STUDY OF IP MULTIMEDIA TRAFFIC RUNNING OVER ABR ATM CONNECTIONS AND USING DIFFSERV POLICIES WITHIN THE INTERNET SEGMENTS

Alejandro Talavera Martínez, José R. Gallardo‡*, Dimitrios Makrakis‡**

*Advanced Communications Engineering Centre
Department of Electrical and Computer Engineering
The University of Western Ontario
London, Ontario, Canada N6A 5B9
martinez@engga.uwo.ca

‡Broadband, Wireless and Internetworking Laboratory
School of Information Technology and Engineering
University of Ottawa
Ottawa, Ontario, Canada K1N 6N5
{gallardo, dimitris}@site.uottawa.ca

ABSTRACT

We analyze the behavior of multimedia IP traffic running over an ATM network following the CIOA (Classical IP and ARP over ATM) approach. We propose to apply the concept of Expedited Forwarding Per Hop Behavior (EF PHB) within the Internet segments. EF is one of the mechanisms proposed by the IETF for Differentiated Services on the Internet. In our proposal, EF is applied within the Internet routers at the IP layer combined with the ABR feedback information from the ATM network. In this study we are using realistic traffic models to achieve more accurate results and conclusions.

1. INTRODUCTION

Today, the accelerated Internet development and the much more demanding services required by the network users have driven us to look for ways to provide a better service than the best-effort service traditionally offered by the Internet. Usually these new services are requested by multimedia applications, which have strict requirements for delay, loss ratio and bandwidth.

In this study we analyze the behavior of multimedia IP traffic, running over an ATM network. We use the approach described in the CIOA (Classical IP and Address Resolution Protocol over ATM) proposal [10]. This work focuses on a Virtual Private Network with users running multimedia applications. Since we are dealing with real-time traffic, we propose to apply the concept of Expedited Forwarding (EF) Per Hop Behavior (PHB) [6] within the Internet segments. EF is one of the mechanisms proposed by the IETF for Differentiated Services (DiffServ) on the Internet. We selected this DiffServ approach in order to build an end to end service for the aggregate traffic with low losses, low latency, low jitter, and assured bandwidth through a private network.

During the development of this work, the first author was a student at CICESE Research Center, Ensenada, B. C., Mexico, and was sponsored by a scholarship from the Mexican Science and Technology Council (CONACYT). He is currently with OPNET Technologies, Inc.

The second author is also with CICESE Research Center.

2. MODEL DESCRIPTION

We are assuming the presence of IP traffic streams with different QoS requirements, represented by voice, video, and data (FTP). EF in our model, as an addition to the CIOA, consists of a mechanism at the IP level that controls what kind of traffic goes first into the ATM network. As CIOA indicates, the traffic between two Logical IP Subnetworks (LIS) needs to go through an intermediate Internet router, which will be the one to perform the address resolution function. This means that the end users cannot use ATM connections alone. This in turn implies that the bandwidth guarantees given by the ATM connections are lost on those non-ATM links. Thus, we are applying EF to alleviate this weakness in the CIOA and to increase our chances of satisfying the multimedia traffic requirements.

ABR-compatible ATM switches work by sending feedback information to the service users notifying of the maximum rate at which they should transmit information (known as the Available Cell Rate – ACR) in accordance with the current switch state [5]. That is, ACR is dynamically changed between Minimum Cell Rate (MCR) and Peak Cell Rate (PCR), according to the available bandwidth in the network after servicing the background (VBR and CBR) traffic. As shown in Figure 1, EF is applied within the router at the IP layer combined with the ABR feedback information from the ATM network.

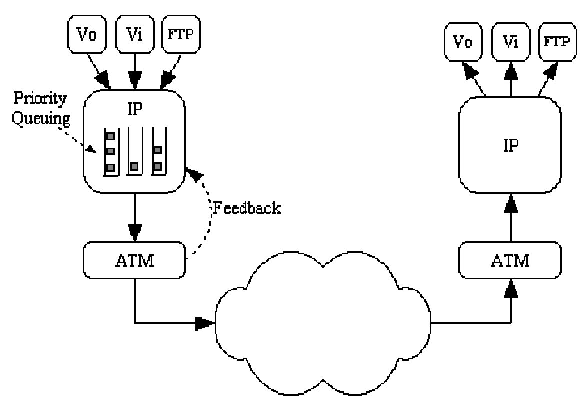

Figure 1. Model implementation

We pursue the following objectives:

- Analyze the behavior of a Classical IP over ATM network when there are multimedia applications running on it.
- Determine the minimum bandwidth requirements for the ATM connections to satisfy the demands from the applications, based on delays, losses and throughput.
- Show how applying the Expedited Forwarding mechanism at the IP layer can improve the service for applications with very strict time constraints and loss ratios, such as those in multimedia applications.
- Use realistic traffic models to achieve accurate results and conclusions. We compare the use of alpha-stable self-similar VBR background traffic against Poisson VBR background traffic.

