1

**PM-1 3 1/2"x4" PHOTOGRAPHIC MICROCOPY TARGET
NBS 1010a ANSI/ISO #2 EQUIVALENT**

PRECISIONSM RESOLUTION TARGETS

"Supporting MPEG-2 Video Communications over ATM Networks"

by

Eldon J. Mellaney, B.Sc.

Laurentian University

Sudbury, Canada

A thesis submitted to the Faculty of Graduate Studies and Research

in partial fulfilment of the requirements

for the degree of

Master of Computer Science

Ottawa-Carleton Institute for Computer Science

Faculty of Science

School Of Computer Science

Carleton University

Ottawa, Ontario, Canada

September 18, 1996

© E. J. Mellaney, 1996

National Library
of Canada

Acquisitions and
Bibliographic Services Branch

395 Wellington Street
Ottawa, Ontario
K1A 0N4

Bibliothèque nationale
du Canada

Direction des acquisitions et
des services bibliographiques

395 rue Wellington
Ottawa (Ontario)
K1A 0N4

Yours / Votre référence

Yours / Votre référence

The author has granted an irrevocable non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of his/her thesis by any means and in any form or format, making this thesis available to interested persons.

L'auteur a accordé une licence irrévocable et non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de sa thèse de quelque manière et sous quelque forme que ce soit pour mettre des exemplaires de cette thèse à la disposition des personnes intéressées.

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

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

ISBN 0-612-19482-5

Canada

Name ELDON JOHN MELLANEY

Dissertation Abstracts International is arranged by broad, general subject categories. Please select the one subject which most nearly describes the content of your dissertation. Enter the corresponding four-digit code in the spaces provided

Applied Sciences - Computer Science

SUBJECT TERM

0984

SUBJECT CODE

U·M·I

Subject Categories

THE HUMANITIES AND SOCIAL SCIENCES

COMMUNICATIONS AND THE ARTS

Architecture 0729
Art History 0377
Cinema 0900
Dance 0378
Fine Arts 0357
Information Science 0723
Journalism 0391
Library Science 0399
Mass Communications 0708
Music 0413
Speech Communication 0459
Theater 0465

EDUCATION

General 0515
Administration 0514
Adult and Continuing 0516
Agricultural 0517
Art 0273
Bilingual and Multicultural 0282
Business 0688
Community College 0275
Curriculum and Instruction 0272
Early Childhood 0518
Elementary 0524
Finance 0277
Guidance and Counseling 0519
Health 0680
Higher 0745
History of 0520
Home Economics 0278
Industrial 0521
Language and Literature 0279
Mathematics 0280
Music 0522
Philosophy of 0998
Physical 0523

Psychology 0525
Reading 0535
Religious 0527
Sciences 0714
Secondary 0533
Social Sciences 0534
Sociology of 0340
Special 0529
Teacher Training 0530
Technology 0710
Tests and Measurements 0288
Vocational 0747

LANGUAGE, LITERATURE AND LINGUISTICS

Language 0679
General 0289
Ancient 0290
Linguistics 0291
Modern 0401
Literature 0294
General 0295
Classical 0297
Comparative 0298
Medieval 0316
Modern 0591
African 0305
American 0352
Asian 0355
Canadian (English) 0593
Canadian (French) 0311
English 0312
Germanic 0315
Latin American 0313
Middle Eastern 0314
Romance 0314
Slavic and East European

PHILOSOPHY, RELIGION AND THEOLOGY

Theology 0422
Philosophy 0318
Religion 0321
General 0319
Biblical Studies 0320
Clergy 0322
History of 0469
Philosophy of
Theology

SOCIAL SCIENCES

American Studies 0323
Anthropology 0324
Archaeology 0326
Cultural 0327
Physical
Business Administration 0310
General 0272
Accounting 0770
Banking 0454
Management 0338
Marketing 0385
Canadian Studies
Economics 0501
General 0503
Agricultural 0505
Commerce Business 0508
Finance 0509
History 0510
Labor 0511
Theory 0358
Folklore 0366
Geography 0351
Gerontology
History 0578
General

Ancient 0579
Medieval 0581
Modern 0582
Black 0328
African 0331
Asia, Australia and Oceania 0332
Canadian 0334
European 0335
Latin American 0336
Middle Eastern 0333
United States 0337
History of Science 0585
Law 0398
Political Science 0615
General 0616
International Law and Relations 0617
Public Administration 0814
Recreation 0452
Social Work
Sociology 0626
General 0627
Criminology and Penology 0938
Demography 0631
Ethnic and Racial Studies
Individual and Family Studies 0628
Industrial and Labor Relations 0629
Public and Social Welfare 0630
Social Structure and Development 0700
Theory and Methods 0344
Transportation 0709
Urban and Regional Planning 0999
Women's Studies 0453

THE SCIENCES AND ENGINEERING

BIOLOGICAL SCIENCES

Agriculture 0473
General 0285
Agronomy 0475
Animal Culture and Nutrition 0476
Animal Pathology 0359
Food Science and Technology 0478
Forestry and Wildlife 0479
Plant Culture 0480
Plant Pathology 0817
Plant Physiology 0777
Range Management 0746
Wood Technology
Biology 0306
General 0287
Anatomy 0308
Biostatistics 0309
Botany 0379
Cell 0329
Ecology 0353
Entomology 0369
Genetics 0793
Limnology 0410
Microbiology 0307
Molecular 0317
Neuroscience 0416
Oceanography 0433
Physiology 0821
Radiation 0778
Veterinary Science 0472
Zoology
Biophysics 0786
General 0760
Medical

Geodesy 0370
Geology 0372
Geophysics 0373
Hydrology 0388
Mineralogy 0411
Paleobotany 0345
Paleoecology 0426
Paleontology 0418
Paleozoology 0985
Polynology 0427
Physical Geography 0368
Physical Oceanography 0415

HEALTH AND ENVIRONMENTAL SCIENCES

Environmental Sciences 0768
Health Sciences
General 0566
Audiology 0300
Chemotherapy 0992
Dentistry 0567
Education 0350
Hospital Management 0769
Human Development 0758
Immunology 0982
Medicine and Surgery 0564
Mental Health 0347
Nursing 0569
Nutrition 0570
Obstetrics and Gynecology 0380
Occupational Health and Therapy 0354
Ophthalmology 0381
Pathology 0571
Pharmacology 0419
Pharmacy 0572
Physical Therapy 0382
Public Health 0573
Radiology 0574
Recreation 0575

Speech Pathology 0460
Toxicology 0383
Home Economics 0386

PHYSICAL SCIENCES

Pure Sciences
Chemistry 0485
General 0749
Agricultural 0486
Analytical 0487
Biochemistry 0488
Inorganic 0738
Nuclear 0490
Organic 0491
Pharmaceutical 0494
Physical 0495
Polymer 0754
Radiation 0405
Mathematics 0605
Physics 0986
General 0606
Acoustics 0608
Astronomy and Astrophysics 0748
Atmospheric Science 0607
Atomic 0798
Electronics and Electricity 0759
Elementary Particles and High Energy 0609
Fluid and Plasma 0610
Molecular 0752
Nuclear 0756
Optics 0611
Radiation 0463
Solid State
Statistics
Applied Sciences
Applied Mechanics 0346
Computer Science 0984

Engineering 0537
General 0538
Aerospace 0539
Agricultural 0540
Automotive 0541
Biomedical 0542
Chemical 0543
Civil 0544
Electronics and Electrical 0348
Heat and Thermodynamics 0545
Hydraulic 0546
Industrial 0547
Marine 0794
Materials Science 0548
Mechanical 0743
Metallurgy 0551
Mining 0552
Nuclear 0549
Packaging 0765
Petroleum 0554
Sanitary and Municipal System Science 0790
Geotechnology 0428
Operations Research 0796
Plastics Technology 0795
Textile Technology 0994

PSYCHOLOGY

General 0621
Behavioral 0384
Clinical 0622
Developmental 0620
Experimental 0623
Industrial 0624
Personality 0625
Physiological 0989
Psychobiology 0349
Psychometrics 0632
Social 0451

EARTH SCIENCES

Biogeochemistry 0425
Geochemistry 0996

"Supporting MPEG -2 Video Communications over ATM Networks"

by

Eldon J. Mellaney, B.Sc.

Laurentian University

Sudbury, Canada

A thesis submitted to the Faculty of Graduate Studies and Research

in partial fulfilment of the requirements

for the degree of

Master of Computer Science

Thesis Supervisor

Thesis Supervisor

Director,

School of Computer Science

Carleton University

September 18, 1996

Abstract

Video services are expected to comprise the most significant portion of the traffic in future BISDNs. An integral part of providing high quality video services over ATM-based BISDNs is the MPEG-2 video coding algorithm. However, the prerequisite of commercially successful video services over ATM network using MPEG-2 video coding is careful planning of the ATM network resources. Part of this planning must include a thorough understanding of the characteristics of the MPEG-2 video traffic in the ATM network in nonhomogeneous scenarios. The thesis uses an experimental Video-on-Demand (VoD) system developed around the PC platform to determine how the characteristics of the MPEG-2 traffic in the OCRInet ATM MAN are affected by the network's intermediate elements and the end-to-end protocols employed in such a system under realistic network scenarios. Performance issues involved in the use of low end workstations for supporting a cost-effective, flexible and scaleable VoD system are also addressed.

Acknowledgements

I wish to thank most sincerely my supervisor Dr. Luis Orozco Barbosa for his contributions towards the completion of this thesis. Dr. John E. Neilson, my cosupervisor, must also be acknowledged for his insights and suggestions. Thanks are also in order to the individuals in the Television Broadcast Technologies, Research Video Systems at the Communications Research Centre, Industry Canada including Gilles Gagnon, André Vincent, Peter Haighton, and Stephane Trinh for the use of their research facilities and for providing me with guidance in the science of video processing. An additional thank you to Les Chan of the BADLab, Communications Research Centre for his help in setting up the experimental VoD system. To Setiadi Yazid, Peter Ashton, Allan Sand, and Susan Li in the COBRANet Laboratory at Nortel Technologies, Ottawa a sincere thank you for the use of your facilities and the time taken to aid me with the completion of the thesis.

To my loving parents, John and Carolyn, and my sister, Cynthia, a very sincere thank you for your support over the last two years — the phone calls were invaluable.

To Tamara, a most heartfelt thank you — you deserve so much. It is my sincerest hope that you get all that you deserve.

Dedication

Like an evening at the orchestra — dulcet,
elegant, august, and graceful — her love of life
and her love of books have inspired me in life and
in study.

This thesis is dedicated in loving memory of
Eileen Natalie Skakoon, B.Sc. M.Sc.

Table of Contents

ABSTRACT	iii
ACKNOWLEDGEMENTS.....	iv
DEDICATION	v
LIST OF FIGURES.....	ix
LIST OF ACRONYMS	xiii
1. INTRODUCTION	1
1.1 BACKGROUND	1
1.2 THESIS OBJECTIVE.....	4
1.3 THESIS ORGANISATION	6
2. MOTION PICTURE EXPERT GROUP VIDEO STANDARD — ISO 13818-2.....	8
2.1 OVERVIEW	8
2.2 TEMPORAL REDUNDANCY REDUCTION	10
2.3 SPATIAL REDUNDANCY REDUCTION	11
2.4 CONSTANT AND VARIABLE BIT RATE VIDEO SEQUENCES	12
2.5 MPEG-2 VIDEO BIT STREAMS STATISTICS	14
3. ASYNCHRONOUS TRANSFER MODE.....	22
3.1 OVERVIEW OF ATM.....	22
3.2 BISDN LAYERED PROTOCOL.....	24

3.2.1 Physical Layer (PL)	25
3.2.2 ATM Layer	25
3.2.3 ATM Adaptation Layer (AAL)	27
3.3 IP OVER ATM	28
3.4 ATM TRAFFIC MANAGEMENT	29
3.4.1 Traffic Contract	29
3.4.2 Traffic Descriptor	31
3.4.3 Usage Parameter Control (UPC)	32
4. EXPERIMENTAL SYSTEM	34
4.1 NETWORK TOPOLOGY	34
4.1.1 IP over ATM -- Encapsulation and Reassembly Processes	36
4.1.2 Impairments in the ATM Subnetwork	38
4.2 MPEG-2 VIDEO SERVER HARDWARE	39
4.3 MPEG-2 VIDEO CLIENT HARDWARE	40
4.3.1 MPEG-2 Hardware Decoder	40
4.4 MPEG-2 VIDEO CLIENT-SERVER SOFTWARE	41
4.4.1 MPEG-2 Video Client Server Protocol	42
4.4.2 MPEG-2 Video Server	46
4.4.3 MPEG-2 Video Client	47
4.5 ATM TESTSETS	48
5. EXPERIMENTAL RESULTS	51
5.1 FIRST CAMPAIGN OF MEASUREMENTS	51
5.1.1 Introduction	51
5.1.2 Measurements	51

5.1.3 Socket Statistics	52
5.1.4 Buffer Occupancy	56
5.1.5 Performance Issues in the MPEG-2 Video Client Server System	60
5.2 SECOND CAMPAIGN OF MEASUREMENTS	65
5.2.1 Introduction	65
5.2.2 Measurements	66
5.2.3 Usage Parameter Control in Network	67
5.2.4 Generating Traffic to Load the ATM Network	67
5.2.5 Unloaded Network Conditions	69
5.2.5.1 Cell Rate Snapshots	69
5.2.5.2 Cell Interarrival Times	74
5.2.5.2.1 CBR MPEG-2 Cell Arrival Snapshots	77
5.2.5.2.2 VBR MPEG-2 Cell Arrival Snapshots	79
5.2.6 Moderately Loaded Network Conditions	86
5.2.6.1 Cell Rate Snapshots	86
5.2.6.2 Cell Interarrival Times	90
5.2.7 Heavily Loaded Network Conditions	91
5.2.7.1 Cell Rate Snapshots	91
5.2.7.2 Cell Interarrival Times	96
6. CONCLUSIONS	99
REFERENCES	104
APPENDIX A MPEG-2 VIDEO STREAM STATISTICS	110

List of Figures

FIGURE 2.1	MPEG-2 GOP STRUCTURE	11
FIGURE 2.2	SNR COMPARISON BETWEEN THE CBR AND VBR MPEG-2 ENCODED AYERS SEQUENCE FOR EQUIVALENT RATES.....	15
FIGURE 2.3	AVERAGE BITS PER FRAME FOR THE VBR MPEG-2 ENCODED PIRATES, AYERS, AND MARTIN VIDEO STREAMS AS A FUNCTION OF THE Q PARAMETER.....	16
FIGURE 2.4	AUTOCORRELATION COEFFICIENTS OF BITS PER FRAME FOR VBR MPEG-2 ENCODED PIRATES VIDEO SEQUENCE ($Q=3$ AND 15)	20
FIGURE 2.5	AUTOCORRELATION COEFFICIENTS OF BITS PER FRAME FOR CBR MPEG-2 ENCODED PIRATES VIDEO SEQUENCE (6 Mb/s)	20
FIGURE 3.1	STRUCTURE OF THE ATM CELL AT THE UNI.....	23
FIGURE 3.2	BISDN ATM PROTOCOL REFERENCE MODEL	24
FIGURE 4.1	TOPOLOGY OF THE EXPERIMENTAL VoD SYSTEM.....	35
FIGURE 4.2	PROTOCOL ARCHITECTURE OF THE VoD EXPERIMENTAL SYSTEM.....	38
FIGURE 4.3	MEASUREMENT POINT AT THE ATM NETWORK AT NT PREMISES	50
FIGURE 5.1	BUFFER OCCUPANCY PLOT DURING THE TRANSFER OF THE AYERS 3 Mb/s CBR STREAM (5 TRANSFERS).....	59
FIGURE 5.2	BUFFER OCCUPANCY PLOT DURING THE TRANSFER OF THE AYERS VBR $Q=9$ (5 TRANSFERS).....	59

FIGURE 5.3	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS CBR 3 Mb/s (100 ms RESOLUTION); UNLOADED CONDITIONS	71
FIGURE 5.4	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS VBR $Q=12$ (100 ms RESOLUTION); UNLOADED CONDITIONS	71
FIGURE 5.5	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS CBR 3 Mb/s (500 ms RESOLUTION); UNLOADED CONDITIONS	72
FIGURE 5.6	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS VBR $Q=12$ (500 ms RESOLUTION); UNLOADED CONDITIONS	72
FIGURE 5.7	CELL INTERARRIVAL TIME DISTRIBUTION — AYERS CBR 3 Mb/s (0.5 μ s).....	76
FIGURE 5.8	CELL INTERARRIVAL TIME DISTRIBUTION — AYERS VBR $Q=12$ (0.5 μ s).....	76
FIGURE 5.9	RELATIVE CELL INTERARRIVAL TIMES OF 4000 CELLS OF AYERS 3 Mb/s CBR MPEG-2 TRAFFIC	79
FIGURE 5.10	RELATIVE CELL INTERARRIVAL TIMES OF 4000 CELLS OF AYERS VBR $Q=12$ MPEG-2 TRAFFIC	79
FIGURE 5.11	RELATIVE CELL INTERARRIVAL TIMES OF CELLS WITHIN ONE 24 KB WINDOW (AYERS VBR $Q=12$; CELLS 4952-5490)	82
FIGURE 5.12	BIT RATES OF ADMITTED TRAFFIC WITHIN ONE 24 KB WINDOW (480 CELL TIME SLOTS OR 1358.4 μ s RESOLUTION)	82
FIGURE 5.13	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS CBR 3 Mb/s (100 ms RESOLUTION)	87
FIGURE 5.14	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS VBR $Q=12$ (100 ms RESOLUTION)	87
FIGURE 5.15	CELL INTERARRIVAL DISTRIBUTION — AYERS CBR 3 Mb/s (0.5 μ s), 40% LOADED	88
FIGURE 5.16	CELL INTERARRIVAL DISTRIBUTION — AYERS VBR $Q=12$ (0.5 μ s); 40% LOADED	88

