

Fair Packet Discarding for Controlling ABR Traffic in ATM Networks

Sammy Chan, *Member, IEEE*, Eric W. M. Wong, *Member, IEEE*, and K. T. Ko, *Member, IEEE*

Abstract—The asynchronous transfer mode (ATM) Forum has chosen rate-based control as the flow control scheme for available bit-rate (ABR) service. However, rate-based schemes can achieve congestion control only if all users act in a cooperative manner. Even a limited number of uncooperative users can cause congestion collapse. In this paper, we propose a mechanism called *fair packet discarding* to provide incentives to users to participate in network congestion control so that the network can operate in a more efficient manner.

Index Terms—ABR traffic, congestion control, packet discarding.

I. INTRODUCTION

ONE OF THE current challenges in asynchronous transfer mode (ATM) networks is the transparent support of the conventional connectionless LAN traffic in ATM's connection oriented environment. Such traffic is best served by the best effort method. That is, they should be allowed to send their traffic whenever they need to, without reserving the bandwidth beforehand, but the network does not provide any strict quality of service (QoS) guarantees. Also, they should be allowed to use as much bandwidth as available, perhaps including the idle bandwidth reserved for other services. However, they will be flow controlled during network congestion. Such a service class has been defined by the ATM Forum [1] as the available bit rate (ABR) service.

Various flow control schemes have been proposed for ABR traffic [2], and they can be classified into two main classes, namely, end-to-end rate-based and link-by-link credit-based schemes. In late 1994, the ATM Forum selected the rate-based control as the flow control scheme for ABR service due to its simplicity. However, as discussed in [3], rate-based schemes can achieve congestion control only if all users act in a cooperative manner. Even a limited number of uncooperative users can cause congestion collapse. Therefore, an additional mechanism is necessary to ensure that all users will participate in congestion control. Fairness discarding [4] is such a mechanism. In each switch, this mechanism ensures that each user gets no more than a fair share of network resources during overload by discarding all excessive cells. The essence of this mechanism is to give incentives to users to cooperate so that the network can operate in a more efficient manner. That is, it discards cells from misbehaving users such that they are

the only ones experiencing congestion. It is then up to each user to perform flow control in order to recover from cell loss and be restored to its normal network usage. Therefore, this mechanism complements the role of rate-based flow control.

Although the efficacy of fairness discarding in congestion control has been demonstrated, discarding excessive cells indiscriminately causes an unnecessary waste of bandwidth [5], [6]. This is because packets from the upper protocol layer are segmented into cells. While a cell is discarded, other cells from the same packet may be transmitted successfully, but will eventually be discarded in the destination as the packet is incomplete. Some capacity is thus wasted to transmit these useless cells.

Rather than maintaining fairness, another discarding mechanism called *packet discard strategy* or *early packet drop* [7], [8] (EPD) has been proposed to reduce capacity waste due to transmitting incomplete packets. In this mechanism, when the buffer occupancy exceeds a predetermined threshold, the switch starts looking for the first cell to arrive belonging to a new packet. It discards the first arriving cells of the new packets, and all of their subsequent cells. Whole packets of cells will continuously be discarded until the buffer occupancy drops below the threshold. This mechanism has good throughput performance, but does not provide protection of the network from misbehaving users. Even though congestion is caused by misbehaving users, new packets of all users are discarded whenever the threshold is reached.

To overcome the above drawbacks, we propose a new discarding mechanism called *fair packet discarding* (FPD). Unlike the fairness discarding scheme, FPD also aims to minimize bandwidth wastage while maintaining fairness. The discarding mechanism is activated only when congestion is detected, and discarding is carried out for complete packets. And unlike EPD, FPD confines packet discarding to sources which have received more than a fair share of bandwidth. Note that, although FPD controls the admission of cells into the buffer, it is not an algorithm for buffer allocation such as that in [9]. Rather, it effectively allocates bandwidth to different connections in a fair manner through the discarding process. The remainder of this paper is organized as follows. Section II describes the details of the FPD algorithm. Section III describes a fairness criterion used to identify sources whose packets are to be discarded. In Section IV, FPD is tested by simulations with TCP used as the upper layer protocol.

II. FAIR PACKET DISCARDING (FPD)

FPD assumes that ATM switches are based on the output-buffered architecture. It operates independently in each output buffer. All ABR cells intending to reach a particular output link are queued in the corresponding buffer. The switch always

Paper approved by P. E. Rynes, the Editor for Switching Systems of the IEEE Communications Society. Manuscript received January 9, 1996; revised December 12, 1996. This work was supported by the Hong Kong Research Grants Council under Grant 9040157. This paper was presented in part at the Australian Telecommunication Networks and Applications Conference, Melbourne, Australia, December 1996.