EF proposes that the minimum departure rate for the high-priority traffic needs to be independent from the intensity of other traffic [6]. In order to accomplish this, and aiming at keeping our proposal as simple as possible, we use Priority Queuing Schedule. To avoid the starvation of the lower-priority traffic (as requested in [6]), we are implementing a threshold mechanism to limit the damage that the EF traffic could inflict on the competing traffic.

Since CIOA proposes the use of the ABR service to transport IP traffic, we follow that indication. This proposal is due to the fact that ABR has a feedback congestion control and minimum-cell-rate guarantees, and because it has a lower cost than VBR and CBR services. The UBR service is not considered as an option because it does not have minimum-bandwidth guarantees or a congestion control mechanism, which can result in a very poor performance.

In order for our study to be more realistic, we are performing tests using background VBR traffic in the ATM segment to see how it affects the behavior of the multimedia traffic that is using the ABR service. ABR is inevitably affected by the intensity of the VBR traffic, since the latter has a higher priority.

First, we are going to determine the minimum amount of resources (bandwidth) that a multimedia application would need from an ATM connection alone to achieve the correct levels of delay, loss ratio (for voice and video), and throughput (for FTP data).

Then, we run simulations with the whole CIOA configuration, adding the EF functionality, and using the parameters obtained in the first part of the experiment. In this step, the network is composed of two Logical IP subnetworks. Our goal here is to test how the presence of the intermediate router affects the behavior of the different applications, and to see if the values for minimum bandwidth obtained in the first part were suitably determined.

We are using realistic traffic models for the multimedia traffic (voice, video and FTP data) and also for the background VBR traffic in order to achieve more accurate results and conclusions. For video we use a source that reproduces real MPEG traces [2][13]. Our voice traffic generator is based on [8][12], and consists of a two-state Markovian model using G.729A compression. For FTP traffic, we implemented a Pareto-modulated ON/OFF data source from [11]. Finally, for aggregate VBR traffic, we are using an alpha-stable self-similar traffic model taken from [3]. For comparison purposes, we also include simulations using VBR traffic modeled as a Poisson process. It is important to say that, both the self-similar and the Poisson VBR background traffic have the same mean rate, 17 Mbps, to make a fair comparison of how both sources affect the performance of the ABR connections.

3. MODEL CONFIGURATIONS

To analyze the first scenario, we are using one of the standard ABR test configurations defined by the ATM Forum. We are using the two-node configuration since we need to have a bottleneck in which the network bandwidth changes dynamically, according to the demands from the background VBR traffic, as shown in Figure 2.

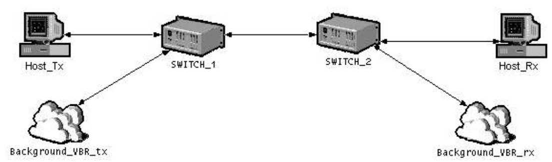

Figure 2. Two node configuration

In this part of our work, we consider the case of a router concentrating traffic from different sources (video, voice, and control and FTP data) and sending them as a single traffic stream through an ATM ABR connection.

The users run video over IP, voice over IP and data over IP. They ask for a minimum service rate, based on the corresponding traffic characteristics.

As mentioned above, the objective of this experiment is to see what is the minimum amount of resources that need to be reserved in order to carry the traffic across the network without violations. At the same time, we want to compare the performance achieved with and without using EF.

We will follow three different approaches:

(a) *All the types of traffic receive the same treatment at the IP layer without any discrimination.*

(b) *The Expedited Forwarding mechanism is applied at the IP layer, giving two different priorities: higher priority for voice and video, and lower priority for FTP traffic.*

(c) *The Expedited Forwarding mechanism is applied at the IP layer, giving different priorities to each traffic flow (voice, video and FTP data).*

We are running the above-mentioned scenarios for three different cases, each assuming respectively:

(I). No VBR background traffic in the network.

(II). Poisson distributed VBR background traffic in the network.

(III). Alpha-stable self-similar VBR background traffic in the network.

