An Analysis of Space Priority Queueing in ATM Switches

Karen L. Reid   Richard B. Bunt
Department of Computational Science
University of Saskatchewan
Saskatoon, Canada S7N 0W0

July 12, 1993

An Analysis of Space Priority Queueing in ATM Switches

Abstract

High-speed packet switched networks such as B-ISDN require algorithms to manage the flow of traffic into and through the network. Each switching node and network access node is susceptible to overload. End-to-end congestion control algorithms will do much to ensure that overload periods are minimized, but intelligent buffering decisions at internal nodes may aid in managing congestion by shifting cell loss to lower priority traffic.

This paper investigates how space priority queueing might aid in supporting acceptable quality of service for different classes of traffic on a high-speed packet switched network. Simulation was used to examine the behaviour of three priority queueing algorithms operating at the output buffer of a switch under realistic traffic models, and showed that the performance of these algorithms depends on the shape and type of the traffic. The results of this research suggest that space priority algorithms, when combined with other congestion control algorithms, can aid in supporting acceptable quality of service for the users of the network.

Keywords: performance evaluation, simulation, ATM, buffer management

1 Introduction

The high-speed integrated service networks of the future will support a wide variety of applications simultaneously. The characteristics of the network traffic generated by these applications will be equally diverse, and this diversity presents new challenges. Several characteristics of the traffic heavily influence decisions about the types of control measures the network can use. Delay-sensitive services require short network delays and, therefore, cannot afford to wait for cells (packets) to be retransmitted. Not every cell need be delivered, however. On the other hand, loss-sensitive services require that every cell arrive at its destination, but some delay might be tolerable. Traffic is expected to be bursty, a feature that aggravates congestion problems. Relying on reserved channels for all traffic, however, leads to inefficient use of bandwidth. The impact of these traffic characteristics on performance stresses the importance of realistic models in studies of network congestion.

Managing congestion is a significant component of meeting quality of service (QOS) demands. Congestion can and will be managed at several levels in future networks. The most prominent approach, end-to-end congestion control, adjusts the rate at which sources transmit to prevent or react to overloading on the network. Access control algorithms and end-to-end congestion control algorithms may not be sufficient to prevent congestion in a high speed integrated services network, and more intelligent decisions about dropping cells inside the network can help to support end-to-end algorithms. This paper examines a method of managing congestion inside the network to support end-to-end control schemes through the use of priority-managed buffers at the output links of the switches.

This research considered three algorithms that use different techniques to determine which cells to drop as the switch buffer fills. Each algorithm reflects a slightly different interpretation of how cell priority information may be used. The performance of the algorithms is affected by the portion of the traffic in each priority class, the total load on the switch, and the shape of the traffic. These algorithms do not reduce the total cell loss. Instead, their function is to shift the cell loss to lower priority cells, thereby providing better protection to the quality of service expected by high priority applications. Shifting cell loss from high priority connections to low priority connections may also allow a higher level of network utilization by encouraging more aggressive end-to-end congestion control algorithms. Simulation was used to examine the behaviour of the various algorithms under a variety of workload mixes.

The CCITT has reserved one bit in the ATM (asynchronous transfer mode) header for cell loss priority [1]. This bit can be used to assign priority to cells belonging to different classes of traffic. Initial thoughts were that two priority classes may not provide sufficient discrimination, and thus sufficient protection for the high priority class. For this reason, the use of four priority classes (two priority bits) was also considered in this investigation.

The paper is organized as follows. Section 2 describes the three priority queueing algorithms considered in this study and briefly summarizes the results of some earlier work. Because of the anticipated impact of workload characteristics on the performance of these networks, care was taken to incorporate realistic workload models into this investigation. Section 3 describes the workload models used in the simulations to represent voice, video, and data traffic. The simulation experiments and the results obtained are discussed in Section 4, and Section 5 presents the conclusions of the work.

2 Priority Queueing and Overload Control

When the capacity of an output link is exceeded by a short-term burst of cells, a buffer (or queue) can be used to reduce the cell loss. The decision of when to drop cells from the buffer (in the event of overflow) and the choice of which cells to drop are the subject of this investigation. If QOS requirements are to be met, the choice should be based on the priority of the cells in the buffer. Three algorithms are considered in this investigation: push out, partial buffer sharing, and selective packet discarding. These are examples of *space priority* queueing algorithms, where *space priority* addresses the management of space in the buffer (i.e. which cells to drop as the buffer fills). *Time priority* queueing algorithms operate on the order in which cells will be transmitted on the network as they leave the switch.

