INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA
313/761-4700 800/521-0600

An Explicit Rate Control Algorithm and A Source Bandwidth Request Algorithm for the Transport of MPEG Video over ATM

by

Sunitha Krishnan, B.Sc. Hons.
York University

A thesis submitted to
the Faculty of Graduate Studies and Research
in partial fulfillment of
the requirements for the degree of

Master of Engineering

Ottawa-Carleton Institute for Electrical Engineering
Department of Systems and Computer Engineering

Carleton University
Ottawa, Ontario
Jan 23 1998
© 1998 Sunitha Krishnan

National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-32402-8

The undersigned recommend to
the Faculty of Graduate Studies and Research
acceptance of the thesis

**An Explicit Rate Control Algorithm and A Source
Bandwidth Request Algorithm for the Transport
of MPEG Video over ATM**

submitted by Sunitha Krishnan, B. Sc. Hons.
in partial fulfillment of the requirements for
the degree of Master of Engineering

Thesis Supervisor

Chair, Department of Systems and
Computer Engineering

Carleton University

Jan 1998
date

Abstract

An explicit rate control algorithm is proposed along with a bandwidth request algorithm for the transport of MPEG video over an ABR-like service over ATM. This form of closed-loop feedback-based control of real-time video traffic offers several advantages over conventional transport of video using the Constant Bit-Rate (CBR) and the Variable Bit-Rate (VBR) service classes of ATM. Detailed knowledge of the source video traffic statistics is not required for negotiating bandwidth at setup which makes it ideal for real-time video.

The proposed bandwidth request algorithm based on the delay experienced at the source buffer, is shown through simulations to compare favorably with a scheme based merely on the prediction of upcoming frame sizes. An explicit rate algorithm which allocates a fixed rate component equal to the negotiated minimum cell-rate and a variable rate component based on the request made and the available bandwidth at the current time, is also proposed. It is shown to offer better performance in terms of lower cell-loss and good frame delays as compared to an adaptation of an explicit rate control scheme proposed for data traffic.

The composite system with our proposed bandwidth request policy at the video sources and our explicit rate algorithm implemented at the network switches is shown to provide good quality of service for heavy network loads and high propagation delays.

Acknowledgements

I wish to thank my thesis supervisor Professor Ioannis Lambadaris for his valuable suggestions, patience and guidance during the course of my thesis. I also wish to thank Professor Mihail Devetsikiotis whose guidance and encouragement at critical times was much appreciated.

I gratefully acknowledge the funding and equipment provided by the Telecommunications Research Institute of Ontario and by the Natural Sciences and Engineering Research Council in support of this research.

I am also deeply indebted to my parents for their encouragement and prayers without which this thesis might not have been possible.

Table of Contents

Acceptance Sheet	ii
Abstract	iii
Acknowledgements	iv
Table of Contents	v
List of Tables	ix
List of Figures	x
List of Acronyms	xi
1 Introduction	1
1.1 Background	1
1.2 Motivation	2
1.3 Thesis Objectives	5
1.4 Organization of Thesis	5
2 Overview of Video over ATM	7
2.1 MPEG Video	7
2.1.1 Basic MPEG Video Compression	8
2.1.2 Quality Requirements for Video	10
2.2 ATM Technology	12
2.2.1 The ATM Protocol Stack	12
2.2.2 ATM Service Parameters	14
2.2.3 ATM Service Categories	17

2.2.4	The ATM Traffic Management	19
2.3	Bandwidth Prediction for Video over ATM	20
2.3.1	Long-Range Prediction for RCBR	20
2.3.2	Prediction for Cell-Rate Smoothing for VBR	22
3	The Available Bit-Rate (ABR) Service Class	26
3.1	Basic Service Model for ABR	26
3.1.1	Resource Management (RM) Cells	27
3.1.2	Source End System (SES) Rules	30
3.1.3	Destination End System (DES) Rules	31
3.2	Categories of Rate Control Schemes	32
3.2.1	Single bit feedback	32
3.2.2	Explicit Rate (ER) feedback	35
3.2.3	Virtual Source/Destination Control	35
3.2.4	Credit-based feedback	36
3.3	Some Explicit Rate Algorithms	37
3.3.1	Enhanced Proportional Rate Control Algorithm (EPRCA)	38
3.3.2	Congestion Avoidance using Proportional Control (CAPC)	39
3.3.3	Explicit Rate Indication for Congestion Avoidance (ERICA)	40
3.3.4	Explicit Rate Allocation Algorithm (ERAA)	41
3.3.5	Fast Max-Min Rate Allocation Algorithm (FMMRA)	41
3.3.6	Suitability of Explicit Rate Control to VBR video	43

4	Source and Switch Algorithms	45
4.1	Source Bandwidth Request Algorithms	45
4.1.1	Hypothetical Perfect Prediction	46
4.1.2	Dual Delay Threshold Algorithm (DDTA)	47
4.2	Switch Algorithm	50
4.2.1	Fast Max-Min Rate Allocation (FMMRA) - 1	50
4.2.2	Minimum Cell-Rate-based Proportional Allocation (MCRPA)	51
4.3	Simulation Models	54
4.3.1	Video Traces	54
4.3.2	Source End System	55
4.3.3	Destination End System	57
4.3.4	ATM Switch Model	59
5	Simulations For Performance Analysis	61
5.1	DDTA vs Perfect Prediction	61
5.1.1	Simulation Setup	62
5.1.2	Observations	63
5.2	FMMRA-1 vs MCRPA	68
5.2.1	Simulation Setup	68
5.2.2	Observations	70
5.2.3	Effect of Minimum Cell Rate	75
5.3	Sensitivity of the DDTA-MCRPA system to propagation delays and loads	77
5.3.1	Simulation Setup	77
5.3.2	Observations for Varying Loads	78
5.3.3	Observations for Varying Propagation Delays	80

6	Conclusions and Future Results	82
6.1	Summary and Conclusions	82
6.2	Suggestions for Future Investigations	83

List of Tables

3.1	ABR parameters signaled at setup	27
3.2	Fields in an RM cell	28
4.1	Statistics for Ghost video - frame sizes in bytes for I,P frames at a quantization level of 3, and B frames at a quantization level of 6 . . .	55
4.2	Statistics for BBC news video - frame sizes in bytes for I,P frames at a quantization level of 3, and B frames at a quantization level of 6 . .	55
4.3	Statistics for NCSU football video - frame sizes in bytes for I,P frames at a quantization level of 3, and B frames at a quantization level of 6	56
5.1	Mean Actual Cell Rate to Mean Allocated Cell Rate ratios for the three videos, using DDTA	66
5.2	95 % Confidence intervals for FMMRA-1, prop delay=10ms, 12 runs .	70
5.3	95 % Confidence intervals for MCRPA, prop delay=10ms, 12 runs . .	73
5.4	Fairness Results for MCRPA with 95% Confidence intervals for prop delay=10ms, 12 runs, 2000cps minimum cell-rate for all	74
5.5	95% Confidence intervals for MCRPA with N identical football sources, prop delay = 10ms and N = 17, 12 runs	78

List of Figures

2.1	MPEG forward and backward prediction	8
3.1	Fields in the ATM cell header	28
3.2	An overview of ABR rate control	33
4.1	Histograms and Auto-Correlation for NCSU football, BBC News & Ghost video traces	57
4.2	System Diagram	58
4.3	Mean fraction of cells per slice for NCSU football, BBC News & Ghost video traces	59
5.1	Ideal Source-Receiver System	62
5.2	Comparing DDTA with Perfect Prediction for football	63
5.3	Comparing DDTA with Perfect Prediction for BBC news	64
5.4	Comparing DDTA with Perfect Prediction for Ghost	65
5.5	Heterogeneous Sources, 3 Switch System	68
5.6	Results using FMMRA-1, MCR = 2000cps for all	71
5.7	Results using MCRPA, MCR = 2000cps for all	72
5.8	Results using MCRPA, MCR = 2000cps for Ghost, 4000cps for BBC news and 6000cps for football	76
5.9	Statistically Identical, Independent Sources and Two Switch System	77
5.10	Varying number of football sources	79
5.11	Varying propagation delay	80

List of Acronyms

AAL	ATM Adaptation Layer
ABR	Available Bit-Rate
ACR	Allowed Cell-Rate
ATM	Asynchronous Transfer Mode
BRM	Backward Resource Management (cell)
CBR	Constant Bit-Rate
CCR	Current Cell-Rate
CDF	Cumulative Distribution Function
CDV	Cell Delay Variation
CDVT	Cell Delay Variation Tolerance
CLR	Cell Loss Ratio
EFCI	Explicit Forward Congestion Indication
ER	Explicit Rate
FRM	Forward Resource Management (cell)
GCRA	Generic Cell-Rate Algorithm
ISDN	Integrated Services Digital Network
MCR	Minimum Cell-Rate
MPEG	Motion Picture Experts Group
PCR	Peak Cell-Rate
QoS	Quality of Service
RCBR	Renegotiated Constant Bit-Rate
RM	Resource Management (cell)
SCR	Sustained Cell-Rate
UBR	Unspecified Bit-Rate
VBR	Variable Bit-Rate
VC	Virtual Circuit

Chapter 1

Introduction

1.1 Background

The concept of Broadband Integrated Services Network (B-ISDN) was developed in order to integrate data, voice and video information into a single all-purpose network [1]. Asynchronous Transfer Mode (ATM) has been chosen by the International Telecommunications Union (ITU) as the switching technology required to implement B-ISDN networks [2]. This is because of its ability to provide high speed switching and multiple service categories suitable for different traffic types. With the growing demand in real-time video services like video-conferencing, support of real-time video over ATM has become an important issue. Video requires high bandwidth of the order of hundreds of Megabytes per second, making it necessary to compress the video streams before their transmission. Several compression standards for video have been proposed which include H.261, motion JPEG, and MPEG [3]. MPEG is emerging as a popular standard owing to its VHS-quality video, its ability to make use of any available bandwidth through trade-offs between quality and bit-rate, and use of bi-directional prediction thereby reducing bandwidth requirements [4]. Compressed real-time video streams are sensitive to both delay and loss since the real-time delay requirements make it infeasible to use retransmission to recover lost or errored cells.

ATM offers two main traffic classes for video [5], namely, Constant Bit-Rate (CBR) and Variable Bit-Rate (VBR). The CBR service as the name suggests, guarantees a

fixed bit-rate for the duration of the connection. The VBR service allows video sources into the network based on statistical characteristics of the video sources such as worst case burstiness and long term mean bandwidth requirements. The Quality of Service (QoS) requested by the source will be guaranteed as long as the source traffic adheres to the statistics negotiated at setup. Both CBR and VBR have their disadvantages for real-time compressed video. The burstiness of compressed video traffic [6] makes it difficult to specify the bandwidth required for its transmission. Allocating a fixed bandwidth to a video stream for the duration of transmission as in CBR, could lead to under-utilization of network bandwidth when the video bit-rate is low. Allocating bandwidth based on statistical bandwidth bounds specified by the video source at the time of setup as in VBR, allows statistical multiplexing gain at the network [7]. However, the use of conventional leaky bucket-based VBR service for allocating bandwidth and policing of the video services has been found to be questionable [8, 9, 10]. The difficulty of determining optimal leaky bucket parameters for a video source and preventing QoS violations even when the sources conform to the leaky bucket parameters considered to be acceptable at setup, are the major disadvantages of the VBR service. Moreover, for real-time video, statistics of the video source are not available in advance making it difficult to specify the source traffic parameters at the time of setup. Satisfying the needs of compressed real-time video over ATM networks is thus a challenge.

1.2 Motivation

Existing traffic classes for transporting compressed video over ATM, which are the Constant Bit-Rate (CBR) and Variable Bit-Rate (VBR) service, are basically open-loop services requiring accurate knowledge of the video traffic at the time of setup. This open-loop nature of these services makes them difficult to use for real-time video, the characteristics of which are not available in advance. A natural path for

the investigation of the transport of real-time compressed video over ATM is therefore closed-loop feedback control schemes. Renegotiated Constant Bit-Rate (RCBR) [11] was proposed as an enhancement to the CBR service. The video source would periodically renegotiate a new bandwidth if it anticipated a change in its current requirements. However, it does not make use of full closed loop control as the switches do not convey congestion information back to the source. A full closed loop feedback control service is provided by ATM for data traffic. This is the Available Bit-Rate (ABR) service which is intended for data traffic that can tolerate delays but is very sensitive to loss. However, previous work done with ABR-like rate control schemes for VBR video have shown promising results. This includes work done with single bit feedback indicating the presence or absence of congestion [12], feedback of buffer-occupancy and service rate computed on a per-connection basis at network switches [13] and feedback of switch port utilization from network switches [14].

Explicit rate control is a recent addition to the ABR rate control schemes. The feedback cell called the Resource Management (RM) cell, containing the desired transmission rate is transmitted by the source at regular intervals. On receiving an RM cell, each switch computes a rate that it can support for the connection, and overwrites the rate already in the cell if its computed rate is below the requested rate and passes the cell forward to the next switch along the transmission path. The smallest rate that can be supported along the path is thus passed back to the source which then updates its rate of transmission to the received rate. This “explicit” specification of the rate by the network thus differentiates it from other rate control schemes. Several explicit rate control schemes have been proposed for ABR data traffic [15, 16, 17, 18, 19]. Explicit rate control offers many advantages over single bit congestion indication schemes, the main being fairness and quick reaction to congestion [20, 21]. The ATM Forum [5] has adopted rate-based control which includes both single bit congestion indication schemes and explicit rate schemes for ABR data traffic. A detailed discussion of ABR rate control schemes can be found in [22, 23, 5].

Very recently, an explicit rate service has been proposed for transporting compressed video over ATM by Lakshman et al [24]. Their simulations show that an explicit rate control scheme with a minimum bit-rate guarantee along with a source bandwidth prediction algorithm can offer good performance for compressed video. Our work differs from theirs in two main ways. The explicit rate scheme that we propose is simpler to implement compared to the weighted fair-queuing scheme proposed in [24]. We do not have to recompute the status and weights of VCs other than the one being currently processed. The rate allocation computed by our explicit rate scheme has a constant component depending on the minimum cell-rate guaranteed to the source and a variable component depending on the source request and the available bandwidth at the current time. We believe that this policy can tightly control the offered load at the network switches thus reducing the probability of over-allocation of bandwidth and therefore cell-loss. The second difference is in our bandwidth request algorithm at the source, which uses delay at the source buffer instead of previous frame sizes, for the computation of the rate to be requested from the network. A concern with using previous frame sizes for computing new requests is due the effect of adaptive quantization. Adaptive quantization is commonly used to lower the source bit-rate when the transmission rate available to the video source is inadequate [25, 26]. When the network is temporarily congested, and is unable to grant the source its requested bandwidth, the source sets its target encoder rate to the available rate, thus being able to reduce its bit-rate at the expense of quality. Future source requests based on previous reduced bit-rate, lower quality frames could thus be lower than what is needed for best quality, even when the network becomes uncongested.

This thesis investigates a simple and effective source bandwidth request policy based on the delay experienced by frames at the source buffer coupled with a switch bandwidth request policy which explicitly includes the minimum cell-rate for computing the rate allocation. It does not require prior knowledge of statistics for the

video source, assumption of video traffic models or complex calculations and offers good QoS for a wide range of propagation delays and network loads.

1.3 Thesis Objectives

1. To implement a simple bandwidth request algorithm for real-time video that effectively captures the bandwidth needs of different video sources
2. To study the suitability of the bandwidth request scheme at the source with an ABR explicit rate control scheme implemented as the switch bandwidth allocation policy for video transmission. And make any adaptations necessary for video traffic.
3. To demonstrate the sensitivity of the composite source-switch system to variations in traffic loads and propagation delays

1.4 Organization of Thesis

1. Introduction
(this chapter)

2. Overview of Video over ATM

This chapter takes a look at MPEG video compression and ATM technology. We describe the requirements for video services and issues pertaining to the transport of video over ATM.

3. The Available Bit-Rate Service Class

In this chapter we go over the ABR service class as specified by the ATM Forum. An important part of this chapter is the description of explicit rate control schemes. This includes schemes that have been approved by the ATM

Forum for data traffic over ATM and two recently proposed improved explicit rate schemes.

