Simulation and Analysis of the Performance of EPRCA in a Wide Area ATM Network Consisting of Both ER and EFCI Switch Mechanisms

R. Walthall and W. Desposorio
Department of Electronic Engineering Technology
Brigham Young University
Provo, UT 84606
Email walthalr@et.byu.edu,
desposow@et.byu.edu

Abstract

ATM networks are seen as the most likely technology for implementing the National Information Infrastructure (NII). Successful ATM deployment relies on effective flow control for ABR traffic. In this paper we evaluate the performance of the Explicit Proportional Rate Control Algorithm (EPRCA) in a Wide Area Network (WAN) made up of both Explicit Rate (ER) and Explicit Forward Congestion Indicator (EFCI) switches. MPEG video traces from MPEG-2 encoded video streams are used to generate VBR traffic. Available Bit Rate (ABR) and Variable Bit Rate (VBR) sources are used to simulate the real world environment in which EPRCA will be required to operate (transient state). The performance of EPRCA is evaluated in terms of throughput, buffer occupancy, fairness, and fast convergence.

1 Introduction

Asynchronous Transfer Mode (ATM) is slated to become the protocol on which the information superhighway will rely to transfer voice, data and video. A great deal of research has already gone into ATM networks including numerous testbeds and trials throughout the United States, Japan, Europe, and Australia. It is probable that these networks will form the platform from which a global communications system will be launched. Successful deployment of ATM networks will require a reliable method of flow control for ABR service. This research evaluates the performance of the flow control mechanism as defined by the ATM Forum (EPRCA) under realistic operating conditions in a WAN.

ABR service is important to successful deployment of ATM networks because it is the means of maximizing the utilization of ATM switches. The UNI 3.1 specification declares that the cell loss probability for VBR and CBR traffic is required to be less than 10^{-8}[4]. Studies have shown that in order to maintain a cell loss probability of 10^{-12} for VBR traffic, a switch sending only VBR traffic would only be able to utilize approximately 30 percent of its capacity [8]. ABR service is not guaranteed to have a cell loss probability of 10^{-8}, but relies instead on a control system to maintain low cell loss. For over 2 years now the world's leading authorities on flow control have been working in the ATM Forum Traffic Management Technical Group in an effort to develop a specification that will meet the demands required of this control system. The result is a specification which defines end system behavior and the behavior for two types of switch control mechanisms: Explicit Rate (ER), and Explicit Forward Congestion Indicator (EFCI) [9]. The objective of the control system is the fair allocation of bandwidth to ABR users while maximizing throughput and minimizing cell loss and the required queue size (delay)[2]. Because of the complexity involved in analyzing flow control in ATM networks, the concepts driving the ATM Forum Traffic Management Specification Version 4.0 are based mostly on simulation. For this study we have used a modified version of the National Institute of Standards and Technology (NIST) ATM Network Simulator which is capable of simulating both EFCI and ER switches [5].

The EFCI switch is a simple mechanism, and was defined in UNI 3.1 [4]. When a switch is congested, a single bit will be set in the cell header on outgoing cells for all virtual circuits using ABR service. The cells will be forwarded to the destinations, and the destinations will send this information to the sources in the next Resource Management (RM) cell. When a source receives notice that the switch is congested it must reduce its transmission rate by a predetermined value which was negotiated at call setup time. There are a number of ATM switches currently on the market that support the EFCI switch mechanism. It was determined that the single bit binary feedback used by EFCI control mechanism did not provide a sufficient amount of control. Thus the focus of the Traffic Management Group shifted to the development of the ER switch mechanism/specification.

The ER switch is a fairly complex control mechanism capable of dynamically allocating bandwidth to ABR sources. The bandwidth is allocated through the use of RM cells. The ABR source is required to send an RM cell at a rate negotiated at call setup time. In the ER switch, the explicit rate at which each source may transmit is calculated and returned to the source in these RM cells. A number of papers have been written that describe this control mechanism in detail[2, 1, 7] . It was decided that the Traffic Management 4.0 specification will declare what must be accomplished by an ER switch, but not necessarily how it is to be accomplished. Thus, a number of methods have been developed that provide the required control. These algorithms vary widely in their complexity and effectiveness. For this study we have used the algorithm developed by researchers at NIST [2].

To date, most of the evaluation of the performance of the control mechanism has been based on steady state performance in networks consisting of either all ER or all EFCI switches. This is not likely to be the case in the real world. Because of the bursty nature of computer traffic, networks usually operate in a transient state. Also, because backwards compatibility was designed into the specification (interoperability between ER and EFCI switch mechanisms), there are a number of EFCI switches now in operation whose owners may not choose to upgrade to the new ER switch mechanism. This research evaluates the performance of the control system in a wide area ATM network consisting of both ER and EFCI switches in terms of throughput, buffer occupancy (delay), fairness, and fast convergence while operating in both the transient and steady state. The transient state is maintained by introducing VBR traffic in the form of MPEG-2 video streams.
In other words, the response of the control system is evaluated in an environment that is as close an approximation to its destined environment as possible.

Questions to be addressed in this paper:

1. How will the EPRCA perform when operating in a network where actual MPEG-2 encoded video streams are used to simulate VBR traffic (forcing transient response from EPRCA).

2. If the Customer Premise Equipment is not upgraded to the ER switch mechanism, will the telco system (assumed ER) be adversely affected.