FIGURE 5.17	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS CBR 4 Mb/s (100 ms RESOLUTION).....	93
FIGURE 5.18	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — PIRATES Q=18 VBR (100 ms RESOLUTION).....	93
FIGURE 5.19	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — AYERS CBR 4 Mb/s (500 ms RESOLUTION).....	94
FIGURE 5.20	MEASURED CELL RATE VERSUS SOURCE RATE STATISTICS — PIRATES Q=18 VBR (500 ms RESOLUTION).....	94
FIGURE 5.21	CELL INTERARRIVAL DISTRIBUTION — AYERS CBR 4 Mb/s (0.5 μ s).....	97
FIGURE 5.22	CELL INTERARRIVAL DISTRIBUTION — PIRATES VBR Q=18 (0.5 μ s).....	97
FIGURE A.1	AYERS CBR 3 Mb/s.....	109
FIGURE A.2	AYERS CBR 4 Mb/s.....	109
FIGURE A.3	AYERS VBR Q=9.....	110
FIGURE A.4	AYERS VBR Q=12.....	110
FIGURE A.5	AYERS VBR Q=18.....	111
FIGURE A.6	MARTIN CBR 3 Mb/s.....	112
FIGURE A.7	MARTIN CBR 4 Mb/s.....	112
FIGURE A.8	MARTIN VBR Q=9.....	113
FIGURE A.9	MARTIN VBR Q=12.....	113
FIGURE A.10	MARTIN VBR Q=18.....	114
FIGURE A.11	PIRATES CBR 3 Mb/s.....	115
FIGURE A.12	PIRATES CBR 4 Mb/s.....	115
FIGURE A.13	PIRATES VBR Q=9.....	116
FIGURE A.14	PIRATES VBR Q=12.....	116

FIGURE A.15	PIRATES VBR $Q=18$	117
-------------	--------------------------	-----

List of Acronyms

AAL	ATM Adaptation Layer
ATM	Asynchronous Transfer Mode
BADLab	Broadband Application and Demonstration Laboratory
BISDN	Broadband Integrated Services Digital Network
BNR	Bell-Northern Research
CBR	Constant Bit Rate
CCITT	Comité Consultatif International Télégraphique et Téléphonique
CLP	Cell Loss Priority
CPCS	Common Part Convergence Sublayer
CRC	Communications Research Centre of Canada
CS	Convergence Sublayer
DCT	Discrete Cosine Transform
GCRA	Generic Cell Rate Algorithm
GDC	General DataComm
GoP	Group of Pictures
IP	Internet Protocol
ITU	International Telecommunications Union
ITU-T	ITU - Telecommunications standardisation sector

LAN	Local Area Network
MAN	Metropolitan Area Network
MBS	Maximum Burst Size
MPEG	Motion Picture Experts Group
NIU	Network Interface Unit
NT	Nortel Limited
PC	Personal Computers
PCR	Peak Cell Rate
PDU	Protocol Data Protocol
QoS	Quality of Service
SAR	Segmentation and Reassembly Sublayer
SCR	Sustainable Cell Rate
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
UDP	User Datagram Protocol
UTP	Unshielded Twisted Pair
VBR	Variable Bit Rate
VC	Virtual Channel
VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VP	Virtual Path

VPC	Virtual Path Connection
VPI	Virtual Path Identifier

1. INTRODUCTION

1.1 Background

Broadband Integrated Services Digital Networks (BISDN) are expected to offer a wide variety of services. In order to facilitate the introduction of new services while providing support for existing ones, the switching and multiplexing technique used in broadband networks must be both efficient and flexible to make the best possible use of limited network resources. As a result, statistical multiplexing is preferred over traditional time division multiplexing (TDM) systems which lack the required flexibility because they divide a communication link into a finite number of channels each with a fixed transmission rate. The statistical multiplexing solution chosen for BISDN is the Asynchronous Transfer Mode (ATM) [1].

Several ATM networks are being deployed worldwide [2, 3]. Although most of these deployments in public networks are technology trials, the commercial viability of ATM public networks is increasingly being recognised. It is well understood that for ATM to be successful it must support voice and video services.

Entertainment video, such as Video on Demand (VoD), is expected to be one of the video services that will be realised with the widespread deployment of public ATM networks. At present, information providers have begun test marketing their VoD services

to select and relatively small groups of subscribers [4, 5]. How successful the providers will be in their respective test markets will depend not only on the variety of programming offered to their subscribers, but also very strongly on the perceptual visual quality of the video at the customers' premises. Providing consistent high quality video over ATM networks requires careful planning of resources allocated in the network as well as at the network boundaries.

An integral part of the providing high quality video over ATM networks on a on-demand type basis is the Motion Pictures Expert Group (MPEG) video coding standard (ISO/IEC 13818-2) [6, 7]. The MPEG standard defines a generic video coding method that addresses a wide range of applications, bit rates, resolutions, qualities and services. The basic requirements of MPEG-2 video coding are a high compression ratio with good image quality and support for a number of optional features, such as random access, fast search, and reverse playback.

By exploiting the immense temporal and spatial redundancies present in natural video sequences, MPEG video encoding allows for a very significant compression of natural video sequences to be realised. This is made possible by the use of a robust hybrid coding algorithm based on temporal and spatial redundancy reduction. A block-based motion compensation algorithm and a two-dimensional Discrete Cosine Transform (DCT) compression algorithm are used for the reduction of temporal and spatial redundancies, respectively, present in the video sequences.

The immediate advantage of the high compression ratios that can be achieved for natural video sequences using the MPEG algorithm is that supporting high quality video services over packet switched networks, such as VoD over ATM, becomes much more feasible due to the significant reduction in the amount of information that must be transmitted over the network.

MPEG VBR encoding produces bit rates for encoded sequences that statistically vary over time. Since statistical multiplexing is used in ATM, a VBR video traffic source which is bursty in nature can be allotted a network connection with specified sustainable and peak cell rates. In this manner, many compressed video streams can be transported in the network even though the sum of their peak bit rates exceed the capacity of the transmission link. In this manner, a statistical gain can be achieved [8, 9]. From the user's point of view, the advantage of using ATM for the transport of VBR MPEG video is that a fixed quality video sequence can be delivered even during periods of high motion and scene changes.

Although there are many benefits to transporting VBR video streams in an ATM network, there are also issues in ATM which make the same task challenging. Cell delay variation and cell loss in the network are sources of quality of service (QoS) degradation. Intermediate network elements may be contributing sources to QoS degradation in that the manner in which they service the video connection may be such that they introduce significant amounts of cell delay variation in the cell stream. If cells sent into the network are subsequently delayed at intermediate network elements, the arrival process of the cells

The autocorrelation coefficients for the compressed bit stream Pirates using Q parameters of 3 and 15 are given in Figure 2.4. As expected, a better quality video sequence (e.g., $Q=3$) exhibits a higher correlation than a lower quality bit stream. The figure also shows that for a better quality encoded video stream, the effect of the encoding scheme (i.e., $N=12$, $m=2$) is smoothed out, indicating that lower compression ratios were achieved in the encoding process. Thus, the correlation is more representative of the contents of the video sequence. It is also interesting to observe the effects of maintaining a constant video quality with the use of a high-valued Q parameter. For instance, the correlation exhibits flat sections when $Q=15$. This characteristic was found to be present in all three video streams considered in the thesis. In the Pirates video sequence, three such periods can be easily identified (Figure 2.4). Figure 2.5 shows the autocorrelation coefficients for the 6 Mb/s CBR Pirates video bit stream. The effects of the video encoding (e.g., compression processes) are more evident in this figure than in the previous one for the VBR Pirates bit streams with $Q=3$ and 15. The autocorrelation decreases monotonically as a result of the compression processes; the oscillations present are due to use of the different picture types (i.e., I, P, and B) by the encoding process.

focused on the study of the performance of the PCs for the purposes of receiving a MPEG-2 CBR or VBR encoded bit stream via a 10 Mb/s Ethernet connection from a video server and decoding it for display. The use of low end workstations such as the ones used in the experimental VoD system are thought to provide a flexible and scaleable solution for delivering high quality video over existing communication networks [12, 13]. The second campaign of measurements focused on the characteristics of the traffic generated in the ATM network by the experimental MPEG-2 VoD system.

Many recent studies have investigated video traffic characteristics and their impacts on ATM network performance and engineering. These studies, however, have considered the video traffic sent to the ATM network to be identical to that generated by a video encoder such as an MPEG-2 (variable bit rate) encoder [14-19]. Additionally, many of these studies have considered only fairly homogeneous scenarios in which the only other traffic present in the ATM network, if any, are other MPEG-2 video sources.

Using the VoD experimental system, a study of the requirements of supporting MPEG-2 video communications over an ATM MAN was undertaken with the following two main objectives:

- a) to evaluate the feasibility of using low-end workstations to implement a VoD service by studying their performance characteristics and identifying bottlenecks that exist in such systems that may be mitigated or eliminated altogether;
- b) to determine the actual characteristics of the MPEG-2 video traffic generated by the VoD application in the ATM network in the presence of other traffic.

1.3 Thesis Organisation

Chapters two and three, respectively, provide the relevant background on the Motion Picture Expert Group (MPEG) video standard and the ATM-based BISDN required for the completion of the thesis. Specifically, chapter two provides an overview of the MPEG video standard (ISO 13818-2) with emphasis given only to the highest levels of the organisation of the MPEG video sequence. With a representative video sequence used in the thesis, the two streams produced from the sequence by constant bit rate (CBR) and variable bit rate encoding used in MPEG are qualitatively compared in terms of the signal-to-noise (SNR) ratios achieved in each stream on a frame-by-frame basis. Chapter three consists of a brief discussion of the BISDN ATM layered protocols and the structures and functions of the various layers within the protocol. Additionally, an overview of some of the proposals for carrying Internet Protocol (IP) over ATM networks is given with emphasis given to RFC 1483 which is the method used in the VoD experimental testbed to carry IP over ATM.

Chapter four details the experimental system designed and implemented to carry out the experiment work conducted for the thesis. The topology of the system and the design of the client-server VoD application are treated with due detail.

Chapter five presents the results of the two campaigns of measurements completed to meet the objectives of the thesis in two subsections of the chapter. The first section presents the results of the first campaign of measurements which focused on the performance of the low end workstations in the terms of their ability to be part of a

flexible and scalable solution of VoD systems. Results of the MPEG-2 video traffic measurements conducted in the ATM network under the various network loading conditions are given in the second section of the chapter.

Lastly, chapter six presents an overview of the major findings of the thesis.

2. MOTION PICTURE EXPERT GROUP VIDEO STANDARD — ISO 13818-2

2.1 Overview

The activities of the Motion Picture Expert Group (MPEG) cover more than just video compression; however, since the thesis concerns itself with only the video portion of the MPEG standard, it will be the only portion of the standard which will be overviewed. The MPEG-2 video coding standard developed by ISO/IEC [6, 7] defines a generic video coding method that addresses a wide range of applications, bit rates, resolutions, qualities and services. These different requirements have been integrated into a single syntax which facilitates bit stream interchange among different applications. The basic requirements of MPEG-2 video coding are a high compression ratio with good image quality and support for a number of optional features such as random access, fast search, and reverse playback.

A high compression ratio can be realised by exploiting the significant temporal and spatial redundancies present in natural video sequences. Consequently, the MPEG-2 video coding standard is based upon a hybrid coding structure of temporal and spatial processing. The degree of compression achieved in a video sequence has a significant impact on the transmission and storage of the MPEG video bit stream data.

To satisfy the requirement of a high compression ratio (best satisfied with pure intraframe coding) while maintaining random access to the video sequence (best satisfied with pure interframe coding), MPEG choose two intraframe coding techniques: predictive and interpolative. The MPEG video compression algorithm [6, 7] uses block-based motion compensation for the reduction of temporal redundancy and transform domain-based compression for the reduction of spatial redundancy. Motion compensated techniques are used in both causal (pure predictive coding) and non-causal predictors (interpolative coding). *Prediction error* signals, the differences between the actual and predicted values, undergo a two-dimensional discrete cosine transform (DCT) to reduce any spatial redundancy present. The motion information is transmitted together with the spatial information and is compressed using variable length codes to further reduce the amount of information that must be stored or sent. There are two major implications as a result:

1. full-motion video becomes a form of computer data
2. motion video (and its associated audio component) can be delivered over existing computer and telecommunication networks

The second implication of the MPEG video standard is the focus of the thesis.

During the standardisation process of MPEG-video, there were suggestions that the MPEG video standard could be used for other applications beyond video conferencing and videotelephony. These suggestions involved relaxing the constraints on delay and bit rates within a transmission network to permit increased visual quality for different

applications. One of these applications is the transmission of compressed television signals for use by broadcasters. The thesis is concerned with broadcast quality MPEG video bit streams which are suitable for transmission over existing and future computer and telecommunication networks.

2.2 Temporal Redundancy Reduction

For temporal processing MPEG-2 defines three main picture types: I, P and B, or intra, predictive, and interpolated or bi-directionally predictive pictures, respectively. I pictures are coded independently without reference to any other picture in the video sequence. They provide the random access points to the coded bit stream where decoding can begin. The P pictures are predictive coded pictures; their references are the previously coded I or P pictures. B pictures are bi-directionally predictive coded pictures with two picture references: one from the past and a second from the future. In general, because I pictures do not exploit the temporal correlation of natural video sequences, the compression ratios achieved for P and B pictures are much higher than those achieved for I pictures. Moreover, B pictures achieve the highest compression ratios as they use two references for prediction although they are never used as a reference for prediction. As a result, they can tolerate more distortion than I and P pictures. The organisation of the three pictures types is very flexible so as to support a wide range of applications.

In MPEG the video sequences are segmented into Groups of Pictures (GoP) containing at least one I picture, a number of P pictures and optionally a number of B pictures between the I and P pictures. The structure of a GoP is described by the variables

N and m . N specifies the total number of pictures in the GoP while m is the distance, measured in pictures, between two successive P pictures in a GoP structure. Figure 2.1 shows the temporal GoP picture structure of a MPEG-2 video sequence.

Figure 2.1 MPEG-2 GoP structure

2.3 Spatial Redundancy Reduction

Since a block-based motion compensation scheme is used in MPEG, a block-based spatial redundancy reduction was also chosen. The DCT was chosen for the spatial redundancy reduction technique because it is the best of the orthogonal transforms with a fast algorithm and is a very close approximation to the optimal one for a large class of images [7].

I pictures and the prediction error signals for P and B pictures have a very high degree of spatial redundancy. As a result, each is divided into blocks of 8×8 pixels. Each block in turn undergoes a two-dimensional DCT. The DCT coefficients are then quantized and variable length encoded, generating the compressed bit stream.

2.4 Constant and Variable Bit Rate Video Sequences

A MPEG-2 encoder is capable of producing a video stream encoded at a constant or variable bit rate, or CBR and VBR, respectively, which are the two basic types of coding used in MPEG in terms of the variation in the bit rates of the compressed video bit streams. Based on the communication channel requirements, the appropriate encoding method will be chosen.

In general, the bit rate of the compressed bit stream varies according to the content of the video sequence being compressed. If a VBR stream is to be transmitted over a fixed rate channel, a buffer is required to smooth out variations in the bit rate of the stream so as to produce a CBR stream. In order to prevent the buffer from overflowing and underflowing, information on the occupancy of this buffer must be fed back to the encoder. The encoder uses this information to adjust the coding parameters, such as the quantization parameter Q , so that the bit rate can be kept constant throughout the sequence. For example, if the buffer is in risk of overflowing, the value of the quantization parameter is increased to reduce the bit rate. If on the other hand, the buffer becomes almost empty, the quantization parameter is decreased, increasing the bit rate. The main disadvantage

with CBR video coding is that maintaining a constant bit rate is usually at the expense of a variable picture quality. In contrast, VBR provides a relatively stable picture quality at a cost of variable bit rates. Stable picture quality is achieved by fixing the value of the Q parameter throughout all the frames in a sequence. Since the content of the video sequence changes from frame to frame, the respective number of bits generated for each frame also varies on a frame to frame basis.

