INFORMATION TO USERS

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

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

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

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

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

UMI

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

A Multiple Access Scheme and An ABR Rate Control Scheme for Satellite ATM Networks

by

Eko Adi Wibowo, B.Sc.

University of New Brunswick

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

Master of Engineering

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

Carleton University
Ottawa, Ontario
September, 1998
© Eko Adi Wibowo 1998

National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

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

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

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

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

0-612-36900-5

Canada

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

**“A Multiple Access Scheme and An ABR Rate
Control Scheme for Satellite ATM Networks”**

submitted by **Eko Adi Wibowo, B. Sc.**
in partial fulfilment of the requirements for
the degree of Master of Engineering

Chair, Department of Systems and Computer Engineering

Rafik A. Goubran, Professor

Thesis Supervisor

Ioannis Lambadaris, Assistant Professor

Carleton University
September, 1998

Abstract

A multiple access scheme for a general multimedia satellite network is proposed. The access scheme utilizes four request handling mechanisms (strategies), each directed toward services having similar requirements. With a flexible association of request strategies to service classes, the multiple access scheme can be suited to the five ATM classes of service, making it an ideal scheme for integrating ATM technology into satellite networks. Performance evaluation using discrete-event simulations shows the excellent performance (throughput and robustness) and features (access fairness and traffic isolation) of the scheme.

An integrated multiple access and rate control scheme for ABR traffic is necessary to achieve efficient bandwidth utilization. This integration requires modifications to the scheduling of ABR transmission and to the switch algorithm. Modifications to integrate the well-known ERICA switch algorithm with the access scheme is described for the case of onboard processing satellites. Efficient uplink utilization is indeed achieved and confirmed through discrete-event simulations. Nevertheless, performance improvement offered by a rate control scheme developed with special consideration to the strengths and limitations of a satellite can be substantial. A new rate control scheme based on the broadcast of congestion information is proposed. The new scheme is shown to offer better performance in many aspects, including bandwidth utilization, fairness, and maximum buffer occupancy.

Acknowledgements

The development of the new ECFDAMA protocol was initiated by Spar Aerospace Ltd. under the Advanced Satcom program, funded by the Canadian Space Agency. The author acknowledges the invaluable support, guidance and encouragement from thesis supervisor Dr. Ioannis Lambadaris, Dr. Mihail Devetsikiotis and internship supervisor Dr. Augustin Iuoras of Spar Aerospace Ltd. The author also recognizes the helpful encouragement from colleagues at Spar Aerospace Ltd., among others, Steve Hatzigeorgiou, Colin Black, Peter Takats and Rocco Digirolamo.

Table of Contents

Abstract.....	iii
Acknowledgements	iv
Table of Contents.....	v
List of Figures	ix
List of Tables.....	xiv
List of Acronyms.....	xv
Introduction	1
1.1 Background.....	1
1.2 Motivation	3
1.3 Thesis Contributions.....	4
1.4 Organization of Thesis.....	5
Satellite ATM Networks	8
2.1 Primary Access Mechanism	11
2.2 Frame Structure	13
2.3 Multiple Access Mechanism	15
Review of Multiple Access Schemes	17
3.1 Fixed Assignment Scheme	17
3.2 Random Assignment Scheme.....	18
3.3 Demand Assignment Scheme.....	20
3.4 Combined Fixed-Demand Assignment Scheme	23

3.5 Combined Random-Demand Assignment Scheme	24
3.6 Combined Free-Demand Assignment Scheme	30
3.7 Summary	32
Enhanced CFDMA for Satellite ATM Networks	34
4.1 The Frame Structure	35
4.2 Request/Assignment Strategies	36
4.2.1 Reserved Capacity (RC) Strategy	37
4.2.2 Rate-Based Dynamic Capacity (RBDC) Strategy	38
4.2.3 Volume-Based Dynamic Capacity (VBDC) Strategy	39
4.2.4 Free Capacity Assignment Strategy (FCA)	40
4.3 Access Protocol Highlights	40
4.3.1 Light-Weight Access Protocol	40
4.3.2 In-Band Request Strategy	41
4.3.3 Out-of-Band Request Strategy	42
4.3.4 Access Fairness through Round-Robin Scheduling	43
4.3.5 Selective Free Capacity Assignment Strategy	44
4.3.6 Absolute Assignment Signaling Strategy	45
4.3.7 Flexible Association of Service Classes to Request/Assignment Strategies ..	52
4.4 Terminal Uplink Access Process	54
4.4.1 Terminal Internal Structure	55
4.4.2 Uplink Access Process	58
4.4.2.1 Request Calculation Method	60
4.4.2.2 Transmission Scheduling Strategy	66
4.5 Example of Data Transfer Operations	72
4.6 Satellite Model	74

Performance Evaluation of ECFDAMA	76
5.1 Simulation Model	76
5.2 Service Association and Traffic Mixing in Ground Terminals	82
5.3 Simulation Test Cases	84
5.4 Simulation Results	88
5.5 Summary	109
Review of Congestion Control Schemes	110
6.1 Congestion Control Schemes for Non-ATM Networks	112
6.2 Congestion Control Schemes for ATM Networks	118
6.2.1 The Rate-Based Congestion Control Framework for ATM Networks	119
6.2.2 Examples of Rate-Based Schemes	123
6.2.2.1 Congestion Control with Explicit Rate Indication (CCERI)	123
6.2.2.2 Explicit Rate Indication for Congestion Avoidance (ERICA)	125
6.2.2.3 Max-Min Rate Control Algorithm (MMRCA)	126
6.2.2.4 Dynamic Max Rate Control Algorithm (DMRCA)	128
6.2.2.5 Fast Max-Min Rate Allocation (FMMRA)	129
6.3 Congestion Control Scheme for Satellite ATM Networks	131
Integration of Multiple Access and Rate Control Schemes for ABR Service	133
7.1 Transmission Scheduling of ABR Traffic in Ground Terminals	135
7.1.1 Capacity Request Calculation Algorithm for ABR Traffic	137
7.1.2 Capacity Assignment Decomposition	139
7.1.3 Transmission of ABR Traffic	140
7.2 The Modified ERICA Switch Algorithm	142
7.3 Broadcast Rate Control Allocation for Onboard Processing Satellites	143
7.3.1 The Port Algorithm	145

7.3.2 The Onboard Network Controller (OBNC) Algorithm	149
7.3.2.1 The Consolidation Process	149
7.3.2.2 The OBNC Algorithm for Input Ports	150
7.3.2.3 The OBNC Algorithm for Output Ports	153
7.3.2.4 The Final Decision and Broadcast Process.....	157
7.3.2.5 Summary of the OBNC Congestion Control Algorithm	158
7.3.3 The Terminal Algorithm for Rate Adjustment	162
Performance Evaluation of An Integrated MAC and ABR Rate Control Scheme .	167
8.1 Persistent Source Simulations	169
8.1.1 Network Loading Configurations.....	170
8.1.2 Simulation Test Cases.....	173
8.1.3 Simulation Results	179
8.2 Bursty Source Simulations.....	199
8.2.1 Simulation Test Cases.....	199
8.2.2 Simulation Results	201
8.3 Summary.....	206
Conclusions and Future Research.....	207
9.1 Summary and Conclusions.....	207
9.2 Future Research Directions.....	211
9.2.1 Call Admission Control.....	211
9.2.2 Scheduling.....	211
9.2.3 Traffic Modeling.....	212
9.2.4 Rate-Based Control	212
References	214

List of Figures

Figure 1. The Typical Satellite ATM Network Architecture.....	10
Figure 2. The MF-TDMA Frame Structure	15
Figure 3. The Frame Structure	36
Figure 4. The Request Minislot Format	42
Figure 5. Consecutive Time Slot Assignment.....	47
Figure 6. The Assignment Message Format	47
Figure 7. The Assignment Process First Pass	50
Figure 8. The Assignment Process Second Pass.....	51
Figure 9. The Assignment Process Second Pass (Cont'd)	52
Figure 10. A Typical Block Diagram of an Earth Terminal	56
Figure 11. A Summary of Terminal's Tasks Performed at Every Frame.....	57
Figure 12. The Terminal Uplink Access Process Algorithm	59
Figure 13. The Pseudo-code for Time Slot Assignment Distribution Algorithm	69
Figure 14. The Pseudo-code for Time Slot Assignment Distribution Algorithm (cont'd)	70
Figure 15. The Pseudocode for Cell Transmission Scheduling Algorithm.....	71
Figure 16. An Example of Data Transfer Operation using the RBDC Strategy	73
Figure 17. A Typical Block Diagram of an Onboard Processing Satellite	74
Figure 18. Satellite Network Configuration for Simulations	77
Figure 19. The IDC of the 2-state MMPP Model.....	79

Figure 20. The IDC of the Self-Similar (WWW) Model.....	82
Figure 21. Test Cases 1 and 2 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	90
Figure 22. Test Cases 3 and 4 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	91
Figure 23. Test Case 5 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	94
Figure 24. Test Case 6 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	95
Figure 25. Test Case 7 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	96
Figure 26. Test Case 8 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	98
Figure 27. Test Case 9 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	99
Figure 28. Test Case 10 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	101

Figure 29. Test Case 11 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	102
Figure 30. Test Case 12 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	105
Figure 31. Test Case 13 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	106
Figure 32. Test Case 14 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	107
Figure 33. Test Case 15 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay	108
Figure 34. Queueing Structure in Ground Terminals.....	136
Figure 35. Flow Charts of Terminal's Capacity Request Calculation Algorithm.....	139
Figure 36. Flow Charts of Terminal Algorithm for Transmission of ABR Traffic	141
Figure 37. Flow Charts for BRCA's Port Algorithm	147
Figure 38. Flow Charts for BRCA's Port Algorithm (cont'd).....	148
Figure 39. The Consolidation Process Algorithm.....	150
Figure 40. Flow Charts of OBNC Algorithm for Input Ports	153
Figure 41. Flow Charts of OBNC Algorithm for Output Ports.....	159
Figure 42. Flow Charts of OBNC Algorithm for Output Ports (cont'd).....	160
Figure 43. Flow Charts of OBNC Final Decision Algorithm	161

