MODELING AND SIMULATION OF A VIDEO-ON-DEMAND NETWORK IMPLEMENTING ADAPTIVE SOURCE-LEVEL CONTROL AND RELATIVE RATE MARKING FLOW CONTROL FOR THE AVAILABLE BIT RATE SERVICE

by

Elvin Lattis Taylor, Jr.

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

in

Electrical Engineering

Scott F. Midkiff, Chair

Nathaniel J. Davis, IV

Charles J. Parry

December, 1997
Blacksburg, Virginia

Keywords: Video-on-Demand, Traffic Management, High-speed Networking

MODELING AND SIMULATION OF A VIDEO-ON-DEMAND NETWORK IMPLEMENTING ADAPTIVE SOURCE-LEVEL CONTROL AND RELATIVE RATE MARKING FLOW CONTROL FOR THE AVAILABLE BIT RATE SERVICE

by

Elvin Lattis Taylor, Jr.

Scott F. Midkiff, Committee Chair

Electrical Engineering

(ABSTRACT)

The Available Bit Rate (ABR) service class for the Asynchronous Transfer Mode (ATM) protocol was originally designed to manage data traffic. ABR flow control makes no guarantees concerning cell transfer delay or cell delay variation. A closed-loop feedback mechanism is used for traffic management. To use this class of service for video transport, the video source will accept feedback from the network and adapt its source rate based on this status information. The objective of this research is to assess the ability of the ATM ABR service class to deliver Moving Picture Experts Group version 1 (MPEG-1) video. Three approaches to source-level control are compared: (i) arbitrary loss or no control method, (ii) selective discard of MPEG B-pictures, and (iii) selective discard of MPEG B- and P-pictures. Performance is evaluated based on end-to-end delay, congested queue occupancy levels, network utilization, and jitter. A description of the investigation, assumptions, limitations, and results of the simulation study are included.

Acknowledgments

I would like to thank my advisor Dr. Scott F. Midkiff for his advice, patience, and commitment to excellence during this research endeavor, including the writing of this thesis. I want to thank committee members Drs. Nathaniel J. Davis and Charles J. Parry for their time and cooperation in this study.

I want to thank the National Science Foundation Graduate Research Fellowship Program for supporting this work.

I would like to thank Farooq-e-Azam, Scott Harper, Rhett Hudson, and David Lee for help with solving UNIX and OPNET problems.

I want to thank my parents Elvin L. Taylor, Sr. and Joyce W. Taylor and my sister Michelle R. Taylor for their love, encouragement, and support.

I want to thank Cheryl D. Earle, Jeff Earley, Edward McPherson, and James Moore III for their support.

I want to thank O. Rose for making publicly available the MPEG data sets used in this study. Data sets are available via anonymous FTP to ftp-info3.informatik.uni-wuerzburg.de in the /pub/MPEG/ directory.

Finally, I want to thank God who has allowed me to triumph through patience and perseverance [Philippians 4:13].

Table of Contents

List of Figures

List of Tables

Chapter 1. Introduction

The Broadband Integrated Services Digital Network (B-ISDN) concept is a world-wide effort to merge data[1] (packet-switched), voice (circuit-switched), and video (e.g., cable television) networks [LAPO94]. This move to produce a single network capable of providing services for computer communications, telecommunications, and video delivery, has been widely accepted. The Asynchronous Transfer Mode (ATM) has been chosen as the connection-oriented, cell-based, switching and (de)multiplexing method for B-ISDN. Of the three services mentioned above, providing video service, such as video-on-demand, video-conferencing, and real-time broadcasts, is the most challenging because of its high bandwidth requirements and fixed delay bounds. To handle video service, ATM designates the Constant Bit Rate (CBR) and the Variable Bit Rate (VBR) service classes. These service categories guarantee strict contract requirements, agreed to by the source and the network, before the network establishes the connection.

The Available Bit Rate (ABR) service class, which was originally designed to handle data traffic, meets less stringent requirements. The ABR class makes no guarantees concerning cell transfer delay or cell delay variation. A closed-loop feedback mechanism is used for traffic management. To use this class for video service, however, requires that the video source interact with the network. More specifically, the video source must accept feedback from the network and adapt its source rate based on this status information. If the network status indicates congestion, then the source must lower its rate.

1.1 Research Objective

Previous research in ABR flow control includes [TSE96], [ZUKE97], [WALT96], and [GUPT95]. In [TSE96], the authors evaluate the performance and fairness of the Explicit Forward Congestion Indication (EFCI). Multiple ABR source-destination pairs transport

[1] In this thesis, data refers to computer network traffic such as electronic mail or file transfers.

data while VBR video traffic is sent in the background. Network congestion status information is fed back to the ABR sources, but not to the VBR source. Therefore the video source is non-adaptive. This research is further extended in [ZUKE97]. The same network configuration is used, however, comparison is made between EFCI and Explicit Rate (ER) flow control.

This researcher is aware of only one study that examines MPEG video over an ATM network implementing ABR service [WALT96]. The study focuses on network performance when switches that support ER and EFCI are present in a network. Source rate adaptation is not addressed. In [GUPT95], the authors suggest that adaptive video coding algorithms seem suitable for high speed networks offering ABR services. Their research focuses on the LAN environment. Also, their simulation is based on a generic network model and not on an ATM network providing ABR service.

The objective of this research is to assess the ability of the ATM ABR service class to deliver video. Three approaches to source-level control are compared: (i) arbitrary loss or no control method, (ii) selective discard of MPEG B-pictures, and (iii) selective discard of MPEG B- and P-pictures. Performance is evaluated based on end-to-end delay, congested queue occupancy levels, network utilization, and jitter.

This research makes the following assumptions.

- An ATM Wide Area Network (WAN), representing an MPEG-1 Video-on-Demand (VoD) backbone network, is simulated.
- Switches that support Relative Rate Marking are used. This type of switch is defined in [ATMF96] as part of Traffic Management Specification Version 4.0.

Limitations of this research are listed below.

- **Interactivity**. A VoD service, by definition, should allow interaction. The customer sends an upstream request for a specific movie title and a start time to the service

provider. This thesis does not address the issue of interactivity, but is concerned with the unidirectional distribution of high data rate video traffic.

- **Set-Top Box Features**. The network interface device through which each customer receives the service is called a set-top box. The study does not model effects of VCR functions such as fast-forward, fast-rewind, and pause/resume.
- **Errored Cells**. In any network, there exists a non-zero probability that the channel will introduce bit errors. In this instance, the network or the destination could discard or correct errors which are detectable. This study, however, does not model the effects of errored cells.
- **Audio**. The associated movie audio, which is much less demanding of network resources than the video, is omitted.

1.2 Thesis Organization

Chapter 2 provides an overview of significant phases in communications networks leading to the development of ATM. Historical highlights include the T1 carrier, optical fiber, packet-switching, narrow and broadband integrated services digital networks, and ATM.

Chapter 3 presents ATM bandwidth management. Quality of Service parameters are discussed. Real-time and non-real-time service categories are also explained. Lastly, explicit forward congestion indicator, explicit rate, and relative rate switch mechanisms are described.

Chapter 4 gives background information about video compression standards and adaptive video compression techniques. Common video compression schemes, Motion-JPEG, H.261, and MPEG-1, are reviewed. Adaptive video encoding methods are presented.

Chapter 5 describes the experimental design and implementation. Explanations are given for each node model, queue or processor module, and process model. Also, this chapter

presents end-to-end delay, queue occupancy, backbone utilization, and jitter, which are the response variables used to evaluate the model.

Chapter 6 presents results of this research. This chapter describes the trace data sets used, the model's backbone data rate, bursty source parameters, and network queue congestion. Also, based on the response variables, this chapter presents an evaluation of the drop mode off, dropping B-pictures only, and dropping B- and P-pictures cases. Alternative representations of ETE delay, such as delay variation (jitter) and distribution of delay, are also discussed.

Chapter 7 contains the conclusions.

Chapter 2. Communications Networks Overview

2.1 Historical Highlights in Communications

Significant phases in communications leading to the development of ATM can be partitioned into decades. Each subsequent phase results from technology that improves upon previous phases.

2.1.1 The T1 Carrier

During the 1960s, the world began to upgrade from its all-analog public switched telephone systems to new systems capable of combining analog and digital systems. In North America, the system culminated into one 1.544 megabits per second (Mbps) signal commonly known as the "T1 Carrier." This signal is formed by taking 24 channels at telephone-quality voice signal rate of 64 kilobits per second (Kbps) and multiplexing them together [VETT95, SAAD94]. Note that 64 Kbps derives from 8-bit Pulse Code Modulation (PCM) samples taken at a rate of 8,000 samples per second.

2.1.2 Optical Fiber Medium and Packet-Switching

In 1970, the first low-loss optical fiber that could achieve less than 20 dB/km loss was introduced by Corning Glass Works. This permitted networks to start using optical fiber as the physical transmission medium for high-speed communications in the wide area.

In 1971, the Advanced Research Project Agency (ARPA) of the U. S. Department of Defense, activated an experimental packet-switching network called ARPANET. These developments led to digital data being transferred more quickly and with fewer bit errors than ever before [VETT95, SAAD94].

2.1.3 Integrated Services Digital Network

As communications networks emerged, they became diverse and specific to their application. For example, cable companies used a dedicated network for transmitting

television signals. Telephone companies used their own networks for carrying voice. Generally, computers had networks which traditionally transferred files, delivered electronic mail, etc. Some argued that these formats should be combined. One advantage to merging these formats was that through large-scale use of a single network, this technology would become more affordable to build and maintain [BERT92].

As a result of the goal to combine traffic types, in 1984, the 6th Plenary Assembly of the International Consultative Committee on Telegraphy and Telephony (CCITT) adopted the first I-series recommendations pertaining to a Narrowband Integrated Services Digital Network (N-ISDN) [SAAD94, HÄND91]. (The CCITT is now known as the International Telecommunications Union, Telecommunications Standardization Sector -- ITU-TSS or ITU-T.) The purpose of such a network was to merge multiple services such as voice, data, and low-grade video at rates between 64 Kbps and 2 Mbps. N-ISDN or ISDN uses the digital voice network as its medium. The integration of these services by ISDN is considered a positive step, but has limited scope. ISDN gives end-users a unified interface to a set of networks and services, but network and service providers still maintain multiple networks. Further, LANs continue to use protocols different from those of ISDN [LAPO94]. Because ISDN is not suited to handle applications requiring high bandwidth, it can be considered as a way to exploit current voice networks for services offered and cost efficiency [DEPR93].

2.1.4 Broadband Integrated Services Digital Network

Recognizing the limitations of ISDN, in 1985, the ITU-T (then called CCITT) formed the Task Group on Broadband Aspects of ISDN (B-ISDN) to study data rates higher than 2 Mbps [SAAD94]. This was the first of several recommendations concerning B-ISDN, which combines not only voice and data, but also video traffic. Subsequent recommendations were published, designating the Synchronous Optical Network (SONET) in North America and Synchronous Digital Hierarchy (SDH) in Europe for the data rates, synchronization, and framing format for B-ISDN. In 1988, the Asynchronous

Transfer Mode (ATM) was chosen as the switching and multiplexing method for B-ISDN [VETT95].

2.1.5 Synchronous Optical NETwork

B-ISDN uses Synchronous Optical NETwork (SONET) in North America. SONET is scaleable in multiples of 1, 3, 12, 48, and 192 times 51.84 Mbps (i.e., 51.84 Mbps, 155.52 Mbps, …, 9.953 Gbps). Data rates are commonly referred to as Optical Carrier 1 (OC1 or 51.84 Mbps), Optical Carrier 3 (OC3 or 155.52 Mbps), et cetera. Also, optical fiber provides better security since it is more difficult to patch into an optical line than it is to tap a copper wire.

2.2 Asynchronous Transfer Mode

The Asynchronous Transfer Mode (ATM), also known as cell switching, cell relay, or fast packet, is the transport mode used for B-ISDN. ATM partitions all of its information, including voice, data, video, and control, into fifty-three-octet[2] long cells, as shown in Figure 2.1. Utilizing fixed-size packets simplifies multiplexing and switching operations.

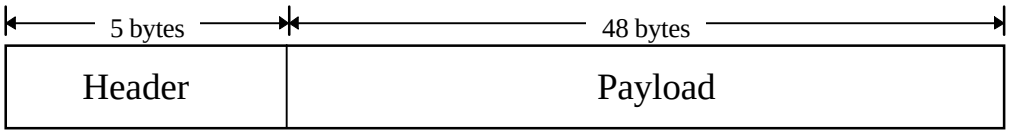

Figure 2.1 ATM Cell Format

The payload size is a compromise between the telephone and computer industries [VETT95]. According to [BERT92], 6 ms are required to sample 48 bytes of voice data read at 64 Kbps. Reconstructing the analog voice signal at the destination must allow for this 6 ms delay plus transmission delay. There is the possibility that this rebuilt voice may be 'partially reflected' at the receiver and added to the voice from destination back to source. This means that there could be a period of as much as a 12 ms plus two times the

[2] Octet and byte are used interchangeably throughout this thesis.

transmission delay before the original sender hears her/his echo. Echo cancellation is available in some areas, but because its use is not ubiquitous, it was decided that this delay needed to be small. Thus telephone companies wanted the cell to be no larger than 32 bytes.

On the contrary, computer companies wanted a cell size no smaller than 64 bytes, since transferring large data files would incur a large performance penalty because of header overhead and additional delay due to segmenting and reassembling. The result is an agreement to average the two sizes, yielding 48 bytes [VETT95].

2.2.1 ATM Protocol Reference Model

Standards created by the ITU are the foundation of the ATM Protocol Reference Model (PRM). The ATM PRM, shown in Figure 2.2, is layered into three levels: the physical layer, the ATM layer, and the ATM Adaptation Layer (AAL).

2.2.1.1 Physical Layer

The physical layer specifies the physical characteristics of the media such as cable and connector types and the electrical carrier signal properties, including voltage and wavelength. This lowest layer consists of two sublayers, the Physical-Medium-Dependent (PMD) sublayer and the Transmission Convergence (TC) sublayer. The PMD sublayer interfaces the particular transport medium. Supported media are single-mode optical fiber, multimode optical fiber, and category five unshielded twisted pair wire. This sublayer is also responsible for bit timing and data encoding and decoding. The TC sublayer is the interface for enabling other physical layers to transfer ATM cells. The TC sublayer handles cell delineation (determination of boundaries), performance monitoring, generating and processing of Header Error Check (HEC), and payload rate matching of the different transport media [SIUJ95, ATMF95].

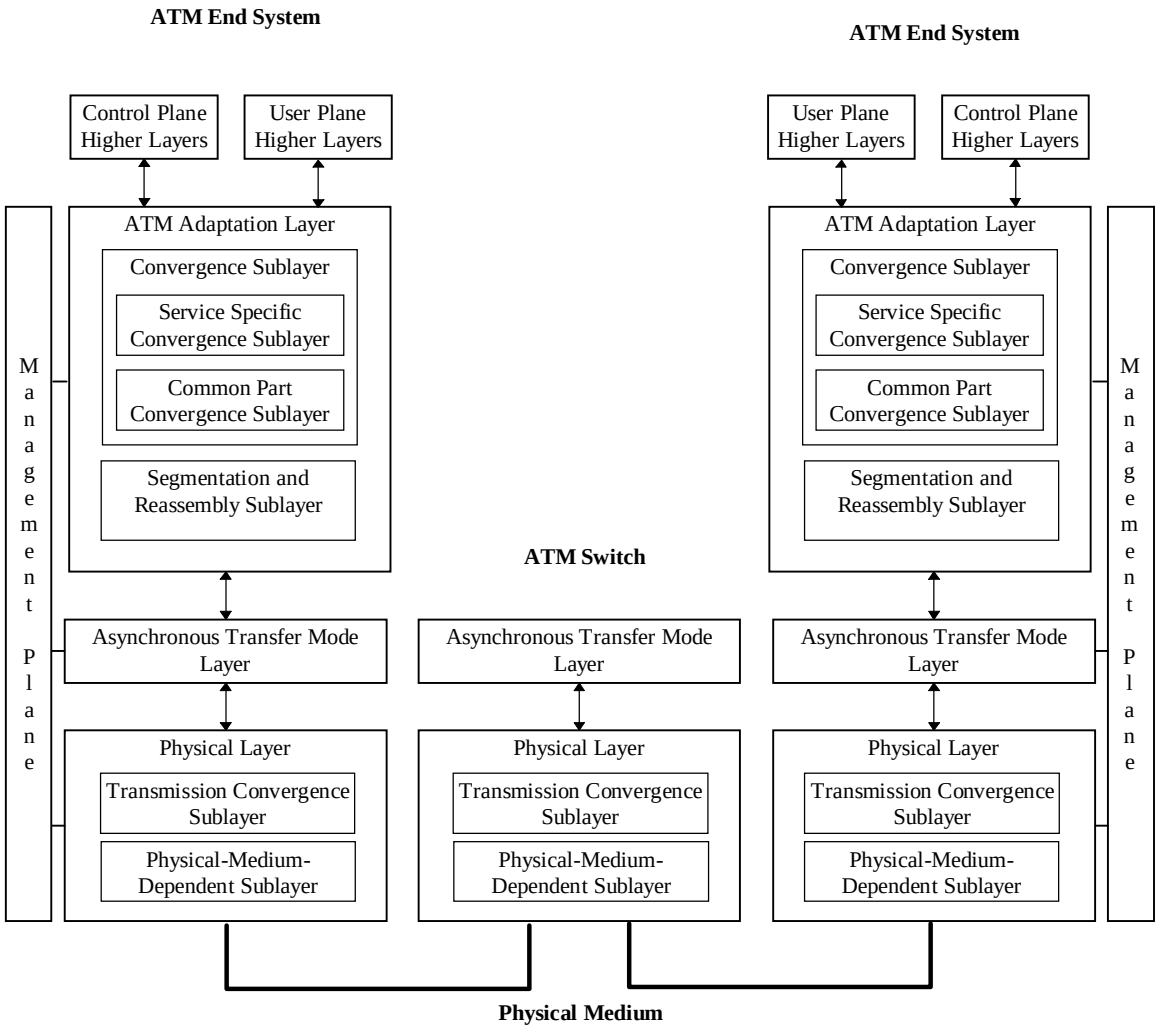

Figure 2.2 ATM Protocol Architecture [LAPO94, ROOH95, SIUJ95]

Synchronous Optical NETwork (SONET) and Synchronous Digital Hierarchy (SDH) are the most often used framing and synchronization methods. Common transmission rates are 155.52 Mbps and 622.080 Mbps.

2.2.1.2 ATM Layer

The ATM layer provides the link between the physical layer and the AAL. When located in a switch, the ATM layer decides to which link the incoming packets should be forwarded, resets the corresponding connection identifiers, and then passes the cells to the

next link. Additionally, the ATM layer is responsible for unassigned-cell[3] insertion and deletion, cell priority processing and scheduling, cell loss priority marking and reduction, pacing and peak rate enforcement of the cell rate, cell payload type marking and differentiation, explicit forward congestion marking and indication, and generic flow control. The ATM layer constructs and verifies the cell header and routes and multiplexes cells onto and demultiplexes cells from a single stream. Figure 2.3 is a diagram of the cell header.

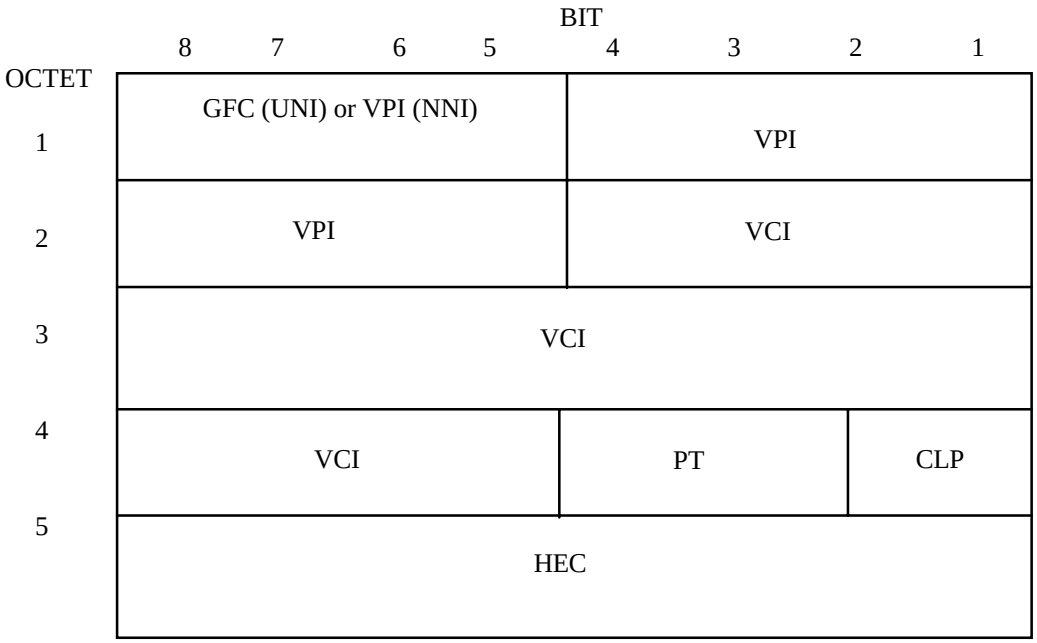

Figure 2.3 ATM Cell Header Format

There are two types of header format for ATM cells, one for the User-to-Network Interface (UNI) and another for the Network Node Interface or Network-to-Network Interface (NNI). For the UNI, the four-bit Generic Flow Control (GFC) field was designed to control the amount of traffic entering the network. This capability is important since it permits the network to limit the volume of data entering the network

[3] Unassigned-cells (empty cells) are sent on certain physical links when there is no data to transmit but something is required to be sent because of the framed structure of the link [DUTT95]. This concept is also known as cell rate decoupling.

during congested periods. For the NNI, there is no GFC field. These four bits, instead, are used for the Virtual Path Identifier (VPI) field.

- **Virtual Path Identifier and Virtual Channel Identifier.** The Virtual Path Identifier (VPI) field and the Virtual Channel Identifier (VCI) field compose the two-level addressing structure. The VPI indicates the physical path used by a group of VCIs, or in other words, a VPI can be considered bundled VCIs with a common endpoint. A VCI is used to set up a connection between two end users through the network. This routing scheme allows for 16.8 million sessions at the UNI (8-bit VPI, 16-bit VCI) and 268.4 million sessions at the NNI (12-bit VPI, 16-bit VCI). Figure 2.4 is a conceptual view of how VPIs and VCIs are related.