2.1 Push Out

The push out algorithm operates only when the buffer is full. When a low priority cell arrives at a full buffer, it is dropped. When a high priority cell arrives at a full buffer and a lower priority cell is in the buffer, then the lower priority cell is "pushed out," and the arriving high priority cell is added to the buffer. This approach is *work conserving* [7]; it drops cells only when the buffer is full. A parameter to the algorithm determines how many buffer cells to search for a lower priority cell. Searching a small number of cells incurs less overhead than searching the whole buffer.

This algorithm has been studied analytically [3, 6] using Poisson traffic models and two priority classes. Lin and Sylvester [7] also take into account possible correlation between the two traffic classes. All three studies show that when two priority classes contribute equally to the load, the cell loss for the higher priority traffic (Class 1) is dramatically reduced while the cell loss for the lower priority traffic (Class 2) increases slightly.

The simulation model used for this study gave similar results for the push out algorithm under a Poisson traffic model, but when bursty traffic models were used cell loss was much higher. One of the central goals of this research was to investigate the effect of more realistic traffic models.

Saito [11] studied the push out algorithm under an MMPP traffic model. His results were limited to a buffer of size 10 but show results similar to [3] and results using MMPP with the simulation model used in this study.

2.2 Buffer Sharing

The buffer sharing algorithm reserves a fixed number of buffer spaces for high priority cells. Cells from all classes share the buffer space until the number of occupied cells reaches a threshold. If the number of cells in the buffer exceeds the threshold then only higher priority cells will be accepted (as long as there is space).

Similar methods as in the push out scheme are used to study buffer sharing approaches. Kröner [6] shows that partial buffer sharing performs only slightly worse than the push out approach. Lin and Sylvester [7] show that show a reduction in cell loss of several orders of magnitude for class 1 traffic comes at the cost of only one or two orders of magnitude increase in cell loss for class 2 traffic (with Poisson arrivals).

Buffer management is simpler with partial buffer sharing (as compared to push out), but the performance of such an approach is dependent on the choice of the threshold and the proportion of the traffic in each priority class. It is not work conserving since cells may be dropped before the buffer is full. Therefore the total cell loss may be higher than if no priority queueing scheme was used. As the number of cells reserved for higher class cells decreases, the total cell loss will be closer to the cell loss with no priority algorithm.

2.3 Selective Packet Discarding

Selective packet discarding allows the system to choose which cells to discard. In this approach, cells may be discarded before the buffer is full. This approach is based on an understanding of the coding and packetizing methods used for video and voice applications. It has been shown [5] that consecutive cell loss produces a more noticeable error than losing the same number of cells distributed randomly throughout the cell stream. Petr and Frost [9] describe a cell discarding method in which a decision of which cells to discard is made every switch cycle. A cost function is associated with each priority class.

A simple algorithm has been chosen to represent selective packet discarding in this study. Two threshold values (B2, B1) define the overload and underload periods respectively (see

Figure 1). When the number of cells in the buffer exceeds B2, an overload control mechanism is activated until the buffer length drops below B1. The overload control mechanism decides which cells to drop based on a fixed probability of dropping a cell of each priority level. These probabilities are related to quality of service requirements and are specified per switch. It is assumed that the probability of dropping a low priority cell is greater than the probability of dropping a high priority cell. Note that the percentage of cells dropped depends on the traffic mix.

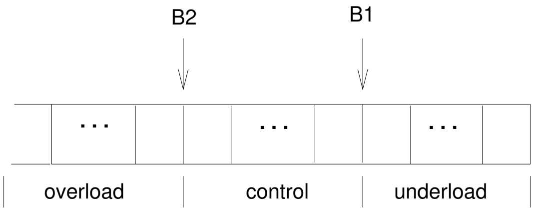

Figure 1: Selective Packet Discarding

3 Traffic Models

Traffic on high-speed networks will be bursty and correlated. Therefore, simple traffic models such as Poisson do not capture important characteristics of the traffic. Furthermore, since the mandate of high-speed networks such as B-ISDN is to integrate a wide variety of services, the traffic characteristics of each class of services will differ. An important goal of this research was to incorporate a wider range of traffic types into the investigation. Results of studies based on simplistic assumptions (such as Poisson traffic models) may have little relevance to practical networks.

Three traffic models were used in this study, representing three broad classes of traffic: voice, video, and data (each model represents an aggregate stream of traffic). Much work has been done in characterizing traffic of each of these three types in different networks. The three traffic models used in this study were drawn from previous work.