For comparative purposes, Figure 2.2 depicts the signal-to-noise ratio (SNR) values versus the frame number for both MPEG-2 CBR and VBR coding for one of the video sequences used in the thesis experiments. The figure also illustrates the quantization parameter Q with respect to the frame number for CBR coding scheme. Since under CBR, the quantization parameter Q is allowed to vary from macroblock to macroblock and frame to frame, the quantization parameter Q depicted in Figure 2.2 is actually the average Q parameter over each frame. On the other hand, for VBR, the quantization parameter Q is fixed at a constant value of 15, generating a relatively stable picture quality. The mean bit rate of the two streams over a period of 708 frames is approximately 3 Mb/s. It should be noted from Figure 2.2 that the quantization parameter Q used for CBR may take a value lower or higher than 15, resulting in a better or worse picture quality than that produced by the VBR encoding process. Moreover, untrained viewers may be unable to tell the difference between a CBR and VBR encoded video sequence of comparable bit rates.

It is important to note that the coding and display rates may differ for a video sequence. The former dictates the rate at which the frames are encoded and decoded while the latter defines the rate at which the frames are displayed (e.g., 30 frames/s). The discrepancy of the coding and display rates depends on the specifics of the material being encoded (e.g., 24 frames/s for motion picture films). During the decoding process, the decoder is able to adapt the decoding rate to the display rate.

2.5 MPEG-2 Video Bit Streams Statistics

One of the first tasks undertaken in the thesis was a study of the statistics of representative video sequences encoded using both the CBR and VBR encoding schemes in MPEG-2. The goal of this task was not to make an exhaustive study of the statistics of MPEG-2 video sequences, but, rather, to gain an understanding of the role played by the encoding process parameters on the statistical characteristics of the encoded bit streams. For this purpose, three video streams which exhibit different content and levels of motion were selected with the aid of members of the Television Broadcast Technologies, Research Video Systems group at CRC. The three selected video streams have been named "Ayers", "Pirates" and "Martin"; the video streams will be referred to by these names throughout the thesis. All three MPEG-2 encoded video bit streams consist of 708 frames and were generated at a coding and a display rate of 24 frames/s and 30 frames/s, respectively, with a picture size of 720×480 pixels. In all cases, the MPEG-2 encoding parameters of N and m were fixed at 12 and 2, respectively.

Figure 2.2 SNR comparison between the CBR and VBR MPEG-2 encoded Ayers sequence for equivalent rates

Figure 2.3 Average bits per frame for the VBR MPEG-2 encoded Pirates, Ayers, and Martin video streams as a function of the Q parameter

Figures A.1 and A.3, respectively, give the number of bits per frame for the Ayers video bit streams generated by CBR encoding at 3 Mb/s and VBR encoding with a quantization parameter Q of 9. In the CBR stream, the trace was taken from the output of the encoder just before being fed into the encoding buffer. In this scheme, the decoder requires the encoded bit stream at a constant rate equal to the nominal coding rate used when encoding the video sequence (e.g., 3 Mb/s). In the case of VBR encoding, the decoder will require the video stream at a variable rate. In studying the traces depicted in Figures A.1 and A.3 note that the use of the CBR encoding scheme eliminates periods of high bit rates. This is done, in particular, at scene changes, in camera panning scenes or in scenes which exhibit a high degree of motion. In the Ayers sequence a scene change

followed by a camera panning scene occurs between frames 50 and 250 resulting in a period of sustained high bit rates during this time period with VBR encoding (Figure A.3). Figure A.1 shows that the CBR encoded video bit stream does not exhibit this characteristic.

For the purpose of the study, the impact of periods in which the bit rate of a VBR encoded video stream approaches or temporarily exceeds the transmission rate of the network were of particular interest. Since, as previously stated, the decoder requires the timely delivery of the encoded bit stream, the use of a buffer (i.e., a buffer in application space — Section 4.4.3) should allow enough information to be stored to compensate for the difference between the coding and transmission rates. In the case of CBR, the peak bit rates are, for the most part, predictable since the encoding scheme leads to a regular pattern of peak rates every time an intra frame is encountered (i.e., every 12 frames since $N=12$). Scene changes or camera panning scenes tend for the most part to be smoothed out at the expense of perceptual video quality. However, CBR does not present major problems to the network since the timely transmission of the video sequence can be achieved by simply matching the network capacity to the nominal mean coding rate of the encoded stream plus a small amount for the overhead (approximately 14% in the VoD system) incurred in transporting the stream in the ATM network. A small buffer to account for delay variation is the minimum requirement at the receiving end.

In the VBR encoded sequences peak rates are less predictable since by keeping the Q parameter constant, scene changes, camera-panning scenes or high-motion scenes will

require a larger amount of information than their counterpart CBR encoded scenes. Furthermore, the same video stream may exhibit very different peak rates depending on the Q parameter used in the VBR encoding process (Figures A.8-A.10). Typically, the lower the value of Q , the less compression that is achieved and, consequently, the higher the peak rates in the encoded stream. Figure A.13 depicts the trace of the video stream Pirates with $Q=9$. Observe from the figure that the bit rate in an encoded stream may suddenly change. For example, there is a very substantial drop in the bit rate around frame 375. The trace of the video sequence Martin ($Q=9$) exhibits similar characteristics (Figure A.8).

Figure 2.3 shows the mean bit rate for all three sequences as a function of the Q parameter used in the VBR encoding process. From the figure it is clear that the number of bits per frame decreases as the value of the Q parameter is increased; however, the decrease in number of bits per frame is less significant as Q is increased beyond 18. From Figure 2.3 observe also that the number of bits per frame steadily increases as the Q parameter is decreased below 9. For this reason, all the experiments conducted in the thesis use VBR encoded video sequences with Q having values in the range of 9 to 18 with the exception of the Pirates sequence that the VoD system was unable to successfully transmit at Q parameters lower than 12. VBR video streams encoded with $Q=18$ still exhibit fairly good qualitative video. In order to gain a better understanding of the impact of the encoding scheme on the statistics of the encoded video stream, the autocorrelation coefficients for the number of bits per frame in all three sequences were considered.

The autocorrelation coefficients for the compressed bit stream Pirates using Q parameters of 3 and 15 are given in Figure 2.4. As expected, a better quality video sequence (e.g., $Q=3$) exhibits a higher correlation than a lower quality bit stream. The figure also shows that for a better quality encoded video stream, the effect of the encoding scheme (i.e., $N=12$, $m=2$) is smoothed out, indicating that lower compression ratios were achieved in the encoding process. Thus, the correlation is more representative of the contents of the video sequence. It is also interesting to observe the effects of maintaining a constant video quality with the use of a high-valued Q parameter. For instance, the correlation exhibits flat sections when $Q=15$. This characteristic was found to be present in all three video streams considered in the thesis. In the Pirates video sequence, three such periods can be easily identified (Figure 2.4). Figure 2.5 shows the autocorrelation coefficients for the 6 Mb/s CBR Pirates video bit stream. The effects of the video encoding (e.g., compression processes) are more evident in this figure than in the previous one for the VBR Pirates bit streams with $Q=3$ and 15. The autocorrelation decreases monotonically as a result of the compression processes; the oscillations present are due to use of the different picture types (i.e., I, P, and B) by the encoding process.

Figure 2.4 Autocorrelation coefficients of bits per frame for VBR MPEG-2 encoded Pirates video sequence ($Q=3$ and 15)

Figure 2.5 Autocorrelation coefficients of bits per frame for CBR MPEG-2 encoded Pirates video sequence (6 Mb/s)

The need to provide the MPEG-2 video bit stream to the hardware decoder at a variable rate has important implications on the communication services required for successfully implementing a client-server video service over an ATM/LAN internetworking system. The requirements for supporting MPEG-2 video communications over such a internetworking system is the focus of the thesis.

3. ASYNCHRONOUS TRANSFER MODE

3.1 Overview of ATM

Broadband Integrated Services Digital Networks (BISDN) are expected to offer a wide variety of services. Since no one can predict what services will be offered in the future, the underlying switching and multiplexing technique used in broadband networks must be both efficient and flexible to make the best possible use of limited network resources. Consequently, statistical multiplexing is preferred over traditional TDM systems which lack the required flexibility because they divide a communication link into a finite number of channels each with a fixed transmission rate. The statistical multiplexing solution chosen for BISDN is the Asynchronous Transfer Mode (ATM) [1].

In ATM short fixed-length packets called *cells* consisting of a small header containing network control information and an information field containing user data are used to divide the transmission bit stream. The sizes of the cell's header and payload were the result of a compromise between the audio and data groups involved in the development of ATM who sought smaller and larger information fields, respectively. However, explicit destination addresses cannot be expressed within a cell header because its small size. As a result, label switching is used to route cells from their source to their destination. Label switching requires the maintenance of translation tables at switching and multiplexing nodes in the network which map incoming cells with a given label to an

outgoing link and, as required, a new label. The label consists of the concatenation of the Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) fields present in the cell header

Figure 3.1 Structure of the ATM cell at the UNI

Since ATM statistically multiplexes cells onto a communication link, sending cells over the link only when there is a cell to send, it is possible to have a number of sources share a link whose capacity is less than that of the sum of the peak rates of the sources. Consequently, there is a statistic gain in bandwidth that can be realised with ATM which is not possible with time division or frequency division multiplexing methods. Additionally, any given source in the network may generate cells at a fixed rate or at statistically varying rate over time. Thus, not only is ATM capable of providing constant bit rate (CBR) services, but it is also able to exploit the inherit burstiness nature of various sources to provide the most efficient possible use of bandwidth resources.

3.2 BISDN Layered Protocol

The BISDN layered protocol reference model based on ATM consists of three planes and four layers (Figure 3.2). *User*, *management*, and *control* are the three planes in the BISDN layered protocol. Only the user plane will be dealt with in the thesis and will therefore be the only plane of the layered protocol described here. The user plane spans the physical layer (PL), the ATM layer, the ATM Adaptation Layers (AALs), which may be null, through to the higher layer [20].

Figure 3.2 BISDN ATM protocol reference model

The functions and formats of the three lower layers of the protocol are now described.

3.2.1 Physical Layer (PL)

The Physical Layer (PL) is responsible for the reliable transmission of bits and for inserting and extracting ATM cells from the ATM layer into and from the physical medium frames, respectively.

The PL consists of two sublayers. The lower sublayer, the Physical Medium (PM), interfaces with the actual physical medium and is responsible for passing extracted bit streams to the Transmission Convergence (TC) sublayer. The TC sublayer extracts and inserts ATM cells from and into the Plesiochronous or Synchronous Digital Hierarchies (PDH or SDH) Time Division Multiplexed (TDM) frame and passes them to and from the ATM layer, respectively [20].

3.2.2 ATM Layer

The ATM layer is responsible for the switching and multiplexing services. These services are based on the ATM cell.

The ATM cell is a fixed-length packet of 53 octets consisting of an 5 octet header carrying control information and a 48 octet payload. There are two standardised coding schemes for the ATM cell structure: one for the User Network Interface (UNI) and another for the Network-Node, or Network-to-Network Interface (NNI). The UNI is the interface between customer premises and the network switch while the NNI is the interface between switches and/or networks. Figure 3.1 illustrates the coding scheme for the ATM cell at the UNI.

At the UNI, the 5 octets in the cell header are divided into the following fields:

Generic Flow Control (GFC) Used to support simple implementations of multiplexing and provides a simple flow control mechanism.

Virtual Channel Identifier (VCI) Identifies a single Virtual Channel (VC) on a particular Virtual Path (VP). Used by switching nodes to establish a Virtual Channel Connection (VCC), the basic type of connection, through the use of lookup tables that map an incoming VCI value to an outgoing link.

Virtual Path Identifier (VPI) Used in a similar fashion as the VCI. The VPI is used to establish a Virtual Path Connection (VPC) for one or more VCs.

Payload Type (PT) Used to differentiate the payloads of cells on the same Virtual Channel Connection (VCC) that are carrying user and management information.

Cell Loss Priority (CLP) Specifies the loss priority of the cell.

Header Error Control (HEC) provides error checking for the ATM cell header which is used by the TC sublayer of the PL.

There are only two differences between the coding schemes used for the ATM cell at the UNI and the NNI. The first is that at the NNI the GFC field does not exist in the cell header. The second is that the 4 bits used by the GFC at the UNI are allotted to the VPI field at the NNI.

3.2.3 ATM Adaptation Layer (AAL)

The ATM Adaptation Layer (AAL) adapts the services provided by the ATM layer to those required by the higher layers. To provide this adaptation, the AAL is subdivided into the Convergence Sublayer (CS) and the Segmentation and Reassembly (SAR) layer. The functionality of the CS is divided into the Common Part (CP) and Service Specific (SS) components, or the CPCS and SSCS. While it is possible that the SSCS may be null (i.e., not implemented), the CP and the SAR sublayers are always mandatory.

The CPCS accepts AAL PDUs for encapsulation into CS-AAL PDUs. In turn, the SAR accepts these CS-AAL PDUs for segmentation into SAR PDUs which are at most 48 bytes long. At the ATM layer, each SAR PDU is mapped into the payload of one ATM cell. In all, there are five AAL types (e.g., AAL1-5). Different formats are used for the SSCS, CPCS, and SAR PDUs in each AAL type so that different services may be provided to the higher layers by the AAL.

At present there are four service classes defined by International Telecommunications Union (ITU) to which all AALs belong (Table 3.1). The service classes are defined in terms of bit rate, connection mode, and the timing relationship that exists between the sender and the receiver.

Table 3.1 ATM ITU ATM/BISDN Service Classes

Attribute	Service Class			
	Class A	Class B	Class C	Class D
Timing Relationship between source and destination	Required		Not required	
Bit Rate	Constant	Variable		
Connection Mode	Connection oriented			Connectionless
AAL(s)	AAL1	AAL2	AAL3/4 or AAL5	AAL3/4 or AAL5

3.3 IP over ATM

The Internet Engineering Task Force (IETF) IP over ATM working group has the task of developing a set of standards for the routing and forwarding of IP datagrams over ATM networks [21]. At present, there are many proposals for carrying IP over ATM. Many of the proposals, however, have the common thread of significantly reducing or eliminating altogether the overhead incurred in carrying the IP headers in the ATM network. "TCP and UDP over Lightweight IP" (TULIP) and "TCP and UDP over a Nonexistent IP Connection" (TUNIC) are two such proposals discussed within the working group. A simpler approach to sending IP over ATM is given in RFC 1483 [22].

In contrast to the approaches taken in TULIP and TUNIC, the "Null Encapsulation" or "Virtual Channel Multiplexing" described in RFC 1438 — "Multiprotocol Encapsulation over ATM" — does not attempt to reduce the overhead

incurred in sending the IP headers across the ATM network. Instead, RFC 1438 simply specifies that IP datagrams be placed in AAL PDUs for transmission and that the end point in an IP environment for a Virtual Channel (VC) is the IP layer. This is achieved through the use of a AAL instance which terminates a VC directly onto a IP layer endpoint. One limitation of this approach is that each pair of IP entities require their own VC.

The "Null Encapsulation" method is used by the CISCO 7000 routers for the transport of IP datagrams over ATM in the experimental VoD system used in the thesis.

3.4 ATM Traffic Management

3.4.1 Traffic Contract

A traffic contract is a set of rules that governs the manner in which a traffic source submits traffic to an ATM network. The contract is an agreement between the traffic source and the network. A separate traffic contract exists for each VPC or VCC in the network. The agreement has four components:

- Quality of Service (QoS) to be provided by the network for compliant cells (e.g., cell lost rate, cell delay variation)
- traffic parameters that specific the characteristics of the source traffic (i.e., cell flow of the source)
- a means of interpreting the traffic parameters to determine the conformance of a source to its contract (e.g., Generic Cell Rate Algorithm, or GCRA)

- the network's definition of a compliant connection

The QoS parameters include cell loss rate, cell error rate, severely errored cell block ratio, cell misinsertion rate, cell transfer delay, and cell delay variation. To simplify matters for end users, five predefined QoS classes are defined by ATM Forum UNI 3.0 [23]. The five QoS classes are given in Table 3.2.

Table 3.2 ATM Forum QoS Classes

QoS Class	QoS Parameters	Typical Application
0	Unspecified	"Best Effort", "At risk"
1	Specified	Circuit emulation
2	Specified	VBR video/audio
3	Specified	Connection-oriented data
4	Specified	Connectionless data

QoS classes are defined minimally by the average delay, CDV, the loss rate for high priority cells (i.e., cell loss priority (CLP) bit = 0), and the loss rate for low priority cells (CLP = 1) parameters. The "Specified" QoS classes guarantee performance in terms of the values set for these QoS parameters for a given ATM connection (VCC or VPC). In the "Unspecified" QoS class no values are specified for the QoS parameters. As a result, the typical use of this class is a "best effort" system. Ongoing activities in the ATM Forum that deal with this class of service commonly refer to it as the Unspecified Bit Rate.

Note that the QoS classes are related to the AAL types in Table 3.1 that are numbered correspondingly.

