

A Bandwidth Control Scheme for Connectionless ATM Traffic with Multiple Traffic Classes

Jörg Liebeherr[†]Ian F. Akyildiz[‡]Debapriya Sarkar^{*}[†] Computer Science Department, University of Virginia, Charlottesville, VA 22903.[‡] School of ECE, Georgia Institute of Technology, Atlanta, GA 30332.^{*} Hughes Network Systems, Hughes Network Systems, Germantown, MD 20876.

Abstract

A bandwidth control mechanism is proposed for ATM networks that can control the usage of bandwidth in the presence of both connection-oriented and connectionless traffic, as well as multiple classes of connectionless traffic. The bandwidth control mechanism operates at three levels. At the topmost level, bandwidth is dynamically regulated between connection-oriented and connectionless traffic based on the utilization of each traffic type. At the next level, bandwidth is controlled between different classes of connectionless traffic, such as real-time traffic, bulk data traffic, and so on. At the lowest level, bandwidth is distributed among flows belonging to the same connectionless traffic class.

1 Introduction

Several different approaches have been proposed for supporting connectionless traffic in ATM networks [2]. In this study, we consider the so-called direct approach [2, 8] where *connectionless servers* (CLS) attached to ATM switches provide routing functions for connectionless traffic. In Figure 1 we illustrate the direct approach where two connectionless LANs access the ATM network through interworking units (IWU). Each IWU maintains a permanent virtual path (VP) to a CLS which is attached to an ATM switch. CLSs are endpoints of virtual paths in the ATM network, and each CLS is responsible for maintaining a permanent VP to at least one other CLS. When a local area network transmits connectionless traffic to a remote network, it actually transmits a series of packets¹ to its IWU. The IWU, in turn, relays the packets to a CLS, which then passes the packets to another CLS closer to its destination.

Any network protocol that implements the direct approach must address the issue of bandwidth allocation to connection-oriented and connectionless traffic. Obviously, a static allocation of bandwidth to VPs

carrying connectionless traffic will result in wastage of bandwidth during periods of low intensity of connectionless traffic, and in congestion during periods of high traffic intensity. A better alternative is to vary the bandwidth allocation dynamically, using the actual connectionless traffic load as a heuristic for determining the need for bandwidth. However, even if sufficient bandwidth is allocated to connectionless traffic, one still has to address the problem of distributing the available bandwidth to a possibly large number of connectionless traffic flows.²




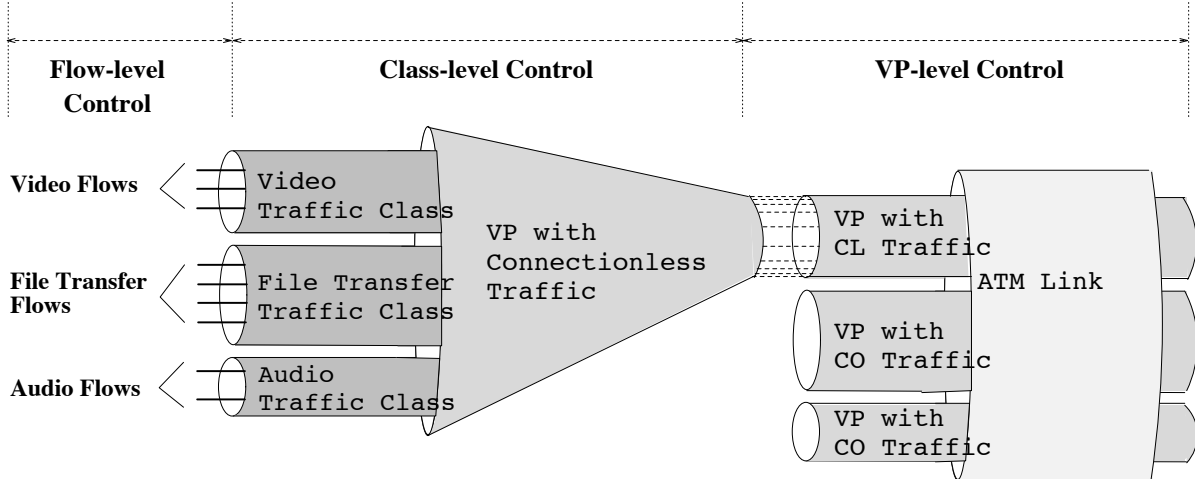
Figure 1: ATM Network with Connectionless Servers.