MMPP (Markov Modulated Poisson Process) has been used to characterize telephone voice traffic [4]. MMPP is an n-state model where the holding times and interarrival times in each state are exponentially distributed. For this study, a two-state model was used. This model captures burstiness still present in multiplexing of many voice sources.

This model was modified to enable it to label cells as high or low priority. Because MMPP models aggregate voice traffic, low priority cells are randomly distributed throughout the traffic stream and the proportion of low priority cells to high priority cells is determined by a parameter to the model.

The model used for data traffic is similar to MMPP in that it is a two-state model, but it has a deterministic rate for cell transmission during the active state and does not transmit cells during the idle state. This model is motivated by the recognition that bursts of data traffic usually consist of a single transfer of a larger piece of information and that cells generated by this transfer will likely be transmitted at a constant rate [13, 2]. The model chosen has an exponentially distributed idle period and the number of cells in a burst is geometrically distributed. During a burst, cells are transmitted at the peak rate.

The characteristics of video traffic depend to a large extent on the encoding algorithm used. The video model used in this simulation study is one developed by Maglaris *et al.* [8] based on measurements. The bit rate of a single source during the nth frame is modeled by a first-order autoregressive Markov process. The length of a burst is based on the number of frames per second and the cell rate during a burst is based on the number of cells (or bits) per frame. By varying these two values, traffic streams of various load levels can be obtained. This does not violate the characteristics of the model since Maglaris *et al.* [8] note that their model assumptions hold for a wide variety of video sources. Although the burst length of a video source is expected to be fairly long (on the order of 40 ms), this length is infeasible for practical simulation lengths. A burst length of 1.7 ms was used for this study. Longer burst lengths lead to smoother traffic at the switches, and if call acceptance algorithms ensure that the total traffic presented to the network does not exceed the network's capacity, then this should lead to lower cell loss.

Figure 2 highlights the different characteristics of each of the traffic models at a load level of 0.1. The mean bit rate for each aggregate traffic stream is 15 Mbps but the peak bit rates vary. The variation in arrival rates is much greater for the data model than for the other two models. The video model shows the least variation in the traffic rate, and this variation occurs over longer periods of time than in the other two models.

Determining which cells should be high or low priority is a difficult problem that depends on the philosophy of cell delivery. The experiments show, as expected, that the algorithms provide the best protection to high priority traffic when the proportion of high priority traffic is low. Therefore, if the chosen philosophy dictates that high priority cells should be guaranteed delivery, the approach to assigning priorities must limit the number of high priority cells. Another possible philosophy is assigning high priority to cells or connections with the lowest loss-tolerance. This would lead to a greater proportion of cells in the high priority class. Recall that the goal of priority queueing is to support end-to-end congestion schemes; therefore either philosophy is feasible and each has its merits. This set of experiments favours the second approach. Table 1 illustrates the priority class divisions used for these experiments.

The traffic on high-speed networks will have a large range of loss-tolerance levels. Yin *et al.* [14] suggest voice traffic could tolerate between 0.02 and 0.2 cell loss. Other authors

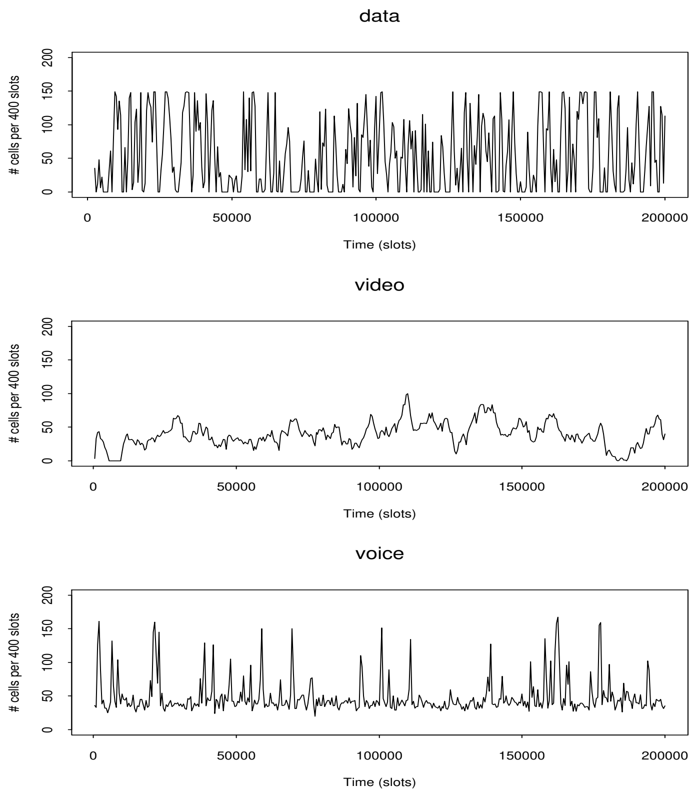

Figure 2: Shape of Three Traffic Models (load level = 0.1)

