Simulation and Analysis of Credit- and Rate-based Switch Interoperability in an ATM Network

A Thesis

Presented to the

Department of Computer Science

Brigham Young University

In Partial Fulfillment

of the Requirements for the Degree of

Master of Science

By

Charles R. Duncan

February 7, 1997

This thesis by Charles R. Duncan is accepted in its present form by the Department of Computer Science of Brigham Young University as satisfying the thesis requirement for the degree of Master of Science.

Mark J. Clement, Committee Chairman

Bryan S. Morse, Committee Member

Parris K. Egbert, Committee Member

Date

Scott N. Woodfield, Graduate Coordinator

Acknowledgments

I would like to express my deep appreciation to several people for their friendship and support throughout my college experience. Foremost, I would like to thank Dr. Mark Clement who not only encouraged me to pursue a Master's Degree when I expressed an interest but also supported me along the way by employing me as a research assistant. He provided me with direction to my research, especially in finding a topic that was not only interesting but useful to the industry. He also provided me with the opportunity to submit a paper containing my research to the 5th International Telecommunication Conference this Spring, which has been accepted.

I would also like to thank my family, especially my wife, for putting up with me while I spent most of my time researching at the school. The patience and consideration that she has shown me and my work has been beyond measure.

I would also like to thank Tasha Slade for her friendship and support, but most of all, her English skills. I don't know how I would have made it through my thesis, let alone five years of college, without her to proofread my papers and reports for my terrible spelling and grammar errors.

And last, but not least, I would like to thank all my professors and friends for an enjoyable college experience. Thank you.

Contents

List of Figures

List of Tables

1 Introduction

1.1 ATM

Asynchronous Transfer Mode (ATM) is slated to become the protocol on which the information superhighway will rely to transfer voice, data and video. The cell-oriented architecture of ATM makes it possible to send these different services over a single link while providing the required Quality of Service (QoS) needed by each (e.g., low cell loss, low jitter) [1].

A great deal of research has already gone into ATM networks, including numerous test-beds 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.

However, with the introduction of ATM comes many new challenging problems for network designers. One of these challenges lies in the efficient management of ATM traffic to allow the communication bandwidth of the network to be optimally utilized [10, 24, 11, 6]. This has proven particularly difficult because of the variety of services that ATM can offer. Because of this challenge, successful deployment of ATM networks will require a reliable method of flow control for Available Bit Rate (ABR) service [26].

ABR service is necessary for high utilization in ATM networks because it was designed to use any bandwidth unneeded by other services. Also, as ATM moves into the local area network (LAN) environment, ABR service is the best choice for LAN traffic. However, ATM is meeting some stiff competition as Fast Ethernet and FDDI gain more acceptance in the data communications industry. ATM has obvious advantages over each of these protocols, but as a newer technology, support for ABR traffic is limited.

1.2 Overview of Research

This research evaluates the performance of two very different types of flow control for ABR service: rate-based and credit-based. For the rate-based flow control, I used the Explicit Proportional Rate Control Algorithm (EPRCA, from here on referred to as ER), defined by the ATM Forum (the standards group developing and standardizing the ATM technology) and implemented by the National Institute of Science and Technology (NIST) in their ATM simulator [9]. The credit-based flow control algorithm used is the Quantum Flow Control (QFC) Specification version 2.0, published by the QFC Alliance [7].

These two different algorithms were studied running independently *and* running together in a realistic transient (dynamic) environment on both metropolis area networks (MANs) and wide area networks (WANs). Three different network configurations are used to show how these two algorithms respond to known network problems. The performance of the ABR flow control is evaluated in terms of throughput, buffer occupancy, fairness, and fast convergence.

Because of the complexity involved in analyzing flow control in ATM networks, the concepts used by the ATM Forum and others are based mostly on simulation. This study was performed using a modified version of the NIST ATM Network Simulator, which is capable of simulating both QFC and ER switches [9].

1.3 Thesis Statement

ATM switches using rate-based and credit-based ABR flow control will interoperate effectively in an ATM network.

1.4 Overview of Thesis Document

The ATM specification defines several different services based on different QoS needs. Section 2 gives an overview of these different services and provides the reader with a basic understanding and justifications of each service.

Because this research is based on the ABR service class and the flow control mechanism that it uses, Section 3 gives an in-depth look at this subject. Credit- and rate-based flow control are explained in detail, along with a little of the history that has shaped what the industry uses today.

The environment used to run the simulations is explained in Section 4. This includes an explanation of MPEG and why it is a candidate for good background traffic. Also, a description of the parameters used in the ER mechanism is provided for those users who wish to duplicate these findings.

Section 5 explains what work was performed as part of this thesis, the problems that arose, and the solutions found. Section 6 will give the reader a basic understanding of the different network configurations used in this research and why they were used. The assumptions behind each configuration are also explained, along with an analysis of what the simulations' results would be if the algorithms worked perfectly.

The results of each of the configurations are presented in Section 7, including conclusions and suggestions. Section 8 provides some ideas for future research in the flow control area.

2 ATM Overview

2.1 Basics of ATM

Asynchronous Transfer Mode (ATM) is a cell-based switching and multiplexing technology designed to be a general-purpose, connection-oriented transfer mode for a wide range of services. ATM is also being applied to the LAN environment and private network technologies, as specified by the ATM Forum.

The primary unit in ATM is the cell. The ATM standard defines the cell with a fixed-size length of 53 octets (bytes), comprised of a 5-octet header and a 48-octet payload. The fixed cell size simplifies the switching and multiplexing process and enables implementations of these at very high-speeds. The fixed cell size also eliminates the problems of short packets being delayed behind larger ones. This allows ATM to provide service to such things as voice and video, where large transmission time variation is unacceptable.

ATM handles both connection-oriented traffic and connectionless traffic. Connection-oriented traffic may be dealt with directly or through the use of adaptation layers, while connectionless traffic can be provided only with adaptation layers. Each ATM cell sent into the network contains addressing information that identifies a virtual connection from origination to destination. All cells are transferred, in sequence, over this virtual connection.

ATM provides either Permanent or Switched Virtual Connections (PVCs or SVCs). PVCs are set up for long periods of time, and remain up even if they are not in use. SVCs, on the other hand, are set up on a per-connection basis and are torn down as soon as they are finished being used. ATM is asynchronous because the transmitted cells need not be periodic. Synchronous Transfer Mode (STM) uses a fixed, periodic time slot for each virtual connection.

ATM offers the potential to standardize on one network architecture, defining the multiplexing and switching method. Synchronous Optical Network (SONET) and STM provide the basis for the physical transmission standard for very high speed rates. ATM also supports multiple QoS classes for differing application requirements on delay and loss performance.

The vision of ATM is that an entire network can be constructed using ATM Application Layers (AALs), and switching and multiplexing principles, to support a wide range of services. In this manner, ATM provides bandwidth-on-demand through the use of SVCs, and also supports LAN-like access to available bandwidth.

2.2 Service Classes

As mentioned above, one of the key attractions of ATM is the variety of services it offers. These services are what enable a network provider to offer voice, video, and data over a single line. Before ATM, two separate infrastructures were needed to provide voice and data; the idea of sending both of these services and video over a *single* infrastructure is very attractive. We can already see the benefits of such technology in the adoption of distant learning.

ATM services are based on a variety of desired communication attributes (see Table 1), such as Cell Loss Rate (CLR), Cell Transfer Delay (CTD), Cell Delay Variance (CDV), Sustainable Cell Rate (SCR), Peak Cell Rate (PCR), and Minimum Cell Rate

	QoS				
Attribute	CBR	VBR-RT	VBR-NRT	ABR	UBR
CLR	Specified			Specified	Unspecified
CTD	Specified			Unspecified	
CDV	Specified		Unspecified	Unspecified	
SCR	N. A.	Specified		N. A.	
PCR	Specified			Specified	
MCR	N. A.			Specified	N. A.
Controlled	No			Yes	No

Table 1: ATM Layer Service Categories

(MCR). By taking different combinations of these attributes, we have four basic service classes:

- Constant Bit Rate (CBR): Used for emulating circuit switching, where the required bandwidth is constant and known in advance (e.g., voice and television).
- Variable Bit Rate (VBR): Allows users to send at a variable rate that is known in advance and can be defined in terms of PCR, SCR, and MCR (e.g., video conferencing). VBR has two sub-categories: real-time (VBR-RT) and non-real-time (VBR-NRT) [8].
- Available Bit Rate (ABR): Designed for normal data traffic that cannot (or is hard to) be predicted in advance and is not time sensitive (e.g., file transfer and email).
- Unspecified Bit Rate (UBR): Designed for those data applications that want to use any leftover bandwidth and are not sensitive to cell loss or delay (e.g., file transfer and email).