The authors are with the Department of Electronic Engineering, City University of Hong Kong, Kowloon, Hong Kong.

Publisher Item Identifier S 0090-6778(97)05745-0.

maintains a record of bandwidth usage for each virtual circuit (VC) carrying ABR traffic through its output ports. From these records, it can be determined which VC has already used more than its fair share of bandwidth (the definition of fairness is given in the following section), and these VC's are called controlled VC's. When the buffer occupancy exceeds a preselected threshold, the switch then looks for the first cell of a new packet. When such a cell is detected, and if this cell belongs to a controlled VC, the switch will discard it and all subsequent cells of the same packet. On the other hand, if the first cell does not belong to a controlled VC, all cells of the new packet will be admitted into the buffer unless the buffer is full. On average, each user thus receives a fair share of available bandwidth. Hence, misbehaving users could not benefit from their aggressiveness.

Implementing FPD is straightforward with AAL5. Since AAL5 does not support the simultaneous multiplexing of packets on a single VC, once the switch discards a cell from a VC, it can continuously discard all the subsequent cells belonging to the same VC until the end-of-packet cell is seen.

III. DETERMINING CONTROLLED VC'S

A simple fairness criterion is equal sharing. That is, whenever the total demand exceeds the available bandwidth, the bandwidth is divided equally among the competing VC's. However, equal sharing is not necessarily the proper way to divide the overdemanded resources. In some cases, demand from VC's may vary in a wide range, and equal sharing may underutilize resources and hence over-suppress resource usage of heavy users. Here, we will use the fairness criterion proposed in [3] which takes into account the demand of each VC.

Let K represent the number of VC's competing for the total available capacity C in the interval T . Also, let r_i be the offered traffic of the i th VC, which is the number of cells that have been offered into the buffer in the interval T , where $i = 1, 2, \dots, K$.

If there is a burst level congestion, the total offered traffic will be greater than the capacity, that is, $\sum_{i=1}^K r_i > C$.

Let us define the concept of *excess load*, denoted E , as the difference between the available capacity and the total offered traffic. That is,

$$E = \sum_{i=1}^K r_i - C. \quad (1)$$

Clearly, when $E > 0$, the apportioned bandwidth for some (maybe all) of the VC's will have to be less than their offered traffic. These VC's are called *controlled*. The other VC's will be called *uncontrolled*. In other words, the allocated bandwidth of a controlled VC is less than its offered traffic, while the allocated bandwidth of an uncontrolled VC equals its offered traffic. The chosen fairness criterion as stated in [3] is as follows.

No VC will enjoy higher bandwidth usage than a controlled VC.

This implies that all controlled VC's should have equal bandwidth apportionment, and an uncontrolled VC's appor-

tionment cannot be more than that of a controlled VC. One of the advantages of this criterion over equal sharing is that the available capacity is fully utilized by the competing VC's.

This criterion uniquely defines a set of controlled VC's for a given set of r_i and C . A simple method for identifying the controlled VC's is as follows. First, the K VC's are ordered according to their offered traffic, namely, the r_i values, such that $r_i \geq r_h$ if $i < h$. If $E \leq 0$, there are no controlled VC's. If $E > 0$, there exists μ (to be determined later) such that VC $1, 2, \dots, \mu$ are controlled. The total capacity available for these controlled VC's, C_c , is given by

$$C_c = C - \sum_{h=\mu+1}^K r_h = \sum_{h=1}^{\mu} r_h - E. \quad (2)$$

According to the criterion, all controlled VC's should have an equal share of bandwidth, which is given by

$$\gamma_i = \frac{C_c}{\mu}, \quad i = 1, 2, \dots, \mu \quad (3)$$

where γ_i is the bandwidth apportionment for the i th VC.