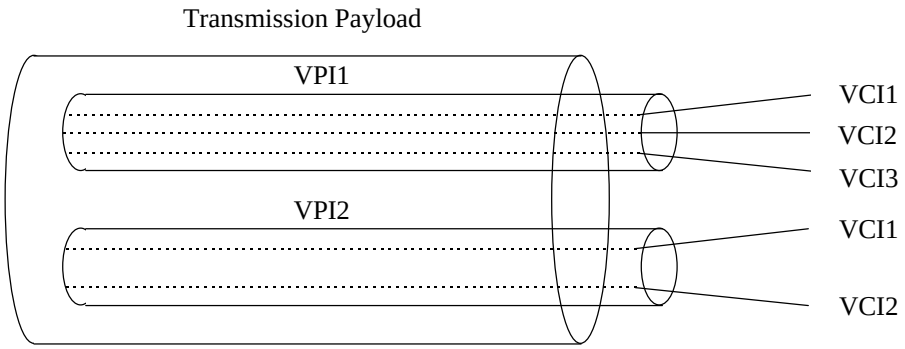

Figure 2.4 ATM Connection Identifier [SAAD94]

- **Payload Type Indicator.** The three-bit Payload Type Indicator (PTI) differentiates between cells containing control information and cells carrying user data. This implies beneficial separation of control information and user data on different subchannels. Additionally, this field can be used to flag cells that have experienced congestion or to identify traffic that is from an ATM Adaptation Layer 5 (AAL5) connection.

- **Cell Loss Priority.** The one-bit Cell Loss Priority (CLP) field denotes high priority versus low priority data. The CLP can be explicitly set by the network or by the user.

The network can drop any cell marked with a high cell loss priority. In other words, a cell with CLP=1 is less critical than a cell with CLP=0. This property is beneficial to the network when it becomes congested. This feature also helps with network policing by allowing the network to discard cells from a connection whose source is exceeding its negotiated contract parameters.

- **Header Error Checksum.** The Header Error Checksum (HEC) uses an 8-bit Cyclic Redundancy Checksum (CRC) to detect an error in the header. Single-bit error correction is sufficient to remove most header errors since ATM most often uses reliable fiber optic lines where the Bit Error Rate (BER) is lower than 10^{-9} [SIUJ95].

2.2.1.3 ATM Adaptation Layer

The ATM Adaptation Layer or AAL provides an interface between the services required by higher layer protocols and the ATM layer. Services provided by the AAL are end-to-end. Several AAL protocols have been defined. All AAL protocols are partitioned into two sublayers, the Segmentation and Reassembly (SAR) sublayer and the Convergence Sublayer (CS). The SAR sublayer is primarily responsible for the segmentation of user data units from higher-layer protocols into ATM cells and reassembling ATM cells into packets of information which have meaning to higher-layer protocols. The CS further subdivides into the Service Specific Convergence Sublayer (SSCS) and the Common Part Convergence Sublayer (CPCS) [LAPO94]. These sublayers, which are not relevant to this research, are mentioned for completeness but are not discussed further. The CS handles error detection, flow control, and other functions. This sublayer depends on the type of service that is being provided [ROOH95, SAAD94]. Table 2.1 contains descriptions of the different AALs and corresponding service classes.

Table 2.1 ATM Service Classes [HÄND91, LEE95]

Attribute	Service Class			
	Class A	Class B	Class C	Class D
Timing relationship between source and destination	Required		Not required	
Bit rate	Constant	Variable		
Connection mode	Connection-oriented			Connectionless
ATM Adaptation Layer	AAL-1	AAL-2	AAL-3/4 or AAL-5	AAL-3/4 or AAL-5

There are five AAL formats and four service classes [DUTT95].

- **AAL-0.** AAL-0 designates a "transparent mode" for AAL. No AAL functionality is specified and cells are relayed between the service interface and the ATM network.
- **AAL-1.** AAL-1 provides functions for traffic Class A, which emulates a continuous or constant-bit-rate leased line. Example applications are 64 Kbps voice and fixed rate, uncompressed video.
- **AAL-2.** AAL-2 provides necessary functions for variable-bit-rate traffic in traffic Class B. Example application areas are transport of compressed voice traffic with silence suppressed and delivery of compressed, packetized video.
- **AAL-3/4.** AAL-3/4 is capable of handling connection-oriented data applications of Class C and connectionless data applications of Class D. Originally, AALs 3 and 4 were defined separately but have been combined since it was realized that the same process could perform both functions.
- **AAL-5.** AAL-5 is similar to AAL-3/4 but with less functionality and thus less complexity. Also, there is error detection with no error recovery. This format typically supports datagram traffic and is known as the Simple and Efficient Adaptation Layer (SEAL). Currently, AAL-5 is the most commonly implemented adaptation layer. This layer is targeted to support traffic Classes C and D, but vendors

are using it to support all of the service classes. A typical service is support of datagram traffic.

2.3 Summary

This chapter provided an overview of significant phases in communications networks leading to the development of ATM. Historical highlights began with the 1.544 Mbps T1 carrier developed in the 1960s. Also, mentioned were optical fiber from Corning Glass Works which made it possible for the Department of Defense to implement an experimental packet-switching network called ARPANET.

As communications networks emerged, they became diverse and specific to their application. ISDN was an attempt to unify data and voice network services. B-ISDN went further in that it incorporated video as well as voice and data services. ATM was chosen as the switching and multiplexing method for B-ISDN. Among the ATM protocol's features is that it packetizes video, voice, and data into 53-byte cells. This enables simplified cell handling. A key property of ATM is bandwidth management, which is discussed in Chapter 3.

Chapter 3. Bandwidth Management

Bandwidth management, also known as traffic management, is critical in high-speed networks such as ATM. In the few milliseconds of time it might take for the network to notice that there is congestion, a large number of cells could be transmitted and then subsequently dropped. If the dropped cells contained data traffic, retransmissions[4] must occur, which further compound the congestion [KRAM95]. (Although video and voice traffic are not retransmitted, their performance may be degraded by compounded network congestion due to retransmissions of data traffic.)

3.1 Quality of Service Traffic Contract

Each ATM connection has a given Quality of Service (QoS) associated with it. A QoS traffic contract is negotiated before the connection is established. This agreement is used to guarantee allocation of network resources to a connection and maximize utilization of network resources. During severely overloaded periods, when the network cannot recover from overload by dropping only non-critical cells with CLP=1 (CLP is defined in Section 2.2.1.2), the network can choose which cells to discard based on the QoS characteristics of the connection [DUTT95, SIU95]. The following section describes the attributes specified in the QoS contract.

3.1.1 Contract Parameters

A source can specify the following parameters during connection set-up [ATMF96, SIU95].

- **Maximum Cell Transfer Delay.** Maximum Cell Transfer Delay (maxCTD) is defined as the elapsed time between a cell entering the network and exiting the network for a certain connection. One can consider maxCTD as the maximum network latency.
- **Peak-to-peak Cell Delay Variation.** Cell Delay Variation (peak-to-peak CDV) is the measure of variation in CTD or jitter experienced by a cell traversing the network.

[4] Retransmission would be implemented by some higher, non-ATM protocol layer.

Transmission of voice or video traffic on a network with a large CDV implies large buffers, which would shape the traffic so that delay experienced by the cells will be more uniform. Large buffers, however, imply large delay.

- **Peak Cell Rate.** Peak Cell Rate (PCR) is the highest instantaneous rate at which the user will send information. The PCR is the inverse of the minimum cell inter-arrival time.
- **Minimum Cell Rate.** Minimum Cell Rate (MCR) is the minimal rate which the network guarantees to provide for the connection.
- **Sustained Cell Rate.** Sustained Cell Rate (SCR) is the average cell rate across a connection during a long time interval.
- **Cell Delay Variation Tolerance.** When cells from multiple connections are merged, cells from one connection might be delayed while cells from another connection are being inserted at the multiplexer's output. As a result, some of the randomness may affect the interarrival time between successive cells of a given connection. Cell Delay Variation Tolerance (CDVT) is specified to minimize the influence of this randomness.
- **Cell Loss Ratio.** Cell Loss Ratio (CLR) is defined for a connection as the ratio of lost cells in the network that were originally sent by the source, to the total number of cells transmitted by the source. Potential causes of cell loss can be unintentional buffer overflow or deliberate cell discarding policies. The CLR can be specified both for low-priority cells (CLP=1) and high-priority cells (CLP=0).
- **Maximum Burst Size.** Maximum Burst Size (MBS) is given by Equation 3.1.

$$MBS = \left\lfloor 1 + \frac{\tau_s}{T_s - T} \right\rfloor \qquad \text{Equation 3.1}$$

where $\lfloor x \rfloor$ represents the integer part of x, τ_s is the Burst Tolerance, T_s is the reciprocal of the SCR, and T is the reciprocal of the PCR [ATMF96]. Burst Tolerance, alternatively denoted in the literature as BT, is the largest burst length of packets at the PCR. BT is the bucket size parameter in the leaky bucket algorithm that controls the traffic entering the network [SIU95]. [ATMF96] discusses in detail

the leaky bucket algorithm as it relates to ATM's Generic Cell Rate Algorithm (GCRA).

The following parameters are not negotiated.

- **Cell Misinsertion Rate.** A misinserted cell outcome occurs when a cell is received for which there is no corresponding transmitted cell. Cell Misinsertion Rate (CMR) is defined as the number of misinserted cells divided by an interval of time. Cell misinsertion is most often due to an undetected error in the header of a cell being sent on a different connection.
- **Cell Error Ratio.** Cell Error Ratio (CER) for a connection is defined as the ratio of errored cells to the sum of errored cells plus successfully transferred cells.

3.1.2 Service Categories

There are five service categories defined by the ATM Forum. They are:

- **CBR** Constant Bit Rate
- **rt-VBR** Real-Time Variable Bit Rate
- **nrt-VBR** Non-Real-Time Variable Bit Rate
- **UBR** Unspecified Bit Rate
- **ABR** Available Bit Rate

These categories can simultaneously exist on the same physical link. Also, any of the five categories may be used on Virtual Channel Connections (VCC) and Virtual Path Connections (VPC) [SCHU95]. These service categories belong to one of two groups, either real-time or non-real-time. Categories CBR and rt-VBR are designed for real-time traffic. The three remaining categories are designed for non-real-time traffic.

3.1.2.1 Real-Time Traffic

Both CBR and rt-VBR negotiate a PCR required for the connection. The PCR value is specified to make bandwidth available during the lifetime of the connection. All switches along the connection must be able to guarantee the PCR or the connection will be refused.

The source may emit cells at or below the agreed upon PCR. For physical connections with only CBR traffic, when sources transmit cells at rates below the PCR, they waste bandwidth [SCHU95].

The rt-VBR service category is characterized by an SCR, an MBS, and a CDVT in addition to a PCR. (Recall that these parameters are described in Section 3.1.1). For brief periods, the source may burst above the SCR (up to the PCR), but the network will maintain the SCR as the average for the connection by adjusting later traffic flow to a lower rate for a corresponding period of time [HUGH95]. Switches along the new connection calculate the probability that all the VBR connections it is currently supporting will burst at their respective PCRs simultaneously. If the probability is too high, then the new connection will not be established. An application transmitting below SCR is potentially wasting bandwidth.

3.1.2.2 Non-Real-Time Traffic

The nrt-VBR service class targets non-real-time applications which have bursty traffic characteristics. The same traffic parameters specified for rt-VBR are specified for nrt-VBR (PCR, SCR, and MBS) [ATMF96].

The UBR service category offers "best-effort" service. This class specifies neither a bit rate nor any traffic parameters and, therefore, does not provide any QoS guarantees regarding cell loss, cell delay, or cell delay variation. The UBR class utilizes excess bandwidth, but does so ineffectively since the network considers these cells least important. Successful throughput or "goodput" can fall to unacceptable levels with this category [HUGH95].

The ABR service class is distinct from other service categories in that it uses a closed-loop feedback mechanism for traffic management. The ABR category utilizes excess bandwidth like the UBR class, but unlike this class, ABR switches monitor network congestion (alleviation) and ABR sources adjust to avoid cell loss (or adjust to use the

excess bandwidth). Parameters PCR and MCR are negotiated. The source agrees not to send traffic faster than PCR and the network agrees to continuously provide MCR or greater. Values for PCR and MCR can default to the available rate and to zero, respectively. Table 3.1 illustrates how ATM service categories are associated with traffic descriptors.

Table 3.1 ATM Service Category Attributes [ATMF96]

Attribute	ATM Layer Service Category				
	CBR	rt-VBR	nrt-VBR	UBR	ABR
Traffic Parameters:					
PCR and CDVT	specified			specified	specified
SCR, MBS, CDVT	n/a	specified		n/a	
MCR	n/a			n/a	specified
QoS Parameters:					
peak-to-peak CDV	specified		unspecified		
maxCTD	specified		unspecified		
CLR	specified			specified	*

* CLR is kept low for sources that adjust cell flow according to control feedback. Network implementation determines whether or not a quantitative value for CLR is specified.

3.2 Switch Mechanisms

There are generally three types of switch control mechanism: the Explicit Forward Congestion Indicator (EFCI), the Explicit Rate (ER), and the relative rate (RR) marking.

The EFCI or binary feedback method is described in UNI 3.1 [ATMF95] as part of the original proposal for an end-to-end rate-based traffic control scheme. This method allows switches to indicate congestion by setting the EFCI bit contained in data cell headers. In general, the EFCI method employs a single FIFO queue, however, per-connection queueing is often used to improve fairness. There are high and low queue thresholds that

determine when congestion exists. When the queue length exceeds the high threshold, data cells passing through the queue have their EFCI bit set. When the queue length drops below the low threshold, the switch returns to an uncongested state [ATMF96].

The ER method has several variations. A popular approach is the Enhanced Proportional Rate Control Algorithm (EPRCA). The EPRCA requires the source to send RM cells, containing a desired ER, at fixed intervals. The switch calculates a Mean Allowed Cell Rate (MACR) for all connections. The switch also observes the queue length. If congested, then when the RM cell returns from its destination, the switch reduces the ER field. If the ER field value is less than 7/8 of the MACR (fair share), then no adjustments are made. On the contrary, if the ER field is greater than the fair share and the switch is congested, then the ER value is lowered to its fair share [ATMF96].

This research uses the RR marking approach. The RR marking method combines characteristics of the EFCI method and the ER method. The RR approach uses RM cells to probe the network like ER. But instead of calculating an explicit rate, switches set the CI and NI bits contained in the RM cells. Congestion is determined based on queue threshold levels as in the EFCI approach. When congestion exists, either switches can discard cells or the source can discard cells or both.

Unlike the EFCI approach, the RR marking technique enables multiple levels of congestion indication: no congestion, light congestion, or heavy congestion. Also, the RR marking method requires lower switch complexity than the ER algorithm.

3.3 Summary

This chapter discussed ATM bandwidth management. Each ATM connection has a QoS traffic contract which is negotiated before the connection is established. The contract guarantees allocation of network resources. A source can specify the following parameters: maxCTD, CDV, PCR, SCR, CLR, and MBS. Parameters which are not

negotiated are CMR and CER. These parameters are used based upon the connection's service category. Service categories are CBR, rt-VBR, nrt-VBR, UBR, and ABR.

Also discussed in this chapter, are three types of switch control mechanisms: EFCI, ER, and RR. EFCI switches set the EFCI bit in data cell headers according to queue threshold levels. ER switches implementing the EPRCA method explicitly calculate a rate to place in the ER field of a special RM cell. RR switches monitor queue levels like the EFCI method, however, congestion is indicated in CI and NI fields within RM cells.

Not only is an understanding of ATM bandwidth management required for this research, but familiarity with video compression schemes is helpful. Chapter 4 provides an overview of video compression techniques.

Chapter 4. Video Compression

Video traffic tends to require high bandwidth as compared with other forms of traffic such as data and voice. For example, delivering uncompressed full-screen (640x480 pixels), full-color (24 bits per pixel), full-motion (30 frames per second) video requires 221 Mbps. This large bandwidth is for only one video stream. Multiple streams would debilitate a network with even a gigabit per second capacity. Thus, there is a need for video compression.

Video compression techniques can be categorized as either lossy or lossless. Lossy compression yields frames that are not quite as detailed as the original frame prior to encoding. The loss of detail is not always visible. On the other hand, the lossless compression process does not discard any information and, therefore, produces an exact copy of the original frame. The following sections discuss compression methods which produce variable bit rate video and constant quality. Lastly, a section lists adaptive video compression techniques.

4.1 Joint Photographic Experts Group

The Joint Photographic Experts Group (JPEG) developed from a merger of the ISO Subcommittee 2, working group 8, with the CCITT (now ITU) study group VIII. The JPEG group issued ISO Draft 10918 in 1991 and an international standard in 1992. The JPEG specification can be used for still image compression and for video sequence compression called Motion-JPEG.

There is a "lossy" mode based on the Discrete Cosine Transform (DCT) and a "lossless" mode based on intraframe (same frame) prediction, entropy and run-length coding. Pictures can range between 1x1 and 65,535x65,535 pixels and each pixel can range between 1 and 255 colors [SCHÄ95].

4.2 H.261

The CCITT (now ITU) developed the H.261 standard, in 1989, for low bit rate color video transmission. The primary use of the H.261 standard is for video delivery over ISDN at p x 64 Kbps, where p=[1,30], for multiple channel transmission. This recommendation specifies a hybrid scheme using DCT and Differential Pulse Code Modulation (DCPM) with motion compensation.

In the H.261 approach, the source encodes the first frame in the intraframe mode to produce an I-picture (low compression ratio). Each subsequent frame is encoded using interframe prediction, based only on the previous frame, to yield a P-picture (high compression ratio). Real-time applications, like video-conferencing, can use H.261 since future frames rely only on adjacent previous frames and the entire video segment does not need to be stored prior to encoding. Finally, the H.261 standard utilizes conditional replenishment; the ability to update a macroblock (8x8 grid of pixels) only if current content differs from previous content in the same macroblock.

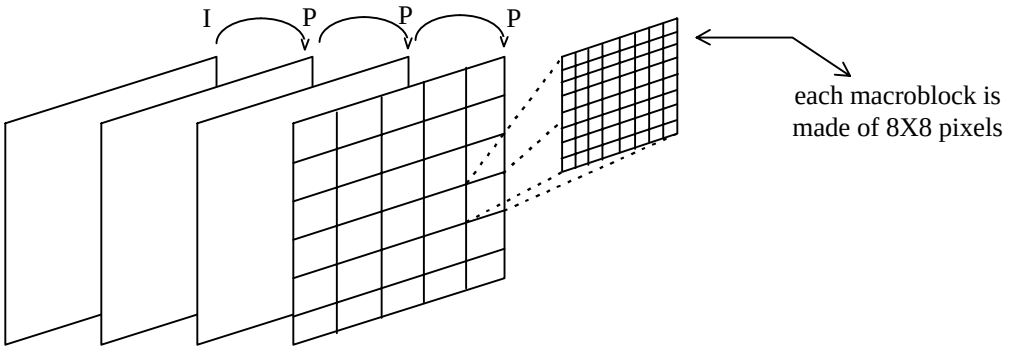

Figure 4.1 Representation of H.261 Video Coding Sequence

4.3 Moving Picture Experts Group

The Moving Picture Experts Group (MPEG) formed under the ISO subcommittee 2 in 1988. The charge of the group was to develop a video coding standard for digital storage media and bit rates at up to about 1.5 Mbps [SCHÄ95]. The MPEG-1 (ISO 11172) standard was released in 1992. Today it is common to see MPEG-1 coding used in video

clips available on the Internet and CD-ROM applications. Features include random access of video, fast-forward/fast-rewind searches, reverse playback, and editing capability.

The MPEG-1 interframe coding method is similar to the H.261 coding scheme. There are I-pictures which are coded independently of any other frame. The I-pictures can serve as reference points for arbitrary access or "step" points for fast-forward and fast-rewind. The P-pictures are predicted from previous I- or P-pictures. The B-pictures are bidirectionally predicted or bidirectionally interpolated from past or future I- or P-pictures.

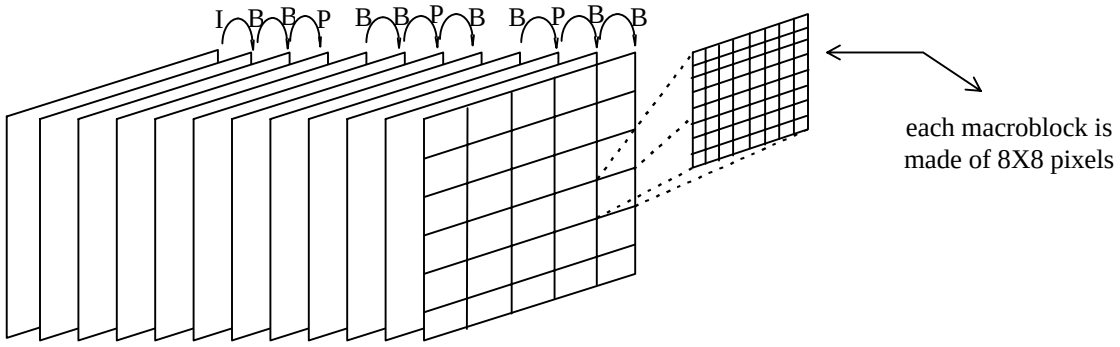

Figure 4.2 Representation of MPEG Video Coding Sequence

4.4 Adaptive Video Compression

Adaptive video compression techniques differ from previously mentioned techniques in that adaptive methods adjust their output based on changing network load conditions. In other words, when the network is congested, the encoding algorithm will decrease its output but when congestion abates, the algorithm will increase its output.

4.4.1 Methods Applicable To Any Algorithm

Gupta and Williamson suggest the following approaches to adaptively compress video [GUPT95].

- **Window Size** - The encoding algorithm can vary the horizontal and vertical picture dimensions. For example, a factor of four load decrease is obtained if the picture size is reduced from 640x480 pixels to 320x240 pixels. The end user at the destination, however, may become annoyed if such spatial variations are frequent.
- **Frame Rate** - The algorithm can decrease the frame rate by lowering the rate that it releases frames into the network. Full-motion video at 30 fps can be scaled down to 24 fps, which is movie quality or further, to 15 fps, for a factor of 2 decrease in bandwidth.
- **Color Depth** - The compression algorithm can decrease the number of colors, e.g., 16.7 million (24-bit) to 256 (8-bit) producing a factor of three decrease.
- **Coding Algorithm** - This approach differs from the methods above in that it involves not just one, but multiple algorithms. The source could progressively switch from using no compression, to algorithms with low, medium, and high compression ratios. This approach increases decoder complexity at the destination.

4.4.2 Methods Applicable To MPEG

- **Group of Pictures Pattern and Size** - The MPEG-1 algorithm encodes video according to a periodic sequence of frames. This repeated pattern is called Group of Pictures (GOP). The GOP size as well as GOP pattern can be dynamically adjusted. For example, let a given GOP pattern be the set of ten I's. This sequence produces the best resolution and random access capability but achieves the lowest compression. The algorithm can replace the second through tenth I-pictures with B-pictures. The resulting GOP sequence (IBBBBBBBBB) sends reference I-pictures every ten frames. Further compression can be achieved if the GOP size is increased, i.e., reference I-pictures are sent less often. In the example above, the GOP size could increase from 10 to 15 to decrease the encoder's output bandwidth.