Figure 44. Summary of the OBNC Rate Control Algorithm	162
Figure 45. Flow Charts of Terminal Algorithm for Rate Adjustment	166
Figure 46. Configuration for Balanced Traffic Loading.....	171
Figure 47. Configuration for Balanced Traffic Loading with Different ABR Source Rates.....	171
Figure 48. Configuration for Non-Balanced Traffic Loading.....	172
Figure 49. Configuration for Concentration Traffic Loading	173
Figure 50. Persistent Source Test Case 1 Results (BRCA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size.....	182
Figure 51. Persistent Source Test Case 1 Results (ERICA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size.....	183
Figure 52. Persistent Source Test Case 2 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size.....	184
Figure 53. Persistent Source Test Case 2 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size.....	185
Figure 54. Persistent Source Test Case 3 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Uplink Load (per traffic classes); (d) Downlink ABR Queue Size	186
Figure 55. Persistent Source Test Case 3 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Uplink Load (per traffic classes); (d) Downlink ABR Queue Size	187
Figure 56. Persistent Source Test Case 4 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size.....	190
Figure 57. Persistent Source Test Case 4 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size.....	191

Figure 58. Persistent Source Test Case 5 Results (BRCA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size.....	192
Figure 59. Persistent Source Test Case 5 Results (ERICA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size.....	193
Figure 60. Persistent Source Test Case 6 Results (BRCA) : (a), (b) ACRs for Downlink 0; (c) Downlink Load; (d) Downlink ABR Queue Size.....	194
Figure 61. Persistent Source Test Case 6 Results (ERICA) : (a), (b) ACRs for Downlink 0; (c) Downlink Load; (d) Downlink ABR Queue Size.....	195
Figure 62. Persistent Source Test Case 7 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size.....	197
Figure 63. Persistent Source Test Case 7 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size.....	198
Figure 64. Bursty Source Test Case 1 Results (No Rate-Based Control) : (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay	202
Figure 65. Bursty Source Test Case 2 Results (No Rate-Based Control) : (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay	202
Figure 66. Bursty Source Test Case 1 Results : (BRCA) (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay; (ERICA) (c) CDFs of Downlink ABR Queue Size; (d) End-to-End Mean Delay	204
Figure 67. Bursty Source Test Case 2 Results : (BRCA) (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay; (ERICA) (c) CDFs of Downlink ABR Queue Size; (d) End-to-End Mean Delay	205

List of Tables

Table 1. The Request Type Encoding.....	42
Table 2. Description of Variables in the Assignment Process Flow Charts.....	49
Table 3. Several Possible Association of Service Classes to Request/Assignment Strategies	54
Table 4. The 2-MMPP Model Parameter Values	78
Table 5. The WWW Model Parameter Values	81
Table 6. The Mapping of ATM Service Classes to Request Strategies	83
Table 7. Various Terminal Types and Configurations for ECFDAMA Performance Evaluation	85
Table 8. Simulation Test Cases for The ECFDAMA Performance Evaluation.....	87
Table 9. Persistent Source Test Cases for ERICA and BRCA Performance Evaluations.....	177
Table 10. ERICA Parameters for Persistent Source Simulations	178
Table 11. BRCA Parameters for Persistent Source Simulations	178
Table 12. Common ABR Rate Parameters for Persistent Source Simulations.....	179
Table 13. Bursty Source Test Cases for ERICA and BRCA Performance Evaluations	200

List of Acronyms

ABR	Available Bit Rate
ACR	Allowed Cell Rate
A-PRMA	Aggressive PRMA
ATM	Asynchronous Transfer Mode
BRCA	Broadcast Rate Control Allocation
BIS	Busy/Idle Signal
BTP	Burst Time Plan
CBR	Constant Bit Rate
CCR	Current Cell Rate
CDMA	Code Division Multiple Access
CDT	Cell Delay Tolerance
CDV	Cell Delay Variation
CDVT	CDV Tolerance
CFDAMA	Combined Free Demand Assignment Multiple Access
CLR	Cell Loss Ratio
C-PRMA	Centralized PRMA
CREIR	Combined Random Explicit then Implicit Reservation
CTD	Cell Transfer Delay
ECFDAMA	Enhanced CFDAMA
ER	Explicit Rate

ERICA	Explicit Rate Indication for Congestion Avoidance
FCA	Free Capacity Assignment
FCFS	First Come First Serve
FDMA	Frequency Division Multiple Access
FRB	Frame Reference Burst
HDTV	High Definition Television
IBR	In-Band Request
ICR	Initial Cell Rate
LAN	Local Area Network
MAC	Multiple Access Control
MASCARA	Mobile Access Scheme Based on Contention and Reservation for ATM
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
MC-CDMA	Multi Carrier CDMA
MF-TDMA	Multi Frequency TDMA
MMPP	Markov-Modulated Poisson Process
nrt-VBR	Non Real-Time VBR
OBNC	Onboard Network Controller
OBR	Out-of-Band Request
PCR	Peak Cell Rate
PRADOS	Prioritized Regulated Allocation Delay-Oriented Scheduling
PRMA	Packet Reservation Multiple Access
PRMA/DA	PRMA / Dynamic Allocation
PRMA-DSV	PRMA Data Steal into Voice
QoS	Quality of Service

R/A	Request/Assignment
R-Aloha	Reservation Aloha
RBDC	Rate-Based Dynamic Capacity
RC	Reserved Capacity
RM	Resource Management
RMA	Reservation Multiple Access
rt-VBR	Real-Time VBR
S-Aloha	Slotted Aloha
SCR	Sustainable Cell Rate
SREJ	Selective Reject
TDD	Time Division Duplex
TDMA	Time Division Multiple Access
UBR	Unspecified Bit Rate
VBDC	Volume-Based Dynamic Capacity
VBDC-HP	High Priority VBDC
VBDC-LP	Low Priority VBDC
VBR	Variable Bit Rate
VOD	Video On Demand
VS/VD	Virtual Source / Virtual Destination
WWW	World Wide Web

Chapter 1

Introduction

1.1 Background

Modern communication networks are now required to support a diversity of applications and provide numerous types of services. Applications such as video conferencing, computer telephony are already common, while applications such as HDTV, VOD and many more are a question of the near future. Likewise, services such as LAN-to-LAN interconnection and wireless personal communications (both mobile and fixed services) are already widespread, with many new narrow-band and wide-band services likely to appear in the future. These bring about demand for advanced techniques for handling the generated traffic, which has been characterized as very bursty and unpredictable. All these trends speed the adoption of ATM technology in the design of communications networks of the future.

ATM has the potential to meet the demands generated by future applications and services. The technology is designed with various types of traffic in mind and is continuously evolving, as researchers gain more in-depth knowledge on how to handle the ATM traffic. The ATM Forum [1] defines five classes of service, each having its own traffic characteristics, depicted by its traffic descriptors and Quality of Service (QoS) requirements.

The first class of service is the Constant Bit Rate (CBR) service. This service has only one traffic descriptor, the Peak Cell Rate (PCR), and is usually used for connections requiring a fixed amount of bandwidth available continuously during the lifetime of the connection. Hence, this service can be used to emulate circuit switching. CBR connections have stringent requirements on cell delay, cell delay variation and cell loss ratio. Typical applications which can benefit from this service are video conferencing, telephony and video-on-demand.

The second class of service is the Real-Time Variable Bit Rate (rt-VBR) service. This service has three traffic descriptors, namely Peak Cell Rate (PCR), Sustainable Cell Rate (SCR) and Maximum Burst Size (MBS). It is intended for connections with tight delay, delay variation and low cell loss ratio requirements, but which do not want a static amount of bandwidth to be continuously available as in the CBR service. As such, rt-VBR service may support statistical multiplexing to efficiently utilize network resources. Any real time applications which can tolerate small cell loss may take advantage of this service.

The next class of service is the Non-Real-Time Variable Bit Rate (nrt-VBR) service. Similar to the rt-VBR service, it also has three traffic descriptors, i.e., PCR, SCR and MBS. There is no requirement, however, on delay and delay variation bounds. Response time critical applications, such as airline reservations and banking transactions, may utilize this service.

The fourth class of service is the Unspecified Bit Rate (UBR) service. This service is intended for connections with no delay, delay variation and cell loss ratio constraints. It has only one traffic descriptor, the Peak Cell Rate (PCR). Applications such as email, facsimile and file transfer may benefit from this service.

The final class of service is the Available Bit Rate (ABR) service. This is a relatively new service and is mainly intended for connections that require stringent cell loss ratio requirement. ABR service has two traffic descriptions, namely Minimum Cell Rate (MCR) and Peak Cell Rate (PCR). Cell delay and delay variation requirements are not specified under this service. In order to maintain a low cell loss ratio, several feedback flow control mechanisms are also specified. UBR applications which require a low cell loss ratio may take advantage of this service.

1.2 Motivation

Extending ATM services via satellite has been a target of extensive studies in recent years [2, 3 4]. The objective of these studies was to seamlessly incorporate ATM technology into the satellite environment. The rationale for this attempt is not difficult to grasp. A satellite network is an excellent candidate for global communication infrastructure due to its wide geographical coverage. Further, as opposed to developing terrestrial networks using fibre optics, a satellite network can be developed relatively quickly and cheaply. In addition, such network can be configured flexibly, as additional users can be accommodated without disrupting network operations.

There are, however, some issues that need to be resolved before seamless integration of ATM technology into satellite networks can be achieved. The root of these issues is in how to maintain the quality of service that exists in terrestrial ATM. Specifically, the core issue is in how to accommodate all five ATM classes of services into such networks while complying to the quality of service associated with each class.