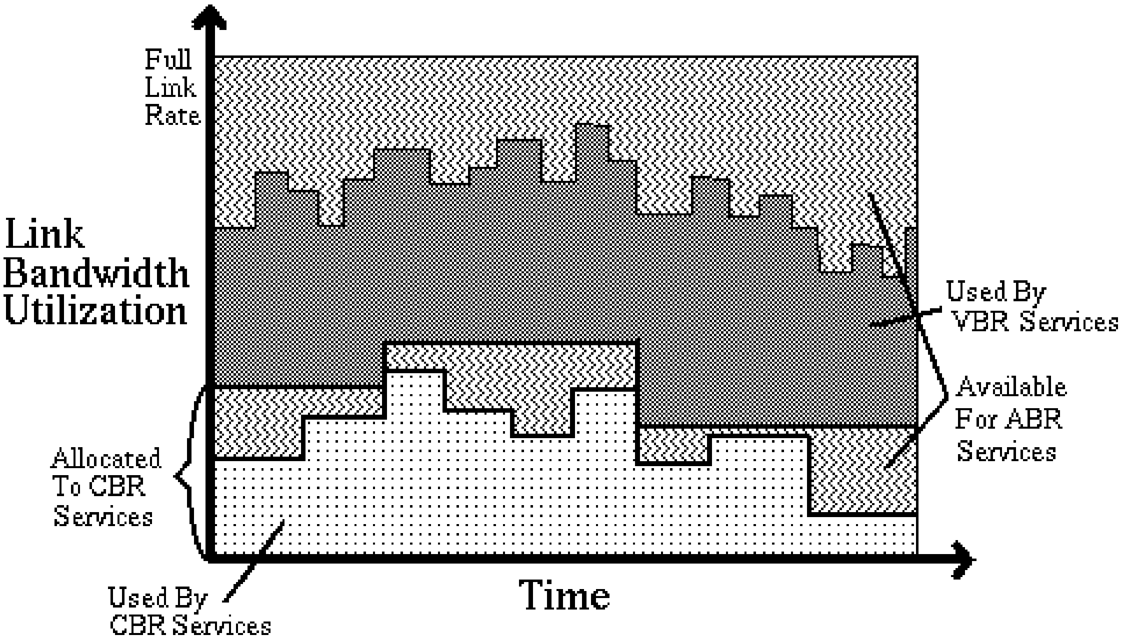

Figure 1: Service Allocations

2.3 How The Services Work Together

A given amount of bandwidth is guaranteed to CBR and VBR connections. Although the entire guaranteed bandwidth is not always used by the connection, the connection has access to all of its reserved bandwidth if necessary. ABR helps fill in this otherwise waisted bandwidth with regular data traffic (Figure 1). UBR was created to do approximately the same thing, except that it had no MCR and was more of a send-and-pray protocol. The only way for a sender to realize his data was not getting through was for the receiver to timeout after a given period of time, then to ask for a resend of a missing cell. This would all have to be done in higher protocols, making it much more inefficient.

ABR, on the other hand, has a flow control mechanism that controls the flow of data and provides not only an MCR but also a low CLR. This service class also emulates the common communications protocols for which many of the industry's applications were built. This makes ABR work well with the current LAN market, while still providing a good migration path to full ATM service.

The down side of ABR, however, is the flow control. The idea of Gigabit networks is still new, and it takes time to figure out how to best design flow control on a medium that runs at gigabits-per-second [5]. The subject of what type of flow control to use has been a highly debated subject in the ATM industry.

3 Flow Control

3.1 ABR Flow Control

Flow control has always been a religious issue in computer communications, but never as much as with ATM's ABR flow control. This is partly due to the fact that the current low-speed legacy LAN applications will co-exist with the future high speed data links for several years. Since ATM has been slated as the standard for supporting gigabit Broadband Integrated Services Digital Networks (B-ISDN), this mismatch of link speeds creates a very difficult problem for congestion control.

A very important part in the selection of a congestion control scheme is traffic patterns. In ATM networks, there are two main classes of traffic: guaranteed and best-effort. Guaranteed traffic requires the network to guarantee a set of conditions to the connection (such as the aforementioned CLR, CTD, CDV, etc.). As can be seen from Table 1, CBR and VBR traffic are guaranteed services. Congestion control for these services is administered through admission control and bandwidth allocation. In other words, at the time of call setup, if the network can provide the requested connection, it is admitted; otherwise, it is rejected.

The UNI 3.1 specification declares that the cell loss probability for CBR and VBR traffic is required to be less than 10^{-8} [6]. 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 the link's capacity [20]. Thus, if only CBR and VBR services existed, the link could never be fully utilized and still maintain the required cell loss ratios.

The best-effort traffic class has been the focus of much recent work by the ATM Forum. This service is not guaranteed to have a cell loss probability of 10^{-8}, but instead it either simply sends the data and hopes it arrives at its destination (UBR) or relies on a control system to maintain low cell loss (ABR).

Originally, only UBR existed, which was used as a no-guarantee protocol that would simply fill in the extra space on the link. However, recent efforts have moved ATM into the LAN, where applications were not designed to predict their own bandwidth requirements, but still needed some type of guarantee. Thus, ABR services were created with a flow control to keep cell loss low and also provide a minimum cell rate (MCR). This, combined with the bursty nature of LAN traffic, has complicated the issue of congestion control [15].

Many different schemes have been proposed to handle the congestion control problem. One of the simplest is the use of large buffers to hold the onslaught of cells during a period of congestion, and to allow a higher-layer protocol to compensate for the cell loss that could occur during long periods of congestion. However, this was shown to be inadequate, and resulted in poor performance [24, 25].

One scheme that has been adopted and is currently in use in many of today's ATM switches is the Explicit Forward Congestion Indicator (EFCI). This is a binary approach, where a congested switch sets a single bit in the ATM cell header that indicates congestion in the network. As the network gets congested, the ATM cell header congestion bit is used to tell the source to slow down its rate of transmission. But this scheme has also been categorized as too simple, and does not allow the switch enough control.

Currently, the industry has settled on one of two flow control schemes: rate-based or credit-based. A rate-based switch uses feedback from the network to govern the rate at which the source can send. A credit-based switch gives the source credits which allows the source to send cells — one cell per credit. Both of these explicitly control the amount of data that the switch will receive on a given link, but they differ significantly on how this is accomplished [14].

3.1.1 Rate-based Flow Control

Rate-based schemes use feedback information from the network to specify the maximum rate at which each source can send data through the network [2]. EFCI is an implementation of a rate-based scheme. When the network is congested, the congestion bit is marked in outgoing cells. The end connection that receives this cell will either slow down the transmission rate or notify the source to slow down.

This scheme, and most other rate-based schemes, use a closed end-to-end control loop (i.e., control information has to complete the entire loop from source to destination and back to the source). Although this scheme is rather easy to implement, adaptation to network conditions often takes longer than the period of congestion, making it ineffective. There are also the questions of how much the source should slow down its transmission rate and when it can speed back up again.

The current congestion control scheme being pushed by advocates of rate-based systems is the Explicit Proportional Rate Control Algorithm (EPRCA, or ER) [12]. In this algorithm, the network tells the source, through the use of resource management (RM) cells, the exact data rate at which it can send. This is very nice for the WAN, where ABR traffic is smoother and not as bursty as it is in the LAN. Also in the WAN, it is easier to bill for services when a provider can give a connection a specific rate at which to run [3]. Thus, most of the large telcos are advocates of the rate-based specification.

A number of papers have been written that describe this control mechanism in detail [4, 2, 12]. Since the Traffic Management 4.0 specification [21] only states what must be accomplished by an ER switch, not necessarily how it is to be accomplished, 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 [4].