- **Selective Discard of B- and/or P-pictures** - There are two locations where selective discarding can be implemented: at the network or at the source. The network approach requires network intelligence. The network must be able to identify cells belonging to B- and P-pictures and then mark these as less important than cells belonging to I-pictures. For ATM networks, this means marking these cells with CLP=1. Also, discarding at the network does not alleviate congestion on the link between the source and the network ingress. A more effective approach would be to have the source discard the entire MPEG picture. In this research, the latter approach is chosen. Entire MPEG frames are discarded at the source. Measurements are made of the effectiveness of dropping B-frames only and P- and B-frames.

4.5 Summary

This chapter discussed video compression standards and adaptive video compression methods. Common video compression standards are Motion-JPEG, H.261, and MPEG-1. Motion-JPEG is generated from a sequence of JPEG compressed images. Because Motion-JPEG uses intraframe coding but not interframe coding, high compression ratios (low bit rates) are not possible. The H.261 standard, developed for low bit rate transmission over ISDN, uses intraframe and interframe coding techniques. Since H.261 P-pictures depend only previous pictures, the standard is well-suited for real-time video transmission. Similar to H.261, MPEG-1 uses intraframe and interframe coding techniques. MPEG-1, however, interpolates bidirectionally. MPEG-1 predicts B-pictures based upon either previous or subsequent pictures. Therefore, unless the GOP pattern is set to encode I- and P-pictures or I-pictures only, MPEG-1 is not suitable for real-time video transport.

Adaptive video compression techniques adjust their output rates according to changing network load conditions. In other words, when the network is congested, the encoding algorithm will decrease its output but when congestion abates, the algorithm will increase its output. Generic methods which might be applied to any compression algorithm are: change the window size, alter the frame rate, vary the color depth, or alternate between

coding algorithms. Methods specific to MPEG-1 are vary the GOP pattern and/or size or selectively discard B- and/or P-pictures.

This chapter ends the discussion of background material. Chapter 5 describes the experimental design and implementation.

Chapter 5. Experimental Design and Implementation

5.1 Experimental Setup

The goal of the simulation experiment is to evaluate the capability of the ATM ABR service class to carry video by using various source control methods. A VoD network transporting MPEG-1 video is built using two geographically separated switches. Figure 5.1 shows the representative network. Video servers are connected to Switch A via multiple 25 Mbps links of an arbitrarily short length. An alternative design using a single video server accessing a video archive is possible, but would introduce issues (e.g., multicasting) that are beyond the scope of this study.

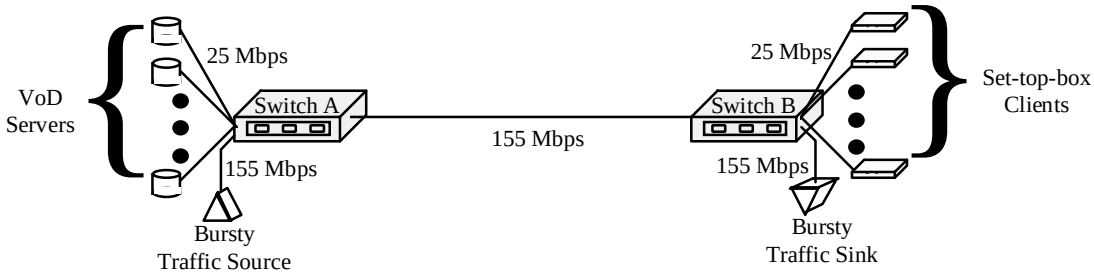

Figure 5.1 Model of Simulated Network

The capacity of the link between the switches is 155 Mbps. The distance between the switches is arbitrarily long. Also, at Switch A, non-video traffic is introduced to the network. A 155-Mbps link connects this bursty traffic source to Switch A. This traffic flow is intended to compete with the video traffic flows.

Switch B directly connects to customer set-top-boxes (clients) via 25-Mbps lines. An alternate approach is to connect to Switch B a local server, which would be responsible for distributing video to the set-top-boxes. Also, a 155-Mbps link connects a sink, which removes the added non-video traffic, to Switch B.

5.2 Response Variables

The response variables used to evaluate the performance of the model are: end-to-end delay, queue occupancy, utilization, and jitter (cell delay variation).

- End-to-end delay is measured. This variable should be minimized for video traffic.
- Queue occupancy is recorded. Small queues are best since queue size affects cell delay and switch cost.
- Utilization is measured since it limits the number of simultaneous connections. A low utilization is best for a fixed number of sources, because it would enable more connections to be added to fill the excess capacity.
- Jitter is important because it affects how the customer perceives the quality of the service. The smaller the jitter, the better.

5.3 Experimental Factors

The simulation parameters varied for this study are: switch queue size, source control method, and seed for the simulation's random number generator.

- Queue size affects delay and cell loss characteristics, as well as the cost of the switch. Queue length is directly proportional to cell delay and switch cost. However, it is inversely proportional to cell loss.
- Source control methods are from none at all, MPEG B-picture discard, and MPEG P- and B-picture discard.
- The seed for the simulation's random number generator is varied to ensure the stability of the network model.

5.4 Model Description

The software environment used for this research is the Optimized Network Engineering Tools Modeler (OPNET Modeler) developed by MIL 3, Incorporated [MIL397]. The software package is used for modeling, simulating, and analyzing the VoD network.

OPNET provides editors in which users can specify models at the network, node, and process levels [MIL396]. A hierarchical description of the network follows.

5.4.1 Network Level

The VoD network consists of ten source/destination pairs (0 to 9), switches at the network ingress and egress, and a bursty source/destination pair, as shown in Figure 5.2. Each source represents a video server consisting of an ATM cell generator node (denoted CELL_GENn) and a traffic management node (denoted TRAF_MGR_SRCn). Each CELL_GENn creates and forwards cells to its respective TRAF_MGR_SRCn via a duplex link.

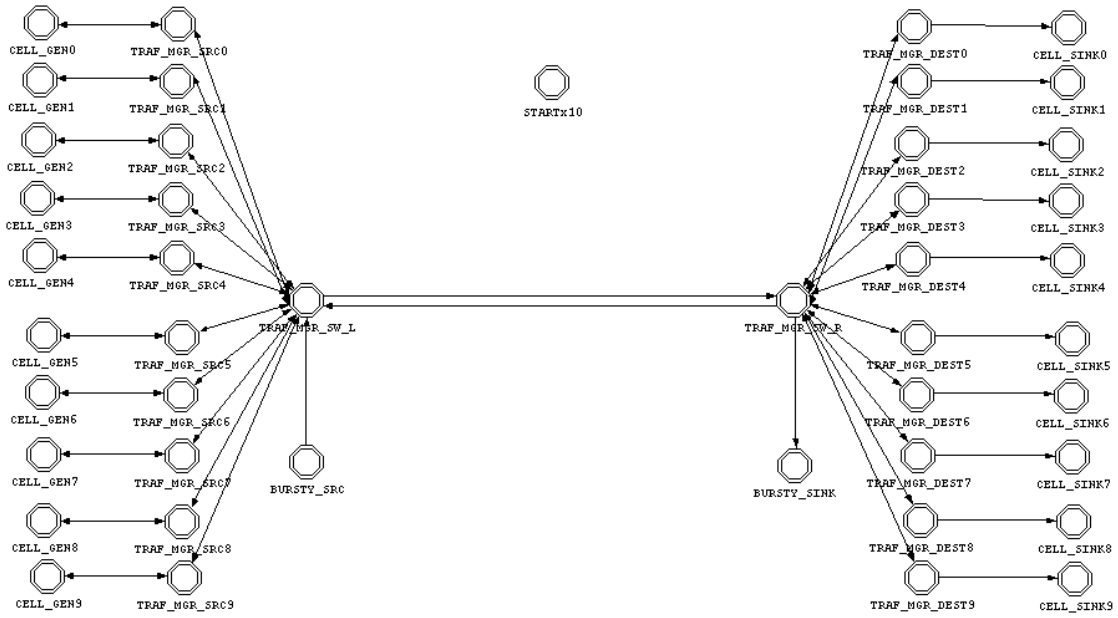

Figure 5.2 Network Model

Each source is connected to the ingress switch node by a duplex link. The ingress switch node, named TRAF_MGR_SW_L, is connected to the egress switch node, named TRAF_MGR_SW_R, via two simplex links. Two simplex links are used instead of a single duplex link so that forward traffic link utilization is more clearly measured.

Data cells and Forward Resource Management (FRM) cells from all connections are multiplexed onto the forward simplex link in the direction from the ingress switch node to the egress switch node.

Each destination is connected to the TRAF_MGR_SW_R via a duplex link. Each destination consists of a traffic management node, denoted TRAF_MGR_DESTn, and a node for discarding cells, denoted CELL_SINKn. Each TRAF_MGR_DESTn node forwards data cells to its respective CELL_SINKn node along a simplex link. Each CELL_SINKn node discards its data cells. Each TRAF_MGR_DESTn turns around FRM cells and sends these cells as Backward Resource Management (BRM) cells in the reverse direction, through TRAF_MGR_SW_R and TRAF_MGR_SW_L, back to the TRAF_MGR_SRCn.

5.4.2 Node Level

The node level describes models consisting of processor, queue, transceiver, and packet stream modules. These models are then instantiated in the network model. Processor and queue modules permit user-specified behavior using process models, which are discussed later. Queues, additionally have special queue and subqueue functions.

5.4.2.1 STARTx10 Node Model

The STARTx10 node model begins the simulation. This node model contains a single node. Figure 5.3 shows this model. Details of the operation of this node are in the STARTx10 process model description.

Figure 5.3 Node Model of STARTx10

5.4.2.2 CELL_GENn Node Model

The CELL_GENn node contains a CELL_GENn processor module, which sends cells to a point-to-point transmitter module through a packet stream. See Figure 5.4.

Figure 5.4 Node Model of Cell Generator

5.4.2.3 TRAF_MGR_SRCn Node Model

The TRAF_MGR_SRCn node, depicted in Figure 5.5, has a DATA_RCV receiver module which feeds a TRAF_MGR_SRCn processor module. This processor forwards data cells and FRM cells to the SRC_TX transmitter module via a packet stream. This processor also receives BRM cells via a packet stream from BRM_RCV.

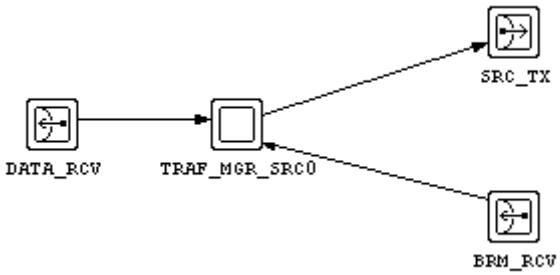

Figure 5.5 Node Model for TRAF_MGR_SRCn

5.4.2.4 BURSTY_SRC Node Model

Figure 5.6 shows the BURSTY_SRC node model, which contains an ideal generator module and a point-to-point transmitter module. A packet stream transfers randomly generated cells from the ideal generator to the transmitter.

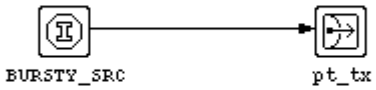

Figure 5.6 Node Model of BURSTY_SRC

5.4.2.5 BURSTY_SINK Node Model

Figure 5.7 shows the BURSTY_SINK node model, which contains a point-to-point receiver module and a BURSTY_SINK processor module. A packet stream transfers randomly received cells from the receiver to the BURSTY_SINK processor.

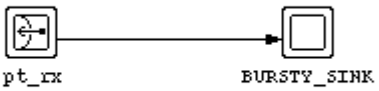

Figure 5.7 Node Model of BURSTY_SINK

5.4.2.6 TRAF_MGR_SW_L Node Model

The TRAF_MGR_SW_L switch node, shown in Figure 5.8, consists of the following modules. The CELL_RCVn receiver and BRM_TXn transmitter pairs, which are logically linked to each other, are connected via packet streams to the TRAF_MGR_SW_L queue module. Logically linking a transmitter and receiver allows duplex links to more readily connect to predetermined logically linked transceiver pairs. Each transceiver pair connects to TRAF_MGR_SRCn node.

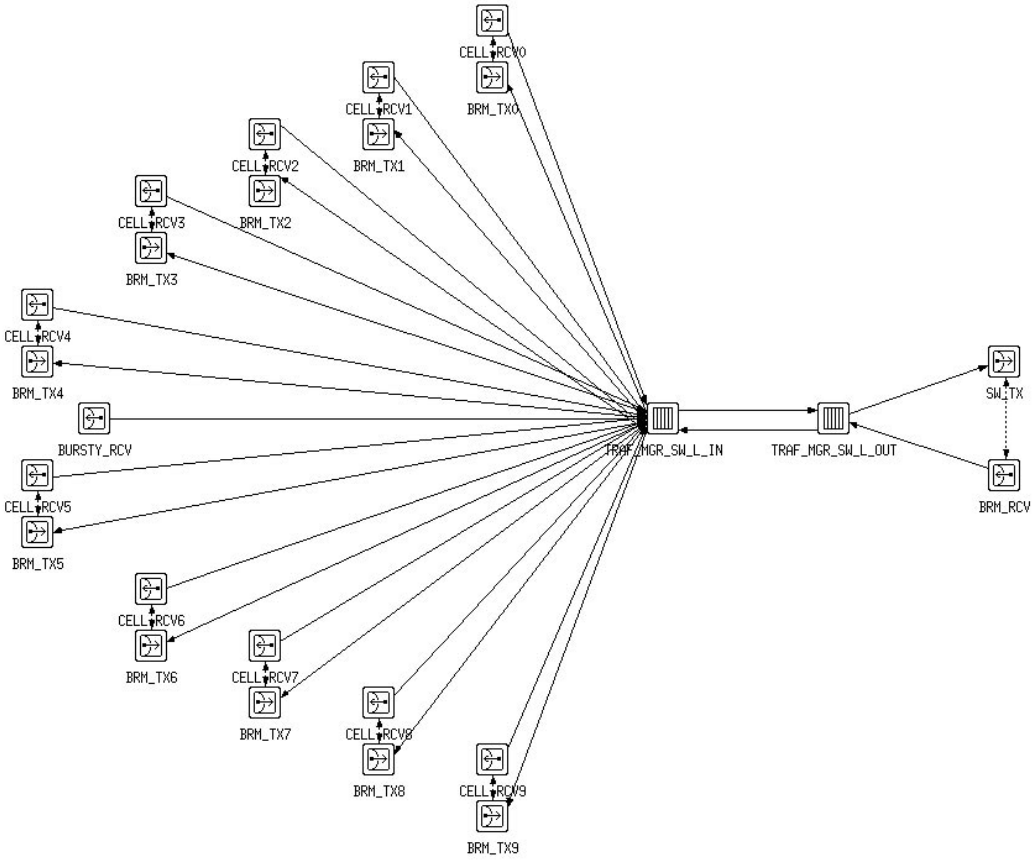

Figure 5.8 Node Model of TRAF_MGR_SW_L

Data cells and FRM cells enter the TRAF_MGR_SW_L switch node via the CELL_RCVn receivers. BRM cells travel the opposite direction and exit via the BRM_TXn transmitters. Cells from the BURSTY_SRC node are received by way of BURSTY_RCV.

The TRAF_MGR_SW_L_IN queue module statistically multiplexes data cells and FRM cells onto a single packet stream and forwards the cells to the TRAF_MGR_SW_L_OUT queue module. The service rate for the TRAF_MGR_SW_L_IN queue module is sufficiently high (10 Gbps or 23,584,490 cells per second), so that there is no congestion at this queue. Congestion for the TRAF_MGR_SW_L node occurs at the

TRAF_MGR_SW_L_OUT queue module. The service rate of this module is set to match the line speed of the out-going simplex link. More details are given in Chapter 6.

A single SW_TX transmitter and BRM_RCV receiver pair is connected by way of packet streams to the TRAF_MGR_SW_L_OUT queue module. This transceiver pair interfaces the network. Data cells and FRM cells are transmitted to the network and BRM cells are received from the network.

5.4.2.7 TRAF_MGR_SW_R Node Model

The TRAF_MGR_SW_R switch node, illustrated in Figure 5.9, consists of the following modules. A single SW_RCV receiver and BRM_TX transmitter pair is connected by way of packet streams to the TRAF_MGR_SW_R queue module. This transceiver pair interfaces the network. Data cells and FRM cells are received from the network and BRM cells are transmitted to the network.

BRM_RCVn receiver and CELL_TXn transmitter pairs, which are logically linked to each other, are connected via packet streams to the TRAF_MGR_SW_R queue module. Each transceiver pair connects to a TRAF_MGR_DESTn node. Data cells and FRM cells exit the TRAF_MGR_SW_R switch node via these CELL_TXn transmitters and BRM cells enter via the BRM_RCVn receivers.

There is also a BURSTY_TX module that forwards cells, generated by the BURSTY_SRC node, to the BURSTY_SINK node.

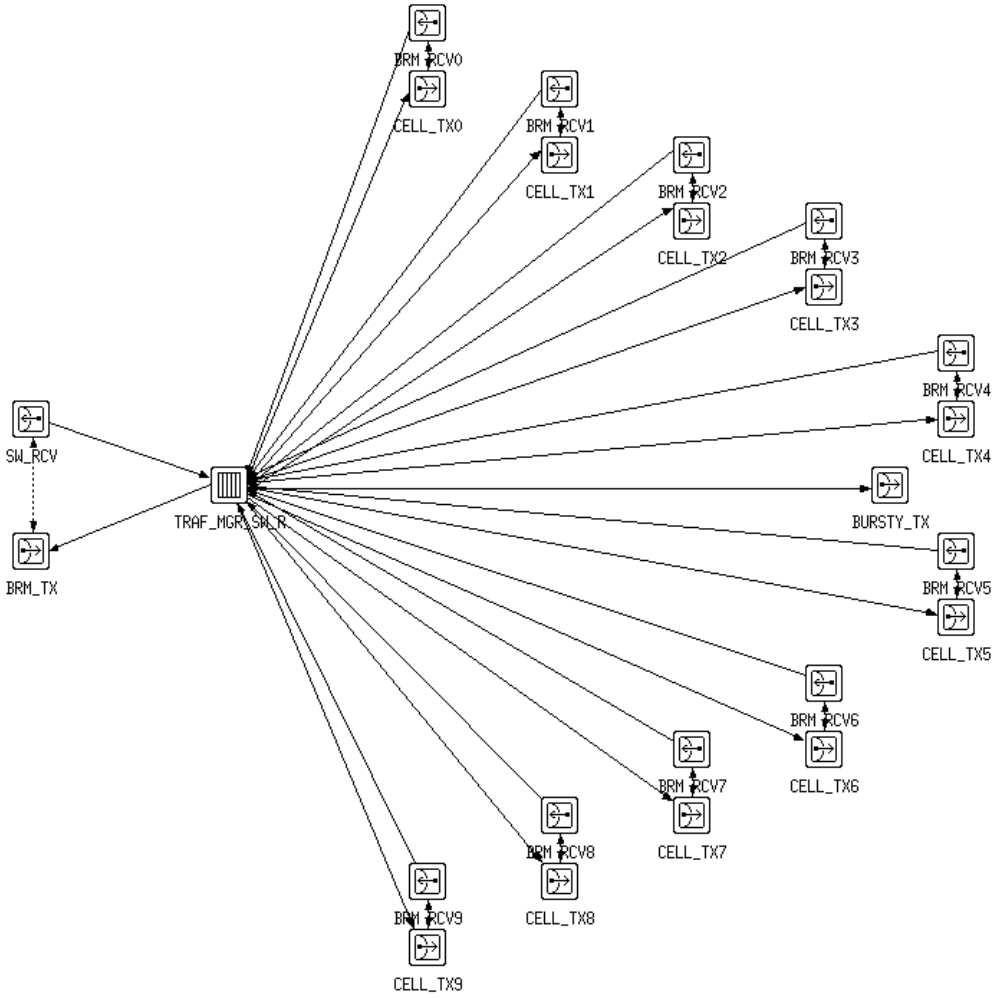

Figure 5.9 Node Model of TRAF_MGR_SW_R

5.4.2.8 TRAF_MGR_DESTn Node Model

The TRAF_MGR_DESTn node model, depicted in Figure 5.10, receives data cells and FRM cells by way of a CELL_RCV receiver module. These cells are sent along a packet stream to the TRAF_MGR_DEST processor module. The TRAF_MGR_DEST processor module forwards cells through a packet stream to the DATA_CELL_TX transmitter module. The TRAF_MGR_DEST processor module turns around FRM cells, and transmits BRM cells via a packet stream to the BRM_TX transmitter module.

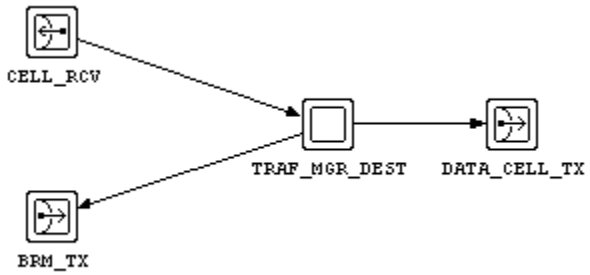

Figure 5.10 Node Model of TRAF_MGR_DEST

5.4.2.9 CELL_SINKn Node Model

The CELL_SINKn node in Figure 5.11 contains a point-to-point receiver which sends packets to the CELL_SINKn.

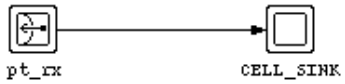

Figure 5.11 Node Model of CELL_SINKn

5.4.3 Process Model Level

Process models specify the behavior of node level processors and queues. Process models are comprised of graphical state-transition diagrams which define the execution flow of a process in response to simulated discrete events. At each state, generic behavior is specified by the user utilizing Protocol C. Protocol C or Proto-C is essentially the C programming language with OPNET Simulation Kernel procedures added [MIL396].

The following is true for all process models. The dark (green on color medium) state represents a Forced State. A Forced State allows both enter executives and exit executives to execute without waiting. Control then transitions to the next state.

A light (red on color medium) state denotes an Unforced State. An Unforced State pauses after executing its enter executives and blocks or waits for an interrupt to occur (i.e., goes to sleep). When the process model is invoked again, the process model receives an interrupt (i.e., awakens), and execution resumes at the exit executives of the Unforced State.