As the second step, we propose the configuration depicted in Figure 3 as our basic CIOA network, according to [10]. The objective here is to test the applications' behavior between two Logical IP Subnetworks with a router in between to see if the overall behavior is improved by applying the Expedited Forwarding mechanism.

The importance of implementing and analyzing this configuration lies on the fact that we are testing all the elements used in the Classical IP and ARP over ATM approach. This is a modified version of the two-node configuration mentioned above, in which the router specified by [10] has been added between the two subnetworks, and the number of users has been increased.

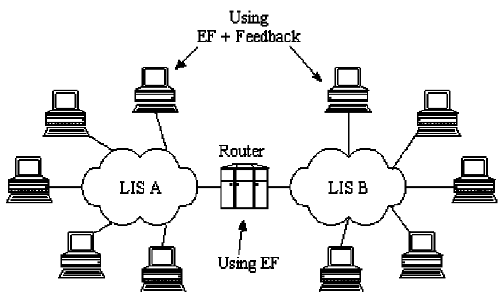

Figure 3. Two-LIS configuration

Similar to above, for this configuration we are going to run three cases:

(d) *All the types of traffic receive the same treatment at the IP layer without any discrimination.*

(e) *The Expedited Forwarding mechanism is applied at the IP layer, giving two different priorities: higher priority for voice and video, and lower priority for FTP traffic.*

(f) *The Expedited Forwarding mechanism is applied at the IP layer, giving different priorities to each traffic flow (voice, video and FTP data).*

4. SIMULATIONS PARAMETERS

We use the values shown in Table I for the requirements on delay and loss rate for voice and video, according to [9][4][8][7].

Table I. Video and voice restrictions

	Delay	Losses
Voice	112 ms	10^{-2}
Video	150 ms	10^{-4}

We assume a 100-kilometer distance between the end systems and the switch, so we apply a delay of 0.0005 seconds. The distance between the switches is assumed to be 1000 kilometers, so there is a delay of 0.005 seconds.

In our simulations, ABR and VBR traffic share a 20-Mbps link. That link capacity is assumed to be available after servicing CBR users.

Since we use long-range dependent traffic sources (such as the Pareto-modulated ON/OFF and the self-similar sources), for which the statistics converge slowly, we use a 1000-second simulation time.

All of our models were created and executed using the OPNET simulation tool.

5. RESULTS

(I). NO VBR BACKGROUND TRAFFIC IN THE NETWORK

First we are going to discuss the results obtained without using background traffic on the network. The next two graphs show the voice and video losses versus the Minimum Cell Rate requested for the ABR connection. The losses in these tests are due to packet rejection at the receiver because of an excessive end-to-end delay. We can see in the case of voice (see Figure 4) that a 1-Mbps MCR is enough to satisfy the voice loss requirements if we apply the three-priority EF mechanism, but we need more than 1.3 Mbps to achieve the required performance without using EF. For video, we can notice in Figure 5 that, due to a stricter requirement for losses, we need 1.2 Mbps to carry the video application through the network without violations when either version of the EF mechanism is applied. We can also observe that, without using EF, we need at least 1.5 Mbps as the MCR for the ABR connection. Thus, the minimum bandwidth required to have both applications (voice and video) work properly in terms of losses is 1.2 Mbps.