4. Source and Switch Algorithms

This chapter provides descriptions of the algorithms that we have proposed and those that we have used to compare them against. This includes the bandwidth request algorithms at the source and the explicit rate bandwidth allocation algorithms at the network switches. We also give an overview of the models used in our simulations.

5. Simulations for Performance Analysis

Here we describe in detail the simulations that were carried out in order to analyze the performance of our algorithms and to evaluate the sensitivity of the system to different traffic loads and propagation delays.

6. Conclusions and Future Results

Finally we present a summary of the thesis with conclusions drawn from this study, and directions for further research.

Chapter 2

Overview of Video over ATM

In this chapter we describe MPEG video compression and quality requirements for real-time video transmission. We also give a brief overview of ATM technology and issues relating to video transmission over ATM.

2.1 MPEG Video

The Moving Picture Experts Group (MPEG), a working group of the International Standards Organization (ISO), has been working on standards for video compression since 1988. The first of these standards is MPEG-1 which was developed for the storage of video and associated audio on digital storage media. The supported media ranged from CD-ROM, DAT, tape drives to writable optical drives and telecommunications channels like ISDN and LANs which could support the target compressed video bit-rate of 1.5Mbps. After the technical foundations of the MPEG-1 standard (ISO/IEC 11172) were laid out in 1990, work on a new standard for television quality transmission at target bit-rates higher than 1.5Mbps and upto 60Mbps, namely MPEG-2, was started. The MPEG-2 standard (ISO/IEC 13818) supports a wide range of features which includes constant bit-rate transmission, variable bit-rate transmission, random access, scalable decoding, fast forward and fast reverse playback. MPEG-2 video streams can be stored in various digital storage media as well

as transmitted over communication networks which include fibre, cable, ISDN/ATM and satellite. It uses two system stream formats, the Program Stream similar to the MPEG-1 video stream for video, audio or other coded bit-streams with a common time base and the Transport Stream for combining one or more Program Streams with independent time bases into a single stream. It also supports different video service hierarchies also known as profiles for different kinds of applications. In order to accommodate the demands of multimedia communications for high interactivity, a new standard namely MPEG-4 is currently being developed [27]. The aim is to support low bit-rate (64Kbps) audio/video conferencing and offer users an ability to compose and manipulate objects within an image.

2.1.1 Basic MPEG Video Compression

Figure 2.1: MPEG forward and backward prediction

While there are differences in the encoder/decoder algorithms between the different MPEG standards, the fundamental algorithm remains the same. A good description of MPEG video compression can be found in [28]. Three types of pictures or frames, namely I, P, and B frames, are used in the compression. I frames or intra-frames are coded based on spatial redundancy within the uncompressed frame and do not have dependency on other frames. P frames or predicted frames are coded with reference to a past frame which may be an I or P frame, thus making use of temporal redundancy. Each P frame is divided up into macro-blocks some of which may be encoded the same way as I frames and others that are encoded predictively using motion estimation. Motion estimation techniques consider future frames to be translations of previous frames. The translation is not constant for the whole frame but computed individually for each macro-block. The array of translation vectors corresponding to individual macro-blocks are called motion vectors. B frames or bi-directionally predicted frames are coded with reference to a past and a future frame and thus offer the highest possible compression of the three types. The frame can be reconstructed by interpolating from a past and future frame by adding a correction term. The correction term which is the difference between the original frame and the interpolated estimation of the frame is the only information that is transmitted. This offers considerable reduction in the number of bits used to encode this frame. A simple pictorial representation of the MPEG prediction described above is shown in Fig 2.1. Next, the information from the macro-blocks, which are the original picture elements for intra-coded macro-blocks, and predicted differences for predicted macro-blocks, are transformed into the frequency domain using Discrete Cosine Transform (DCT). This is followed by quantization of the DCT coefficients, which is a lossy process, and entropy coding to further increase the compression.

Usually, I and P frames are separated by two or three B frames with the sequence IBBPBBPBBPBB or IBBBPBBBBPBBBBPBBBB repeated, where the sequence is called a group of pictures (GOP). Each frame is divided into segments called macro-blocks,

several of which make up a slice. Several slices make up a frame and a sequence of frames beginning with and containing a single I frame forms a GOP. MPEG video frames are partitioned into the above layers for ease of applying compression techniques, and for separating the information required at each layer.

2.1.2 Quality Requirements for Video

The three main factors that affect the quality of transmitted video are end-to-end delay, variation in the end-to-end delay also known as jitter, and cell-loss. The perceptual quality of the video at the receiver depends on the delay, jitter and loss suffered by the received video frames. Below, we present a brief discussion of each of these factors.

Delay

Real-time applications of an interactive nature have low tolerance to delay so as to provide good perceptual quality to the users. Little or no degradation in quality is observed for a one way delay of 150ms while serious degradation is observed for delays greater than 400ms [29]. The major sources of delay in video transmission over a network are encoding and decoding delays for the video, propagation delays and queuing delays at the end systems and network switches, packetization and re-assembly delays, transmission delays and switching delays [2]. The propagation delay can be a major bottleneck since it depends on the distance between the source and destination end systems and the speed of the transmission medium. The video coding delay depends on the chosen codec algorithm and the speed of the hardware used. The queuing delays can be controlled by appropriate negotiation of bandwidth and the priority and type of the service chosen for the transmission. The packetization and re-assembly delays depend on the protocols being used, while the switching delays depend on the offered traffic load and processing power of the network switches.

Jitter

Jitter is another major problem in video transmission. It is the variation in the end-to-end delay that leads to a variability in the arrival times of frames at the receiver while they are produced at a constant rate at the sender (i.e 30 frames per second or 25 frames per second). This can cause three potential problems [30], namely loss of frames, gaps in the video at the receiver end, and high display latency. Loss of frames occur if many frames arrive at the receiver in a burst overloading the available resources (buffer/processor capacity) at the receiving end, resulting in frames being discarded. Gaps occur in the display of frames at the receiver if a frame is not available as soon as the preceding frame has been displayed. Since frames are displayed at a fixed rate, varying arrival times at the receiver makes it difficult for a smooth display. High display latency refers to an increase in the time spent by a frame at the receiver display queue which can grow without bound owing to the inequality between the fixed service rate and the varying arrival rate. Jitter can be reduced by choosing network transmission schemes that can guarantee either an absolute or statistical bound on the maximum delay experienced in the network queues. At the receiver, an adaptive algorithm can be used to compensate for the jitter by choosing an appropriate display latency so that most frames can be displayed as soon as they arrive. An extensive treatment of handling jitter can be found in [30] and references therein.

Cell Loss

Cell-loss occurs as a result of overflowing network or receiver buffers, and corruption of frames due to bit errors during transmission. Frames can also be discarded if they are delayed beyond their play-out time. The impact of the loss depends on the number of cells lost in a region, its shape and the burstiness of the loss [29]. The degradation was found to be higher for random losses than clustered losses in a study by Iai, Kitawaki and Kurita [31]. Loss of entire frames is perceived as jerkiness in the video displayed at the receiving end. Loss of I frames cause more serious

degradation in image quality than loss of B frames. The activity level of the video also determines the impact of the loss. Loss during high activity scenes can be much more disruptive than that during lower activity scenes. Several schemes are used to detect and recover from loss. Bit errors and loss of video frames can be detected by cyclic redundancy checks (CRCs) and sequence numbers respectively, for packets arriving at the receiver. Recovery from detected losses is possible through forward error correction (FEC) and loss concealment. Forward error correction schemes however add additional overhead, increased delay and complexity. An alternative is cell-loss concealment which works by extrapolating from the pixels around the area of loss (spatial) or from corresponding pixels in previous frames (temporal). A drawback is that any residual errors after concealment will propagate into future frames through temporal prediction.

2.2 ATM Technology

Asynchronous Transfer Mode supports data, video and voice over local or wide area networks, under a single switching fabric. Its fixed size cells of 53 bytes each, allow switching at high speeds and constant processing delays making it possible to support real-time applications. This differentiates it from variable-length frame technologies like Frame Relay and X.25. In the sub-sections that follow, we first describe the ATM protocol stack, followed by a brief overview of ATM service categories for video. Next we discuss how the ATM traffic management functions enforce the traffic contract between the network and its users.

2.2.1 The ATM Protocol Stack

The ATM protocol stack includes specifications for communication between different layers from the transmission medium at the lowest level all the way to the service adaptation layer. The layers include the physical layer, the ATM layer, and the adaptation layers. The physical layer includes the underlying transmission medium (eg.

SDH/SONET) and a convergence layer that allows a medium-independent communication with the upper layers. Next comes the ATM layer which takes care of flow control, Virtual Path/Circuit Identifier (VPI/VCI) mapping, multiplexing and demultiplexing of signals and the QoS management. It supports two main interfaces, namely the User-Network Interface (UNI) and the Network-Node Interface (NNI). The UNI is used for connecting an ATM end-system with a public ATM network, while the NNI is used between switches in a public ATM network. At the next higher level is the ATM Adaptation Layer (AAL). This layer is subdivided into a Convergence Sub-layer (CS) and the Segmentation and Re-assembly (SAR) sub-layer. The SAR segments application-level frames into ATM cells and reassembles ATM cells into application frames. The CS has a common part that is responsible for packet framing and error detection and a service specific part that provides for connection-less or connection-oriented service. Different classes of AALs corresponding to different QoS classes are currently available. These classes are described below:

- AAL 1 supports Constant Bit-Rate connection-oriented service over ATM. It provides timing recovery information through Synchronous Residual Time-stamps (SRTS) and sequence numbering for ATM cells to facilitate detection of lost or inserted cells.
- AAL 2 is optimized for Variable Bit-Rate service. In addition to providing sequence numbering, it also marks the beginning, middle and end of higher layer frames.
- AAL 3/4 is suited for Available Bit-Rate type data services. It allows for both connection-less and connection-oriented services. It includes additional overhead to support the multiplexing of multiple users sharing the same Virtual Circuit (VC) and allows for bi-directional point to multi-point connections. The overhead takes up 4 additional bytes in a 53 byte ATM cell leaving 44 bytes for the actual payload.

- AAL 5 has been developed to provide a reduced overhead alternative for point to point and uni-directional point to multi-point data service. It eliminates the sequence numbering and multiplexing identification fields provided by AAL3/4. It expects application level time-stamping and does not provide this functionality.

The ATM AAL layer interfaces with higher layer non-ATM protocols which include the Ethernet or Token Ring protocols for LAN Emulation, the Transport Control Protocol (TCP) and Internet Protocol (IP) at the network level and application level protocols like Real-time Transport Protocol (RTP).

2.2.2 ATM Service Parameters

Traffic sent over ATM service categories are commonly described by two sets of parameters - traffic parameters and quality of service parameters. Traffic parameters describe the characteristics of a traffic stream which include Peak Cell Rate (PCR), Sustainable Cell Rate (SCR), Burst Tolerance (BT), Minimum Cell Rate (MCR) and Cell Delay Variation Tolerance (CDVT). Quality of service parameters describe the quality of service (QoS) required by a traffic stream which the network is expected to guarantee for the duration of the connection. These include peak-to-peak Cell Delay Variation (CDV), mean and maximum Cell Transfer Delay (CTD) and Cell Loss Ratio (CLR). Not all traffic or quality of service parameters need to be specified for all the ATM service categories. Real-time streams like video and audio are usually very sensitive to CTD and CDV and to a certain extent on CLR as well. Non real-time streams which are data traffic are usually very sensitive to CLR and more tolerant to CTD and CDV. Bursty traffic is usually specified by a combination of PCR, SCR and BT. The above parameters as defined by the ATM Forum [5] are further described below.

Peak Cell Rate (PCR) and Minimum Cell Rate (MCR)

Peak Cell Rate is the upper bound on the bandwidth requirement for a connection. It is expected that the cell-rate of the traffic source will not exceed this value. The Minimum Cell Rate is the lowest cell-rate that the traffic source can tolerate. The network is always expected to allocate rates greater than or equal to the MCR for each connection.

Sustainable Cell Rate (SCR)

The Sustainable Cell Rate is commonly used to describe bursty traffic whose instantaneous bandwidth requirement varies over a wide range of values. The long term average rate for such a connection is called the SCR. This value can then be used by the network for bandwidth allocation so that the long term needs of the connection are met.

Burst Tolerance (BT)

Burst Tolerance is another parameter used in addition to SCR and PCR to describe bursty traffic. It is related to the Maximum Burst Size (MBS) which is the maximum number of contiguous cells that may be transmitted at the PCR. The formula relating them is:

$$BT = \frac{(MBS - 1)}{(1/SCR - 1/PCR)}$$

Cell Loss Ratio (CLR)

This is defined as the ratio between the cells lost in the network and the total number of cells transmitted. Cell-loss occurs due to errors at the physical layer and dropping of excess cells when the network switch buffers get congested.

$$CLR = \frac{\text{cells lost in transmission}}{\text{total cells transmitted}}$$

The ATM network allows negotiation of CLR for CBR, r.t and n.r.t VBR and guarantees it for the duration of the connection. At the time of setup, each switch along the connection either accepts or rejects the requested CLR based on its ability to support it.

Cell Transfer Delay (CTD)

The total delay experienced by a cell as it traverses the network from the entry point to the exit point is called cell transfer delay. It includes network queuing delays and processing delays at the switches and the propagation delay between the source and the destination. In ATM networks, the maximum CTD can be negotiated for the CBR and r.t VBR traffic categories. The maximum CTD is defined as the “ $1 - \alpha$ ” quantile (for a small α) of the CTD experienced during the life of a connection. At the time of connection setup, switches along the path add their contribution to the accumulated maximum CTD so far. Thus signaling protocols used during the connection setup are able to estimate the maximum CTD for the path from the source to the destination in a cumulative fashion across the network.

Cell Delay Variation (CDV)

The variation in the CTD is termed Cell Delay Variation. This can be measured in several ways. The most common measure is peak-to-peak CDV. The term “peak-to-peak” refers to the difference between the best CTD (lowest value) and the worst CTD (highest value). The lowest value corresponds to the fixed propagation delay for the connection, while the worst CTD corresponds to the value above which the CTD increases only with a very small probability.

Cell Delay Variation Tolerance (CDVT) is related to the CDV and is used by the network to take into consideration cells being received at rates exceeding the minimum inter cell arrival time ($1/\text{PCR}$) while performing policing operations.

2.2.3 ATM Service Categories

The ATM Forum's TM 4.0 specifications [5] identify five major traffic classes which have been derived from the original Bearer Capability (BC) service model for Broadband ISDN (B-ISDN) [32]. The main focus of the TM 4.0 classes is whether the service provided is real-time or non real-time and connection-oriented or connection-less service. For real-time traffic, Constant Bit-Rate (CBR) and real-time Variable Bit-Rate (r.t VBR) service are the two classes specified. For non-real-time traffic, the three available classes are non-real-time VBR (n.r.t VBR), Available Bit-Rate (ABR) and Unspecified Bit-Rate (UBR) service. Each of these services are further elaborated below.

Constant Bit-Rate (CBR)

This service class provides connection-oriented, circuit emulation service over ATM. At the time of connection setup, a constant bandwidth (usually the PCR) is negotiated between the source and the network for the duration of the connection. Quality of service parameters like CTD, CDV and CLR are also negotiated and guaranteed for the connection. CBR cells are given the highest priority over other ATM traffic classes during their transfer across the network. AAL-1 adaptation layer is used for the CBR service since it provides timing recovery information for synchronization.

Real-time Variable Bit-Rate (r.t. VBR)

While CBR is appropriate for traffic sources with close to constant bandwidth requirements and with very stringent quality of service demands, it is less efficient for bursty sources. Bursty traffic sources like video show a wide variation between the peak and mean cell generation rates. Thus when the peak bandwidth requirement is reserved for the duration of the connection, the utilization of the network bandwidth drops during transmission at the mean or lower rates. The VBR service allows efficient use of network bandwidth as it exploits statistical multiplexing gain. At the

time of connection setup, the PCR, MBS and SCR of the source traffic along with quality parameters such as CTD, CDV and CLR to be provided by the network are negotiated. During the transmission, the quality parameters are guaranteed as long as the source traffic conforms to a long term mean cell-rate of SCR and sends no more than MBS cells at the PCR. AAL-2 adaptation layer is used for VBR transport though AAL-5 may be used when timing recovery is provided by higher level protocols.