Arrows indicate transitions. Arrows with solid lines represent unconditional transitions. This means that immediately following the execution of the exit executives, control will flow to the state pointed to by the arrowhead. Arrows with dashed lines denote conditional transitions. For this type of transition, control flows to the state pointed to by the arrowhead after the completion of the exit executives as with unconditional transitions. However, an additional criterion must be true. A Boolean expression must evaluate to TRUE. This expression appears in parentheses close to the corresponding dashed-line arrow.

5.4.3.1 STARTx10 Process Model

Simulation starts when the STARTx10 process model, shown in Figure 5.12, receives a begin simulation interrupt. At the START_SIM state, remote interrupts, which invoke each TRAF_MGR_SRCn process model, are randomly scheduled. The random time is generated from a uniform distribution between 0 and 1/25 seconds.

The TRAF_MGR_SRCn processes are invoked first to ensure that the first cell is an FRM cell. After scheduling all interrupts, execution flows to the END_SIM state.

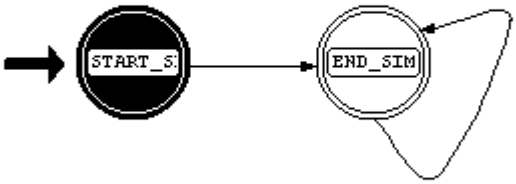

Figure 5.12 STARTx10 Process Model

5.4.3.2 TRAF_MGR_SRCn Process Model

The TRAF_MGR_SRCn process model is the most important and the most complex of the process models. The state machine is shown in Figure 5.13. This process model implements the source end system behavior that is defined in [ATMF96].

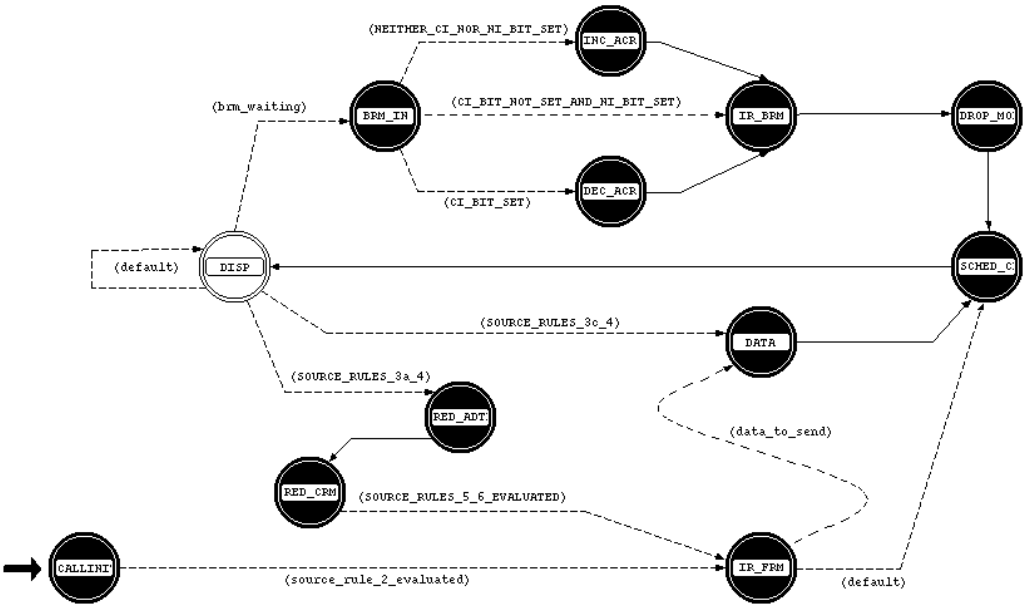

Figure 5.13 Process Model of TRAF_MGR_SRCn

Execution begins at the CALLINIT&FIRSTDATA state. At this state, the Peak Cell Rate (PCR) is retrieved since it is set during simulation time in this research. The Initial Cell Rate (ICR) is set equal to PCR. The parameter representing the lowest frequency of FRM cells (Nrm) is set to its default value of 32. This means that every 32^{nd} cell will be an FRM cell.

The Rate Decrease Factor (RDF) is set to 1. The Minimum Cell Rate (MCR) is initialized to 0 cells per second. The Rate Increase Factor (RIF) is set to 1. Setting RIF = 1 effectively removes the ability of the source to gradually increase its rate when there is no congestion (NI = 0 and CI = 0). Instead, the increase is instant. As indicated in Table 5.1, for NI = CI = 0, if RIF = 1, ACR will never get the value of ACR + RIF*PCR since this quantity is never less than ER or PCR.

Table 5.1 NI and CI Bit Settings and Their Associated Actions

NI	CI	Action
0	0	ACR <= Min(ER, ACR + RIF*PCR, PCR)
0 or 1	1	ACR <= Min(ER, ACR - ACR*RDF)
1	0	ACR <= Min(ER, ACR)

The minimum number of data cells that must be sent between RM cells or the maximum frequency at which RM cells may be sent (Mrm) is set to 2 cells. The ACR Decrease Time Factor (ADTF) is set to 10 ms. Trm is set to 100 ms. The fixed round trip time (FRTT) is set maximally to 16.7 seconds. The Transient Buffer Exposure (TBE) is set to 16,777,215 cells. The count of missing RM cells (Crm) is set to:

$$Crm = ceil[TBE/Nrm] \qquad \text{Equation 5.1}$$

The Cut-off Decrease Factor (CDF) is set to 1. The Tagged Cell Rate (TCR) is initialized to 10 cells. The following Boolean variables are initialized to FALSE: data_in_queue,

ir_frwd_rm_cell_sent, ir_bkwd_rm_cell_sent, data_cell_sent, brm_waiting, and data_to_send. The parameter num_ir_cells_sent_since_last_ir_frm is set to 0.

Also, at the INIT state, the first RM cell is created. The Allowed Cell Rate (ACR) is set equal to ICR according to Source Behavior 2. After setting the fields PTI = RM_CELL_PATTERN, DIR = FORWARD, CCR = ACR, CLP = IN_RATE, and ER = PCR (according to Source Rules 7 and 10), execution transitions to the IR_FRM state, if the source_rule_2_evaluated Boolean variable is TRUE.

At the IR_FRM state, if source_rule_2_evaluated is FALSE, a new RM cell is created and CCR, PTI, DIR, CLP, and ER are set like in the initial state. Next, the RM cell is transmitted on to CELL_STRM_OUT. The time_of_last_ir_frm is recorded. For more detailed descriptions of these parameters see [ATMF96]. The next state is the DATA state if data_to_send is TRUE. This transition prevents the gradual build-up of data cells whenever it is time to send an FRM but there is also a data cell waiting to be sent. Otherwise, when data_to_send is FALSE, the next state is SCHED_CELL.

The next state is SCHED_CELL. This state is an empty forced state which immediately transitions to the DISP (dispatch) state. At this state, a condition is tested to see if the first RM cell was just sent. If this is the case, this process model sends a remote interrupt to invoke the CELL_GENn processor to begin sending data cells. (Recall that Source Rule 2 specifies that at the beginning of a connection, the first cell must be an in-rate FRM). Next, the DISP state pauses, waiting for an interrupt. If a stream interrupt occurs (indicating a new cell arrival), a determination is made as to which input has the new cell. If the stream is BRM_CELL_STRM_IN, a Boolean flag brm_waiting is set to TRUE. If the stream is DATA_CELL_STRM_IN, a Boolean flag data_to_send is set to TRUE. Lastly, a variable time_since_last_ir_frm stores the elapsed time since the most recent FRM cell.

Transition from the DISP state can take one of three routes. If there is a BRM cell waiting (brm_waiting = TRUE), the next state is BRM_IN. However, the next state is DATA if the following compound condition is TRUE.

(data_to_send && (!brm_waiting) &&
(num_ir_cells_sent_since_last_ir_frm < (Nrm - 1)) &&
((num_ir_cells_sent_since_last_ir_frm < Mrm) || (time_since_last_ir_frm < Trm)))

In words, first there must be a data cell awaiting transmission. Additionally, a BRM should not be waiting and the number of data cells sent since the last FRM cell must be less than Nrm - 1. Also, either the number of data cells sent since last FRM has to be less than Mrm or the time since the last FRM must be less than Trm. These conditions are according to Source Rules 3c and 4.

Finally, transition from the DISP state to the RED_ADTF occurs if the following compound condition is TRUE.

(!brm_waiting && ((num_ir_cells_sent_since_last_ir_frm >= (Nrm - 1)) ||
((num_ir_cells_sent_since_last_ir_frm >= Mrm) && (time_since_last_ir_frm >= Trm))))

In words, a BRM must not be waiting. Also, either the number of data cells sent since the last FRM must be not be less than Nrm - 1 or both the number of data cells sent since the last FRM is greater than or equal to Mrm and the time since the last FRM is greater than or equal to Trm.

At the BRM_IN state, the waiting BRM is acquired from BRM_CELL_STRM_IN. The brm_waiting flag is reset to FALSE. Values are read from the BRM cell's CI, NI, and ER fields. As shown in Table 5.1, if the CI bit indicates congestion, (set to 1), the next state is DEC_ACR. If both the CI bit and the NI bit equal zero, the next state is INC_ACR. If the CI bit is zero and the NI bit is set to 1, meaning that there is mild congestion, the next

state is IR_BRM. Lastly, if the CI bit equals 1, the next state is DEC_ACR, regardless of the value of NI.

At the INC_ACR state, ACR is increased by the quantity RIF times PCR. The RIF parameter is used to prevent abrupt, negative changes to the network. The incremental increase in the load decreases the likelihood of the sudden creation of queues. Afterwards, to ensure that the new ACR conforms to PCR of the contract, ACR is set to the minimum of ACR and PCR. Execution then flows to the IR_BRM state.

At the DEC_ACR state, ACR is decreased by the value ACR times RDF. To make sure that the new ACR meets the contracted specifications of the MCR, ACR is set to the maximum of ACR and MCR.

At the state IR_BRM, the ER is obtained from the ER field of the current BRM cell. The ACR is set to the minimum of ACR and ER. The ACR is then set to the maximum of ACR and MCR. The BRM cell is destroyed. At the DATA state, the data cell is acquired from DATA_CELL_STRM_IN. The data_to_send flag is reset to FALSE. The data cell is transmitted on the output stream CELL_STREAM_OUT.

See Source Rules 8 and 9 in [ATMF96] for details on the behavior of the BRM_IN, INC_ACR, DEC_ACR, and IR_BRM states.

At the DROP_MODE state, a special packet called the “INFO_PACKET” is created. The purpose of the INFO_PACKET is to indicate the congestion status of the network to the CELL_GENn process model. Network congestion status is determined by the CI and NI fields of the most recent BRM cell. Switches set CI when there is heavy congestion and set NI when there is light congestion.

The status that the INFO_PACKET indicates depends upon the simulation’s mode of operation. The modes are: drop mode off, drop B-pictures only, and drop both B- and P-

pictures. If the current simulation is to examine network behavior with the drop mode off, every INFO_PACKET will indicate that the drop mode is off.

If the current simulation is to evaluate network performance with dropping only B-pictures under heavy congestion, the status is set to drop B-pictures when the CI bit is set. When neither the CI bit nor the NI bit is set, the status indicates that the drop mode is off. Also, if CI is not set and NI is set (indicating light congestion), the status is not explicitly set. The status remains unchanged from its default value, which is drop mode off.

Next, if the current simulation is to view network response to dropping both B- and P-pictures under heavy congestion, the status is set to drop B- and P-pictures when the CI bit is set. When neither the CI bit nor the NI bit is set, the status indicates that the drop mode is off. Also, if CI is not set and NI is set (indicating light congestion), the status is not explicitly set. The status remains unchanged from its default value, which is drop mode off.

Lastly, the BRM cell is destroyed and the INFO_PACKET is sent to CELL_GENn.

The RED_ADTF state implements Source Behavior 5. This behavior is significant in that it address the issue of ACR retention. A source could be assigned a high rate when the network is lightly loaded and then retain this rate during a time when the network is heavily loaded. Thus, this connection would have an unfair share of the bandwidth. After several proposed solutions, the ATM Forum agreed to a simple “Use It or Lose It” (UILI) policy. If the time between consecutive FRM cells is greater than the ADTF and ACR is greater than ICR, the ACR for the connection is reduced to ICR.

The RED_CRM state implements Source Rule 6. If the number of unacknowledged FRM cells sent, i.e., outstanding FRM cells which have been transmitted without corresponding BRM cells having been received, is greater than Crm, then ACR is reduced by ACR times CDF. The ACR is then set to the maximum of ACR and MCR.

5.4.3.3 CELL_GENn Process Model

Figure 5.14 illustrates the state transition diagram for the CELL_GENn process model.

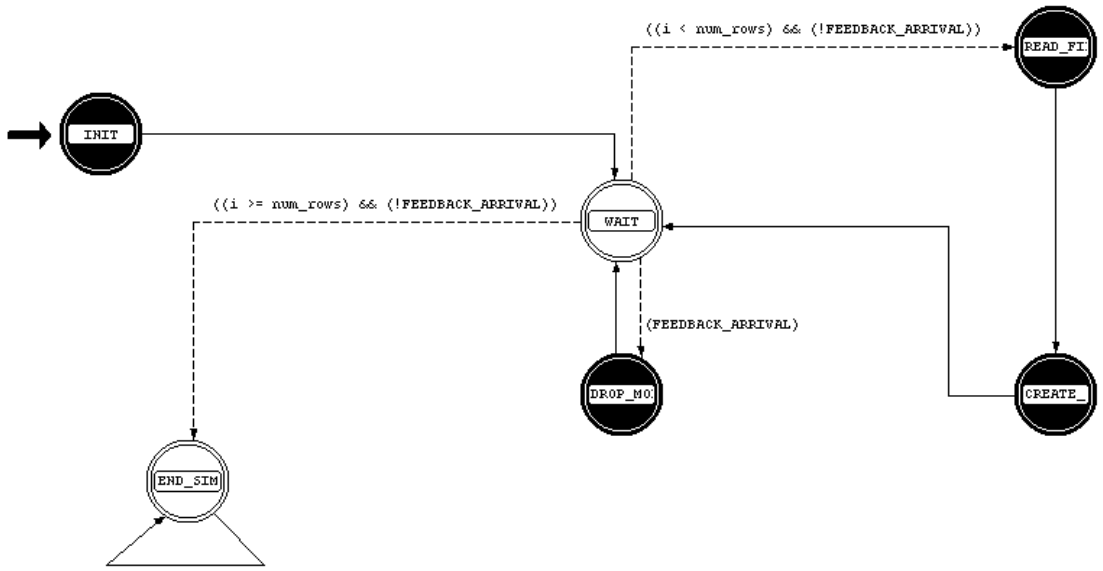

Figure 5.14 Process Model of CELL_GENn

Execution begins at the INIT state. At this state, the trace data file is opened and read. The number of bits per MPEG frame is stored in a variable. The total number of 48-byte payloads required to transport this MPEG picture is calculated. The time between each cell is computed according to:

inter_cell_delay = (1/frame_rate) * (1/num_payloads), Equation 5.2

where the frame_rate is defined to be 25 frames per second. This computation ensures that the entire frame is transmitted in 1/25 seconds and the cells belonging to this frame are evenly spaced. Next, the ATM cells for the first MPEG picture are generated. For each cell, the PTI field is set to DATA_CELL_PATTERN and CLP is set to IN_RATE. The cells are then transmitted. A self-interrupt is scheduled to occur after 1/frame_rate seconds.

Control then unconditionally flows to the WAIT state. The enter executives and exit executives are empty at this state. When this process model receives an interrupt, this state is exited. If there are more MPEG frames (i < num_rows) and an INFO_PACKET has not arrived (!FEEDBACK_ARRIVAL), control flows to the READ_FILE state. If an INFO_PACKET has arrived (FEEDBACK_ARRIVAL), execution immediately flows to the DROP_MODE state. If there are no more MPEG frames (i >= num_rows), execution flows to the END_SIM state.

At the DROP_MODE state, the first action is to obtain the INFO_PACKET from the input stream. Next, the status is read. Boolean variables drop_mode_on, drop_b_pictures, and drop_b_and_p_pictures are set or reset accordingly.

At the READ_FILE state, the next MPEG picture is read from the trace data file. The picture type is then determined via modular arithmetic on a sequence number maintained in this state. If either of the drop modes is on and the current picture type matches the type that should be dropped, the drop_mode_on Boolean variable is set TRUE. If neither of the drop modes is on, then just as in the INIT state, the total number of 48-byte payloads is computed, and the inter-cell delay is calculated.

Execution continues unconditionally to the CREATE_CELL state. If the drop_mode_on Boolean variable is true, a counter for the number of dropped frames is incremented. Otherwise, if the drop_mode_on variable is false, ATM cells are created, just as in the INIT state. Each cell has its PTI and CLP fields set and is transmitted. A self-interrupt is scheduled to occur after 1/frame_rate seconds. Control transitions to the WAIT state.

After every MPEG frame has been read from the trace data file (i >= num_rows), the END_SIM state is entered. No more cells are created and sent. At this state, memory is deallocated.

5.4.3.4 ACP_FIFO Process Model

All queue modules used in this research incorporate a variation of the basic acp_fifo queue process model built into OPNET. The basic process model is shown in Figure 5.15.

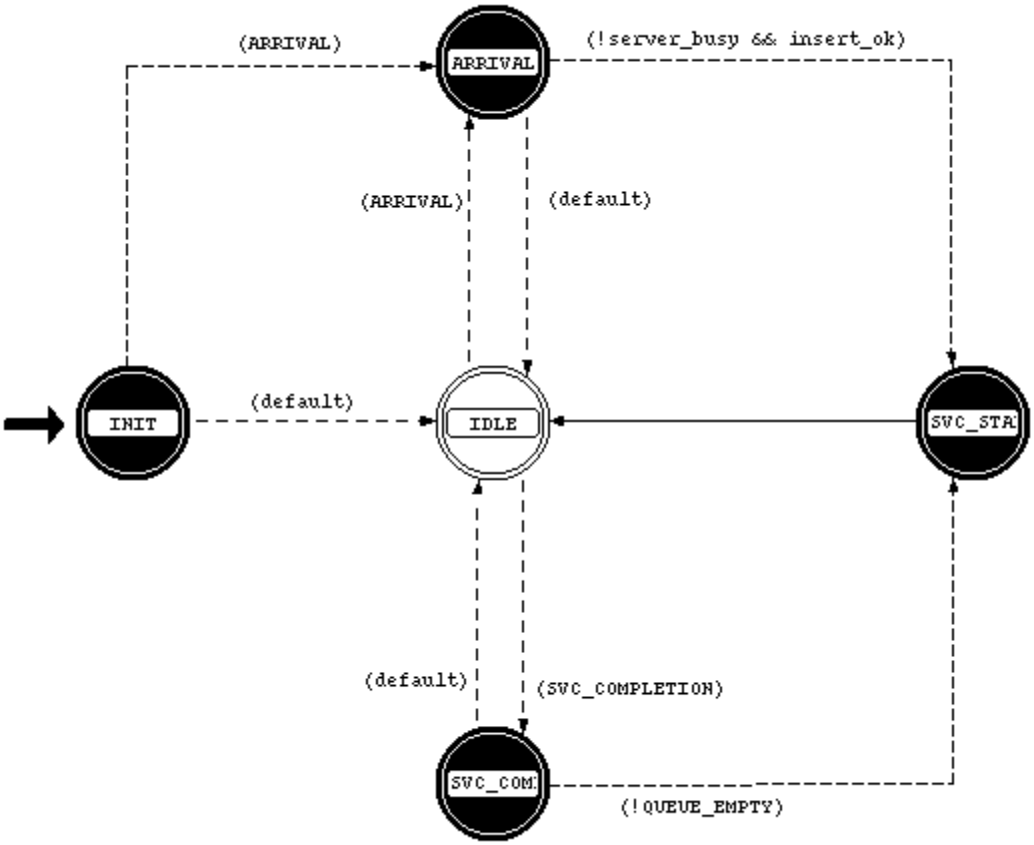

Figure 5.15 Process Model of ACP_FIFO

The letters ACP_FIFO stand for Active, Concentrating, Packet-oriented, First In, First Out. Active implies that the model autonomously dequeues and transmits cells according to built-in criteria such as a service rate. Requests for cells from other modules are not needed since the service mechanism is built-in. Concentrating means that all incoming cells are merged onto a single output stream. Packet-oriented means that the queue service time is based on packets instead of being based on the number of bits. This

distinction is not as important with ATM's constant 53-byte cell size as it would be if the packet size were variable. The First In, First Out discipline means that cells exit the queue in the same order they enter the queue.

The ACP_FIFO_SW_L_IN, ACP_FIFO_SW_L_OUT, and ACP_FIFO_SW_R process models define the behavior for the TRAF_MGR_SW_L_IN, TRAF_MGR_SW_L_OUT, and TRAF_MGR_SW_R queue modules, respectively. For the following description, recall that TRAF_MGR_SW_L is the INGRESS switch and TRAF_MGR_SW_R is the EGRESS switch.

The process model begins at the INIT state. First a server_busy flag is initialized to idle. Next, the service_rate and SWITCH_TYPE attributes, which are set at simulation runtime, are obtained. If there is a cell arrival, the next state will be ARRIVAL, otherwise by default, the next state will be IDLE.

At the ARRIVAL state, the cell is obtained. The PTI field is read from the cell to determine the cell type. In the ACP_FIFO_SW_L_IN process model, if the cell type is DATA_CELL, the cell's VCI field is set to its arrival_stream. Also, in the ACP_FIFO_SW_L_IN process model, if the cell type is RM_CELL and if the cell's DIR field indicates that the cell is a FORWARD RM cell, the cell's VCI field is set to its arrival_stream. The cell is enqueued and statistics are gathered.

From the ARRIVAL state, if the server is available, the next state is SVC_START, Otherwise, the next state is IDLE.

At the SVC_START state, cell service time is set to the reciprocal of the service_rate. A self interrupt is scheduled to occur at the end of the service time and the server_busy flag is set to busy. From the SVC_START state, there is an unconditional transition to the IDLE state.

The IDLE state has neither enter executives nor exit executives. From the IDLE, control flows to the ARRIVAL state if there is a cell arrival. Control flows to the SVC_COMPL state if a self interrupt indicating the end of service time has occurred.

At the SVC_COMPL state, the cell is dequeued and its VPI, VCI, and PTI fields are read. The following statements are true for the INGRESS switch. In the ACP_FIFO_SW_L_IN process model, if the cell is either a FRM or a DATA_CELL, the cell is sent to the TRAF_MGR_SW_L_OUT queue module. BRM cells are sent onto the output port indicated by its VCI. In the ACP_FIFO_SW_L_OUT process model, FRM and DATA_CELLS are transmitted onto the virtual path indicated by its VPI. BRM cells are sent to the TRAF_MGR_SW_L_IN queue module.

If the switch_type is EGRESS and if the cell is a FRM or a DATA_CELL, the cell is transmitted onto the output port indicated by its VCI. If the cell is a BRM, it is sent onto the virtual path indicated by its VPI. The server_busy flag is reset to idle.