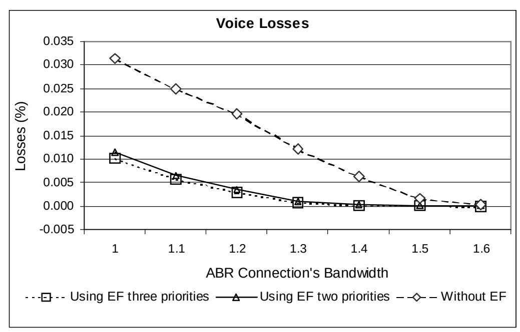

Figure 4. Voice losses vs. bandwidth

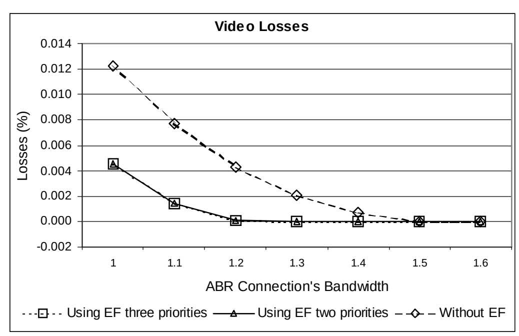

Figure 5. Video losses vs. bandwidth

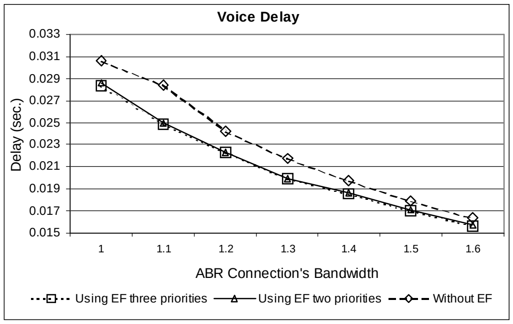

Figure 6. Voice mean delay vs. bandwidth

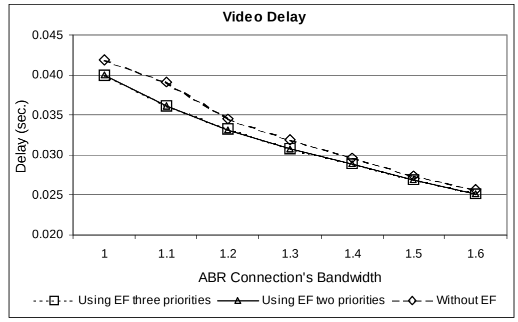

Figure 7. Video mean delay vs. bandwidth

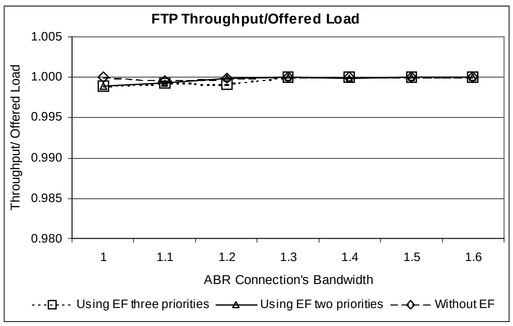

Figure 8. FTP Throughput / offered load vs. bandwidth

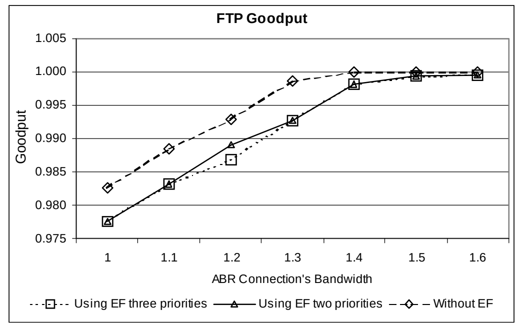

Figure 9. FTP goodput vs. bandwidth

The Expedited Forwarding mechanism, as shown in Figures 6 and 7, decreases the mean delay by up to ~12% for voice and up to ~4% for video packets. This improvement in the mean delay may not be particularly great, but the advantage of using ER, as described above, is the obtained reduction in the number of packet losses by abating the existence of individual packets with excessive delay.

A very important result is that FTP (data) traffic is not being dramatically affected by the application of the EF mechanism. The metrics used to evaluate the service given to the data traffic are: *i)* the ratio of the throughput to the offered load, and *ii)* the goodput, defined as the ratio of the number of useful packets (excluding retransmissions) to the total number of successfully transmitted packets. Thus, we can see in Figures 8 and 9 that, despite the preferential treatment given to the real-time traffic, the data performance is not greatly degraded.

From here we can observe that the overall number of packets transmitted, with and without EF, is roughly the same (a few packets are actually lost because of buffer overflow). What changes, mostly, is the order in which packets are serviced, which introduces an additional delay to the data packets. The data traffic is essentially insensitive to these extra delays, except for the (fortunately few) cases in which they cause TCP timeouts and retransmissions.

Another conclusion that can be drawn from these results is that there is no significant difference between using the EF mechanism with two or three priorities.

(II). POISSON DISTRIBUTED VBR BACKGROUND TRAFFIC.

Based on the previous results, we select a MCR of 1 Mbps and a PCR of 1.5 Mbps for the ABR connection. We ran this simulation adding VBR background traffic modeled as a Poisson process with a mean arrival rate of 17 Mbps.

As shown in Figure 10, the performance of all three packet-forwarding policies was acceptable for the voice users in terms of losses, as expected. However, we can also see that the use of EF dramatically improved the voice performance in terms of this metric. Under these conditions, no video losses were observed with any of the packet-forwarding policies.