Non-real-time Variable Bit-Rate (n.r.t VBR)

This service is similar to real-time VBR but does not provide stringent real-time quality of service. Only the mean CTD is guaranteed for the connection as long as it conforms to the SCR. It is thus suitable for bursty data service that requires a bounded mean delay. Two such applications include transaction processing and Frame Relay transport.

Available Bit-Rate (ABR)

The Available Bit-Rate (ABR) service as the name suggests is a best effort service which is also closed-loop. At the time of setup, the source specifies the PCR, MCR and the maximum CLR that it requires. All that the network guarantees is the MCR while the CLR is minimized as long as the source responds to the presence or absence of congestion in the network by reducing its transmission rate when the network is congested and increasing when allowed. Several ABR rate control schemes have been proposed which aim to minimize cell-loss, provide fairness between competing ABR sources and allow complete utilization of the available network bandwidth. This service is suitable for data traffic which is highly sensitive to loss but can tolerate delay. AAL-3/4 is currently used for ABR, though AAL-5 may soon replace this. A detailed discussion of the ABR service class is provided in the next chapter.

Unspecified Bit-Rate (UBR)

The Unspecified Bit-Rate Service is completely best-effort and open-loop. No guarantees are provided by the network. It is thus unsuitable for any real-time applications. Data sources with vague service requirements that can detect and handle loss are most appropriate for this type of service. Examples of such applications include E-mail, Telnet, fax, file transfers, and LAN interconnection. While it might appear to be highly risky to use UBR, ATM congestion control schemes which include traffic shaping and packet discard policies may make UBR service a fairly reasonable option. AAL-3/4 or AAL-5 maybe used for the UBR service.

2.2.4 The ATM Traffic Management

Connections are provided different QoS guarantees depending on the type of ATM service category they chose. We have discussed the different kinds of services that ATM provides, particularly for real-time video. For each service category, the traffic source is expected to conform to the negotiated behavior and the network is expected to guarantee the negotiated QoS. In order for this contract between the source and the network to be upheld, ATM provides certain traffic management functions. The most important functions include Connection Admission Control (CAC) and Usage Parameter Control (UPC).

Connection Admission Control (CAC) is used by the network to decide whether to accept or reject a new call based on its current service commitments and the parameters requested by the new call. A new call is accepted only if the network resources can support it while maintaining the QoS for existing connections. Thus CAC schemes need to determine whether the available resources are adequate for the new connection and also gauge the effect of admitting a new connection on existing connections.

Usage Parameter Control (UPC) is used to monitor existing connections for violation of the traffic parameters negotiated at setup. This is to ensure that sources

do not exceed their quota of network usage. This is performed through the Generic Cell-Rate Algorithm (GCRA). The algorithm implements a leaky-bucket controlled by two parameters - increment (I) and limit (L) - often referred to as GCRA(I,L). The size of the bucket is limited to $I+L$. Every time a cell arrives, the bucket counter is incremented by I. The bucket is emptied at regular intervals by 1 unit. For CBR, a single leaky- bucket is used with $I = 1/PCR$ and $L = CDVT$. Cells arriving earlier than $1/PCR - CDVT$ of the previous cell are tagged as non-conforming while the rest are passed as conforming. Thus the value of the CDVT determines the leeway in accepting a cell that might arrive earlier than expected due to the presence of jitter. For VBR, two leaky buckets are required to police the PCR, MBS and SCR. The first leaky bucket is the same as that for CBR, while the second bucket implements a GCRA($1/SCR, BT+CDVT$). The GCRA is the algorithm of choice of the ATM Forum and is widely implemented. However, studies [9], [10] indicate that the leaky-bucket approach might not be efficient in policing long term average rates like the SCR.

2.3 Bandwidth Prediction for Video over ATM

An important area in the study of video transmission over ATM is the use of prediction schemes for long range bandwidth prediction and smoothing of burstiness of compressed video. Below we present some such prediction schemes.

2.3.1 Long-Range Prediction for RCBR

Compressed video shows burstiness at multiple time-scales. Long-Range prediction schemes must be able to filter out the short-range burstiness information so that only the long-range trends are captured. Bandwidth renegotiation schemes for Renegotiated Constant Bit-Rate (RCBR) [11] require an estimation of future bandwidth requirements over a long time range in order to avoid frequent renegotiations due to

fluctuations in bandwidth caused by short-term burstiness of MPEG video streams. We present two such schemes proposed for the RCBR service. The first is a causal scheme for real-time video for which prior information is not available while the second is suitable for stored video.

A Causal Auto-Regressive Prediction Scheme

The Causal Auto-Regressive (AR) rate predictor scheme [11] uses an AR(1) prediction with buffer thresholds to determine the new rate required. Every time-slot (i), the rate for the next time-slot ($\hat{r}_{i+1}$) is predicted from the current arrival rate (r_i) of video cells at the source buffer. A time constant T reflecting the long-term rate of change of the rate function and the current buffer occupancy (b_i) in excess of the upper buffer occupancy threshold (B_h) are used in the prediction as follows:

$$\hat{r}_{i+1} = (1 - T^{-1})\hat{r}_i + T^{-1}(r_i + \max(b_i - B_h, 0))$$

The terms T^{-1} and $1 - T^{-1}$ are used to weight the actual incoming (arrival) rate plus the amount of bandwidth needed to flush out the source buffer at the i th time-slot and the predicted rate for the i th time slot based on earlier time slot calculations. A new bandwidth (s_{new}) is computed as:

$$s_{new} = \lceil r_{i+1} / \Delta \rceil \Delta$$

where Δ is the bandwidth allocation granularity.

If ($b_i > B_h$) and ($s_{new} > s$) or ($b_i < B_l$ and $s_{new} < s$) then the new bandwidth is requested.

Double Threshold Moving Window (DTMW) prediction

A technique for optimal renegotiation for RCBR so that the frequency of renegotiations is minimized, has been proposed by Huang et al [33]. A double-threshold moving window predictor has been used successfully, for predicting long-term bandwidth needs for RCBR service. It makes use of four thresholds, one for the arrival

rate of video cells at the source buffer ($T1$), another for the upper and lower limits of buffer occupancy at the source buffer. and a moving window size threshold ($T2$) to determine when the change in request is to be made. The DTMW module sets consecutive bits in an array of size equal to the chosen window size (N) in a round-robin fashion. At every time-slot 't', if the arrival rate ($Y(t)$) at the source buffer is greater than or equal to the threshold $T1$, the bit corresponding to this time-slot is set. The total number of set bits is counted and if this value is greater than or equal to $T2$, the DTMW outputs 1 else it outputs 0. This is expressed for the i th time-slot computation as follows:

$$S_N(i) = \sum_{j=0}^{N-1} I(Y(i-j) \geq T1); \text{ if } S_N(i) \geq T2 \text{ then output 1 else 0}$$

The inverse DTMW (IDTMW) module checks for arrival rates less than or equal to $-T1$ to take care of falling rates. A High/Low buffer occupancy controller (HLBC) monitors the buffer occupancy and sets two bits namely the HBC and the LBC. When the buffer occupancy falls below the low buffer threshold, the LBC is set to 1 else it is set to zero. Similarly, when the buffer occupancy increases above the upper buffer threshold, the HBC is set to 1 else it is set to 0. Every frame-time (33ms), the DTMW, IDTMW (Inverse DTMW), HLBC (High/Low Buffer Occupancy Control) determine a rate change depending on the values of the DTMW, IDTMW, HBC and LBC. An increase equal to $T1$ is requested if either the DTMW or the HBC are 1, while a decrease of the same magnitude is requested if either the IDTMW or the LBC are 1. The value $T1$ is determined from the stored video sequence apriori, by minimizing the number of renegotiations required for different magnitudes of bandwidth granularity. The values of $T2, N$ and HBC are chosen so as to respond mostly to long-range burstiness as opposed to short-range burstiness.

2.3.2 Prediction for Cell-Rate Smoothing for VBR

Video transmitted over the Variable Bit-Rate (VBR) service of ATM is provided statistical guarantees as described earlier in this chapter. Submitting a smoother

video stream decreases the probability of cell-loss at the network and allows increased statistical multiplexing gains [34, 35]. Source video smoothing comes at the expense of increased delay and buffer requirements. A common method for smoothing VBR traffic is through the introduction of delay. The frames from the video source are delayed for a small period of time and then transmitted at the average arrival rate during this period. This enables the averaging of high bit-rate frames with low bit-rate ones resulting in a smooth average rate. Work done by Reibman et al [36] has shown the effect of introducing different delays corresponding to integer multiples of frame times. The results show that a delay corresponding to two frame times offers a good smoothing to delay tradeoff. Another method is the use of prediction for estimating upcoming frame-sizes so that the transmission rate can be set to the average of the predicted frame sizes. Below we discuss two techniques for predicting future frame sizes from current frame sizes.

A Simple Prediction Scheme

A simple scheme for predicting frame-sizes of upcoming frames from previous frame-sizes. long term mean frame-size and standard deviation of the distribution of frame-sizes was proposed by Pancha and El Zarki [7]. They have used the assumption that the probability of the frame bit-rate (or cell-rate) changing by more than one standard deviation in a subsequent frame of the same type is negligible. This has been based on empirical observations of statistical properties of compressed video presented in the paper. The predicted cell generation rate for the next frame (c_n) is obtained from the long term mean (μ) and standard deviation (σ) of the cell generation rates as below:

$$\hat{c}_n = \max(\mu, c_{n-1} + \sigma)$$

The long term mean and standard deviation are computed apriori from the original video before being transmitted. Thus its application to real-time video is limited.

Smoothing with Rate Bounding

An extensive VBR traffic smoothing scheme along with an on-line prediction of the mean frame size was proposed in [34] . The aim was to keep the transmission rate between a lower bound (so as to eliminate delay violations at the source) and an upper-bound (so as to prevent complete emptying of the source buffer leading to large idle periods). The average between the maximum of a sequence of lower bounds and the minimum of a sequence of upper bounds determined at regular intervals, is used as the transmission rate. Forecasts are made from current frame-sizes (X_n) to obtain the frame sizes of upcoming frames (X_{n+k}). The formula used to predict the size of the k th frame following the current frame n is based on the high correlation between successive frame sizes for frames of the same type (namely intra/inter/bidirectional frames) . A correlation coefficient R of around 0.98 is applied as shown below:

$$X_{n+k} = \mu_n + R^k(X_n - \mu_n)$$

The mean (μ_n) is computed on-line using an exponential averaging with a weighting constant C_p (set to 0.99 for the results presented) as:

$$\mu_n = C_p \mu_{n-1} + (1 - C_p) X_n$$

The delay constraints at the source are computed with the basic idea that all bits arriving at the source buffer at the n th time-slot i.e during the interval $[(n-1)D, nD]$ must be sent out by $nD + \Delta$ at the latest. D is the time interval between consecutive time-slots and is usually one frame time and Δ is the allowed delay for transmission. The maximum of the lower bounds are computed using the arrival rate into the buffer $Y(t)$, frame-sizes X_{n+k} and buffer occupancy $C(t)$ at the current time slot t for $j=0,1,2$... as:

$$L_j(t) = \max_{0 \leq l \leq j} \left(\frac{C(t) + Y(t) + \sum_{k=1}^l X_{n+k}}{n - t + l + \Delta} \right)$$

And the minimum of the upper bounds are computed as:

$$U_j(t) = \min_{0 \leq l \leq j} \left(\frac{C(t) + Y(t) + \sum_{k=1}^l X_{n+k}}{n - t + l} \right)$$

The average between the lower and upper bounds is chosen as the transmission rate.

Further enhancements to the algorithm are provided in [34]

Chapter 3

The Available Bit-Rate (ABR) Service Class

In this chapter, we describe the Available Bit-Rate (ABR) service category supported by ATM networks for data traffic. First, we present the framework for ABR traffic management. Next we give a brief overview of the available congestion control schemes proposed for controlling ABR traffic. Finally, a few explicit rate control schemes are presented.

3.1 Basic Service Model for ABR

Data traffic has vague requirements for bandwidth and delay but is very sensitive to loss [23]. This makes feedback of the network status to the transmitting sources of vital importance in order to prevent loss due to network congestion. The Available Bit-Rate (ABR) service class provides a closed-loop feedback control for data traffic. The objective of ABR is to prevent data loss on the one hand and network under-utilization on the other. The network switches monitor their port buffers for congestion and provide feedback to the traffic sources to indicate their current congestion status. Data traffic sources are expected to reduce their usage when the network is heavily loaded and increase their usage if they so desire, when the network can support it. Below, we discuss the structure of the basic units for the exchange of

feedback information, namely the resource management (RM) cells, followed by a brief description of rules that the source end systems (SES) and the destination end systems (DES) are expected to follow. A detailed treatment of rules for source and destination end systems can be found in [37]. As a further reference, the ATM Forum Traffic Management Specification [5] provides a complete guide-line for ABR traffic management over ATM networks.

3.1.1 Resource Management (RM) Cells

<i>Name</i>	<i>Description</i>	<i>Default</i>
PCR	Peak Cell Rate - the maximum rate at which the source can transmit at	none (mandatory)
MCR	Minimum Cell Rate - the lowest rate that the source can tolerate	0
ICR	Initial Cell Rate - the rate at which a source is allowed to send initially after an idle period	PCR
TBE	Transient Buffer Exposure - the negotiated number of cells that the source rate is expected to be limited to, during startup periods and before receiving the 1st BRM	16,777,215
FRTT	Fixed Round Trip Time - accumulated sum of all the RM cell fixed delays plus propagation delays in the round-trip call path	none (accumulated)
RDF	Rate Decrease Factor - the amount by which the source is allowed to decrease its transmission rate upon receipt of a BRM cell with the CI bit set	1/16
RIF	Rate Increase Factor - the amount by which the source is allowed to increase its transmission rate upon receipt of a BRM cell with the CI and NI bit unset	1/16

Table 3.1: ABR parameters signaled at setup

Resource Management cells are the units of information exchange between the

<i>Field</i>	<i>Octet</i>	<i>Bit(s)</i>	<i>Description</i>
Header	1-5	all	ATM header
ID	6	all	Protocol identifier (ABR)
DIR	7	8	Direction
BN	7	7	BECN cell
CI	7	6	Congestion Indication
NI	7	5	No Increase
RA	7	4	Acknowledge
Reserved	7	3-1	Reserved
ER	8-9	all	Explicit Rate
CCR	10-11	all	Current Cell Rate
MCR	12-13	all	Minimum Cell Rate
QL	14-17	all	Queue Length
SN	18-21	all	Sequence Number
Reserved	22-51	all	Reserved
Reserved	52	8-3	Reserved
CRC-10	52 53	2-1 all	Cyclic Redundancy Check-10

Table 3.2: Fields in an RM cell

Figure 3.1: Fields in the ATM cell header

sources and the network. The term Forward RM cells (FRM) refer to RM cells traveling in the forward direction from the source to the destination. Backward RM cells (BRM) refer to RM cells traveling in the reverse direction from the destination to the source. They have special fields for different parameters some of which can be modified by the network. Some of the parameters (Table 3.1) are negotiated at the time of setup of each ABR connection, using signaling protocols. Table 3.2 shows the various fields of an RM cell. Of these fields, the most important are further elaborated

below:

- **Header** The header is the standard ATM cell header (Figure 3.1). It contains sub fields for various information. The payload type identifier (PTI) field indicates whether the cell is an RM or data cell. The payload type identifier is a 3-bit field, the middle bit of which is used for Explicit Forward Congestion Indication (EFCI). It is set to 1 when congestion is detected in the network and 0 otherwise. This is often used in ABR rate control as will be seen later in this chapter. The cell loss priority (CLP) bit indicates the relative importance of the cell. When a switch is congested and has to drop cells in order to relieve the congestion, cells with a lower priority are dropped before cells with a higher priority. The connection is identified by the virtual circuit identifier (VCI) and the virtual path identifier (VPI) fields.
- **DIR** This bit is set to 0 or 1 depending on whether the RM cells are traveling toward the destination or toward the source. Forward RM cells traveling from the source to the destination in the same direction as the data cells have the DIR bit set to 0. Backward RM cells traveling in the reverse direction from the destination to the source have the DIR bit set to 1.
- **BN** This bit indicates whether the RM cell is a Backward Explicit Congestion Notification (BECN) cell or a regular RM cell. BECN cells are generated by the destination or a network switch to inform the source of the presence of congestion. Non-BECN cells are regular RM cells generated by the source.
- **CI** This is the congestion indication bit which can be set either by network switches or the destination to indicate the presence of congestion. Switches monitor their port buffers and signal the presence of congestion by setting the CI bit to 1. At the destination, the CI bit of the RM cell is set to 1 if the most recent data cell it received had the EFCI bit in the cell header set to 1. When

an RM cell is received back at the source with the CI bit equal to 1, it decreases its allowed cell-rate.