In this paper we propose a multi-level bandwidth control scheme for managing connectionless traffic in an ATM network that operates at three levels as shown in Figure 2: at the lowest level (*flow level*), we control the capacity available to a single connectionless traffic stream (=flow) between a pair of IWUs; at the next level (*class level*), we control the bandwidth made available to flows from the same traffic class; and at the highest level (*VP level*), we control the allocation of bandwidth to VPs carrying connectionless traffic. In Figure 2 the bandwidth at an ATM link is divided among VPs carrying connectionless and connection-oriented traffic. This stage of the allocation is governed by VP-level bandwidth control. The bandwidth allocated to the VP for connectionless traffic is then divided among connectionless traffic classes, such as video, file transfer, and audio traffic classes, by class-level bandwidth control. Finally, the bandwidth allo-

⁰The work of Debapriya Sarkar and Jörg Liebeherr was supported in part by the National Science Foundation under Grant No. NCR-9309224.

¹We use the term 'packet' to denote a *CL-PDU* (=Connectionless Protocol Data Unit).

²We use the term *flow* to refer to a stream of connectionless traffic between a pair of IWU. There can be an arbitrary number of flows between the same pair of IWUs.



(CO = connection-oriented; CL = connectionless)

Figure 2: Multi-level Bandwidth Control at an ATM Link.

cated to each traffic class is divided into flows belonging to that class. This stage of bandwidth allocation is determined by flow-level bandwidth control.

Even though several different schemes have been suggested for allocating bandwidth to connectionless traffic in ATM networks, including *bandwidth renegotiation* [5] and *fast reservation protocol* [1], our work is the first attempt to provide a multi-level bandwidth regulation of connectionless traffic in ATM networks.

2 Multi-level Control of Connectionless Traffic

In this section, we formally develop the proposed multi-level bandwidth control scheme. In the following description we use “network” to denote a virtual network of VPs carrying connectionless traffic between ATM switches with a CLS. When we need to refer to the actual physical ATM network, we shall state so explicitly.

We consider a network as shown in Figure 1 where connectionless servers (CLS) are connected by VPs. The overlay network is accessed via *interworking units (IWUs)* which are connected to a CLS by virtual paths. A *flow* in this network is a single traffic stream between two IWUs (see Figure 1) and each flow belongs to a *traffic class*. We need the following notation:

- \mathcal{L} Set of all VPs in the network.
- C_l Capacity of VP l .
- \mathcal{R}_i Route of flow i ; $\mathcal{R}_i = (l_{i_1}, l_{i_2}, \dots, l_{i_K})$ where l_{i_k} is the k th VP on the route of flow i .
- P Number of traffic classes.
- \mathcal{F}_p Set of all class- p flows.
- Δ_{lp} Set of flows in class p with VP l on their route ($\Delta_{lp} = \{i \mid l \in \mathcal{R}_i, i \in \mathcal{F}_p\}$).

We assume that all traffic in the network can be accurately described in terms of traffic rates. The traffic rate which describes the bandwidth demand of a flow, say flow i , is referred to as the *offered load* and de-

noted by λ_i . The rate of actual data transmission is called the *throughput* of the flow and denoted by γ_i .

At each VP l , there is a *class guarantee* G_{lp} which indicates the minimum bandwidth that the flows in class p can use for transmission (with $\sum_{p=1}^P G_{lp} = C_l$). In addition, at each VP l , a *class capacity* C_{lp} , gives the maximum capacity that can be used by all class- p flows, i.e., $\sum_{i \in \Delta_{lp}} \gamma_i \leq C_{lp}$. Note that the class guarantees do not specify how the bandwidth is distributed to individual flows.