3. What level of performance can the user expect if he/she chooses not to upgrade from the EFCI to ER switch mechanism.

4. How much of the degradation of performance is due to the VBR traffic forcing a transient response.

We expected to find that users of an EFCI edge switch would experience a loss of performance to in terms of: increased time to converge on the fair share, decreased throughput (switch utilization), increased delay due to increased buffer occupancy, and unfair allocation of rate to greedy ABR sources. Our results indicate that this is the case for both the transient and steady state response. We also expected to find that accepting an EFCI edge switch would have a negative effect on the users of the ER system in terms of increased buffer occupancy and unfairness of rate allocation. Our results indicate that this is not the case. In all simulations run so far, the buffer occupancy of the interior switches remained low. And while the rate allocation to users of the network interior (ER switches) did not remain fair, the result from the unfairness caused by the EFCI edge switch was beneficial to ER switch users in that their allowed rate was increased. These results are of interest to end users operating EFCI switches, switch manufacturers, researchers, and network engineers planning on providing ABR service in the WAN environment.

2 Simulation Environment

The simulation was configured for both the transient and steady state response. For the transient state, VBR traffic was generated in each of the configurations by ten VBR sources. A high action video was chosen in order to simulate the worst case environment in which the control mechanism will be required to perform. VBR traffic generators ramp up from 1 to 10 sources during the first second of simulation. The sources start at an interval of approximately 100 msecs. One VBR source active at time = 0, two at t = 100 msecs, and so forth. The test setup was intended to simulate a Video On Demand environment. The resulting VBR traffic pattern is shown in Figure 20. It is interesting to note that this traffic pattern is not a poisson distribution. For this reason, traditional methods of mathematical analysis, using Markov processes, are not applicable to real MPEG traffic in ATM networks.

It was determined by the ATM Forum that greedy, persistent sources present the worst case application for rate based flow control. ABR sources in all simulations are therefore greedy, persistent sources. All links have a capacity of 150Mbps. Backbone links (BB1, BB2) have a length of 1km. All other links are 0.1km. The ER control mechanism was configured for optimum performance, and is shown in Table 1.

2.1 MPEG Video Compression

MPEG compressed video is quickly becoming the standard form for video transmission. It is currently prevalent in PC video applications, CD-ROMs, and satellite video transmissions. MPEG-1 is designed for 1.5 Mbps audio and video transmission. A frame size of 320 by 240 pixels is typical. It can send at a rate of 30 frames per second, and supports CD quality sound. Motion compensation allows data to be shared between pictures, which results in high compression ratios for MPEG encoded video. MPEG uses discrete cosine transform algorithms that can decompose data sets into wave frequencies. The less important frequencies are discarded which allows even further compression. Compression ratios of 50:1 are typical.

MPEG-2 has been developed to present an even higher standard of video processing. It was designed to be capable of reproducing full-screen, broadcast quality video and sound. The MPEG-2 standard defines a systems transport stream which is capable of carrying and maintaining synchronization for multiple video and audio programs. Current speculation from industry is that the ability to broadcast audio and video through the Internet (Video On Demand) will cause the demise of broadcast television as we now know it [3]. For the transient response, traces from MPEG-2 transport streams (which were used to carry MPEG-1 encoded video sequences) were used to generate VBR traffic. The video was a simulation of a jet flying through canyon lands which presents a high action (worst case) environment for the control mechanism (EPRCA) to perform.

2.2 Edge Switch Assumptions

Early claims made by FORE Systems representatives were that the upgrade required to make their switches EFCI compliant would be a software upgrade (free of charge). History shows it to be a \$1000.00 per port hardware upgrade. In like manner, there are now speculations that the ER switch upgrade will be very inexpensive. Should this prove not to be the case there are many users who currently own EFCI compliant switches that may not choose to upgrade. It is unlikely however that the ATM service providers will choose not to upgrade their network interiors. It is for this reason that the test configurations are run with the edge switch using the EFCI mechanism while the network interior consists of all ER switches.

2.3 Test Configurations

Three network configurations have been developed by researchers at Ohio State University to evaluate specific characteristics of the control mechanism [7]. They are: the Fast Convergence, Parking Lot, and Upstream Bottleneck configurations. The transient configurations have been modified to include VBR traffic generators, and the edge switch is alternated first ER then EFCI.

2.3.1 Fast Convergence

The "Fast Convergence" configuration shown in Figure 22 tests for the time required for ABR sources to converge on the fair share of the available bandwidth. It consists of two ABR sources one of which begins transmitting at time = 0. The second ABR source begins transmitting some time later and continues transmitting until t = 2000 msecs. If no VBR traffic were present the fair share when both sources are transmitting would be 75 Mbps each.

Figures 1 - 6 show the results from replacing the edge switch with an EFCI switch. Figures 1 and 2 show a significant increase in the buffer occupancy of the edge switch which translates to increased delay.. Figure 3 and Figure 4 show a significant increase in time required to converge on the fair share with EFCI edge switches (slow response). Figures 5 and 6 show a decrease in the overall switch utilization (lower throughput).

2.3.2 Parking Lot

This test was devised largely because in the original (EFCI) scheme the probability that an ABR source would be asked to decrease its rate increased exponentially by the number of hops the Virtual Circuit (VC) was required to make [6]. This is called the beat down effect. This configuration tests for fairness of rate allocation.

The configuration shown in Figure 21 tests for fairness in that each source should be allocated 50 Mbps of bandwidth (assuming no VBR traffic) no matter how many hops are made by the VCs. The results from this test were unexpected. The performance variation from changing the edge switch to EFCI is almost unnoticeable. This result indicates that fair bandwidth allocation can occur even with EFCI edge switches.