- **NI** This is the “no increase” bit which is set by the network to signal an impending congestion. On receiving an RM cell with the NI bit set to 1, the source is only expected to not increase its allowed cell-rate . It does not have to decrease its rate.
- **ER** The Explicit Rate field in an RM cell is initialized at the source to the rate that it requires from the network. Network switches that use explicit rate congestion control schemes are allowed to decrease this rate but not increase it. Eventually, when the RM cell is received back at the source, it sets its allowed cell rate to the value in this field.
- **CCR** The Current Cell Rate field carries the current allowed cell-rate at the source at the time of transmitting the RM cell. Though this value is not altered as the RM cell passes through the network, it could be used by network switches in making computations for explicit rate or congestion indication.
- **MCR** The Minimum Cell Rate field is set at the source to the minimum rate that was negotiated for the connection at the setup time. Again, this value is not altered by the network but is used by switches to ensure that the source allocation does not go below this value.

3.1.2 Source End System (SES) Rules

- Sources are expected to always transmit at rates less than or equal to their computed allowed cell-rate. They need not however, transmit at rates below their negotiated Minimum Cell Rates (MCRs). They are expected to follow the appropriate policy for computing the allowed cell-rate depending on the values of the CI bit, NI bit and ER field value of returning RM cells.

- At the beginning of transmission, the source has to first send out an RM cell. The initial transmission rate is the negotiated Initial Cell Rate (ICR). Once transmission has begun, if the source is idle for more than 500ms. it is expected to send out an RM cell and revert to its ICR, until it is allowed a larger transmission rate by the network.
- Sources can send out two kinds of RM cells - in-rate cells and out-of-rate cells. In-rate RM cells are transmitted every $\text{NRM}-1$ data cells, where NRM is set to a default of 32. Such in-rate RM cells are sent with their cell loss priority (CLP) bit set to 0. Out of rate RM cells are sent by the source. switches or destinations in exceptional circumstances. When the time elapsed since transmission of the last RM cell exceeds 100ms, and a minimum number of data cells have been transmitted. the source is required to send an RM cell regardless of the allowed cell-rate i.e out-of-rate. All such out-of-rate RM cells must have their CLP bit set to 1. This allows the network to discard them if there is congestion.

3.1.3 Destination End System (DES) Rules

- Destination end systems are expected to turn around any RM cells they receive back to the source after carrying out the permitted modifications. The ER field may be reduced if the destination is experiencing congestion. CI bit and NI bits may be set depending on the presence of congestion. The DIR field is set to indicate the backward flow of the RM cell.
- When EFCI marking is used by the network to indicate congestion, destinations are expected to monitor the EFCI bits on incoming data cells and store the most recent value. This value is used to set the CI bit of RM cells routed to the source.
- When the destination wishes to issue warnings of congestion back to the source even though they have no RM cells to turn around to the source, they may send

out-of-rate RM cells known as BECN cells to the source. Such out-of-rate cells are to have their CLP bit set to 1 so that they may be discarded by the network in the event of congestion.

3.2 Categories of Rate Control Schemes

A variety of schemes for feedback-based rate control have been proposed over the last few years. A comprehensive survey of recent advances in ABR congestion control, can be found in [22]. Discussions of the evolution of the ABR Rate Control framework can be found in [38] and [23]. A general frame-work for classifying ABR rate-based schemes can be found in [39]. Below, we discuss four main types of congestion control schemes - namely, single-bit feedback, explicit rate control, virtual source/destination control and credit-based feedback. Figure 3.2 provides an overall view of ABR rate control schemes.

3.2.1 Single bit feedback

Data cell headers contain a bit field called the Explicit Forward Congestion Indication field (EFCI) as seen earlier, which can be set to 0 or 1 indicating the absence or presence of congestion respectively. The EFCI bit in each data cell is set to 0 at the source and as the cell passes through congested switches, the bit is set to 1. The switches determine their congestion status depending on whether the queue-length at the port buffer goes beyond a threshold or not. If it exceeds the threshold, the EFCI bit of VCs passing through the switch port is set. At the destination, the EFCI bit value of the most recent data cell is remembered. Every periodic interval, the destination monitors the most recent EFCI bit value and if it is 1, an RM cell is sent to the source, with the CI bit set. The source uses an additive increase/ multiplicative decrease policy to modify its transmission rate in response to the feedback. When an RM cell with the CI bit set is received, the source decreases its allowed cell-rate

Figure 3.2: An overview of ABR rate control

(ACR) by a multiplicative rate reduction factor. As long as it does not receive an RM cell, it increases its ACR by an additive increase factor. This particular policy of sending RM cells only to decrease the rate is called “negative polarity”. If these RM cells were lost due to congestion in the reverse path, the sources would simply keep increasing their rate resulting in overloading the network. A solution to this problem was the proportional rate control algorithm.

The Proportional Rate Control Algorithm (PRCA) [40] uses “positive polarity” instead of “negative polarity”. Sources set the EFCI bit on every cell except the N th cell. When the destination receives any data cells with the EFCI bit off, it sends an RM cell to the source, which can then increase its ACR on receiving this RM cell. The source keeps decreasing its ACR in proportion to its current cell-rate, until it receives an RM cell. The main problem with this scheme was what is commonly referred to as the “beat down” problem [22], where VCs traveling through more hops suffered from an unfair higher probability of having the EFCI bit set than those VCs traveling through fewer hops. The solution to this problem is called “intelligent congestion marking”.

Intelligent congestion marking schemes selectively set the CI bit of the RM cells belonging to VCs whose cell-rates are above a certain value. Thus VCs with low bit-rates (and therefore unlikely to cause the congestion) are not affected. Sources periodically send out an RM cell indicating their current ACR in the CCR field of the RM cell, with the CI bit off. A scheme suggested in [41] determines a rate called the “fair-share” for each VC at each switch port. This rate is calculated as a moving average over the rates carried in the CCR field of the RM cells of each VC passing through a switch port. When the switch detects congestion, it sets the CI bit only in the RM cells belonging to VCs whose rates exceed their computed “fair-share”.

3.2.2 Explicit Rate (ER) feedback

Explicit rate control schemes allow a switch to specify a transmission rate for each VC passing through it. This rate depends on different factors depending on the type of explicit rate scheme chosen. Some of the important explicit rate schemes are elaborated in section 3.3. RM cells are used for the communication of the rate between the VC source and switches. Every NRM-1 data cells, each VC transmits an RM cell with the explicit rate field set to the bandwidth it desires. This field can only be marked down by switches along the VC path, to the rate each switch is able to support. When the RM cell finally reaches the source of the VC, the ER field contains the lowest rate supported by the switches along the path. The VC updates its transmission rate known as the allowed cell-rate (ACR) to this rate. Explicit rate schemes show better performance compared to single bit rate control schemes due to their quick, oscillation free response to network congestion and fairness to VCs ([20], [21]). The ATM Forum has included explicit rate control schemes in its specifications for ABR traffic.

3.2.3 Virtual Source/Destination Control

While congestion control schemes are usually end-to-end, hop-by-hop congestion control, also known as virtual source/destination control, is also allowed for ABR traffic by the ATM Forum. The term “virtual source” refers to any component along the path of transmission of a VC including the source end system and switches that forward cells for the VC to the next component. “Virtual destination” refers to any component along the transmission path that receive cells for a VC from another component. Each virtual destination is allowed to send feedback to the virtual source from which it is receiving cells to indicate congestion. Thus the feedback loop is on a per hop basis. A disadvantage of this form of control is the need for per VC queuing per switch port in order to control the rate of transmission of each VC. The advantage

is that it allows VCs to be isolated from each other, enabling a high degree of fairness. It also cuts down the time for the feedback control loop since this loop is now over a hop as opposed to the entire round trip path. The interested reader may refer to [42] and references therein.

3.2.4 Credit-based feedback

Credit-based rate control schemes [43] use credits or tokens to limit the transmission rate of VCs. The receiver periodically informs the sender about the buffer space available for the data being sent. This is done by allocating a certain number of credits for each sender according to some buffer space allocation policy at the receiver. Every time the sender transmits data cells to the receiver, it decrements its current credit balance by the number of data cells being transmitted and every time it receives credits from the receiver, it increments its credit balance. The sender is allowed to transmit only as much data as allowed by the number of credits in its current credit balance. The term “sender” refers to any component including the source end system and switches that forward data to another component along the transmission path for a VC. The term “receiver” refers to any component including the destination end system and switches that receive data being sent along the transmission path for a VC. Credit based schemes require the implementation of per-VC queuing which could add considerable complexity to the switch as the number of VCs passing through the switch increases. This was the major argument against credit-based schemes for ABR rate control as compared to explicit rate control schemes. Jain [22] discusses the prolonged rate versus credit debate in the ATM Forum and provides several references to ATM Forum contributions in that area. The ATM Forum has eventually chosen the rate-based approach over credit-based schemes for ABR traffic over ATM.

3.3 Some Explicit Rate Algorithms

In this section we present a few explicit rate control schemes which mainly differ in the way the bandwidth allocation for each data source (VC) is computed. While the ATM Forum [5] does not make the implementation of any particular explicit rate algorithm mandatory, it has approved some of them. These algorithms include EPRCA, CAPC and ERICA each of which are described in the sections that follow. We also describe two recently proposed explicit rate algorithms, namely ERAA and FMMRA. These algorithms reallocate unused bandwidth from the share of sources that are unable to use it, to sources that require additional bandwidth. Before proceeding, we define some common terms that occur in the explicit rate computations:

- *ABRCapacity* This refers to the share of bandwidth leftover for ABR connections after subtracting out the amount of bandwidth used by higher priority service like CBR and VBR. Thus it is a variable quantity for each switch port that depends on the high priority traffic commitments at a given time.
- *LoadFactor* This refers to the ratio of the total incoming cell-rate (*TIInRate*) of ABR traffic to the *ABRCapacity* of the switch port. This is also a variable quantity for each switch port that depends on the current ABR traffic and the *ABRCapacity* at the switch port.
- *EqualShare* This is the amount of bandwidth that each ABR VC is allowed if all the allocations were equal. It is equal to the *ABRCapacity* divided by the total number of ABR VCs passing through the switch port.
- *FairShare* This quantity is defined differently for different explicit rate schemes. It is sometimes defined to be the same as the *EqualShare*.
- *RM.ER*, *RM.CCR* Terms with *RM.x* indicate the value of the “x” field of the RM cell. *RM.ER* and *RM.CCR* therefore refer to the ER and CCR fields of the RM cell under consideration.

3.3.1 Enhanced Proportional Rate Control Algorithm (EPRCA)

This scheme is considered to be an enhancement to the Proportional Rate Control Algorithm (PRCA) proposed in [40] and further enhanced by Siu and Tzeng [15]. Sources are only allowed to transmit at rates less than or equal to the *FairShare* allocation computed by network switches. RM cells are sent at a rate proportional to the current rate of the source. The ER field of the RM cell is set to the desired rate by the source, the CCR field of the RM cell is set to the current allowed cell-rate (ACR) at the source, and the CI (congestion indication) bit is set to 0. As the RM cell passes through switches along the path to the destination, each switch updates a value called the mean allowed cell-rate (MACR) for the VC as an exponential average using a weighting factor α which is typically 1/16. The *FairShare* is then set to a fraction *swdpf* (typically 7/8) of the MACR.

$$\begin{aligned} MACR_{new} &= (1 - \alpha)MACR + \alpha RM.CCR \\ FairShare &= swdpf \times MACR \end{aligned} \tag{3.1}$$

When the RM cell is turned around by the destination and passes through the switches, each switch reduces the ER field of the RM cell to *FairShare* if it is in a state of congestion, (indicated by the queue length at the switch port) and if the value of *RM.ER* is greater than the *FairShare*.

The destinations check the EFCI bit in all the received data cells and if the most recent one had this bit set to 1 then the CI field of the RM cell is set to 1 else it is left as is. The sources meanwhile keep decreasing their rates by a fraction of their current transmission rate. When the source receives an RM cell, if the CI bit is not set, then it is allowed to increase its rate by a constant factor but limited by the both the value in the ER field and the negotiated peak cell-rate. If the CI bit is set, then the source is allowed to transmit at a minimum of its current transmission rate

(ACR), the value of the ER field and the peak cell-rate. [22] provides references for recent developments to this algorithm with the congestion indication at the switch based on the rate of growth of port buffer queue as opposed to queue length.

3.3.2 Congestion Avoidance using Proportional Control (CAPC)

This algorithm described in [16] uses the *LoadFactor* at the switch to indicate whether to increase or decrease the maximum rate that the switch grants to any VC. This rate is called the *FairShare*. It is initialized to the *EqualShare* and is increased when there is an under-load ($LoadFactor < 1$) and decreased when there is an over-load ($LoadFactor > 1$) at the switch port. The *LoadFactor* is periodically measured at the switch port.

$$\begin{aligned}
 & \text{if } (LoadFactor < 1) \text{ then} \\
 & \quad FairShare_{new} = FairShare \times \min(ERU, 1 + (1 - LoadFactor) \times Rup) \\
 & \text{if } (LoadFactor > 1) \text{ then} \\
 & \quad FairShare_{new} = FairShare \times \max(ERF, 1 - (LoadFactor - 1) \times Rdn)
 \end{aligned} \tag{3.2}$$

where *Rup* is a slope parameter in [0.025,0.1]

ERU is the maximum increase factor allowed and is set to 1.5

Rdn is a slope parameter in [0.2,0.8]

ERF is the minimum decrease required and is set to 0.5

A queue threshold is also used by the switch to determine if it is congested or not. If the threshold is exceeded, the congestion indication (CI) bit of the RM cells from all VCs passing through the switch, is set. When the source receives an RM cell with the CI bit set, it does not increase its rate, thus allowing the switch queue to reach a controllable size.

3.3.3 Explicit Rate Indication for Congestion Avoidance (ERICA)

The ERICA switch bandwidth allocation algorithm allocates bandwidth according to the current cell-rate of the VC and the offered load at the switch. It computes a rate called the $VCShare$ which is a fraction of the total available bandwidth $ABRCapacity$. This fraction is the ratio of the current cell-rate (CCR) of the VC to the incoming cell-rate ($TIInRate$) at the switch port.

$$RM.ER_{new} = \min(RM.ER, \min(\max(FairShare, VCShare), ABRCapacity)) \quad (3.3)$$

where $FairShare = \frac{ABRCapacity}{\#vcs}$ i.e the *EqualShare*

$VCShare = \frac{RM.CCR \times ABRCapacity}{TotalInrate}$, where $\frac{TotalInrate}{ABRCapacity}$ is the *LoadFactor*, so the

$VCShare$ can be written as $\frac{RM.CCR}{LoadFactor}$

Several innovations to the basic algorithm have been outlined in [17]. These include measures for ensuring max-min fairness, avoiding transient overloads at the switch, measuring the CCR at the switch on a per VC basis to avoid inaccuracies in the CCR field of the RM cell ($RM.CCR$) and exponential averaging of the *LoadFactor* to prevent sharp short-term fluctuations of this value.

An enhancement to ERICA called ERICA+ is also proposed in [17]. ERICA+ takes into account the switch port buffer queue length. When the queue length is long, the rates are allocated conservatively till the buffer has been drained and the queue length is under control. The target rate is computed as a function of the queue length, target queue delay, the link rate, and the VBR traffic rate. A case study of the ERICA and ERICA+ algorithms for LAN and WAN environments can be found in [44].