3.4.2 Traffic Descriptor

A traffic descriptor is the set of parameters whose values capture the behaviour of a traffic source in such a manner that enforcement of the source is possible. The key traffic contract parameters as specified by the ATM Forum include

- Peak Cell Rate (PCR) = T^{-1} — maximum cell rate where T is the minimum space between consecutive cells measured in seconds
- CDV tolerance = τ — measured in seconds and is usually specified by the network; used to specify the maximum number of cells, N , that can be sent back-to-back at the transmission medium line rate (e.g., 155 Mb/s for OC-3) where $N = \tau / T + 1$.
- optional Sustainable Cell Rate (SCR) — maximum average rate (in cells/s) at which a bursty traffic source can submit at the peak rate ($SCR \leq PCR$)
- Maximum Burst Size (MBS) — maximum number of cells that can be sent at the peak rate

The specification of the SCR and the MBS is similar to that of the PCR and were added to ITU Recommendation I.371 [24] to better model bursty traffic [20].

Having specified a descriptor for a given traffic source, a means must exist by which the network can determine the conformance of the source to its traffic contract. The conformance checking rule used at the UNI is the Generic Cell Rate Algorithm (GCRA)

which is a Virtual Scheduling Algorithm or a continuous-state *Leaky Bucket* algorithm ATM Forum UNI 3.0 [23].

3.4.3 Usage Parameter Control (UPC)

To ensure that the traffic contract for a given source or sources (e.g., a VCC or a VPC) is upheld, a mechanism called the Usage Parameter Control (UPC) exists at the UNI. The role of the UPC is to ensure fairness in the allocation of the network resources amongst the various traffic sources. A corresponding mechanism called in the Network Parameter Control (NPC) is used at the NNI to patrol traffic entering the network from other networks.

While the UNI 3.0 standard [23] does not stipulate any one implementation that must be used for the UPC (or NPC), the performance of the implementation of the UPC (or NPC) must be specified in relation to the aforementioned leaky bucket conformance checking rule. Example implementations of the UPC include the leaky bucket UPC, and the sliding and jumping window UPC [20]. Allowing different implementations to be used for the UPC may result in a cell stream from a traffic source being enforced in a slightly different manner from one UIC to another. For example, a cell stream with noncompliant cells as defined by the leaky bucket, may have different cells and maybe even fewer cells declared as noncompliant by different UPC implementations. An example of such a situation is given in [20] for the sliding and the jumping window UPC. A pathological example is also given in [20] in which the leaky bucket UPC, the sliding window UPC,

and the jumping window UPC declare 20%, 40%, and 30% of the cell stream as being noncompliant, respectively.

Additionally, there is latitude permitted in what the action the UPC will take if a traffic source violates its traffic contract. Nonconforming cells from a traffic source may be dropped at the UNI (i.e., not transmitted) or tagged. When a cell is tagged, the value of its CLP bit is set to 1 (i.e., low priority) but it is still forwarded onto the network. Alternatively, a UPC implementation may delay nonconforming cells in such a manner that the outgoing cell stream is compliant in that the nominal interval spacing (i.e., T) is never violated [39, 40].

4. EXPERIMENTAL SYSTEM

The experimental VoD system was built around the OCRInet ATM MAN [25] (Figure 4.1). Using the ATM network, a system was designed and implemented so that the MPEG-2 video data sent from the server would traverse the ATM network in such a manner that an ATM traffic analyser could be used in a controlled environment to observe the behaviour of the video traffic in the network under various network loading conditions. Additionally, the feasibility of using low end workstations for the purposes of providing a client with high quality CBR and VBR encoded MPEG-2 video bit streams was a matter of study.

4.1 Network Topology

The experimental VoD system has been implemented following a client-server paradigm with an ATM/local area network (LAN) internetworking system used as the underlying communication facility. The MPEG-2 video server¹ is physically located at CRC. It is connected to the OCRInet ATM MAN via a 10 Mb/s Ethernet on unshielded twisted pair (UTP) and then via an OC-3 fibre connection to a CISCO 7000 router. The router is located at the Broadband Application and Demonstration Lab (BADLab) at CRC. The

¹ To facilitate describing the experimental Video on Demand (VoD) system, the video server will be fixed at CRC although the roles of the server and the client are reversed for the second campaign of measurements.

interconnection between the UTP and the OC-3 uses an electrical to optical/optical to electrical adapter. Traffic destined for OCRInet from the BADLab goes through a Newbridge VIVID 12-port Workgroup switch and a Newbridge ATMnet Mainstreet 36150 16-port switch which are connected in series and are all located at the BADLab. The UNI and the NNI at CRC premises are located at the router and the Mainstreet switch, respectively. The MPEG-2 client was physically placed at Nortel (NT). It is connected directly to a CISCO 7000 router, the UNI at NT premises, via a 10 Mb/s Ethernet on UTP which is in turn connected to two FORE ForeRunner ASX-200 switches and a General DataComm (GDC) APEX DV-2 switch connected in series. The APEX DV-2, the NNI at NT, completes the client's path from NT to the OCRInet ATM MAN.

Figure 4.1 Topology of the experimental VoD system

4.1.1 IP over ATM — Encapsulation and Reassembly Processes

MPEG-2 video data is sent from the server to the client as follows. First, 8 KB of the video stream data is encapsulated in a UDP datagram at the transport layer. Each UDP datagram is subsequently segmented and encapsulated into six IP datagrams in the IP layer for transmission over the Ethernet link. Each IP datagram is small enough to be sent in one Ethernet physical frame (i.e., ≤ 1492 bytes). After reaching the CISCO 7000 router (the boundary between the IP and ATM networks) via a 10 Mb/s dedicated Ethernet, each IP datagram is extracted from its Ethernet physical frame and is mapped into one Common Part Convergence Sublayer (CPCS) ATM protocol data unit (PDU) of the ATM Adaptation Layer Type 5 (AAL5) using the "Null encapsulation" or "VC multiplexing" method specified in RFC 1483 [22]. The maximum transmission unit (MTU) for the CISCO router given that the input interface is Ethernet with a MTU of 1500 bytes and the output interface is ATM is 4470 bytes. Each AAL5 PDU, containing 1 IP datagram, is fragmented into ATM cells by the SAR sublayer of AAL5. The resulting ATM cells are sent across the ATM network by the router according to the Usage Parameter Control (UPC) specified for the MPEG-2 connection where they are received at the second CISCO 7000 router which acts as a gateway.

At the second router the cells are passed to the AAL5 for reassembly into IP datagrams as per RFC 1483. Each of these IP datagrams is encapsulated into one Ethernet physical frame and sent to the destination IP node across the Ethernet. At the client, the

Ethernet frames are received and the IP datagrams are extracted and sent to the IP layer. In the IP layer the UDP datagrams are reassembled from their constituent IP datagrams and sent up to the transport layer. Once reassembled, the checksum contained in the UDP datagram header is computed to determine whether the UDP datagram data has arrived intact at the correct destination [26]. If the checksum computed does not match the one contained in the header of the UDP datagram, the message is discarded. Having verified the checksum, the UDP datagram header is removed from the UDP datagram in the transport layer and the user level data is sent up to the application layer for the user level software to interpret. This encapsulation process occurs for UDP datagrams sent from both the server and the client. The protocol architecture is illustrated in Figure 4.2. A simple analysis of the encapsulation process shows that the total overhead is 14%.

Figure 4.2 Protocol architecture of the VoD experimental system

4.1.2 Impairments in the ATM Subnetwork

For the purposes of the thesis the ATM network was assumed to be error and loss free. Since the hardware decoder solution used in the client is not robust to data loss or errors in the video bit stream (Section 4.3.1) and since UDP will discard any UDP datagram with an incorrect checksum based its the header and user portions [26], it was assumed that any unsuccessful transfer of a video bit stream was the result of an error in the data bit stream or the loss of a message containing part of the bit stream. (This having verified that buffer underflow had not occurred in the application layer.) Therefore, a bit error rate and data loss rate of 0 were assumed for all experiments.

Issues of cell delay and cell delay variation were studied in order to characterise the MPEG-2 traffic generated by the VoD application. While important issues in the ATM network, the cell delay and cell delay variation experienced by cells carrying MPEG-2 video data did not have any impact on the perceived visual quality of the displayed video sequence at the client. Thus, delayed cells did not scale to *jitter*² observable to the viewer in the displayed video sequence. This was due to the large buffer used in the application layer at the client which smoothes the variations present in the cell arrival/IP datagram arrival processes. Moreover, this buffer precedes the hardware decoder's device driver's buffer from which the decoder receives the video stream data to decode. As a result, any variations in the IP/UDP datagram arrival processes will not appear in the 8 KB data movements to the device driver's buffer.

4.2 MPEG-2 Video Server Hardware

The video server has been implemented using an Intel 90 MHz Pentium PCI bus personal computer (PC). The server is equipped with 32 MB of RAM and a 10 Mb/s Ethernet adapter. For the first campaign of measurements, the video server was the machine known as "crc_mpeg2" which is located at CRC. crc_mpeg2 is equipped with a high-performance 10 MB/s 1 GB Seagate SCSI hard drive where the MPEG-2 video bit streams are stored. When the second campaign of measurements were completed, the roles, but not the

² The term *jitter* is often used for CDV in a communication network. In the thesis the term will refer to impairments in the smooth playback of the displayed video sequence (i.e., inconsistent, jumping, or freezing of the motion).

hardware, of the video server and client were reversed. For the second campaign of measurements, the video server was the machine known as "bnr_mpeg2" located at NT. It is equipped with a high-performance Western Digital hard drive. The video streams stored on both machines are identical (simply copied from one machine to the other using *ftp*). Note that identical MPEG-2 video hardware decoders [11] are used in both *crc_mpeg2* and *bnr_mpeg2* making the reversal of the roles of client and server simply a matter of running the server and client programs on the desired machines. The hardware decoder used in the two machines for the experiments is described in full detail in the next section.

4.3 MPEG-2 Video Client Hardware

The video client has been implemented using an Intel 90 MHz Pentium PCI bus PC. It is equipped with 32 MB of RAM and a 10 Mb/s Ethernet adapter.

4.3.1 MPEG-2 Hardware Decoder

A MPEG-2 hardware decoder developed by SGS Thomson Limited [11] is used in the client. The hardware decoder is capable of decoding encoded MPEG-2 video bit streams (H.262) encoded at rates up to 10 Mb/s and with a picture size of 720×480 and a display rate of 30 frames/s (i.e., a video rate up to $720 \times 480 \times 60$ Hz). For the thesis, a display rate of 30 frames/s with a picture size of 720×480 was used. The rates of the streams used in the experiments are given in Section 2.5.

The error concealment functions of the decoder are minimal. The decoder implements automatic macroblock error concealment. If the decoder detects a syntax error

in the video bit stream, it will replace the erred macroblocks with the corresponding macroblocks from the preceding picture. The error concealment will stay in effect until the next slice start code is detected in the bit stream. If the slice in which the error occurred is the last slice in the picture, concealment will continue until the end of the picture [11].

Initially, the hardware decoder was designed to run as a programmed input/output device under the DOS operating system; however, the decision was made to port the client to the LINUX operating system for performance gains (see section 5.1.5). This necessitated porting the DOS device driver code for the hardware decoder to LINUX where the same code, modified as required, provided the abstraction of a UNIX device.

Currently the decoder exists as an interrupt-driven device under LINUX. When the occupancy of the decoder's buffer falls below a fixed threshold, the decoder generates an interrupt and its device driver transfers 1 KB of MPEG-2 video data to the decoder's buffer. In this way, the decoder's buffer can never suffer from overflow; however, the decoder is susceptible to data starvation. If the MPEG-2 video data is not delivered to the decoder in a timely fashion, buffer underflow will occur in the decoder and it will fail to decode the bit stream.

4.4 MPEG-2 Video Client-Server Software

LINUX was used as the operating system for both the video client and the server. It was chosen after having verified that it offered much better performance than DOS/Windows. In fact, after conducting extensive tests, file transfer under LINUX was found to perform roughly 4 times faster than under DOS/Windows [27]. The application software of the

MPEG-2 VoD system is written in C and C++ using the socket interface available in the LINUX operating system.

4.4.1 MPEG-2 Video Client-Server Protocol

Under LINUX the User Datagram Protocol (UDP) was chosen at the transport level to explore the use of a simple protocol for the transfer, decoding and display of video streams over an ATM network. The objective of this task was therefore to evaluate the feasibility and performance of using UDP supplemented by simple control mechanisms for supporting both CBR and VBR MPEG-2-based video communications. The use of UDP, or UDP supplemented by another layer above it (e.g., real-time protocol, or RTP), seems to be a natural choice for supporting continuous media applications [28], [29]. UDP offers unreliable, connectionless service to transport datagrams from one host to another by using the underlying Internet Protocol (IP) layer which is also unreliable and connectionless [26]. As a result, UDP messages may be lost, duplicated, delivered out of order, or sent too quickly for the receiver to handle without loss. Therefore, measures were taken to ensure that data would not be lost unnecessarily due to messages arriving too quickly at the client to handle without loss.

At the application level, a simple protocol has been put in place. Synchronisation between the client and the video server is achieved by means of feedback messages. Feedback messages consisting of a single UDP datagram containing only a single byte in the user portion are used to implement a window flow control mechanism on top of UDP in the application layer. The messages sent back to the server convey no information about

the client's current state (e.g., the hardware decoder has suffered from underflow). Flow control was necessitated by the fact that if IP datagrams are sent continuously from the server to the client, IP datagram loss occurs. To avoid the unnecessary loss of IP datagrams, a simple flow control protocol was proposed and implemented.

In the application protocol implemented, the client starts the process by issuing a request for a particular video stream to the server. Upon receiving this request, the server sends the video stream by repeatedly reading the video bit stream file from its hard disk and invoking the socket interface of LINUX. With each invocation of the socket interface, a UDP datagram containing 8 KB of the encoded MPEG-2 bit stream and 12 bytes of overhead is sent to the awaiting client. The size of the user portion of the UDP datagram was decided upon based on the throughput results obtained in repeated experimental trials. Each UDP datagram is segmented into six IP datagrams in the IP layer. Each of the IP datagrams is at most 1492 bytes in length so there exists a one-to-one mapping of IP datagrams to Ethernet physical frames. At the client the IP layer must have all six IP datagrams which comprise a given UDP datagram before it may reconstruct the datagram and pass it up to the transport layer.

Since each UDP datagram corresponds to one *receive* system call on the client's side (and one *sendto* system call on the server's side), there exists a tradeoff between the size of the UDP datagram and the number of system calls required to transmit the bit stream over the network. Testing showed that if the UDP datagram size is halved, many video sequences encoded with VBR and CBR MPEG-2 that were previously received and

successfully decoded at the client could no longer be decoded without failure. This was due to the fact there were twice as many system calls being made and therefore twice as many kernel to user space and user to kernel space copies taking place during the transfer and decoding of the bit stream. If, on the other hand, the size of the UDP datagram is increased above 8 KB (e.g., 16 KB), the Ethernet adapter's ability handle the larger amount of data back-to-back without loss prevents the successful decoding of the encoded video stream being transferred.

The size of the window, W , used by the flow control mechanism was set to three after experiments revealed that this is the largest number of UDP datagrams that can be received back-to-back at the client without loss. Consider that there are two obvious factors that will limit the size of the window used in the flow control implemented. The first is the number of Ethernet physical frames that can be received successfully back-to-back. Call this quantity E . The second limitation comes from the size of the receive buffer allotted in kernel space to the client's socket. Consider that N UDP datagrams can be stored in this buffer. As a result, the size of the window, W , used by the flow control mechanism should be a value near E or N UDP datagrams.

The value of E is a function of the arrival process of the Ethernet frames at the client (i.e., the packet train) together with the ability of the adapter to process the arriving frames. Experiments conducted by sending UDP datagrams back-to-back from the server to the client revealed that well in excess of 100 UDP datagrams could be received before loss occurred. Thus, the adapter is able to handle at least three UDP datagrams back-to-

back. Consider now the size (N UDP datagrams) of the receive buffer for the socket. Since each UDP datagram is sent as six IP datagrams, the number of IP datagrams that can be stored in the receive socket buffer is $I = 6N$. Then, the two situations that must be considered with regard to the maximum size of the window are $E < I = 6N$ and $E > I = 6N$.

Consider the first situation when $E < I = 6N$. In this situation, there is sufficient resources to store N UDP datagrams in the receive socket buffer, but the number of Ethernet frames that can be received back-to-back successfully at the adapter, E , is less than $I = 6N$. In the VoD implementation, the *getsockopt* system call returns 60 KB as the size of the receive socket buffer. Assuming that the return value of the call is correct, then $W = 3 < N = \lfloor 61440 \text{ bytes} / 8204 \text{ bytes} \rfloor = 7$ (8204 bytes is the size of the user portion of the UDP datagram in the VoD system) indicates that the window size is limited by the Ethernet adapter's ability to receive physical frames that arrive closely together and not by the size of the receive socket buffer. But this contradicts the statement that the number of UDP datagrams that could be received back-to-back at the Ethernet adapter was well in excess of 100. This is true only when the data received is not written to the MPEG-2 hardware decoder. When the hardware decoder is used, there exists the potential for deadlock, but more significant is the load that the decoder places on the CPU for data to be transferred it from its device driver's buffer. Since the decoder is an interrupt driven device, as the rate of CBR encoding (or as the value of Q decreases for VBR streams), the number of cycles consumed in the CPU handling the decoder's interrupts increases

(linearly with the rate of encoding for the rate r_i for CBR streams). Thus, the size of the window W used in the system is neither E nor N , but rather something much smaller than either value. Issues related to the design of high performance network interface units are discussed in several research papers [30-35]. Others [28] note that often the kernel resources allocated for the receive socket buffer (for various network protocols including TCP/IP) is insufficient for properly supporting real time continuous media applications.