Figures 11 and 12 show the mean delay for voice and video applications, respectively. It can be observed that the values are notably below the respective maximum allowed values. It can also be seen that there is no significant improvement when EF is used, similar to what happened when no background traffic was running through the ATM network.

As regards the FTP (data) traffic, our results revealed that there is no important difference whether we apply the EF mechanism or not, as shown in Figures 13 and 14. The conclusion from these tests is that the IP real-time applications are not significantly affected by the assumed load of Poisson VBR background traffic. To be more specific, our results show that, in the worst case, the ABR bandwidth (ACR) was reduced from 1.5 Mbps (requested PCR) to 1.475 Mbps, which shows that the Poisson background traffic only takes a very small amount of the bandwidth from the ABR connections.

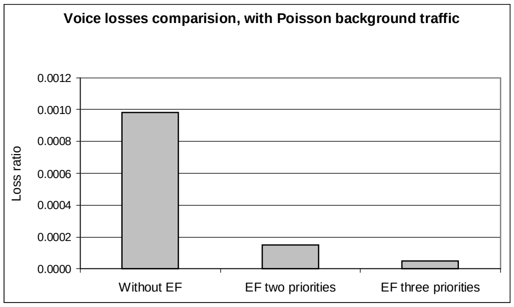

Figure 10. Voice losses using Poisson background traffic

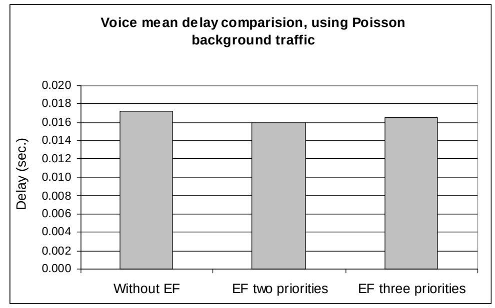

Figure 11. Voice delay using Poisson background traffic

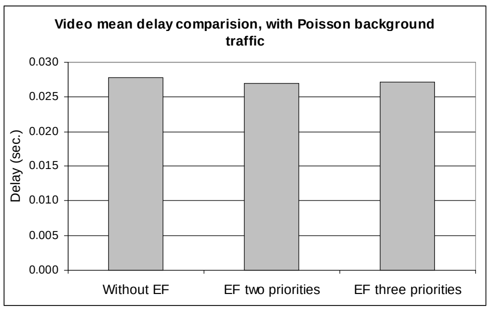

Figure 12. Video delay using Poisson background traffic

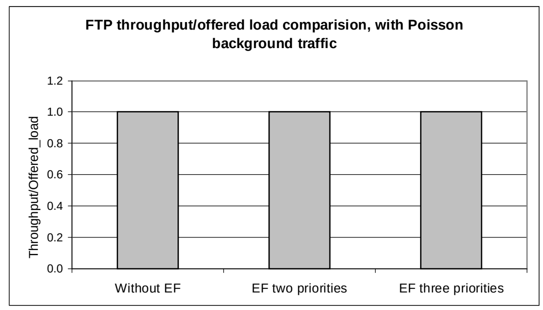

Figure 13. FTP throughput / offered load using Poisson background traffic

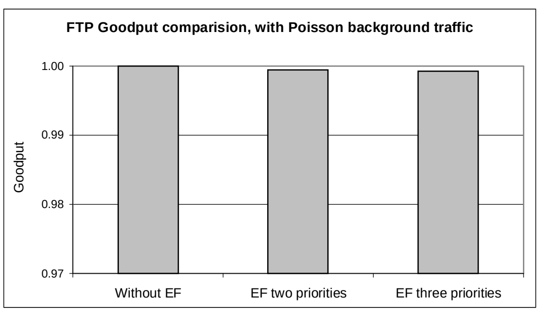

Figure 14. FTP goodput using Poisson background traffic

(III). ALPHA-STABLE SELF-SIMILAR VBR BACKGROUND TRAFFIC IN THE NETWORK.

In this part of our experiment, we assume that the VBR background traffic behaves as an alpha-stable self-similar process with the same mean arrival rate as the one assumed in part II (17 Mbps). We also use the same values assumed in part II for MCR and PCR (1 and 1.5 Mbps, respectively).

As Figure 15 shows, the use of EF dramatically improved the performance of the voice applications, decreasing the losses. However, in this case the voice application has an acceptable performance even without the use of the EF mechanism. The voice mean delay is marginally reduced by the use of EF, as shown in Figure 16.

