ACTIVE CELL DISCARD MECHANISM IN ATM NETWORKS[*]

Josep Mangues-Bafalluy and Jordi Domingo-Pascual
e-mail: jmangues@ac.upc.es and jordid@ac.upc.es
Universitat Politècnica de Catalunya. Departament d'Arquitectura de Computadors.
Campus Nord. Mòdul D6. Jordi Girona 1-3. 08034 Barcelona.
Phone: +34 3 4016981 Fax:+ 34 3 4017055

Abstract

When dealing with packet oriented communications over an ATM network, some problems arise. The most important one is that of having to discard an entire packet when only one ATM cell of that packet has been lost. Thus, we can see that the important parameter in this type of communications is not the Cell Loss Ratio (CLR), but the Packet Loss Ratio (PLR).

This paper presents a general and efficient congestion control mechanism for packet oriented communications over ATM. The proposed mechanism is called Active Cell Discard (ACD) and it derives from the Partial Packet Discard (PPD) and Early Packet Discard (EPD). Its objective is to use a threshold in the buffer to control that there is always enough free space to queue the set of initial cells of every packet. The usefulness of this mechanism becomes clear if we think that critical information of the connection (like synchronization or control) is located in that set of initial cells.

Introduction

Maybe the most important characteristic of an ATM network is its flexibility, i.e. the ability of the network to accept all kinds of traffic: real-time and non-real time, interactive and non-interactive, connection-oriented or connectionless oriented ... In fact, the goal of ATM is to offer a unique bearer service able to deal with heterogeneous traffic.

In a general ATM communication, one of the important parameters that controls the performance of the network is the Cell Loss Ratio (CLR). However, when packets are transmitted, the unit of information that must be preserved by the network is no longer the cell but the packet, because once one cell of the packet is lost, that packet is not useful for the end-user. But in that case, the remaining cells of the lost packets are wasting network resources unnecessarily.

The objective of selective cell discarding mechanisms is either reacting or preventing congestion by making the waste of network resources as low as possible. This is accomplished by selectively eliminating those cells belonging to a lost packet.

As a continuation of the work presented in (DOM 1995a) and (DOM 1995b), this article presents a congestion control mechanism which also derives from the PPD and EPD mechanisms (ROFL 1995). This mechanism is called Active Cell Discard (ACD).

The main objective of this article is to show the behavior of the ACD mechanism without the influence of upper layer protocols (like TCP) and to compare it with other mechanisms in the same simulated conditions.

In the next section, some of the results obtained in ROFL 1995 are stated. Then, we present the Active Cell Discard Mechanism. And the source model implemented is discussed in the following section along with the considerations made in our simulations. We also show the results of our simulations and finally we present our conclusions and the work to be done.

---

[*] This work has been funded by the Spanish Ministry of Education (CICYT) under grant number TIC95-0982-C02-02 and by the Catalonian Research and Universities Council (CIRIT 1995SGR-00464).

Cell discarding policies

Some works dealing with discarding policies have been developed by Romanow and Floyd (ROFL 1995). In that paper, some simulation results are presented for the Partial Packet Discard (PPD) and Early Packet Discard (EPD) mechanisms. They show the performance of TCP over plain ATM and compare it with TCP working over a packet switched network (this latter case will be referred as TCP packet).

The results obtained for TCP/ATM are very poor in comparison to those obtained for TCP packet and their performance depends on the number of TCP connections, the TCP window size, the buffer size and the packet size. But TCP packet maintains a high effective throughput even if some of these parameters are changed. The PPD and EPD mechanisms are presented as possible solutions to solve this unwanted behavior.

Their work showed that both mechanisms help to attain a much higher effective throughput than in TCP over ATM without discarding policies. The solution that presented the best behavior of both was EPD which guarantees a performance similar to that of TCP packet except when short buffers are considered.

Active Cell Discard (ACD) Mechanism

If our mechanism has to cope with congestion in any condition all kinds of traffic must be considered for its evaluation. Real-time communications don't allow either the retransmissions or the loss of packets. Therefore, it could be interesting to guarantee that at least a portion of the packet arrives at the receiver and the packet is not completely discarded. The objective is to maintain the connection within the minimum QoS required by the end-user using the information carried by this portion of the packet. In that way, a minimum cell rate is always guaranteed.

Some reorganization of the information may be needed to put the most relevant information at the beginning of the packet.

Besides, today's protocols present a bi-modal behavior, they usually use long packets for user information and short packets for signaling and control purposes (e.g. IP). The user information being important it is not as important as the control one, which is vital to maintain the connection within the contracted QoS. The mechanism proposed here would be very appropriate to guarantee this information to reach the end even with high congestion levels as the results show.

Another potential use would be to maintain that at least the synchronization packets in a video connection are received. In that way, the user would not have to spend any time trying to recover the lost image.