2.1 Flow-Level Bandwidth Control

In this subsection we assume that all class capacities are fixed and given by the class guarantee G_{lp} , that is, $C_{lp} \equiv G_{lp}$ for all traffic classes and all VPs. With fixed class capacities, the bandwidth that is left unused by some traffic class cannot be made available to other traffic classes.

A flow-level bandwidth control scheme dictates how to distribute the class capacity C_{lp} to the class- p flows on a VP l . We assume that the goal of flow-level bandwidth control is to achieve *max-min fairness*, a scheme that was within the context of ABR traffic control [7]. Max-min fairness enforces that flows in the same class p have an identical bandwidth limit $\alpha_p(l)$. The value $\alpha_p(l)$ is referred to as the *fair share* of traffic class p on VP l . Obviously, in a network that enforces fair shares, the end-to-end throughput of flow i is limited by $\alpha_p(l_i^*) = \min_{l \in \mathcal{R}_i} \alpha_p(l)$, that is, the smallest fair share

on the route of the flow. For each flow i , we use l_i^* to denote the VP with the smallest bandwidth limit on the route, referred to as the *bottleneck VP* of flow i . Note that a flow-level control scheme that is based on *shares* enforces an intuitive notion of *fairness*, in the sense that all flows with the same bottleneck VP have identical throughput limits [3, 9].

If the network guarantees that each flow i can obtain the share at its bottleneck VP l_i^* then we obtain:

$$\gamma_i = \min(\lambda_i, \alpha_p(l_i^*))$$

We use $\alpha_p^*(l)$ to denote the largest possible value for the fair share of class- p on VP l , referred to as the *maximal share*. With maximal shares the entire class capacity on the bottleneck VPs of those flows which cannot transmit their entire load is utilized. In other words, if there exists a class- p flow i flow with $\gamma_i < \lambda_i$ and bottleneck VP k , then maximal shares ensure that $\sum_{j \in \Delta_{k,p}} \gamma_j = C_{k,p}$. Next we discuss how to determine the values of the maximal fair shares.

Given the fair share on each VP we can partition the class- p flows on a VP l into three groups: *underloaded*, *overloaded*, and *restricted*. The set of underloaded flows, denoted by $\mathbf{U}_{l,p}$, contains all class- p flows on VP l that can satisfy their end-to-end bandwidth demand, i.e., $\gamma_i = \lambda_i$. All other flows have bandwidth requirements larger than their throughput, i.e., $\gamma_i < \lambda_i$; these flows are classified as overloaded or restricted. Flows that are ‘overloaded on VP l ’ have VP l as their bottleneck. Flows on VP l that are ‘restricted’ have their bottleneck on some VP k on their route with $k \neq l$. For VP l , let $\mathbf{O}_{l,p}$ be the set of overloaded class- p flows, and let $\mathbf{R}_{l,p}(k)$ be the set of restricted class- p flows that have their bottleneck on VP k . The sets $\mathbf{U}_{l,p}$, $\mathbf{O}_{l,p}$ and $\mathbf{R}_{l,p}(k)$ are specified as follows:

$$\begin{aligned} \mathbf{U}_{l,p} &= \left\{ i \in \Delta_{l,p} \mid \alpha_p(l) \geq \lambda_i, i \notin \bigcup_{k \in \mathcal{C}} \mathbf{R}_{l,p}(k) \right\} \\ \mathbf{O}_{l,p} &= \left\{ i \in \Delta_{l,p} \mid l = l_i^*, \alpha_p(l) < \lambda_i \right\} \\ \mathbf{R}_{l,p}(k) &= \left\{ i \in \Delta_{l,p} \mid k = l_i^*, \alpha_p(k) < \lambda_i, k \neq l \right\} \end{aligned}$$

Now we characterize the maximal fair shares in a network. Recall, that with maximal shares the entire class capacity $C_{l,p}$ on a VP l is utilized if there is at least one flow in class p that is overloaded on this VP. We obtain:

$$\begin{aligned} C_{l,p} &= \sum_{i \in \Delta_{l,p}} \gamma_i = \sum_{i \in \Delta_{l,p}} \min(\lambda_i, \alpha_p^*(l_i^*)) \\ &= |\mathbf{O}_{l,p}| \cdot \alpha_p^*(l) + \sum_{i \in \mathbf{U}_{l,p}} \lambda_i + \sum_{k \in \mathcal{C}} |\mathbf{R}_{l,p}(k)| \cdot \alpha_p^*(k) \end{aligned}$$