provide much lower figures for high quality video, around 10^{-10} [12]. Two priority classes may not be sufficient to capture this range. When four priority classes are introduced, it is possible to combine the philosophies discussed above. Class 1 traffic can be limited to a small proportion of the total load, virtually guaranteeing no cell loss, while the remaining three classes can absorb varying degrees of loss. Table 2 provides an example of a division of priority classes. Although data traffic can tolerate no cell loss, it is given lower priority than the video cells since data cells may be retransmitted if lost.

Table 1: Two Priority Classes

Class		Traffic types
1	delay and/or loss sensitive	Data High priority video High priority voice
2	best effort	Low priority video Low priority voice

Table 2: Four Priority Classes

Class		Traffic types
1	delay and loss sensitive	High priority video
2	loss sensitive	Data
3	delay sensitive	High priority voice
4	best effort	Low priority video Low priority voice

4 Simulation and Results

A single 8x8 output-buffered ATM switch is modeled. The input and output links carry 53 byte ATM cells and have a total bit rate of 150 Mbps. This data rate corresponds to proposals for Broadband ISDN networks [1]. The traffic stream from each of the 8 inputs consists of cells from each of the three traffic models destined for a single output. Modeling only one output port yields all the interesting information when separate buffers are attached to each output since the priority queueing algorithms operate at the point where cells enter the output buffer.

The buffer size was fixed for all the experiments. In a real system, the choice of buffer size will depend on the expected load level of the network and the expected level of burstiness in the traffic. Both the traffic characteristics and the size of the buffers affect the performance measure of interest, average cell loss. For this investigation the load level (0.87) and the buffer size (50 cells) were chosen to correspond to a cell loss on the order of 1% to 4%. Although this cell loss would be unacceptable in a real network, it made it possible to study the relative behaviour of the algorithms with a reasonable simulation length. Confidence intervals were in the range of $\pm 5\%$ to $\pm 10\%$, but these intervals were small with respect to the differences between results.

A number of combinations of the three traffic models were selected to determine the effect of the priority queueing algorithms when each of the models dominated the traffic stream. As expected, when the video model dominated, the traffic flow was smoother and experienced less cell loss. The same was also true to a lesser extent when the voice model dominated. The total cell loss was highest when the three models were equally represented. The overall shape of the traffic in this mix is smoother than when the data model dominates but the buffer is more likely to be partially full when the sharper data bursts arrive. For this presentation, two traffic mixes were chosen to illustrate the results of the priority queueing algorithms in more detail (complete results are given in [10]). Table 3 shows the parameters of these two mixes.

Table 3: Composition of Traffic Mixes (traffic types and two priority classes)

Name	% video	% voice	% data	Load	% Class 1	% Class 2
mix1	33.3	33.3	33.3	0.87	76	24
mix2	25	50	25	0.87	57	43

4.1 Experiments with Two Priority Classes

Figures 3, 4, and 5 illustrate the behaviour of the three algorithms using two priority classes. Cell loss for each class is measured as a fraction of the cells in that class that were lost. This gives an indication of how well each class is meeting loss requirements since the total cell loss is measured as a fraction of the total number of cells.

Because the push out algorithm (Figure 3) is work conserving, the total cell loss remains constant (0.026 for mix1 and 0.009 for mix2) when the number of cells to search is varied. As more cells are searched, a low priority cell (Class 2) is more likely to be found and replaced by a high priority cell (Class 1). The cell loss curves for both classes flatten as the number of cells to search increases, demonstrating that the overhead of searching the whole buffer for a cell to replace is unnecessary. Mix2 shows a lower total cell loss than mix1 because the traffic stream is smoother. The class loss curves flatten more quickly in mix2 because the proportion of high priority traffic is lower. Although these levels of cell loss are unacceptable in a real network, they serve to describe the relative behaviour of the algorithms while maintaining reasonable simulation lengths.

