A Packet Discard Scheme for the Assured Forwarding in the Differentiated Services Per-Hop-Behavior *

Maurizio Casoni

Dept. of Engineering Sciences
University of Modena and Reggio Emilia
Via G.Campi 213/B - 41100 Modena - I
casoni@dsi.unimo.it

D.E.I.S.
University of Bologna
Viale Risorgimento, 2 - 40136 Bologna - I
mcasoni@deis.unibo.it

Abstract

In this paper two packet discard schemes are proposed with the goal to provide incoming data flows different levels of quality of service in a network node. In particular, the assured forwarding in the Differentiated Services per-hop-behavior, as reported in the Request For Comments 2597, is considered as reference scenario.

The different levels of quality of service imply different portions of bandwidth assigned to data flows which share that same output link. Good performance in terms of throughput and fairness in bandwidth sharing among equal priority flows are also considered as the requirements to meet.

The dropping schemes are shown to provide very good results and to represent a possible solution for dealing with the most demanding data flows during congestion intervals. Numerical results are reported and discussed for homogeneous and non homogeneous data sources.

*This work has been partially presented at IEEE Int. Conference on Computer Communications and Netowrks 1999 and at IEEE Int. Conference on Networks 1999.

1 Introduction

TCP/IP networks do not guarantee any type of quality of service (QoS) since their delivery service, "best effort", implies that each IP data flow can have an indeterminate level of packet loss, delay, out-of-sequence. Since many multimedia applications have quite stringent requirements in terms of delay, loss and miminum bandwidth, a certain number of proposals have come out to provide some QoS to applications running over TCP/IP. According to [1] a possible classification of models for service differentiation is the following: relative priority marking, service marking, label switching [2], integrated services [3]/RSVP [4] and static per-hop classification.

On the other hand, Asynchronous Transfer Mode (ATM) [5], chosen by ITU-T as the basic transfer mode for the Broadband Integrated Services Digital Network (B-ISDN), networks are supposed to carry many different types of traffic with different bandwidth requirements. Applications range from continuos ones transmitting informations at fixed rates to highly bursty which send large blocks of data at high peak rates but with long-term average rates relatively low and the different and changing characteristics of service requirements of current and emerging Internet applications push for a flexible network architecure.

In ATM all applications are supposed to provide a set of traffic descriptors at the virtual circuit setup to describe the traffic sent with sufficient accuracy for routing and call acceptance. While many applications (audio, video, low speed data) are predictable enough for this strategy to work, many recent data applications involve interactive information retrieval with large data files (images) retrieved across the network on demand by a human user. The unpredictability of the human user and the high peak transmission rate required to provide acceptable interactive response imply that one must accept either a low network utilization or the possibility of overload periods during which the traffic sent to some network links exceeds their capacity. The fact that end-to-end protocols, such as TCP/IP, fragment their data in packets containing many ATM cells makes the impact of overload periods very dangerous, since a single lost cell can lead to the loss and retransmission of an entire transport level packet. Available bit rate (ABR) services are controlled by rate-based mechanisms but they work on round trip time intervals and so some other mechanisms must be implemented to cope with congestion occurring in short time intervals [6]. Packet discard schemes [7] [8] are ATM buffer management techniques to ensure high end-to-end throughputs [9] and fairness [10] [11] for bursty data applications during overload periods. Packet discard schemes are implemented in output buffered ATM switches, either output queued or combined input/output queued [12], and perform their task by selectively discarding cells from the incoming information flows.

An important issue is then to study how a router must treat packets belonging to different classes of service. The main approaches so far proposed can be summarized as packet or threshold dropping and priority scheduling. With the former, routers have one FIFO shared buffer among packets of all classes per output link while with the latter routers have separate logical queues for each class and some algorithm drives the scheduler for packet forwarding. Thus, whereas threshold dropping realizes service differentiation by means of thresholds, priority scheduling does it by properly serving packets of separate queues.

For instance, several approaches to congestion control and fairness in ATM switches are based on the implementation of per-VC queueing and several algorithms and designs for the output controller have been proposed [13] [14]. On the other hand, some recent works have underlined possible weaknesses of traditional packet discard techniques, such as random early detection [15], when employed in real networks [16]. Other works, such as [17], confirm the effectiveness of early packet discard to improve the throughput. In conclusion, more efforts must be made to better understand the performance of these techniques and, in particular, to develop ever improved solutions.