Figure 7 and Figure 8 show that the maximum buffer showed no significant change. Figure 9 and Figure 10 show that fairness of rate allocation did not change. And Figures 11 and 12 show that the switch utilization remained the same.

2.3.3 Upstream Bottleneck

One of the criterion for max/min fairness is that all ABR sources should be allocated their fair share of the available bandwidth; provided they can use it. The third configuration "Upstream Bottleneck" (Figure 23)consists of three ABR sources at the edge switch two of which are single hop connections. With no VBR traffic these three sources would share the 150 Mbps available on BB1. ABR1, ABR2, and ABR3 should get 50 Mbps each; however, ABR4 and ABR2 share BB2. The fair share for BB2 is 75 Mbps each, but ABR 2 is constrained to 50 Mbps on BB1. Therefore the test is to see if the control mechanism will reallocate the bandwidth so that ABR4 gets 100 Mbps while ABR1 remains at 50 Mbps.

Figure 13 and Figure 14 show a significant increase in the buffer occupancy of the edge switch which corresponds to increased delay for ABR traffic. Figure 15 shows the rates for all ABR sources in the all ER switch network. It can be seen at time = 0 that ABR1, 2, and 3 get approximately 50 Mbps each, while ABR 4 gets roughly 90 Mbps. This is expected, because there is one VBR source (10 Mbps) active at this time. Two graphs were required to show the unfairness resulting from replacing the edge switch with and EFCI switch. In Figure 16 it can be seen that ABR1 and ABR3 get approximately 50 Mbps while ABR2 gets only 40 Mbps (Figure 17). ABR4 gets approximately 100 Mbps. The desired rate distribution is 50 Mbps for ABR1, 2, and 3. Link utilization was lower with the EFCI switch in place, as can be seen by comparing Figures 18 and 19.

Changing the edge switch resulted in unfair allocation of bandwidth and increased delay (buffer occupancy). This is undesirable, and we predict that similar results will occur on larger networks.

3 Conclusions

We found that when all switches are using the ER mechanism the control system performs well in all three test configurations with respect to low buffer occupancy, fast convergence, high throughput, and fairness even when MPEG-2 encoded video streams are used to force EPRCA into transient mode operation. However, replacing the edge switch with an EFCI switch resulted in significant degradation of performance for both the transient and steady state response.

- The Fast Convergence configuration showed:
  1. Increase in time required to converge on the fair share
  2. Increased delay due to increased buffer occupancy
  3. Lower throughput (switch utilization)
- The Parking Lot configuration showed:
  1. No significant change
- The Upstream Bottleneck configuration showed:
  1. Unfair allocation of bandwidth
  2. Increased delay (buffer occupancy)

The performance difference between using the ER and EFCI switch mechanism under steady state response was not significant enough to justify spending the money to upgrade. However, when VBR traffic forces a transient response the performance improves dramatically in terms of buffer occupancy and switch utilization. For instance, the buffer occupancy when using ER switches in the bottleneck configuration under transient response reaches 55 cells, compared with 250 when using the EFCI edge switch. Unfairness of bandwidth allocation was observed in the bottleneck configuration for both the transient and steady state response when the EFCI edge switch was implemented.
In some circumstances the increase in delay resulting from increased buffer occupancy could be significant. For instance, if a user's connection were a T-1 line (1.5Mbps), video applications would experience a noticeable loss in performance when buffers reach 250 cells.

We note that a customer may choose not to upgrade from EFCI to ER switch mechanism provided he/she is willing to accept loss of performance in terms of: increased time to converge on the fair share, decreased throughput (switch utilization), increased delay due to increased buffer occupancy, and unfair allocation of rate to greedy ABR sources. This may well be acceptable depending on the types of data applications being run by the user. For instance, if ABR service is to be used for email programs or web browsing it is unlikely that the loss in performance would be noticeable.

Finally, we note that the telco system is compromised in terms of reduced utilization of the backbone switches and links when and EFCI edge switch is used to control ABR traffic. Also, when using the EFCI edge switch for the Upstream Bottleneck configuration the control mechanism did not perform very well in that the source that should have been given 100Mbps received an average of 90Mbps (ABR7). This is a result of the algorithm that handles the case when a source does not utilize its fairshare of the bandwidth (ABR3). The result is lower utilization of both the switch and the backbone links. In this same configuration we note that ABR3 gets an unfair allocation of the available bandwidth. This unfairness occurred in the largest network configuration that was simulated. We predict that as the size of the network increases, and consequently the number of EFCI edge switches increases, the amount of unfairness will increase as well.