From the SVC_COMPL state, control transitions to the SVC_START state, if there is a cell awaiting service, i.e., the queue is not empty.

5.4.3.5 TRAF_MGR_DEST Process Model

Figure 5.16 illustrates the TRAF_MGR_DEST process model. First, it should be noted that there is only one copy of the TRAF_MGR_DEST process model. Unlike the TRAF_MGR_SRC process model, which is replicated and then altered slightly for each connection, the TRAF_MGR_DEST process model does not exist on a “per connection” basis.

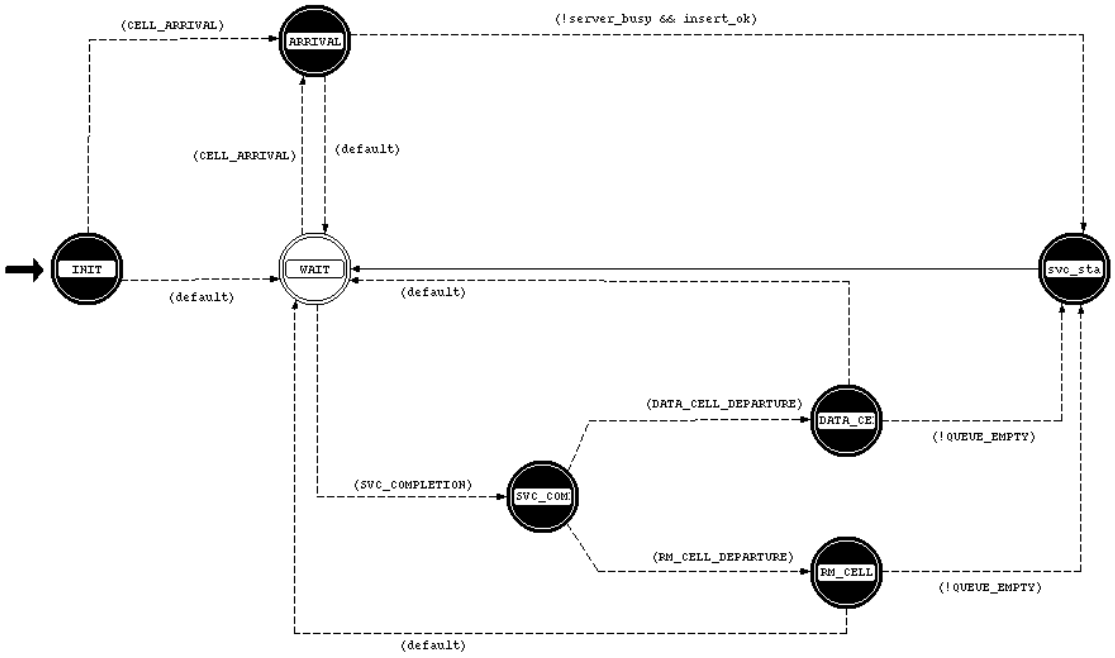

Figure 5.16 Process Model of TRAF_MGR_DEST

5.4.3.6 CELL_SINK Process Model

The CELL_SINK process model is shown in Figure 5.17. At the INIT state, the cell count variable is initialized. At the ARRIVAL state, statistics such as end-to-end delay are calculated. At the DISCARD state, a cell counter is incremented and the DATA_CELL is destroyed.

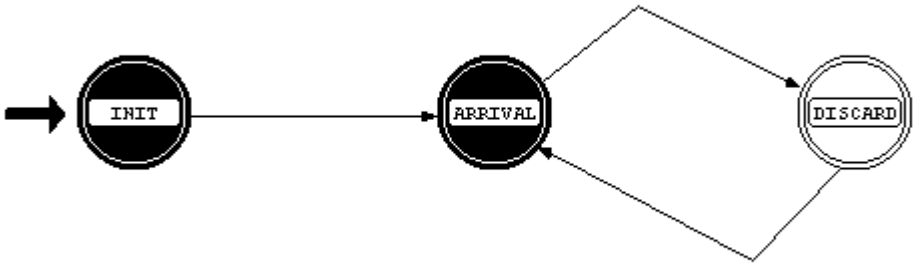

Figure 5.17 Process Model of CELL_SINK

5.5 Verification and Validation

The quality of a simulation model is measured by how closely the model approximates the real system. Model verification and model validation are two steps which are used to ensure the model's implementation and correctness. Model verification involves utilizing debugging methods which guarantee that the simulation operates as the model design specifies. Model validation involves ensuring that the assumptions made during the design of the model are reasonable.

Both verification and validation steps are necessary because they are independent concepts. For example, it is possible for a model to be properly implemented (verified) but not representative (invalid) of a real system [JAIN91].

The following sections describe the verification techniques and validity tests which were used in this research.

5.5.1 Verification Techniques

The verification techniques used in this study are listed below.

1. Top-down modular design was automatic because of OPNET's hierarchical modeling approach.

2. Examples of anti-bugging and debugging methods are: "printf" statements were inserted into various portions of code, animation was used extensively to verify network packet flows and process model state transitions, cells generated at the source were compared to those received at the destination, and OPNET's debugger was used.

3. During the model development stage, the model was explained to colleagues. Network, node, and process model-level behavior, including lines of code, were discussed.

4. Simplified cases were run. Representative "dummy" trace data was created to flood the network for short simulations and long simulations.

5.5.2 Validation Methods

Table 5.2 summarizes the validation methods used in this study.

Table 5.2 Validation Methods

	Validity Test		
Key Aspect	*Expert Intuition*	*Real System Measurements*	*Theoretical Results*
Assumptions	met with advisor often to discuss		
Input parameter values and distributions		parameter values are based on real switches	most parameter values are set to default values from [ATM96]
Output values and conclusions	no major anomalies or unexpected behavior		

5.6 Summary

This chapter discussed the experimental setup used for this research. A VoD network transporting MPEG-1 video is built using two geographically separated switches. Ten video servers and a single bursty traffic source are connected to the ingress or left-edge switch. This switch multiplexes the cells and then forwards the traffic flow to the egress or right-edge switch. This switch sends the video traffic to the set-top box clients. This switch also transmits cells created by the bursty traffic source to the bursty traffic sink.

Also, this chapter presented the response variables. End-to-end delay, queue occupancy, backbone utilization, and jitter are used to evaluate the model.

Next, the network model was described in detail. Explanations were given for each node model, queue or processor module, and each process model.

Chapter 6 presents the simulation results.

Chapter 6. Performance Results

6.1 Simulation Performance

Simulation was done to evaluate the concepts developed in this research and to demonstrate the utility of the simulation model. This chapter describes values for network parameters and compares the results.

6.2 Parameter Settings

This section gives details on the trace data sets, backbone data rate, bursty source parameters, and network queue congestion.

6.2.1 Trace Data

Two publicly available data sets, *Formula 1 Car Race at Hockenheim 1994* and *Soccer World Cup 1994 Final: Brazil - Italy* were obtained from [ROSE95]. These two video sequences were chosen because of their high mean bit rates, 0.77 Mbps and 0.63 Mbps, respectively, relative to the other available trace data.

According to [ROSE95], the video clips were encoded using the following parameter set. Each frame consists of one slice. The GOP pattern is IBBPBBPBBPBB (12 frames). The quantizer scales are 10 (I), 14 (P), and 18 (B). The motion vector search is logarithmic/simple, half pixel, 10. The reference frame is original. The encoder input is 384 x 288 pixels with 8-bit resolution. The number of frames per sequence is 40000, equivalent to 1600 seconds or about half an hour of video.

Test simulations were run to determine an appropriate run length, since simulating for the entire 1600 seconds would be expensive in terms of time and computing resources. Also, simulating for the entire 1600 seconds would not be necessary since the aggregated traffic does not fluctuate greatly during the long term, achieving steady state in the short term. However, to ensure that characteristics of both video sequences are captured in their entirety, each data set was divided into 5 partitions of approximately 8000 frames. More

specifically, because of the GOP pattern and size, each partition contained 8005 frames consisting of 667 complete GOPs and one I frame. This produces a one-frame overlap between each partition, except for partition 5. Partition 5 overlaps 21 frames of partition 4. This was necessary since the final complete GOP pattern ends at frame number 39,996. Figure 6.1 shows the video clip segmentation.

6.2.2 Backbone Data Rate

Since the network model consists of only ten source-destination pairs, the backbone data rate was scaled downward. An actual VoD network could potentially support hundreds or even thousands of simultaneous source-destination connections. Multiple test simulations were run. As a result, 6.48 Mbps was chosen because it provided a suitable level of congestion. This scaled data rate is 93% of the aggregate mean data which was produced by all ten sources. Also, this scaled data rate is one eighth OC1 or one twenty-fourth OC3. Recall from Section 2.1.5 that OC1 equals 51.84 Mbps and OC3 equals 155.52 Mbps.

6.2.3 Bursty Source Parameters

The purpose of the bursty source is to create background traffic which contends with video traffic for backbone bandwidth. Additional test simulations were run to determine the interarrival time for the bursty source. An interarrival time of 327.16 μs was chosen. This interarrival time produced 3,056.608 cells per second or 1.296 Mbps or one fortieth OC1. This resultant data rate raised the network load so that the operation of the drop mode feature was activated.

An exponential distribution was used. The exponential distribution is continuous and memoryless. The memoryless property means that the interarrival times are independent. Knowledge of the time since the most recent event does not aid in predicting the next event.

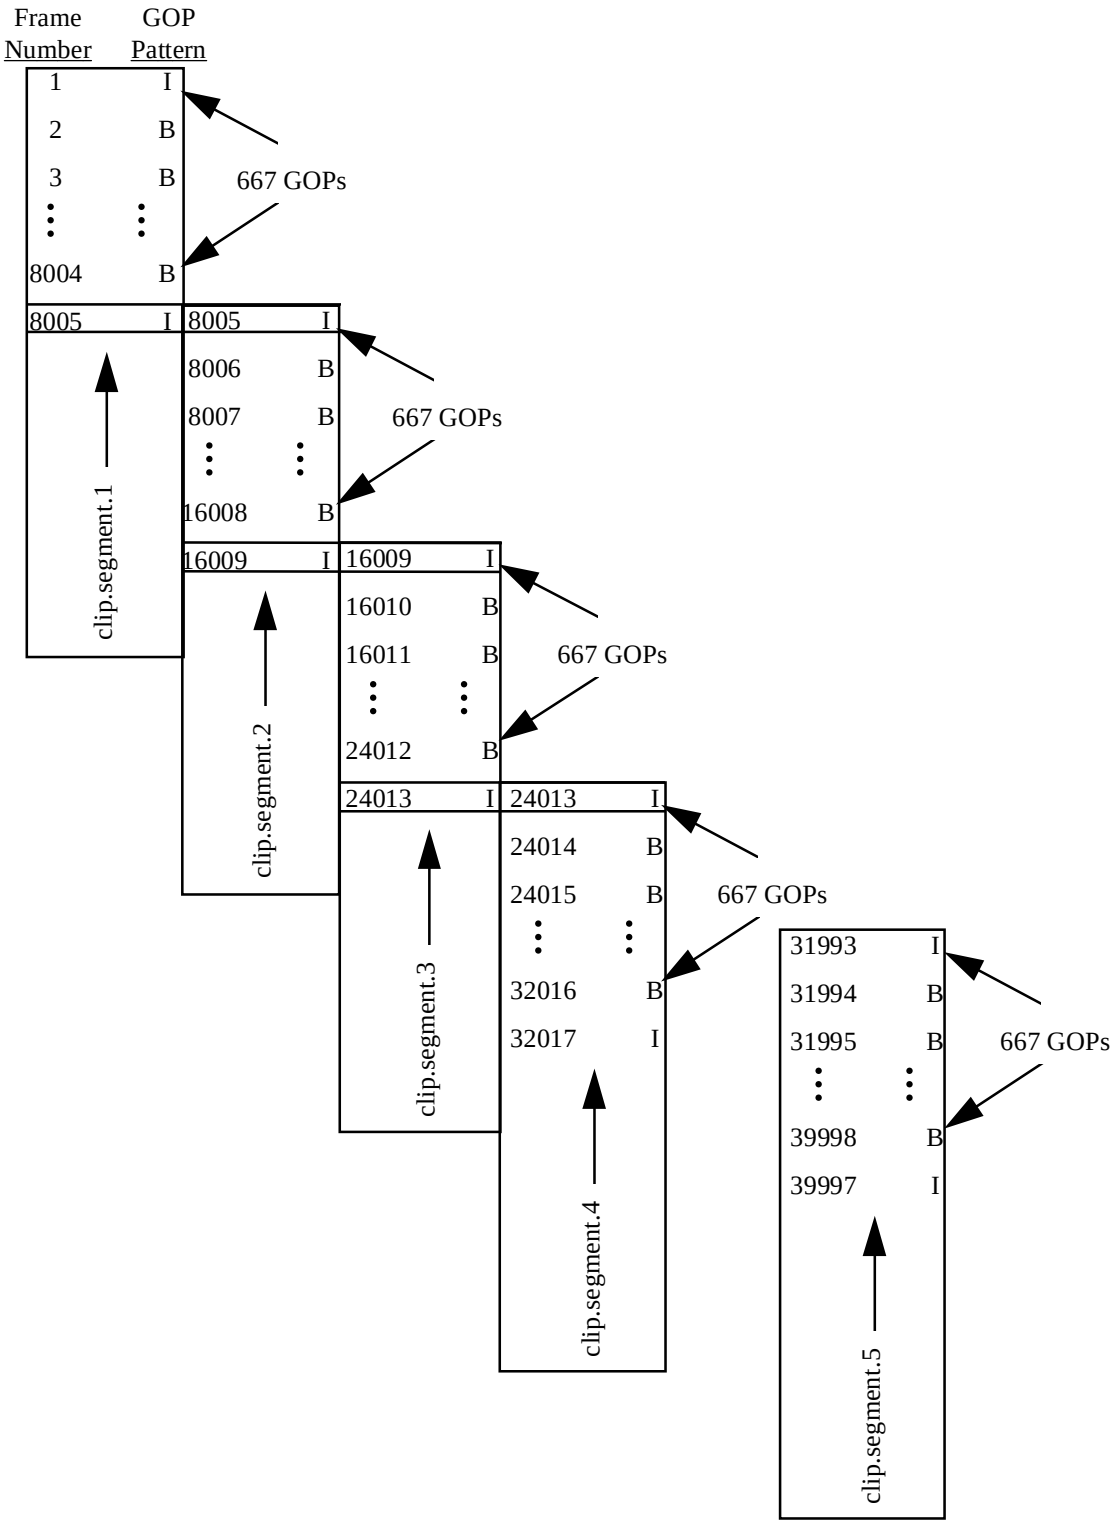

Figure 6.1 Video Sequence Segmentation

6.2.4 Network Congestion Points and Queue Congestion Thresholds

Test simulations were run to locate points of network congestion. The primary bottleneck occurred at the ingress (left-edge) switch. More specifically, the congestion point was the output buffer. This was because the service rate of the output queue was matched to the data rate of the outgoing link. Although the input buffer was the point where all ten streams were multiplexed, it had no congestion because the service rate was 10 Gbps.

No congestion occurred at the egress (right-edge) switch since the queue's service rate was matched to the incoming line speed. Also, it was assumed that the destination set-top-boxes were equipped to handle traffic at the source's transmission rate.

Congestion was indicated when the number of queued cells was greater than the congestion thresholds. The low congestion threshold was 60% of the queue size and the high congestion threshold was 80% of the queue size.

6.3 Analysis of Source Level Control Methods

In this section, the "drop mode off," "drop B-frames only," and "drop B-frames and P-frames" schemes are evaluated. For each drop mode, simulations were run by varying three seeds (1,2, and 3) and two queue sizes (200 cells and 300 cells). Different seeds, chosen arbitrarily, were necessary to ensure that the network model is stable. Different queue sizes should directly impact delay. Additionally, in an actual implementation, queue size affects cost.

The following sections describe the results of simulation runs. Each section compares the effects of source level control on end-to-end delay, congested queue depth, and backbone utilization. All data reflects the average of the three major data sets obtained from the three seeds.

6.3.1 Effects on End-to-End Delay

End-to-end delay is measured from the time a cell is created in CELL_GENn until the cell reaches CELL_SINKn. The model calculates ETE delay per cell. However, to limit output file size, an average of the delay experienced by cells received during a one-second interval is recorded each second. Figures 6.2 through 6.4 show delay graphs for drop mode off, drop B-pictures only, and drop B- and P-pictures, respectively. Also, the maximum queue size is 200 cells.

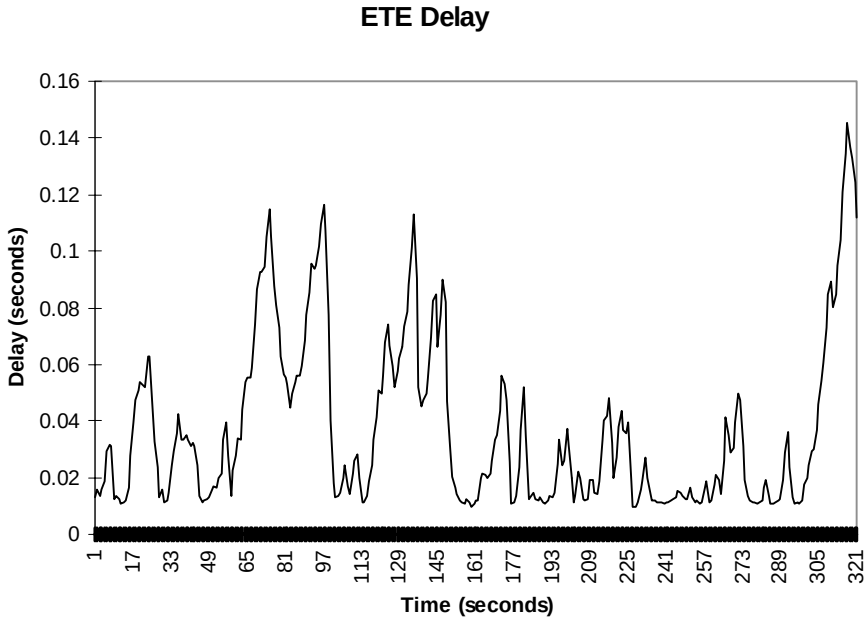

Figure 6.2 Drop Mode Off, End-to-End Delay, 200-Cell Queue Capacity

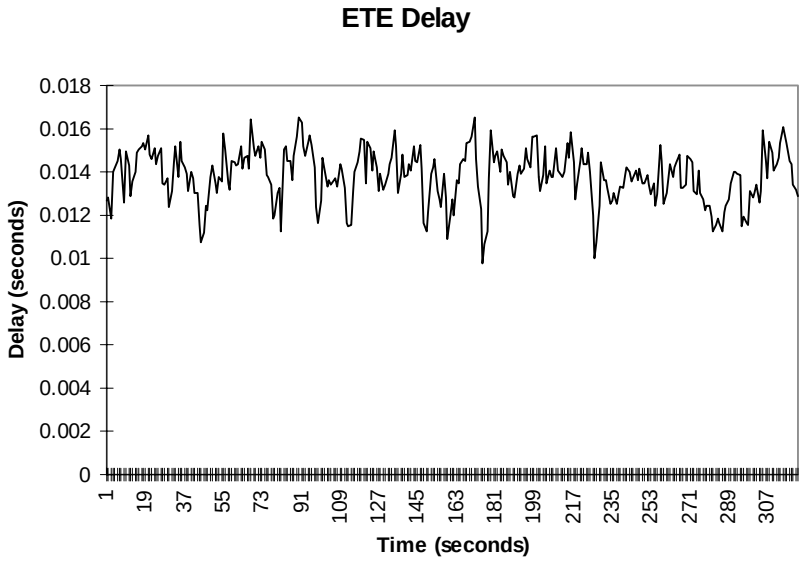

Figure 6.3 Drop B-Pictures Only, End-to-End Delay, 200-Cell Queue Capacity

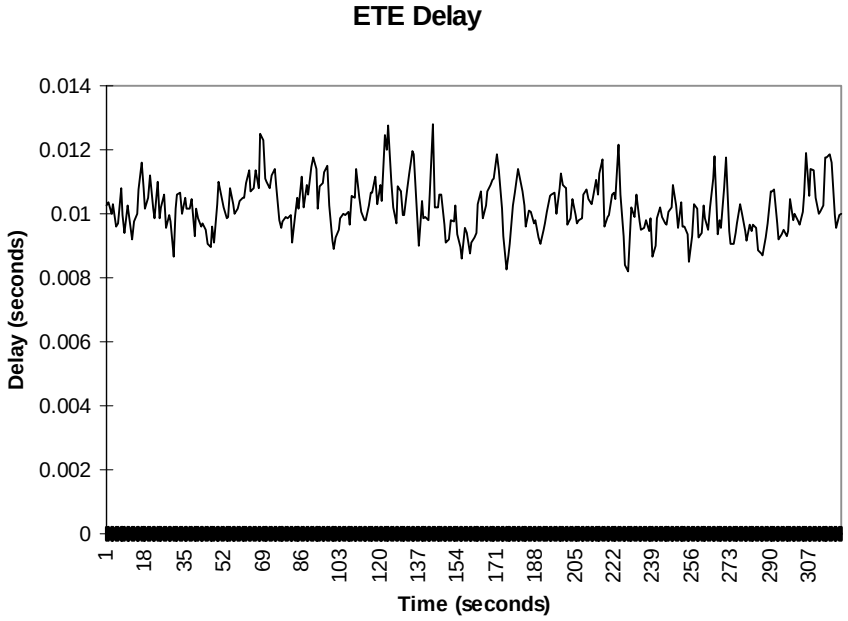

Figure 6.4 Drop B- and P-Pictures, End-to-End Delay, 200-Cell Queue Capacity

Figures 6.2, 6.3, and 6.4 show that delay is worst when the drop mode is off. The mean ETE delay is 0.037 seconds. Delay improves when B-pictures are dropped. The mean delay becomes 0.014 seconds. The delay is minimized when dropping both B- and P-pictures. The average delay is 0.010 seconds.

When the queue capacity is increased from 200 cell to 300 cells, ETE delay increases. As expected, the cell queueing delay contributes to the ETE delay. Figures 6.5 through 6.7 depict ETE delay for a maximum queue size of 300 cells.