In this paper the reference architecture is the one described in [1] ed in particular the focus is on the study of an effective mechanism to provide different levels of drop precedence in Differentiated Services (DS) nodes. The Assured Forwarding (AF) Per-Hop-Behavior (PHB) [18] aims at providing delivery of IP packets in four independently forwared AF classes. Each AF class is in each DS node allocated some portion of the forwarding resources, i.e. buffer space and bandwidth. Within each AF class IP packets are marked with one of three possible drop precedence values. If congestion occurs in a DS node (which may be equipped with an ATM switching architecture), this precedence determines how a packet has to be treated within the AF class. A congested DS node then must try to discard packets in accordance to these precedence values.

In this paper two packet discard schemes are proposed to properly manage in congested DS nodes the flows of packets depending on their drop precedence level within their AF class. One of these schemes can be employed in enterprise networks where two drop precedence levels are enough while the other can be used in DS domains so as to provide three drop levels. Their goal is also to mantain high throughput and fair exploitation of the bandwidth of the output link among equal-priority flows.

The reported results, obtained from the simulation of a single switch output controller, show that these goals are met so that these schemes can be considered possible solutions to implement service differentiation in congested DS nodes.

The paper is organized as follows: section 2 shows the reference system model, section 3 describes the first proposed discarding scheme suitable for enterprise networks; in section 4 the second scheme to be implemented in DS node belonging to large domains is explained. Section 5 reports some numerical results about the two proposed discarding schemes for different scenarios and traffic patterns; finally section 6 closes the paper with a drawing of the achieved outcomes.

2 System Model

The reference architecture is the one described in [1] and the goal of this paper is to study a possible candidate as packet dropping scheme for providing different levels of drop precedence in DS nodes. The assured forwarding PHB aims at providing delivery of IP packets in four independently forwared AF classes. Each AF class is in each DS node allocated some portion of the forwarding resources, i.e. buffer space and bandwidth. Within each AF class IP packets are marked with one of three possible drop precedence values. If congestion occurs in a DS node (which may be equipped with an ATM switching architecture), this precedence determines how a packet has to be treated within the AF class. A congested DS node then must try to discard packets in accordance to these precedence values.

It is worthwhile to underline that packets in a AF class must be forwarded independently from packets of other AF classes, meaning that aggregation of packets belonging to different classes is not allowed. Also, a DS node must be able to deal with three drop precedence levels but in some cases they must yield at least two levels, e.g. in enterprise networks where congestion is rare, whereas three levels are desirable in DS domains.

The implementation of the AF class must avoid long-term congestion within each class, while it is allowed short-term congestion resulting from bursts. This implies the employment of an active queue management algorithm for providing the three drop precedence levels within each AF class. This algorithm can be thought to work in cooperation with a priority scheduling algorithm which has the task to manage the four AF classes.

The three different drop precedence levels within each AF class are identified by different codepoints. The drop precedence levels are defined as Low Drop Precedence (LDP), Medium Drop Precedence (MDP) and High Drop Precedence (HDP), listed in decreasing order of importance, i.e increasing discarding "probability".

In the following two packet discard schemes are proposed to properly manage in congested DS nodes the flows of packets depending on their drop precedence level within their AF class. One of these schemes can be employed in enterprise networks where two drop precedence levels are enough while the other can be used in DS domains so as to provide three drop levels. Their goal is also to mantain high throughput and fair exploitation of the bandwidth of the output link among equal-priority flows.

3 Packet Discard for Enterprise Networks

Some packet discard schemes have already been presented [7] [8] and evaluated [9] [10] [11] but none of them can discriminate among different classes of service. Here we assume to divide incoming flows, associated to VCs, between low and high priority (two classes of service) which can be respectively mapped in HDP and LDP levels of a DS node.

The proposed scheme employs two thresholds, b_l (low) and b_h (high). The buffer controller tries to keep the occupancy level between b_l and b_h by activating or deactivating the VCs at packet boundaries according to their priority level. The goal of this packet discarding scheme is to provide higher priority VCs (i.e. cells, i.e packets) with most of the bandwidth, to optimize the bandwidth exploitation and to reach at the same time a good fairness in bandwidth sharing among equal priority VCs. VC de/activation depends on the number of successfully transmitted cells during a time interval, called the current window, which is compared with the number transmitted by the other VCs. Since the buffer occupancy level has a cyclic behaviour [9], two thresholds aim at improving the performance by increasing the number of packets successfully transmitted during each cycle. A virtual circuit is defined to be active if the buffer controller has decided to accept the cells of the current packet; otherwise, it is inactive. In addition, if any cell of the packet must be discarded due to queue overflow, the remainder of that packet is discarded (tail packet discard).