3.3.4 Explicit Rate Allocation Algorithm (ERAA)

This algorithm proposed by Kalampoukas et al [18, 45] distributes the bandwidth unused by VCs that cannot use it, to those that can use this bandwidth. The VCs that require a bandwidth less than that which can be allocated to them are termed “non-bottle-necked” or “satisfied”, while those that require more than that which they can be allocated at the switch are termed “bottle-necked” or “unsatisfied” VCs. The maximum bandwidth that can be allocated to a bottle-necked VC is:

$$A_{max}(t) = \frac{ABRCapacity(t) - \sum_{i \in S_u(t)} A_i(t)}{N_b(t)} \quad (3.4)$$

where t refers to the time an RM cell is received,

$S_u(t)$ refers to the set of satisfied VCs,

$A_i(t)$ refers to the allocation given by the switch to the i th VC (in the set S_u),

$N_b(t)$ refers to the number of bottle-necked VCs at time t

The request of a VC is computed as $\min(RM.ER, RM.CCR)$ of the received RM cell. A detailed algorithm is provided in [18]. When the source receives an RM cell returning from the destination, it updates its allowed cell-rate as :

$$ACR_{new} = \min(ACR + NRM \times AIR, RM.ER)$$

where $NRM - 1$ is the number of data cells between the current RM cell to the next,

AIR is a constant that limits the amount of increase allowed

3.3.5 Fast Max-Min Rate Allocation Algorithm (FMMRA)

This is another algorithm proposed recently [19, 46] that performs an allocation of bandwidth in proportion to the current cell-rate of each VC to the current load at the switch port. The allocation that each switch port can provide to VCs passing through the port is called the advertised rate (γ_l) for the switch port. This rate is computed

by determining the bandwidth leftover from VCs using less than their allocated rate (bottle-necked VCs) and dividing by the number of VCs that could use more than their allocated rate (non bottle-necked VCs). [19] shows the steps that finally lead to the following computation of the advertised rate:

$$\gamma_{l \text{ new}} = \gamma_l + \frac{\gamma_l \Delta \beta - \Delta \lambda}{N_l - (\bar{N}_l + \Delta \beta)} \quad (3.5)$$

where N_l is the total number of connections passing through the port l

$\bar{N}_l$ is the number of bottle-necked connections, i.e connections that do not make use of what the switch port is able to offer it.

$\Delta \beta$ is the change in the bottleneck status of the VC (bottle-necked to unbottle-necked or vice-versa)

$\Delta \lambda$ is the change in the bottleneck usage (allocation) of the VC, which is the difference between the current allocation and the previous allocation.

A VC gets a maximum of the advertised rate and the computed rate proportional to its CCR. The final allocation is only limited by its actual request (ER).

Every time a forward RM cell (FRM) is received by a network switch, the explicit rate (ER) field of the cell is adjusted as follows:

$$\begin{aligned} & \text{if} \quad (\text{port buffer occupancy} < \text{queue threshold}) \\ & \text{then} \\ & \quad RM.ER_{new} = \min \left(\max \left(\frac{RM.CCR(1 - \beta_i)}{LoadFactor_l}, \gamma_l \right), RM.ER \right) \\ & \text{else} \\ & \quad RM.ER_{new} = \min(RM.ER, \gamma_l) \end{aligned} \quad (3.6)$$

where i appearing as a subscript refers to the VC for which the computation is being performed

l appearing as a subscript refers to the switch port through which the VC i passes

port buffer occupancy refers to the current buffer occupancy at the switch port

queue threshold refers to threshold on the buffer size to indicate the onset of congestion and is set to a conservative value of 50 cells in our simulations.

$RM.ER_{new}$ is the newly calculated value to overwrite the value in the explicit rate field of the forward RM cell

β_i is the bottle-neck status of connection i - 1 or 0 depending on whether the connection is bottle-necked or not

$LoadFactor_l$ is the last estimated value of the offered load of the port (l)

γ_l is the advertised rate at the port (l)

A detailed description of this algorithm can be found in [19], while a comparison of different explicit rate algorithms with the FMMRA algorithm is presented in [46]. A modification of this algorithm for video transmission along with our proposed algorithm, Minimum Cell-Rate-based Proportional Allocation (MCRPA), is presented in the following chapter.

3.3.6 Suitability of Explicit Rate Control to VBR video

Explicit rate algorithms have several advantages over services like constant bit-rate (CBR) and variable bit-rate (VBR) for video transmission. These advantages are listed below:

- Explicit rate control schemes do not require the sources to specify statistics of the video in advance as opposed to CBR and VBR services. The only parameter negotiated in advance is the minimum cell-rate which is guaranteed for the duration of the connection. This feature makes it attractive for real-time video for which source video statistics are not available in advance.
- Explicit rate schemes allow flexibility to the Source End Systems (SES) for making bandwidth requests to the network based on their current needs. This rate may be reduced by network switches along the path that are unable to

support the request. On receiving the supportable rate from the network, the SES need to simply ensure that their rates do not exceed this value.

- The closed loop characteristic of explicit rate control allows sources and switches to communicate continuously with each other during the transmission. This allows sources to vary their transmission rate depending on the network feedback. Thus sources are able to respond quickly to network congestion and bandwidth availability.

These characteristics make ABR-like explicit rate control highly suitable for real time video traffic which do not require absolute bandwidth guarantees. In the following chapter we propose algorithms for carrying video over ABR-like explicit rate service.

Chapter 4

Source and Switch Algorithms

In this chapter we describe a bandwidth request algorithm and a switch bandwidth allocation algorithm for carrying MPEG video over an ABR-like explicit rate control service over ATM. We also describe control algorithms with which we compare the performance of our proposed algorithms. An overview of the models used in simulations to analyze the performance of our algorithms is given in the last section of this chapter.

4.1 Source Bandwidth Request Algorithms

Explicit rate algorithms, as we have seen in the last chapter, require sources to communicate a desired rate to the network switches. The strict delay requirements and the burstiness of video traffic necessitates an effective bandwidth request policy at the source for determining the desired transmission rate. A few schemes for predicting the required bandwidth for compressed video were outlined in section 2.3. These include Auto-Regressive (AR) causal prediction, Dual Threshold Moving Window (DTMW) prediction, simple prediction using the mean and standard deviation of the frame size distribution and a rate bounding scheme using AR-1 (1st order auto-regression) for predicting frame sizes. The AR causal prediction depends on the proper setting of buffer thresholds which are not intuitive parameters for real-time video transmitted

at varying rates. The DTMW algorithm requires prior knowledge of the video source bit-rate in order to determine an optimal value for the arrival rate threshold which is also the granularity for requesting the increase/decrease of the rate. The simple prediction scheme requires the mean and standard deviation of the video source to be known accurately in advance, while the rate bounding scheme requires complex computations at the source. In addition to the above drawbacks, we believe that schemes using previous frame sizes in predicting future bandwidth needs could make lower requests than that required for best video quality. This could happen when the source lowers its encoder bit-rate through adaptive quantization in response to network congestion. Even when the network is no longer congested, the source could continue to make smaller requests based on previous lower bit-rate, lower quality frames.

Below, we describe two bandwidth request algorithms. A hypothetical perfect prediction algorithm is followed by the Dual Delay Threshold Algorithm (DDTA) which we propose as a simple, causal, bandwidth request algorithm. The hypothetical perfect prediction serves as a control against which the performance of the DDTA is evaluated.

4.1.1 Hypothetical Perfect Prediction

In a recent paper by Lakshman et al [24], a bandwidth request algorithm based on the forecasting rule in [34] has been proposed for explicit rate feedback for video transmission. More accurate video source models based on the type of video are used to predict upcoming frame sizes. The authors in [24] propose averaging the predicted frame sizes for the next N frames that occur in the round-trip-time (RTT) for the source-to-receiver-then-back-to-the-source path. This average divided by the duration of a frame (33ms for 30 frame per second video) is then used as the rate to be requested from the network. Based on this, we implement a hypothetical perfect prediction algorithm in which perfect knowledge of upcoming frame sizes is available

at any given time. This is possible by storing the entire video sequence on disk and performing look-aheads. We also include the current frame-size in this averaging to accommodate small propagation delays and cases where the current frame is just beginning to be transmitted. The hypothetical perfect prediction is used as an ideal request algorithm, against which to compare the algorithm we propose next.

4.1.2 Dual Delay Threshold Algorithm (DDTA)

We now propose a bandwidth request algorithm based on violations of delay thresholds at the source buffer which is being drained at the current transmission rate (ACR) allowed by the network. An upper and lower delay threshold are used to control the delay experienced by video frames waiting at the source buffer, to be transmitted. The intuition behind choosing the threshold values is to exploit the smoothing effect of introducing delay, on the required transmission rate for the video source. This has been observed in the work by Reibman et al [36]. It was shown that the introduction of a constant delay of N frame times per frame resulted in smaller peak to mean cell-rate ratios and a smaller peak cell-rate for the transmission rate of video over the VBR service class of ATM. It was found that the smoothing effect was improved considerably by increasing the delay from one frame-time to two frame-times. Smaller improvements were observed for larger values of N . The delay cannot be made too large since this will affect the real time QoS of the video transmission. Using delay threshold violations rather than buffer threshold violations is more intuitive for delay sensitive traffic like video. In the DDTA, we use violation of delay thresholds to indicate the need for a change in the required bandwidth. The amount of change is automatically suggested by the difference between the required allocation satisfying the delay threshold bounds and the current allocation. A detailed description of the algorithm is provided next.

Every time a forward RM cell has to be sent by the source, the required rate ($req(t)$) at the current time (t) is calculated depending on whether the delay-thresholds

are violated. The future buffer occupancy at the time of transmitting the next RM cell (at Δ from the current time t), $\hat{b}(t + \Delta)$, is estimated using the average rate of growth $\frac{\overline{\delta b(t)}}{\delta t}$ of the buffer.

Let us first define the following terms:

$b(t)$ = current buffer occupancy in cells

$acr(t)$ = current transmission rate at the source in cells per second

$NRM = 32$ by default for ABR where $NRM - 1$ is the number of data cells after which an RM cell is sent

$\Delta = NRM/acr(t)$ which is the time between the transmission of 2 consecutive RM cells at the current transmission rate

n = number of RM intervals over which the average growth of buffer is computed and is chosen to be 5. Take for eg. a transmission rate of 6000cps. Five RM intervals correspond to a time of around 26ms for an NRM value of 32 cells. A large value for n will delay the reaction to fast growing buffer occupancy while a small value will result in rapid fluctuations based on short-term variations in buffer occupancy.

We compute the buffer occupancy at $t + \Delta$ as follows:

$$\begin{aligned} \frac{\overline{\delta b(t)}}{\delta t} &= \frac{\sum_{k=0}^{n-1} \frac{b(t-k\Delta) - b(t-(k+1)\Delta)}{\Delta}}{n} \\ \hat{b}(t + \Delta) &= b(t) + \frac{\overline{\delta b(t)}}{\delta t} \Delta \end{aligned} \quad (4.1)$$

The violation of the delay bounds is determined by computing the time taken for the source buffer to be drained at the current transmission rate considering both the current buffer size ($b(t)$) and the estimated future buffer size ($\hat{b}(t + \Delta)$) at the time Δ seconds later when the computation is repeated for the next RM cell. When the time for draining the source buffer at the current rate for both the current buffer occupancy and the estimated future buffer occupancy is below the lower threshold, the source is willing to request a reduced rate given by the current buffer occupancy divided by the mean delay threshold. When the drain-time for either the current

or estimated future buffer occupancy, exceeds the upper delay threshold, the source would prefer to request a higher rate, again given by the current buffer occupancy divided by the mean delay threshold.

We first define the following terms:

$req(t)$ = request made to the network through setting the ER field of the RM cell

$l.d.t$ = lower delay threshold

$u.d.t$ = upper delay threshold

$m.d.t$ = mean delay threshold = $\frac{u.d.t + l.d.t}{2}$

$$\begin{aligned}
 & \text{if } (b(t)/acr(t) < l.d.t) \text{ and } (\hat{b}(t + \Delta)/acr(t) < l.d.t) \\
 & \text{or } (b(t)/acr(t) > u.d.t) \text{ or } (\hat{b}(t + \Delta)/acr(t) > u.d.t) \\
 & \quad \text{then } req(t) = b(t)/m.d.t \\
 & \quad \text{else } req(t) = acr(t)
 \end{aligned} \tag{4.2}$$

In the unlikely event of the buffer occupancy becoming zero, the request will then become zero. To prevent this from happening and also to prevent the source rate from dropping too low too fast, due to the short range burstiness of the MPEG video frame sequence (IBBPBBP..), a damping factor is now introduced. This results in gradual ramping down of the rate when the computed request decreases, while the increases are in direct response to the buffer. A damping factor of 0.95 was found through trial and error, to work well in our simulations, so as to provide good tradeoff between the delay experienced by the frames and the utilization of bandwidth.

$$RM.ER = \max(damping_factor * acr(t), req(t)) \tag{4.3}$$

An important advantage of the DDTA algorithm is that no prior knowledge of the source is required. The source is expected to perform a simple computation described above, based on its current buffer size and transmission rate. Thus no statistical model or prior information of the video being transmitted is required, making it especially suited for real-time video.

4.2 Switch Algorithm

The network switches require a bandwidth allocation algorithm in order to divide up the available network bandwidth between the VCs passing through them. Several explicit rate algorithms have been proposed for the ABR service class, the most relevant of which have been described in the previous chapter. These algorithms have been shown to be effective for the rate control of data traffic. In [19, 46], the Fast Max-Min Rate Allocation algorithm (FMMRA) has been compared with several existing explicit rate algorithms and shown to offer performance advantages like low computational complexity, quick response to congestion, fairness and robustness over other algorithms. We have chosen this algorithm for investigating explicit rate control of video traffic. A modification of this algorithm, which we call FMMRA-1 for clarity, is presented first. Next, we propose an algorithm which attempts to solve some of the drawbacks of FMMRA-1 for video traffic.

4.2.1 Fast Max-Min Rate Allocation (FMMRA) - 1

The FMMRA algorithm described in the previous chapter, uses the CCR field of the RM cell which indicates the current rate at which the source is transmitting to the network, to compute $RM.ER_{new}$. It does not make use of the requested bandwidth ($RM.ER$) present in the ER field sufficiently. The allocation is in proportion to the current cell-rate ($RM.CCR$) of the source as opposed to the explicit rate request ($RM.ER$) of the source. Since our bandwidth request algorithm (DDTA) makes requests in proportion to the delay at the source buffer, it is critical for the network to grant rates in proportion to the requested rates. A later version of the FMMRA algorithm [46] uses $RM.ER/LoadFactor$ in place of $RM.CCR/LoadFactor$. This however does not take into account the weight of a VC's request with respect to the other requests at the current time. Hence, we have modified FMMRA so that the allocation is in proportion to the ratio of the explicit rate to the sum of the most

recent explicit rates from all VCs and call it FMMRA-1 in our simulations. The formula in equation 3.6 for the rate allocation when the port buffer queue size is below the threshold, now becomes:

$$RM.ER_{new} = \min(\max(\frac{RM.ER(1 - \beta_i)ABRCapacity}{\sum_{j=1, j \neq i}^{N_i} er_j + RM.ER}, \gamma_l), RM.ER) \quad (4.4)$$

where $\sum_{j=1, j \neq i}^{N_i} er_j$ is the sum of the most recent requests made by all VCs excluding the current VC i

4.2.2 Minimum Cell-Rate-based Proportional Allocation (MCRPA)

As seen in equation 4.4, the FMMRA-1 algorithm allocates at least the advertised rate to all sources, which is only limited by the actual request made by the source. The advertised rate is the share of bandwidth available to sources that are unbottle-necked and can use this bandwidth. Thus, it is at least the equal-share ($\underline{\underline{\Delta \frac{link_capacity}{N_i}}}$) or greater than it. Sources get the larger of the advertised rate and the proportional rate, limited only by the source request. The sum of the allocated rates could thus often exceed the total available bandwidth. This over-allocation of the network bandwidth could lead to cell-loss which is detrimental to the quality of the compressed video.

An alternative is to make the allocation strictly proportional i.e eliminate the advertised rate. This however, does not allow for the reallocation of unused bandwidth to sources that could use more bandwidth. Also, it does not protect sources whose proportional allocation is small due to the high requests made by other sources. All explicit rate algorithms guarantee the minimum cell-rate (MCR) negotiated by each source. Switches admit new calls as long as the total of all the negotiated minimum cell-rates including that for the new call is less than or equal to the available network bandwidth. When the computed proportional allocation is smaller than the minimum cell-rate, the source is given its negotiated minimum cell-rate. Thus

low bit-rate sources can be protected from high bit-rate sources if they negotiate a sufficiently high minimum cell-rate when a strictly proportional allocation scheme is used. However, this creates a potential for over-allocation of network bandwidth since the net allocated bandwidth is a sum of proportional cell-rates and minimum cell-rates, which could exceed the available network bandwidth. Thus even if all sources obey the rates allocated to them by the network, there is a possibility of congestion. Factors like the variation in the available bandwidth due to higher priority traffic like CBR, and the time-lag for information transfer between sources and switches could further aggravate the congestion. These problems can be prevented by our proposed policy outlined next.