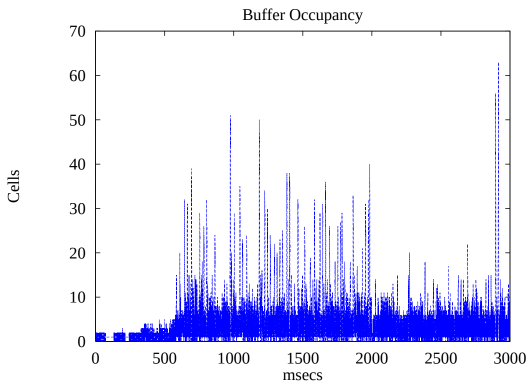

Figure 1: ER Buffer

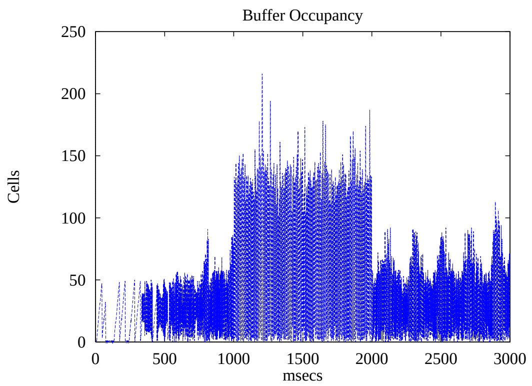

Figure 2: EFCI Buffer

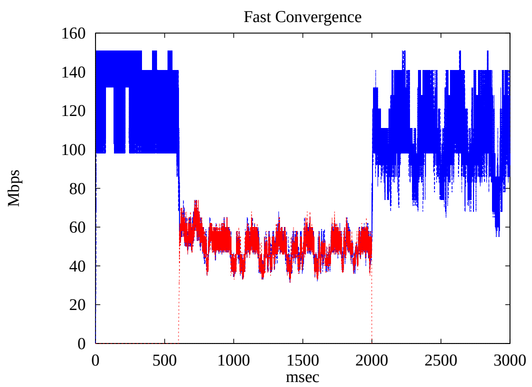

Figure 3: ER Convergence

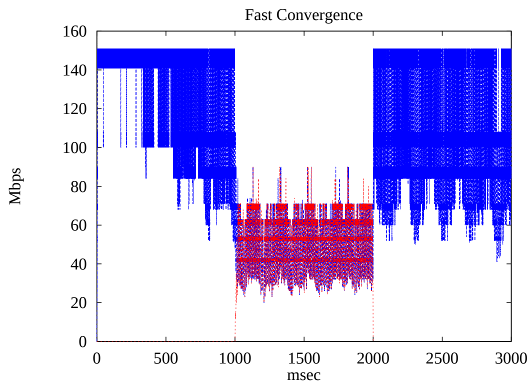

Figure 4: EFCI Convergence

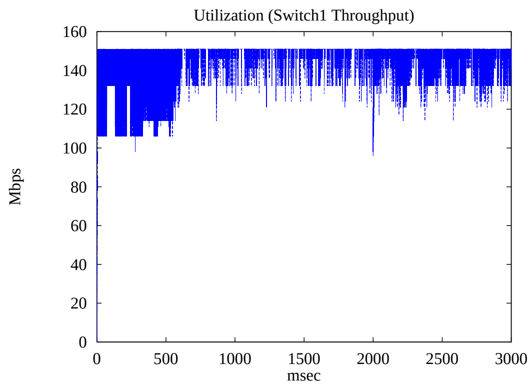

Figure 5: ER Utilization

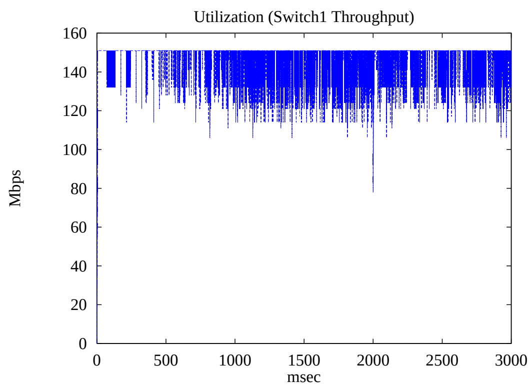

Figure 6: EFCI Utilization

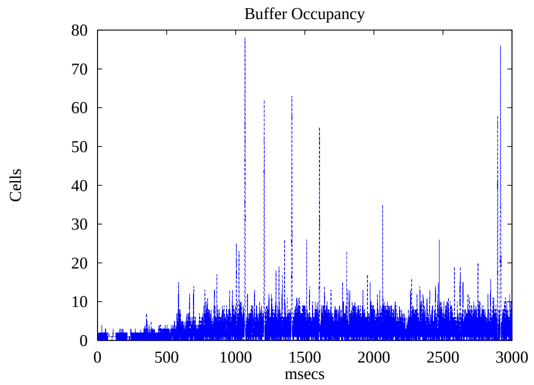

Figure 7: ER Buffer

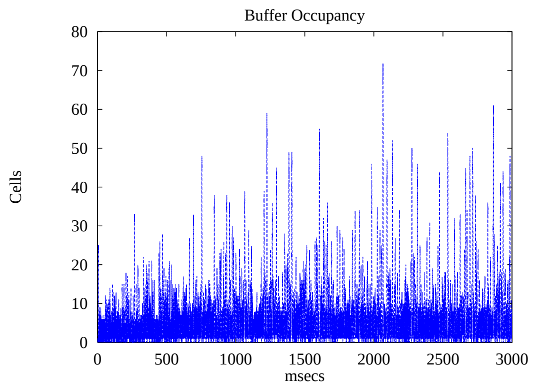

Figure 8: EFCI Buffer

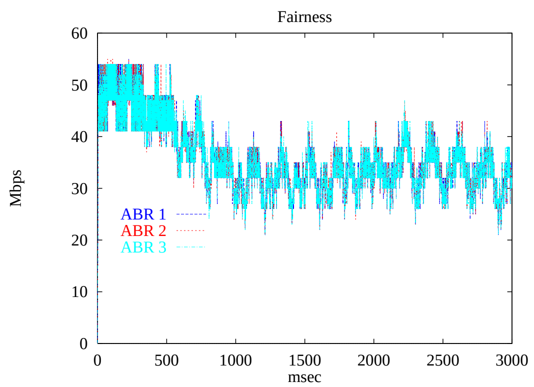

Figure 9: ER Fairness