Let n_i be the number of transmitted cells by VC i in the current window; A_m (I_M) be the smallest (largest) n_i from the set of currently active (inactive) VCs. The scheme here proposed is similar to the one presented in [10] with some variants to take priorities into account. Keeping in mind that during congestion all VCs send packets continually one beside the other, the scheme works as follows:

1. if the current buffer level exceeds b_h, at the packet boundaries of VC J make

J inactive and always discard HDP cells;

2. if the current buffer level is below b_l, at the packet boundaries of VC J make J active;

3. if the current buffer level is between b_l and b_h, at the packet boundaries of VC J the buffer level is falling, J is inactive, $n_j < A_m$ and J is high priority make J active;

4. if the current buffer level is between b_l and b_h, the level is rising, J is active, $n_j > I_M$ and J is low priority make J inactive;

5. in all other cases do not change the state of J.

First two rules try to avoid that buffer overflows and underflows for optimizing the throughput and to save buffer space for high priority data. Third and fourth ones try to obtain fairness for equal priority VCs by turning on and off the VCs that have sent fewer or more cells in the current window, with respect to A_m and I_M. To handle fairness, the proposed scheme must keep the information regarding bandwidth usage of each VC in the current window. Also, it must maintain a record of the evolution of the buffer occupancy level in time in order to be able to decide if it is rising or falling.

Considering a homogeneous situation in which r VCs transmit data continually at a normalized rate of λ (that is, λ is the fraction of the link bandwidth required by a single virtual circuit), $r\lambda$ is defined as the overbooking ratio and for an overloaded link $r\lambda > 1$. It is assumed that all packets of all virtual circuits contain ℓ cells and let $k = \lfloor 1/\lambda \rfloor$ be the maximum number of virtual circuits that the link can handle without loss. Let B denote the number of cells the buffer can contain. The values of B, b_l, b_h largely determine the goodput, defined as the fraction of the link's capacity used to carry complete packets, that it is possible to achieve on a link that has a buffer managed by this packet discard scheme.

If b_l and b_h are properly set to avoid overflow and underflow and the difference $(b_h - b_l)$ is large enough, it is reasonable to foresee a great difference in bandwidth usage between high and low priority data and a quite good degree of fairness among equal priority VCs. The reason is that in this condition the scheme can turn on and off the VCs depending on priorities and bandwidth usage without penalizing throughput.

4 Packet Discard for DS Domains

In addition to the parameters n_i, A_m and I_M previously introduced, here one more is defined: A_{m-LDP}. It is used in a similar way as A_m but it regards packets belonging to the Low Drop Precedence level only; the goal is to manage the fairness of the low drop precedence level data flows separately from the others.

The second scheme adds some variants to the previous in order to manage a three levels of priorities. The scheme works as follows:

1. if the current buffer level exceeds b_h, at the packet boundaries of VC J make J inactive and always discard HDP cells;
2. if the current buffer level is below b_l, at the packet boundaries of VC J make J active;
3. if the current buffer level is between b_l and b_h, the level is rising, J is active, $n_j > I_M$ and J is either HDP or MDP make J inactive;
4. if the current buffer level is between b_l and b_h, at the packet boundaries of VC J the buffer level is falling, J is inactive, $n_j < A_m$ and J is MDP make J active;
5. if the current buffer level is between b_l and b_h, at the packet boundaries of VC J, J is inactive, $n_j < A_{m-LDP}$ and J is LDP make J active;
6. in all other cases do not change the state of J.

The main characteristics here are that LDP data flows are never made inactive when the buffer level is below the threshold b_h and that the fairness is managed separetely from MDP and LDP flows.

5 Numerical results

A simulation program of the algorithms previously described has been developed to study the performance in terms of per VC bandwidth usage, output link goodput and fairness in bandwidth usage. The bandwidth usage is the figure of merit for the level of quality assigned to each VC in presence of congestion and it must respect the drop precedence level of the related user data flow. The discarding schemes will then be mainly tested on this parameter. Besides, it is worth reminding that the proposed schemes employ two thresholds which aim at increasing the number of succesfully transmitted packet per cycle, thus improving the goodput.

The simulation requires to set up several parameters: the simulation duration is 1000 packet times long, the current window is 200 packet times. All numerical results assume $\ell = 171$ to represents one ATM adaptation layer (AAL) 5 packet.

Assuming a 600 Mbit/s output link, the bandwidth usage is considered in the overbooking ratio range up to 3 Gbit/s which is a likely scenario, which means $r\lambda$ in the range [1..5]. The scheme has been evaluated in different operating scenarios and traffic patterns, namely, with homogeneous sources sending data at the same rate, or normalized rate λ, and with non homogeneous sources which have been grouped in two different sets characterized by different data rates, λ_1 and λ_2 respectively.