Each source is allocated the minimum cell-rate (MCR) plus a proportion (the ratio of the difference between the source request and its negotiated MCR to the difference between all requests and all negotiated minimum cell-rates) of the left-over bandwidth after the minimum cell-rates are subtracted from the available bandwidth. The switch maintains an array of the most recent allocations ($alloc_{vc}$) made to each VC and an array of the most recent requests made by each VC (er_{vc}). Since the negotiated minimum cell-rate of each VC (mcr_{vc}) is available in the RM cell, it is optional to store an array of these values in the switch port. Every time a new call is admitted, a new entry is made in the MCR array for the new VC. The explicit rate array and the allocated rate array are updated every time an RM cell is processed at the switch port. The admission policy is simply:

$$\sum_{j=1}^{N_l} mcr_j \leq ABRCapacity \quad (4.5)$$

Let us consider what happens when an RM cell from the i th VC is processed. The request $RM.ER$ contained in the RM cell is compared with the MCR (mcr_i). If the request is less than the MCR, then the request is granted right away, since every source is guaranteed at least its MCR. Otherwise, the allocation is computed as:

$$proportional\ alloc = \frac{RM.ER - mcr_i}{\sum_{j=1, j \neq i}^N er_j + RM.ER - \sum_{j=1}^N mcr_j} \times (ABRCapacity - \sum_{j=1}^N mcr_j) \quad (4.6)$$

$$RM.ER_{new} = \min(RM.ER, \max(mcr_i + \max(0, proportional\ alloc), ABRCapacity - \sum_{j=1, j \neq i}^N alloc_j)) \quad (4.7)$$

where $alloc_j$ is the most recent allocation given to the j th VC.

er_j is the request contained in the most recent RM cell for the j th VC.

N is the number of VCs passing through the switch port

After the above processing, the array element $alloc_i$, corresponding to VC i in the array of allocated rates is now updated to the value of $RM.ER_{new}$. The value of the array element er_i is updated to the value originally contained in the ER field of the RM cell, i.e $RM.ER$. In order to speed up the computation of $\sum_{j=1, j \neq i}^N er_j$, $\sum_{j=1}^N mcr_j$ and $\sum_{j=1, j \neq i}^N alloc_j$, three variables sum_er , sum_mcr and sum_alloc can be maintained in the switch port and updated as follows. Every time a new call is admitted, its negotiated MCR is added to sum_mcr which now replaces $\sum_{j=1}^N mcr_j$. When an RM cell is received for VC i , the previous request (er_i) is subtracted from sum_er after which er_i is over-written by the new request $RM.ER$. Now, sum_er can replace $\sum_{j=1, j \neq i}^N er_j$ in equation 4.6. Similarly, the previous allocation ($alloc_i$) is subtracted from sum_alloc before the computations in equations 4.6, 4.7 and this value of sum_alloc can replace $\sum_{j=1, j \neq i}^N alloc_j$. After the computations in equations 4.6, 4.7 the new value of the allocation ($RM.ER_{new}$) is added to sum_alloc , which now contains the sum of the most recent allocations made to all VCs passing through the switch port. The original request $RM.ER$ is added to sum_er so that it now contains the sum of the most recent requests.

Notice that the term $ABRCapacity - \sum_{j=1, j \neq i}^N alloc_j$ allows the network to allocate at least what is currently unallocated i.e leftover from the available bandwidth, allowing complete utilization of network resources at the current time. This ensures that a source that needs a high bit-rate can make use of what is leftover from other sources some of which might be transmitting at rates below their proportional allocation. In order to prevent variations in available bandwidth due to high priority traffic from affecting the guaranteed minimum cell-rates, we suggest that the sum of minimum cell-rates be limited to a fixed proportion of the total network bandwidth that is always available to the ABR video service class. The remaining variable bandwidth can then be used for the proportional allocation component for each VC.

4.3 Simulation Models

A number of simulations were performed to evaluate the performance of the algorithms that we have proposed in this chapter. A detailed account of these simulations is presented in the next chapter. In this section, we describe the video traces, the end system and switch models, and the calculations performed during the simulation.

4.3.1 Video Traces

In our simulation experiments we have used traces from three MPEG-1 encoded videos. These include a low bit-rate trace of scenes with slow movements from the movie Ghost, a medium bit-rate trace of a talking head interspersed with high action scenes from a BBC newscast, and a high bit-rate trace of high activity scenes from a North Carolina State University (NCSU) football game. Statistics for the three traces are presented in Tables 4.1, 4.2, 4.3. The composite mean cell-rate for the three clips computed over the entire sequence (IBBPBBPBBIBBP....) and considering AAL-3 packetization (44 byte cell payload) were 2545cps for Ghost, 5641cps for BBC news and 8990cps for football. Graphs for the marginal distribution and autocorrelation

for the three traces are presented in Figure 4.1. The three traces have been chosen as typical of real-time video sequences and for the variety they present. In the simulations described in the following chapter, the original traces are used for single source experiments, while Transform Expand Sample (TES) models of the original traces are used for experiments using more than one video trace of the same type. TES ([47], [48]) models capture both the auto-correlation and the marginal distribution of the empirical data. Details of TES modeling of variable bit-rate video can be found in [49] and [50].

<i>frame</i>	<i>mean</i>	<i>peak</i>	<i>min</i>	<i>stddev</i>	<i># frames</i>
I	8860.4	14240	6322	1890.3	167
P	5763.3	11533	3372	1192.5	500
B	2324	4760	1542	487	1330

Table 4.1: Statistics for Ghost video - frame sizes in bytes for I,P frames at a quantization level of 3, and B frames at a quantization level of 6

<i>frame</i>	<i>mean</i>	<i>peak</i>	<i>min</i>	<i>stddev</i>	<i># frames</i>
I	15352	23006	7815	3420.8	188
P	13443	30034	7231	3186.1	563
B	5454	12950	2304	1424.9	1500

Table 4.2: Statistics for BBC news video - frame sizes in bytes for I,P frames at a quantization level of 3, and B frames at a quantization level of 6

4.3.2 Source End System

Our video source is modeled using an encoder that generates video frames at the rate of 30 frames per second. Each frame is broken up into 15 slices each of which is enqueued in a buffer that is drained at the rate allowed by the network (allowed cell-rate). The proportion of bytes per slice for each frame is the average value computed over all the frames for each of the three video traces (Figure 4.3). The conversion of

<i>frame</i>	<i>mean</i>	<i>peak</i>	<i>min</i>	<i>stddev</i>	<i># frames</i>
I	23360	36266	11525	6096.4	188
P	21849	35864	12946	5059.2	563
B	8658	16603	1435	2895.3	1500

Table 4.3: Statistics for NCSU football video - frame sizes in bytes for I,P frames at a quantization level of 3, and B frames at a quantization level of 6

bytes of video data into ATM cells is based on the AAL-3 adaptation layer, in which each 53 byte ATM cell carries 44 bytes of payload and 9 bytes of overhead. We have used a simple quantization model based on the PVRG MPEG-1 codec [25], which we have modified to include the upper delay threshold ($u.d.t$). When the delay at the local buffer ($\frac{b(t)}{acr(t)}$) exceeds the upper delay threshold, the encoder generates a lower bit-rate frame by using a higher quantization index (q) which corresponds to a coarser image quality.

$$q = QOffset + QDFactor \times (\max(0, \frac{b(t)}{acr(t)} - u.d.t)) \quad (4.8)$$

where q = quantization index level for encoding the frame

$QDFactor$ = a constant multiplier

$QOffset = 3$, the minimum quantization index level (maximum quality)

Adaptive quantization works simultaneously with our DDTA source algorithm. The video encoder computes the quantization at the start of encoding each frame and the rate controller computes the request ($req(t)$) to be made to the network, every time an RM cell is to be sent. We use the ATM Forum ABR default for explicit rate control of 31 data cells for every RM cell (i.e $NRM = 32$). The rate controller also sets the current transmission rate ($acr(t)$) of the source to what the network allows it ($RM.ER$), every time a BRM (Backward RM) cell is received by the source. The quantization index for B frames is twice that for I and P frames as in the PVRG codec. A graphical representation of the system is shown in Figure 4.2.

Figure 4.1: Histograms and Auto-Correlation for NCSU football, BBC News & Ghost video traces

4.3.3 Destination End System

The destination end system receives data and RM cells sent by the source and re-routes RM cells back to the source with the DIR bit set to 1. The end-to-end frame delay is computed per frame at the destination using a message-in message-out (MIMO) calculation [51]. Assuming that the source indicates the time of creation of each cell in a time-stamp field within each cell, and that the source and destination clocks are synchronized, the end-to-end frame delay (see Figure 4.2) is computed at the destination as :

$$Frame\ Delay = Min (FILO - nominal\ frame\ duration, LILO) \quad (4.9)$$

Figure 4.2: System Diagram

where *FILLO* (First-In Last-Out) is the difference in time between the reception of the last cell of the frame at the destination and the creation time of the first cell of the frame

nominal frame duration is the duration of a frame, which is 33ms for 30 frames per second video

LILLO (Last-In Last-Out) is the difference in time between the reception of the last cell of the frame at the destination and its creation time

Figure 4.3: Mean fraction of cells per slice for NCSU football, BBC News & Ghost video traces

In our simulations, the frame delay includes only the propagation delay between the source and the destination and queuing delays in the source and network queues. Considering a maximum tolerance for end-to-end delay for real-time interactive video to be 350ms, and subtracting out the encoding delay, packetization delay and decoding delay for the video, the tolerated end-to-end delay is conservatively determined in [14] to be 180ms. We consider frame delays upto 190ms to be within acceptable QoS bounds for real-time video. The cell-loss is also measured at the destination. It is calculated as the ratio between the cells lost and the total cells transmitted by the source on a per frame basis. The average cell-loss-ratio is the value of the ratio averaged over all the received frames.

4.3.4 ATM Switch Model

Our ATM switches are modeled with FIFO (First-In First-Out) output port buffers of a maximum capacity of 256 cells. Incoming cells are routed to the appropriate output port buffers, where they are queued till they can be transmitted. The time for processing each cell is considered negligible compared to the time spent in the

output port buffer queues. The switch discard policy to protect against congestion starts discarding low priority cells i.e cells with the Cell Loss Priority (CLP) bit in the header set to 1. when the buffer occupancy is more than 50% of its maximum capacity. When the buffer is full, both high and low priority cells i.e cells with CLP bit set to either 0 or 1, are discarded. The video sources set the CLP bit of I and P frame cells and RM cells to 0 which indicates high priority (HP) and that of B frames to 1 which indicates low priority (LP). Thus B frame cells are discarded before I and P frame cells which carry more important information for MPEG compressed video.

The ABR rate control is also implemented at each switch, and is performed on the traffic passing through each output port. We consider only video traffic for all our simulations assuming that the entire link capacity is available to the ABR-like explicit rate service for video. We assume that in real networks, delay sensitive video sources should not have to compete with ABR data traffic, which should be treated as a lower priority traffic class.

Chapter 5

Simulations For Performance Analysis

In this chapter, we describe in detail, the three major sets of simulation experiments that were carried out in order to analyze the performance of our algorithms. All our simulations were carried out using OPNET-2.5 [52]. The first set of simulation experiments compare the Dual Delay Threshold Algorithm (DDTA) with the hypothetical Perfect Prediction outlined in the previous chapter for effectiveness as a source bandwidth request algorithm. The second set of simulation experiments compare the two switch bandwidth allocation policies, Fast Max-Min Rate Allocation-1 (FMMRA-1) and Minimum Cell-Rate-based Proportional Allocation (MCRPA), using the DDTA at the source. The third set of simulations analyze the sensitivity of the DDTA-MCRPA system for varying loads and propagation delays.

5.1 DDTA vs Perfect Prediction

Our proposed bandwidth allocation algorithm, Dual Delay Threshold Algorithm (DDTA) is compared with the hypothetical Perfect Prediction in order to evaluate its effectiveness. The setup, observations and conclusions for this set of simulations are presented below.

5.1.1 Simulation Setup

Figure 5.1: Ideal Source-Receiver System

To isolate the performance of the bandwidth request algorithms from the network allocation algorithms, we consider a simple scenario of a single source-receiver system. We assume that the network always grants the source its requested bandwidth. Thus the RM cells always convey the requested bandwidth allocation to the source after traveling over the forward and reverse paths back to the source. Figure 5.1 shows the configuration for the simulations. The link-rate is 45Mbps both to and away from the source and the Initial Cell Rate (ICR) for the video source is chosen to be 2000cps. The bit-rate of the source is not adaptively controlled but fixed at a constant quality (quantization index level of 3). Thus, the bandwidth request policy alone controls the delay experienced at the source buffer. The experiment is repeated using each of the three original video traces, namely NCSU football, BBC news and Ghost, each run for a period of 60 seconds for four different propagation delays ranging from .125ms to 40ms. The above experiments are carried out for each of two algorithms, the hypothetical perfect prediction and the DDTA, implemented at the source. For the DDTA, the experiment is carried out for two choices of delay threshold bounds, [10ms,30ms] and [30ms,60ms]. The bounds have been chosen to be close to half, one and two frame-times (1 frame time = 33ms).

The measurements from each experiment includes the mean actual cell-rate at the source, the mean allocated cell-rate at the source and the frame-delay for each frame received at the destination.

Figure 5.2: Comparing DDTA with Perfect Prediction for football

5.1.2 Observations

The graphs in Figures 5.2, 5.3, 5.4, show the cumulative distribution function (CDF) for the frame delays experienced by each of three video clips for propagation delays of .125ms, 10ms, 20ms and 40ms. Table 5.1 shows the ratio between the mean actual cell-rate and the mean allowed cell-rate, measured at the source. This ratio indicates the utilization of the requested bandwidth by the source.

- The DDTA algorithm gives good frame-delays for the three video clips that we have used, namely football, Ghost and BBC news. The smaller delay threshold bounds [10ms,30ms] show smaller frame-delays than the higher delay threshold bounds [30ms,60ms] as expected. 90% of the frame-delays are within $u.d.t + \rho$, (where $u.d.t$ is the upper delay threshold and ρ is the propagation delay). This implies that the upper delay threshold serves to provide a statistical bound

Figure 5.3: Comparing DDTA with Perfect Prediction for BBC news

on the frame delay measured at the receiver, when all the source requests are granted by the network.

- Comparing the CDF graphs for DDTA with perfect prediction, an interesting result can be found for the football video clip. The hypothetical perfect prediction produces much higher delays (of the order of hundreds of milliseconds) for all the four propagation delays chosen. The football video has higher frame sizes and standard deviation between frame sizes than that for BBC news and Ghost video traces. For such video, the cell accumulation at the source buffer will be higher and faster since the hyp. perfect prediction request policy is based only on upcoming frame sizes as opposed to the occupancy level at the local buffer. The frame-delays at the local buffer can thus gradually increase to very large values (of the order of seconds), especially when the network is unable to

Figure 5.4: Comparing DDTA with Perfect Prediction for Ghost

grant the source its request. Considering the BBC news and Ghost traces, the frame delay for lower propagation delays for perfect prediction is much higher than that for the DDTA. For higher propagation delays, the perfect prediction performs slightly better than the DDTA with the [30ms,60ms] delay thresholds. The difference in the performance of the perfect prediction for lower and higher propagation delays is because the number of frames used in the averaging depends on the round-trip time (RTT). More frames are used in the averaging for longer RTT, smoothing out the variation in frame-sizes between consecutive frames (the reader may refer to section 4.1.1 for details on the hypothetical perfect prediction).