4.4.2 MPEG-2 Video Server

The video server is single threaded and is therefore only able to support a single client at one time. For the purposes of the thesis, support for more than one client was not required; however, the capabilities of the video server will improve when it is reworked by the individuals who will continue the work within the larger MPEG-2 over ATM project.

The server protocol is as follows. The video server waits for a request from a client and responds with the list of available MPEG-2 video bit streams on its hard disk. Once the user at the client has made his choice of the stream he wants transferred from the server, the client sends a message containing the user's selection back to the server. Having received the user's selection, the server begins sending the MPEG-2 video bit stream to the requesting client by repeatedly reading 8 KB of the encoded bit stream from its hard disk and invoking the socket interface in LINUX. There is a small overhead of three 4 byte integers sent with each 8 KB of MPEG video data. This overhead is used in the implementation of the ring buffer in application space at the client. The server will send three such messages (i.e., the 24 KB window; $W=3$) and then wait for an

acknowledgement from the client before sending the next 24 KB of the video stream. During its wait, the server will prefetch the next 8 KB of the video data that is to be sent. Because the server is single-threaded, it cannot prefetch more than one message before it must handle the acknowledgement from the client.

4.4.3 MPEG-2 Video Client

At the video client a ring buffer is defined in user space which is used to store portions of the encoded video stream before they are required by the hardware decoder. The ring buffer contains 500 elements each large enough to store one message sent from the server (i.e., 8 KB of the MPEG-2 video bit stream plus the 12 bytes of overhead). Two processes are used in the video client: the "decoding" process and the "replenishing" process. The decoding process reads 8 KB of the video bit stream data stored in a ring buffer and writes it to the MPEG-2 device (i.e., `write("/dev/mpeg", ...)`) while the replenishing process reads data from the network (i.e., `receive(sd, ...)`) and sends the acknowledgement messages back to the server requesting more data according the window flow control protocol detailed in section 4.4.

Initially, the ring buffer is prefilled to 2/3 of its capacity to avoid starvation once the decoding process begins. This prefill stage was necessitated by the fact that if the decoding and replenishing process begin simultaneously, certain VBR video streams caused buffer underflow to occur because of the high demands placed on the client's CPU for data to be transferred from the device driver's kernel buffer to the hardware decoder's buffer. Since this is an interrupt-driven process, with each interrupt causing 1 KB of data

to be copied from the device driver's buffer to the decoder's buffer, as the instantaneous bit rates of the VBR sources increase, the ability of the client to avoid starvation at the decoder's buffer diminishes. Performance issues in the VoD system are addressed more fully in section 5.1.5.

Synchronised access to the ring buffer by the decoding process and the replenishing process is achieved through the use of System V semaphores. If there is at least one empty ring buffer element, the replenishing process may run; otherwise, it will wait on the semaphore *EMPTY* until it has a value of at least 1. Similarly, the decoding process can only run if the value of the semaphore *FULL* has a value of at least 1; otherwise, it will wait on that semaphore.

The replenishing process is responsible for sending the acknowledgement messages to the server to implement the simple window flow control in the application layer. Thus, after having copied three (i.e., $W=3$) messages sent from the server from kernel space (i.e., from the receive socket buffer) into the application space ring buffer via the *receive* system call, the replenishing process sends an acknowledge to the server to request more data. Once all the video bit stream data has been written to the decoder, the decoding process (i.e., child process) ceases to exist (i.e., exits) and the remaining process (i.e., the replenishing process and parent process) terminates the session with the server.

4.5 ATM Testsets

The Adtech AX/4000 ATM testset [38] was used exclusively to obtain ATM cell level measurements of the MPEG-2 traffic in the ATM network. The AX/4000 testset allows

for the collection of a wide range of statistics on a virtual path identifier/virtual channel identifier (VPI/VCI) pair basis. ATM cell counts, cell throughput together with the corresponding PDU counts and rates of the CPCS of AAL5 are provided by the testset. The testset was used to obtain measurements of the MPEG-2 traffic under various network loading conditions to determine the characteristics of the MPEG-2 traffic in the ATM network. Cell rate snapshots with resolutions of 100 and 500 ms and cell interarrival time measurements with a resolution of 0.5 μ s were considered for this purpose. To load the output port of a switch so as to create the various network loading conditions, the AX/4000 testset traffic generator was used to generate traffic to be submitted to the ATM network. In all cases the testset was the source and sink of the generated traffic.

In the absence of any other traffic (i.e., interference from generated traffic) in the ATM network, no significant differences were observed between the ATM cell level measurements taken at the UNI and those taken at the NNI of the ATM MAN at NT premises (Figure 4.3). This was in spite of the fact that there are three intermediate shared ATM switches between the UNI of the CISCO 7000 router and the NNI of the ATM MAN. To this end, all cell level measurements in the network were taken at the output of the first FORE switch after the UNI at NT premises.

Figure 4.3 Measurement point at the ATM network at NT premises

5. EXPERIMENTAL RESULTS

5.1 First Campaign of Measurements

5.1.1 Introduction

A principle focus of the thesis, and the larger project of which it is part, was to gain an understanding of the requirements of supporting MPEG-2 video communications over an ATM network. Further, since the MPEG-2 server and client are general purpose workstations attached to dedicated Ethernet LANs and then to the ATM network and not connected directly to the ATM network, there was interest in studying the performance of the VoD system developed around the PC platform running the LINUX operating system using the UDP/IP internetworking protocols. To this end, the analysis of the performance and operation of the end systems in the VoD system constitutes a significant portion of the study of the requirements for supporting MPEG-2 video communications over an ATM network in a non-homogenous, perhaps more realistic, network topology. Consequently, a thorough study of the VoD system was completed in order to achieve to this objective.

5.1.2 Measurements

Measurements were taken at different points of the protocol architecture at the video client. The occupancy of the client's ring buffer was recorded over the duration of the

transfer and decoding process. Statistics were also collected for the socket interface. These statistics were written to separate disk files. The statistics pertaining to the socket interface provide a means of determining the socket (i.e., application layer) throughput required for the successful transfer of a given bit stream and permit an estimate to be made on the load placed on the client's CPU for the transfer of a given MPEG-2 video stream.

5.1.3 Socket Statistics

Statistics were collected for the socket interface at the video client and include the number of socket calls per second and the mean socket throughput of the client's software for the transfer of a given MPEG-2 encoded video stream. To obtain the mean socket throughput value for the transfer of a given MPEG-2 video stream, it was transferred back-to-back five times. This trial was repeated four times so that the bit stream was transferred a total of 20 times. Further, each group of five transfers of the stream of interest was separated by at least five transfers of a different stream (the maximum separation distance was 5 groups of 5 transfers, or 25 stream transfers.) To calculate the mean socket throughput, the last socket throughput value obtained in each group of five transfers of the stream was considered and the average of these five values was computed to obtain the mean socket throughput during the transfer of the stream. This was done to minimise the effect of prefilling the ring buffer to $2/3$ its capacity during the first transfer of the stream. This method was used to obtain the socket statistics for all the stream transfers (e.g. all rates of CBR and all values of Q for VBR streams). Note that the data considered at the socket interface does not include the overhead of the transfer (approximately 14% in the VoD

system), but rather only the application layer data which consists of 8204 bytes (8 KB of video data and 12 bytes of overhead) in the user portion of each UDP datagram.

The results show that for MPEG-2 CBR encoded bit streams, the mean socket throughput is approximately that of the rate of encoding (Table 5.1). The variance reported for each throughput value is the variance of the five throughput values used to compute the average throughput presented. Although the variances vary significantly from one stream to the next, the size of the variance in all cases is extremely small.

Table 5.1 Mean socket throughput for CBR encoded MPEG-2 bit streams

Sequence	3 Mb/s		4 Mb/s	
	Mean (Mb/s)	Variance	Mean (Mb/s)	Variance
Pirates	2.978067	5.24717E-08	3.969979	1.03247E-10
Ayers	2.978802	1.59350E-09	3.971063	4.23319E-09
Martin	2.978774	1.69001E-10	3.970351	1.14350E-09

The small differences between the encoding rates and the mean socket throughputs recorded can be attributed to the fact that although the rate of the encoding is reported as 3.0 or 4.0 Mb/s per second, the relatively short length of the sequences (approximately 30 seconds) causes the rate of encoding to be never exactly 3.0 or 4.0 Mb/s. These mean

socket throughput measurements show that a bandwidth equal to the rate of encoding, plus a small amount for the overhead of the transport of the MPEG-2 video data in the network, is sufficient to support the transfer and successful decoding of a CBR encoded bit stream over the ATM MAN.

For MPEG-2 VBR encoded video sequences, the mean socket throughput and the mean number of socket calls per second varies according to the video sequence content and the quantization parameter Q selected for the encoding process and (Table 5. 2).

Table 5.2 Mean socket throughput for VBR encoded MPEG-2 bit streams

Sequence	$Q=18$		$Q=9^{\dagger}$	
	Mean (Mb/s)	Variance	Mean (Mb/s)	Variance
Pirates	3.024796	2.27480E-11	4.400340	2.01266E-07
Ayers	2.329289	4.19994E-11	4.747456	3.09502E-10
Martin	1.656552	1.18749E-12	3.760072	2.52619E-09

For the transfers of the VBR encoded bit streams the small differences observed between the mean bit rates of the sequences (i.e., total number of bits in the sequence /

[†] Except for Pirate, $Q=12$.

length of the sequence) and the mean socket throughputs can again be attributed to the length of the video sequences being considered. The Martin video sequence encoded with $Q=18$ and 9 produces mean socket throughput values of 1.66 and 3.76 Mb/s, respectively, using 25.3 and 57.3 socket calls per second, respectively, during the transfer. The Ayers video sequence with the same Q parameters produces mean socket throughput values of 2.33 and 4.74 Mb/s, respectively, using 35.5 and 72.2 socket calls per second, respectively, during the transfer. Thus, the load placed on the client's CPU is significantly higher for the transfer and decoding of the Ayers video sequence than for the Martin video sequence. Section 5.1.5 deals with the performance issues of the VoD system in detail.

As expected, the mean socket throughput is clearly a function of both the Q parameter selected and the content of the video sequence. Consequently, the bandwidth required to support the transfer and successful decoding of a VBR encoded bit stream is a function of both the Q parameter selected and the video sequence's content. Therefore, the bandwidth required to support the transfer of a VBR sequence may not necessarily be very easy to determine. A similar result was found in [14]. While it was known that certain VBR bit streams, for example those containing periods in which the rate of encoding exceeds that of the transmission medium (i.e., 10 Mb/s), would pose a challenge to the VoD system, it was not known what CBR streams could be successfully transferred.

When CBR encoded video sequences with higher rates of encoding (e.g. 6 Mb/s) were considered, limitations of the VoD system become apparent as these sequences could not consistently be transferred and decoded successfully. For the 6 Mb/s CBR encoded bit

streams the mean socket throughput measured was much less than the nominal bit rate required by the decoder. In the case of the CBR Ayers sequence encoded at 6 Mb/s, a mean socket throughput of 5.93 Mb/s was observed during the transfer and decoding of the sequence. Thus, the ring buffer in application space at the client would be depleted after a sufficiently long period of time causing starvation to occur at the hardware decoder's buffer. It is simple observation: the rate of depletion (6 Mb/s) exceeds the arrival rate of the data in application space at the client (i.e., 5.93 Mb/s). These throughput results confirm that a bottleneck exists at the client in the VoD system as it can not maintain throughput values of 6 Mb/s while decoding the CBR stream encoded at the same rate. This bottleneck in the client was also evident when attempts to transfer and decode VBR encoded video sequences with a low Q parameter with comparable mean bit rates were made. Issues relating to the performance of the system are dealt with separately in more detail in Section 5.1.5.

5.1.4 Buffer Occupancy

The occupancy of the ring buffer in application space at the client was recorded over the duration of the transfer and decoding process. Observing the behaviour of the buffer over time is useful in properly dimensioning the client's application layer buffer to support the various MPEG-2 CBR and VBR stream transfers within the end-to-end flow control implemented in the VoD system.

As expected, during a CBR transfer the buffer drains and fills at nearly equal rates so the application layer buffer remains very near full for the entire duration of the transfer.

In other words, the buffer requirements for the transfer of a CBR encoded bit stream transfer are minimal. A typical buffer occupancy plot over time during the transfer of a CBR stream is given in Figure 5.1. A small buffer is required to allow for the overhead that exists in transferring the encoded video bit stream from the server to the client (i.e., various levels of encapsulation and reassembly) and any delays that may be incurred in the network. However, the application layer buffer requirements at the client are substantial for the transfer of VBR streams.

Depending on the Q parameter selected and the video sequence being transferred, the buffer requirements will vary. For high quality (i.e., low Q parameter) VBR streams, there typically exists periods in the stream with high sustained rates of decoding that cause the client's buffer to be consumed at a high rate. There are two reasons which will cause the client's buffer to be consumed at a high rate.

First, as previously stated, there may at times be periods in which the rate of decoding of the bit stream approaches or temporarily exceeds the transmission rate of the physical medium. Further, these periods, as well as others, may exist in a VBR stream in which the rate of decoding exceeds that of the maximum throughput of data traffic in the network imposed by the end-to-end flow control protocol. In other words, if the rate of decoding exceeds the maximum sustainable throughput of the system, then the application layer buffer will be depleted if it is not dimensioned properly.

Second, there may be periods in which acknowledgement messages from the client to the server are delayed as a result of the hardware decoder generating interrupts at very

high rates so that the client's CPU is consumed with handling the interrupts. If acknowledgements are not sent, then data from the server will not come. In this situation, if the decoder becomes very demanding, then buffer depletion will ensue. Consequently, to prevent starvation, the client's buffer must be defined to be sufficiently large so that despite the sustained high rates of decoding in which the buffer is depleted rapidly, the replenishing process may sustain the transfer.

Figure 5.2 shows the buffer occupancy when the VBR Ayers bit stream with $Q=9$ was transferred five times in a row. The figure depicts the buffer occupancy from the moment the decoding process is started after having filled up the ring buffer to $2/3$ of its capacity. As seen from the figure the storage (buffering) and bandwidth requirements of the VoD application are so demanding that a buffer of 1.5 MB can be easily drained in a few seconds depending on the bit rate of the source and length of this period.

Figure 5.1 Buffer occupancy plot during the transfer of the Ayers 3 Mb/s CBR stream (5 transfers)

Figure 5.2 Buffer occupancy plot during the transfer of the Ayers VBR Q=9 (5 transfers)

5.1.5 Performance Issues in the MPEG-2 Video Client-Server System

Performance bottlenecks were discovered to exist in the client's host operating system (LINUX 1.2.8) during the design and implementation of the MPEG-2 video client-server system. The bottlenecks found to exist include the device drivers used for the Ethernet adapter and the MPEG-2 hardware decoder and the kernel resources allocated to the client's socket.

In the client the Ethernet adapter and the hardware decoder exist as interrupt-driven devices under LINUX each with its own device driver. Consequently, the load placed on the CPU by the interrupts generated by device drivers for the transfer of data increases as the rate of encoding increases for CBR streams (e.g. to 6 Mb/s) or as the value of the Q parameter is decreased for VBR streams (e.g. to 9 for sequences like Pirates which have lesser amounts of temporal redundancy in certain groups of sequential frames). Since the hardware decoder generates an interrupt whenever the occupancy of its hardware buffer falls below a set threshold and since the Ethernet adapter generates an interrupt every time an Ethernet frame arrives at its network interface, it is clear that as the rate of encoding/decoding of the MPEG-2 bit stream rises, the ability of the client to support the successful transfer and decoding of the stream is reduced.

For example, with a 6 Mb/s CBR bit stream, Ethernet frames will arrive at the Ethernet adapter at a nominal mean rate of 6 Mb/s. This translates into handling 500

transfers of 1500³ bytes from the Ethernet adapter to a kernel buffer (i.e., socket's receive buffer) in main memory per second. Concurrently, the decoder will be generating interrupts at a rate of 733 per second (6 Mb/s/1024 bytes/interrupt) for data to be transferred to its hardware buffer from its device driver's buffer. Thus, during the transfer and decoding of a 6 Mb/s CBR encoded bit stream interrupts will occur at a rate of 1233 per second. Each interrupt generated will cause two context switches further increasing the overhead. Note that the number of interrupts generated per second increases linearly with the rate of encoding in a CBR stream. Therefore, the device drivers limit the client's ability to successfully receive and decode the MPEG-2 streams transferred from the server over the ATM MAN.