In our perspective, one answer to this issue lies in the multiple access protocol. Recent studies have pointed out its importance in the integration of ATM technology into

the wireless medium [5,6]. The sought after protocol should be capable of ensuring bandwidth for real-time types of applications. It should also be able to provide bandwidth on-demand to allow efficient multiplexing of variable-rate connections. Further, such protocol should have a built-in mechanism for handling priorities that exist among the five classes of service. Finally, an ideal protocol should make the best use of valuable network resources under any or most circumstances.

Further, we believe that the rate-based congestion control scheme currently adopted by the ATM Forum should not be ported directly, without any modification, into the satellite environment. The reason is that it has been conceived for terrestrial network environment where transmission capacity is abundant, frame structure is much less rigid and multiple access protocol plays a less important role in affecting network performance. As a result, when applied unmodified to the satellite environment, such scheme may perform poorly. We believe that an ideal congestion control scheme should work in complement to the multiple access protocol in order to maintain optimal use of bandwidth. We also believe that it should take advantage of satellite broadcast capability.

1.3 Thesis Contributions

1. A demand-based multiple access scheme (ECFDAMA) capable of handling different priorities and quality of service requirements defined for ATM is proposed and the performance is shown to be very satisfactory.
2. A simple technique to integrate the ECFDAMA scheme to the adopted ABR rate-based control scheme is proposed, and the performance of the integrated ECFDAMA and ERICA scheme is evaluated.

3. A rate control scheme (BRCA) that takes into considerations the strengths and limitations of a satellite is proposed and its superior performance demonstrated.

1.4 Organization of Thesis

1. Introduction

This chapter provides a brief background on the five service categories supported by ATM, describes the motivation behind extending ATM services via satellite and the accompanying issues that need to be resolved, and explains how the thesis contributes to reaching the goal of seamless integration of ATM technology into the satellite environment more closely.

2. Satellite ATM Networks

This chapter describes a possible architecture of future satellite ATM networks, which will be used as a reference architecture for the thesis. The chosen frame structure and primary access mechanism for the proposed demand-based multiple access scheme is described.

3. Review of Multiple Access Schemes

This chapter introduces the six general categories of multiple access protocols and reviews several access protocols from each category. Assessment will be made with respect to the suitability of each of the six categories for integrating ATM technology into satellite networks. Assessment criteria include performance, ability to support heterogeneous traffic and to provide quality-of-service guarantee and complexity.

4. Enhanced CFDAMA for Satellite ATM Networks

In this chapter the Enhanced Combined Free Demand Assignment Multiple Access (ECFDAMA) scheme, its three request handling mechanisms and various features and refinements with respect to the original CFDAMA scheme, is described in detail. Further, since each earth terminal has a significant role in achieving efficient uplink utilization, the terminal uplink access process is also explained.

5. Performance Evaluation of ECFDAMA

In this chapter, the performance of the ECFDAMA scheme is evaluated using discrete-event simulations. The simulation model and the miscellaneous test cases employed to assess various performance aspects of the ECFDAMA scheme are described. Simulation results are presented in terms of several important performance measures, namely throughput, cell transfer delay and cell transfer delay variation.

6. Review of Congestion Control Schemes

This chapter reviews many burst-level congestion control techniques which have been developed for both terrestrial and satellite networks in the past and which may have potential for implementation in satellite environment.

7. Integration of Multiple Access and Rate Control Schemes for ABR Service

In this chapter, the modifications required to integrate the currently adopted rate-based control framework with the ECFDAMA protocol will be explained. The well-known ERICA switch algorithm is used to illustrate the required modifications to the many developed switch algorithms. Finally, a new rate control technique for ABR traffic based on the broadcast of congestion information is described.

8. Performance Evaluation of an Integrated MAC and ABR Rate Control Scheme

In this chapter, the performance of the integrated multiple access and rate control scheme is evaluated through discrete-event simulations. The modified ERICA scheme is compared to the Broadcast Rate Control Allocation (BRCA) scheme, and performance improvement offered by the BRCA scheme is confirmed.

9. Conclusions and Future Research

This chapter presents the summary of the thesis and describes some unresolved problems for further studies.

Chapter 2

Satellite ATM Networks

The high potential for satellite to be a significant element in global information infrastructure is well recognized. For decades, satellites have been used as one of the key technologies in the area of telecommunication. They enable people otherwise separated over continents to communicate easily and cost effectively. They are also well known for their wide coverage, making them the preferred choice to distribute information to a large number of people. Not until recently, however, has the maximum potential of satellites been realized.

ATM technology has been present for quite some time. The technology enables seamless integration of voice, video, and data services over a network. Until recently, its use has only been exploited on terrestrial or wired network. Introducing the ATM technology in the wireless or satellite networks is an active research area. Much has still to be investigated before the technology can be implemented efficiently on the wireless medium. Many promising results have recently been published, in the form of new concepts and technologies [7]. It appears that, with the current trend and pace of advancement, the dream of ubiquitous multimedia communication will become a reality in a not too distant future.

Current research results [2, 9] have shown that the incorporation of ATM technology into satellite networks will lead to significant changes in the way the

communications over satellite are performed. A conventional satellite carries simple onboard functions. Little processing is performed upon reception of payloads, before they are forwarded to another satellite or to ground terminals. Hence, traditional satellite can be alternatively called bent-pipe satellite, as its function is to funnel signals from one destination to another.

Future satellite networks will have to be much more sophisticated in order to support the various types of applications and services. A possible architecture for satellite ATM networks, which will be used as a reference architecture for the thesis, is depicted in Figure 1 [2, 3]. New functionality added to the satellite (onboard processor) may include scheduling, queueing, switching and flow control [3]. The uplink scheduler, which may reside in the front-end of the onboard processor and is in charge of immediate allocation of bandwidth to ground terminals, should be very efficient and effective in handling priorities and variation in network traffic. There may also be a need for an onboard network controller, which will be responsible for managing satellite network capacity. Alternatively, all scheduler and network controller functions can be located on-ground (for example, within the ground controller) in order to reduce hardware complexity onboard the satellite. This will decrease the responsiveness of the system to the variation in network traffic, however, since additional delay is incurred for terminal messages to reach the ground controller and for the corresponding controller's responses to reach the satellite. In any case, queueing onboard the satellite is inevitable, since uplink payloads will be sharing the same switching and downlink transmission resources. Flow control-related functions (such as measurements) may also be located onboard the satellite in order to reduce the negative impact of long propagation delay on congestion management. Additional complexity is also unavoidable in the ground terminals. This is partly due to the

additional processing of ATM cells and partly due to the more sophisticated multiple access scheme. Finally, the master control station is responsible for the general operations of the satellite network, including satellite tracking and position control, network management, etc.

The need for a satellite ATM layer can be rationalized as follows [11]. In preparing ATM cells for uplink transmission, additional information bits for signaling, routing and error control, may need to be incorporated into the ATM cells. Further, in order to satisfy the QoS requirements of each service class, it is crucial that close interaction is established between the ATM layer and the multiple access protocol. In essence, the access protocol should also be able to support the QoS requirements of each service class, which implies that it should become part of a standard for ATM transport over satellite. Establishing close interaction between the ATM layer and multiple access protocol and performing additional processing on ATM cells are functions of the satellite ATM layer.

Figure 1. The Typical Satellite ATM Network Architecture

2.1 Primary Access Mechanism

The primary access method for satellite ATM network should be selected such that it allows for the most efficient use of bandwidth and most cost effective implementation, and is capable to support variable transmission rate. There are three basic primary access methods for wireless communications, namely frequency division multiple access (FDMA), time division multiple access (TDMA) and code division multiple access (CDMA). Each of these basic methods has its advantages and disadvantages, and individually, none of them satisfies the requirements of an ATM satellite network.

In FDMA, transmission capacity is divided into usually equal-size frequency bands or subchannels. A ground terminal has its own frequency band(s) to transmit its payloads, hence collision can be avoided. Another advantage of FDMA is inexpensive implementation, due to its relative simplicity compared to TDMA and CDMA technologies. Nevertheless, FDMA also has several disadvantages. First, the number of ground terminals is limited to the number of subchannels. Second, sufficient guard bands are required to avoid interference between subchannels, which translates into some loss of efficiency in bandwidth utilization. In addition, there is little flexibility in handling traffic variability.

In TDMA, transmission capacity is divided into normally equal-size time slots. In the most simplistic case, a ground terminal is assigned one or more time slots on a fixed periodic basis. In a more advanced case, a ground terminal is allowed to vary its transmission requirement on an as needed basis, resulting in a bandwidth-on-demand scheme. Therefore, a large number of terminals can be supported by the network. Further, TDMA technology provides flexibility in handling variability of network traffic. There are, however, problems associated with it. First, higher transmission capacity results in bigger

antenna size, which finally has a negative impact on the cost of ground terminals. Second, there is a need for accurate time synchronization for ground terminals, which results in additional hardware complexity. Third, there is still a need for guard bands (in terms of time), to give sufficient rooms to tolerate variable propagation and multipath delays [12]. Nevertheless, TDMA technology is considered a good match for Integrated Service Digital Network (ISDN) [13].

In CDMA, the whole transmission capacity is available to all terminals at all time. The concept is in contrast to the previous two access mechanisms, where a ground terminal is given only a portion of the transmission capacity. In this technique, signals from one terminal are allowed to interfere with those from other terminals. Hence, multiple access capability is built-in, and, due to its random access nature, adding users to the network is a simple process. Each terminal is assigned a unique code. Coded signals can be extracted in the receiver side by using a complementary code. Each code is designed to be orthogonal to any other code, that is, each code is made to be as different as any other [14]. As a result, signal interference is minimal. Other advantages of CDMA technology include built-in security mechanism and robustness to various impairments in wireless channels or systems, such as multipath fading, frequency selectivity and synchronization errors [15]. The major problem associated with CDMA technology is its low bit rate. Additionally, the requirement for orthogonal codes for each connection places a constraint on the possible connection set-up among terminals.