The contract between the user and the network includes the Peak Cell Rate (PCR), the Cell Delay Variation (CDV)(depending on whether it is a delay variation-sensitive service) and the number of cells of the packet that must be preserved (Pmin).

As in the EPD and PPD, the buffer occupancy (q) must be monitored by the switch, and it should keep the state for each VC. The switch needs a counter per VC. And some additional hardware is needed to implement the different cases that are found when applying this mechanism.

ACD is implemented by using AAL5 like PPD and EPD. The counter of each VC is not reset until the switch receives the last cell, which is not dropped. This cell marks the end of the packet and can be recognized examining the ATM-layer-user-to-user (AUU) parameter in the ATM cell header. When it is set, the present cell is the last of the packet.

The threshold (Q_t) is needed to always have enough space for the Pmin packets and to prevent congestion. And $Q_t=s*Q$ ($0<s<1$), where Q is the buffer size.

For connections with high Delay*Transmission_rate product, the amount of cells crossing the network is very important. As a consequence, the advantage of preventing congestion rather than acting after it appears is significant. In some way, ACD uses its threshold to prevent congestion from appearing.

The mechanism works as follows:

Be n the packet cell counter value per VC.

1. If $0<n<=Pmin$ and $q<Q$ the cell is not discarded. A Pmin cell is only discarded when there

is no space in the buffer. In this case the remaining cells are discarded until AUU indicating the end of the packet is found.

2. If n>Pmin and q<=Qt the cell is queued.

3. If n>Pmin and q>Qt the cell is discarded, and all the remaining cells, too.

As it can be deduced from these three points, the threshold is used to reserve the space between the threshold and the buffer capacity for Pmin cells only. The selection of the threshold value is a trade-off between the amount of user information that crosses the switch and the guarantee that none of the Pmin packets is lost. This value depends on the traffic parameters with which the switch will work. More work has to be done in that direction.

An estimation of the required buffer can be approximated analytically. We have N active sources. If we define K_i , which characterizes the transmission rate of the source in the following way:

$$K_i=C/PCR_i \qquad \text{for } 1<=i<=N$$

were PCR_i is the Peak Cell Rate for the i-th source and C is the capacity of the link. We establish that the size of the minimum packet to be preserved for the i-th source is $Pmin_i$ cells. Then, the minimum contention period, i.e. the time during which the cells coming from the sources can arrive simultaneously to the switch, is:

$$T = \min_i (P\min_i * K_i) \quad \text{for } 1<=i<=N$$

and the buffer size to accommodate all this traffic and not loosing any of these cells is:

$$Q = T * \left(\sum_i \frac{1}{K_i} - 1 \right) \quad \text{for } 1<=i<=N$$

Evaluation of the ACD mechanism

Our working environment is composed by one reference connection and some background sources. The number of background sources is determined by their traffic characteristics and the global load that we want to attain in the simulated link.

The characteristics of the reference connection can be entered independently from the background sources which are treated as a group. The traffics generated by the reference and the background connections are mixed and generates a cell flow which arrives to the output queue of the switch where the discarding mechanism is applied.

Traffic model

Each source is modeled by an ON-OFF model. The parameters that characterize each source are:

- P: The packet length, i.e. the number of ATM cells per packet.
- K: Measures the transmission rate of the source as defined in the preceding section.
- a: The load offered by the source referred to its maximum capacity (a=1) that would be obtained if we transmitted one cell every K during all the connection. This latter case corresponds to a CBR source. Then, this parameter measures the burstiness of the source.

As it can be deduced from the above parameters, the length of the burst -the time during which the source is in busy state- is P*K cells.

Other parameters that determine the switch behavior are:

- Q: The size of the buffer.
- Q_t: The threshold.
- r: Is the probability of transition from idle state to busy state. It is calculated based on the parameters entered for each source to obtain the desired global load and to assure that the mean sojourn time in idle state follows an exponential distribution corresponding to a Poisson arrival Process.

ACD in IP over ATM

As we explained in the preceding sections, the aim of the ACD mechanism is to work well with all kinds of traffic crossing an ATM network. In particular, if we consider data traffic, the most currently used protocol is IP.

A study to characterize the traffic traveling through the Spanish academic ATM network (RedIRIS) is reported in (CAST 1996). This network uses IP over ATM to interconnect Ethernet LANs at both ends. The measurements show that most part of the information transfer is done using short packets. In particular, as it can be observed in figure 1, almost the 60% of the traffic travels in packets that can be mapped into 3 ATM cells.

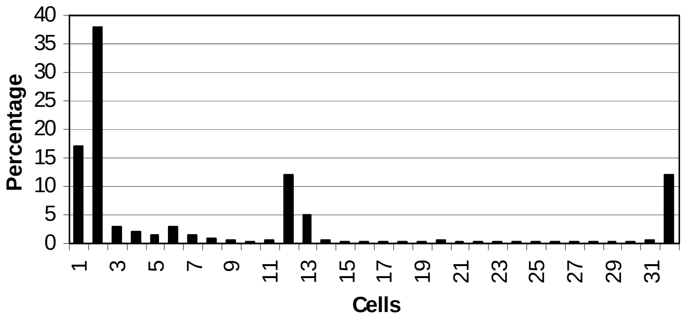

Fig. 1: Distribution of the length of IP packets mapped into ATM cells

If we analyze the nature of the IP protocol, we see that the control packets are among these short packets, and hence, the interest of maintaining them with a mechanism that guarantees very low loss probability for them, like ACD.