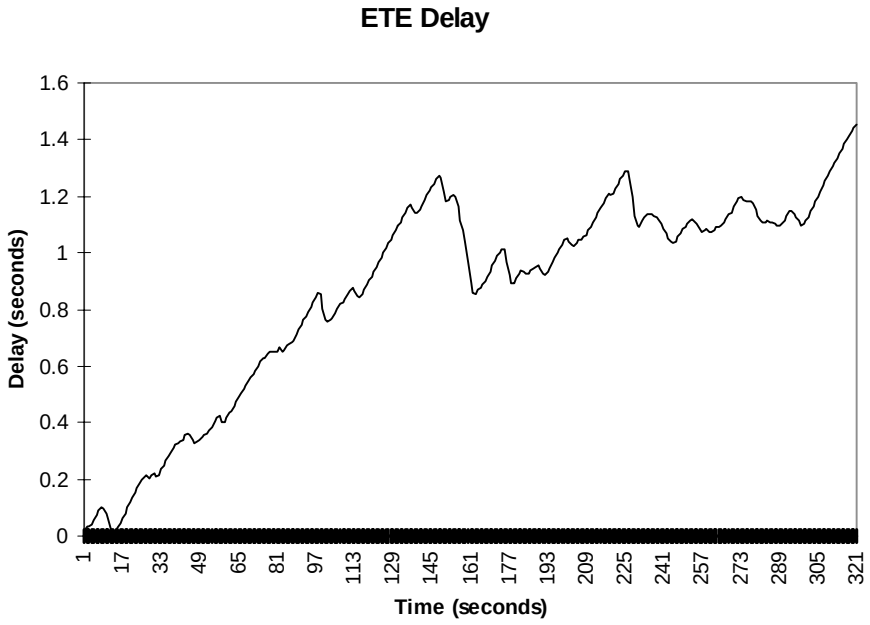

Figure 6.5 Drop Mode Off, End-to-End Delay, 300-Cell Queue Capacity

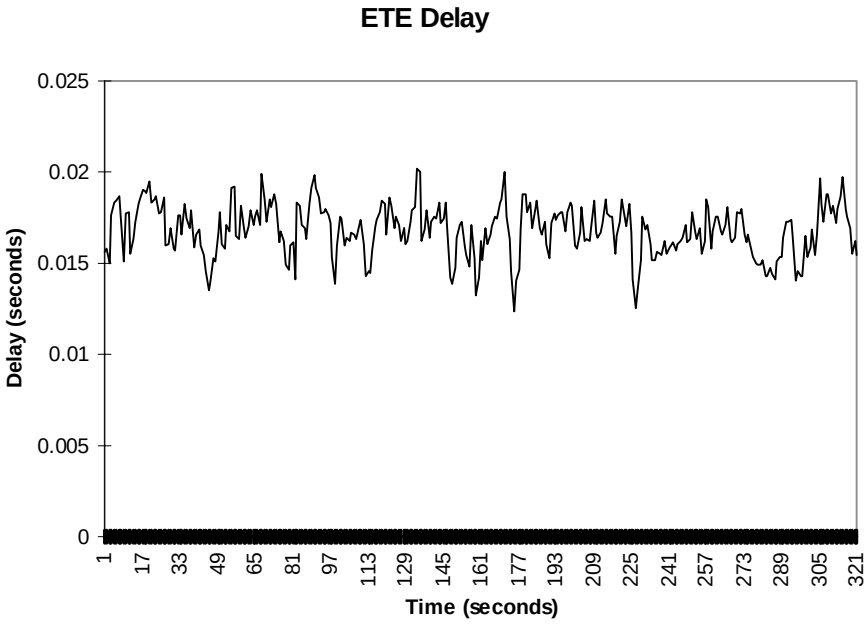

Figure 6.6 Drop B-Pictures Only, End-to-End Delay, 300-Cell Queue Capacity

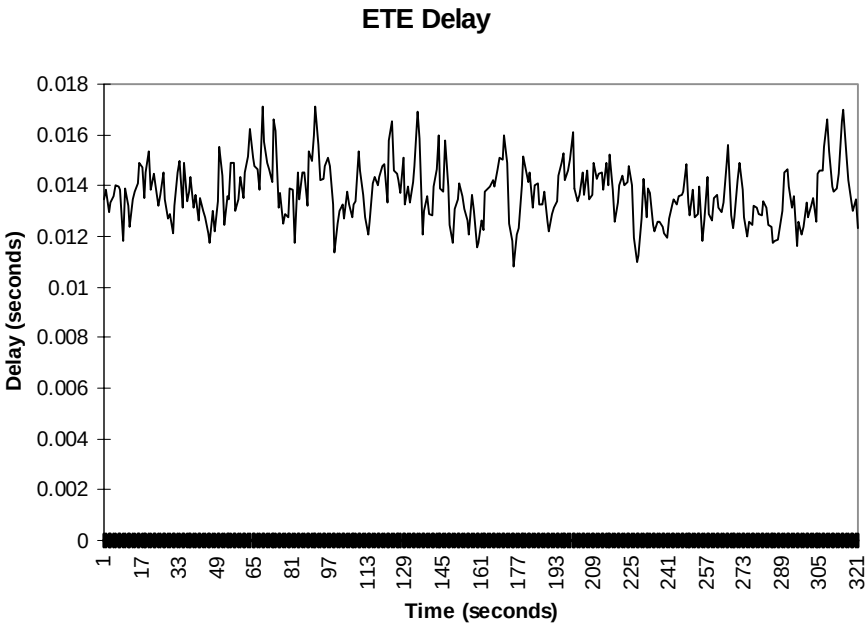

Figure 6.7 Drop B- and P-Pictures, End-to-End Delay, 300-Cell Queue Capacity

Figures 6.5, 6.6, and 6.7 indicate that delay is worst, again when the drop mode is off. The mean ETE delay is 0.087 seconds. Delay improves when B-pictures are dropped. The mean delay becomes 0.017 seconds. The delay is minimized when dropping both B- and P-pictures. The delay is 0.014 seconds.

6.3.2 Effects on Congested Queue Occupancy

As discussed in Section 6.2.4, the primary point of congestion occurs at the output buffer of the ingress switch. Therefore, queue depth is measured here. To keep the output file size manageable, the queue length is not written to file for every cell arrival. Instead, every half-second, the average value of queue levels during the current half second interval is written to the output file. Figures 6.8 through 6.10 show queue depths for drop mode off, drop B-pictures only, and drop B- and P-pictures, for a maximum queue size of 200 cells. Since the frequency of recording queue occupancy values is twice that of the other output data, there are twice the number of data points.

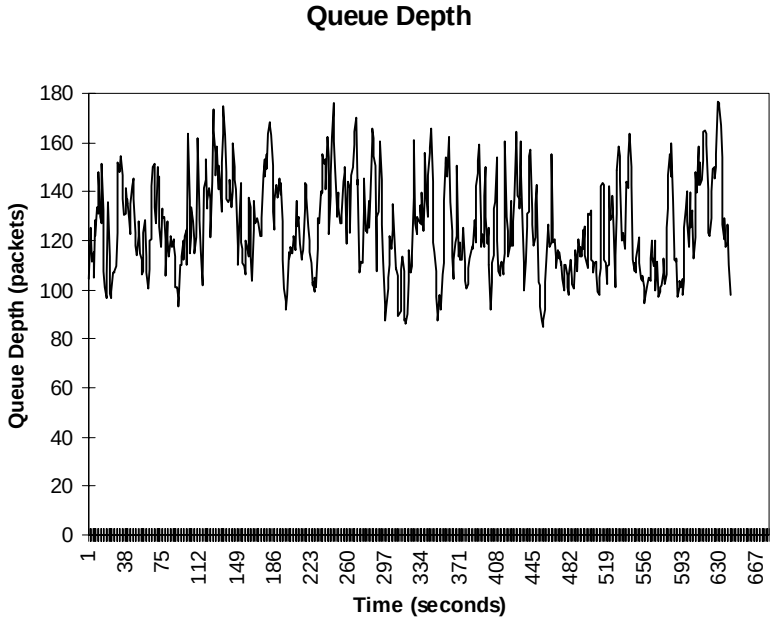

Figure 6.8 Drop Mode Off, Queue Depth, 200-Cell Queue Capacity

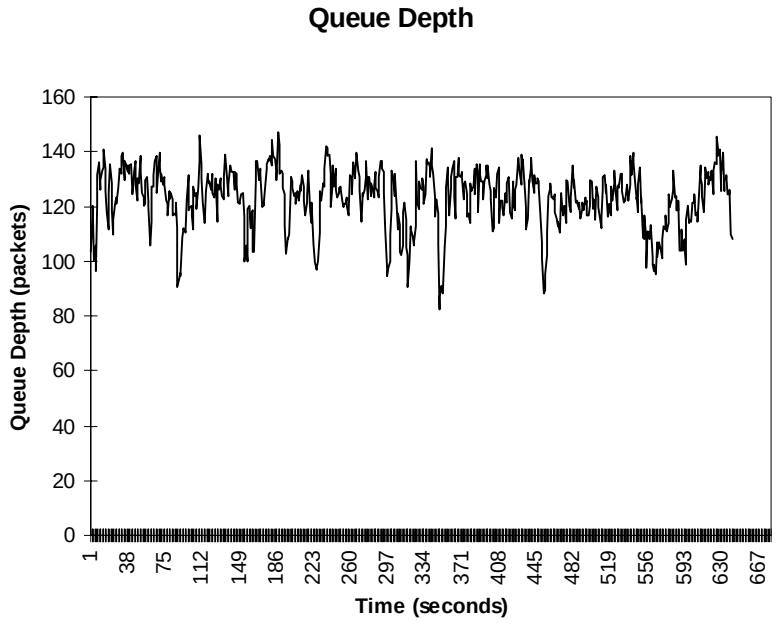

Figure 6.9 Drop B-Pictures Only, Queue Depth, 200-Cell Queue Capacity

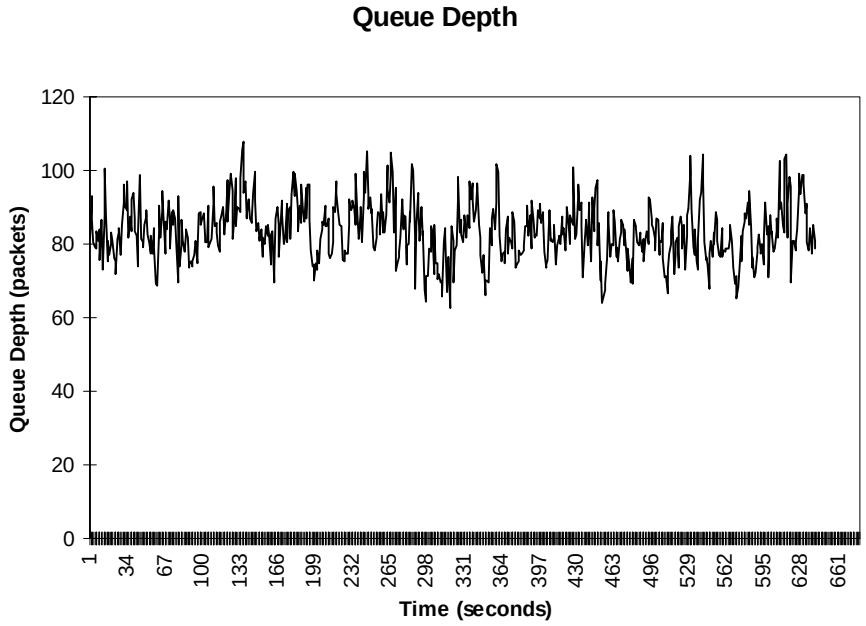

Figure 6.10 Drop B- and P-Pictures Only, Queue Depth, 200-Cell Queue Capacity

As Figures 6.8 through 6.10 show, the largest queue size occurs when the drop mode is off. The mean queue depth is 126 cells or 63% full. When dropping only B-pictures, the mean queue length is 123 cells or 61% occupied. The smallest average queue depth of 83 cells or 42% full, occurs when dropping both B-and P-pictures.

Figures 6.11 through 6.13 show queue occupancy levels when the drop mode is off, dropping B-pictures only, and dropping B- and P-pictures, respectively. The maximum queue size is 300 cells.

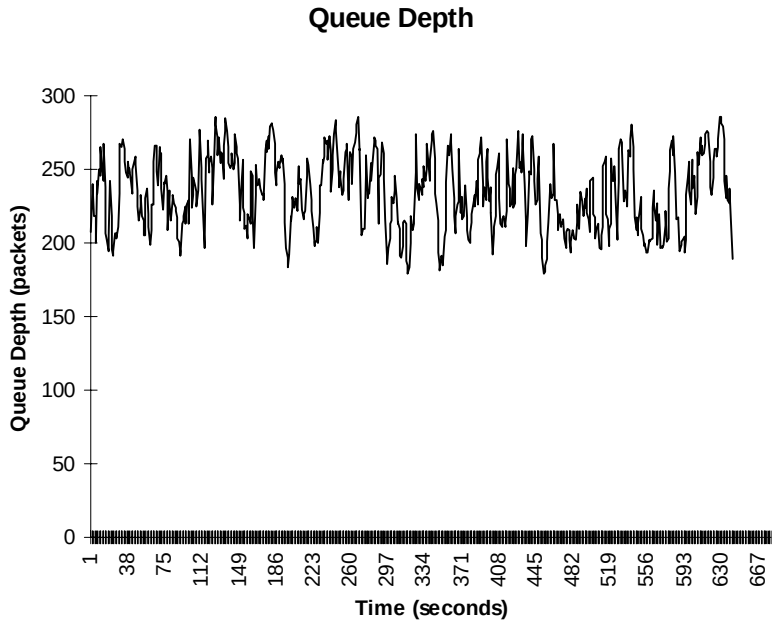

Figure 6.11 Drop Mode Off, Queue Depth, 300-Cell Queue Capacity

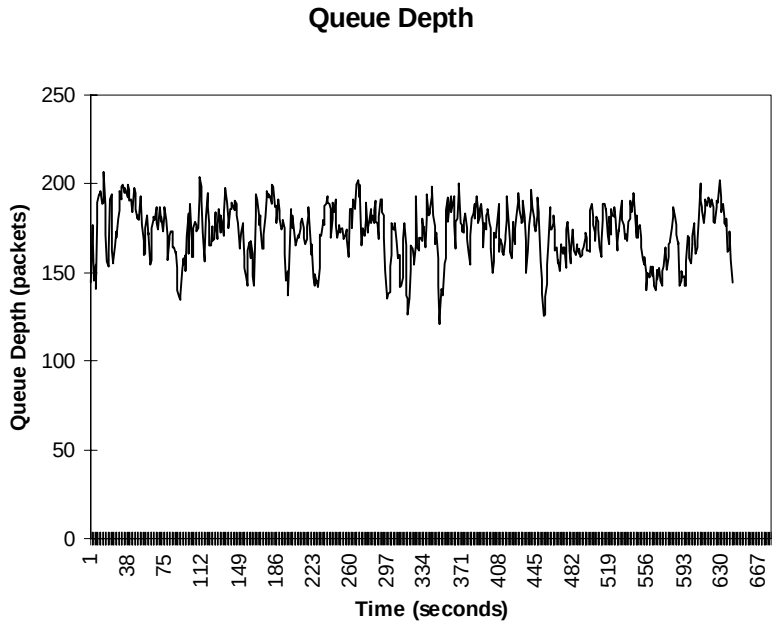

Figure 6.12 Drop B-Pictures Only, Queue Depth, 300-Cell Queue Capacity

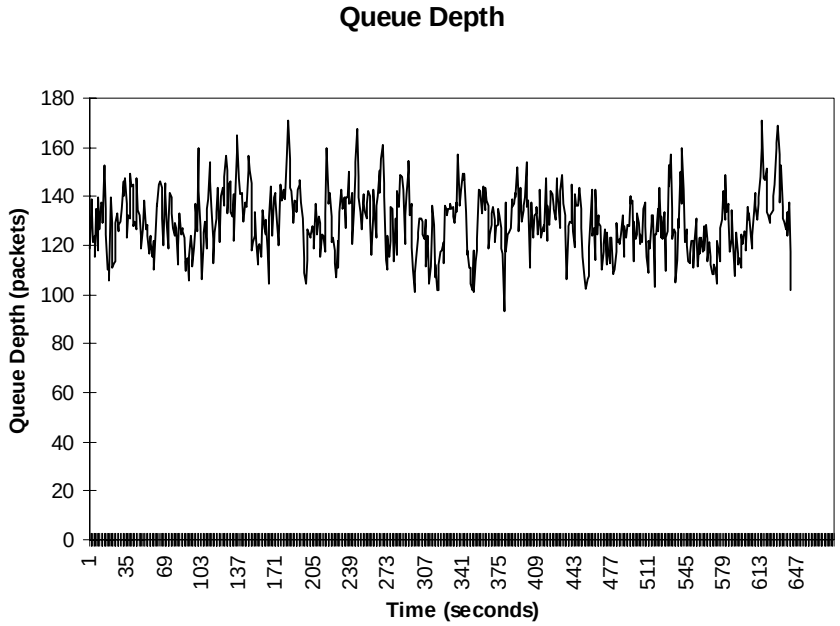

Figure 6.13 Drop B-and P-Pictures Only, Queue Depth, 300-Cell Queue Capacity

Figures 6.11 through 6.13 indicate that the largest queue depth is when the drop mode is off. The mean queue occupancy is 234 cells or 78%. When dropping B-pictures only, the average queue depth decreases to 172 cells or 57% full. When dropping B- and P-pictures, the mean queue depth drops to 130 or 43% full.

6.3.3 Effects on Network Backbone Utilization

Network backbone utilization is measured on the simplex link from TRAF_MGR_SW_L to TRAF_MGR_SW_R. Traffic flow in the reverse direction, from TRAF_MGR_SW_R to TRAF_MGR_SW_L, is not measured. Backbone utilization is important because it limits the number of simultaneous connections. Figures 6.14 through 6.16 show backbone utilization for drop mode off, dropping B-pictures only, and dropping B- and P-pictures, respectively, when the maximum queue size is 200 cells.

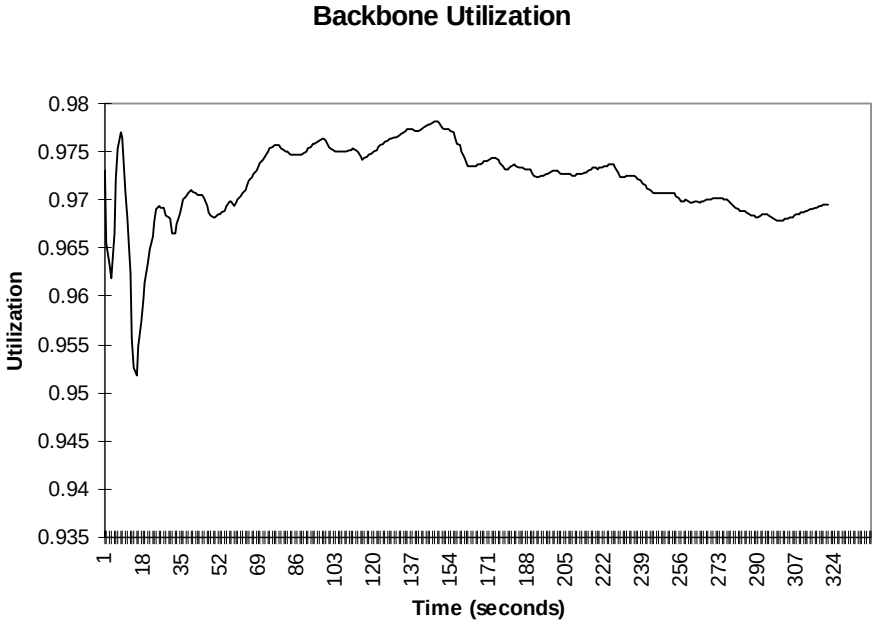

Figure 6.14 Drop Mode Off, Backbone Utilization, 200-Cell Queue Capacity

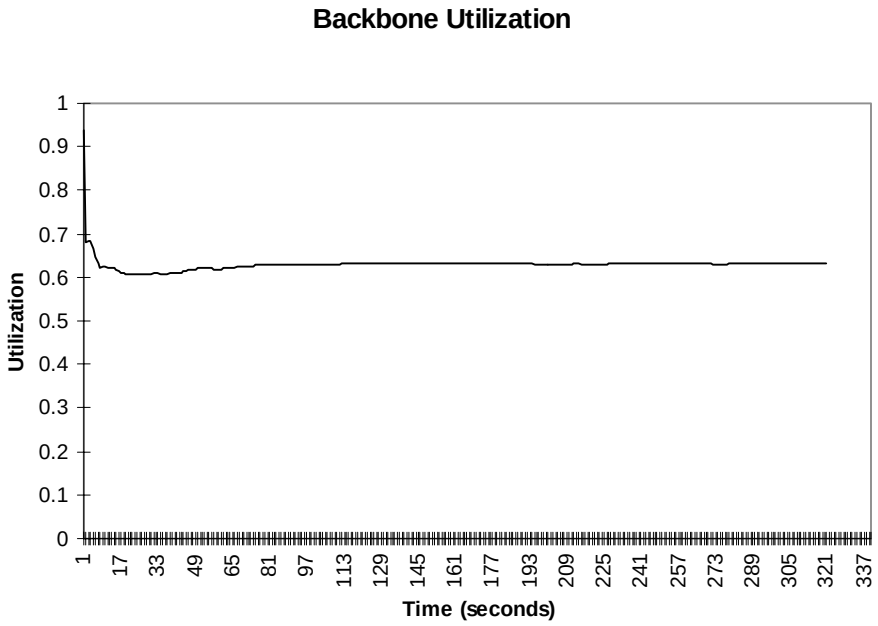

Figure 6.15 Drop B-Pictures Only, Backbone Utilization, 200-Cell Queue Capacity

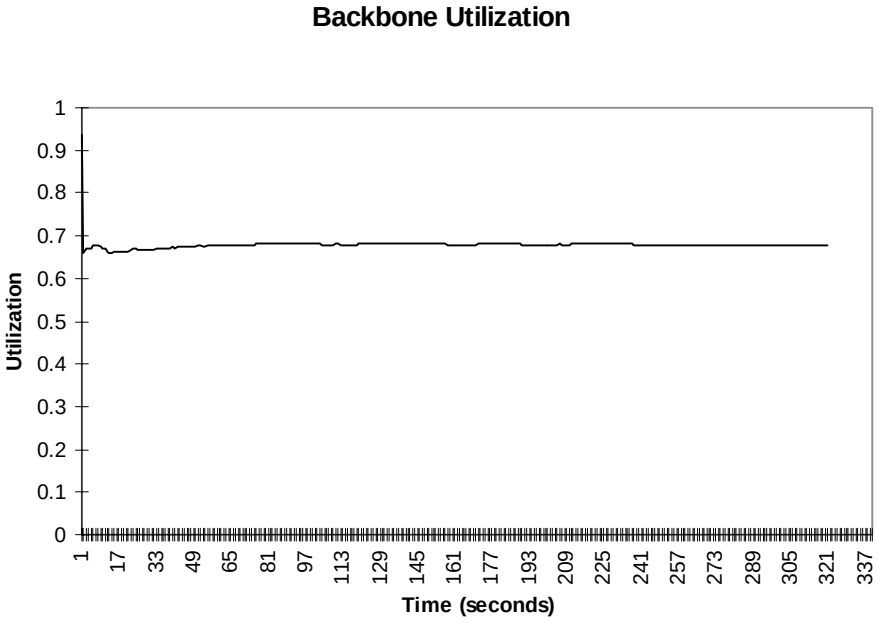

Figure 6.16 Drop B- and P-Pictures, Backbone Utilization, 200-Cell Queue Capacity

Figures 6.14 through 6.16 show that the network is utilized most when the drop mode is off. The mean backbone utilization is 0.97. The average backbone utilization drops to 0.63 when dropping B-pictures only. The mean backbone utilization rises slightly to 0.68 when dropping B- and P-pictures. This rise is likely due to more of the large I-pictures actually getting through the network. Figures 6.17 through 6.19 illustrate when the drop mode is off, dropping B-pictures only, and dropping B- and P-pictures, respectively, for a maximum queue size if 300 cells.