As none of these three primary access mechanisms possesses all the qualities of an excellent access mechanism, researchers attempt to combine them in order to gain as many advantages and eliminate as many disadvantages as possible. Examples of such

combinations are Multi-Carrier CDMA (MC-CDMA) and Multi-Frequency TDMA (MF-TDMA).

MC-CDMA aims to increase network capacity by adding multiple carriers. The increase in network capacity is proportional to the number of additional carriers. The additional enhancement of MC-CDMA over the basic CDMA technology comes, however, at the expense of additional hardware complexity.

In MF-TDMA, transmission capacity is divided into both frequency and time slots. A ground terminal is allowed to transmit at any frequency and any time slot, but not at more than one frequency simultaneously. This imposes a limit on the rate at which the terminals can transmit (i.e., carrier rate), with positive impact on antenna size and transmission power. An additional enhancement to increase bandwidth efficiency includes the introduction of preamble-less MF-TDMA [16]. One drawback of MF-TDMA is the more sophisticated demodulation process required onboard the satellite (compared to that required by TDMA). Nevertheless, due to its advantages, MF-TDMA is chosen as the primary access mechanism for the thesis.

2.2 Frame Structure

TDMA is a frame-based scheme. This means that time slots are organized into a regular structure. Slots in a frame are grouped according to their functions. These functions relate to how data transfer process is conducted in a TDMA satellite network. Before a terminal is allowed to use any of the network resources, it has to register itself to the network. Next, it has to constantly synchronize itself to the network timing so that it does not disrupt or interfere with any other terminals. Then, if demand-based multiple access is employed, the terminal has to signal its request for uplink capacity to the process

in charge of uplink capacity allocation (the uplink scheduler). The terminal may then transmit its data once the assignment message arrives. Based on this transmission procedure, time slots can be classified into five main groups, namely log slots, acquisition slots, synchronization slots, request slots and payload slots.

Log slots are used in the registration or log-in process, and are usually randomly accessed by ground terminals (e.g. S-Aloha). Acquisition slots are used to achieve synchronization prior to any data transfer. Synchronization slots are used by terminals to maintain themselves synchronized to the network frame timing. Request slots carry capacity request information from each terminal, and consequently, are only useful for demand-based multiple access schemes. These slots may be grouped distinctly in the frame, or they may be piggy-backed into slots from other groups, such as synchronization or payload slots. Finally, payload slots are used to transfer the actual network traffic. A MF-TDMA frame structure is depicted in Figure 2 below [3].

In Figure 2 the request slots are not shown as a group. These slots are piggy-backed in both synchronization and payload slots in order to reduce the transmission overhead. Synchronization and acquisition slots can also be viewed as overhead, since they are not actually used to carry network traffic. Consequently, to minimize overhead, the number of payload slots should exceed by far the number of synchronization and acquisition slots.

Figure 2. The MF-TDMA Frame Structure

- ♦ Log slots not shown, as they are part of a separate carrier

2.3 Multiple Access Mechanism

The performance of a satellite ATM network will rely to a great extent on its multiple access mechanism. Consequently, an advanced, high performance, multiple access mechanism is critical to such network. This is due to the following reasons. First, satellite's wide coverage and generally expensive network resources lead to extensive resource sharing. Potentially a large number of terminals will share the limited transmission bandwidth. Second, the long propagation delay inherent in satellite networks would reduce the performance of otherwise excellent multiple access protocol for terrestrial networks. We believe that a suitable access protocol should avoid randomness as one of its components as much as possible. Combined with the need to handle

prioritized, bursty traffic with variable quality of service guarantee, this creates a requirement for a sophisticated, efficient, yet flexible demand-based multiple access protocol.

In the next chapter, a review of multiple access protocols will be presented. Each of these access protocols will be assessed from the point of view of its suitability for integrating ATM technology into satellite communication networks. This chapter will also serve as an introduction to the progeny of the proposed access protocol for satellite ATM networks described in Chapter 4.

Chapter 3

Review of Multiple Access Schemes

In [17] J.I. Mohamed classified multiple access protocols into six general categories, namely fixed, random, demand, combined fixed-demand, combined random-demand, and combined free-demand assignment schemes. In this chapter, we will review these six categories and, when available, provide brief descriptions of several newly reported access protocols in each category. Further, we will assess each of these categories with respect to its suitability for integrating ATM technology into satellite networks. The assessment criteria include performance, ability to support heterogeneous traffic and to provide quality-of-service guarantee, complexity, etc. This assessment is by no means exhaustive, however, as the main objective of this chapter is to introduce the plethora of ideas which have been conceived in various studies on the subject.

3.1 Fixed Assignment Scheme

In fixed assignment scheme, each earth terminal is allocated a portion of uplink capacity according to its request in the terminal connection establishment phase, and the allocated capacity stays fixed for the lifetime of the connection. This category of access scheme is the simplest among the six and, predictably, requires low hardware complexity for implementation. This scheme is especially efficient for a network with a low population

of constant-traffic terminals. Since a terminal does not have to request or compete for uplink capacity every time it needs to send traffic, fast access time (the elapsed time between the moment traffic arrives at the terminal and the moment it is sent) and low packet delay are experienced. Additional processing is needed, however, to handle a large number of terminals, for example by using the HRR scheduling as described in [18]. Further, such access mechanism is highly inefficient when terminal's traffic is bursty, since there is no mechanism for a terminal to vary its capacity allocation following its active (traffic burst) and silence periods. Therefore, the benefit of inter-terminal multiplexing of variable bit rate traffic cannot be realized. For all the above reasons, this category of access scheme is not suitable for satellite ATM networks.

3.2 Random Assignment Scheme

In random assignment scheme, an earth terminal may utilize the transmission channel as it wishes. There is, therefore, a random probability of collision when two or more terminals attempt to transmit in overlapping periods. When packets do collide, each corresponding terminal may need to reschedule transmission of its collided packet.

This category of multiple access scheme has its root in the Aloha system described in [19], which was originally conceived for radio communication. The strengths of access protocols belonging to this category include fast access time and low packet delay at low traffic load, built-in support for bursty traffic (each terminal may attempt to transmit each packet as it arrives) and large terminal population, and a relatively simple access algorithm (medium complexity) in ground terminals. Transfer delay performance, however, degrades quickly as traffic load increases due to an increase in collision probability. The long delay in transmission acknowledgement (due to the inherently long propagation delay of a

satellite network) and the additional traffic from retransmission further degrade the performance. In fact, it has been shown in various textbooks that the maximum stable load in an Aloha system is approximately 18 %.

The Aloha system has been improved significantly since its introduction. In Slotted Aloha (S-Aloha), the maximum stable load is increased to 36 %, by restricting the start of transmission of each new packet to the beginning of a time slot, i.e., a terminal has to wait until the beginning of the next time slot before it can attempt to transmit its newly arrived packet. The Selective Reject (SREJ) Aloha (referenced in [20]) is capable of reaching a maximum stable load similar to the Slotted Aloha without the need of accurate timing synchronization between terminals as in the Slotted Aloha, by subpacketizing traffic messages and re-transmitting only the collided parts of messages. A further enhancement of the SREJ Aloha is the SREJ Aloha / FCFS (also referenced in [20]), which is capable of reaching a maximum stable load of 70 % by using the time information from the order of the arrival of the subpackets (which is broadcasted within the subpacket acknowledgements) to organize the retransmission of collided packets to avoid further collisions. Various schemes enhancing the SREJ Aloha or the SREJ Aloha / FCFS system in order to support more variety of services are described in [20]. One of the suggested modifications to integrate voice and data services in a VSAT network is the addition of fields in each subpacket header to indicate that no re-transmission will be attempted and that a synchronous retransmission interval, similar to that in the SREJ Aloha / FCFS scheme, is desired. Then, the first voice subpacket will be transmitted in random access mode, while subsequent subpackets will be transmitted in synchronous mode within the retransmission interval. The paper, however, does not provide details on how such scheme

would provide different quality-of-service guarantees for different classes of service in ATM.

3.3 Demand Assignment Scheme

Access protocols belonging to the category of demand assignment scheme aim to maximize uplink capacity utilization by requiring each earth terminal to reserve capacity for its traffic. The scheduler then allocates bandwidth in response to the terminal's demand. As a result, bandwidth is not wasted, as may be the case for fixed assignment scheme (by reserving more bandwidth than is needed) or random assignment scheme (by re-transmitting collided packets). Further, since the scheduler has full control over bandwidth allocation, the network can be designed to be very robust and fairness can be easily incorporated into the scheme. Finally, stable operation up to maximum load is achievable, with graceful degradation in delay performance at high load.

Demand assignment schemes do have several disadvantages. Since each terminal needs to send requests in order to obtain capacity allocation, a portion of uplink capacity has to be set aside for request signaling. Terminal's request is transmitted in a request minislot. Request minislots typically have their own identity, but the requests can also be piggy-backed in data slots. In the former case, request minislots must be allocated to earth terminals either by random access scheme or by fixed assignment. In the latter case, terminals with capacity allocation may transmit their requests along with their data.

When request minislots are allocated using random access scheme, the network should be designed such that the load corresponding to request signaling is less than 18 % of the maximum request load. When request load is low, terminals experience a small system access delay (defined as the delay experienced by a terminal to successfully transmit its request). When request minislots are fixed-assigned to earth terminals, the

maximum terminal population in the network will be limited to the number of request minislots in a frame and the interval between successive assignment of request minislots. In each of the two schemes (random access and fixed assignment), there is a trade-off between communication overhead for request area and the maximum number of terminals. Fairness may be an issue with piggy-backed request schemes, since only terminals with capacity allocation have the opportunity to request more capacity. Therefore, a demand assignment scheme should not rely solely on the piggy-backed requests.