For VPs without overloaded class- p flows, $\mathbf{O}_{l,p} = \emptyset$, we set $\alpha_p^*(l) = \infty$. Then we obtain the following values for the maximal shares:

$$\alpha_p^*(l) = \begin{cases} \infty & \text{if } \mathbf{O}_{l,p} = \emptyset \\ \frac{C_{l,p} - \sum_{i \in \mathbf{U}_{l,p}} \lambda_i - \sum_{k \in \mathcal{C}} |\mathbf{R}_{l,p}(k)| \cdot \alpha_p^*(k)}{|\mathbf{O}_{l,p}|} & \text{otherwise} \end{cases}$$

Thus, the maximal share is obtained by subtracting the throughput of the flows that are *not* overloaded

from the class capacity, and by dividing the remaining bandwidth by the number of overloaded flows. Note that for the calculation of maximal shares, no knowledge is required on individual flows. Precisely, two pieces of information must be available to determine the maximal share for class p on a VP: (1) the throughput of class- p flows that are not overloaded on the VP, and (2) the number of overloaded class- p flows on the VP.

2.2 Class-Level Bandwidth Control

The bandwidth control scheme described so far has one major drawback: if the flows in the same class, say p , do not consume their class guarantee $G_{l,p}$ on a VP l , this unused bandwidth cannot be utilized by other traffic classes. Class-level bandwidth control overcomes this drawback by adapting the available capacity $C_{l,p}$ to the actual traffic demand.

If a class p does not utilize its class guarantee, i.e., $G_{l,p} - \sum_{i \in \Delta_{l,p}} \gamma_i$, then class-level bandwidth control makes the unused bandwidth temporarily available to other traffic classes q by increasing $C_{l,q}$. Two issues remain to be resolved for a class-level bandwidth control scheme. First, one must determine the amount of unused bandwidth that can be given to needy traffic classes. Second, one must devise a policy for distributing the ‘surplus’ bandwidth. We will first address the second issue.

We consider a policy where the amount of additional bandwidth that is temporarily given to a class q is proportional to its class guarantee $G_{l,q}$. Consider a traffic class p with a high demand for bandwidth and a class guarantee of $G_{l,p}$ on VP l . Then we let $\rho_l G_{l,p}$ be the maximum bandwidth that class p can ‘borrow’ from the guarantees of other classes. We refer to ρ_l as the *class capacity surplus (or C-surplus) factor* of VP l . With the C-surplus factor, the maximum class capacity of a class p on a VP l is given by $(1 + \rho_l)G_{l,p}$. On the other hand, if a traffic class does not utilize its class guarantee, then the available class capacity $C_{l,p}$ is reduced. More formally, the class capacity $C_{l,p}$ on VP l for class p is set to:

$$C_{l,p} = \min \left(\sum_{i \in \Delta_{l,p}} \min(\lambda_i, \alpha_p(l_i^*)), G_{l,p}(1 + \rho_l) \right)$$

The above equation assumes the enforcement of fair shares $\alpha_p(l)$ for all flows, as discussed in Subsection 2.1.

A goal of a bandwidth control scheme is to select the C-surplus factors ρ_l as large as possible. If the C-surplus factor on a VP l is chosen maximally, denoted by ρ_l^* , then the entire VP bandwidth can be made available for transmission if there is an overloaded flow in some class on this VP. Formally, we have $\sum_{p=1}^P \sum_{i \in \Delta_{l,p}} \gamma_i = C_l$ whenever $\sum_{i \in \Delta_{l,q}} \gamma_i = G_{l,q}(1 + \rho_l)$ for at least one traffic class q on this VP.