Figure 4 shows that the buffer sharing algorithm is more effective at protecting Class 1 traffic from cell loss than the push out algorithm, even when only a few cells are reserved for Class 1 traffic. However, this benefit comes at a high cost for Class 2 traffic. Also, note that the total cell loss decreases as the number of cells shared by both classes increases. This is most pronounced in mix2. The total cell loss becomes statistically indistinguishable from the total cell loss when no queueing algorithm is used when the sharing limit is 35 cells. Again, the steeper curves in mix2 are the result of a lower proportion of high priority cells.

Preliminary experiments with the selective discarding algorithm showed that the performance of the algorithm was much more sensitive to the probability parameters than to the

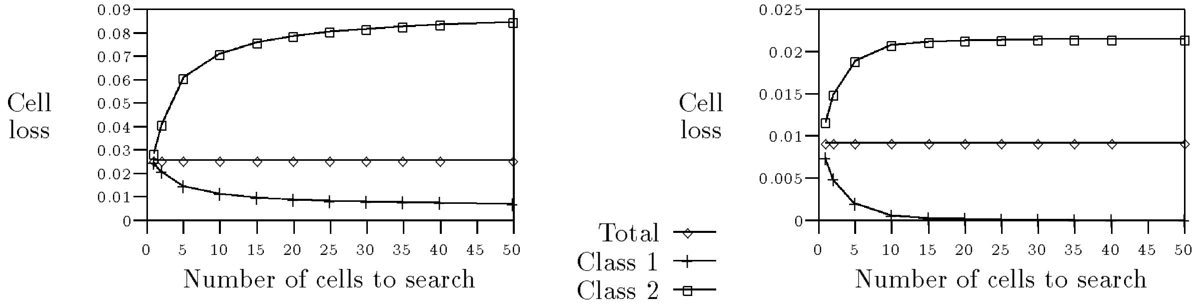

a) mix1 b) mix2

Figure 3: Push Out (two classes, fraction of class)

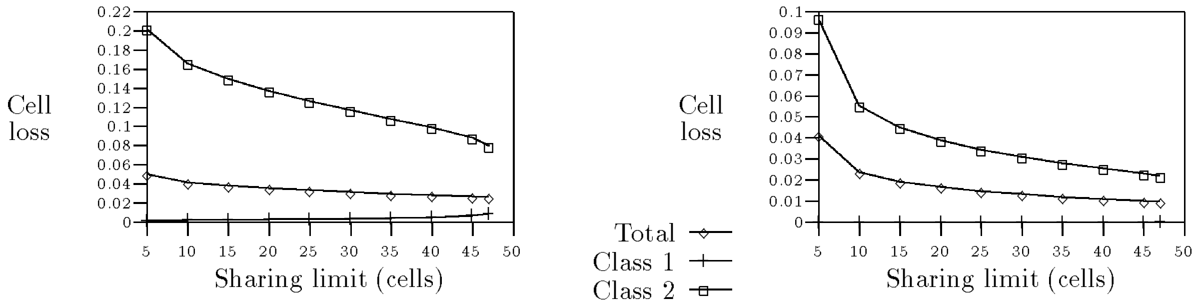

a) mix1 b) mix2

Figure 4: Partial Buffer Sharing (two classes, fraction of class)

threshold parameters. Figure 5 presents the cell loss for each class using the probability parameters shown in Table 4. As the probability of dropping a cell decreases, the cell loss for each class becomes indistinguishable from the total cell loss. The most protection for Class 1 traffic is achieved when the probability of dropping cells from Class 2 is high.

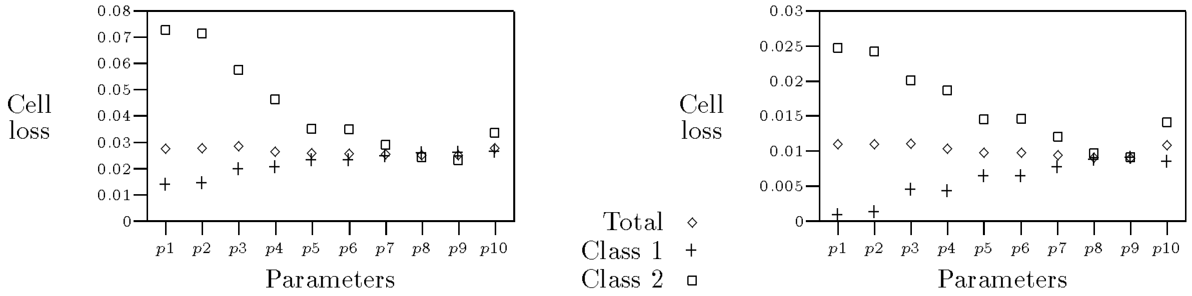

a) mix1 b) mix2

Figure 5: Selective Packet Discarding (two classes, fraction of class, B1=30, B2=45)