One of the main downfalls of the ER scheme is its complexity. It costs twice as much to produce a Network Interface Card (NIC) to handle the ER scheme (as compared to a credit-based scheme) because of the rate control that needs to be added [17]. But again, there are no NICs in the WAN, only in the LAN. Even in the switch the complexity is high, but this is outweighed by the large number of ports on a WAN switch, and by the supposed reduced number of buffers required to run a rate-based scheme vs. a credit-based scheme [19].

3.1.2 Credit-based Flow Control

Credit-based schemes use a link-by-link window flow control approach [23]. In other words, the flow control is between two switches, with a single link between them, instead of between the source and the destination. Because this allows for a much smaller Round Trip Time (RTT), the network can respond to congestion much more quickly.

The most popular credit-based scheme is Quantum Flow Control (QFC). The way this and most other credit-based schemes work is that a source (meaning any switch in a connection's path that needs to send a cell to the next switch) can only send a cell if given permission (a credit) by the destination switch [13]. A switch sends credits to the upstream node, who then uses these credits to send cells downstream. This scheme can respond to congestion effectively because if a node gets congested, it does not send any credits upstream. As soon as the upstream node uses all of its credits, it will no longer be able to send any more cells.

Some of the benefits of this scheme are

- It adapts very well to extremely bursty traffic.
- It guarantees no cell loss due to congestion.
- It can guarantee full bandwidth use at all times.
- It is simple.

The biggest advantage with QFC is that it has overcome one of the biggest problems with most credit-based schemes: per Virtual Connection (VC) buffering. It accomplishes this by using a system of dynamic buffer allocation and shared buffers.

Per VC buffering will not work in a WAN switch since we can have hundreds of ports, with thousands of VCs per port, and links on the order of one hundred kilometers [19]. As an example, consider a WAN switch that has 100 ports, each with 1000 VCs, with link lengths of 1000 km each running at 150 Mbps — this switch would require 445 Mbytes of memory to handle just the flow control. One single 622 Mbps, 1000 km link with 5000 VCs would require 222 Mbytes. This amount of memory is simply too expensive for a single flow control mechanism.

QFC is attractive because it overcomes per VC buffering problems, while maintaining its no-cell-loss guarantee and full bandwidth usage. It accomplishes this by incorporating dynamic memory allocation, which is based on the idea that not all VCs can possibly use all the bandwidth of the link at the same time.

3.2 Mixed Flow Control Networks

For several years, there was a great debate over what type of flow control would work best for ATM (rate- or credit-based). Both of these explicitly control the amount of data that the source can send on a given link, but they differ significantly on how the control is accomplished [14]. In 1995, the ATM forum voted (narrowly) to endorse ER (a rate-based approach), and thus it seems rate-based algorithms had won.

But since that time, several very powerful credit-based implementations have been published [7], and many of the smaller LAN switch manufacturers are turning to credit-based schemes in their products due to its simple, inexpensive, and robust nature [7, 16, 13].

The telcos are sticking to ER for now due to the ease of billing and other factors. They also see it as a more secure future because it has been accepted by the ATM Forum as its ABR flow-control scheme. Thus, we can see a future of both types of flow control coexisting in a single network. This research examines how well these two schemes will interoperate.

4 Simulation Environment

4.1 The NIST Simulator

The National Institute of Science and Technology (NIST) developed an ATM simulator in mid-1994, and provided the simulator and the code to the public. The initial version provided basic ATM networking capabilities but had few ABR or flow control options. Version 2.0 was made available in 1995 and included EFCI and ER ABR flow control options.

The simulator allows one to set up a network configuration, including applications (traffic sources), end stations, switches, and links. Applications include TCP/IP, MPEG, CBR, VBR, and ABR data sources, all of which are completely configurable. All the components can be configured to sample network status at different time periods or at specific events.

Flow patterns can be defined from application sources through the network to their destinations, and every component can be configured as to what type of data should be logged during a simulation. This data can be throughput, link utilization, buffer usage, cell loss, cells received/transmitted, etc. The simulations can then be executed and the data analyzed.

I added the QFC flow control mechanism following the QFC specification version 2.0, and modified the simulator to allow more than one flow control in a single configuration. Others in the networking group at BYU added MPEG support, which I've used to provide a dynamic environment in which to run the simulations.

4.2 MPEG Video Compression

MPEG compressed video is quickly becoming a 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. MPEG can send at a rate of 30 frames per second and supports CD quality sound.

Motion compensation allows data to be shared between frames, which results in higher compression ratios for MPEG encoded video. MPEG uses discrete cosine transform algorithms that can decompose data sets into spatial frequencies. The less important (higher-frequency) coefficients are more coarsely quantized, 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 that is capable of carrying and maintaining synchronization for multiple video and audio programs.

For the simulations, traces from MPEG-2 transport streams (which were used to carry MPEG-1 encoded video sequences) were used to generate VBR traffic. The video is a simulation of a jet flying through a canyon which presents a high-action (worst case) environment for the flow-control mechanism to perform.

4.3 Providing a Dynamic Environment

VBR background traffic was generated in each of the configurations by 10 MPEG sources. VBR traffic generators ramp up from 1 to 10 sources during the first second of simulation. The sources start at intervals of approximately 100 msecs. The resulting VBR traffic pattern is shown in Figure 2. The periodic burstiness exhibited in the figure would most likely be the result of key encoded sequences (I-frames) being transmitted.

It is interesting to note that the MPEG traffic pattern is not a Poisson distribution. Other research has shown that traditional methods of mathematical analysis, using Markov processes, are often not applicable to actual traffic in ATM networks [22].

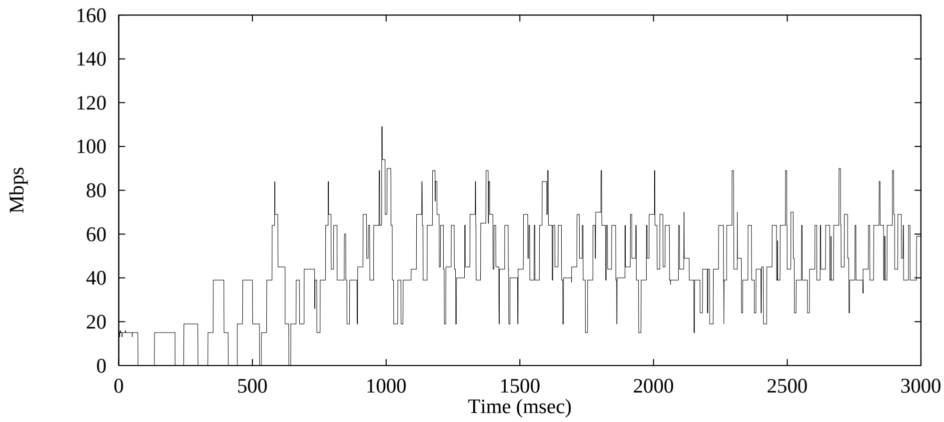

Figure 2: MPEG VBR Background Traffic

5 Work Performed

For this thesis, I first began studying ATM ABR flow control, both rate- and credit-based, and trying to understand the ER implementation in the NIST simulator. I then obtained the QFC standards document and began studying it. After I felt comfortable with both the QFC document and NIST's ER implementation, I started adding QFC to the ATM simulator.

After evaluating the necessary parts of the QFC specitication that would be required for the ATM simulator, I created two new menu-options in the ATM switch to allow the user to choose QFC or a QFC/ER combination as the ABR flow control for that switch. I added the necessary code to both the BTE and switch components to allow the user to easily switch between flow control mechanisms.

I ran into some problems with fundamental differences between the way ER and QFC buffer cells and the way the NIST simulator implemented input buffering in the BTEs. I thus changed the BTE's input buffering mechanisms to test the ABR flow control at simulation run time to see which mechanism was being used and to adapt the buffering technique appropriately. After running several complex simulations to test my additions, the ATM simulator was ready for use in proving my thesis statement.

After running simulations on several configurations and in several different environments, I found that the QFC implementation needed to be changed to watch for ER RM cells and to modify them appropriately according to the switch's congestion status. This is explained in more detail in Section 7.2.

After adding these changes and testing them, I again ran my simulations as described in the next section. I gathered the data generated from these simulations and used it to analyze my results. These results are explained in Section 7.