Note that only the traffic classes with overloaded flows on VP l have a bandwidth demand that exceeds

their allotted bandwidth. Assuming that at least one such class exists on a VP l , i.e., $|\mathbf{O}_{lq}| > 0$, and that the maximal fair shares $\alpha_p^*(l)$ are available for VP l , then a bandwidth control scheme which enforces the maximal fair shares and maximal C-surplus factors satisfies the following equation:

$$C_l = (1 + \rho_l^*) \cdot \sum_{\mathbf{O}_{lq} \neq \emptyset} G_{lq} + \sum_{\mathbf{O}_{lq} = \emptyset} \left(\sum_{i \in \mathbf{U}_{lq}} \lambda_i - \sum_{k \in \mathcal{L}} |\mathbf{R}_{lp}(k)| \cdot \alpha_p^*(k) \right)$$

If the VP does not have any overloaded flows, that is, $|\mathbf{O}_{lq}| = 0$ for all traffic classes, the C-surplus factor can be selected arbitrarily large, i.e., $\rho_l^* = \infty$. Then we obtain that a bandwidth control scheme with maximal C-surplus factors $\rho_l^* > 0$ satisfies:

$$\rho_l^* = \begin{cases} \infty & \text{if } \bigcup_{q=1}^P \mathbf{O}_{lq} = \emptyset \\ \frac{1}{\sum_{\mathbf{O}_{lq} \neq \emptyset} G_{lq}} \left(C_l - \sum_{\mathbf{O}_{lq} \neq \emptyset} G_{lq} - \sum_{\mathbf{O}_{lq} = \emptyset} \left(\sum_{i \in \mathbf{U}_{lq}} \lambda_i + \sum_{k \in \mathcal{L}} |\mathbf{R}_{lp}(k)| \cdot \alpha_p^*(k) \right) \right) & \text{otherwise} \end{cases}$$

2.3 VP-Level Bandwidth Control

So far we have been concerned only with the virtual network of VPs that carry connectionless traffic between the CLSs. Next we consider the dependence of C_l , the bandwidth that is allocated for connectionless traffic on a VP l , on the bandwidth allocation to connection-oriented traffic³. Let C_l^{atm} denote the capacity of the ATM link l , and let λ_i^{co} denote the capacity that is currently being allocated to connection-oriented traffic. The bandwidth allocated to the VP carrying connectionless traffic is given by C_l . We introduce C_l^{min} as the lower bound for the bandwidth that is allocated to VP l with connectionless traffic. We adapt the capacity of the connectionless VP, C_l , to the bandwidth allocated to connection-oriented traffic. The following assignment of capacity to the VP with connectionless traffic implements this objective:

$$C_l = \min \left(\sum_{i \in \Delta_{lp}} \min(\lambda_i, \alpha_p(l_i^*)), \max(C_l^{min}, C_l^{atm} - \lambda_i^{co}) \right)$$

³For the sake of a simplicity we make the somewhat simplifying assumption that there is only one VP l for connectionless traffic on each ATM link.

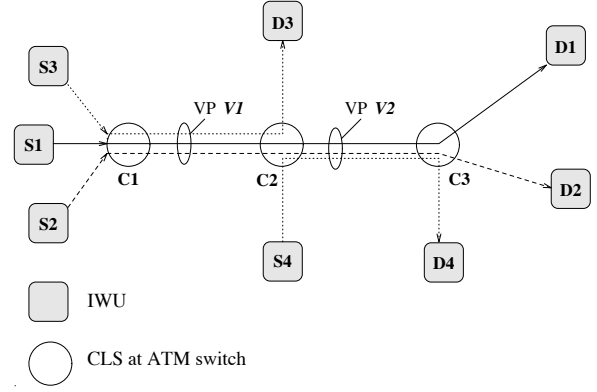


Figure 3: Simulated Network.

From	To	Route	Traffic Class	Offered Load (Mb/s)	Start Time (sec)
$S1$	$D1$	(V1, V2)	0	70 Mb/s	$t = 10$
$S2$	$D2$	(V1, V2)	II	65 Mb/s	$t = 20$
$S3$	$D3$	(V1, V2)	II	70 Mb/s	$t = 30$
$S4$	$D4$	(V2)	I	80 Mb/s	$t = 40$

Table 1: Flow Parameters.