Figure 17 indicates that there are violations to the acceptable number of video losses when we do not apply EF. With the EF mechanism, on the other hand, either with two or three priorities, the number of video losses falls dramatically (down to 0), allowing the transport of the video application without violations. As for the video mean delay, we notice a slight decrease when we apply the EF mechanism (see Figure 18).

Once again it is clear that, without the EF mechanism, when a large FTP burst goes into the network, it affects directly the video and voice packets that follow by greatly increasing their delay and causing losses. The EF mechanism, when used, causes a reduction in the voice and video peak delay values at the expense of enlarging FTP delays. The FTP packet transmission, however, is not noticeably affected by this procedure, as described below.

Regarding the FTP data packets, the achieved throughput is equal to the offered load (ratio is 1), as shown in Figure 19. We also obtained very high values for the Goodput, as seen in Figure 20. Again, these results are due to the resilience of FTP to longer delays, as compared to video or voice applications, as long as they do not cause too many TCP timeouts. This is also an indication that, under these conditions, there is a friendly association between TCP and the EF/ABR mechanism.

If we compare the results from this section and section II, we see that the realistic alpha-stable self-similar VBR background traffic interacts much more strongly with the ABR traffic than the Poisson traffic. This shows us that the common tendency of simplifying the traffic models to achieve easier and faster results can cause for the obtained conclusions to be completely inaccurate.