6 Configurations and Assumptions

6.1 Basic Assumptions

It was determined by the ATM Forum that greedy, persistent sources present the worst-case application for flow-control mechanisms [26]. ABR sources in all simulations are therefore greedy, persistent sources. Simulations were run in both MAN and WAN configurations, with backbone links of 20 km and 1000 km, respectively. All other links are 0.1 km. All links have a capacity of 150 Mbps. Switch buffers are set to 3000 per link, and the broadband terminal equipment (BTEs) has 500 buffers per link. The ER control mechanism was configured as shown in Tables 2 and 3 in each switch and BTE, respectively.

Parameter	Definition	Value
MIR	MACR Additive Increase Rate	0.5 MB/s
TR	Target Rate	142 Mbps
ERF	Explicit Reduction Factor	0.95
N	Measurement Interval	10 cells
TAU	Congestion tolerance	0 cells
IMACR	Initial MACR	142 Mbps
VCS	VC Separator	0.875
AV	Average Factor	0.0625
MRF	MACR Reduction Factor	.95
DPF	Down Pressure Factor	0.875

Table 2: ER Tuning Parameters for the Switches

Parameter	Definition	Value
ACR	Allowed Cell Rate	Dynamically Updated
PCR	Peak Cell Rate	149.76 Mbps
NRM	Number of RM cells	10
ICR	Initial Cell Rate	7.49 Mbps
MCR	Minimum Cell Rate	0.149 Mbps
AIR	Additive Increase Rate	0.212 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	0.0625
MRM	Controls bandwidth for RM and data cells	2
TRM	Max Time between forward RM cells	100000
TDF	ACR Time Decrease Factor	256

Table 3: ER Tuning Parameters for the BTEs

6.2 Configurations

Three network configurations have been developed by researchers at Ohio State University to evaluate specific characteristics of the control mechanism [12]: Fast Convergence, Parking Lot, and Upstream Bottleneck (Figures 3, 4, and 5 beginning on page 21). The configurations have been modified to include VBR traffic generators and to allow each individual switch's flow control to be changed between credit- and rate-based schemes.

In the figures, each configuration is labeled for easy understanding of flow patterns: each source and receiver are named the same except for the last letter being either an 's' or an 'r' (source and receiver, respectively). These are shown as ovals in the network configurations and generate either ABR or MPEG (VBR) traffic. The BTE simulates the computer that the source/receiver is running on. These names also follow the same rules as explained above.

Four tests were run on each configuration, which are referred to as ER, QFC, ER/QFC, and QFC/ER (named for the positioning of the different flow control mechanisms). In the ER simulations, all the switches run the ER flow control mechanism. Likewise, QFC simulations use all QFC switches. In the ER/QFC simulations, ER runs on the edge switch and QFC runs in the MAN/WAN. Likewise, QFC/ER is QFC at the edge and ER in the MAN/WAN.

For each configuration, the results of the simulations are shown in a table. The data presented is average utilization of each of the sending links (L1s, L2s, L3s, and L4s), average utilization of each of the backbones (BB1 and BB2), average queue occupancy of the significant switches (Q1 and Q2), and peak queue occupancy of significant switches (Q1 Pk and Q2 Pk). Only those components that are significant, or appear in the configuration, will be presented.

6.2.1 Fast Convergence Configuration

The "Fast Convergence" configuration shown in Figure 3 tests for the time required for ABR sources to converge to the fair allocation of the available bandwidth. It consists of two ABR sources, one of which begins transmitting at time t = 0 msecs. The second ABR source begins transmitting at t = 1000 msecs and continues transmitting until t = 2000 msecs. During this time period, the MPEG source is transmitting at an average rate of 52.4 Mbps, which would leave a fair share of 48.8 Mbps for each ABR source.

As the second ABR source begins transmitting at t = 1000 msecs, the first source should slow down quickly enough to avoid filling up or overflowing the switch's buffers. When the second source stops transmitting, the first source should ramp up quickly

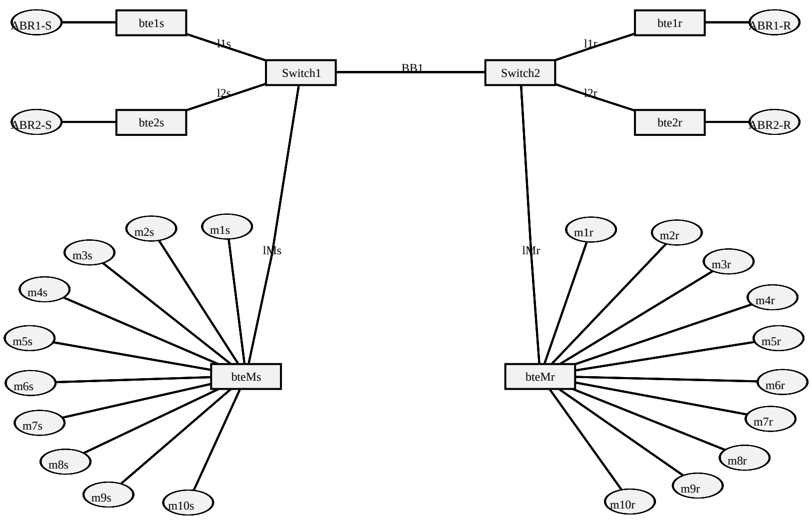

Figure 3: Fast Convergence Configuration

enough to keep the links fully utilized.

The results of these simulations are found in Tables 4 and 5 in Section 7.

6.2.2 Parking Lot Configuration

The "Parking Lot" configuration, shown in Figure 4, was devised to ensure that all sources receive a fair share of the bandwidth, even though some may be routed through more switches than others, each of which may be using a different flow control mechanism. Statistics were recorded starting at 1000 msecs (after the MPEG VBR sources finished ramping up) and lasted for 2000 msecs. The average MPEG data rate was 51 Mbps during this time period, which leaves a fair share of 33 Mbps for each of the three sources ((150-51)/3). Note that because each source should get equal bandwidth allocation, BB1 should not be fully utilized.

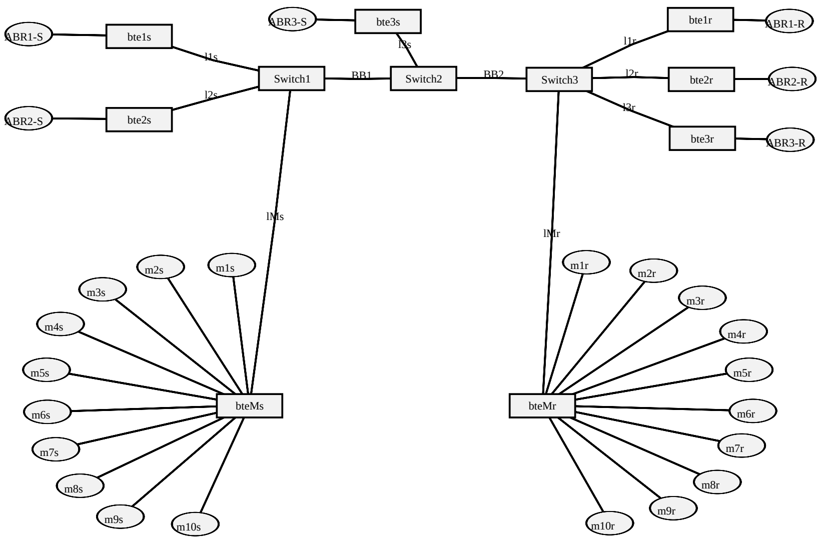

Figure 4: Parking Lot Configuration

The results of these simulations are found in Tables 6 and 7 in Section 7.

6.2.3 Upstream Bottleneck Configuration

One of the criteria 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 5), consists of three ABR sources at the edge switch and a fourth ABR source joining in at Switch2. ABR1 and ABR3 drop out at Switch2 causing ABR1 and ABR3 to share BB1, and ABR2 and ABR4 to share BB2.