Next paragraphs will report some numerical results for these two main scenarios in enterprise networks and DS domains.

5.1 Packet Discard for Enterprise Networks

In enterprise networks, where congestion is supposed to be rare and very short, two drop precedence levels may be enough so that the reported results refer to the scheme described in section 3.

overbooking ratio	2	3	4	5
Case 1	0.9772	0.9825	0.979	0.9786
Case 2	0.9977	0.9993	0.9995	0.998

Table 1: Goodput as a function of the overbooking for the two analyzed cases.

5.1.1 Homogenous sources

Two situations are here analyzed. Both refer to an homogeneous situation meaning that all VCs have the same rate. Case 1 has $k = 12$, $B = 512$, $b_h = 256$, $b_l = 64$ and 1/3 of the VCs are LDP. Case 2 corresponds to $k = 10$, $B = 812$, $b_h = 512$, $b_l = 128$ and 1/4 of the VCs are LDP. In all cases the LDP level VCs have been randomly chosen.

Figure 1 refer to case 1 and show the bandwidth usage for each VC for an overbooking ratio equal to 2. In Figure 1 LDP (high priority) VCs belong to the number set $(3, 6, 9, 11, 15, 18, 20, 23)$: most of the bandwidth is assigned to higher priority VCs and within the two classes of services a significant level of fairness is obtained.

Figure 2 refer to case 2 and show the bandwidth usage for each VC for an overbooking ratio equal to 3. In Figure 2 LDP VCs belong to the number set $(2, 5, 8, 11, 15, 18, 22, 26)$ and the bandwidth is again assigned mostly to high priority VCs with a reasonable fairness.

Table 1 shows the goodput values reported for the two considered cases and underlines the effectiveness of the proposed scheme in optimizing the output link.

The results obtained mainly match the proposed target of assigning the output link to LDP VCs first, keeping the goodput as high as possible and trying to provide an equal sharing of the link among equal priority VCs.

5.1.2 Non homogeneous sources

In this section the packet discard scheme is studied in a different scenario. Let us assume that VCs have two different rates, λ_1 and λ_2 with $\lambda_1 > \lambda_2$. It is also assumed that VCs still transmit packets of the same length, $\ell = 171$. As a consequence, the packet transmission time is different for VCs with different rates and $\ell/\lambda_1 < \ell/\lambda_2$. Therefore, during the transmission time $T_2 = \ell/\lambda_2$ more than one packet boundaries per each VC can occur for faster VCs. This means that faster VCs ask the buffer controller for more accesses to the buffer than slower ones and this may lead to some unfairness between fast and slow VCs.

In the following cases, high priority is assigned to faster VCs, i.e. they have a LDP level, assuming that they are used for time critical applications or because users simply pay for that. Of course, other assumptions and mixes can be made but they will not be discussed here.

In figure 3 $\lambda_1 = 0.2$, $\lambda_2 = 0.1$, $r_1 = 5$ and $r_2 = 30$ and the overbooking ratio is then $r_1\lambda_1 + r_2\lambda_2 = 4$; the high priority VCs $(2, 13, 23, 28, 33)$ are $1/7$ of the total. Again, as desired, they get most of the bandwidth and HDP (low priority) VCs can take small portions of it.

Basically, all the bandwidth is taken by the r_1 sources and a quite good level of fairness is achieved among them. On the other hand, the r_2 sources get almost nothing but, actually, this is just what we want: priorities have been set up to assign resources, i.e. bandwidth, to specific users, in particular during the short congestion intervals.

5.2 Packet Discard for DS Domains

In large DS domains it is likely that congestion occurs more often in DS nodes than in enterprise networks. Thus, a more accurate packet discard scheme must be employed to take into account more than two drop precedence levels.

The following results have been obtained by using the discarding scheme detailed in section 4, which can treat in a different way data flows belonging to three different priority levels, namely LDP, MDP and HDP levels.

5.2.1 Homogenous sources

Two situations are here analyzed as well. Both refer to an homogeneous situation, again meaning that all sources have the same rate. Case 1 has $k = 12$, $B = 512$, $b_h = 256$, $b_l = 64$. Case 2 corresponds to $k = 10$, $B = 812$, $b_h = 512$, $b_l = 128$.

Figure 4 refer to case 1 and show the bandwidth usage for each VC for an overbooking ratio equal to 2. In this figure LDP (high priority) VCs belong to the number set $(2, 5, 10, 13, 17, 22)$ and have been randomly chosen. Also, the set $(3, 6, 9, 11, 15, 18, 20, 23)$ corresponds to MDP VCs.