There is also the drawback of long end-to-end packet delay associated with demand assignment schemes. End-to-end packet delay is defined as the delay experienced by a packet from the time it arrives at the source terminal to the time it is completely received (or reassembled) by the destination terminal, and includes, among others, the system access delay, packet transmission delay, queueing delay, etc. Depending on where the scheduler is located, the minimum end-to-end packet delay can be either two round-trips (with on-board scheduler) or three round-trips (with ground scheduler). For GEO satellite networks, this delay can be 500 ms or 750 ms at minimum. As a final disadvantage, demand assignment scheme requires a more complex hardware, both in the controller (scheduler) and in ground terminals. Yet, despite all these drawbacks, demand assignment scheme is considered to be suitable for wireless or satellite ATM networks, mainly because it permits optimum use of network resources, is robust and is very flexible (one may incorporate algorithms to achieve some desired performance relatively easily).

Demand assignment scheme has received considerable attention in recent years. In [21] the performance of Time Division Duplex (TDD) Aloha - Reservation MAC protocol for integrated variable bit rate video and data in wireless LAN was evaluated. Although the protocol was designed for wireless LAN, it should work for satellite networks without

modifications. In a TDD channel, uplink and downlink transmission is conducted on the same carrier. The time axis is divided into frames, a portion of which is intended for the uplink channel while the rest is used for the downlink channel. An uplink frame can be further broken down into two subframes, a control subframe and a message subframe. A control subframe consists of several acknowledgement and request minislots, while a message subframe is composed of several data slots. The protocol operation depends on the type of traffic being exchanged. For data traffic, terminals have to contend for request slots every time they need to acquire access to the network. For video traffic, however, only the first reservation is obtained through random access. Subsequent reservations can be performed by using terminals' fixed-assigned request minislot - a request minislot allocated by the network for the terminal's exclusive use.

The performance of Load-Adaptive / TDMA Systems for integrated service satellite networks was evaluated in [22]. The key concept in the paper is the incorporation of fairness into the scheduler's channel allocation algorithm. The paper describes two channel allocation algorithms, one designed to give preference to terminals generating small capacity requests (i.e., low-load terminals), while the other gives preference to high-load terminals. The paper shows that one may take advantage of the control the scheduler has in allocating bandwidth to earth terminals by incorporating fairness criteria to the scheduler's bandwidth allocation algorithm in order to create a fair demand-based multiple access protocol.

Multiservice Dynamic Reservation (MDR) TDMA was proposed for wireless ATM networks in [23, 24]. In MDR TDMA protocol, time is subdivided into frames of variable length. An uplink frame is broken down into message and request subframes. The message subframe is further broken down into fixed allocation subframe intended for CBR

traffic (including voice) and dynamic allocation subframe for VBR and data traffic. Request minislots are intended for terminals' initial entry into the system in contention mode. The access scheme is a combination of circuit mode reservation for CBR traffic and dynamic mode assignment for VBR and data traffic. Each terminal is granted a fixed number of time slots in each frame for its CBR traffic. In contrast to the fixed assignment scheme, fixed assignment of time slots in MDR TDMA stays only for the duration of the terminal's CBR traffic connection. Further, the scheduler allocates capacity for VBR traffic based on capacity requests to VBR traffic and the usage parameter control (UPC) values declared in the phase of connection establishment. Next, capacity allocation for data traffic may be performed on a simple first-come-first-serve basis or some other more advanced techniques. Finally, requests for CBR traffic rank the highest in priority and those for data traffic the lowest, that is, all CBR requests will be serviced before the scheduler services requests for VBR, and finally, for data traffic.

3.4 Combined Fixed-Demand Assignment Scheme

The long system access and end-to-end packet delays prevent the use of demand assignment scheme for real-time and interactive applications. Therefore, access protocols belonging to the category of combined fixed-demand assignment scheme use the combination of fixed and demand assignment to handle both real-time and non-real-time (jitter-tolerant) applications. These protocols, however, cannot avoid the drawbacks of fixed-assignment scheme, which include less than optimal uplink utilization and support for only a limited number of terminals.

The combined fixed-demand assignment scheme is realized using TDMA. An uplink frame is divided into a control subframe, a fixed assignment subframe and a demand

assignment subframe. The control subframe consists of acknowledgement and request minislots. A terminal has a fixed-assigned slot in the fixed assignment subframe, and possibly one or more slots in the demand assignment subframe depending on the traffic load in the terminal. The result is that end-to-end packet delay performance is increased at low load, since most traffic is transmitted using fixed-assigned slots. This compares to the minimum delay of two or three round-trips with demand assignment scheme, even at low load. An example of such scheme is the Reservation Multiple Access (RMA) scheme described in [25]. In RMA, each terminal has a fixed-assigned request minislot and a fixed-assigned data slot in every frame. Therefore, in addition to an improved end-to-end packet delay performance, fast system access delay is achieved.

3.5 Combined Random-Demand Assignment Scheme

The objective of access protocols in the category of combined random-demand assignment scheme is similar to that in the previous category, that is to gain all of the advantages of demand assignment scheme (such as efficiency, flexibility and robustness) while avoiding the poor delay performance of demand assignment scheme at low traffic load. In the combined random-demand assignment scheme, a terminal usually sends the first packet from a data message using an unreserved data slot in random access mode. Subsequent packets from the same message, however, will be transmitted in slots reserved for the terminal. Hence, since only the first packet of a data message may experience collision, the probability of collision (and re-transmission) is reduced. Next, since an inactive terminal can send its request almost immediately as it becomes active, a fast system access is achieved. Further, at low load most traffic is transmitted in randomly accessed slots, while at high load, most traffic will be sent in reserved slots. Therefore,

end-to-end delay performance should be comparable to random assignment scheme at low load, and to demand assignment scheme at high load.

Reservation-Aloha (R-Aloha) [26] is an example of combined random-demand assignment scheme, originally conceived for satellite networks. In R-Aloha, a terminal initially accesses TDMA slots in random access mode. The terminal has two options to obtain reservations in subsequent frames, namely by-use (implicit reservation) and by-use with indication of continuing use. In the first option, a terminal that successfully transmits a packet in random access mode has reservation over the same time slot position in subsequent frames. The terminal will be eligible for its reserved slot as long as it can still transmit a packet through the slot. Therefore, an unused reserved slot will be empty for one frame before it is made available to other terminals. The second option is aimed at avoiding this inefficiency. In this option, at every frame a terminal with reserved slots has to declare its intention to use its reserved slots in the next frame. Hence, an unused reserved slot can be immediately made available for contention. R-Aloha is a distributed access scheme and provides an expectedly good delay performance at low traffic load. As the load increases, however, collisions are becoming more frequent, as terminals are competing for smaller number of unreserved slots. As a result, R-Aloha gives poor performance at high load. A limitation of R-Aloha is that the frame duration has to be longer than the round-trip delay. This limitation is necessary to avoid terminal's ambiguity in identifying its reserved slots. Another limitation is that the protocol is not appropriate for voice-data integration, since it does not differentiate between voice and data terminals.

Packet Reservation Multiple Access (PRMA) is an adaptation of the R-Aloha protocol to the cellular environment. The protocol and its variants have received a lot of attention in the past [27-31]. PRMA takes advantage of the negligible propagation delay

in cellular communication to improve upon R-Aloha, which makes it inappropriate for satellite networks. However, many of the concepts incorporated in its variants may also be suitable for R-Aloha. The Centralized PRMA (C-PRMA) [28] assigns the role of scheduling the transmission by terminals to the base station, rather than to each terminal. By letting the scheduler have control over transmission scheduling, one has flexibility to implement more complex scheduling algorithm than in PRMA. Further, C-PRMA dedicates unreserved slots for reservation only. It is the scheduler's task to poll terminals, based upon reservation messages it receives and their priorities, in order to let them send their traffic. The bottleneck problem found in PRMA due to retransmission of data traffic is therefore alleviated. Further, C-PRMA provides the capability to integrate various classes of service into the network.

A technique called Data Steal into Voice (DSV) is proposed to adapt PRMA to an integrated voice and data system [29]. PRMA-DSV gives full (higher) priority to voice terminals. A two-bit Busy/Idle Signal (BIS) field for each slot is transmitted (broadcasted) in downlink messages and is used by the base station (or the scheduler in the case of satellite networks), to indicate whether a slot is reserved by voice terminals, reserved by data terminals, or unreserved. A voice terminal ready to send its first voice packet (of a talkspurt) has to randomly choose a slot from those unreserved by other voice terminals. Once the packet is transmitted, the terminal checks the BIS field of the same slot position in the next frame to see whether it can continue transmission. If the packet does not collide with other packets and if the slot is not reserved by any data terminals, the voice terminal may then continue transmission and it does not have to worry about the BIS field until it enters its next talkspurt state. Otherwise, it has to stop and wait for another slot unreserved by voice terminals. Data terminals, however, have to examine the BIS field of

any slot they intent to use. Further, they are only allowed to send their traffic on slots unreserved by both voice and other data terminals. Hence, even though both voice and data terminals have to contend for available slots for their first voice packets (of talkspurts) or data packets (of data messages), voice terminals are given preference over data terminals. As far as voice terminals are concerned, the communication channel will always look like a speech-only channel, irrespective of the data traffic load. The DSV technique, however, does not permit more than two level of priorities in the network.

Aggressive PRMA (A-PRMA) [30] is another method proposed to integrate both voice and data services into a wireless network. Its operational principle is quite similar to PRMA-DSV, in that voice terminals may compete for slots unreserved by other voice terminals, while data terminals may only compete for slots unreserved by both voice and data terminals. A-PRMA lets a data terminal which was successful in reserving a slot and whose message has experienced a delay exceeding a certain threshold put the slot in the “data lock” state, which makes the slot unavailable for both voice and data terminals until the terminal has completely emptied its message buffer. A-PRMA, however, like the PRMA-DSV technique, is designed to integrate only voice and data traffic.