Table 4: Probability Parameters for Selective Packet Discarding

	p1	p2	p3	p4	p5	p6	p7	p8	p9	p10
P1	0.001	0.01	0.1	0.001	0.001	0.00001	0.0001	0.0001	0.0001	0.1
P2	0.5	0.5	0.5	0.2	0.1	0.1	0.05	0.01	0.001	0.2

In summary, the results of the experiment with the three space priority queueing algorithms with two priority classes show that protection of Class 1 traffic from excessive cell loss can be achieved to some degree. The performance of Class 2 traffic necessarily degrades, absorbing cell loss from Class 1. The goal of priority queueing algorithms is to ensure that the higher priority traffic does not suffer too much cell loss while the performance of the lower priority class degrades as little as possible.

The buffer sharing algorithm most closely adheres to the philosophy of guaranteeing almost no cell loss for Class 1; it drops lower priority cells before the buffer is even full. The push out algorithm applies this philosophy more conservatively; it allows the buffer to fill before dropping lower priority cells. The selective discarding algorithm follows the philosophy of trying to maintain QOS requirements for all classes. It assumes that all types of traffic can tolerate some cell loss so it tries to control the cell loss for each class by dropping cells before the buffer fills.

The push out algorithm and the selective packet discarding algorithm were limited in how much they could protect Class 1 against loss. Buffer sharing was able to provide the most protection for Class 1 traffic, but these gains were costly; Class 2 loss and the total cell loss was much higher than for the work conserving algorithm.

4.2 Experiments with Four Priority Classes

The division of the traffic mixes into 4 priority classes is shown in Table 5.

Table 5: Composition of Traffic Mixes (four priority classes)

Name	% Class 1	% Class 2	% Class 3	% Class 4
mix1	27.7	29.9	18.3	24.1
mix2	13.5	24.3	19.5	42.8

When five cells are searched using the push out algorithm (Figure 6), the fraction of cells lost in each class is spread over a wide range. Class 1 traffic experiences almost no cell loss (when 5 cells are searched, cell loss is 1.6×10^{-5} for mix1 and $< 10^{-7}$ for mix2) and Class 2 traffic has a cell loss lower than the total cell loss. In other words, four priority classes made it possible to protect Class 1 traffic from loss more thoroughly than with only two priority classes without increasing the lowest priority cell loss any more than when two priority classes are used. Figure 6b shows that even when the proportion of higher priority cells is lower, four priority classes still provide better protection for Class 1 cells than two priority classes for a small number of buffer slots searched.

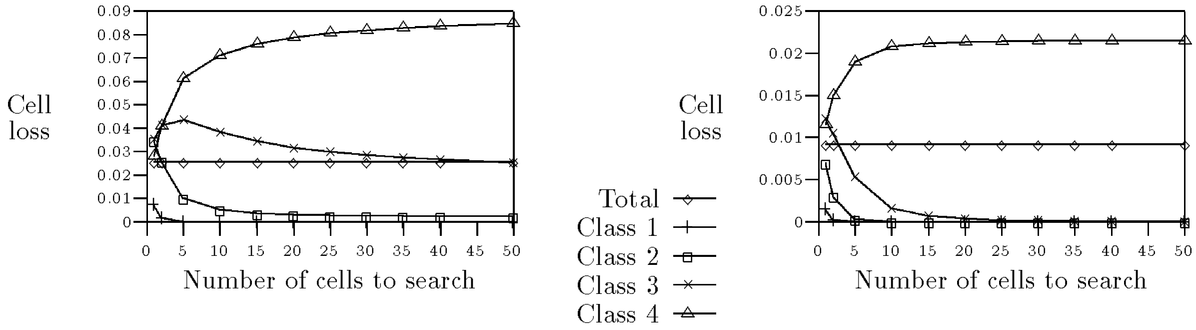

a) mix1 b) mix2

Figure 6: Push Out (four classes, fraction of class)

By contrast, the buffer sharing algorithm gained no benefits from the extra priority classes (Figure 7). Classes 1, 2, and 3 displayed statistically indistinguishable cell loss for both mixes. The cell loss experienced by Class 4 traffic was identical to the cell loss of Class 2 traffic in the 2 priority class experiment.

The selective discarding algorithm results are varied (Figure 8). More parameters are required for four priority classes (Table 6). The best performance for all classes is achieved when the probability parameters for the lower two classes are greater than 0.1.