3 Simulation Experiments

We have designed a protocol that implements the multi-level bandwidth control scheme from the previous section in a distributed fashion. The protocol describes a feedback loop that involves the interworking units (IWUs) and the connectionless servers (CLSs) of the ATM network. The details of the protocol are reported elsewhere [6]. Here we wish to present experiments from a simulation program implemented using the REAL (version 4.0) network simulator [4].

We make the following assumptions in the simulations. The packet sizes are constant for all flows and set to 1056 bytes which are split into twenty-four 44-byte chunks. Propagation delays are small and set to 1 ms. We assume that IWUs are the sources of a flow, that the offered load by each flow is known, and that each IWU generates packets after fixed time intervals. Cell losses due to transmission errors or buffer overflows at CLSs do not occur. The latter is achieved by selecting the buffer sizes at the CLSs and the ATM switches sufficiently large. Also, end-to-end window flow control mechanisms are not used in the simulation. Finally, the scheduling discipline at all CLSs is assumed to be FIFO.

As shown in Figure 3, the topology of the simulated network consists of eight IWUs, $S1 - S4$ and $D1 - D4$, and three CLSs, $C1 - C3$. The capacities of the ATM link, and VPs V1 and V2 are given by 622 Mb/s and 155 Mb/s, respectively.

We simulate the behavior of four flows from three different traffic classes: 0 , I , and II . The class guarantees of the traffic classes are set such that classes 0 , I , and II obtain guarantees on 25%, 30%, and 45%, respectively, of the total VP bandwidth. Thus, we have $G_0 = 38.75$ Mb/s, $G_I = 46.5$ Mb/s, and $G_{II} = 69.75$ Mb/s.

All flows start out with an offered load of $\lambda = 10$

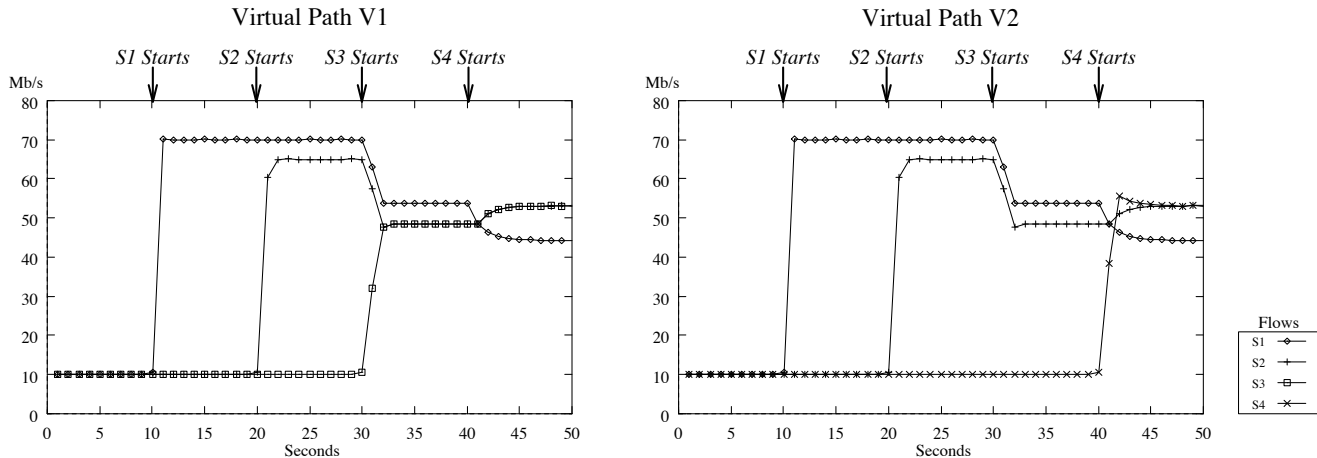


Figure 4: Flow-level Bandwidth Control (Fixed Class Guarantees).