Whereas the previous three flavours of PRMA are created with voice and data traffic in mind, PRMA / DA (PRMA / Dynamic Allocation) [31] is designed for wireless ATM networks. In PRMA/DA time on an uplink channel is divided into frames of equal duration and number of time slots. A frame is segmented into two regions, the available slot region and the reservation slot region. Slots in the available slot region are open for contention for all types of service classes. Reserved slots in the reservation slot region are grouped according to the classes of service and arranged in the following order : CBR slots, VBR slots, and ABR slots. PRMA/DA allows terminals to explicitly declare the

number of slots they required. It is up to the central scheduler, however, to allocate slots to reserving terminals. CBR reserving terminals are guaranteed allocations of slots at every frame, equal to their declared peak rates. Next, VBR reserving terminals are guaranteed at least the number of slots equal to their declared average rates. Finally, ABR reserving terminals will be allocated a number of slots equal to the minimum of the number of slots requested for ABR traffic and the number of slots unreserved by CBR and VBR terminals and not made available for contention. A unique feature of PRMA/DA is that the number of slots made available for contention is dynamically adjusted by the central controller, based on the intensity of network access demand, with the objective of providing the maximum attainable bandwidth to reserving stations while minimizing network access delay for contending terminals. One limitation in PRMA/DA is that it only guarantees assignment for VBR reserving terminals up to their average rates. This may require the network to treat real-time VBR terminals as CBR terminals, since only CBR terminals are guaranteed transmission rates equal to their peak rates. This, however, will lead to inefficiencies since rt-VBR traffic may be quite bursty.

Combined Random with Explicit then Implicit Reservation (CREIR) [32] is another access protocol belonging to this category and was originally conceived for satellite networks with onboard processing (scheduling) capability, although the concept can also be applied for satellite networks with ground-based scheduling. In CREIR, time on the uplink channel is divided into frames of equal length. Each frame is further broken down into two groups, a group of reservation minislots and another for data slots. Reservation (request) minislots are open for contention, while a data slot can be unreserved (in which case it is also open for contention) or reserved. A contending terminal, i.e., one with no reserved slots, randomly chooses one of the unreserved slots for

its first packet. At the same time the packet is sent, the terminal also sends a request for the number of packets it has in its buffer minus one (since one packet is already being sent). There are four outcomes for this action : both data and request packets may be transmitted successfully, or either one or both of them may collide. If both data and request packets are received successfully, the terminal will have one reserved slot at every frame for up to the previously requested number of time slot. The location of the slot, however, is determined by the scheduler. If only the data packet is received, the scheduler will assume that the terminal implicitly requires that one slot be made available for its use in subsequent frames. The scheduler will release the slot and put it back to the pool of contention slots, if it detects that the terminal does not use the slot in a frame. Next, if only the request packet is successfully received, the scheduler will assign one slot per frame to the terminal for up to the requested number of slots plus one (since the first slot collided and needs to be re-transmitted). If none of the packets is received free of collision and error, the terminal has to repeat the contention process for its traffic. CREIR attempts to reduce the probability of collision by sending data and request packets separately. The result is that it manages to provide a good and relatively stable delay performance up to a moderately high load, but performance degrades significantly in very high load situations. Further, even though [32] does not specifically indicate how CREIR (which is a centrally controlled protocol) will handle multiple classes of service, one can easily incorporate a more sophisticated scheduling algorithm to address this issue.

MASCARA/PRADOS [5, 33] is yet another multiple access protocol designed to integrate all five ATM service classes to wireless networks. MASCARA (Mobile Access Scheme Based on Contention and Reservation for ATM) is a medium access control protocol developed within the magic WAND (Wireless ATM Network Demonstrator)

project, while PRADOS (Prioritized Regulated Allocation Delay-Oriented Scheduling) is the channel allocation algorithm part of MASCARA. In MASCARA, similar to CREIR, time on the uplink channel is divided into two groups of request minislots and data slots. MASCARA, however, provides another method for request signaling, that is by piggy-backing requests in data slots. Unlike CREIR, the numbers of request minislots and data slots are variable from frame to frame. Similarly, the number of downlink data slots is also variable. A disadvantage of setting frame duration variable is that it adds complexity to both the base station (or the satellite) and terminals. The resulting capability to reach maximum uplink utilization, while maintaining each connection's negotiated QoS, may or may not outweigh this added complexity. PRADOS, the scheduling algorithm of MASCARA, is quite sophisticated in that it not only allocates bandwidth dynamically to satisfy all requests while taking into consideration the QoS requirement of each connection, but also arranges the location of slots allocated to each serviced request such that the cell delay tolerance (CDT) parameter of each serviced connection is met. A disadvantage of this technique would be the high signaling overhead of assignment messages, since slot assignment for each connection in a terminal may have to be separately encoded, rather than combined and encoded as total slot assignment for the terminal.

3.6 Combined Free-Demand Assignment Scheme

The combined free-demand assignment category was first introduced in [17]. In this category, unlike any other protocols described above, unreserved slots are freely distributed by the scheduler to terminals. This will simplify hardware configurations both in the satellite and ground terminals. Several techniques may be used to control the

distribution of these free slots. One algorithm may treat all terminals as equally eligible for free assignment of time slots (or free capacity assignment), while others may give preference to one type of terminals over another, based on a given criteria.

In [34], a version of CFDAMA with piggy-backed request signaling was described. In this flavour of CFDAMA, a terminal with no requests queued in the scheduler request queue has to wait until it receives free assignment of a time slot to send its traffic and imbed its request in the header of its data packet. Since this may create unfavourable conditions for terminals with no queued requests in high load situation, [34] suggests that such terminals be given higher priority in receiving free slot assignment, that is, all terminals with no queued requests will be considered for free capacity assignment before those with queued requests.

In [35], a version of CFDAMA with fixed-assigned request signaling was described. Here, a separate signaling channel is used, in which each terminal has one reserved request minislot for its use. This would solve the potential unfairness problem of the piggy-back version at a cost of more signaling overhead and a limited terminal population. Further, the authors suggest to integrate voice and data applications by adding implicit reservation for voice traffic, in the sense that an assigned slot for a voice terminal will be implicitly kept in subsequent frames until the terminal sends a release request.

Both versions of CFDAMA share many of the disadvantages of demand assignment scheme, including non-trivial signaling overhead and long end-to-end packet transfer delay. Furthermore, both versions do not address the issue of handling different ATM classes of service as well as providing QoS guarantee to each of the services. In addition, the issue of access fairness has not been sufficiently dealt with, especially in the environment of mixed traffic. Nevertheless, the potential of CFDAMA for satellite ATM

networks is quite substantial. Due to its centralized scheduling process, mechanisms to handle multiple classes of service and to increase access fairness can be incorporated relatively easily. Furthermore, since the uplink scheduler has total control over the distribution of unreserved slots to ground terminals, packet transfer delay can be reduced, while performance stability is still maintained.

3.7 Summary

Six general categories of multiple access schemes have been described in this chapter. Fixed assignment scheme is the simplest among the six categories. In this scheme, each earth terminal is allocated a fixed portion of uplink capacity. The strengths of fixed assignment scheme include fast access time and low packet transfer delay. The drawbacks include its inability to handle variation in network traffic efficiently and to support a large number of terminals.

Random assignment scheme allows an earth terminal to utilize the transmission channel as it wishes. It has many advantages, namely fast access time and low packet transfer delay at low traffic load, built-in support for bursty traffic and large terminal population, and a relatively simple access algorithm. The performance of random assignment scheme, however, degrades quickly at higher loads due to the increase in collision probability.

Demand assignment scheme aims to maximize uplink capacity utilization by requiring each earth terminal to request for transmission capacity to send its traffic. By avoiding collision, demand assignment scheme is able to achieve robust and stable performance up to very high loads. The primary drawbacks of this scheme are the non-trivial signaling overhead and the long transfer delay experienced by each packet.

The combined fixed-demand assignment scheme aims to accommodate both real-time and non real-time applications. In this scheme, a terminal has one or more fixed-assigned slots for its real-time traffic, and perhaps one or more slots obtained through requests for its non real-time traffic. This scheme, however, cannot completely avoid the main disadvantage of fixed assignment scheme, namely inefficient capacity utilization.

The combined random-demand assignment scheme attempts to combine many of the advantages of both random and demand assignment schemes, such as fast access time and low transfer delay at low traffic load and stable performance at high traffic load. Due to the non-zero probability of collision at any traffic load, however, this scheme may not be suitable for satellite networks, where long delay is inherent.

The combined free-demand assignment scheme distributes uplink capacity primarily based on demands. Furthermore, the uplink scheduler has full control over any unreserved capacity. The main advantages of this scheme is its stable and robust performance up to very high loads and its free capacity assignment feature. This scheme still shares the main drawbacks of demand assignment scheme, namely the non-trivial signaling overhead and the long packet transfer delay.

One objective of the thesis is to propose a multiple access scheme appropriate for satellite ATM networks. To this end, we will base our research on the combined free-demand assignment scheme. In the next chapter, we will explain thoroughly the methods developed to address the above issues. Furthermore, in Chapter 5, various performance aspects of the resulting scheme will be assessed using software (discrete-event) simulation.

Chapter 4

Enhanced CFDAMA for Satellite ATM Networks

The Enhanced CFDAMA (ECFDAMA) scheme has been presented briefly in [36], along with some performance results. The scheme attempts to improve upon the general CFDAMA protocol by introducing three new request handling mechanisms, or request/assignment (R/A) strategies, and incorporating several refinements. Each request handling mechanism receives different priority for service from the uplink scheduler. As a result, the access protocol has a built-in capability to handle different priorities that exist among the five ATM service classes. Further, a connection may utilize more than one R/A strategy during its lifetime. Since each request strategy aims to provide a different class of guarantee for quality of service, the access protocol has the capability to support flexible QoS matching and guarantee. Finally, the R/A strategies are quite different from each other and, as a whole, are designed to provide integrated support for both circuit-switched and packet-switched services, for both connection-oriented and connectionless services, and for multimedia applications.