Since the MPEG sources are averaging 51 Mbps, each backbone is left with 99 Mbps. The test is to see if the control mechanism will reallocate the bandwidth on BB2, enabling ABR4 to receive 66 Mbps while ABR2 remains at 33 Mbps. ABR1, ABR2, and ABR3 all share BB1, thus each should receive 33 Mbps. Since only ABR2 and ABR4

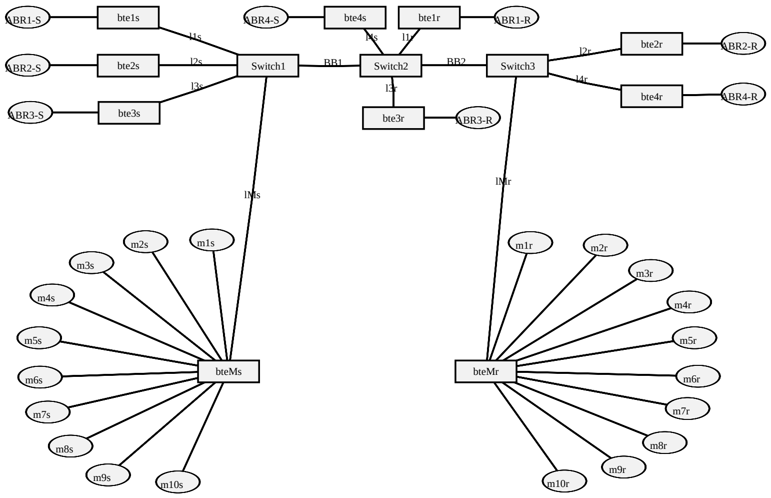

Figure 5: Upstream Bottleneck Configuration

occupy BB2, and ABR2 is restrained to 33 Mbps (by Switch1), ABR4 should receive the 66 Mbps left over.

The results of these simulations are found in Tables 8 and 9 in Section 7.

7 Results and Conclusions

7.1 Fast Convergence

Table 4 gives a summary of average rates and queue occupancy of the four simulations run in a WAN environment. The results were obtained during the one second time period from 1000-2000 msecs. As can be seen from the table, all of the simulations were quite close to operating optimally with no major unfairness surfacing, though some of the results were surprising.

name	L1s	L2s	LMs	BB1	Q1	Q1 Pk
ER	48.4	47.4	52.4	149.2	82.9	1442
QFC	48.7	48.7	52.4	150	37	135
ER/QFC	48.5	47.5	52.4	149.4	86.9	1390
QFC/ER	48.8	48.4	52.4	150	52	248
Optimal	48.8	48.8	52.4	150		

Table 4: Results of the Fast Convergence Configuration in a WAN Environment

When QFC is in control of the edge switch, the backbone remains completely full (150 Mbps), although ER and ER/QFC averages only slightly lower (149.2 and 149.4, respectively). QFC was able to do this while keeping the average and peak queue occupancy lower than ER. QFC's average queue occupancy of 37 is really not too much less than 82.9 for ER, but QFC's peak of 135 compared to ER's 1442 is a significant difference. ER/QFC and QFC/ER also follow this same pattern.

The higher average queue occupancy for ER is surprising since QFC generally uses more buffers to fulfill its full bandwidth guarantee. The lower peaks are due to QFC being able to respond to congestion quicker than ER. Although the difference between ER and QFC bandwidth usage (convergence) is not significant, ER does maintain a much higher peak buffer occupancy in this simulation. This will cause cell loss and jitter and will increase with heavier congestion or burstier background traffic.

Figures 6 and 7 show, in graphic form, how much faster QFC is able to converge than ER. QFC's backbone remains completely full, while ER is unable to converge quickly enough to keep its backbone fully utilized.

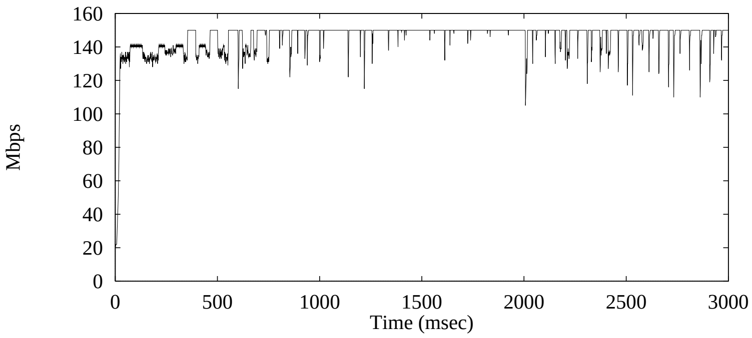

Figure 6: Fast Convergence BB1 Utilization for ER

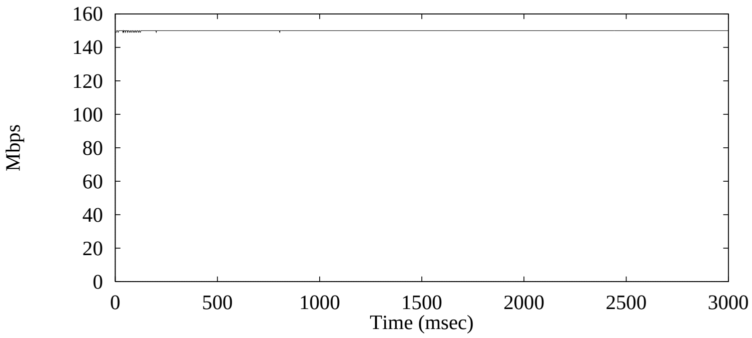

Figure 7: Fast Convergence BB1 Utilization for QFC

This can also be seen from how square the left edge of the graph is for QFC versus ER's slightly slower response. The table shows that the average data rates of these two schemes differ by less than a megabyte, but from the graph we can see that we are losing more bandwidth than the numbers would indicate.

Table 5 provides a summary of average rates and queue occupancy of the four simulations run in a MAN environment. The results were taken during the same time period as above. Note that the results are very close to the WAN results: all the simulations were quite close to operating optimally with no major unfairness resulting. It is surprising to see that some of the average and peak buffer occupancies actually increase slightly. Also, QFC is still able to keep the average and peak buffer occupancy lower than ER.

name	L1s	L2s	LMs	BB1	Q1	Q1 Pk
ER	48.5	47.6	52.4	149.5	72.3	946
QFC	48.7	48.7	52.4	150	37.1	135
ER/QFC	48.5	47.5	52.4	149.4	88.4	1399
QFC/ER	48.7	48.5	52.4	150	53	296
Optimal	48.8	48.8	52.4	150		

Table 5: Results of the Fast Convergence Configuration in a MAN Environment

Note that the peak buffer occupancy for ER is 946 and the peak for ER/QFC went up to 1399. This is because the QFC mechanism operating on Switch2 responded to congestion faster than Switch1. Switch2 would slow down Switch1's data rate sooner than Switch1 was able to tell the source to slow down. This is because ER is an end-to-end mechanism and takes longer to get information back to the source, thus causing the buffers to peak higher in Switch1 than they did in a straight ER environment.

This could cause problems in a situation where the telephone companies or major backbone carriers decided to switch to QFC (or any link-by-link control mechanism), and the subscribers to this service stick to ER (or an end-to-end control mechanism). The subscribers' edge switches to the MAN would experience abnormal burstiness in their buffer occupancy. This will cause problems any time the MAN becomes congested, and could not be controlled by the subscriber. In other words, there wouldn't be much the subscribers could do about the problem except increase buffer space on the edge switch, or change to a link-by-link control mechanism.

7.2 Parking Lot

Table 6 gives a summary of average rates and queue occupancy of the four simulations run in a WAN environment. The results were taken during the two-second time period from 1000-3000 msecs. As shown in the summary table, all simulations performed well except ER/QFC.

name	L1s	L2s	LMs	BB1	L3s	BB2	Q1	Q1 Pk	Q2	Q2 Pk
ER	31.7	31.7	51	114.9	31.7	147.3	2.4	162	283.7	2489
QFC	32.7	32.8	51	116.9	32.8	150	49.4	193	1435	2059
ER/QFC	48.9	48.9	51	149.8	1	150	81.5	1879	2417	3000
ER/QFCM	32.6	32.8	51	116.7	32.9	150	58	178	438.2	1925
QFC/ER	31.7	31.7	51	115	31.7	147.2	2.3	153	268.6	2822
Optimal	33	33	51	117	33	150				

Table 6: Results of the Parking Lot Configuration in a WAN Environment

In our implementation, the QFC switch passes ER RM cells on without modification; thus, Switch2 cannot slow down upstream sources. When our QFC implementation is modified to look like an ER switch to ER RM cells (but still use QFC on its links), the unfairness disappears. This is shown as 'ER/QFCM' ('M' for Modified) in the table.