The main result is that the proposed scheme assignes the output link bandwidth respecting the drop precedence levels: almost 50% of the bandwidth is given to LDP flows and fairly shared among them, roughly 8.1% each; the eight MDP flows get approximately 4.2% each, 33.6% globally; finally, also HDP flows get some bandwidth, roughly 1.6% each, and 16% is the amount for the whole HDP level.

Figure 5 refer to case 2 and show the bandwidth usage for each VC for an overbooking ratio equal to 5. In this figure LDP (high priority) VCs belong to the number set $(3, 7, 11, 14, 17, 20, 24, 28, 30, 33, 38, 42)$ and again have been randomly chosen. Also, the set $(1, 4, 8, 10, 13, 16, 23, 26, 31, 35, 37, 40, 44, 46, 48, 50)$ corresponds to MDP VCs. Also in this case, the proposed scheme respects the drop precedence levels: 52.8% of the bandwidth is given to the twelve LDP flows and fairly shared among them, roughly 4.4% each; the sixteen MDP flows get approximately 2.9% each, 46.4% globally also here HDP flows get some bandwidth, roughly 0.04% each, with 0.8% the amount for the whole HDP level. Again, it is important to underline the remarkable level of fairness obtained within each level.

5.2.2 Non homogeneous sources

As in subsection 5.1.2, VCs may have two different rates, λ_1 and λ_2 with $\lambda_1 > \lambda_2$, and transmit packets of the same length, $\ell = 171$.

Figure 6 refers to the case with $k_1 = 5$, $k_2 = 10$, $r_1\lambda_1 + r_2\lambda_2 = 4$, $B = 512$, $b_h = 256$, $b_l = 64$, $(3, 6, 9, 13, 17, 20, 26, 30, 32)$ are the LDP VCs, randomly chosen again,

VC	1	2	3	4	5	6	7	8	9	10	11	12
Gdput	0.0012	0.0345	0.0602	0.0012	0.0346	0.0582	0.0006	0.0032	0.0593	0.0364	0.0008	0.0369
VC	13	14	15	16	17	18	19	20	21	22	23	24
Gdput	0.0398	0.0006	0.0382	0.0007	0.0573	0.0364	0.0016	0.0578	0.0351	0.0050	0.0023	0.0356
VC	25	26	27	28	29	30	31	32	33	34	35	Total
Gdput	0.0010	0.0614	0.0350	0.0293	0.0010	0.0579	0.0365	0.0584	0.0022	0.0384	0.0037	0.9623

Table 2: Goodput per VC.

and it shows the bandwidth usage for each VC. The $r_1 = 5$ fast sources are in the set $(2, 13, 23, 28, 33)$: one of them has LDP level (number 13), two have a MDP level (number 2 and 28) and two a HDP level (23 and 33). It is worthwhile noting that the bandwidth has been assigned according to the drop precedence levels also in this case meaning that the proposed scheme can effectively support the presence of different rates sources. Actually, with respect to the results with homogenous sources, here there is more fluctation in the amount of assigned bandwidth within each level but, still, the difference among them is clear. As a matter of fact, considering the LDP sources, five get 6.3% the bandwidth, three the 5.8% and one the 3.9%, and the sum is roughly 52.8%; as regards the twelve MDP sources, they get on average the 3.8%, globally the 45.6%; finally the HDP sources share the remaining 1.6%, which is roughly 0.1% each.

Thus, basically the output link bandwidth is assigned to LDP and MDP sources, even if a small portion is given also to HDP level ones, and between them most of it is taken by LDP sources that is exactly what we want. In addition, a quite good level of fairness is achieved among flows belonging to the same drop precedence level.

In table 2 the total goodput, equal to 96%, is shown with contributions from all VCs. Again these values confirm the effectiveness of the proposed scheme in optimizing the output link as well.

6 Conclusions

In this paper a pair of selective packet discard schemes has been proposed for managing congestion in output buffered ATM switches taking into account the class of service, or drop precedence level, of the incoming data flows. They have been employed in network nodes in two different scenarios, namely enterprise networks and differentiated service domains. Their goal is to provide output bandwidth to data streams in accordance to their priority and two or three priority levels have been considered.

The analyzed schemes try to keep the buffer occupancy level between two thresholds as long as possible by de/activating the VCs depending on the number of successfully transmitted packets in a time window and on the drop precedence level, in order to provide high goodput and good fairness in bandwidth sharing among equal priority VCs. In other words, the goal is not to avoid congestion, in fact congestion already occurs being $r\lambda > 1$, but at least to avoid overflow, i.e cell losses, and underflow, i.e. output link inactivities, providing a reasonable fairness among equal priority VCs.