As the $(\mu + 1)$ th VC is not controlled, we have that $\gamma_{\mu+1} = r_{\mu+1}$, and by the fairness criterion, its share of bandwidth cannot be greater than that of a controlled VC, hence,

$$r_{\mu+1} \leq \frac{C_c}{\mu}. \quad (4)$$

Accordingly, by (2) and (4), we can easily find the set of controlled VC's by obtaining μ as the smallest value for ω such that

$$r_{\omega+1} \leq \frac{\sum_{h=1}^{\omega} r_h - E}{\omega}. \quad (5)$$

IV. SIMULATION SETUP AND RESULTS

Simulation models in this study are built based on the ATM module of OPNET, which is a commercial simulation package for communications networks [10]. In order to reduce simulation time, we have chosen a small network which consists of four source nodes sending packets toward the same destination node via an intermediate switch (Fig. 1). There are two types of sources—cooperative and uncooperative. Cooperative sources use the transmission control protocol (TCP) window flow control in the packet level and rate-based flow control in the ATM level, respectively. TCP uses adaptive window flow control, and retransmits lost packets. Uncooperative sources do not use any flow control schemes. They continuously send traffic into the network regardless of the congestion status of the network, and they do not retransmit lost packets. In the source and destination nodes, AAL 5 is used in the adaptation layer. Packets from the transport layer are first encapsulated as CS-PDU's (protocol data units). Then the SAR sublayer breaks CS-PDU's into 48 byte SAR-PDU's which are sent to the destination in ATM cells. The capacity of each link of the network is 155 Mbits/s. The propagation delay between adjacent nodes is 3 μ s. This is to model a LAN environment. The simulated time of all simulations is 5 s.

In this study, all sources are assumed to have an infinite supply of data. Also, each TCP connection has a window

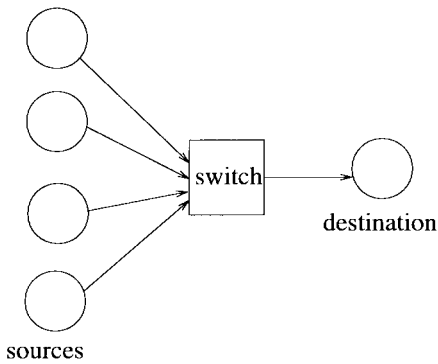


Fig. 1. Network model.

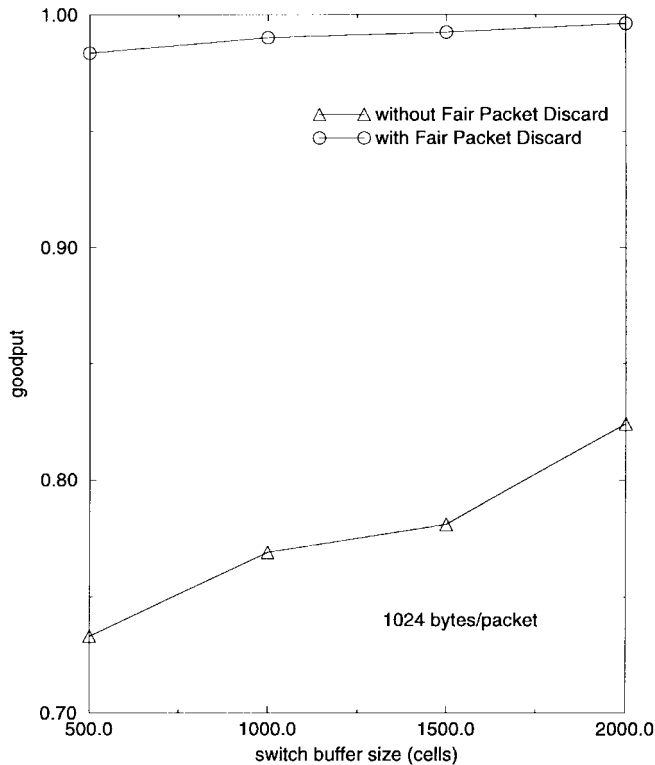


Fig. 2. Goodput against switch buffer size (packet size = 1024 bytes).

size of 64 kbytes. The time window for recording bandwidth usage of VC's is 50 ms. All simulations were run for a range of packet sizes and switch buffer sizes. The throughput performance of FPD is investigated under two scenarios.

Let us consider the first scenario, where all users are cooperative. First, a packet size of 1024 bytes is used. Fig. 2 compares the normalized goodput of the network with and without FPD for various switch buffer sizes. Goodput is defined here as the throughput that does not include cells that are part of a retransmission or an incomplete packet. It can be seen that with FPD, the goodput is improved dramatically and brought very close to 1. Similar performance is observed when the packet size is increased to 4352 and 9180 bytes. Therefore, it demonstrates that FPD performs as well as EPD with cooperative users.