- Table 5.1 shows the ratios between the mean actual cell-rate of the source and the mean allocated cell-rate for the three videos for different propagation delays

<i>video</i>	<i>delay threshold bounds</i>	<i>propagation delay</i>	<i>actual/allocated</i>
football	[10ms,30ms]	.125ms	0.98
		10ms	0.82
		20ms	0.82
		40ms	0.84
	[30ms,60ms]	.125ms	1.00
		10ms	0.99
		20ms	0.96
		40ms	0.96
BBC news	[10ms,30ms]	.125ms	0.95
		10ms	0.77
		20ms	0.78
		40ms	0.80
	[30ms,60ms]	.125ms	1.00
		10ms	0.98
		20ms	0.95
		40ms	0.96
Ghost	[10ms,30ms]	.125ms	0.83
		10ms	0.61
		20ms	0.57
		40ms	0.62
	[30ms,60ms]	.125ms	0.99
		10ms	0.92
		20ms	0.88
		40ms	0.90

Table 5.1: Mean Actual Cell Rate to Mean Allocated Cell Rate ratios for the three videos, using DDTA

and choice of two sets of delay threshold bounds for DDTA. We can see that for the [30ms,60ms] delay bounds, the utilization is above 0.95 for the football and BBC news videos, and close to 0.90 for the Ghost video for the four propagation delays considered. The utilization is lower for the [10ms,30ms] delay bounds. This can be explained by the smoothing effect of delay on the output cell-rate of the video source as mentioned in the previous chapter. For the [30ms,60ms] bounds, the source tolerates upto 60ms delay which is close to two frame times

(66ms) and tends to smooth out the difference between consecutive frame sizes which can be quite high due to the IBBPBBP pattern. The smoothing results in lower variation in the requested rates, thereby making the scheme less sensitive to the time-lag for the feedback cell to reach the source. Thus the requested rates match the actual rates to a high degree. Another effect observed is that for low propagation delays, the utilization is higher than for high propagation delays. When there is a lower need for bandwidth at the source, the new reduced allocation is received by the source one round trip time later. Thus the allocated rate tends to be greater than the actual cell-rate at the source due to the latency introduced by the round trip time. The utilization is lower for the lower bit-rate Ghost video clip than the higher bit-rate BBC and football video clips. As we have seen in chapter 3, RM cells are sent every $NRM - 1$ data cells. The time for $NRM-1$ data cells is longer for the low bit-rate sources than for the high bit-rate sources. This slows down the ramping down effect produced by the damping factor (equation 4.3) when there is a reduced need at the source. The above results imply that lower frame delays for the [10ms,30ms] bounds come at the expense of lower bandwidth utilization. The [30ms,60ms] bounds provide a good trade-off between delay and bandwidth utilization for the different propagation delays for the three video traces considered.

- Overall, the DDTA algorithm yields consistent results for all the three video clips and provides controlled frame-delay depending on the delay thresholds chosen. The DDTA compares favorably with perfect prediction and in some cases performs better than the hyp.perfect prediction since it is able to track accumulations in the source buffer effectively and does not depend on factors like number of frames used in the averaging like the hyp. perfect prediction does. The [30ms,60ms] delay threshold bounds might be a more appropriate choice for the DDTA since they produce a high utilization of the requested bandwidth

for the three video traces. Considering the performance and the simplicity of implementation of the DDTA at the source, we believe that a scheme like the DDTA, based on the actual delay at the source buffer, may prove to be an effective alternative to schemes based on prediction of upcoming video frames.

5.2 FMMRA-1 vs MCRPA

We have compared the bandwidth allocation algorithms implemented at the source in the previous section, and found the DDTA to give good performance. In this section, the focus is on the bandwidth allocation algorithm implemented at the switch, with the DDTA implemented at the source. We compare the performance of the MCRPA algorithm with the FMMRA-1 algorithm, both of which have been described in the previous chapter. Below, we describe the series of simulations performed in order to do this comparison.

5.2.1 Simulation Setup

Figure 5.5: Heterogeneous Sources, 3 Switch System

The setup chosen for the following series of simulations involves multiple connections competing with each other for network resources. Figure 5.5 shows a three switch configuration, with traffic entering the first and second switches merging at the second switch, which thus becomes our bottleneck switch. The link-rate of the inter-switch links is 45Mbps. In this series of simulations, unlike the previous set, TES models of the three original video clips are used, so that our study is extended to sources with identical first and second order statistics. Twelve video connections, four of each of the three types of video, namely, football, BBC news and Ghost, enter at the first switch and are considered to be our main traffic. The cross-traffic comprises an aggregate of eleven video sources, two each of Ghost, BBC and football traces, entering the first switch and two Ghost sources, two BBC sources and one football source entering the second switch forming around 48% of the total traffic offered at the bottleneck. The mean bandwidth requirement of the 23 video sources at the bottleneck switch port, estimated at best quality for each video trace (quantization index level of 3) is 1.2 times the output link capacity. Thus the switch relies on the congestion control mechanism of the bandwidth allocation policy to keep the cell losses low especially during times of peak rate transmission. The DDTA implemented at the source uses the [30ms,60ms] delay thresholds that have been shown in the previous section to provide both good frame-delays and network utilization. The propagation delay between the switches is 5ms, and that for each of the links connecting the end systems to the network is 1 μ s. The net end-to-end propagation delay across the network is 10ms.

In these simulations, the network will not be able to provide the source with its requested rate at all times since there is competition between 23 video sources for the network resources. Thus, adaptive quantization is used at the source to alter the bit-rate of the encoded video to what the network is able to provide. Simulations are performed with each of the two switch bandwidth allocation algorithms, FMMRA-1 and MCRPA, implemented at the switches. The remaining factors are kept constant

between the two simulation environments. Each source starts at a random time chosen from a uniform distribution in the interval $[0,100\text{ms}]$ from the start of the simulation. The Initial Cell Rate and the Minimum Cell Rate are both chosen to be 2000 cps for all the sources. The measurements obtained from a single simulation for both experiments are frame-delay, quantization index level and cell-loss-ratio for each frame transmitted at the source; and the load factor and buffer occupancy at the bottleneck switch port sampled every 10ms. Each of the two experiments is run for a duration of 45s and is repeated 12 times to compute the 95% confidence intervals for the average values of our measurements. This number of repetitions was computed based on the t-distribution [53] for an accuracy in the frame delay of 5ms.

5.2.2 Observations

<i>statistic</i>	<i>value</i>	<i>confidence</i>
football		
mean frame delay (s)	0.043	(+/-) 0.003
mean qil (unit)	3.1	(+/-) 0.06
mean clr	0.192	(+/-) 0.013
prob(clr > 0.001)	0.343	(+/-) 0.032
BBC news		
mean frame delay (s)	0.040	(+/-) 0.002
mean qil (unit)	3.0	(+/-) 0.001
mean clr	0.189	(+/-) 0.014
prob(clr > 0.001)	0.334	(+/-) 0.033
Ghost		
mean frame delay (s)	0.031	(+/-) 0.002
mean qil (unit)	3.0	(+/-) 0
mean clr	0.200	(+/-) 0.015
prob(clr > 0.001)	0.37	(+/-) 0.032
Bottleneck Switch		
mean buffer occ.(cells)	122	(+/-) 3
mean util.	1.2	(+/-) 0.010

Table 5.2: 95 % Confidence intervals for FMMRA-1, prop delay=10ms, 12 runs

Figure 5.6: Results using FMMRA-1, MCR = 2000cps for all

Figure 5.7: Results using MCRPA, $MCR = 2000\text{cps}$ for all

<i>statistic</i>	<i>value</i>	<i>confidence</i>
football		
mean frame delay(s)	0.080	(+/-) 0.001
mean qil (unit)	4.0	(+/-) 0.4
mean clr	0.0006	(+/-) 0.0008
prob(clr > 0.001)	0.001	(+/-) 0.001
BBC news		
mean frame delay (s)	0.075	(+/-) 0.001
mean qil (unit)	3.6	(+/-) 0.2
mean clr	0.0007	(+/-) 0.0009
prob(clr > 0.001)	0.001	(+/-) 0.001
Ghost		
mean frame delay (s)	0.0595	(+/-) 0.002
mean qil (unit)	3.0	(+/-) 0.02
mean clr	0.0007	(+/-) 0.001
prob(clr > 0.001)	0.001	(+/-) 0.002
Bottleneck Switch		
mean buffer occ.(cells)	3.1	(+/-) 1.1
mean util.	0.96	(+/-) 0.003

Table 5.3: 95 % Confidence intervals for MCRPA. prop delay=10ms. 12 runs

The results for the simulations proposed above are shown in Figures 5.6, 5.7, and Tables 5.2, 5.3, 5.4. Figures 5.6 and 5.7 show the quality of the video received at the destination for three reference connections (one each of Ghost, BBC and football) and the bottleneck switch port conditions for each of the switch algorithms, FMMRA-1 and MCRPA. Tables 5.2 and 5.3 provide 95% confidence intervals for the mean of the statistics for 12 repetitions with each of the switch algorithms. Table 5.4 shows the mean allocated rates for four video sources for each of the three video traces, which form our main traffic.

- Figure 5.6 shows a high offered load at the bottleneck switch port averaging at 1.2, and a high buffer occupancy . This leads to high cell-loss ratios for the three video sources averaging at 0.2, which implies that an average of 20% of the

<i>video</i>	<i>mean acr (cells/second)</i>	<i>confidence</i>
football		
ref1	7004	(+/-) 188
ref2	6757	(+/-) 280
ref3	6940	(+/-) 265
ref4	7065	(+/-) 221
BBC news		
ref1	4434	(+/-) 208
ref2	4286	(+/-) 162
ref3	4275	(+/-) 126
ref4	4374	(+/-) 248
Ghost		
ref1	2607	(+/-) 74
ref2	2646	(+/-) 70
ref3	2616	(+/-) 69
ref4	2617	(+/-) 65

Table 5.4: Fairness Results for MCRPA with 95% Confidence intervals for prop delay=10ms. 12 runs. 2000cps minimum cell-rate for all

cells for the received frames at the destination are lost, which is unacceptable for compressed video. This can be explained by observing that the FMMRA-1 over allocates since it gives proportional allocation to high bit-rate sources, while it gives upto the advertised rate for requests with small proportional allocations. Considering the results for the MCRPA algorithm, in Figure 5.7, we find the mean offered load at the bottleneck switch port is high (96%) and kept very close to 100%. The buffer occupancy is low and the mean cell-loss ratio is around 0.0007 for the three video sources. Comparing the cumulative distributions of the frame-delays for the two Figures (5.6, 5.7), we find that the lower delays for FMMRA-1 come at the expense of cell-loss. We find the cumulative distributions for MCRPA are well within our target of 180ms for acceptable quality.

The MCRPA results in close to zero cell-loss and frame-delays less than 150ms,

while keeping the average load at the switch at a high value of 0.96 (96%). Thus it is effectively able to optimize between delay and loss, keeping the average offered load at the switch close to the maximum capacity without exceeding it.

- An effective switch bandwidth allocation algorithm must also provide fairness between competing sources. In Table 5.4, we show the allocated cell-rates for our main traffic which comprises 4 sources for each of the three types of video. It can be seen that the allocated cell-rates for the same type of video are very close, with higher bit-rates being given to the football sources (which have high mean bit-rates), and smaller bit-rates being given to the BBC and Ghost sequences.

5.2.3 Effect of Minimum Cell Rate

In order to see the performance for a choice of different minimum cell-rates for the three different types of video for the MCRPA, we then repeat the simulation with MCRPA at the switch, with other factors remaining the same. The minimum cell-rates are chosen as 2000cps for Ghost video sources, 4000cps for BBC video sources and 6000cps for the football video sources. These values were based roughly on the composite mean cell-rate for the three clips, 2545cps for Ghost, 5641cps for BBC news and 8990cps for football, all computed considering AAL-3 packetization.

Figure 5.8 shows the results for the above simulation for a reference connection for each type of video and the bottleneck switch port load and buffer occupancy. The cumulative distributions for the frame delays for the football and BBC traces show slight improvement (which can be seen clearly from the quantization curves) to those obtained for the previous choice of minimum cell-rates (Figure 5.7) of 2000cps for all video sources. This is expected since the football and BBC news traces are now guaranteed a higher minimum bandwidth. The quality of the video frames for the Ghost video trace is however lower than the previous case. Again, this can be explained by the fact that larger minimum cell-rates are guaranteed for the football

Figure 5.8: Results using MCRPA, MCR = 2000cps for Ghost, 4000cps for BBC news and 6000cps for football

and BBC traces, while the total available bandwidth remains the same. This leaves a smaller bandwidth for the Ghost sources. The utilization at the bottleneck switch buffer is high and the buffer occupancy is low leading to zero cell-loss-ratios. Notice again that the cumulative distributions of the frame delays for the three reference connections, for each type of video, show frame delays within our 180ms quality of service bound.

5.3 Sensitivity of the DDTA-MCRPA system to propagation delays and loads

In the previous sections, we have first analyzed the performance of two source bandwidth request algorithms and found the DDTA to be an effective bandwidth request algorithm. Next, we analyzed the performance of two switch bandwidth allocation algorithms for multiple competing video sources, with DDTA at the sources. We concluded that the MCRPA had definite performance advantages over the FMMRA-1. In this section, we further analyze the sensitivity of the MCRPA algorithm to increasing load at the switch and varying propagation delays.

5.3.1 Simulation Setup

Figure 5.9: Statistically Identical, Independent Sources and Two Switch System

The setup for the simulations is illustrated in Figure 5.9. A two switch configuration is used with the inter-switch propagation delay varying over a range of values to

measure the sensitivity of the algorithms to propagation delay. N statistically identical and independent football sources, modeled using TES provide the input traffic, which is varied by changing the number of sources, N , over a range of values. The minimum cell-rate chosen for all the football sources is 6000cps. The start times for all the sources are uniformly distributed between the 1st 100ms. The delay threshold bounds of [30ms,60ms] is used for the DDTA at the video sources, and MCRPA is implemented at the switches. Each experiment was run for a duration of 45s. Repeating simulations with each of the different combinations of N and the propagation delay to obtain confidence intervals was infeasible due to considerable time taken for each simulation. Thus, we have obtained 95% confidence intervals for a single case of 17 connections at 10ms propagation delay repeated 12 times, as representative of the entire set of simulations.

<i>statistic</i>	<i>value</i>	<i>confidence</i>
football		
mean frame delay (s)	0.088	(+/-)0.0015
mean qil (unit)	3.8	(+/-)0.1
mean clr	0.0	(+/-)0.0
Bottleneck switch		
mean util.	0.978	(+/-)0.0002

Table 5.5: 95% Confidence intervals for MCRPA with N identical football sources. prop delay = 10ms and $N = 17$, 12 runs

5.3.2 Observations for Varying Loads

With the inter-switch propagation delay being fixed at 10ms, the number of football sources N , is varied over 11, 13, 15 and 17 connections. As the number of connections goes from 11 to 17, the equal-share for each connection goes down from 9665.55 to 6254.18 cps, while the composite (over I,B,P frames) mean football source bit-rate at the highest quality (quantization of 3) is 8990 cps. Figure 5.10 shows the cumulative

Figure 5.10: Varying number of football sources

distributions for the frame delay and quantization index for each transmitted frame for a reference football connection. The 95% confidence intervals for the mean frame delay, quantization level, cell-loss ratio and the bottleneck switch load can be found in Table 5.5

- As the number of connections is increased, the frame-delay CDF curves show that the probability of frames having higher delays increases. As the number of connections is increased from 11 to 17, the the frame delay CDF curves and the quantization CDF curves shift more to the right. When the delay at the source buffer is beyond the upper threshold (60ms) the adaptive quantization kicks in, resulting in a reduced quality, which is reflected in the cumulative distributions for the quantization index for the frames being shifted to the right.
- The load-factor at the bottleneck switch takes values of .88, .94, .95, and .98 as the load is increased over the range of 11, 13, 15 and 17 connections. The probability that the quantization index is higher than 15 is practically zero and of the frame delay exceeding 180ms is negligible upto a maximum of 17 connections. The cell-loss ratio though not graphed is practically zero for the range of connections used above. Thus we observe that the quality and the

frame delay experienced by the video is acceptable even as the offered load at the bottleneck switch increases upto 0.98, with negligible cell-loss.