Additionally, the amount of data movement within the client system is also a significant performance factor. First, the data whose ultimate destination is the decoder, must be copied from the socket's receive buffer to the application layer ring buffer using the *receive* system call. Next, this same data must then be copied back to kernel space using the *write* system call to put the data into the decoder's device driver's buffer. Since the unit of transfer for both the *receive* and *write* system calls is a ring buffer element of 8 KB, 184 system calls per second must be completed to successfully decode the 6 Mb/s CBR encoded bit stream. Effectively then, to decode 8 KB of video data sent across the network to the client, two system calls must be made and a total of 6 and 8 hardware

³ Actually, 1492 bytes when the SNAP header is used as per IEEE 802.3. 1500 bytes is used to simplify calculations.

interrupts must be handled, for the Ethernet adapter and the hardware decoder, respectively.

The management of the ring buffer at the client also introduces some overhead in that the buffer exists in a shared memory segment shared by two processes (i.e., replenishing and decoding processes) that use semaphores for interprocess communication (IPC). Since IPC is used, system calls (e.g., *semop*) required to evaluate the value of a semaphore add to the load placed on the CPU during the transfer and decoding process. The management of the shared memory segment by the operating system, given its size, is also likely to increase the demand on the CPU. Changing the implementation to be thread-based (i.e., lightweight processes) would eliminate the overhead involved in the management of a shared memory segment since the threads would share a common address space; however, the IPC overhead would still exist.

Note that [30] points out that a send operation to write data to the network may not be allowed to complete until all incoming data has been received. While this situation will never occur in the VoD system, a similar delay in a send operation may occur when the hardware decoder is generating interrupts at a very high rate when a stream is being decoded (e.g., Ayers VBR $Q=9$). Consequently, acknowledgements used in the implementation of window flow control in the system may be delayed by the servicing of the decoder's interrupts. Hence, underflow may occur as the decoding process will be consuming the ring buffer contents when there is no incoming data from the network as a result of an acknowledgement message being delayed.

Prior to this work, the client's hardware consisted of a 486 66 MHz PC running the same version of LINUX. Using the current system implementation on the old hardware, but with a 8 KB window size (i.e., 1 UDP datagram), the transfer and successful decoding of various streams sent from the server to the client across a back-to-back Ethernet connection could not be realised. Simply upgrading the client's hardware to the current 90 MHz Pentium PCI bus system allowed the realisation of successful transfers of these same streams. This increase in performance is most likely a function of the faster CPU providing improved interrupt response coupled with the fact that the PCI system bus used in the Pentium system is superior to the one used in the 486 machine. A similar point is argued in [13] for Sun Sparc machines.

Another performance bottleneck identified in the VoD system is the size of the UDP receive socket buffer in kernel space allocated to the client. Others [36, 37] have found similar findings for TCP. In the VoD system implementation effective throughput at the application layer is a problem; however, it is a function of the simple window flow control that has been implemented in the application layer. Recall that the size of the receive buffer allocated to the client's socket is 60 KB. As a result, at most seven UDP datagrams can be stored in the buffer before loss ensues. Therefore, the maximum window possible with these resources is seven UDP datagrams or 60 KB (assuming improvements to the device driver for the decoder that would permit the Ethernet adapter to receive seven UDP datagrams back-to-back).

Currently, the VoD system uses a 24 KB window (three UDP datagrams) which makes use of only 40% of the buffering resources available. However, even if all 60 KB of the receive socket buffer could be used (i.e., for a 60 KB window), it still may not be sufficient to properly support multimedia communication over an ATM MAN [36, 37]. As a result, the link interconnecting the client and the server is underutilised placing an artificial limit on which MPEG-2 streams that can be successfully transferred and decoded across the ATM MAN.

This point also illustrates the synchronisation problem that exists between different layers in the protocol architecture [41]. In the VoD system the occupancy of the ring buffer at the client is used to drive the transfer of data from the server. Consider the situation in which there is only one free element in the ring buffer and the replenishing process has just copied the last of the three UDP datagrams of the current window into the free space. At this point the end of the current window has been reached and so the client sends an acknowledgement message to the server for the next three UDP datagrams. In this situation the client requests data from the server even though, at the current time, it has no room in the application space ring buffer to store it. Consequently, when the data arrives at the client, the data must remain in the receive socket buffer in kernel space until such a time that there is an available ring buffer element and the replenishing process runs to complete a *receive* call. This, then, limits the size of the window that can be used in the VoD system — the size of the window is limited to the size of the receive socket buffer. However, if the client only generates an acknowledgement when there is sufficient space

in its ring buffer to store all the requested data, the efficiency of the system decreases and underflow at the decoder may occur for certain demanding streams (e.g., Pirates $Q=9$).

One immediate consideration is the use of direct memory access (DMA) for the transfer of data to the decoder. With a DMA device driver, the number of interrupts, and, consequently, the amount of time spent by the processor handling interrupts, would be greatly reduced. Increasing the size of the kernel buffer used by the client's socket is another immediate consideration. Once the decoder's device driver's code is optimised so that the performance penalty that the Ethernet adapter takes is eliminated, increasing the size of receive socket buffer in the kernel would allow the use of a larger window size in the system. This task, however, would entail modifying the kernel socket data structures and code. The obvious benefit of the use of a larger window size is higher throughput as a result of improved use of bandwidth resources in the network.

5.2 Second Campaign of Measurements

5.2.1 Introduction

In this section of the thesis, the focus of study was on how intermediate network elements and the end-to-end flow control implemented on top of UDP affect the characteristics of the MPEG-2 traffic generated by the experimental VoD system in the ATM network. Moreover, there was a strong interest in understanding how the presence of generated bursty traffic affected the characteristics of the video traffic. These experiments were

conducted after having gained sufficient knowledge of the performance of the VoD system and the various bottlenecks that exist in it.

Many recent studies [14-19] have been conducted which involve the transport of constant and variable bit rate MPEG-2 video in ATM networks; however, many of these studies which have been completed have considered fairly homogeneous scenarios in which the only other traffic present in the ATM network, if any, are other MPEG-2 video sources. The goal, then, of conducting these measurements was to determine how the end-to-end flow control in the VoD system and the intermediate network elements affect the characteristics of the MPEG-2 traffic in the network under the various network loading conditions in which the generated traffic sources are bursty in nature, and are not other MPEG-2 sources.

5.2.2 Measurements

To determine the characteristics of the MPEG-2 traffic in the ATM network the experimental VoD system was employed. Without changing any hardware or network elements, the roles of the MPEG-2 video server and client were reversed (i.e., the server process ran on the old client (e.g., `bnr_mpeg2`) and the client became the old server (e.g., `crc_mpeg2`)). In this configuration, cell interarrival time and cell rate snapshots measurements were obtained under no load, moderately ($\rho = 0.4$) and heavily (0.8) loaded network conditions using the Adtech AX/4000 ATM testset [38]. The AX/4000 analyser was placed at the output of the first switch at NT premises (Figure 4.3). Here, the testset was used to obtain cell interarrival time measurements with a resolution of 0.5 μ s and cell

rate snapshots with resolutions of 100 and 500 ms of the MPEG-2 traffic in the ATM network.

5.2.3 Usage Parameter Control in Network

At the UNI at NT premises (Figure 4.3) the Usage Parameter Control (UPC) for the MPEG-2 connection was set to enforce a sustained cell rate and peak cell rate corresponding to 10 Mb/s. The maximum burst size was set to 32 cells. Since the physical connection to the router (the UNI at NT premises) is 10 Mb/s Ethernet, the specified sustainable and peak rates of 10 Mb/s were thought to be rational choices for the UPC parameters. Consequently, all MPEG-2 traffic entering the ATM network was compliant according to the UPC for the connection. However, since the UPC is not used in the first FORE switch at NT premises (Figure 4.3) where the generated traffic is applied, it is possible that at the output of the switch the MPEG-2 traffic may no longer be compliant in terms of the UPC for the connection specified at the UNI. How a generated load when multiplexed with the MPEG-2 video traffic at a shared output port of a switch could affect the characteristics of the MPEG-2 traffic leaving the output port of the switch was, of course, a matter of interest and study.

5.2.4 Generating Traffic to Load the ATM Network

The AX/4000 ATM testset traffic generator was used to generate traffic to load the output port of the first FORE ASX-200 switch in the ATM network at NT premises used by the MPEG-2 client-server system. Therefore, at the output port of the ATM switch the

generated traffic and the MPEG-2 traffic were multiplexed and sent on to the second ATM switch in the network. The cell rate snapshots and the cell interarrival time measurements were taken at the output of the first switch (Figure 4.3).

The "Randomly spaced bursts" traffic model in the traffic generator was used to generate bursty traffic to load the output port of the switch at the desired percentage of OC-3 rates (e.g. 40 or 80%). The randomly spaced bursts traffic model is defined in terms of three parameters: *average burst gap*, *average burst length*, and *average burst cell rate*. The first two parameters are measured in cell interarrival time units (2.83 μ s at OC-3 rates) and the third is measured in cells/s. The foreground traffic source enters the "send burst" state after a random time period based on the average burst gap. In the "send burst state", the foreground source sends traffic at a defined average cell rate with Poisson cell arrivals. The foreground source returns to the idle state after a random amount of time based on the average burst length [38].

To ensure that the results obtained when the output port was loaded at 80% could be compared to those obtained when the port was loaded at 40%, the average burst gap and the average burst length were fixed at 250.0 and 333.33 cell interarrival times, respectively, and the average cell burst rate was varied to achieve the desired traffic load. Seven (the maximum number available) independent traffic sources were used in the AX/4000 traffic generator for generating the bursty traffic; the same parameters for the randomly spaced bursts traffic model were used for all seven traffic sources. Thus, the sum

of the average bandwidths of the seven traffic sources consumed the desired percentage (i.e., 40 or 80%) of the bandwidth at OC-3 rates.

5.2.5 Unloaded Network Conditions

5.2.5.1 Cell Rate Snapshots

The best case scenario was considered in which the only traffic on the ATM network was the MPEG-2 video bit stream being sent from the server to the client. During the transfer of the bit stream (actually, the same stream concatenated four times in order to create a longer sequence) cell rate snapshots with resolutions of 100 and 500 ms were taken.

With a resolution of 100 ms, the cell rate snapshots were plotted against the MPEG-2 video source rate statistics. Figures 5.3 and 5.4 compare the bit rates of the 3 Mb/s CBR and $Q=12$ VBR Ayers MPEG-2 streams, respectively, with the bit rates (converted from the cell rates) measured at the output of the FORE ASX-200 switch [42] to the ATM network (Figure 4.3). To compute the MPEG-2 video source rates, the average bit rate was computed over groups three consecutive frames in each stream (i.e., each frame is $1/24$ s in duration, so $3 \cdot 1/24 \text{ s} = 125 \text{ ms} \cong 100 \text{ ms}$). The cell rate snapshots are evenly spaced at 100 ms intervals. Thus, each point in the cell rate snapshot plot represents the average bit rate seen on the ATM network in the past 100 ms while every data point of the MPEG-2 video source rate plot represents the source rate over a 125 ms epoch; however, the two intervals are sufficiently close for comparison purposes.

The results indicate that with a resolution of 100 ms, the end-to-end flow control implemented at the application level has a noticeable impact on the video traffic characteristics at the ATM level. Figure 5.3 shows the CBR encoded video bit rate to have minor fluctuations around the average bit rate of 3 Mb/s. The measured bit rate, on the other hand, has more significant fluctuations around a higher average bit rate due to the end-to-end flow control and an encapsulation overhead of approximately 14% for the video data transported in the ATM network. Figure 5.4 shows similar behaviour for the $Q=12$ VBR Ayers source. Further, there are many 100 ms epochs in the trace in which the measured rate is 0 during the transfer of the VBR sequence.

These intervals correspond to periods in which the client is not requesting data from the server either because its application space ring buffer is full or the replenishing process is not running. These periods may also correspond to delays introduced at the server (Section 5.2.5.2.1). Additionally, it is also known that the hardware decoder generates a large numbers of interrupts to request data from its device driver's kernel buffer during the decoding of the Ayers VBR streams causing a high utilisation of the client's CPU. It is therefore possible that the 100 ms epochs in which the measured cell rate is 0 correspond to the times when the client's CPU is consumed with handling the decoder's interrupts.

Figure 5.3 Measured cell rate versus source rate statistics — Ayers CBR 3 Mb/s (100 ms resolution); unloaded conditions

Figure 5.4 Measured cell rate versus source rate statistics — Ayers VBR $Q=12$ (100 ms resolution); unloaded conditions

Figure 5.5 Measured cell rate versus source rate statistics — Ayers CBR 3 Mb/s (500 ms resolution); unloaded conditions

Figure 5.6 Measured cell rate versus source rate statistics -- Ayers VBR Q=12 (500 ms resolution); unloaded conditions

The same transfers as above were repeated having changed only the resolution of the cell rate snapshot to 500 ms. With a resolution of 500 ms, the time period of over which the cell rate snapshots are completed is sufficiently long to record the entire transfer of all four streams over the ATM network. Again, the cell rate snapshots were plotted against the video source rate statistics having converted the two sets of values to bit rates. This time, the average bit rate over consecutive groups of 12 frames for each stream was computed (Each frame is $1/24$ s in duration, so $12 \cdot 1/24 \text{ s} = 500 \text{ ms}$). Thus, each point in the cell rate snapshot plot represents the average bit rate seen on the network in the past 500 ms and every data point of the MPEG-2 video source rate plot represents the average source rate over a 500 ms epoch

Figures 5.5 and 5.6 contrast the measured bit rate in the network with the source rate statistics of the Ayers 3 Mb/s CBR and $Q=12$ VBR streams, respectively. As before, there are more significant fluctuations in the measured bit rate in the network than in the video source rate statistics; however, the fluctuations are not as pronounced due to the resolution used for the measurement. In other words, the use of a longer averaging interval (i.e., 500 versus 100 ms) causes the fluctuations to be less visible. Note that the overhead (approximately 14%) of the encapsulation of the video data for transport in the ATM network is clearly evident in these plots. Similar results were found for the CBR and VBR Pirates and Martin video streams.

5.2.5.2 Cell Interarrival Times

The distributions of the cell interarrival times obtained for the 3 Mb/s CBR and $Q=12$ VBR Ayers video streams under no load network conditions are given in Figures 5.7 and 5.8, respectively. The cell interarrival time distributions obtained under no load network conditions were found to be nearly identical for all the MPEG-2 video streams considered (CBR at r_i and VBR with $Q = Q_i$). The mean cell interarrival time was found to be approximately 42.9 μ s for all the video sequences. In the distributions there are a number of distinct modes each separated by large gaps. Note that the range of the distribution is limited to 255 μ s when a resolution of 0.5 μ s is used for the cell interarrival time measurement in the AX/4000 analyser.

At the UNI the UPC was set to enforce a mean and a peak bit rate of 10 Mb/s for the MPEG-2 traffic with a maximum burst size of 32 cells. The router uses the "Null encapsulation" or "VC multiplexing" [22] where the end points of a VC is the Internet Protocol layer. Therefore, bursts of 32 cells (1 IP datagram — 20 cells in the last of the six IP datagrams of the UDP datagram used in the VoD system) will be forwarded to the ATM network at a time. The measured mean cell interarrival time of 42.9 μ s corresponds to the sustained and peak cell rates as specified by the UPC for the connection (i.e., mean rate = 1 cell / 42.9 μ s = 424 bits / 42.9 μ s $\approx$ 10 Mb/s). Considering that the MPEG-2 cells are sent into the ATM network at a mean rate of 10 Mb/s when the arrival rate of the incoming IP datagrams must be lower because of the transmission medium (i.e., 10 Mb/s Ethernet), the router must be servicing the IP datagrams and subsequently forwarding the

encapsulated IP datagrams to the ATM network at the mean rate specified by the UPC for the connection. As a result, the router determines the characteristics of the MPEG-2 traffic in the ATM network under no load conditions.

Thus, the manner in which the CISCO 7000 router services the MPEG-2 connection causes the cell interarrival time distributions to be nearly identical for all the CBR and VBR video sequences. As a result, the largest mode is around the mean cell interarrival time of 42.9 μ s. The remaining modes, all of much smaller and approximately equal amplitude, are a result of various delays incurred at the server (section 5.2.5.2.1).

In order to determine more precisely the manner in which the router services the MPEG-2 connection, the AX/4000 analyser was used to capture 8192 cells at a time on the connection. (The analyser is capable of capturing at most 8192 complete cells at one time) The cells have been given identifiers relative to their position in the cell capture completed with the analyser; the cells are numbered from 0 to 8191 with the first cell in the capture identified by 0. Having captured the cells on the MPEG-2 connection, their relative interarrival times were examined.

