Cell Loss Clustering Prevents Avalanches in ATM Networks

Andrew Shaw*
NE43-257
shaw@mit.edu
Computation Structures Group

1 Introduction

ATM networks handle congestion by discarding cells that overflow switch queues, leaving the responsibility of reconstructing damaged protocol data units (PDU) to higher-level protocols. The loss of a single ATM cell from a PDU can require the retransmission of the entire PDU, which can be disastrous because the loss of a few ATM cells can cause the loss of a much larger amount of PDU data – this is the so-called "avalanche effect" [3]. The avalanche effect has caused much concern in the networking community regarding the effectiveness of ATM for transporting data network traffic because of the short cell length of ATM cells relative to the long packet lengths of typical PDUs.

Our simulation results and those of others [2] [1] indicate that cell losses are highly correlated, and tend to occur in clusters. Cell losses which occur closely in time are likely to be from the same PDUs, thus alleviating the avalanche effect.

Hardware support has been proposed to handle avalanche effect, including switches which are aware of the higher level protocols, and which drop cells of the same PDU subsequent to the first lost cell – this strategy is called *Partial Packet Discard* (PPD) [4]. We simulate PPD and conclude that because of the high degree of cell loss clustering, PPD only marginally improves performance over plain ATM switches.

2 Simulation Results

To simulate the relationship between ATM cell losses and PDU losses, we used `netsim`, an event-driven network simulator developed in the Advanced Network Architecture group at MIT. We simulated the effect of a n-input, n-output ATM switch on PDUs arriving at Poisson intervals segmented into ATM cells. The topology of the simulated network is shown in Figure 1. The switch is output queued, and the speed of the output links is equal to the speed of the input links. We measured the number of ATM cells and PDUs lost, and the time-dependent behavior of those losses.

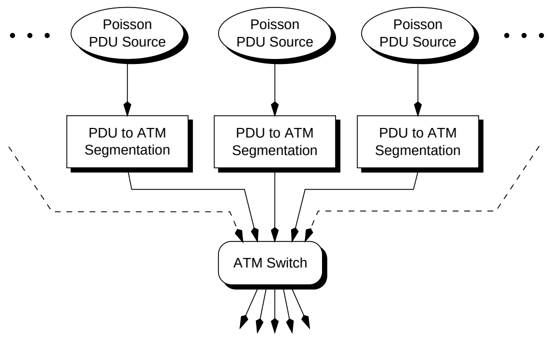

Figure 1: *Simulation setup. The Poisson sources produce PDUs which are then segmented into ATM cells and sent to the ATM switch.*

Clustering of Cell Losses

We looked at the intervals between ATM cell losses due to queue overflow. The simulation parameters we considered were for PDUs of length 700, and a mean utilization rate of the switch of 90%. Figure 2 shows the distribution of the interarrival periods of cell losses. *About 85% of the cell loss interarrival periods were less than the output period of the PDU to ATM Segmentation Unit*. This indicated that cell losses occurred in clumps, and were likely occuring from the same PDUs.

Architectural Support for Avalanche Effect

We simulated the effectiveness of PPD architectural support for reducing the avalanche effect. The plain ATM switch discards cells only on queue overflow, whereas the PPD switch drops subsequent cells in the same PDU to of first lost cell, except for the last cell. The arrival of the last cell allows the receiver to realize that the PDU has been corrupted and resynchronize.

Our measure of effectiveness is the *partial packet discard rate*, which is the percentage of the ATM cells in the lost PDUs which were actually discarded (due to queue overflow and/or PPD discard) by the respective switches. For each simulation, we varied the size of the PDUs and the length of the queues. Following [4], we tried PDU lengths

*Faculty Supervisor: Prof. Arvind. The research described in this paper was funded in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-92-J-1310. This paper describes joint work with Prof. David Tennenhouse.