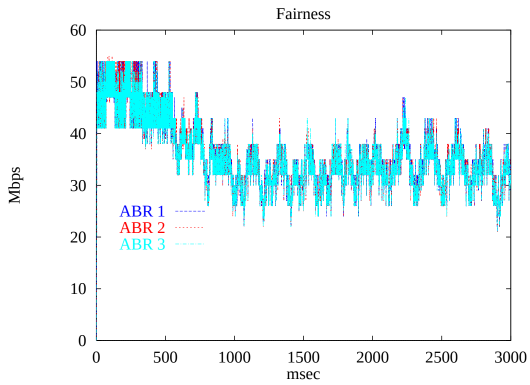

Figure 10: EFCI Fairness

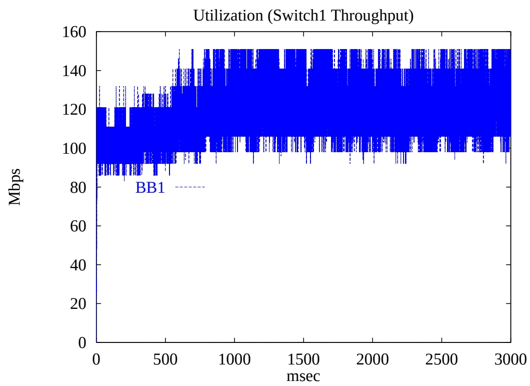

Figure 11: ER Utilization

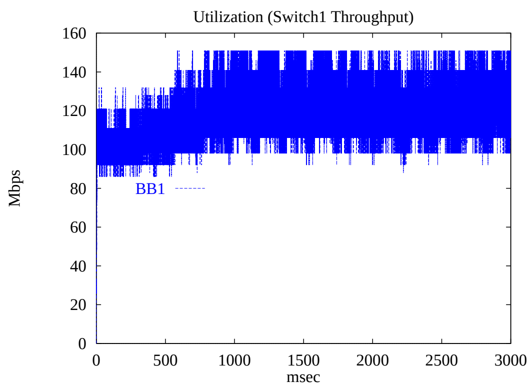

Figure 12: EFCI Utilization

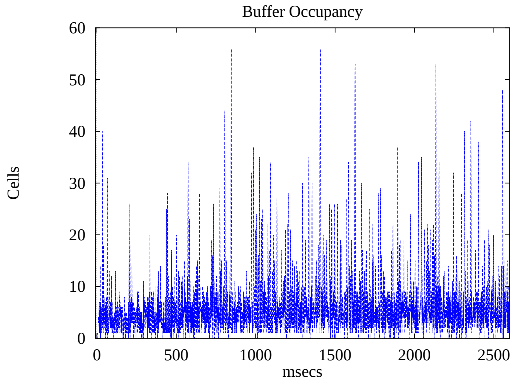

Figure 13: ER Buffer

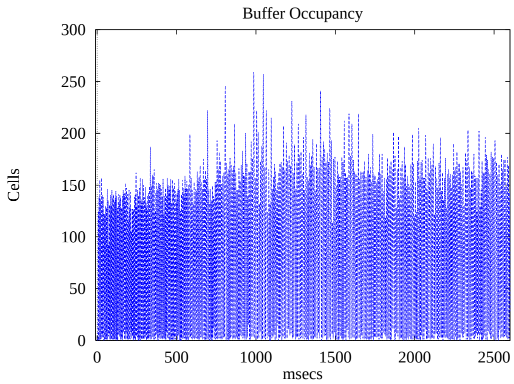

Figure 14: EFCI Buffer

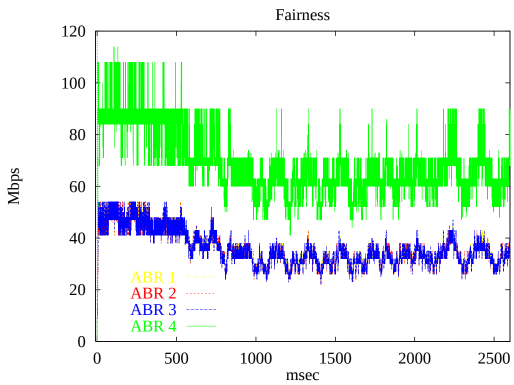

Figure 15: ER Fairness

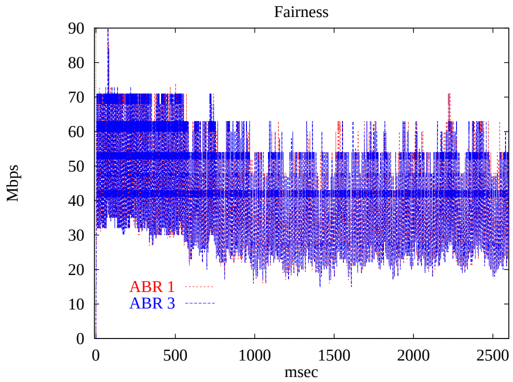

Figure 16: EFCI Fairness1

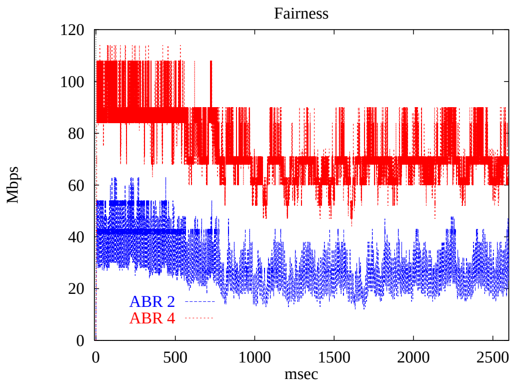

Figure 17: EFCI Fairness2

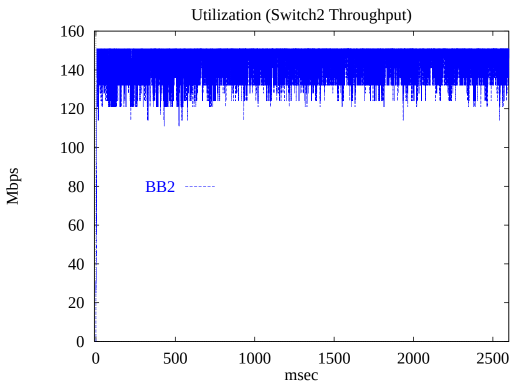

Figure 18: ER Utilization

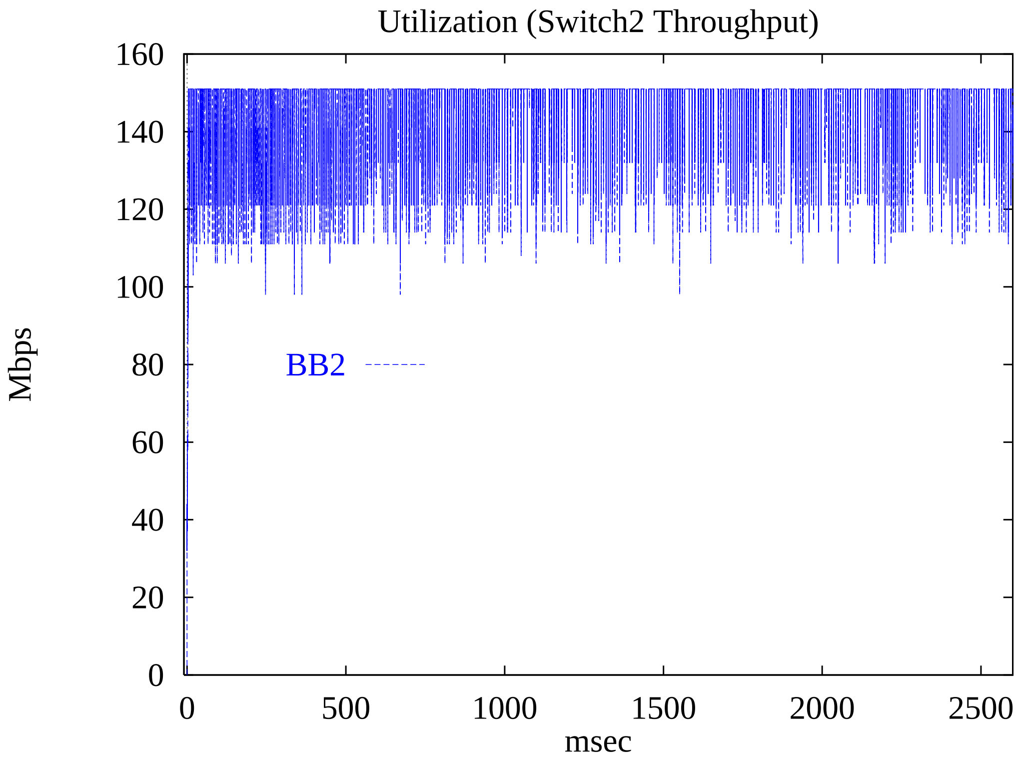