Figure 5.7 Cell interarrival time distribution — Ayers CBR 3 Mb/s (0.5 μs)

Figure 5.8 Cell interarrival time distribution — Ayers VBR Q=12 (0.5 μs)

5.2.5.2.1 CBR MPEG-2 Cell Arrival Snapshots

Consider first the case of CBR MPEG-2 video streams encoded at a bit rate r_i . An efficient VoD system implementation should produce rates of $r_i + Cr_i$ within the transport network, where Cr_i represents some allowable amount for the encapsulation overhead. When the high resolution measurements (i.e., 0.5 μ s) of the MPEG-2 traffic in the ATM network are considered, however, the mean rate of the traffic was found to be 10 Mb/s.

Simply matching the network capacity to the nominal rate r_i of encoding for a CBR stream while allowing for a fixed overhead for the transport of the data in the network (i.e., $r_i + Cr_i$ where $C = 0.14$ in the VoD system) is sufficient to support the transfer. This claim is based on measurements taken at the socket interface in the client which showed that the rate of data transfer into application space was that of the nominal rate of encoding of the stream being transferred and decoded. This is a function of the end-to-end flow control implemented in the VoD system. Thus, the measured mean rate of 10 Mb/s for the transport of r_i CBR MPEG-2 streams, where $\max(r_i) = 4$ Mb/s, indicates a very significant waste of network bandwidth.

Figure 5.9 shows the relative interarrival times for cells 4001 to 8000, inclusive, for the Ayers 3 Mb/s CBR stream. Note the regular peaks in the interarrival times of the cells in the figure. These peaks correspond to cells that are from consecutive 24 KB windows in the VoD application; each 24 KB window contains 540 cells. The first cell in window W_i will always have a large interarrival time compared to all the other cells. This is due to the fact that the server sends the video data in bursts of 24 KB (three UDP

datagrams) before stopping and waiting for an acknowledgement before sending the next 24 KB of data. Consequently, the end-to-end flow control implemented at the application layer is responsible for the fluctuations seen in the cell rate snapshot measurements presented in the previous section when the averaging interval for the measured cell rate is 100 ms (Figure 5.3 and 5.4).

There are no cells in this capture that have interarrival times larger than those found in the first mode (e.g., $35 \leq t \leq 50 \mu s$) of the cell interarrival time distribution presented in Figure 5.7. This is because the capture was taken during the transfer of a CBR stream once the buffer has already been prefilled.

Before decoding of the first stream begins, the ring buffer in application space at the client is filled to 2/3 its capacity (approximately 2.5 MB). During this time video stream data is sent from the server to the client as quickly as the end-to-end flow control permits. However, since the server is single-threaded, there is a limit to the amount of data that can be prefetched before an acknowledgement arrives from the client. The result is that there is a delay incurred at the server in sending the data to the client (section 5.2.5.2.2). The modes of the cell interarrival distributions, other than the one in which the mean appears, are a result of these delays.

Therefore, the other modes which appear in the cell interarrival time distribution in Figure 5.7 are not visible in the relative cell interarrival times seen in the cell capture of Figure 5.9.

5.2.5.2.2 VBR MPEG-2 Cell Arrival Snapshots

To determine in finer detail the fashion in which the router services the MPEG-2 connection during a VBR stream transfer, several captures of 8192 cells were taken during the transfer of the $Q=12$ Ayers stream.

Figure 5.9 Relative cell interarrival times of 4000 cells of Ayers 3 Mb/s CBR MPEG-2 traffic

Figure 5.10 Relative cell interarrival times of 4000 cells of Ayers VBR $Q=12$ MPEG-2 traffic

Figure 5.10 presents the relative interarrival times of 4000 cells (cells 4001 to 8000, inclusive, in the capture) of one such capture. As in the CBR 3 Mb/s stream transfer (Figure 5.9), there are regular peaks (of the order of 10 000 μs) in the interarrival times of the captured cells that correspond to cells that are from consecutive 24 KB windows in the VoD application. In contrast to the CBR 3 Mb/s relative cell interarrival time plot, Figure 5.10 shows cells which have interarrival times above the range of the first mode (e.g., $35 \leq t \leq 49.5 \mu\text{s}$) in the cell interarrival time distribution for the Ayers $Q=12$ VBR stream (Figure 5.8) and below those of the cells which belong to separate windows. The interarrival times of these cells constitute the remaining modes which appear in the cell interarrival time distribution for the stream. The delays such cells experience within a

2 OF/DE 2

FM-1 3 1/2" x 4" PHOTOGRAPHIC MICROCOPY TARGET
NBS 1010a ANSI/ISO #2 EQUIVALENT

PRECISIONSM RESOLUTION TARGETS

given 24 KB window during the transfer of the $Q=12$ VBR Ayers video stream are now explored in greater detail.

The interarrival times of 540 cells within one 24 KB window during the transfer of the Ayers $Q=12$ VBR stream are given in Figure 5.11. The window consists of cells 4952 to 5490 in the capture of 8191 cells. Note that the first data point, corresponding to the first cell (cell 4951) in the 24 KB window is omitted so that the reader may see more clearly the interarrival times of the remaining cells in the window as a result of the reduced range of the y-axis (For completeness, the first cell has a relative interarrival time of 58 917 μs). In Figure 5.11 there are four instances where the measured cell delay is very high relative to the other cells within its window. The four instances occur at the cells having the identifiers 5015, 5047, 5195, and 5311. In all cases, the cell prior to the one whose interarrival value is significantly higher than the cells preceding has a payload type value of 1. The last cell of an AAL5 PDU has a payload type of 1 [20].

To further examine these delays experienced by the cells within the same cell capture of 540 cells from the 24 KB window presented above is used and the rate at which cells are admitted to the network in subsequent epochs of time are considered. At OC-3 rates (i.e., 155.45 Mb/s), one cell interarrival time slot is 2.83 μs in length. Recall, that the mean cell interarrival time under no load conditions was found to be 42.9 μs (corresponding to a mean/peak rate of 10 Mb/s). Therefore, on average 1 cell should be sent out to the network by the router every 15 time slots at OC-3 rates according to the mean rate specified by the UPC. ($N_s = 42.9 \mu\text{s} / 2.83 \mu\text{s/slot} \approx 15$ slots). The AAL5 PDU is

32 cells in length. The number of cells sent in an interval spanning $32 \cdot 15$ slots = 480 slots (i.e., 1358.4 μ s) can be found by considering the relative interarrival times of the cells captured by the analyser. The maximum number of cells permitted to enter the network in an interval of 480 cell time slots is 32 as specified by the UPC. Using the relative interarrival times of the cells given in Figure 5.11 a count of the number of cells sent out to the network by the router during each epoch of 480 slots can be found. Taking the sums of the number of cells admitted to the network in each epoch and dividing each of the sums by the length of the epoch in seconds, a plot of the admitted cell rates to the network in each epoch can be derived. This cell rate plot is given in Figure 5.12. The admitted cell rates have been converted to bit rates so as to facilitate comparison with the parameters of the UPC for the connection.

Figure 5.11 Relative cell interarrival times of cells within one 24 KB window (Ayers VBR $\rho=12$, cells 4952-5490)

Figure 5.12 Bit rates of admitted traffic within one 24 KB window (480 cell time slots or 1358.4 μ s resolution)

5.2.6.2 Cell Interarrival Times

With a load factor of 0.4, the cell interarrival time distributions differ significantly from those obtained under unloaded network conditions. The cell interarrival time distributions obtained under moderately loaded network conditions for the 3 Mb/s CBR and $Q=12$ VBR Ayers video streams are given in Figures 5.15 and 5.16, respectively. The mean cell interarrival time was found to be approximately $43.1 \mu\text{s}$ for all the streams sent through the network. Compared to unloaded network case, the mean cell interarrival time is slightly larger. This increase can be attributed to the change in the cell interarrival time distribution caused by the presence of the generated bursty traffic.

As in the unloaded network case, the cell interarrival time distributions, multimodal with large gaps between successive modes, are nearly identical for all the MPEG-2 video streams considered; however, the shape of the modes has changed. The modes present in the interarrival time distributions show a much greater dispersion under 40% network load conditions than under no load conditions. Moreover, there are now non-zero cell counts in the range of 0 to $25 \mu\text{s}$ which were not present in the interarrival time distributions for the unloaded network case. The presence of these cells indicates that "cell clumping" is occurring at the output of the FORE switch. Note that since the UPC in the FORE switch was not used for these experiments, cells from the generated traffic sources and the video stream will be serviced and sent onto the network as the switch tries to work under no loss conditions. Since cell clumping of the MPEG-2 traffic was not observed in the unloaded network condition experiments, it must be a direct result of the generated load.

the hard disk before being written to the network. So, the disk read and network write operations are as follows

{ prefetch D_i , wait for acknowledgement, $W.D_i$, $R.D_i$, $W.D_i$, $R.D_i$, $W.D_i$ }

It would therefore be expected that the cells corresponding to IP datagrams 7 and 13 (beginning of the second and third UDP datagrams) should have interarrival times considerably larger than those of the other cells within the same 24 KB window. Results show, however, that this is not the case. Rather, the cells corresponding to IP datagrams 8 and 14 are more likely to be delayed than IP datagrams 7 and 13. This is most likely a function of the contention for the system bus in the server that might exist for the above order of read and write operations coupled with how the send operation from the network interface unit is carried out. That is, $R.D_i$ causes the second IP datagram sent as part of the $W.D_{i-1}$ to be delayed at the network interface unit.

From the high resolution cell interarrival time measurements presented it is clear that the characteristics of the MPEG-2 traffic sent onto the ATM network in the absence of any other traffic in the network are determined by the end-to-end flow control implemented in the application layer of the VoD system and the manner in which the router services the MPEG-2 connection with some allowance for delays incurred in the send operation at the server.

5.2.6 Moderately Loaded Network Conditions

5.2.6.1 Cell Rate Snapshots

Cell rate measurements with 100 and 500 ms resolutions of the MPEG-2 traffic in the network were taken at the output of the first FORE switch having loaded the output port at 40% of OC-3 rates. The cell rate snapshots obtained of the Ayers 3 Mb/s CBR and $Q=12$ VBR MPEG-2 video traffic in the presence of a moderate traffic load are given in Figures 5.13 and 5.14, respectively. Further to the changes caused by the end-to-end flow control to the characteristics of the MPEG-2 traffic carried in the network, as evidenced by the cell rate and cell interarrival time measurements obtained under no load network conditions, the generated traffic applied to the output port of the first FORE switch has a noticeable impact on the video traffic characteristics at the ATM level.

Figures 5.13 and 5.14 show that there are much more significant fluctuations present in the measured video traffic than in the video source rate statistics; however, the fluctuations present in the measured cell rates under moderately loaded network conditions are much more significant than those observed under no load conditions. This remark is especially true of the $Q=12$ VBR Ayers sequence. These changes in the characteristics of the MPEG-2 traffic in the network are a direct result of the generated traffic applied to the output port of the switch shared by the MPEG-2 connection and are in addition to the changes introduced by the end-to-end flow control and the manner in which the router services the MPEG-2 connection. Evidence to support this claim is given in section 5.2.6.2 when the cell interarrival time measurements are considered.

The same transfers and measurements as above were performed having changed only the resolution of the cell rate snapshot 500 ms. Because of the decreased resolution used for the measurement, the cell rate snapshots of the video traffic obtained appear unchanged from those obtained in the unloaded network condition experiments. For the streams considered, there are greater fluctuations in the measured cell rates in the network which are more significant than those found in the source rate statistics of their respective streams.

It is evident, as in the unloaded network case, that the end-to-end flow control protocol employed by the VoD application determine the characteristics of the MPEG-2 video application's traffic in the ATM network under moderately loaded network conditions. Moreover, when the generated traffic, with an average bandwidth of 40% of OC-3 rates, is multiplexed with the MPEG-2 traffic at the same output port on the switch, the amount of fluctuation present in the measured cell rate of the MPEG-2 traffic in the ATM network increases.

Figure 5.13 Measured cell rate versus source rate statistics —
Ayers CBR 3 Mb/s (100 ms resolution)

Figure 5.14 Measured cell rate versus source rate statistics —
Ayers VBR $Q=12$ (100 ms resolution)

Figure 5.15 Cell interarrival distribution — Ayers CBR 3 Mb/s (0.5 μ s); 40% loaded

Figure 5.16 Cell interarrival distribution — Ayers VBR $Q=12$ (0.5 μ s); 40% loaded

5.2.6.2 Cell Interarrival Times

With a load factor of 0.4, the cell interarrival time distributions differ significantly from those obtained under unloaded network conditions. The cell interarrival time distributions obtained under moderately loaded network conditions for the 3 Mb/s CBR and $Q=12$ VBR Ayers video streams are given in Figures 5.15 and 5.16, respectively. The mean cell interarrival time was found to be approximately $43.1 \mu\text{s}$ for all the streams sent through the network. Compared to unloaded network case, the mean cell interarrival time is slightly larger. This increase can be attributed to the change in the cell interarrival time distribution caused by the presence of the generated bursty traffic.

As in the unloaded network case, the cell interarrival time distributions, multimodal with large gaps between successive modes, are nearly identical for all the MPEG-2 video streams considered; however, the shape of the modes has changed. The modes present in the interarrival time distributions show a much greater dispersion under 40% network load conditions than under no load conditions. Moreover, there are now non-zero cell counts in the range of 0 to $25 \mu\text{s}$ which were not present in the interarrival time distributions for the unloaded network case. The presence of these cells indicates that "cell clumping" is occurring at the output of the FORE switch. Note that since the UPC in the FORE switch was not used for these experiments, cells from the generated traffic sources and the video stream will be serviced and sent onto the network as the switch tries to work under no loss conditions. Since cell clumping of the MPEG-2 traffic was not observed in the unloaded network condition experiments, it must be a direct result of the generated load

placed on the output port of the FORE switch in the absence a UPC for the MPEG 2 video connection. Thus, the switch introduces very significant amounts of cell delay variation. Note that this effect is much more evident in the Ayers 3 Mb/s CBR video transfer than in the Ayers VBR $Q=12$ video transfer. This can be attributed to the fact the CBR sequence uses the connection at a fixed rate for the duration of the transfer whereas the VBR sequence does not. For example, during the transfer of the VBR sequence there are 100 ms epochs where the measured cell rate drops to 0 on a number of occasions (Figure 5.14).

Clearly, under moderately loaded conditions, the router, the load placed on the ATM switch and the end-to-end flow control implemented in the application layer determine, for the most part, the characteristics of the MPEG-2 traffic sent onto the ATM network. The presence of cell clumping of the MPEG-2 traffic in the network, a direct result of the load placed on the ATM switch, is further evidence that the intermediate network elements largely influence the characteristics of the video traffic submitted to the ATM network.

5.2.7 Heavily Loaded Network Conditions

5.2.7.1 Cell Rate Snapshots

To load the output port of the FORE switch at exactly twice the factor used in the moderately loaded in network condition experiments, the burst cell rate in the "Randomly spaced bursts" traffic model was doubled from that used previously. Having loaded the

output port at 80% of OC-3 rates, cell rate snapshots with 100 and 500 ms resolutions of the MPEG-2 traffic were taken.

The cell rate snapshots at 100 ms resolution of the Ayers 4 Mb/s CBR stream and the Pirates $Q=18$ VBR stream are given in Figures 5.17 and 5.18, respectively, as representative examples of the results of this experiment. Upon careful examination, the fluctuations present in the two figures are found to be more significant than those present in the cell rate snapshots conducted under the previous network loading conditions. This behaviour is even more apparent when the same measurements are conducted at a lower resolution (e.g., 500 ms).

Consider Figures 5.19 and 5.20 in which the cell rate snapshots at 500 ms resolutions are given for the same two streams considered above. Figure 5.19 shows that there are very significant oscillations present in the measured cell rate of the Ayers 4 Mb/s stream in the ATM network. The second figure which depicts the measured cell rates for the $Q=18$ VBR Pirates stream does not show the same degree of oscillation as is present in measured cell rates of the 4 Mb/s CBR Ayers stream. This is because the 4 Mb/s CBR Ayers stream uses the network bandwidth resources at a fixed rate for the duration of the transfer whereas the $Q=18$ VBR Pirates sequence does not despite the fact the mean bit rate over the entire sequence is comparable to 4 Mb/s. Additionally, with a CBR source, any deviation from the source rate statistics will be easy to see. In the case of a VBR source this is less true. Nonetheless, the oscillations are present but due to the resolution used for the measurement and the nature of the VBR source, they are less discernible

compared to the CBR source. To this end, an explanation of the behaviour of the CBR video traffic in the network will be given with the knowledge that similar reasoning applies to the VBR sources as well. Evidence that the same behaviour is present in both the CBR and VBR streams considered here will be given in section 5.2.7.2 when the cell interarrival time measurements are given.

There are periods in the cell rate snapshots of the 4 Mb/s Ayers stream in which the measured bit rate for the traffic in the network falls dramatically and then rebounds so as to produce a very high measured bit rate. The drop in the measured bit rate is attributable to the load placed on the output port that is used by the MPEG-2 connection. Since the MPEG-2 traffic and the generated traffic are multiplexed together at the same output port of the ATM switch, it is highly probable that during the transfer of the stream that bursts from several of the seven independent sources will occur very close together. In this case, the switch attempts to work under no loss conditions while handling the large volume of traffic from the eight different sources (including the video source). Since the UPC at the FORE was disabled, cells carrying the MPEG-2 traffic may be delayed, causing the drop in the measured bit rate. Following the delay, the switch will service the delayed cells very quickly sending them into the network at rates very near the OC-3 line rate. As a result, the drop in the measured bit rate is followed by a equally sharp rise in the rate as the delayed cells are serviced very quickly. The durations of these two periods are approximately equal and the difference between in the measured rates (i.e., low and high) is approximately zero.