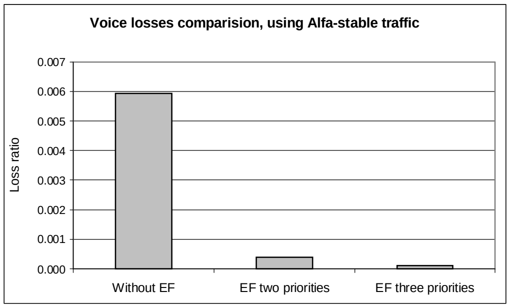

Figure 15. Voice losses using alpha-stable background traffic

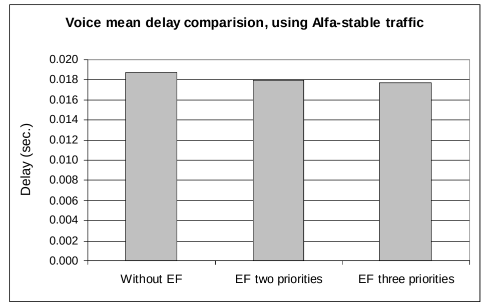

Figure 16. Voice delay using alpha-stable self-similar background traffic

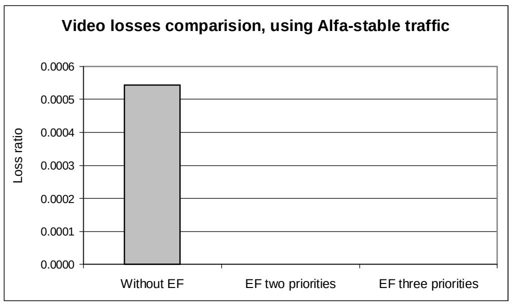

Figure 17. Video losses using alpha-stable background traffic

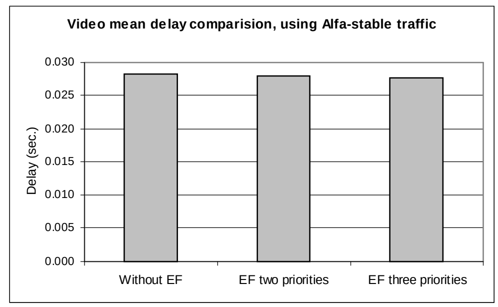

Figure 18. Video delay using alpha-stable background traffic

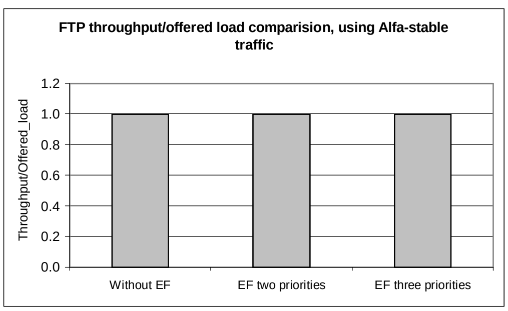

Figure 19. FTP throughput / offered load, using alpha-stable background traffic

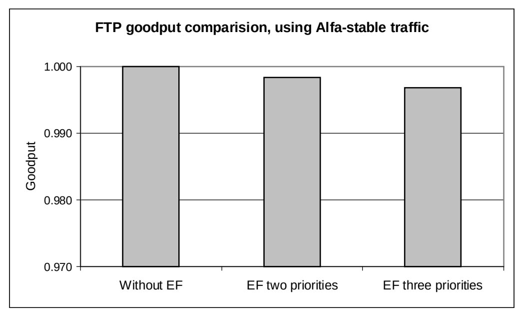

Figure 20. FTP goodput using alpha-stable background traffic

(IV). RESULTS WITH TWO LOGICAL IP SUBNETWORKS

These results have been obtained in our last scenario, running five multimedia users on the network. Once again we can see how the requirements for voice are well achieved, regardless of whether or not we apply the EF mechanism. In Figures 21 and 22 we can see that the voice losses and voice delay are within the permitted limits, and we can also notice that there is an important decrease in voice losses when EF is applied.

On the other hand, we notice that the EF mechanism is essential to avoid violations of video requirements. Without EF, the video losses exceed the limit for a good performance (see Figure 23). We can also see that, in this case, using two or three priorities in the EF mechanism works equally well for video applications.

In this case, as in most of the others, there is only a small reduction in video delay when we use EF, as can be appreciated in Figure 24.

FTP traffic has a very good performance in terms of the ratio of the throughput to the offered load because of the TCP error-recovery actions, but we can see how the *goodput* parameter decreases to 96% (see Figures 25 and 26). Thus, we can say that EF affects marginally the FTP goodput, but still within very acceptable limits.

6. CONCLUSIONS

- We found that the EF mechanism, which is applied at the IP layer, improves the performance of voice and video applications running on an ATM ABR connection, decreasing packet losses.
- The EF mechanism does not reduce significantly the mean delay for voice or video applications, but avoids instead the existence of excessive delay for individual packets. This is the main reason why voice and video losses are decreased, since it is much more likely for