Mb/s initially and increase their loads at specified time instants. The parameters of the four flows in Figure 3 and the start times of the load increases are summarized in Table 1. Since each IWU is the source or destination of at most one flow, we will use the source IWU to identify a flow. The values for the shares and the class capacities are updated each second.

We show the results of only one simulation experiment that demonstrates the effects of both flow-level and class-level bandwidth control. Refer to [6] for a more extensive set of experiments. The simulation results are summarized in Figure 4. The figure depicts two graphs that show the bandwidth (in Mb/s) utilized by each flow on the two VPs, V1 and V2. Each data point in the graph corresponds to the amount of data that is transmitted during an interval of 1 second. The experimental results were found to match the theoretically expected values from Section 2. Next we discuss the outcome of the simulation.

- At $t = 0$, all flows $S1-S4$ start transmission with $\lambda = 10$ Mb/s each. Each flow is underloaded and can send its entire load.
- At $t = 10$, class- 0 flow $S1$ increases its load to 70 Mb/s which exceeds the guarantee of class 0 . As none of the other classes have utilized their full guarantees, class 0 can ‘borrow’ extra bandwidth from the other classes. This allows $S1$ to transmit at its offered load.
- At $t = 20$, class- II flow $S2$ increases its offered load to 65 Mb/s. This causes class II to exceed its guarantee. However, class II can ‘borrow’ sufficient bandwidth from class I to allow $S2$ to transmit at its offered rate.
- At $t = 30$, flow $S3$ from class II increases its transmission to an offered load of 70 Mb/s. Then, classes 0 and II require all of their respective class guarantees on VP V1. Since there is no class- I traffic on VP V1, classes 0 and II can split the available bandwidth in the ratio of their respective guarantees. As $S1$ is the only flow in class 0 , it gets the entire bandwidth available to its class. However, the bandwidth available to class II on VP V1 is split evenly between $S2$ and $S3$.
- At $t = 40$, flow $S4$ from class I increases its load to 80 Mb/s. Since flow $S4$ now requires its entire class guarantee of 46.5 Mb/s on VP V2, all other classes are forced to reduce transmissions to their respective guar-

antees. The reduced ‘surplus’ of VP V2 decreases the throughput available to $S1$, and causes a shift of flow $S1$ ’s bottleneck from V1 to V2. This in turn, makes bandwidth available for the class- II flows on VP V1, yielding a throughput increase for $S2$ and $S3$. This, however, reduces the bandwidth available to flows $S1$ and $S4$, which have their bottlenecks on V2, causing their throughputs to drop. The drop in throughput of $S1$ causes another, smaller increase in the throughputs of $S2$ and $S3$. Note from Figure 4 that the protocol requires a few iterations before settling at the correct bandwidth allocation.

References

- [1] P. E. Boyer and D. P. Tranchier. A Reservation Principle with Applications to the ATM Traffic Control. *Computer Networks and ISDN Systems*, 24:321–334, 1992.
- [2] CCITT. CCITT Recommendation I.327: B-ISDN Functional Architecture, 1991.
- [3] J. M. Jaffe. Bottleneck Flow Control. *IEEE Transactions on Communications*, 29(7):954–962, July 1981.
- [4] S. Keshav. REAL: A Network Simulator. Technical Report 88/472, Computer Science Department, University of California, Berkeley, December 1988.
- [5] L. Mongiovi, M. Farrell, and V. Trecordi. A Proposal for Interconnecting FDDI Networks Through B-ISDN. In *Proc. IEEE Infocom’91*, pages 1160–1167, April 1991.
- [6] D. Sarkar. Distributed Bandwidth Regulation Mechanisms for Multiple Traffic Classes in Wide Area Networks, August 1995.
- [7] S. S. Sathaye. Draft ATM Forum Traffic Management Specification Version 4.0. ATM Forum/95-0013R8, November 1995.
- [8] B. J. Vickers and T. Suda. Connectionless Service for Public ATM Networks. *IEEE Communications Magazine*, pages 34 – 42, August 1994.
- [9] M. Zukerman and S. Chan. Fairness in ATM Networks. *Computer Networks and ISDN Systems*, 26(1):109–117, September 1993.