Figure 19: EFCI Utilization

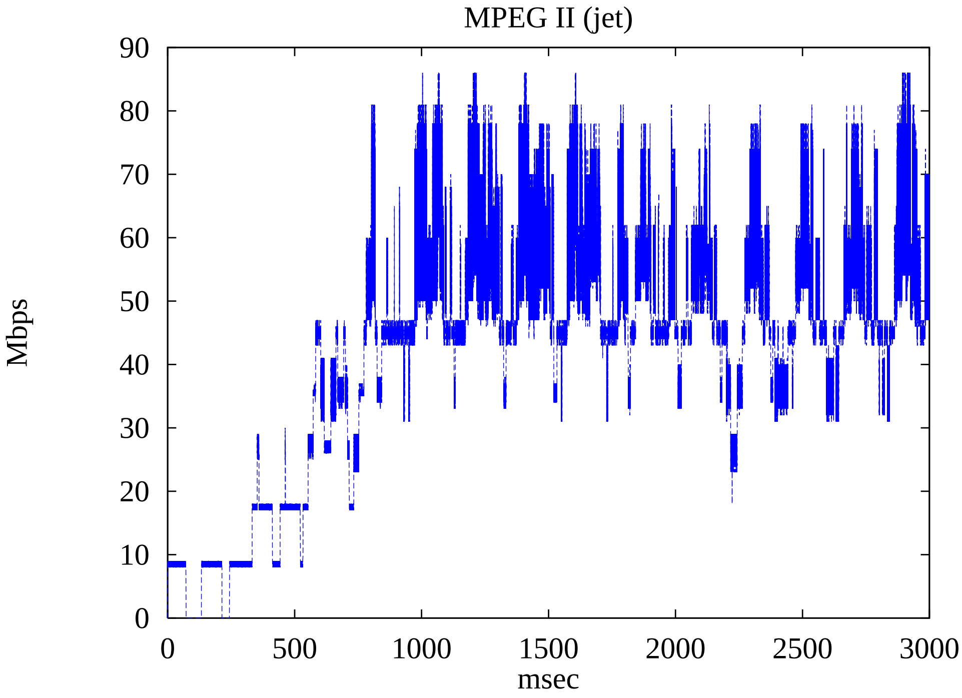

Figure 20: MPEG VBR Traffic

Parameter	Definition	Value
NRM	Number of RM cells	10
LCR	Line Cell Rate	149.76 Mbps
ACR	Allowed Cell Rate	Dynamically Updated
ICR	Initial Cell Rate	7.49 Mbps
PCR	Peak Cell Rate	149.76 Mbps
MCR	Minimum Cell Rate	0.149 Mbps
AIR	Additive Increase Rate	.212 Mbps
RDF	Multiplicative Decrease Factor	512
TOF	Time Out Factor	2
Xrm	X Forward RM cells before return RM cell	32
XDF	Xrm Decrease Factor	.0625
MAIR	MACR Additive Increase Rate	0.5 Mbps
MRF	MACR Reduction Factor	.95
TR	Target Rate	142 Mbps