As can be seen from the table, when ER controls the edge switch, the average buffer occupancy is lower, but with a much higher peak. This is seen even better in Figures 8(a) – 8(b). In the ER case, even though ER is keeping the average buffer occupancy low, it still has very high peaks. Going from low to high occupancy in short periods of time can cause a lot of variance in the amount of time it takes a cell to traverse through the switch; in other words, it can cause jitter.

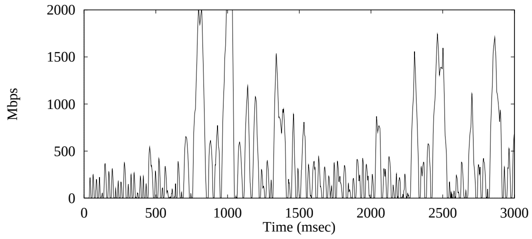

(a) ER: Switch2 Queue Occupancy

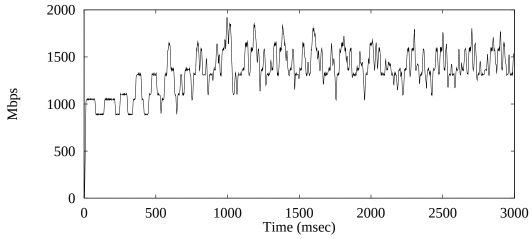

(b) QFC: Switch2 Queue Occupancy

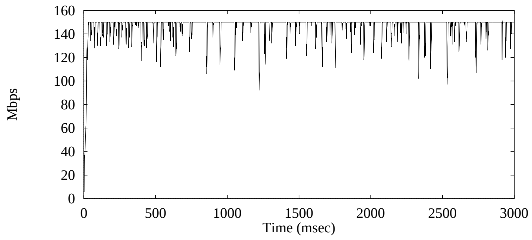

(c) ER: BB2 Utilization

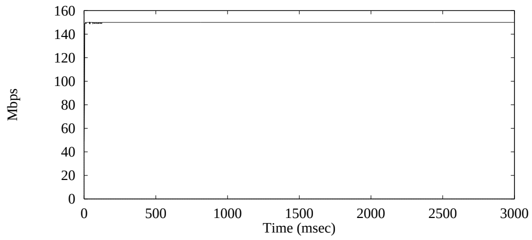

(d) QFC: BB2 Utilization

Figure 8: Parking Lot Graphs

In the QFC case, the average occupancy is much higher, causing larger average delays than ER, but it is more stable (much less jitter). From Figures 8(c) – 8(d), the difference between ER's and QFC's backbone usage is clear. Even though both ER and QFC are doing well at keeping the backbone full, ER is unable to converge fast enough to keep it 100% full, 100% of the time.

Both of these results can be found in all three of the configurations. As link distances decrease (to around 10 km or less) or the burstiness of the background traffic drops, the ER buffers stabilize. But if either of these parameters increase, ER begins to lose cells and does a poor job of keeping the backbone full. This problem is magnified when the ABR sources use traffic with data dependencies (TCP, MPEG, etc.) rather than the greedy, persistent sources. This causes retransmissions and acknowledge (ACK) wait times that cause the utilization to drop.

Table 7 provides a summary of average rates and queue occupancy of the four simulations run in a MAN environment. The results were taken during the two second time period from 1000-3000 msecs. In this set of simulations, only the modified ER/QFC is used.

The MAN environment shows increased throughput with ER and a decrease in

name	L1s	L2s	LMs	BB1	L3s	BB2	Q1	Q1 Pk	Q2	Q2 Pk
ER	32.2	32.1	51	115.9	32.3	148.8	.8	51	51.3	1531
QFC	32.8	32.8	51	116.9	32.8	150	45.8	214	129.7	226
ER/QFCM	32.7	32.8	51	116.8	32.7	150	46	98	284.1	1523
QFC/ER	32.1	32.1	51	115.7	32.1	148.4	1.1	31	84.7	1462
Optimal	33	33	51	117	33	150				

Table 7: Results of the Parking Lot Configuration in a MAN Environment

average and peak buffer occupancy by all. Yet, buffer occupancy does not drop as drastically as expected. Figures 8(a) – 8(d) look nearly identical to those for the MAN environment. ER maintains a lower average buffer occupancy, but QFC has much lower peak buffer usage.

In both the WAN and the MAN environment, ER shows the signs of causing problems with jitter and cell loss. Telephone companies and major backbone carriers should take note of this since they will be the ones affected most by this problem. WAN or MAN switches that connect LAN traffic to the backbone will experience low average buffer occupancy, yet will be prone to cell loss because of the high peaks in buffer usage.

7.3 Upstream Bottleneck

Table 8 gives a summary of average rates and queue occupancy of the four simulations run in a WAN environment. Results were taken during the two second time period from 1000-3000 msecs. As shown in the summary table, only QFC operates optimally, with the ER configuration giving the worst results. ER is not able to maintain high utilization of the backbone (BB2). This might be resolved by setting ER to be more aggressive, although this would increase the chances of cell loss.

It is difficult to try to find an optimal setting for ER in this kind of situation since the mechanism has so many configurable parameters. Unless the mechanism is well understood, finding optimal settings for ER is a hit-and-miss situation and should be avoided. In this simulation, the ER settings shown in Tables 2 and 3 are used, and modification to this did not result in any significant increase in performance without causing problems in other areas (such as peak buffer occupancy or cell loss).

name	L1s	L2s	L3s	LMs	BB1	L4s	BB2	Q1	Q1 Pk	Q2	Q2 Pk
ER	32.7	32.2	32.7	51	149	64.5	148.2	82.4	1195	17.1	134
QFC	32.9	32.9	32.9	51	150	65.8	150	71.8	313	14.1	88
ER/QFC	32.5	32.5	32.5	51	149.5	66	150	89.6	1807	17.9	144
QFC/ER	32.9	32.9	32.9	51	150	65.3	149.8	112.8	657	42.6	543
Optimal	33	33	33	51	150	66	150				

Table 8: Results of the Upstream Bottleneck Configuration in a WAN Environment

ER's average of 148.2 Mbps on BB2 could be argued as insignificant since it is quite close to optimal. But again, graphically compare ER's average of 148.2 Mbps in Figure 9 to QFC's average of 150 Mbps in Figure 10. It becomes apparent that we really are losing a lot of bandwidth.

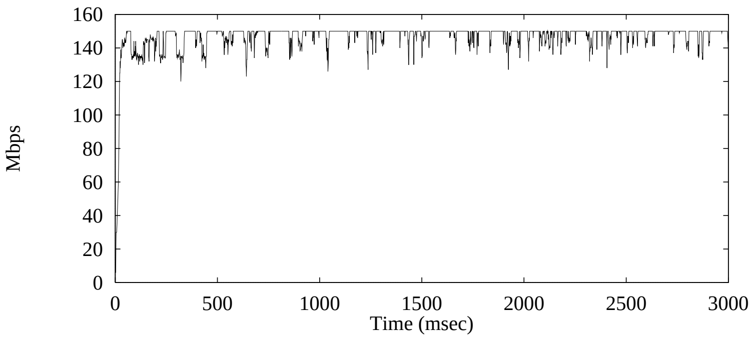

Figure 9: Upstream Bottleneck BB2 Utilization for ER

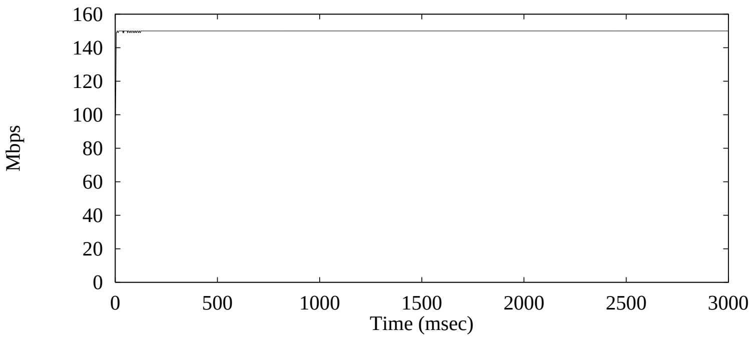

Figure 10: Upstream Bottleneck BB2 Utilization for QFC

Table 9 provides a summary of average rates and queue occupancy of the four simulations run in a MAN environment. Results were taken during the two second time period from 1000-3000 msecs. Again, it is surprising to see how similar the results are between the WAN and the MAN simulations.