In the following sections, we will describe first the frame structure chosen for the implementation of the access protocol. Next, we will explain the three new R/A strategies and refinements to the protocol. We will also describe how the protocol is implemented,

both onboard the satellite and in ground terminals. Finally, an example of data transfer operation will be given.

4.1 The Frame Structure

The MF-TDMA frame structure is chosen as the primary access scheme. MF-TDMA provides many advantages, which include simple integration into the reservation-based multiple access scheme to handle variation in traffic, efficient bandwidth utilization, small antenna size and low power requirement. Figure 3 below illustrates in detail a possible frame structure for a carrier group.

Frame duration is assumed fixed and equal to 24 ms. Each frame is comprised of multiple non-overlapping carriers and each carrier is subdivided in time into equal size slots. Each frame is composed of three types of time slots, namely synchronization slots, acquisition slots, and payload slots. Each synchronization slot consists of four frame reference bursts (FRB) for synchronization purposes and four out-of-band request (OBR) minislots. Each ground terminal is assigned an FRB and an OBR every 16 frames, or 384 ms. Therefore, with the frame structure shown in Figure 3 (i.e., 32 carriers per frame and 4 synchronization slots per carrier), a maximum of 8192 ground terminals can be maintained in synchronization in each carrier group. Acquisition slots are used by ground terminals to register to the network. Payload slots can be used by any type of connections (real-time and non real-time) to transmit data. Data is transmitted in a format similar to the ATM cell format with an additional 3-byte satellite header. Finally, an in-band request (IBR) minislot can be piggy-backed in each payload slot.

Figure 3. The Frame Structure

4.2 Request/Assignment Strategies

The three new R/A strategies are as follows : reserved capacity (RC) request strategy, rate-based dynamic capacity (RBDC) request strategy and volume-based dynamic capacity (VBDC) request strategy. Each of the above R/A strategies and the free capacity assignment strategy are described more elaborately below.

In the following description, the terms connection and terminal are alternatively used although they are not necessarily identical. A terminal may establish more than one connection during its login period. When more than one connection is established, the terminal will request capacity for all of its connections. As an advantage, the uplink scheduler is not required to maintain a database of all connections sharing the satellite network transmission capacity under its management. From this point on, we will refer to the portion of satellite network transmission capacity managed by a particular uplink scheduler as uplink capacity.

4.2.1 Reserved Capacity (RC) Strategy

In reserved capacity strategy, a terminal states its need for uplink capacity at connection setup time in terms of a fixed number of time slots per frame (i.e., a request of fixed-rate assignment). If the request is granted by the network controller, the terminal will have the requested number of uplink time slots assigned to it at every frame. The terminal essentially reserves a portion of network capacity. The request remains effective for the duration of the connection. Consequently, the network controller must ensure that the total reserved uplink capacity does not exceed the total uplink capacity. For this reason, the network controller maintains a database of terminals having reserved uplink capacity. The uplink scheduler also maintains a similar database, which is synchronized to that in the controller's database.

This request strategy is aimed toward real-time connections, those that cannot tolerate or have strict constraints on delay and delay variation (such as voice connections) and is especially suitable for applications with smooth traffic characteristic. CBR and rt-VBR service classes may take advantage of this request strategy. The network should avoid queueing of packets or data cells belonging to these service classes. Therefore, in

granting a fixed-rate request, the network controller must also take into consideration the available downlink transmission capacity.

There is no access delay for connections with fixed-rate request, except in the connection establishment phase. Henceforth, only transmission delay is experienced, and this amounts to one round-trip time (a round-trip time is defined as the length of the transmission path from the ground terminals to the satellite and back).

4.2.2 Rate-Based Dynamic Capacity (RBDC) Strategy

In rate-based dynamic capacity, a terminal may request a variable number of time slots in each frame (i.e., variable-rate request). This implies that the request is explicit and is expressed as a number of time slots per frame. The request is also implicit, in the sense that it remains effective for a number of frames. The terminal may need to update its request periodically. The duration of the periodic update, called the RBDC refresh duration, reflects the compromise between uplink scheduler's capacity to process requests and system responsiveness.

The uplink scheduler guarantees that the request, if received without error, will be granted up to a maximum value negotiated in the connection establishment phase. In effect, the terminal reserves a portion of network capacity, although it may not require all of this reserved capacity at all frames. Therefore, the network controller must also ensure that the total fixed-rate reserved capacity and the maximum variable-rate reserved capacity do not exceed the total available uplink and downlink capacity.

Similar to the RC strategy, the RBDC strategy is also directed toward real-time connections. However, as opposed to the RC strategy, it allows for statistical multiplexing among ground terminals. This should result in a more efficient use of network resources, with a little compromise in the quality of service. With this strategy, there is an access

delay of one round-trip associated with each request. Hence, the total delay experienced by a packet is two round-trips (see Chapter 4.5, Figure 16 for an illustration of data transfer operation using the RBDC strategy). This strategy is suited to applications with considerably bursty traffic characteristics and to those with a more relaxed constraint on delay. Certain rt-VBR types of applications, such as video distribution, may take advantage of this request strategy.

4.2.3 Volume-Based Dynamic Capacity (VBDC) Strategy

In volume-based dynamic capacity strategy, the network does not provide a guarantee on uplink capacity availability. The uplink scheduler will only try its best to satisfy the VBDC request. The request is explicit and a terminal states its request in terms of the total number of uplink time slots required to empty its queue. The request remains effective as long as not all of the requested time slots have been granted, or as long as the request has not timed out. The introduction of request time-out is aimed at reducing transmission overhead by removing the need for request acknowledgement. This aspect will be further explained in a later section. Finally, since there is no guarantee on uplink capacity, the constraint placed to the network controller in the connection establishment phase is relaxed.

This strategy is directed toward non-real-time connections. It is especially suited to bursty traffic. It allows for a great deal of statistical multiplexing, hence it contributes to a very efficient use of network resources. The strategy is further broken down into two types, namely high priority strategy (VBDC-HP) and low priority strategy (VBDC-LP). ABR and nrt-VBR service classes may take advantage of the high priority type, perhaps in combination with any of the first two R/A strategies. This is because nrt-VBR service has the SCR as one of its traffic parameters while ABR service has the MCR. Both SCR and

MCR imply a need for bandwidth guarantee, although the timeliness response to this need is not obligatory to the network. As far as UBR service class is concerned, it may utilize the low priority type, hence it will receive the lowest priority among the five ATM service categories. Access delay is variable for these connections, the minimum being one round-trip. Consequently, the minimum packet transmission delay is two round-trips.

4.2.4 Free Capacity Assignment Strategy (FCA)

Free capacity assignment strategy is a defining feature of the CFDAMA protocol. It is employed to maximize uplink capacity utilization. The uplink scheduler attempts to do this by distributing the available uplink capacity to all registered terminals. This strategy ranks the lowest among all four, the scheduler will satisfy all fixed-rate requests and both classes of variable-rate requests before it distributes unused uplink capacity through free capacity assignment. Further, unlike the other request strategies, ground terminals do not have any control in obtaining uplink capacity through this strategy.

4.3 Access Protocol Highlights

In addition to the four R/A strategies, the access protocol also has some refinements. These refinements are aimed toward improving throughput or uplink utilization, delay performance and fairness. Each of the additional features is explained below.

4.3.1 Light-Weight Access Protocol

The access protocol is designed to be light-weight. Error notification and data transfer acknowledgement is not built into the access protocol and is left for higher communication layers (network and upper layers). Since this approach is similar to that in

ATM technology, it is expected to have the same positive effects, such as reduced hardware complexity, both onboard the satellite and on the ground, and high data transfer capacity.

4.3.2 In-Band Request Strategy

As illustrated in Figure 3 above, a request minislot is piggy-backed in each payload slot. The terminal that has slot assignment in a particular frame owns the request slot piggy-backed in its assigned payload slot(s). Requests are then transmitted along with data, hence they are called in-band requests (IBR).

This technique of requesting has a few advantages :

- Improvement in system access delay

At low network load, a terminal is likely to have assignment at every frame (obtained through requests or free capacity assignment), hence a chance to transmit its request. At high network load, there is a risk of unfairness, since only terminals with slot assignments can send their requests. Nevertheless, terminals with slot assignments will experience a reduced access delay. Those that are not granted assignments are given opportunities to send their requests through out-of-band request (OBR) minislots, which will be described later. The OBR strategy compares to the case where ground terminals have fixed assignment of request minislots, in which access delay is not affected by network load but by terminal population (which means that high terminal population results in higher access delay).

- Reduced hardware complexity

The owner of a request can be identified simply by cross-referencing the current frame with the burst time plan sent one round-trip delay earlier. There is no possibility of collision as in a random access request strategy. Since collision detection onboard the

satellite requires a considerably sophisticated hardware operation, the reduction in hardware complexity is quite substantial.

The format of the in-band request minislot is given in Figure 4. The Request Type (RT) field is used to indicate what type of request strategy is being used. The encoding for this field is shown in Table 1.

Figure 4. The Request Minislot Format

Value	Request Type
00	no request
01	rate-based dynamic capacity strategy
10	volume-based dynamic capacity strategy, high priority
11	volume-based dynamic capacity strategy, low priority

Table 1. The Request Type Encoding

4.3.3 Out-of-Band Request Strategy

As explained earlier, there is a risk of unfairness with piggy-backing requests in payload slots. To decrease the impact of unfairness in accessing network resources, each terminal also has a fixed assignment of request minislot every few number of frames. These requests are transmitted separately from data, hence they are called out-of-band requests (OBR).

The distance between subsequent fixed request assignment is usually more than one frame and is a compromise between the designed maximum terminal population to the tolerable risk of unfairness. OBRs are also piggy-backed, in this case, in synchronization

slots (Figure 3). Each terminal has a fixed assignment of FRB in a synchronization slot, and the terminal that owns the FRB also owns the corresponding OBR. Hence, identifying the owner of an OBR can be performed simply by cross-referencing the database of FRB assignment.