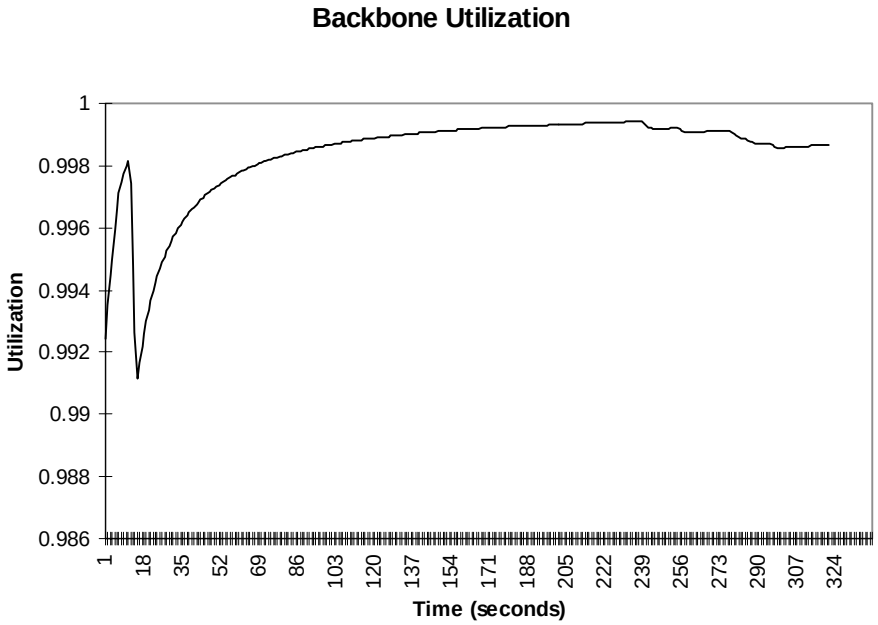

Figure 6.17 Drop Mode Off, Backbone Utilization, 300-Cell Queue Capacity

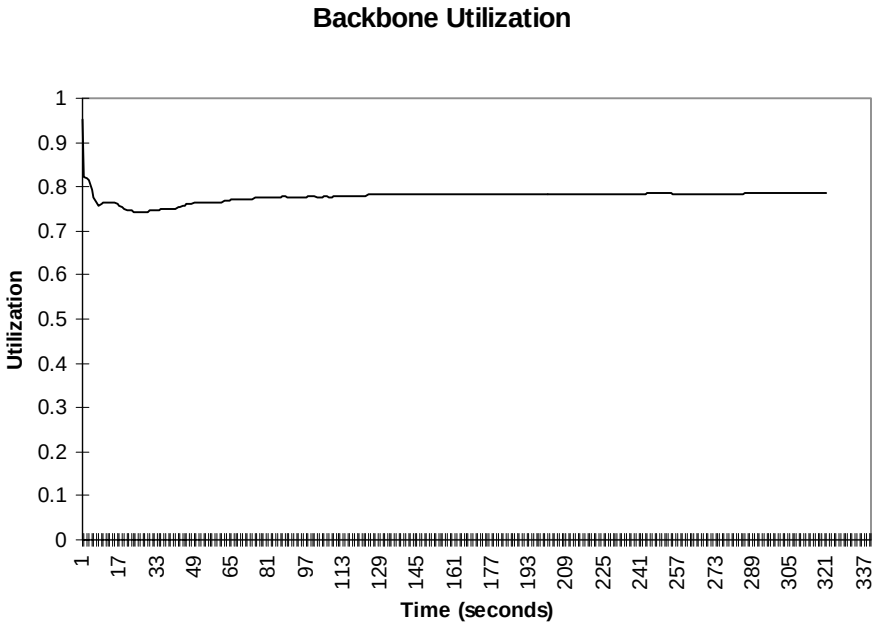

Figure 6.18 Drop B-Pictures Only, Backbone Utilization, 300-Cell Queue Capacity

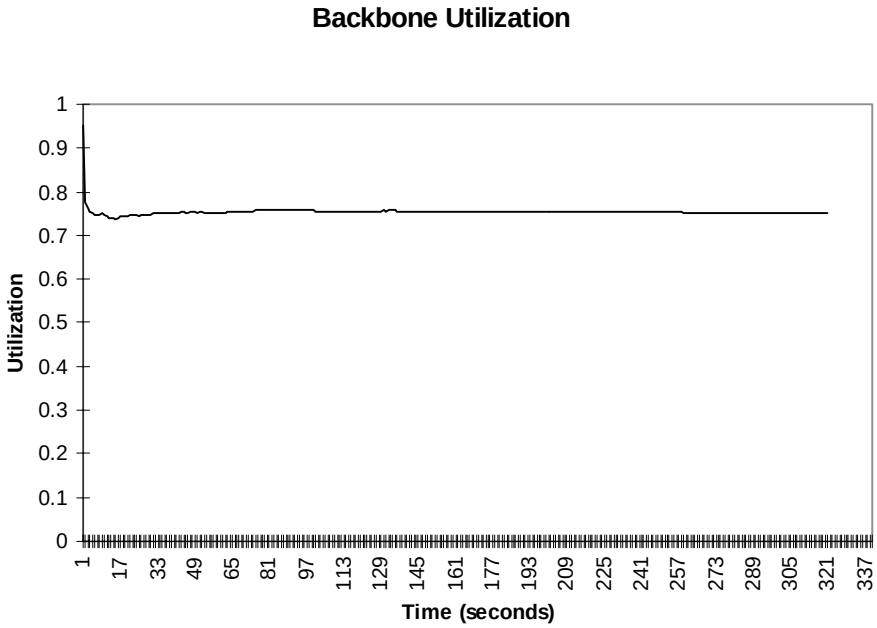

Figure 6.19 Drop B- and P-Pictures, Backbone Utilization, 300-Cell Queue Capacity

Figures 6.17 through 6.19 indicate that network is almost saturated when the drop mode is off. The mean backbone utilization is 0.998. The average backbone utilization drops to 0.78 when dropping B-pictures only. The mean backbone utilization decreases slightly to 0.75 when dropping B- and P-pictures.

Table 6.1 summarizes the statistics for the results discussed above.

Table 6.1 Summary Statistics

Drop Mode	Queue Capacity (cells)	End-to-End Delay (seconds)		Queue Depth (cells)		Backbone Utilization	
		Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Off	200	0.037	0.030	126	20	0.97	0.0039
B-only	200	0.014	0.0012	123	11	0.63	0.019
B and P	200	0.010	0.00081	83	7.9	0.68	0.015
Off	300	0.87	0.37	234	25	0.998	0.0015
B-only	300	0.017	0.0014	172	16	0.78	0.015
B and P	300	0.014	0.0012	130	13	0.75	0.012

6.3.4 Variation in Delay (Jitter)

In this research, jitter is defined to be the variation in ETE delay. Jitter is not calculated during the simulation as are the other response variables. Instead, jitter is calculated after the entire data set is obtained. Jitter is the array of data points equal to the mean ETE delay minus acquired ETE delay values. Cells that arrive late or whose delay is greater than the average ETE delay have a negative jitter. Cells that arrive early or whose delay is less than the average ETE delay have a positive jitter. Figures 6.20 through 6.22 show jitter for drop mode off, drop B-pictures only, and drop B- and P-pictures, when the maximum queue size is 200 cells. In these figures, the jitter is represented in seconds as distance from the mean.

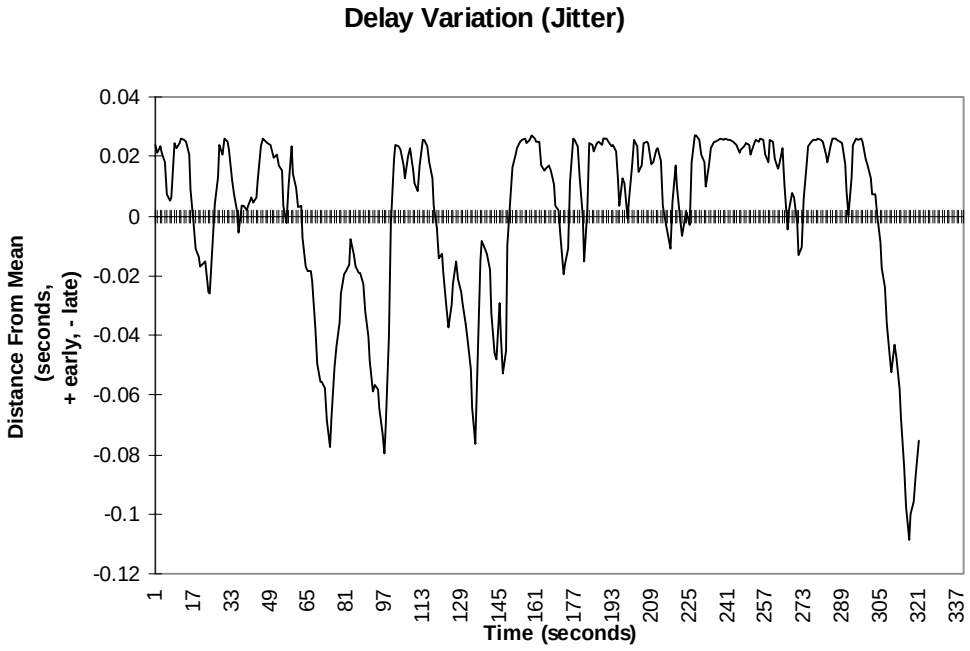

Figure 6.20 Drop Mode Off, Jitter, 200-Cell Queue Capacity

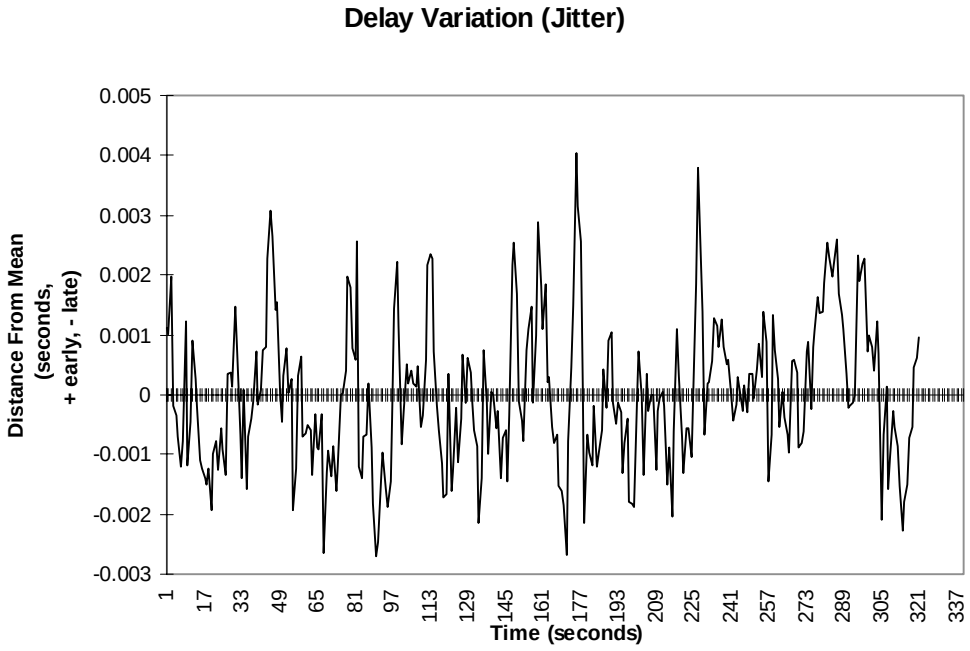

Figure 6.21 Drop B-Pictures Only, Jitter, 200-Cell Queue Capacity

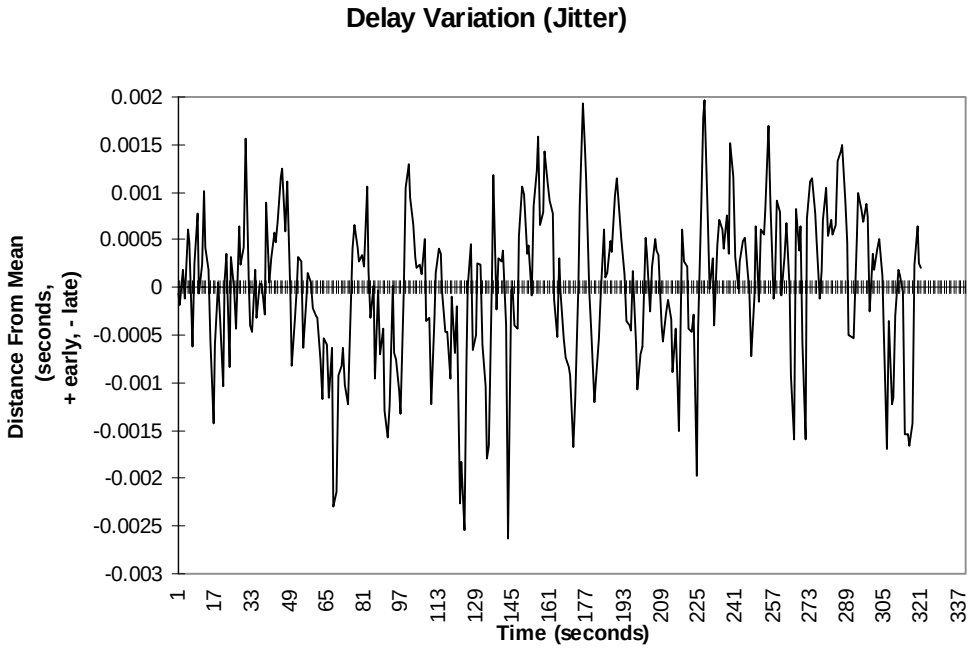

Figure 6.22 Drop B- and P-Pictures, Jitter, 200-Cell Queue Capacity

Figures 6.20 through 6.22 show that the worst jitter occurs when the drop mode is off. Jitter for this case is misleading because there are many time instances that appear to indicate early cell arrival while only a few indicate very late cells. The fact is that the occurrences of extremely delayed cells greatly influence the mean ETE delay. Jitter graphs resulting from dropping B-pictures only and B- and P-pictures are more reasonable.

Figures 6.23 through 6.25 show jitter for drop mode off, drop B-pictures only, and drop B- and drop B- and P-pictures, when the maximum queue size is 300 cells.

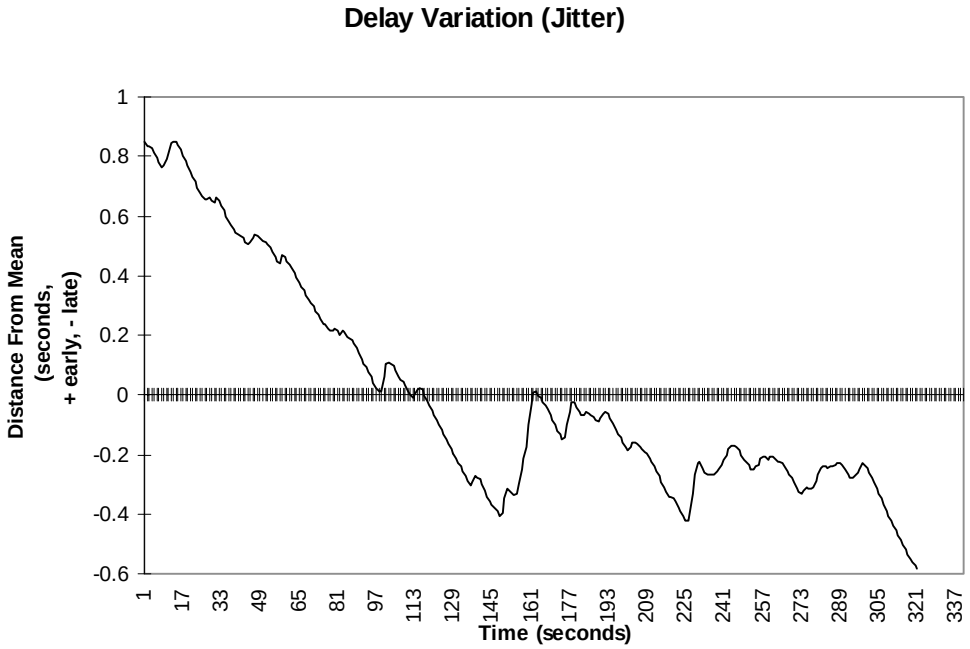

Figure 6.23 Drop Mode Off, Jitter, 300-Cell Queue Capacity

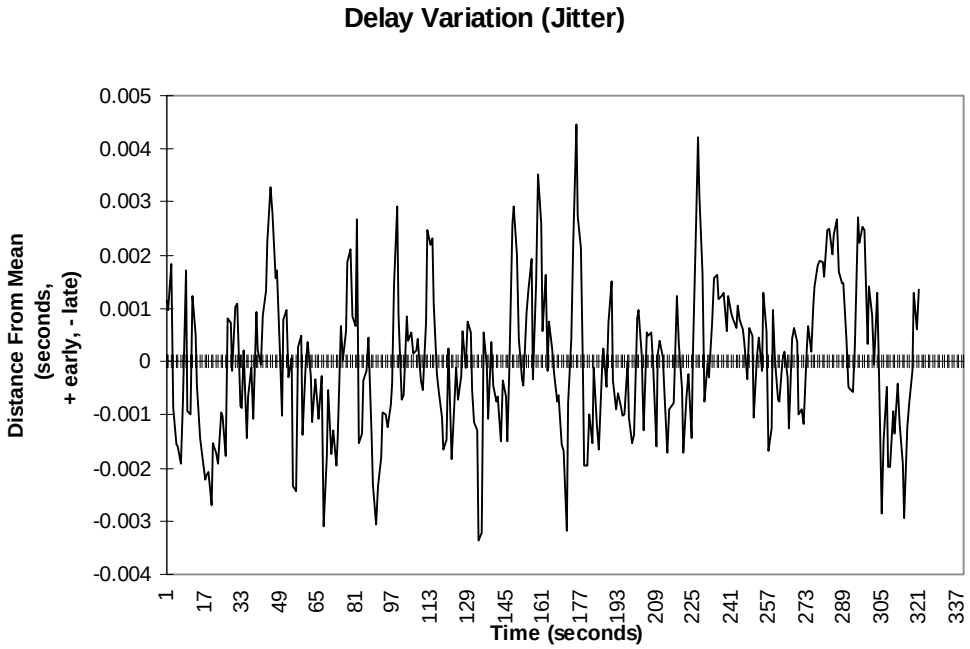

Figure 6.24 Drop B-Pictures Only, Jitter, 300-Cell Queue Capacity

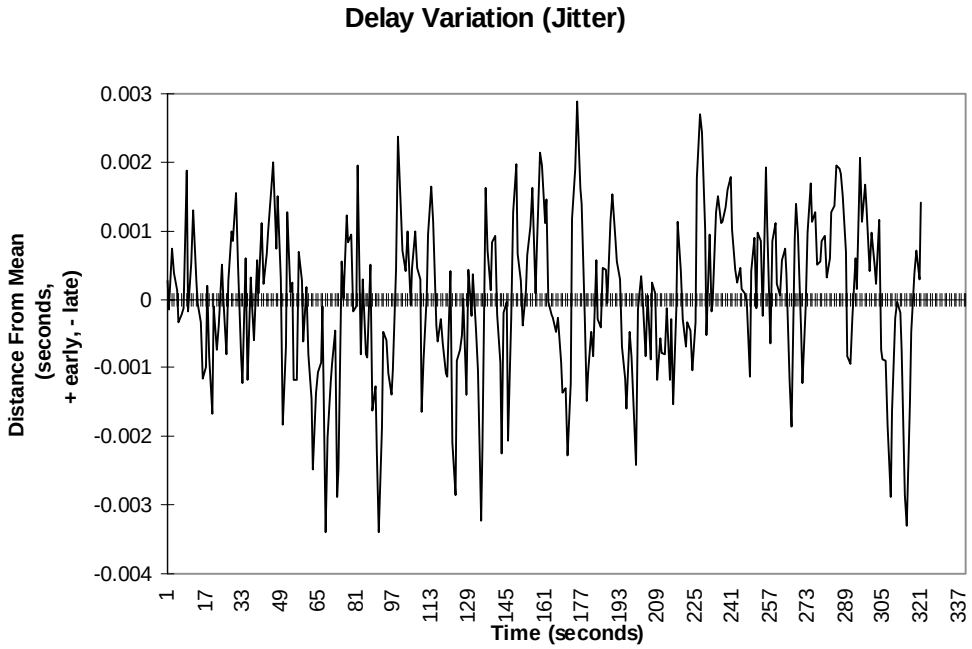

Figure 6.25 Drop B- and P-Pictures, Jitter, 300-Cell Queue Capacity

Figures 6.23 through 6.25 illustrate that the worst jitter again occurs when the drop mode is off. Jitter graphs resulting from dropping B-pictures only and B- and P-pictures are more reasonable.

6.3.5 Distribution of Delay

An alternate way to examine ETE delay is to view its distribution. To generate these graphs, bins were created representing percentages of the mean ETE delay. For example, the number of occurrences of ETE delay greater than 0 but less than or equal to 10% of the mean ETE delay is shown along the X-axis as 0.1. The number of occurrences of ETE delay greater than 10% but less than or equal to 20% of the mean ETE delay is shown along the X-axis as 0.2. The graphs illustrate 10% intervals up to 300% of the mean ETE delay. Figures 6.26 through 6.28 show frequency for drop mode off, drop B-pictures only, and drop B- and drop B- and P-pictures, when the maximum queue size is 200 cells.

