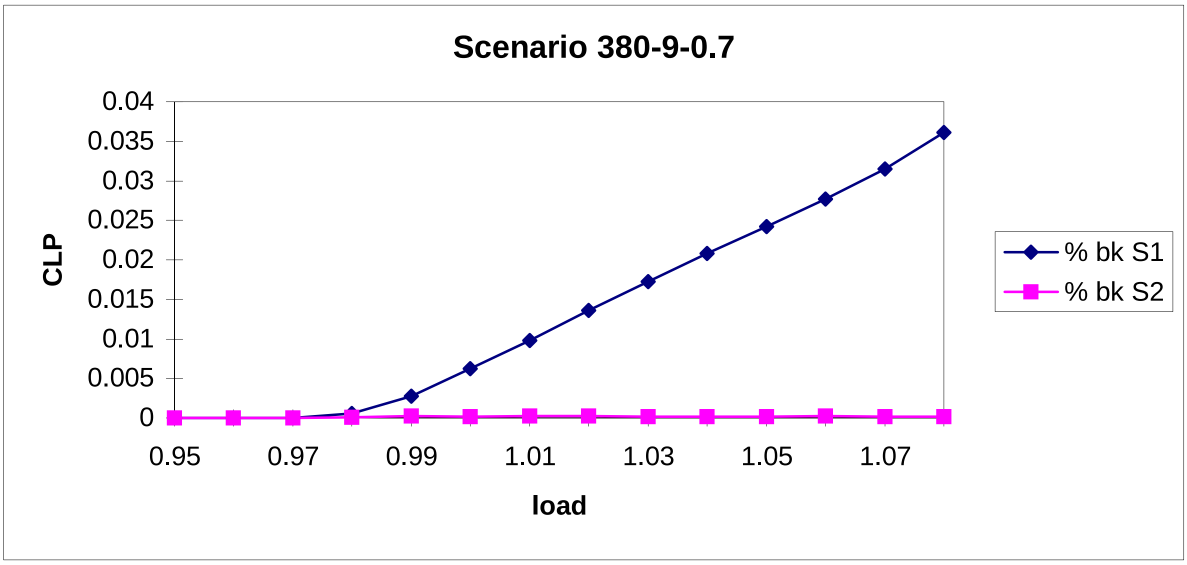

Figure 13. Background lost cells for P = 380, K = 9, a = 0.7

Looking at the results obtained for the Token Ring scenarios, it can be seen that when the packet is long and the rate of the connection is high the mechanism does not provide a significant improvement reducing the number of lost packets but, on the other hand, it always reduces significantly the number of background lost cells, alleviating the effect of the congestion on the other connections that share the link.

5. CONCLUSIONS

We can note that the application of the discarding proposal reduces the number of lost packets because the buffer occupancy is cut down. Best results are obtained for short packets and high load reference connections. On the other hand, this discarding mechanism performs a good congestion control mechanism, and many less background cells are lost when it is applied. It is remarkable the reduction of the number of lost cells corresponding to the connections that share the multiplexing buffer with the reference connection when the packets are long.

Results show that the application of the proposed discarding algorithm does not increase the number of lost packets, but, in fact, reduces it because the buffer occupancy is cut down. The mechanism proposed in this paper is to be applied at the ATM Level and must be implemented in hardware.

In next studies we take into account the fact that that the background traffic is also packet oriented. The reduction of background cell losses will improve significantly the performance of the system, reducing the Packet Loss Probability of the connections.

6. REFERENCES

1 [I.121] ITU-T Recommendation I.121, Broadband aspects of B-ISDN. Geneva 1991.

2 [I.361] ITU-T Recommendation I.361, ATM Layer Specification. March 1993.

3 [I.363] ITU-T Recommendation I.363, B-ISDN ATM Adaptation Layer (AL) Specification. March 1993.

4 [ATMF] The ATM Forum, ATM User-Network Interface Specification, Version 3.0, September, 1993.

5 [ROFL94] Allyn Romanov and Sally Floyd. Dynamics of TCP Traffic over ATM Networks. ACM SIGCOMM London, September 1994.