Figure 5.17 Measured cell rate versus source rate statistics — Ayers CBR 4 Mb/s (100 ms resolution)

Figure 5.18 Measured cell rate versus source rate statistics — Pirates $Q=18$ VBR (100 ms resolution)

Figure 5.19 Measured cell rate versus source rate statistics
— Ayers CBR 4 Mb/s (500 ms resolution)

Figure 5.20 Measured cell rate versus source rate statistics
— Pirates $Q=18$ VBR (500 ms resolution)

Thus, very significant changes to the characteristics of the video traffic result from the heavier traffic load multiplexed at the output port of the FORE switch used by the MPEG-2 connection. Further, in addition to the large oscillations present in the measured cell rates, there other smaller fluctuations present in the measured cell rates of the video traffic which are larger than those of the video source rate statistics. It is clear then, that the intermediate network elements and the load placed on the shared resources (e.g., link, output port of a switch) will greatly influence the characteristics of the video traffic carried in the network. To gain a further understanding of the behaviour of the system under heavily loaded network conditions, cell interarrival time measurements of the MPEG-2 traffic were conducted.

5.2.7.2 Cell Interarrival Times

The cell interarrival time distributions obtained under heavily loaded network conditions for the 4 Mb/s CBR Ayers and $Q=12$ VBR Pirates video streams are given in Figures 5.21 and 5.22, respectively. The cell interarrival time distributions found with a load factor of 0.8 differ significantly from those obtained under no load and moderately loaded network conditions. As in the two previous cases, the cell interarrival time distribution, with a mean cell interarrival time of approximately 44.0 μ s, is nearly identical for all the MPEG-2 streams considered.

Under heavily loaded network conditions, the load placed on the ATM switch by the generated traffic is the primary factor in determining the characteristics of the MPEG-2 traffic when high resolution measurements are considered. Video traffic leaving the

router is compliant according to the UPC for the connection; however, the load placed on the output port of the FORE switch cause the shape of the MPEG-2 traffic to change dramatically.

The most obvious change to cell interarrival time distribution under heavily loaded network conditions is that the modes in the distribution have merged. Additionally, the cell counts in the range of 0 to 25 μ s are three orders of magnitude larger than those measured under moderately loaded network conditions. Thus, there is much more significant amounts of "cell clumping" occurring in the ATM network as a result of the increased generated traffic load. As a result, some cells of the MPEG-2 traffic are delayed, causing lower cell rates to be observed on the network, while others are sent out to the ATM network at extremely high rates (i.e., very near OC-3 rates). This behaviour is also evident in cell rate snapshot measurements of the video traffic at resolutions one million times less (Figure 5.19).

Thus, under heavily loaded network conditions the characteristics of the MPEG-2 traffic sent onto the ATM network are determined by the load placed on the output port in the ATM switch used by the MPEG-2 connection, the router, and the end-to-end flow control implemented in the application layer.

Figure 5.21 Cell interarrival distribution — Ayers CBR 4 Mb/s (0.5 μs)

Figure 5.22 Cell interarrival distribution — Pirates VBR $Q=18$ (0.5 μs)

6. CONCLUSIONS

Video services including entertainment video, such as Video on Demand (VoD), are expected to constitute the most significant source of traffic in future BISDN deployments. Improvements in network and video technologies will drive the growth of the video services industry. The success of such ventures, however, will depend on the careful planning of the allocation of resources in the ATM networks as well as at the network boundaries used to provide VoD to customers' homes. The use of MPEG video coding for the compression of the video source material greatly facilitates the use of ATM-based BISDNs for providing high quality video services to residential customers.

The thesis addressed a number of the planning issues involved in the transport of MPEG-2 video over the OCRI net ATM MAN using an experimental VoD testbed implemented on low end workstations. The conclusions reached in the thesis with regard to the transport of the video traffic in the ATM network are by no means restricted to the specific implementation of the VoD system. There are, however, conclusions given that are implementation-specific, but even many of these claims can be applied a more general fashion. Consequently, the conclusions reached with regard to the characteristics of the MPEG-2 video traffic in the ATM network should be considered to be implementation-specific free.

The primary conclusions of the research and the results of the thesis can be summarised as follows:

1. Low end workstations can be used to provide flexible, scaleable and cost-effective VoD solutions, but there are some inherent performance drawbacks in doing so:
 - limited resources in the operating system for supporting continuous multimedia applications place an artificial limit on the system's ability to provide high quality video streams over the network
 - limitations placed on the quality of the video streams that can be transferred across the network are more a function of the performance of the workstation's network interface unit when used in conjunction with an interrupt-driver MPEG-2 video hardware decoder;
2. The use of lightweight internetworking protocols like UDP supplemented with software that implements a simple flow control protocol with an adequately dimensioned buffer resource in the application layer are sufficient for providing MPEG-2 video services over an ATM MAN;
3. The buffering resources required at the client to support the successful transfer and decoding of MPEG streams is a function of the type of MPEG-2 encoding used in the compression of the source material:
 - for CBR streams the buffering requirements are minimal with only a small buffer required to smooth out any variations in the arrival process of data to the client

- for VBR streams the buffering requirements are very significant. Since there exists the possibility of having periods within the encoded stream in which the rate of encoding approaches or exceeds the transmission rate of the network, a sufficiently large buffer must be provided to avoid data starvation in the hardware decoder;
4. Intermediate network elements and the end-to-end flow control implemented in VoD systems will determine the characteristics of the MPEG-2 video traffic in the ATM network:
- The manner in which the routers in the VoD experimental system handle the incoming IP datagrams is consistent with forwarding the encapsulated IP over the ATM network at a mean rate equal to that of the sustained cell rate specified as part of the connection's UPC when a high time scale resolution is considered.
 - When a lower time resolution is considered (e.g., 100 or 500 ms), the cell rate plot of the video traffic in the ATM network looks very much like that of the video source rate statistics, but with variations due to the manner in which the connection is serviced by the router and the end-to-end flow control employed.
 - Loading the output port of an ATM switch used by the MPEG-2 connection with bursty traffic results in cell clumping of the outgoing video traffic when no UPC is used at the switch for the MPEG-2 connection. In a real deployment that the resultant (non-compliant) MPEG-2 traffic sent into the ATM network

would likely be enforced at a NNI at a public ATM network by a NPC causing cell delay variation (CDV) to scale to cell loss. Since the decoder is not tolerant data loss, the CDV caused by the loading of the output port has important implications on the end viewer. To mitigate the severity of the CDV introduced by the loading of the output port the switch, and to reduce the possibility of cell loss due to the violation of a negotiated traffic contract, a UPC should be used for the MPEG-2 connection at the output port of the loaded switch;

5. VBR MPEG-2 streams with different Q parameters exhibit different behaviours when the autocorrelation of the number of bits per frame are considered:
 - the effects of the MPEG algorithm are smoothed out when the Q parameter is decreased (i.e., for better visual quality)
 - the use of a high valued Q parameter produces an autocorrelation in which contains several flat sections
 - this has important implications when VBR MPEG-2 traffic is multiplexed in an ATM network if the same demanding streams (or similarly demanding streams) are requested by numerous individuals from the same video server (e.g., a motion picture newly available on the VoD system)
6. CBR MPEG-2 streams exhibit consistent behaviour when the autocorrelation of the number of bits per frame are considered:

- there are regular peaks in the autocorrelation which correspond to the encoding process (i.e., $N=12$, $m=2$) and the degree of compression achieved for the various picture types (e.g., I, P, B)

Therefore, a thorough understanding of the impacts of the intermediate network elements and the end-to-end protocols employed in the system on the characteristics of the MPEG-2 traffic in the nonhomogeneous scenarios is the prerequisite of commercially successful video services over public ATM networks.

REFERENCES

- [1] Minzer, S. "Broadband ISDN and Asynchronous Transfer Mode (ATM)". *IEEE Communication Magazine*. September 1989 pp 17-24.
- [2] Parker, M., Robinson, P., Wade, R., Alley, D., Adam, P., Le Moine-Py, B., and Saint-Hilaire, P. B. "The European Pilot". Proceedings of the 15th International Switching Symposium (ISS '95). Berlin. April 23-28, 1995. pp 146-150.
- [3] Grovenstein, L., Pittman, C., Simpson, J., and Spear, D. "NCIH services, architecture, and implementation" *IEEE Network Magazine*. November/December 1994. pp 19-22.
- [4] Adams, M. "Real time MPEG asset delivery over ATM", 1995 NCTA Technical Papers. pp 315-326.
- [5] Perry, T. S. "The trials and travails of interactive tv". *IEEE Spectrum*. April 1996. pp 22-28.
- [6] ISO/IEC 13818-2. "Generic coding of moving pictures and associated audio". March 1994
- [7] Le Gall, D. "MPEG: A video compression standard for multimedia applications". *Communications of the ACM*. Vol 34. No 4. April 1991. pp 46-58.
- [8] Kishino, F. et al., "Variable bit-rate coding of video signals for ATM networks", *Journal on Selected Areas in Communications*. Vol 7. No 6. June 1989. pp 752-760.

- [9] Reininger, D., Raychaudhuri, D., Melamed, B., Sengupta, B., and Hill, J
 "Statistical multiplexing of VBR MPEG compressed video on ATM networks"
 IEEE INFOCOM '93. March 28-April 3, 1993 pp 919-926
- [10] Roberts, J. and Guillemin, F., "Jitter in ATM networks and its impact on peak rate
 enforcement". *Performance Evaluation*. No 16. 1992. pp 35-48.
- [11] SGS-Thomson Microelectronics. "STI3500 MPEG-2/CCIR601 Video Decoder,
 Preliminary Data Sheet". April 1994.
- [12] Civanlar, M. R. and Cash, G. L. "A practical system for MPEG-2-based video-on-
 demand over ATM packet networks and the WWW". *Image Communication* No
 8. 1996. pp 221-227.
- [13] Guha, A., Pavan, A., Liu, J., Rastogi, A. and Steeves, T. "Supporting Real-Time
 and Multimedia Applications on the Mercuri Testbed". *IEEE Journal on Selected
 Areas in Communications*. Vol 13. No 4. May 1995. pp 749-763.
- [14] Pancha, P. and El Zarki, M. "Bandwidth allocation schemes for Variable-bit-rate
 MPEG sources in ATM networks". *IEEE Transactions on Circuits and Systems
 for Video Technology*. Vol 3. No 3. June 1993. pp 190-193.
- [15] Heyman, D. P. and Lakshman, T.V. "Source Models for VBR Broadcast-Video
 Traffic". *IEEE/ACM Transactions on Networking*. Vol 4. No 1 February 1996.
 pp 40-48.
- [16] Marafih, N. M., Zhang, Y-Q and Pickholtz, R.L. "Modelling and Queueing
 Analysis of Variable-Bit Rate Coded Video Sources in ATM Networks" *IEEE*

Transactions on Circuits and Systems for Video Technology. Vol 4. No 2. April 1994. pp 121-128.

- [17] Izquierdo, M. R. and Reeves, D. R. "Statistical Characterization of MPEG VBR Video at the Slice layer" *Proceedings of SPIE Multimedia Computing and Networking 1995*. San Jose, CA. February 1995. pp 268-279.
- [18] Reinginger, D., Melamed, B., and Raychaudhuri, D. "Variable bitrate MPEG video characteristics, modeling and multiplexing". *Proceedings of 14th International Teletraffic Congress*. Antibes Juan-les-Pins, France. June 6-10. 1994.
- [19] Kanakia, H., Mishra, P. P., and Reibman, A. R. "An Adaptive Congestion Control Scheme for Real Time Packet Video Transport". *IEEE/ACM Transactions on Networking*. Vol 3. No 6. December 1995. pp 671-682.
- [20] McDaysan, D. E., and Spohn, D. L. "ATM: theory and practice". MacGraw-Hill, Inc. NJ. 1994.
- [21] Cole, R. G., Shur, D. H., Villamizar, C. "IP over ATM: A framework document". Internet Engineering Task Force. Internet Draft. October 2, 1995.
- [22] Heinanen, J. "Multiprotocol encapsulation over ATM Adaptation Layer 5". Request for Comments (Proposed Standard) RFC 1483. Internet Engineering Task Force. July 1993.
- [23] ATM Forum. "ATM User Network Interface specification — Version 3.0". Prentice-Hall, Inc. NJ. 1993.

- [24] ITU-T. "Traffic control and congestion control in B-ISDN" ITU-T Recommendation I.371. 1993
- [25] OCRInet Worldwide Web site. <http://www.ocrinet.ca/ocrinet.html>. August 1996.
- [26] Comer, D. E., "Internetworking with TCP/IP Volume 1 Principles, protocols, and architecture". 2 Ed. Prentice-Hall, Inc. NJ 1991
- [27] Gagnon, G. e-mail communication June 15, 1995
- [28] Blair, G., Campbell, A., Coulson, G., Garcia, F., Hutchinson, D. Scott, A., and Shepherd, D. "A Network Interface Unit to Support Continuous Media" *IEEE Journal on Selected Areas in Communications*. Vol 11. No 2. February 1993. pp 264-275.
- [29] Schulzrinne, H., Casner, S., Frederick, R., Jacobson, V. "RTP: A transport protocol for real-time applications". Network Working Group. Audio-Video Transport Working Group. RFC 1889 January 1996.
- [30] Ramakrishnan, K. K. "Performance Considerations in Designing Network Interfaces". *IEEE Journal on Selected Areas in Communications* Vol 11 No 2 February 1993. pp 203-219.
- [31] Cooper, E., Menzileioglu, O., Sansom, R., and Bitz, F. "Host interface design for ATM LAN's". Proceedings 16th Conference on Local Computer Network. Minneapolis, MN. October 1991.

- [32] Davie, B. S. "A host-network interface for ATM". Proceedings Symposium Communication Architecture and Protocols. September 1990. Vol 21. No 4.
- [33] Kanakia, H., and Cheriton, D. "The VMP network adapter board (NAB): High performance network communication for multiprocessors". Proceedings Symposium Communication Architecture and Protocols. September 1988. Vol 18. No 4.
- [34] Lumley, J. "A high-throughput network interface to a RISC workstation". Proceedings IEEE Workshop Architecture. and Implementation of High-Performance Communication Subsystems. Tucson, AZ. February 17-19, 1992.
- [35] Traw, C. B. S., and Smith, J. M. "A high performance host interface for ATM networks". Proceedings Symposium Communication Architecture and Protocols. September 1991. Vol 21. No 4.
- [36] Clark, D., Jacobson, V., Romky, J., and Salwen, H. "An analysis of TCP processing overhead". *IEEE Communications Magazine*. Vol 27. No 6. June 1989. pp 23-29.
- [37] Jacobson, V., Braden, R., and Borman, D. "TCP extensions for high performance". Network Working Group - Request for comments (RFC) 1323, May 1992.
- [38] AX/4000 Series ATM Generator/Analyser Operating Manual. Software Version 1.2. Adtech Inc. Honolulu, HI. February 1996.

- [39] Boyer, P., Serval, M., and Guillemin, F. "The SPACER-CONTROLLER: An efficient UPC/NPC for ATM networks". Proceedings of the 14th International Switching Symposium. Yokohama, Japan. 1992.
- [40] Danthine, O. and Boyer, P. "Benefits of a spacer/controller in an ATM WAN: preliminary traffic measurements". EXPLOIT Traffic Workshop. Basel. September 1994.
- [41] Hui, J. Y., Karazan, E., Li, J. and Zhang, J. "Client-Server Synchronisation and Buffering for Variable Rate Multimedia Retrievals". *IEEE Journal on Selected Areas in Communications*. Vol 14. No 1. January 1996. pp 226-237.
- [42] FORE Systems, Inc. ForeRunner™ FORE ASX-200 ATM Switch User's Manual. December 1993.

APPENDIX A MPEG-2 VIDEO STREAM STATISTICS

Figure A.1 Ayers CBR 3 Mb/s

Figure A.2 Ayers CBR 4 Mb/s

Figure A.3 Ayers VBR $Q=9$

Figure A.4 Ayers VBR $Q=12$

Figure A.5 Ayers VBR $Q=18$

Figure A.6 Martin CBR 3 Mb/s

Figure A.7 Martin CBR 4 Mb/s

Figure A.8 Martin VBR $Q=9$

Figure A.9 Martin VBR $Q=12$

Figure A.10 Martin VBR $Q=18$

Figure A.11 Pirates CBR 3 Mb/s

Figure A.12 Pirates CBR 4 Mb/s

Figure A.13 Pirates VBR $Q=9$

Figure A.14 Pirates VBR $Q=12$

Figure A.15 Pirates VBR $Q=18$

END

1 8 0 6 9 7

FIN