5.3.3 Observations for Varying Propagation Delays

In order to observe the effects of varying propagation delay between the source and the destination, the inter-switch delay is varied over the range of .125ms, 5ms, 10ms, 20ms and 40ms (LAN to large WAN propagation delays). N, the number of football sources, is fixed at 17. This results in an equal-share of 6254 cps, which is below the average football bit-rate at the highest quality (QIL=3) of 8989 cps. Figure 5.11 shows the effect of varying the propagation delays on the cumulative distributions of the frame delays and quantization index for the reference football connection. The 95% confidence intervals can be found in Table 5.5.

Figure 5.11: Varying propagation delay

- As the propagation delay is increased, it can be seen that the CDF graph for frame-delay shifts further to the right, showing that the frame delay increases. The frame delays are within our target value of 180ms for propagation delays upto 10ms, while the probability of the delay exceeding 180ms is very small (about 0.025) for higher propagation delays of 20ms and 40ms, with the maximum delay being around 190ms.

- It is interesting to note that the curves for CDF of quantization index for a propagation delay of 40ms shows lower frame delays compared to that for a propagation delay of 20ms. For the 40ms propagation delay, the round-trip-time (RTT) is around 3 frame-times (since 80ms is close to 99ms i.e 3 frame-times). MPEG high bit-rate frames are repeated every 3 frame-times for the frame sequence IBBPBBPBBP, while feedback cells are received back at the source at approximately 1 RTT after having been transmitted by the source. The bandwidth request made 80ms earlier now becomes useful for the next high bit-rate frame, thus slightly reducing the delay at the source buffer. The graph of the CDF of the quantization is reflective of the delay at the source buffer (as seen in equation 4.8), while the frame delay includes the propagation delay from the source to the destination. The quality of the video indicated by the quantization index level (QIL) is below a value of 15 for propagation delays upto 10ms, while that for a propagation delay of 20ms is below 15 more than 95% of the time. This indicates a good image quality.

Chapter 6

Conclusions and Future Results

6.1 Summary and Conclusions

The Dual-Delay Threshold Algorithm (DDTA) has been proposed as a bandwidth request algorithm for estimating the bandwidth requirement for video transmission using ABR-like explicit rate control. An upper and a lower delay threshold bound for the time taken to drain the source frame buffer at the current transmission rate, have been used to indicate the need for a change in the bandwidth allocation for the video source. The use of parameters like delay thresholds for real-time video is more intuitive than buffer bounds. The algorithm also avoids the drawback of requiring prior knowledge of statistics for the video, making it appropriate for real-time video. The DDTA has been shown to provide an effective alternative to a scheme based simply on prediction of upcoming frame sizes for the three video clips of different characteristics considered for our study.

An Explicit Rate Algorithm, the Minimum Cell-Rate based Proportional Allocation (MCRPA), has been proposed for transporting video over an ABR-like service over ATM. The MCRPA algorithm allocates a rate with a fixed component equal to the negotiated minimum cell-rate for the VC and a variable component which is a proportion of the bandwidth left-over after removing the total minimum cell-rates guaranteed to all the VCs passing through the switch port. The proportion depends

on the ratio of the difference between the current request made by the VC and its negotiated minimum cell-rate to the difference between the sum of all the current requests and sum of all the minimum cell-rates at the switch port. Thus the minimum cell-rate is used explicitly in the rate computation and prevents over-allocation of the network bandwidth. The performance of this algorithm has been compared with an adaptation of the Fast Max-Min Rate Allocation (FMMRA) algorithm, one of the recent proposals for Explicit Rate control. With the DDTA as the source request algorithm, the two switch algorithms, MCRPA and FMMRA-1, were compared in terms of quality of the transmitted video. The MCRPA has been shown to offer better performance in terms of optimizing between cell-loss and frame delay, and being able to control the average offered load at the switch port close to but rarely exceeding 100%.

The sensitivity of the DDTA-MCRPA system to varying propagation delays and network loads has been investigated. It has been shown to provide acceptable QoS in terms of the quality of the frame received and the frame-delay experienced at the receiver, for a wide range of propagation delays and network loads. Upto 17 statistically identical, independent football video sources were supported over a 45Mbps line with good video quality and frame delays, upto a propagation delay of 40ms.

In summary, a simple source bandwidth request algorithm and a switch explicit rate control algorithm for the transport of compressed MPEG video over ATM were developed and studied. The DDTA and MCRPA algorithms which we have proposed provide good QoS for video, even for heavy network loads and high propagation delays. Thus, it has been demonstrated that it is possible to use ABR-like explicit rate-control schemes for effectively transporting real-time video over ATM.

6.2 Suggestions for Future Investigations

We now present a few directions for further investigation in the use of explicit rate control schemes for transporting compressed video.

Scheduling Policy

It will be interesting to see the effects of different scheduling policies like simple exhaustive round-robin or weighted fair-queuing when traffic of higher priority like CBR and lower priority like ABR data are multiplexed along with the ABR video service over the same switch output port. Issues like buffer space allocation and a minimum fixed reserved bandwidth for the ABR video service need to be investigated.

Higher Layer Protocols

While it has been assumed in this study that the video frames are directly transported over the ATM layer, in reality higher layer protocols may be used to encapsulate the video frames before being passed to the ATM adaptation layers. Such protocols include the Internet Protocol (IP), Transport Control Protocol (TCP) or User Datagram Protocol (UDP) at the network level, and Real-time Transport Protocol (RTP) at the application level. The effect of the additional overhead and delay produced by these layers on the real-time QoS for video, needs to be investigated.

Multicast Video Scenarios

Multicasting involves the transmission from one or many sources to multiple receivers. Video transmission applications are likely to include many receivers as is often the case with video-conferencing. Thus, the extension of this study to multicast video scenarios may prove to be very useful for exploring the scalability and robustness of explicit rate control for multicast video transmission.

References

- [1] W. Stallings. *ISDN and Broadband ISDN*. Macmillan, 1992.
- [2] R. O. Onvural. *Asynchronous Transfer Mode Networks: Performance Issues*. Artech House, 1994.
- [3] R. Steinmetz and K. Nahrstedt. *Multimedia: Computing, Communications and Applications*. Prentice Hall, 1995.
- [4] D. Ginsburg. *ATM: solutions for enterprise internetworking*. Addison-Wesley, 1996.
- [5] ATM Forum Traffic Management Specification (Version 4.0), Apr. 1996. af-tm-0056.000.
- [6] M. Nomura, T. Fujii, and N. Ohta. Basic Characteristics of Variable Rate Video Coding in ATM Environment. *IEEE Journal on Selected Areas of Communication*, 7(5):752–760, June 1989.
- [7] P. Pancha and M. El Zarki. Bandwidth-Allocation Schemes for Variable-Bit-Rate MPEG Sources in ATM Networks. *IEEE Trans. on Circuits and Systems for Video Technology*, 3(3):190–198, Jun. 1993.
- [8] C. R. Kalmanek, S. Keshav, W. T. Marshall, S. P. Morgan, and R. C. Restrick. Xunet 2: Lessons from an Early Wide-Area ATM Testbed. *IEEE Trans. Networking*, 5(1):40–55, Feb. 1997.

- [9] M. Butto, E. Cavallero, and A. Tonietti. Effectiveness of the Leaky Bucket Policing Mechanism in ATM Networks. *IEEE Journal on Selected Areas of Communication*, 9(3):335–342, Apr. 1991.
- [10] N. Yamanaka, Y. Sato, and K. Sato. Performance Limitation of Leaky Bucket Algorithm for Usage Parameter Control and Bandwidth Allocation Methods. *IEICE Transactions on Commun.*, E75-B(2):82–86, Feb. 1992.
- [11] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic. *ACM Sigcomm 95, Comp. Comm. Review*, 25(4):219–230, Oct. 1995.
- [12] T. Suda and B. J. Vickers. An ATM Service Architecture for the Transport of Adaptively Encoded Live Video. In *ICCCN'96*, pages 179–186, Oct. 1996.
- [13] H. Kanakia, P. P. Mishra, and A. Reibman. An Adaptive Congestion Control Scheme for Real-Time Packet Video Transport. *ACM Sigcomm 93*, pages 20–31, Sep. 1993.
- [14] S. Karademir, I. Lambadaris, M. Devetsikiotis, and A. R. Kaye. Dynamic Rate Control of VBR MPEG Video Transmission over ATM Networks. In *IEEE Global Telecomm. Conf.*, pages 1509–1515, Westminster, London, Nov. 1996.
- [15] K. Y. Siu and H. Y. Tzeng. Intelligent Congestion Control for ABR Service in ATM Networks. *ACM Sigcomm 96, Comp. Comm. Review*, pages 81–106, Oct. 1996.
- [16] A. W. Barnhart. Explicit Rate Performance Evaluations, Oct. 1994. ATM Forum Documentt (94-0983).
- [17] R. Jain, S. Fahmy, S. Kalyanaraman, R. Goyal, and R. Vishwanathan. The ERICA Switch Algorithm: A Complete Description. Aug. 1996. ATM Forum Document (96-1172).

- [18] L. Kalampoukas, A. Varma, and K. K. Ramakrishnan. Dynamics of an Explicit Rate Allocation Algorithm for Available Bit Rate Service in ATM Networks. Technical Report TR UCSC-CRL-95-54. University of California, Feb. 1996.
- [19] A. Arulambalam, X. Chen, and N. Ansari. An Intelligent Explicit Rate Control Algorithm for ABR Service in ATM Networks. In *ICC. Proceedings*, Montreal, June 1997.
- [20] A. Kolarov and G. Ramamurthy. Comparison of ER and EFCI Flow Control Schemes for ABR Service in Wide Area Networks. In *Proceedings of the IFIP-IEEE Conference on Broadband Communications*, pages 123–134, Canada, Apr. 1996.
- [21] G. Ramamurthy and Q. Ren. Analysis of the Adaptive Rate Control for ABR Service in ATM Networks. In *Proceedings of the IEEE Infocom*, pages 1083–1088, May 1995.
- [22] R. Jain. Congestion Control and Traffic Management in ATM Networks: Recent Advances and A Survey. *Computer Networks and ISDN Systems*, 28(13):1723–1738, 1996.
- [23] K. W. Fendick. Evolution of Controls for Available Bit Rate Service. *IEEE. Comm. Magazine*, pages 35–39, Nov. 1996.
- [24] T. V. Lakshman, P. P. Mishra, and K. K. Ramakrishnan. Transporting compressed video over ATM networks with explicit rate feedback control. In *IEEE. Infocom Proceedings*, Apr 1997.
- [25] Portable Video Research Group, Stanford University. *PVRG-MPEG Codec 1.1*, June 1993.

- [26] C. M. Sharon, M. Devetsikiotis, I. Lambadaris, and A. R. Kaye. Rate Control of VBR H.261 Video on Frame Relay Networks. In *Proc. IEEE ICC '95*, Seattle, June 1995.
- [27] L. Chiariglione. MPEG and Multimedia Communications. *IEEE Trans. on Circuits and Systems for Video Tech.*, 7(1):5–18, Feb. 1997.
- [28] D. LeGall. MPEG: A Video Compression Standard for Multimedia Applications. *Communications of the ACM*. 34(4):47–58, Apr. 1991.
- [29] G. Karlsson. Asynchronous Transfer of Video. *IEEE. Communications Magazine*, pages 118–126. Aug. 1996.
- [30] D. L. Stone. *Managing the Effect of Delay Jitter on the Display of Live Continuous Media*. PhD thesis, University of North Carolina, Chapel Hill, 1995.
- [31] T. Kurita, S. Iai, and N. Kitawaki. Effects of Transmission Delay in Audiovisual Communication. *Electronics and Communications in Japan - Part I*, 77(3):63–73. 1994.
- [32] M. W. Garrett. A Service Architecture for ATM: From Applications to Scheduling. *IEEE. Network*, pages 6–14, May/June 1996.
- [33] C. Huang, I. Lambadaris, M. Devetsikiotis, and A. R. Kaye. DTMW: A New Congestion Control Scheme for Long Range Dependent Traffic . In *Proc. of 15th International Teletraffic Congress, ITC 15*, Washington D. C., June 1997.
- [34] T. Ott, T. V. Lakshman, and A. Tabatabai. A Scheme for Smoothing Delay-Sensitive Traffic Offered to ATM Networks. In *IEEE. Infocom Proceedings*, pages 776–785, Florence, Italy, 1992.

- [35] P. Skelly, M. Schwartz, and S. Dixit. A Histogram-Based Model for Video Traffic Behaviour in an ATM Multiplexer. *IEEE Transactions on Networking*, 1(4):446–459, Aug. 1993.
- [36] A. R. Reibman and A. W. Berger. Traffic Descriptors for VBR Video Teleconferencing Over ATM Networks. *IEEE/ACM Trans. on Networking*, 3(3):329–339, Jun. 1995.
- [37] R. Jain, S. Kalyanaraman, S. Fahmy, and R. Goyal. Source Behavior for ATM ABR Traffic Management: An Explanation. *IEEE Comm. Magazine*, pages 50–57, Nov. 1996.
- [38] F. Bonomi and K. W. Fendick. The Rate-Based Flow Control Framework for Available Bit Rate ATM Service. *IEEE Network*, pages 25–39, Mar./Apr. 1995.
- [39] R. Krishnan. Rate Based Control Schemes for ABR Traffic - Design Principles and Performance Comparison . In *IEEE Globecom*, pages 1231–1235, London, Nov. 1996.
- [40] L. Roberts, A. W. Barnhart, and M. Hluchyj et al. Closed-Loop Rate-Based Traffic Management. July 1994. ATM Forum Document(94-0438R1).
- [41] K. Y. Siu and H. Y. Tzeng. Adaptive Proportional Rate Control for ABR service in ATM Networks. Technical Report 07-01-94, Electrical and Computer Engineering, University of California, Irvine, July 1994.
- [42] P. P. Mishra, H. Kanakia, and S. K. Tripathi. On Hop-by-Hop Rate-Based Congestion Control. *IEEE/ACM Transactions on Networking*, 4(2):224–239, Apr. 1996.
- [43] H. T. Kung and R. Morris. Credit-Based Flow Control for ATM Networks. *IEEE Network*, pages 40–48, Mar./Apr. 1995.

- [44] R. Jain, S. Fahmy, S. Kalyanaraman, and R. Goyal. ABR Switch Algorithm Testing: A Case Study with ERICA. Oct. 1996. ATM Forum Document(96-1267).
- [45] L. Kalampoukas, A. Varma, and K. K. Ramakrishnan. An Efficient Rate Allocation Algorithm for ATM Networks providing Max-Min Fairness. In *Proceedings of 6th IFIP International Conference on High Performance Networking, HPN'95*, pages 143–154, Sep. 1995.
- [46] A. Arulambalam, X. Chen, and N. Ansari. Allocating Fair Rates for Available Bit Rate Service in ATM Networks. *IEEE. Communications Magazine*, pages 92–100, Nov. 1996.
- [47] J. R. Hill and B. Melamed. TESTool: A visual Interactive environment for modeling autocorrelated time series. *Performance Evaluation*, 24:3–22, 1995.
- [48] J. R. Hill, B. Melamed, and C. J. Walsh. User Guide and Reference Manual for TESTool (Version 1.0), Sep. 1993. NEC USA, Inc.
- [49] M. R. Ismail, I. Lambadaris, M. Devetsikiotis, and A. R. Kaye. Simulation and Modelling of Variable Bit Rate MPEG Video Transmission over ATM Networks. In *Proceedings of the IEEE Infocom*, Boston, Apr. 1995.
- [50] I. Lambadaris, M. Devetsikiotis, A. R. Kaye, M. R. Ismail, C. M. Sharon, Y. Fang, and C. Huang. Traffic Modelling and Design Methodologies for broadband networks. *Canadian Journal of Electrical and Computer Engineering*, 20(3):105–115, Apr. 1996.
- [51] R. Jain. ATM Forum Performance Testing Specification - Baseline Text, Dec. 1996. Ohio State Universty, BTD-TEST-TM-PERF.00.00.
- [52] MIL-3 Inc. *Opnet Modeler Manuals - Release 2.4*, 1993.

- [53] J. Banks and J.S. Carson. *Discrete Event System Simulation*. Prentice Hall, 1984.

IMAGE EVALUATION TEST TARGET (QA-3)

APPLIED IMAGE, Inc
1653 East Main Street
Rochester, NY 14609 USA
Phone: 716/482-0300
Fax: 716/288-5989

© 1993, Applied Image, Inc., All Rights Reserved