The results obtained in our simulations show that the difference between the entire packets' loss probability for any of the mechanism considered (EPD and ACD) and the Pmin loss probability could be around five orders of magnitude.

In figure 2, the packet loss probability for entire packets as well as for Pmin packets is presented for P=35, Pmin=3, K=10, Q=20 and a threshold of 0.8*Q. It must be remarked that the simulation considered very high load levels ranging from 88% to 100% of the total capacity.

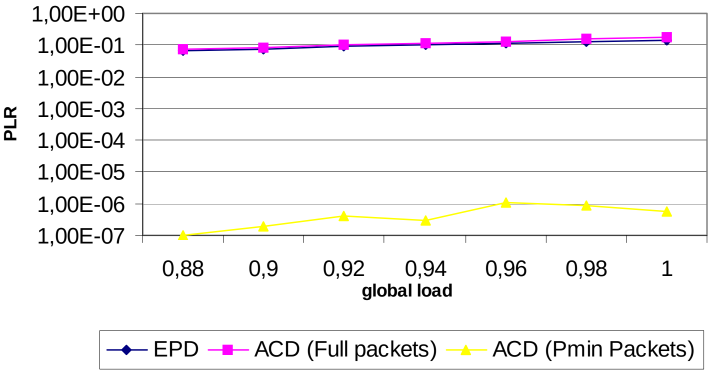

Fig. 2: Packet Loss Ratio (P=35,K=10,a=0.5)

Simulation environment

The packet size chosen for our simulations was is P=35 cells, which roughly corresponds to an Ethernet MTU (1500 bytes).

The intercell space for a connection (K) is chosen to be 10, meaning that in busy state, the source generates one cell every ten. That is, the Peak Cell Rate of our connection is one tenth of the capacity of the link.

And in general, a is taken to be 0.5, which indicates that it is a bursty source generating at the 50% of its capacity.

The buffer size (Q) is chosen quite small (20 cells) to have enough space to deal with cell level congestion (due to multiple cells instantaneously arriving at the same cell interval). But at the same time, it is not very large, because the delay (and cell delay variation in case of CDV-sensitive communications) would be unacceptable as it is the case for long buffers.

The threshold (Q_t) is fixed at 80% of the capacity of the buffer (16 cells). It was not fixed at half the buffer size as in (ROFL 1995) because their objectives are not the same. In our mechanism, the threshold is in charge of preserving enough free space in the buffer to accommodate all the Pmin packets.

Anyway, these values -Q and Q_t- are a function of the kind of traffic being transmitted and further work has to be done to determine the best relationship between the threshold and the traffic parameters.

Another consideration about these simulations is that the link is submitted to a very high load (ranging from 88% to 100% of the link capacity) and as a consequence the congestion is very important. In a real network, this situation will be very unusual, but the objective of this study is to show the efficiency of the ACD mechanism when the worst possible conditions appear.

Conclusions

The packet losses obtained with this mechanism are quite a bit worse than those obtained with PPD and EPD, but the most significant result is that even in highly congested networks, this mechanism guarantees a minimum flow of cells in the form of Pmin packets, and as a consequence a minimum QoS. The figures show that the losses of these packets are some orders of magnitude below the losses of entire packets.

Further investigations to relate the size and threshold of the buffer, and Pmin with the traffic crossing the network must be done to obtain the maximum efficiency of this mechanism. It must be also studied with heterogeneous traffic, which is the closest to real traffic.

References

(ATMF 1993) The ATM Forum, ATM User-Network Interface Specification, Version 3.0, September 1993.

(DOM 1995a) Domingo J and Sánchez J. ‘Quality of Service of Packet Data Streams over ATM Networks’. International Teletraffic Seminar. ITC Sponsored. Editors: Goldstein B., Koucheryavy A., Scheneps-Schneppe M. St. Petersbourg (Russia): 404-414, July 1995.

(DOM 1995b) Domingo J. and Sánchez J. ‘Selective Cell Discarding Mechanism for a Packet Oriented Service on ATM Networks’. Computer Architecture Department. Universitat Politècnica de Catalunya (UPC). Research Report UPC-DAC-1995-47. December 1995.

(ROFL 1995) Romanow A. and Floyd S. ‘Dynamics of TCP Traffic over ATM Traffic’. IEEE Journal on Selected Areas in Communications 13 (4): 633-641, May 1995.

(CAST 1996) CASTBA Report. Polytechnic University of Madrid (UPM). April 1996.