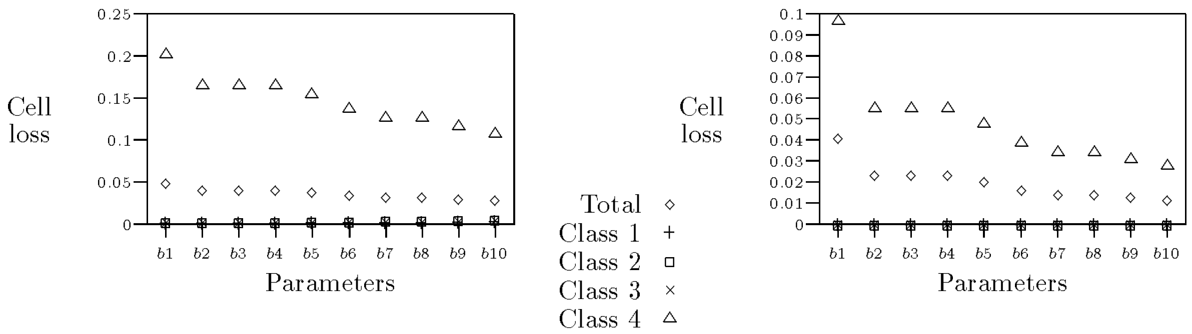

a) mix1 b) mix2

Figure 7: Partial Buffer Sharing (four classes, fraction of class)

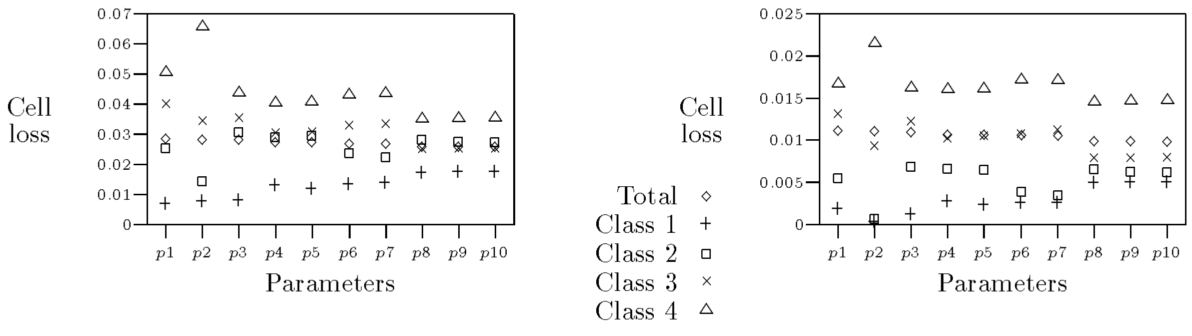

a) mix1 b) mix2

Figure 8: Selective Packet Discarding (four classes, fraction of class)

Table 6: Parameters for Selective Packet Discarding

	p1	p2	p3	p4	p5	p6	p7	p8	p9	p10
P1	0.1	0.001	0.01	0.01	0.001	0.0001	0.00001	0.0001	0.0001	0.00001
P2	0.2	0.01	0.1	0.05	0.05	0.01	0.001	0.005	0.001	0.0001
P3	0.5	0.2	0.2	0.1	0.1	0.1	0.1	0.01	0.01	0.001
P4	0.7	0.5	0.3	0.2	0.2	0.2	0.2	0.1	0.1	0.1

5 Conclusions

The goal of this research was to examine priority queueing algorithms for use at ATM switch buffers as a means of supporting quality of service (QOS) requirements within the network, using realistic traffic models.

The results show that space priority queueing algorithms can be effective in protecting loss sensitive traffic from excessive cell loss by moving the cell loss to a lower priority class. Priority queueing cannot stand alone as a congestion control mechanism, however. It depends on admission control algorithms and traffic shaping to keep the load at a manageable level. If the load on a switch becomes high for an extended period of time, the priority queueing algorithms will be ineffective since the buffer will continually be full.

Although the CCITT has already determined that only two loss priority classes will be allowed, the possibility of extending this to four priority classes should not be ignored. Four priority classes allow a finer control over the level of performance that different service types can expect. To guarantee extremely low cell loss to some services when only two priority classes are allowed, only a small subset of connections can use high priority cells. This forces other services that can afford to lose some cells to be treated the same as services that can tolerate higher cell loss and lower priority services. By placing these services in different priority classes, it may be easier to accommodate the QOS requirements for a wide variety of services.