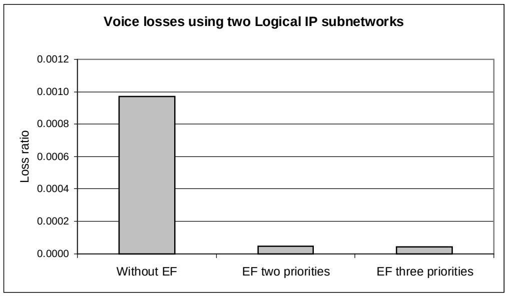

Figure 21. Voice losses using two LIS

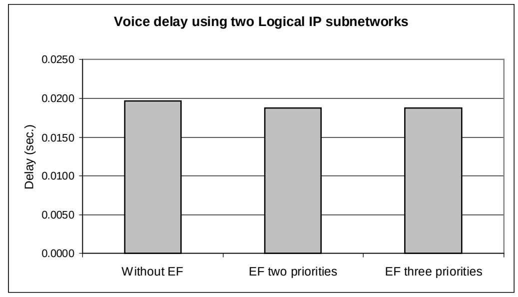

Figure 22. Voice delay using two LIS

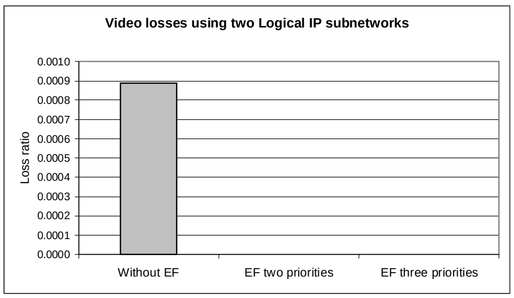

Figure 23. Video losses using two LIS

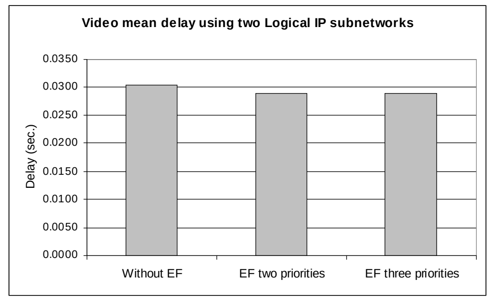

Figure 24. Video delay using two LIS

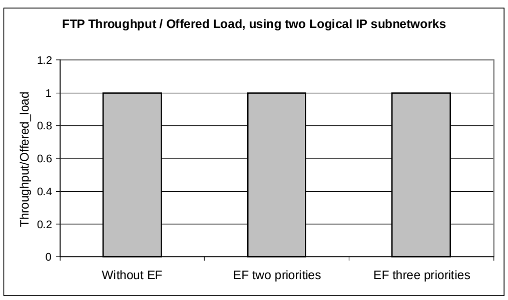

Figure 25. FTP throughput / offered load, using two LIS

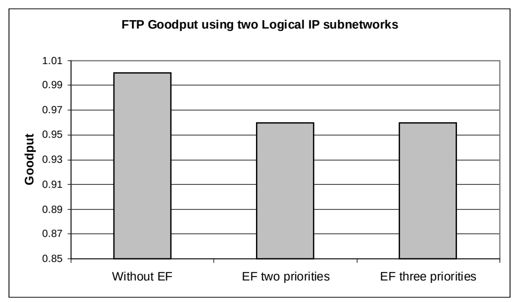

Figure 26. FTP goodput using two LIS

the packets to arrive in time at the destination, thus avoiding being discarded.

- As expected, the EF mechanism marginally affects in a negative way the FTP goodput, but keeping an acceptable FTP performance. We can say that there is a friendly association between TCP and the ABR/EF mechanism.
- There is no significant difference between using EF with two or three priorities. Thus, we recommend the use of two priorities, because it is simpler.
- The overall conclusion is that EF saves bandwidth and, at the same time, improves the network performance when transporting multimedia traffic.
- By comparing the system performance when the background traffic behaves as an alpha-stable self-similar process and when it behaves as a Poisson process, we can see that there is a significant difference. This in turn implies that simplifying our models in order to obtain easier and faster results can cause for the conclusions drawn to be completely inaccurate.
- In the Classical IP over ATM approach, in which separate logical IP subnetworks need to use an intermediate router in order to communicate, the proposed ABR/EF mechanism still maintained a good performance. The goodput for non-real-time applications, such as FTP, was marginally affected, but still within acceptable limits.

REFERENCES

[1] J. D. Cavanaugh, "Protocol Overhead in IP/ATM Networks", Technical Report 1994-08-12, Minnesota Supercomputer Center, Inc., 1994.

[2] ftp-info3.informatik.uni-wuerzburg.de/pub/MPEG/. MPEG traces, University of Würzburg, Germany.

[3] J. R. Gallardo, D. Makrakis, and L. Orozco-Barbosa, "Use of Alpha-Stable Self-Similar Stochastic Processes for Modeling Traffic in Broadband Networks". Performance Evaluation, Vol. 40, No. 1-3, pp. 71-98, March 2000.

[4] P. Goyal, A. Greenberg, R. Kalmanek, W. T. Marshall, P. Mishra, D. Nortz, and K. K. Ramakrishnan, "Integration of Call Signaling and Resource Management for IP

Telephony", IEEE Network, Vol.13, No.3, pp.24-32. May-June 1999.

[5] R. Goyal, R. Jain, S. Fahmy, S. Narayanaswamy, "Modeling Traffic Management in ATM Networks with OPNET", Proceedings of OPNETWORK'98, Washington DC, May 1998.

[6] V. Jacobson, "An Expedited Forwarding Per Hop Behavior", Internet RFC2598, June 1999.

[7] G. Karlsson, "Asynchronous Transfer of Video", IEEE Communications Magazine, Vol.34, No.8, pp.118-126. August 1996.

[8] T. J. Kostas, M. S. Borella, I. Sidhu, G. M. Shuster, J. Grabiec, and J. Mahler, "Real-Time Voice Over Packet-Switched Networks", IEEE Network, Vol.12, No.1, pp.18-27. Jan./Feb. 1998.

[9] M. Krunz, "Bandwidth Allocation Strategies for Transporting Variable-Bit-Rate Video Traffic", IEEE Communications Magazine, Vol.37, No.1, pp.40-46. Jan. 1999,

[10] M. Laubach and J. Halpern, "Classical IP and ARP over ATM", Internet RFC2225, April 1998.

[11] V. Paxson and S. Floyd, "Wide Area Traffic: The Failure of Poisson Modeling", IEEE/ACM Trans. Networking, Vol. 3, No. 1, pp. 223-244, June 1995.

[12] E. P. Rathgeb, "Modeling and Performance Comparison of Policing Mechanisms for ATM Networks", IEEE J. Select. Areas Commun., Vol. 9, No. 1, pp. 325-334, April 1991.

[13] O. Rose, "Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems". Proc. of the 20th Annual Conference on Local Computer Networks, Minneapolis, MN, 1995.