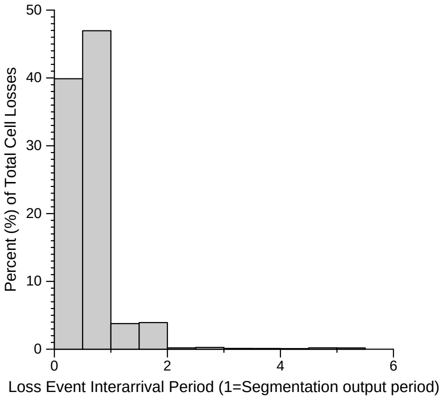

Figure 2: *Histogram of ATM cell loss interarrival times.*

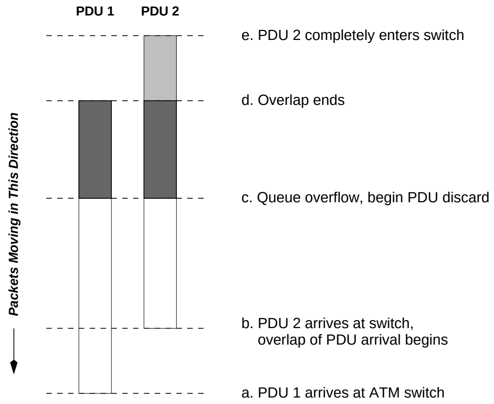

Figure 3: *Scenario for PDU arrival at the switch. Discard regions are in grey.*

of 512, 1500, 4352 and 9180 bytes and we tried queue lengths of 256, 1000, 2000, and 5000 cells. The results are summarized in the following table:

Partial Packet Discard Rate	
plain ATM switch	PPD ATM switch
$.333 \pm .050$	$.500 \pm .047$

Note that the results are very counter-intuitive, in that all of the different simulation parameters gave us similar Partial Packet Discard rates, modulo switch architecture.[1] *Although PPD does discard additional ATM cells which would not have been discarded in the plain ATM case, 66% of the cells which are discarded in PPD would have been discarded anyway.* In the next section, we provide a very simple analysis of the behavior shown in the simulations.

3 Analysis

In our analysis, we consider the overlapping arrival of two PDUs at the switch, and we assume that the cell arrivals are consecutive and occur at the link speed; this will provide a first-order approximation of the behavior at the initial switch from the Segmentation Unit.

Figure 3 shows the anatomy of a potential cell loss. Queue overflow can only happen between point b and point d. In general, it can happen at any point between the two, and the probability is uniform over that region. On average, it will begin at the midpoint of point b and point d.

For the PPD switch case, all of the remaining cells in PDU 1 or PDU 2 will be discarded (except the last cell). However, it is not clear which PDU will be discarded – half the time, it will be PDU 1, half the time, it will be PDU 2. On average, that will result in a 50% Partial Packet Discard rate for the PPD switch.

[1] The ATM cell loss rate and the PDU loss rate do decrease significantly as the queue lengths increase and as the PDU lengths decrease, but the Partial Packet Discard rate is constant across different queue lengths and PDU lengths.

For the plain ATM switch case, note that queue overflow is equally likely to happen anywhere during an overlap, implying that overlapping regions of longer duration are more likely to cause queue overflow than overlapping regions of shorter duration. However, distribution of lengths of overlapping regions is uniform over all possible overlaps. For a PDU length of n, the distribution of overlaps of length l is $P(l) = \frac{2l}{n^2}$. The median loss will occur at the midpoint of point b and point d, so the median length of the loss region will be approximately $\int_0^n \frac{l}{2} P(l) = \frac{n}{3}$, or 33% for the plain ATM switch case.

4 Conclusion

Our simulation results indicate that because of the natural clustering of ATM cell losses, *avalanche effect* is a much less serious problem than has has been proposed in analyses where cell losses were presumed to be uncorrelated [3]. Furthermore, architectural support in the form of switches supporting Partial Packet Discard do not significantly improve Partial Packet Discard rates.

References

[1] Amit Bhargava and Michael G. Hluchyj. Frame Losses Due to Buffer Overflows in Fast Packet Networks. In *Proceedings of IEEE Infocom '90*, pages 132–139, June 1990.

[2] Andre B. Bondi and Wai-Sum Lai. The influence of cell loss patterns and overhead on retransmission choices in broadband ISDN. *Computer Networks and ISDN Systems*, 26(5):585–598, January 1994.

[3] Israel Cidon, Jeff Derby, Inder Gopal, and Bharath Kadaba. A Critique of ATM from a Data Communications Perspective. *Journal of High Speed Networks*, 1(4):315–336, November 1992.

[4] Allyn Romanow and Sally Floyd. Dynamics of TCP Traffic over ATM Networks. In *Proceedings of ACM SIGCOMM '94*, pages 79–88, 1994.