The overhead of each algorithm should also be taken into consideration. The buffer sharing algorithm is the least expensive since it requires only one or three counters, depending on the number of priority classes. The push out algorithm requires pointers to keep track of lower priority cells in the buffer. The selective packet discarding algorithm requires a counter, a state variable, and a processor to compute the function that determines whether or not to drop a cell. The feasibility of implementing these schemes was beyond the scope of this study.

The choice of a buffer size and the behaviour of the priority queueing algorithms depend to a large extent on the shape of the traffic. This makes it difficult to establish appropriate algorithms for priority queueing until the traffic on the network can be measured and characterized. It also emphasizes the need for realistic traffic models in studies of the network. It is possible that the traffic found on individual ATM networks and on segments of larger ATM networks will vary dramatically according to the applications being run on each segment. To obtain good performance on each network or network segment, the buffer size and choice of queueing algorithm and its parameters should be tunable.

Different algorithms for the selective packet discarding approach should be considered. For example, the performance of higher priority classes would improve if lower priority cells in the buffer could be replaced. The selective packet discarding algorithm showed some success in reducing consecutive cell loss [10] but other methods should be considered.

High demands will be placed on future integrated services networks and the challenge for network designers will be to provide the mechanisms to ensure good performance for all of the applications using the network. This study demonstrates the importance of using realistic traffic models since simpler models such as the Poisson model give overly optimistic figures for cell loss. The results of this research suggest that priority queueing in combination with end-to-end congestion control algorithms and access control algorithms can play a significant role in supporting quality of service requirements and should be considered for use in high speed packet switches.

6 Acknowledgments

The research reported in this paper was supported by the Natural Sciences and Engineering Research Council of Canada, through Research Grant OGP0003707 and the Postgraduate Scholarships program, and by Telecommunications Research Laboratories (TR*Labs*).

References

[1] CCITT, Blue Book, Geneva, Switzerland. *CCITT Recommendation I.121, Broadband Aspects of ISDN*, 1989.

[2] G. Gallassi, G. Cappellini, L. Fratta, G. Rigolio, and F. Rossi. Performance evaluation of flow control mechanisms in multi-throughput class networks. In *Proceedings GLOBECOM 88*, pages 16.3.1–6, 1988.

[3] Gérard Hebuterne and Annie Gravey. A space priority queueing mechanism for multiplexing ATM channels. *Computer Networks and ISDN Systems*, 20(1-5):37–44, December 1990.

[4] Harry Heffes and David M. Lucantoni. A markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance. *IEEE Journal on Selected Areas in Communications*, SAC-4(6):856–868, September 1986.

[5] Jay M. Hyman, Aurel A. Lazar, and Giovanni Pacifici. Real-time scheduling with quality of service constraints. *IEEE Journal on Selected Areas in Communications*, 9(7):1052–1063, September 1991.

[6] Hans Kröner. Comparative study of space priority mechanisms for ATM networks. In *Proceedings IEEE INFOCOM '90*, pages 1136–1143, 1990.

[7] Arthur Y. M. Lin and John A. Silvester. Priority queueing strategies and buffer allocation protocols for traffic control at an ATM integrated broadband switching system. *IEEE Journal on Selected Areas in Communications*, 9(9):1524–1536, December 1991.

[8] Basil Maglaris, Dimitris Anastassiou, Prodip Sen, Gunnar Karlsson, and John D. Robbins. Performance models of statistical multiplexing in packet video communications. *IEEE Transactions on Communications*, 36(7):834–843, July 1988.

[9] David W. Petr and Victor S. Frost. Optimal packet discarding: An ATM-oriented analysis model and initial results. In *Proceedings IEEE INFOCOM '90*, pages 537 – 542, 1990.

[10] K. L. Reid. *Priority Queueing in Asynchronous Transfer Mode Switches*. Master's thesis, Department of Computational Science, University of Saskatchewan, 1992.

[11] Hiroshi Saito. Queueing analysis of cell loss probability control in ATM networks. In *ITC-13*, pages 19–24, 1991.

[12] Willem Verbiest, Luc Pinnoo, and Bart Voeten. The impact of the ATM concept on video coding. *IEEE Journal on Selected Areas in Communications*, 6(9):1623–1632, December 1988.

[13] Gillian Woodruff and Rungroj Kositpaiboon. Multimedia traffic management principles for guaranteed ATM network performance. *IEEE Journal on Selected Areas in Communications*, 8(3):437–446, April 1990.

[14] Nanying Yin, San-Qi Li, and Thomas E. Stern. Congestion control for packet voice by selective packet discarding. *IEEE Transactions on Communications*, 38(5):674–683, May 1990.