For the operating scenarios, homogeneous and non homogeneous, here taken into account the proposed schemes have been shown to meet the mentioned target and, in particular, to be capable of managing successfully different priority data flows without penalizing throughput or fairness.

The above results have been obtained through simulation so that it is not straightforward to foresee the exact behaviour of the scheme in a real network. As a matter of fact, realistic scenarios can be more complex than the ones here presented and analyzed. Combinations of different packet lengths with many different data rates and priorities are possible and this is current work.

The subject of packet discarding schemes is still open to discussion about their effectiveness, as recent literature has underlined, so that many more investigations and possibly test-beds are necessary to better clarify this issue. However, these first results show that is possible to set up a packet discard scheme whose task is to provide class-dependent bandwidth assignment and high throughput with an acceptable fairness.

References

[1] S. Blake, D. Blake, M. Carlson, E. Davies, Z. Wang, W.Weiss, "An Architecture for Differentiated Services", *Request for Comments 2475*, December 1998.

[2] R. Callon, G. Swallow, N. Feldman, A. Viswanathan, P. Doolan, A. Fredette, "A Framework for Multiprotocol Label Switching", Internet Draft, draft-ietf-mpls-framework-02.txt, August 1997.

[3] B. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an overview", *Request for Comments 1633*, June 1994.

[4] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) - version 1 functional specification", *Request for Comments 2205*, September 1997.

[5] M. De Prycker, "Asynchronous Transfer Mode: a Solution for B-ISDN", *Ellis-Horwood*, II edition, 1993.

[6] R. Jain, "Congestion Control and Traffic Management in ATM Networks: Recent Advances and a Survey", *Computer Networks and ISDN Systems*, vol. 28, no. 13, pp. 1723-1738, October 1996.

[7] S. Floyd, R. Allyn, "Dynamics of TCP Traffic over ATM Networks", *Computer Communication Review*, vol. 24, no. 4, 10/94.

[8] J.S. Turner, "Maintaining high throughput during overload in ATM switches", Proc. of Infocom '96, 24-28 March 1996, San Francisco, CA (USA).

[9] M. Casoni, J.S. Turner, "On the Performance of Early Packet Discard", *IEEE Journal on Selected Areas in Communications*, special issue on *Advances in ATM Switching Systems for B-ISDN*, vol. 15, no. 5, pp. 892-902, June 1997.

[10] M. Casoni, "Fair Packet Discard as Improvement of Early Packet Discard", Proc. of Globecom'98, 8-12 November 1998, Sydney (AUS).

[11] S. Chan, E.W.M. Wong, K.T. Ko, "Fair Packet Discarding for Controlling ABR Traffic in ATM Networks", *IEEE Transactions on Communications*, vol. 45, no. 8, pp. 913-916, August 1997.

[12] S. Chuang, A. Goel, N. McKeown, B Prabhakar, "Matching Output Queueing with a Combined Input/Output-Queued Switch", *IEEE Journal on Selected Areas in Communications*, vol.17, no.6, pp.1030-1039, June 1999.

[13] L. Benmohamed, Y.T. Wang, "A Control-Theoretic ABR Explicit Rate Algorithm for ATM Switches with Per-VC Queueing", Proc. of Infocom'98, 29 March - 2 April 1998, San Francisco, CA (USA).

[14] Y. Chen, J.S. Turner, "Design of a Weighted Fair Queueing Cell Scheduler for ATM Networks", Proc. of Globecom'98, 8-12 November 1998, Sydney (AUS).

[15] S. Floyd, V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance", *IEEE/ACM Transactions on Networking*, pp.397-413, 1993.

[16] M. May, J. Bolot, C. Diot, B. Lyles, "Reasons not to deploy RED", Proc. of 7th International Workshop on Quality of Service, 31 May-4 June 1999, London (U.K.).

[17] R. Goyal, R. Jain, M. Goyal, S. Fahmy, B. Vandalore, "Traffic Management for TCP/IP over Satellite ATM Networks", *IEEE Communications Magazine*, pp.56-61, March 1999.

[18] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, "Assured Forwarding PHB Group", *Request for Comments 2597*, June 1999.