In the second scenario, one user is uncooperative while the other three users are cooperative. Table I compares the

TABLE I
AGGREGATED GOODPUT OF COOPERATIVE SOURCES

packet size (bytes)	1024		4352		9180	
	FPD	no FPD	FPD	no FPD	FPD	no FPD
1000	0.7539	0.5345	0.7721	0.5230	0.7785	0.5213
2000	0.7787	0.5665	0.7756	0.5650	0.7392	0.5641
3000	0.7233	0.5874	0.7903	0.5811	0.7131	0.5803

TABLE II
GOODPUT OF THE UNCOOPERATIVE SOURCE

packet size (bytes)	1024		4352		9180	
	FPD	no FPD	FPD	no FPD	FPD	no FPD
1000	0.2330	0.2828	0.2127	0.2332	0.2088	0.2142
2000	0.2103	0.3693	0.2108	0.3051	0.2498	0.2376
3000	0.2658	0.3812	0.1964	0.3235	0.2747	0.2635

TABLE III
TOTAL GOODPUT OF ALL SOURCES

packet size (bytes)	1024		4352		9180	
	FPD	no FPD	FPD	no FPD	FPD	no FPD
1000	0.9869	0.8173	0.9848	0.7652	0.9873	0.7355
2000	0.9890	0.9358	0.9864	0.8701	0.9890	0.8017
3000	0.9891	0.9686	0.9867	0.9046	0.9878	0.8438

aggregated goodput of the cooperative users for the network with and without FPD. The results show that without FPD, the aggregated goodput of the cooperative users is adversely affected by the aggressive user. It is significantly less than the ideal value of 75%. However, with FPD, the aggregated goodput of the cooperative users is close to the ideal value. Table II shows the impact of FPD on the uncooperative user. When there is no FPD and the packet size is small, the uncooperative user can achieve more than its fair share of goodput. On the other hand, when the packet size is large, it seems that the uncooperative user also suffers from congestion caused by himself. In both cases, FPD is able to maintain its goodput close to its fair share of bandwidth—25%. So, the misbehaving user would experience packet loss when submitting excessive traffic. Note that the goodput would be even lower if the go back N retransmission scheme is used to recover lost packets. Therefore, FPD discourages users to be uncooperative. Finally, Table III shows that the total goodput of all users is significantly increased by FPD. These results demonstrate that FPD can prevent misbehaving users from monopolizing the capacity and causing throughput degradation.

V. CONCLUSION

In this paper, we have proposed a mechanism called *fair packet discarding*. This mechanism bears the good characteristics of two other congestion control schemes—fairness discarding and packet discarding. That is, it maintains fair bandwidth usage while, at the same time, achieving high network throughput. By simulations, we have shown that users who employ end-to-end flow control can almost obtain their fair shares of available bandwidth, irrespective of whether there are misbehaving users in the network. FPD attempts to penalize only the misbehaving users who cause congestion. It is then up to each user to perform flow control in order to recover from cell loss and be restored to its normal network usage.

REFERENCES

- [1] ATM Forum, "Baseline text for traffic management group," *Specification of the ABR Service*, 1994.
- [2] R. Jain, "Congestion control and traffic management in ATM networks: Recent advances and a survey," *Comput. Networks and ISDN Syst.*, to be published.
- [3] M. Zukerman and S. Chan, "Congestion control by maintaining fairness in high speed networks," in *Proc. IEEE GLOBECOM'94*, Dec. 1994.
- [4] C. O'Neill, "Fairness discarding for congestion control in ATM networks," in *Proc. Australian Broadband Switching and Services Symp.'92*, vol. 1, Melbourne, Australia, July 1992, pp. 185–192.
- [5] S. Chan, E. Wong, and K. T. Ko, "Throughput performance at upper layer in ATM networks with fairness discarding," in *Proc. Regional Int. Teletraffic Seminar*, South Africa, Sept. 1995, pp. 46–51.
- [6] C. Fang, H. Chen, and J. Hutchins, "Simulation analysis of TCP performance in congested ATM LAN using DEC's flow control scheme and two selective cell-drop schemes," *ATM Forum/94-0119*.
- [7] J. Hullett, T. Wong, G. Mercankosk, and A. Budrikis, "Packet discard strategy for congestion control of ATM data networks," in *Proc. Australian Telecommun. Networks Appl. Conf.*, Melbourne, Australia, Dec. 1994, pp. 633–638.
- [8] A. Romanow and S. Floyd, "Dynamics of TCP traffic over ATM networks," in *Proc., SIGCOMM'94*, London, England, Sept. 1994, pp. 79–88.
- [9] G. L. Wu and J. W. Mark, "A buffer allocation scheme for ATM networks: Complete sharing based on virtual partition," *IEEE/ACM Trans. Networking*, vol. 3, pp. 660–670, Dec. 1995.
- [10] *OPNET 2.5.B User Manual*, Vol. 8.1.0, MIL3 Inc.