The only significant change is the surprising increase in the peak buffer occupancy of the ER controlled switches. The average buffer occupancy increased 3.1% and 5.8% for Switch1 and Switch2, respectively, and the peak buffer occupancy increased 25% and 12% for the same. This is a very unexpected result, to have the peak buffer occupancy increase when the link distances decrease.

name	L1s	L2s	L3s	LMs	BB1	L4s	BB2	Q1	Q1 Pk	Q2	Q2 Pk
ER	32.6	32.4	32.6	51	149.5	64	148.2	85	1592	16.1	152
QFC	32.9	33	32.9	51	150	65.8	150	71	313	14	95
ER/QFC	32.5	32.5	32.5	51	149.8	65.6	150	104.9	2428	17.8	145
QFC/ER	32.9	32.9	32.9	51	150	65.4	149.9	113.1	566	45.5	536
Optimal	33	33	33	51	150	66	150				

Table 9: Results of the Upstream Bottleneck Configuration in a MAN Environment

When QFC is in control of the edge switch, average and peak buffer occupancy remain the same or decrease when we move to the MAN environment. Yet, when ER is in control, they actually increase. This could be a result of the congestion period being the same amount of time, on average, as the RTT, or a multiple of it. If this is the case, the flow control information the source is receiving would be opposite of what it should be doing.

In other words, the time at which the source finally receives the information to slow down because of congestion, the congestion would be clearing. Thus, the source would slow down right when it should be speeding up. The opposite would also be true — the source would be ramping up right when congestion would be starting. This, of course, should be a rare occurrence, but even a small period of this type of oscillation would push buffer occupancy up very quickly.

Since ER is an end-to-end flow control, there is no real way around this problem. This analysis is supported by the fact that QFC does not exhibit this kind of problem. Since QFC is a link-by-link mechanism, it responds more quickly and accurately to network congestion.

7.4 Summary

We find that when all switches are using either QFC or ER mechanisms, the control system performs reasonably well in all three test configurations. QFC obviously does a better job of utilizing the backbone's full bandwidth (with fast convergence), and surprisingly keeps the buffer occupancy more stable, and sometimes lower.

ER usually keeps the average buffer occupancy lower, but has higher peak occupancy, which can cause jitter. Since those services that require low delay and jitter use CBR or VBR service classes, ABR service does not value them as highly as fast convergence and bandwidth utilization. Nevertheless, this could change as audio and video becomes more prevalent on our communication infrastructure.

Since ER is an end-to-end flow control, it can fail when the congestion period is the same length, on average, as the RTT (or a multiple of it). Although rare, this can cause a very large peak in buffer occupancy in a short period of time, which could result in cell loss.

In the case where both ER and QFC are operating, there is no significant degradation in performance. The fact that QFC switches cannot ignore ER RM cells is important to note. The switch must look like an ER switch to all the sources using the ER mechanism; this allows it to have some control over these sources.

When QFC is in control of the edge switch (with ER in the WAN), a nice combining effect is found. Usually, this would result in lower average and peak buffer occupancy than ER alone, and would still keep the backbones running at near optimal.

When this situation is reversed (QFC is in the WAN and ER is on the edge switch), the edge switch will have problems with frequent large peaks in buffer occupancy and an overall increase in average buffer usage. This is due to the fact that ER is an end-to-end control mechanism and QFC is controlled link-by-link; thus, QFC is able to slow its control loops down quicker than ER can slow the source down.

Other than this, the two mechanisms operate well together, only occasionally resulting in larger average and peak buffer occupancy. Seldom did we find that the combination of these two algorithms gave worse results than we found with ER alone.

To summarize each configuration:

- The Fast Convergence configuration showed

  1. Full bandwidth usage (faster convergence) and lower average and peak buffer usage when QFC was on the edge switch
  2. All performed near, or at, optimal in fairness and fast convergence
  3. When ER was on the edge switch and QFC in the MAN, the edge switch experienced increased peak buffer occupancy due to the differences in the two

mechansims control loop size

- The Parking Lot configuration showed

  1. All simulations ran on par with one another (using the modified QFC switch)
  2. ER had lower average but higher peak buffer occupancy, which would cause more jitter
  3. QFC switches must be able to handle ER RM cells or unfairness will result

- The Upstream Bottleneck configuration showed

  1. Only QFC was able to keep the backbone full
  2. ER showed increased average and peak buffer occupancy when moving from the WAN to the MAN

As the burstiness of our MPEG traffic increases, the ER mechanism begins overflowing its buffers and losing cells. If TCP was used as our ABR traffic, much lower utilization by ER would result, because of the retransmission and acknowledgment wait period used by the TCP mechanism. Since QFC guarantees no cell loss due to congestion, it never had this problem.

It is also important to note that when link distances are around 10 km or less, ER buffers stabilize and QFC has higher average and peak buffer occupancy than ER. Thus, in the LAN, ER would normally maintain a lower average and peak buffer occupancy, but it would never be able to match QFC's full bandwidth and zero cell loss guarantee.

When deciding what mechanism would be best for the LAN, consideration of what is needed, full bandwidth and zero cell loss guarantee versus lower average and peak switch buffer occupancy, is a must. Additionally, the size of control loop (end-to-end versus link-by-link) provided in the WAN/MAN being connected to must be considered.

In the WAN/MAN, QFC will provide the most optimal flow control in terms of average and peak buffer occupancy and backbone utilization. There is also no negative impact to the WAN/MAN when subscribers, using a rate-based flow control mechanism in their LAN, connect to the backbone.

When using ER in the WAN/MAN, the backbone will experience problems with jitter and cell loss. There is no negative impact on the WAN/MAN when QFC subscribers hook-up to ER backbones. Often, QFC in the LANs is able to decrease the percent of jitter and cell loss in an ER WAN/MAN environment because of its ability to slow down quickly when the backbone becomes congested.

8 Related and Future Research

8.1 Related Research

Related research in flow control is being done by Nyirenda Towela from Kansas University under the direction of Bill Edwards of Sprint. Towela is studying the effects of the flow control mechanisms on TCP/IP data flows. She is using Quantum Flow Control extensively in her research and believes that QFC has no negative impact on TCP/IP data flows.

Also at the University of Kansas, Quingyang Hu is experimenting with combining fuzzy logic with a rate-based flow control mechanisms [18]. Thus, the control method can learn the traffic patterns and adapt appropriately.

Simulations are also being run at the Illinois Institute of Technology by Professor Graham Campbell, and at Harvard University under Professor H. T. Kung (father of QFC). These simulations evaluate the performance of different types of flow control in different environments. As of the writing of this thesis, no results were available to the general public concerning this research.

There is much research going on in the area of combining or modifying different flow control algorithms to try to combine the advantages of each. Much of this research is being conducted at the University of California at Irvine under the direction of Professor Suda.

I was unable to find any other research concerning the interoperability of these two flow control mechanisms. There are, however, several papers concerning the interoperability of different flow control mechanisms, such as EFCI and ER.

8.2 Areas for Future Research

Several topics for future research were mentioned in earlier chapters. Some of the ideas are listed here.

1. Instead of using greedy persistent sources, use TCP/IP data flows for the ABR sources. This would give a better idea of how the flow control affects ABR throughput.

2. Combine the two mechanisms so both per-link flow control and an end-to-end flow control exist. Evaluate the amount of overhead for this mechanism and compare it to the overhead/performance of the individual mechanisms.

3. Evaluate how often ER exhibits the isolations spoken of in Section 7.3 and investigate methods of reducing and/or eliminating these isolations.

4. Evaluate the performance of the two mechanisms and the combination of them in catastrophic conditions (such as a link going down or a switch failing).