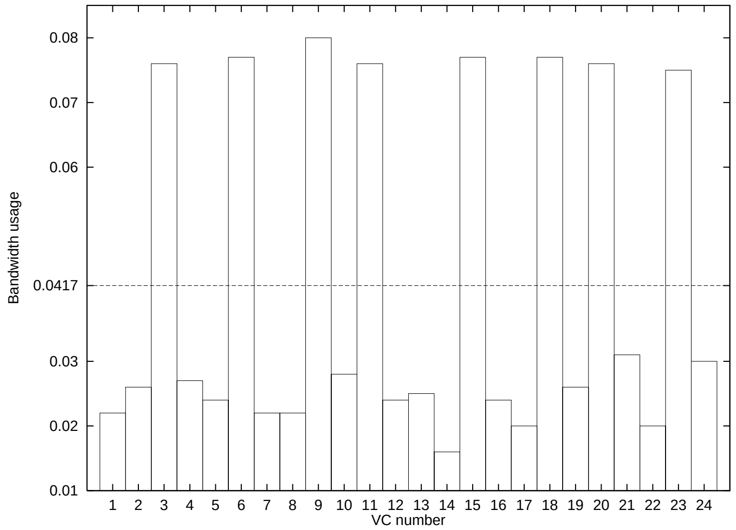

Figure 1: Bandwidth usage for each virtual circuit; $k = 12$, $r\lambda = 2$, $B = 512$, $b_h = 256$, $b_l = 64$; dotted line represents the ideal sharing corresponding to 1/24 of the output link.

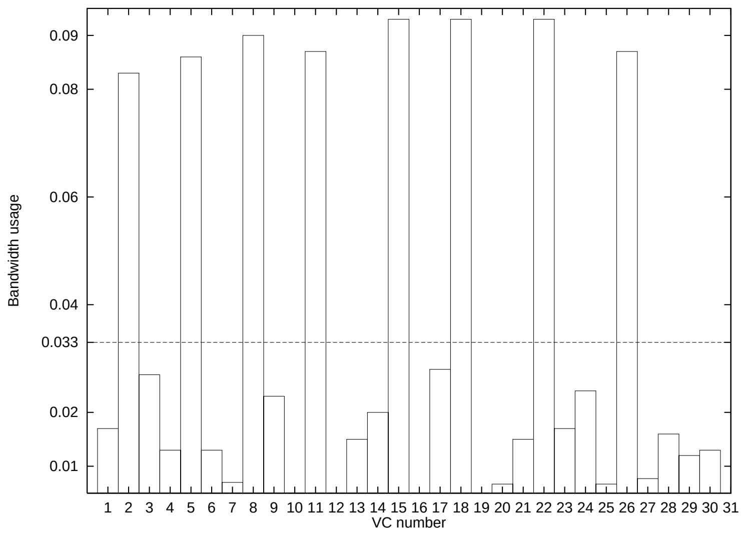

Figure 2: Bandwidth usage for each virtual circuit; $k = 10$, $r\lambda = 3$, $B = 812$, $b_h = 512$, $b_l = 128$; dotted line represents the ideal sharing corresponding to 1/30 of the output link.

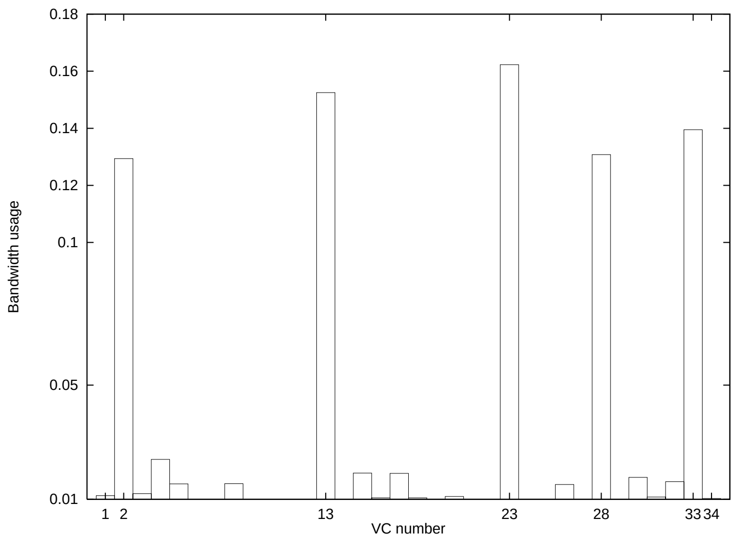

Figure 3: Bandwidth usage for each virtual circuit; $k_1 = 5$, $k_2 = 10$, $r_1\lambda_1 + r_2\lambda_2 = 4$, $B = 512$, $b_h = 256$, $b_l = 64$; (2,13,23,28,33) are the high priority VCs.