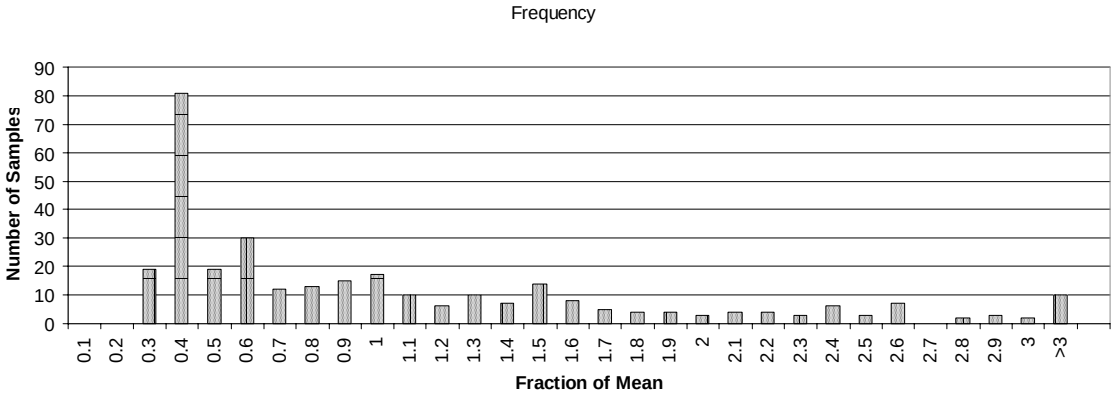

Figure 6.26 Drop Mode Off, Frequency, 200-Cell Queue Capacity

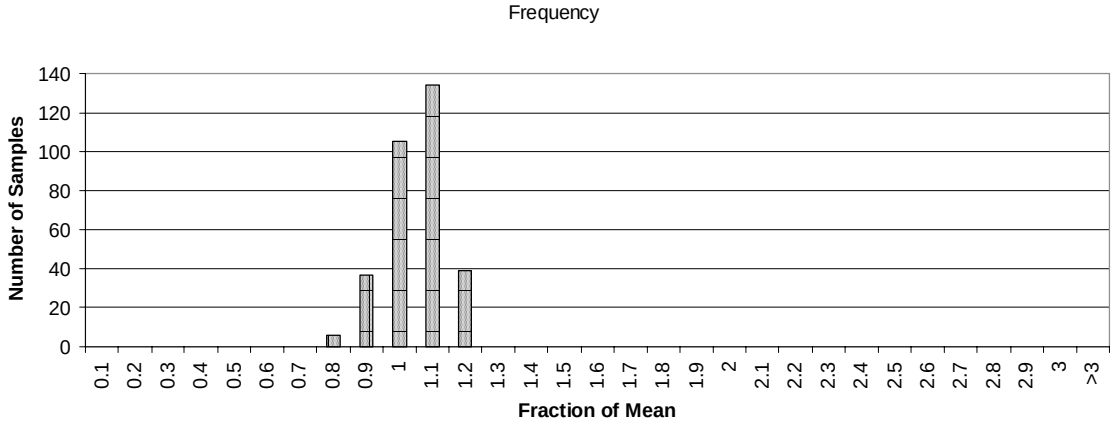

Figure 6.27 Drop B-Pictures Only, Frequency, 200-Cell Queue Capacity

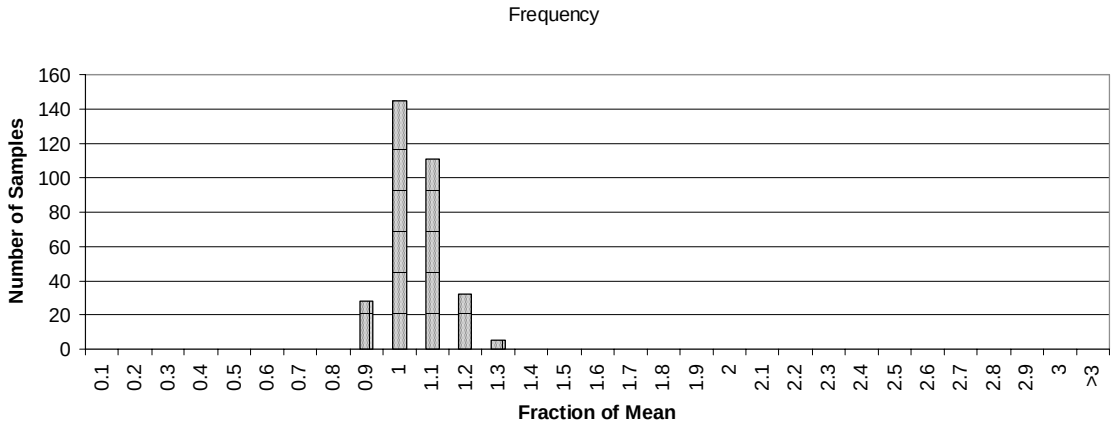

Figure 6.28 Drop B- and P-Pictures, Frequency, 200-Cell Queue Capacity

Figures 6.26 through 6.28 illustrate that when the drop mode is off, delay varies over a wide range. When dropping B-pictures only and dropping B- and P-pictures, the delay range is much tighter. Figures 6.29 through 6.31 show Frequency for drop mode off, drop B-pictures only, and drop B- and drop B- and P-pictures, when the maximum queue size is 300 cells.

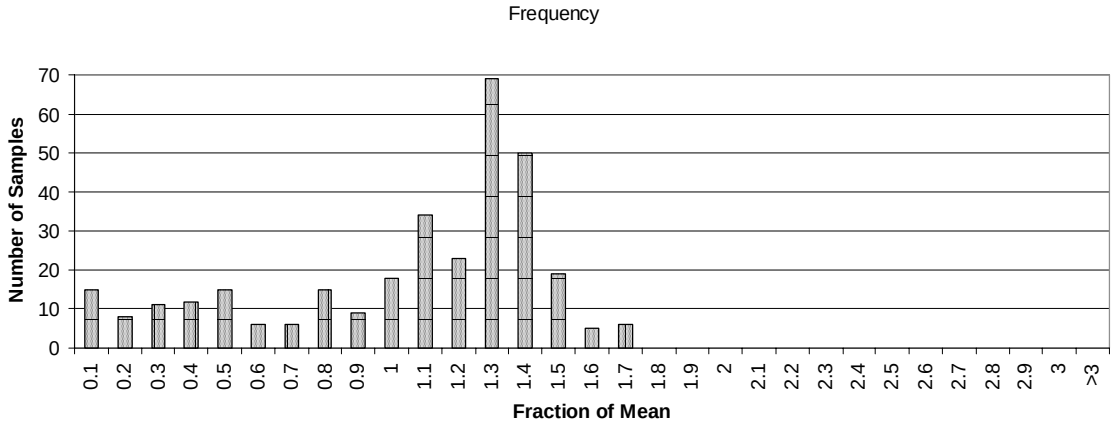

Figure 6.29 Drop Mode Off, Frequency, 300-Cell Queue Capacity

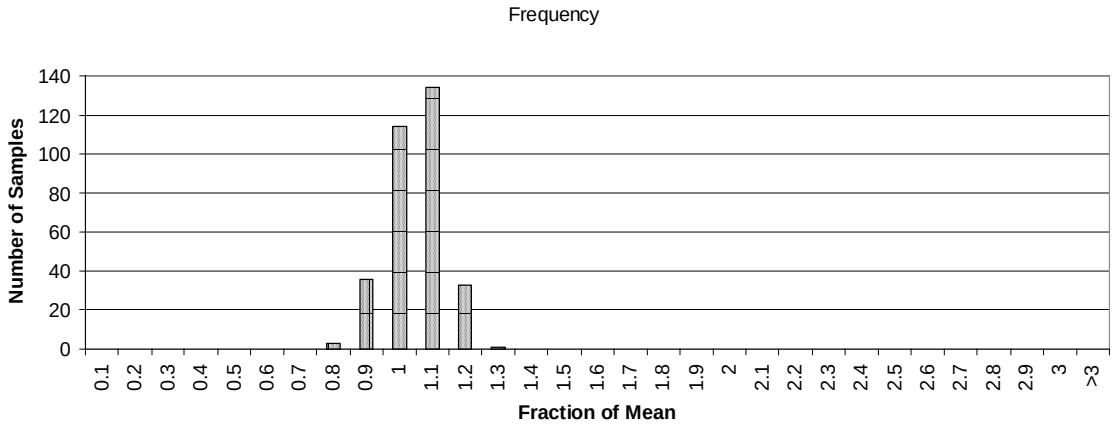

Figure 6.30 Drop B-Pictures Only, Frequency, 300-Cell Queue Capacity

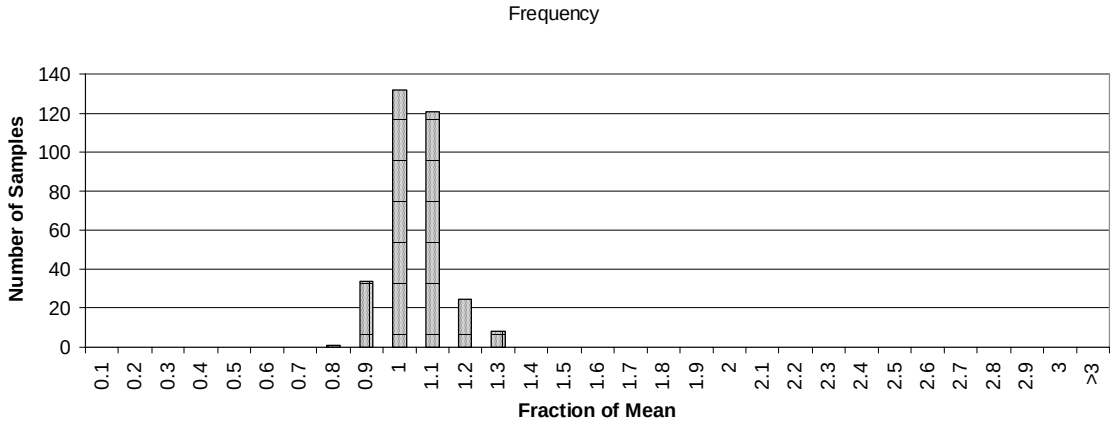

Figure 6.31 Drop B- and P-Pictures, Frequency, 300-Cell Queue Capacity

Figures 6.29 through 6.31 illustrate that when the drop mode is off, delay varies over a wide range. When dropping B-pictures only and dropping B- and P-pictures, the delay range is much tighter.

6.4 Summary

This chapter evaluated the concepts in this research and demonstrated the utility of the simulation model. This chapter described the trace data sets used, the model's backbone data rate, bursty source parameters and network queue congestion.

Response variables such as ETE delay, queue occupancy, and backbone utilization were discussed. Based on the response variables, performance was evaluated for the drop mode off, dropping B- and P-pictures, and dropping B-pictures only. Alternative representations of ETE delay, such as delay variation (jitter) and distribution of delay, were discussed.

For queue capacities of 200 and 300 cells, the delay was worst when the drop mode was off. Delay improved when B-pictures were dropped. The delay was minimized when dropping both B- and P-pictures. When the queue capacity was increased from 200 cell to 300 cells, overall ETE delay increased, as expected because the cell queueing delay contributed to the ETE delay.

For queue capacities of 200 and 300 cells, queue occupancy levels were highest when the drop mode was off. The levels decreased slightly when dropping B-pictures only. The lowest queue occupancy levels occurred when dropping both B-and P-pictures.

For a queue capacity of 200 cells, the network was most utilized when the drop mode was off. The average backbone utilization decreased when dropping B-pictures only. The mean backbone utilization increased slightly when dropping B- and P-pictures. This rise was likely because more of the large I-pictures were actually getting through the network.

For a queue capacity of 300 cells, the network was almost saturated when the drop mode was off. The average backbone utilization dropped when dropping B-pictures only. The mean backbone utilization decreased slightly when dropping B- and P-pictures.

For queue capacities of 200 and 300 cells, the worst jitter occurred when the drop mode was off. Jitter graphs resulting from dropping B-pictures only and B- and P-pictures were more reasonable.

For queue capacities of 200 and 300 cells, when the drop mode was off, delay varied over a wide range. When dropping B-pictures only and dropping B- and P-pictures, the delay range was much tighter.

These results indicate that even though ABR flow control is designed for data traffic, it is suitable for transporting MPEG-1 video streams. However, ABR must be coupled with an effective method of rate-adjustment. Dropping B-pictures and dropping both B- and P-pictures provides higher performance than disabling the drop mode feature.

Chapter 7 reviews information presented in this thesis, presents conclusions, and describes opportunities for future work.

Chapter 7. Conclusions

7.1 Summary of Research

A brief review of the information presented in this thesis is provided before the conclusions and recommendations sections. Chapter 1 offered motivation and objectives for this study. This chapter also listed the assumptions and limitations of this research.

Chapter 2 gave an overview of significant phases in communications networks leading to the development of ATM. Historical highlights consisted of the T1 carrier, optical fiber, packet-switching, narrow and broadband integrated services digital networks, and ATM.

Chapter 3 explained ATM bandwidth management. Quality of Service parameters were discussed. Real-time and non-real-time service categories were also presented. Lastly, explicit forward congestion indicator, explicit rate, and relative rate switch mechanisms were described.

Chapter 4 gave background information on video compression standards and adaptive video compression techniques. Common video compression schemes, Motion-JPEG, H.261, and MPEG-1, were reviewed. Adaptive video encoding methods were presented.

Chapter 5 described the experimental design and implementation for this study. This chapter presented the response variables. Chapter 5 also offered details for each node model, queue or processor module, and process model.

Chapter 6 described the trace data sets used, the model's backbone data rate, bursty source parameters, and network queue congestion. This chapter also presented results of this research by discussing effects on ETE delay, queue occupancy, and backbone utilization.

7.2 Conclusions from Analysis

Even though ABR flow control is designed for data traffic, it is suitable for transporting MPEG-1 video streams. However, ABR flow control must be coupled with an effective method of rate-adjustment. It is important to remember that the ABR flow control mechanism provides congestion status feedback. It is necessary for the either the network switches or VoD servers to take action based on this feedback. As reported in Chapter 6, dropping B-pictures only and dropping both B- and P-pictures provides higher performance than disabling the drop mode feature.

7.3 Recommendations for Future Research

There are several opportunities for further study. In this study, relative rate marking was used. More research could determine if explicit rate marking would enhance performance. Also, would explicit rate computations incur complexity costs? Furthermore, since several variations of explicit rate marking exist, comparison between them might provide further insight into ABR video transport.

The switch buffer management policy for this study was a cell-based FIFO queue which indicated congestion when the occupancy level exceeded a specified threshold. It is possible to monitor the rate of change in queue depth. Congestion can be indicated whenever the queue fills too quickly. A comparison of the two approaches might provide further enhancements.

References

[APTE95] R. T. Apteker, J. A. Fisher, V. S. Kisimov, and H. Neishlos, "Video Acceptability and Frame Rate," *IEEE Multimedia*, Volume 2, Number 3, pp. 32-40, Fall 1995.

[ATMF95] ATM Forum. *ATM User-Network Interface Specification, Version 3.1*, September 1994, Prentice Hall PTR, Upper Saddle, New Jersey, Also, available via anonymous ftp to: ftp.atmforum.com/, subdirectory: pub/UNI/ver3.1/, filename: all_pdf.tar.Z.

[ATMF96] ATM Forum. *Traffic Management Specification, Version 4.0*, April 1996, anonymous ftp to: ftp.atmforum.com/, subdirectory: pub/approved-specs/, filename: af-tm-0056.000.doc.zip.

[BERT92] D. Bertsekas and R. Gallager, *Data Networks, Second Edition*, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1992, pp. 128-140.

[CHEN96] T. M. Chen, S. M. Liu, and V. K. Samalam, "The Available Bit Rate Service for Data in ATM Networks," *IEEE Communications Magazine*, Volume 34, Number 5, pp. 56 - 71, May 1996.

[DEPR93] M. De Prycker, R. Peschi, and T. Van Landegem, "B-ISDN and the OSI Protocol Reference Model," *IEEE Network*, Volume 7, Number 2, pp. 10-18, March 1993.

[DEPR95] M. De Prycker, *Asynchronous Transfer Mode: Solution for Broadband ISDN, Third Edition*, London, Prentice Hall International (UK) Limited, 1995.

[DUTT95] H. J. R. Dutton and P. Lenhard, *High-Speed Networking Technology: An Introductory Survey*, Upper Saddle River, New Jersey, Prentice-Hall, Inc., 1995, pp. 9-41.

[GUPT95] S. Gupta and C. Williamson, "A Performance Study of Adaptive Video Coding Algorithms for High Speed Networks," February 1995, anonymous ftp to: ftp.cs.usask.ca/, subdirectory: pub/discus/, filename: paper.95-4.ps.Z.

[HÄND91] R. Händel and M. N. Huber, *Integrated Broadband Networks: An Introduction to ATM-Based Networks*, Wokingham, England: Addison-Wesley Publishing Company, Incorporated, 1991, p. 2.

[HUGH95] D. Hughes and K. Hooshmand, "ABR Stretches ATM Network Resources," *Data Communications*, Volume 24, Number 5, pp. 123-126, 128, April 1995.

[JAIN91] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement Simulation, and Modeling*, John Wiley and Sons, Incorporated, New York, 1991.

[JAIN96] R. Jain, S. Kalyanaraman, S. Fahmy, R. Goyal, S.-C. Kim, "Source Behavior for ATM ABR Traffic Management: An Explanation," *IEEE Communications Magazine*, Volume 34, Number 11, pp. 50 - 57, November 1996. Also, available via: http://www.cis.ohio-state.edu/~jain/papers/src_rule.htm.

[KALY97] S. Kalyanaraman, R. Jain, R. Goyal, S. Fahmy, and S.-C. Kim, "Performance of TCP over ABR on ATM Backbone and with Various VBR Background Traffic Patterns," *Conference Record for 1997 IEEE*

International Conference on Communications, Montréal, Québec, Canada, pp. 1035 - 1041, June 1997. Also, available via: http://www.cis.ohio-state.edu/~jain/papers/tcp_vbr.htm

[LAPO94] T. F. La Porta, M. Veeraraghavan, E. Ayanoglu, M. Karol, and R. D. Gitlin, "B-ISDN: A Technological Discontinuity," *IEEE Communications Magazine*, Vol. 32, No. 10, pp. 84-97, October 1994.

[LEE95] T.-E. Lee, "A Study and Analysis of a Transmission Scheduling and Discard Algorithm for ATM Networks," Master's Thesis, 1995.

[MIL396] OPNET Modeler Documentation and Users Manuals, 1996.

[MIL397] Mil 3, Incorporated, Maker of OPNET Modeler/Planner, http://www.mil3.com.

[ROBE95] L. G. Roberts, "Can ABR Service Replace VBR Service in ATM Networks," *Proceedings of the COMPCOM '95 Conference*, Digest of Papers, COMPCOM, IEEE Computer Society International Conference, San Francisco, California, pp. 346-348, March 1995.

[ROBE97] L. G. Roberts, "Explicit Rate Flow Control: A 100 Fold Improvement over TCP," Web Document, April 1997. Available via: http://www.ziplink.net/~lroberts/Explicit-Rate/Explicit-Rate-Flow-Control.html

[ROSE95] O. Rose, "Statistical Properties of MPEG Video Traffic and Their Impact on Traffic Modeling in ATM Systems," *Proceedings of the 20th Annual Conference on Local Computer Networks*, Minneapolis, Minnesota, pp. 397-406, October 1995.

[ROOH95] R. Rooholamini and V. Cherkassy, "ATM-Based Multimedia Servers," *IEEE Multimedia*, Volume 2, Number 1, pp. 39-52, Spring 1995.

[SAAD94] T. N. Saadawi, M. H. Ammar, and A. El Hakeem, *Fundamentals of Telecommunications Networks*, New York: John Wiley & Sons, Inc., 1994, pp. 4-5, 367-395.

[SCHÄ95] R. Schäfer and T. Sikora, "Digital Video Coding Standards and Their Role in Video Communications," *Proceedings of the IEEE,* Volume 83, Number 6, pp. 907-924, June 1995.

[SCHU95] D. Schuster, "Available Bit Rate (ABR) Traffic Management," *Connectware, Inc. Technology White Paper*, August 1995, available via: http://www.connectware.com/wh-paper/abr.html.

[SHEL94] T. Sheldon, *LAN Times: Encyclopedia of Networking*, Berkeley, Osborne McGraw-Hill, 1994.

[SIU95] K.-Y. Siu and R. Jain, "A Brief Overview of ATM: Protocol Layers, LAN Emulation, and Traffic Management," *Computer Communication Review*, Volume 25, Number 2, pp. 6-20, April 1995.

[SIU94] K.-Y. Siu and H.-Y. Tzeng, "Intelligent Congestion Control for ABR Service in ATM Networks," *Computer Communication Review*, ACM SIGCOMM, Volume 24, Number 5, pp. 81-106, October 1994.

[TSE96] P. W. Tse, M. Zukerman, and F. Le Faucheur, "Performance and Fairness Issues Related to the EFCI Rate-Based Flow Control for the ABR Service

in ATM Networks," *Computer Systems Science and Engineering*, Volume 11, Number 6, pp. 383-392, November 1996.

[VETT95] R. J. Vetter, "ATM Concepts, Architectures, and Protocols," *Communications of the ACM*, Volume 38, Number 2, pp. 30-38+, February 1995.

[WALT96] R. Walthall and M. Clement, "Simulation and Analysis of the Performance of EPRCA in a Wide Area ATM Network Consisting of Both ER and EFCI Switch Mechanisms," *Fourth International Conference on Telecommunications Systems, Modeling and Analysis*, March 1996. Also available at: http://nebo.cs.byu.edu/~walthalr/congestion.html.

[ZUKE97] M. Zukerman and P. W. Tse, "Performance and Fairness Comparison between EFCI and ERF Flow Control Schemes for the ABR Service in ATM Networks," *Conference Record for 1997 IEEE International Conference on Communications*, Montréal, Québec, Canada, pp. 220 - 224, June 1997.

Appendix. ITU-T and ATM Forum

The body primarily responsible for ATM standardization is the International Telecommunications Union - Technical Standards Division or the ITU-T. The organization was previously known as the International Consultative Committee on Telegraphy and Telephony (CCITT). A group known as the ATM Forum, which is an organization consisting of almost every data communications telecommunications equipment supplier in the world, is defining "standards" called implementation agreements for ATM that go beyond the development done by the ITU. In general, the ITU is defining the WAN standards for ATM and the ATM Forum is defining the LAN standards. The ATM Forum was consciously set up to develop its standards quickly [DUTT95].

Vita

Elvin Lattis Taylor, Jr. was born on March 25, 1971 in Durham, North Carolina. He attended Hillside High School from which he graduated in May 1989. He graduated from North Carolina State University with Bachelor of Science degrees in Electrical Engineering and Computer Engineering in 1994. He interned for five summers with International Business Machines Corporation in Research Triangle Park, North Carolina and Boca Raton, Florida. He graduated with a Master of Science in Electrical Engineering from Virginia Tech in December 1997.