5. Compare QFC with other rate-based mechanisms, such as EFCI.

6. Investigate the effectiveness of using the QFC mechanisms in a multicast environment.

7. Perform the same tests set forth in this paper in a LAN or CAN (Campus Area Network) environment.

8. Test the effectiveness of using QFC or ER in LAN Emulation.

9 Glossary

- Asynchronous Transfer Mode (ATM) - A high-speed connection-oriented multiplexing and switching method specified in international standards that utilizes fixed-length cells to support multiple types of traffic. It is asynchronous in the sense that cells carrying user data need not be periodic.
- ATM Application Layer (AAL) - A set of internationally standardized protocols and formats that define support for circuit emulation, packet video and audio, and connection-oriented and connectionless data services.
- Available Bit Rate (ABR) - Designed for normal data traffic that cannot (or is hard to) be predicted in advance and is not time sensitive.
- Campus Area Network (CAN) - A network which operates over a region the size of a campus, usually on order of 5 km or less.
- Cell Delay Variance (CDV) - The variance in the rate at which the source is receiving data.
- Cell Loss Rate (CLR) - The rate at which the network loses cells.
- Cell Transfer Delay (CTD) - The total delay that the data experiences from source to destination.
- Constant Bit Rate (CBR) - Used for emulating circuit switching, where the required bandwidth is constant and known in advance.
- Control Loop - The round trip that network congestion information must traverse in a given flow control mechanism.

- ER - see EPRCA

- Explicit Forward Congestion Indicator (EFCI) - A rate-based, binary, end-to-end flow control mechanism used by the industry to control ABR traffic.

- Explicit Proportional Rate Control Algorithm (EPRCA) - A rate-based end-to-end flow control mechanism that governs the rate at which a source can send ABR data.

- Fiber Distributed Data Interface (FDDI) - Fiber optic LAN operating at 100 Mbps.

- Jitter - When the rate at which data is being received is highly variable.

- Local Area Network (LAN) - A MAC level data and computer communications network confined to short geographic distances, usually on order of .5 km or less.

- Medium Access control (MAC) - IEEE 802 defined media-specific access control protocol.

- Metropolis Area Network (MAN) - A network which operates over a region the size of a metropolis, usually on order of 25 km or less.

- Minimum Cell Rate (MCR) - The smallest rate at which the source can be reduced to by the network.

- Motion Picture Encoding Group (MPEG) - A standard from the ISO Motion Picture Expert Group using fully encoded key frames, motion prediction, and difference encoding to compress video.

- National Institute of Science and Technology (NIST) - An institute of research in high speed networking.
- Network Interface Card (NIC) - A card placed inside a computer that communicates with a network.
- Peak Cell Rate (PCR) - The largest rate at which the source can send.
- Permanent Virtual Connection (PVC) - A logical dedicated circuit between two user ports in a point-to-point configuration.
- Quality of Service (QoS) - A set of parameters (e.g., CLR, CTD, CDV) that specifies the type of transmission that the application needs or desires.
- Quantum Flow Control (QFC) - A credit-based link-by-link flow control mechanism that controls the flow of ABR data by the use of credits and windowing.
- Resource Management Cell (RM) - The cell used by ABR flow control to send network condition information.
- Round Trip Time (RTT) - The time it takes for one bit of information to go from source to destination and back to the source.
- Sustainable Cell Rate (SCR) - The rate at which the source will be able to sustain a data rate.
- Switched Virtual Connections (SVC) - Virtual circuits similar to PVCs, but established on a call-by-call basis.

- Synchronous Optical Network (SONET) - A United States high-speed fiber optic transport standard for a fiber optic digital hierarchy (speeds range from 51.84 Mbps to 2.4 Gbps).

- Synchronous Transfer Mode (STM) - The T1 carrier method of assigning time slots as channels within a T1 or E1 circuit.

- Transmission Control Protocol/Internet Protocol (TCP/IP) - The combination of a network and transport protocol developed by ARPANET for internetworking IP-based networks.

- Unspecified Bit Rate (UBR) - Designed for those data applications that want to use any leftover bandwidth and are not sensitive to cell loss or delay.

- Variable Bit Rate (VBR) - Allows users to send at a variable rate that is known in advance and can be defined in terms of PCR, SCR, and MCR. VBR has two sub-categories: real-time (VBR-RT) and non-real-time (VBR-NRT).

- Wide Area Network (WAN) - A network which operates over a large region, usually from 25 km to several thousand kilometers.

10 Bibliography

References

[1] A. Alles. *ATM Internetworking*. Cisco Systems, May 1995.

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

[3] John David Cavanaugh and Timothy J. Salo. Internetworking with ATM WANs. *Minnesota Supercomputer Center, Inc.*, 1992. Found at http://ftp.univ-valenciennes.fr/ATM/doc.html.

[4] Y. Chang, N. Golmie, and D. Su. Study of Interoperability Between Various Rate-based Flow Control Mechanisms for Available Bit Rate Traffic in ATM Networks. In *Proceedings of the Fourth International Conference on Computer Communications and Networks*, September 1995.

[5] D. Do. A framework for comparison of congestion control techniques for gigabit ATM networks. *WWW*, 1994.

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

[7] Mike Gaddis. *Quantum Flow Control Version 2.0 FCC-SPEC-95-1.* QFC Alliance, July 1995.

[8] Mark W. Garrett and Walter Willinger. Analysis, modeling and generation of self-similar VBR video traffic. *Computer Communication Review*, 24(4):269–280, October 1994.

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

[10] Duke Hong. Congestion control and prevention in ATM networks. 1991. University of California, Irvine.

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

[12] 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.

[13] H. T. Kung and Robert Morris. Credit-Based Flow Control for ATM Networks. *IEEE Network*, April 1995.

[14] H.C. Lauer. On the Duality of Rate-Based and Credit-Based Flow Control. In *Proceedings of the ACM SIGCOMM '95 Conference*, August 1995.

[15] Mengjou Lin, Jenwei Hsieh, David H.C. Du, and Joseph P. Thomas. Distributed network computing over local ATM networks. *IEEE Journal on Selected Areas in Communications Special Issue of ATM LANs: Implementations and Experiences with an Emerging Technology*, 1995.

[16] Bruce McClure. ER: ATM flow-control mechanism gone awry. *LAN Times*, Febuary 1996.

[17] Ascom Nexion. Bringing efficiency and reliability to ATM. White paper found at http://www.nexen.com/wpqfc.html, 1996.

[18] Cameron Braun Qingyang Hu, David W. Petr. Self-tuning Fuzzy Traffic Rate Control for ATM Networks. Technical Report icc.96b, University of Kansas, 1996.

[19] K.K. Ramakrishnan. ATM flow control: Inside the great debate. *Data Communications*, June 1995.

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

[21] S. Sathaye. *The ATM Forum Traffic Management Specification Version 4.0 af-tm-0056.000.* ATM Forum, April 1996.

[22] M. Schwartz. *Broadband Integrated Networks.* Prentice-Hall, Inc., May 1996.

[23] Robert Simcoe. Reliable and Efficient Hop-by-Hop Flow Control. *IEEE Journal on Selected Areas in Communications*, April 1995.

[24] Kai-Yeung Siu and Hong-Yi Tzeng. Intelligent congestion control for ABR service in ATM networks. 1993. University of California, Irvine.

[25] Brett Vickers. Congestion control and resource management in diverse ATM environments. *IEICE*, 1993.

[26] Robert Walthall. Simulation and Analysis of the Performance of EPRCA in a Wide Area ATM Network Consisting of Both ER and EFCI Switch Mechanisms. In *Proceedings of the Fourth International Conference on Telecommunications*, September 1995.

Simulation and Analysis of Credit- and Rate-based

Switch Interoperability in an ATM Network

Charles R. Duncan

Department of Computer Science

M.S. Degree, Febuary 1997

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 thesis, I evaluate the performance of ABR service in a Wide Area Network (WAN) and a Metropolis area Network (MAN) made up of both Explicit Rate (Rate-based) and Quantum Flow Control (Credit-based) switches. MPEG video traces from MPEG-2 encoded video streams are used to generate background VBR traffic. Variable Bit Rate (VBR) sources are used to simulate the real world environment in which ABR service will be required to operate (transient state). The performance of ABR service is evaluated in terms of throughput, buffer occupancy, burstiness, fairness, and fast convergence.

COMMITTEE APPROVAL:

Mark J. Clement, Committee Chairman

Bryan S. Morse, Committee Member

Parris K. Egbert, Committee Member

Scott N. Woodfield, Graduate Coordinator