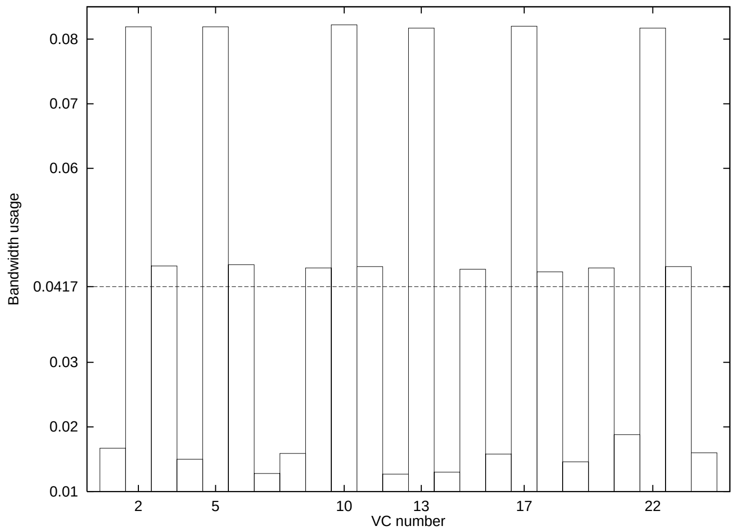

Figure 4: Bandwidth usage for each virtual circuit; $k = 12$, $r\lambda = 2$, $B = 512$, $b_h = 256$, $b_l = 64$; (2,5,10,13,17,22) are the low drop precedence VCs.

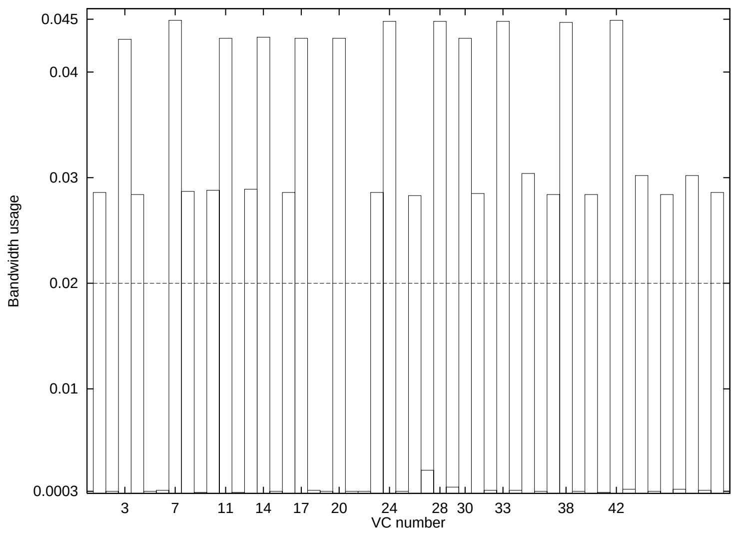

Figure 5: Bandwidth usage for each virtual circuit; $k = 10$, $r\lambda = 5$, $B = 812$, $b_h = 512$, $b_l = 128$; (3,7,11,14,17,20,24,28,30,33,38,42) are the low drop precedence VCs.

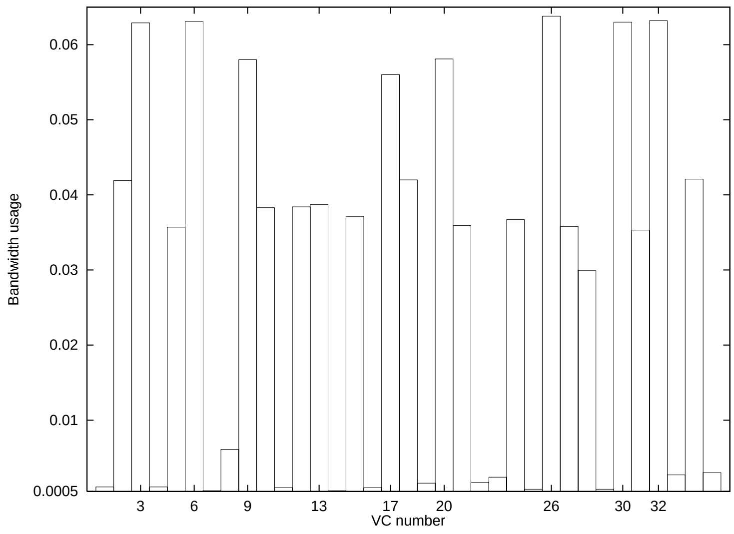

Figure 6: Bandwidth usage for each virtual circuit; $k_1 = 5$, $k_2 = 10$, $r_1\lambda_1 + r_2\lambda_2 = 4$, $B = 512$, $b_h = 256$, $b_l = 64$; (3,6,9,13,17,20,26,30,32) are the low drop precedence VCs.