Table 1: Tuning Parameters

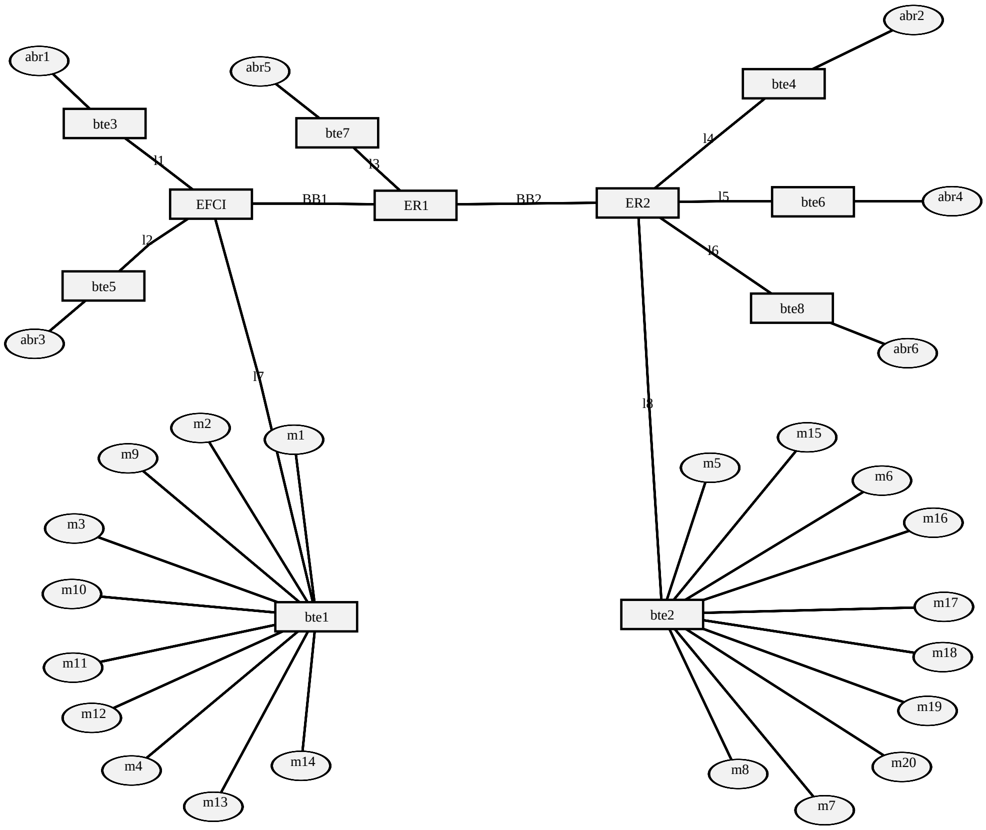

Figure 21: Parking Lot

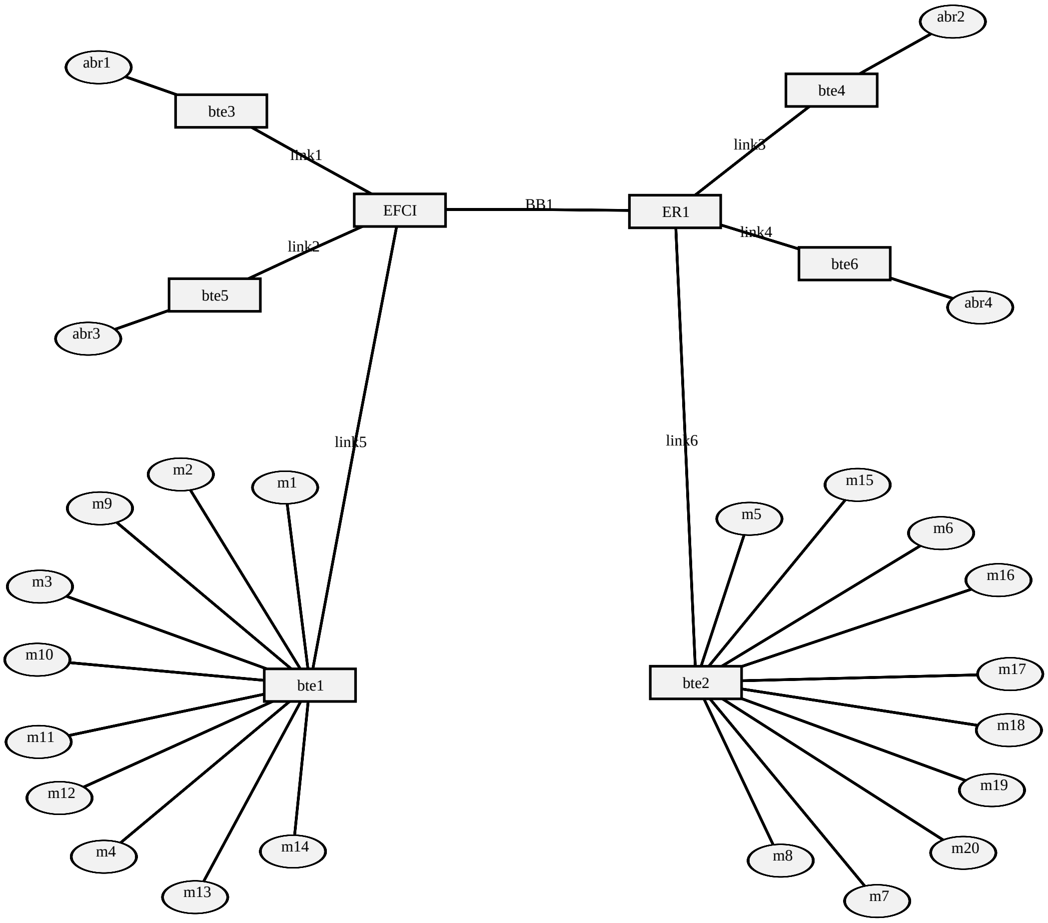

Figure 22: Fast Convergence

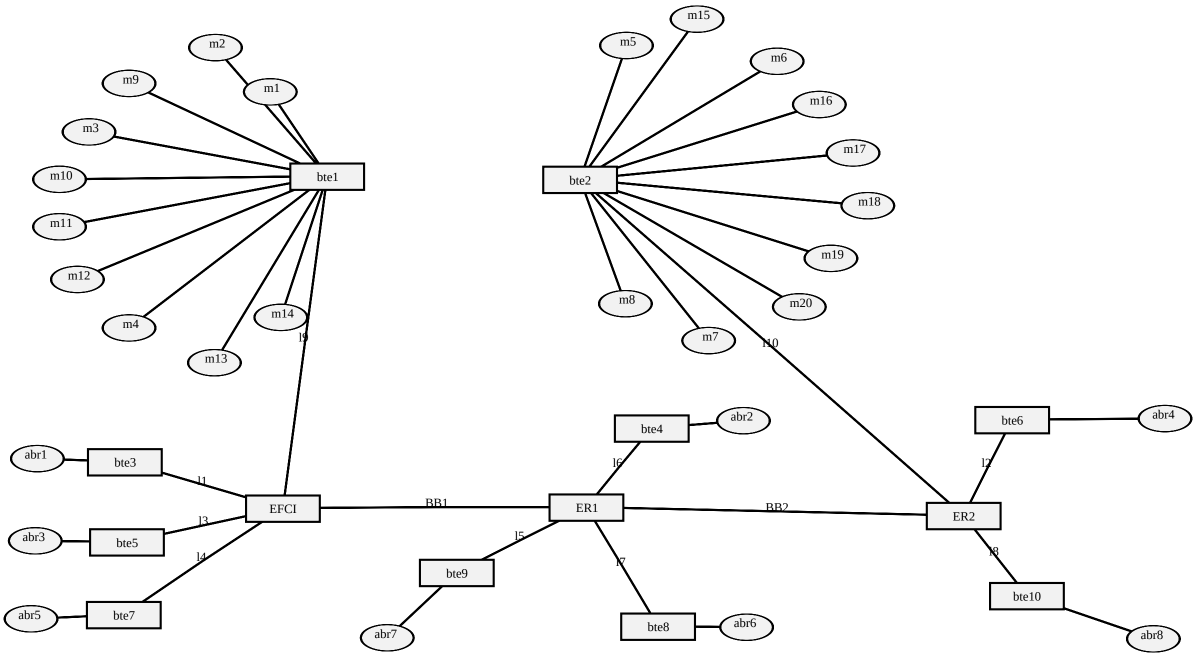

Figure 23: Upstream Bottleneck

References

[1] F. Bonomi and K. Fendick. The Rate-Based Flow Control Framework for the Available Bit Rate ATM Service. *IEEE Network*, April 1995.

[2] Y. Chang, N. Golmie, and D. Su. A Rate Based Flow Control Switch Design for ABR Service in an ATM Network. In *Proceedings of the Twelfth International Conference on Computer Communication*, August 1995.

[3] E. Dejesus. How the Internet will Replace Broadcasting. *Byte Magazine*, January 1996.

[4] ATM Forum. *ATM User-Network Interface Specification Version 3.1.* Prentice Hall, 1995.

[5] N. Golmie, A. Koenig, and D. Su. The NIST ATM Network Simulator. Technical Report NIST Internal Report 5703, NIST, August 1995.

[6] R. Jain. Congestion Control and Traffic Management in ATM Networks: Recent Advances and a Survey. *Submitted to Computer Networks and ISDN Systems*, February 1995.

[7] R. Jain, S. Kalyanaraman, and S. Viswanathan. The OSU Scheme for Congestion Avoidance in ATM Networks Using Explicit Rate Indication. In *Proceedings WATM'95 First Workshop on ATM Traffic Management*, December 1995.

[8] L. Roberts. Can ABR Service Replace VBR Service in ATM Networks. In *Proceedings of IEEE Computer Society International Conference (COMCON)*, March 1995.

[9] S. Sathaye. *Draft ATM Forum Traffic Management Specification Version 4.0 95-0013R5.* ATM Forum, June 1995.