The format of the out-of-band request minislot is similar to the in-band request minislot as illustrated in Figure 4.

4.3.4 Access Fairness through Round-Robin Scheduling

The uplink scheduler maintains a terminal table, which is a database of ground terminals currently registered in the network. The table is arranged in a circular linked-list fashion. Each entry in the table contains information related to a specific ground terminal, its configured RC, its maximum and currently requested RBDC, and two pointers to terminal's VBDC-LP and VBDC-HP queues. The scheduler guarantees capacity assignment for each terminal's RC and RBDC. Hence, access fairness is not relevant for the first two R/A strategies. However, there is no such guarantee for the last two R/A strategies, the VBDC and the FCA. The scheduler controls the distribution of the remaining uplink capacity to ground terminals, and it conducts this task in a round-robin fashion to ensure fairness.

The process of distributing uplink capacity to ground terminals, called the assignment process, involves a few passes through the terminal table. In the first pass, the scheduler counts uplink capacity requested by each of the request strategies (other than FCA). In the second pass, the scheduler inspects each entry in the table starting from the head entry of the linked list. In each inspection, the scheduler will satisfy the entry's RC and its RBDC request. It also attempts to satisfy each entry's VBDC-HP and VBDC-LP requests in full, as long as uplink capacity is still available. After each inspection, the entry

will be moved to the end of the linked list. During the second pass, the scheduler may place a tag on one entry for each of the non-guaranteed-bandwidth R/A strategies, i.e., the low and high priority VBDC and the FCA strategies. A tagged entry corresponding to one of these R/A strategies is the last entry receiving capacity assignment through that particular request strategy. Then, at the end of the second pass, the next entry after the highest-priority tagged entry will be made the head of the linked list. As a result, the two-pass assignment process assures that the entry (terminal) with the highest-priority, ungranted request is placed at the head of the terminal table linked list and is inspected first in the next assignment process (cycle).

4.3.5 Selective Free Capacity Assignment Strategy

The free capacity assignment strategy was originally intended to increase uplink capacity utilization and reduce packet transfer delay. Further, since requests can be piggy-backed in payload slots, this strategy can also be used to reduce system access delay. Care should be taken, however, to avoid unfairness situation, in which a few terminals dominate the use of uplink capacity. Fixed assignment of request minislots, as explained earlier, should reduce access unfairness by guaranteeing access opportunity every 16 frames. The selective free capacity assignment strategy is aimed at reducing access unfairness even further.

In the free capacity assignment strategy, the scheduler freely assigns one time slots to each ground terminal, as long as there are unused slots at each assignment process. In addition, the scheduler selectively assigns unused capacity, that is, terminals with no assignment in a frame is given a higher priority to receive free assignment than those with assignment.

One might argue that free assignment of one time slot is insufficient, especially for ground terminals with no assignment. The expected argument is that a request minislot can only be used by one request strategy. Yet, there are three types of requests, the RBDC strategy and the two flavours of VBDC strategy. The terminal may opt to send request of the first type, since it ranks the highest among the three. Hence, there is no opportunity of the other two request strategies to be used. The counter argument is as follows. Assignment of more than one time slots per terminal may benefit some terminals, especially those that intend to use more than one of the three strategies, but may be unfair for others that do not receive any assignment. However, the extra processing required to make this assignment strategy as fair as possible will increase the assignment processing time or add complexity to the onboard hardware. Further, terminals that utilize the RBDC as one of their strategies have a guaranteed access to the system with the arrival of their requested capacity. They may take advantage of this for other request strategies. In addition, terminals that use the high priority VBDC strategy for their nrt-VBR or ABR service (for which the strategy is intended for) will most probably employ either the RC or the RBDC strategy to meet the SCR or the MCR traffic parameter. Hence, terminals with only UBR traffic may be denied fair access to the network. The selective free capacity assignment strategy is aimed toward providing fair access to these terminals.

4.3.6 Absolute Assignment Signaling Strategy

There are two alternative strategies in signaling the assignment map or the burst time plan (BTP), namely relative (delta) assignment strategy and absolute assignment strategy. ECFDAMA uses the latter. In relative (delta) assignment strategy, the scheduler signals only the relative change in the map, while in absolute assignment strategy, the scheduler signals the whole new assignment map. The former strategy is aimed at reducing

communication overhead (due to the signaling of BTP), while the latter is aimed at simplifying the assignment message so that the process of reconstructing the BTP in ground terminals is trivial.

When traffic is bursty, as it is likely to be, the advantage of relative assignment strategy may not be realized. The reason is that the assignment map in one frame can be very different from the previous one, due to the variability in capacity requested by each terminal. Since for each slot in the relative assignment map, the slot location (in terms of frequency and time slot) and its owner terminal ID must be indicated, the amount of signaling involved will be significant. On the other hand, the signaling overhead associated with the absolute assignment strategy can be reduced by packing terminals' assignments to form the assignment message (as shown in Figure 6). Considering the frame structure shown in Figure 3 and the assignment message format shown in Figure 6, the maximum signaling overhead of the ECFDAMA scheme is 147 packets (4096 13-bit assignment messages per downlink frame and 28 assignment messages per downlink data slot), which corresponds to approximately 3.6 % of the uplink capacity.

Another reason for a significant relative change in the assignment map is related to the build process of a new map at the end of each assignment cycle. As described earlier, the assignment process is performed in a round-robin fashion for fairness reason. Then, as capacity assignment for each terminal is determined during the second pass, the position of time slot assignment for the terminal is also immediately set in a consecutive arrangement, as shown in Figure 5. The assignment message is incrementally constructed during the second pass, with the format shown in Figure 6.

Figure 5. Consecutive Time Slot Assignment

Figure 6. The Assignment Message Format

The disadvantage of such arrangement of time slot position is that there is an inherent variability in a terminal's slot position from one frame to another. A terminal with requests that is located at the head of the terminal table linked list will be assigned uplink capacity in early time slots of the first carrier. In the next frame, due to the round-robin

scheduling strategy previously described, it may be located at the end of the terminal table and be assigned capacity in different time slots of possibly one or two different carriers.

An alternative strategy for time slot positioning is to pre-assign time slot position for each terminal's reserved capacity (which is obtained through the RC or the RBDC strategy). Hence, the position of time slot assignment for real-time traffic does not change from one frame to another. This, however, may place an unnecessary restriction on the amount of RC or RBDC that can be granted by the scheduler. This is due to the limitation where a terminal cannot transmit on more than one carrier simultaneously. However, due to the pre-assignment of time slots for registered terminals, non-concurrent slot position for new (registering) terminals with reserved capacity may not be found. Hence, a call is blocked unnecessarily. In the spirit of maximizing uplink capacity utilization, we choose not to adopt this alternative strategy.

Figures 7 to 9 depict the two-pass assignment process previously described in detail. The selective free capacity assignment and absolute assignment strategies are also illustrated in the flow charts. Table 2 contains some descriptions of important variables used in the flow charts.

Variable	Explanation
Total Capacity	total uplink payload capacity
C_{carrier}	total payload capacity per carrier (frequency)
Table	pointer to terminal table data structure, e.g., Table.First Entry is the pointer to the first terminal entry in the terminal table
T	pointer to a terminal entry, e.g., T.Assign is the final capacity assigned to the pointed terminal
Total BG Count, Total HP Count, Total LP Count	total uplink capacity requested by each corresponding strategy
Total HP Avail, Total LP Avail, Total FCA Avail	total uplink capacity available for each corresponding strategy
HP Tagged Term, LP Tagged Term, FCA Tagged Term	pointer to a tagged terminal for each corresponding strategy
HP Tagged, LP Tagged, FCA Tagged	flag to indicate whether a terminal has been tagged for each corresponding strategy

Table 2. Description of Variables in the Assignment Process Flow Charts

Figure 7. The Assignment Process First Pass

- ♦ note that $(a,b,c)^-$ means the minimum of a , b , and c , while $(a,b,c)^+$ means the maximum of a , b , and c .

Figure 8. The Assignment Process Second Pass

Figure 9. The Assignment Process Second Pass (Cont'd)

4.3.7 Flexible Association of Service Classes to Request/Assignment Strategies

Even though each of the four R/A strategies is directed toward traffic with different characteristics, users do have flexibility in combining the strategies to match the characteristics and quality of service requirement of their traffic. For example, to meet the stringent delay and delay variability performance of rt-VBR traffic, a user may reserve capacity for rt-VBR traffic at its PCR using the RC strategy. However, if for economic reasons, one decides to tolerate longer delay yet maintain delay variability constraints, one can reserve capacity using the RC strategy at the SCR and set the maximum RBDC value

at the difference between the PCR and the SCR. Even further, if one decides to transmit its rt-VBR traffic more inexpensively, the user can reserve capacity using only the RBDC strategy with the maximum RBDC value set at the PCR. As another example, to maximize statistical multiplexing, one may choose to use the high priority VBDC strategy to transmit ABR traffic. Another choice is to use the RBDC strategy with the maximum RBDC value set at the traffic's PCR and the current value continually updated at the signalled Allowed Cell Rate (ACR). This choice limits the maximum population of terminal with ABR traffic, however, since the total capacity reserved (using both the RC and the RBDC strategies) by terminals must not exceed the uplink capacity. Nevertheless, for ABR connections with non-zero MCR, one should still reserve capacity at the MCR, using either the RC or the RBDC strategy, so that the guaranteed MCR can be respected. Table 3 shows several possible associations of service classes to R/A strategies.

Class	RC	RBDC	RC + RBDC	VBDC - HP	RC + VBDC - HP	RBDC + VBDC - HP	VBDC - HP
CBR	✓ ¹	✓ ²					
rt-VBR	✓ ¹	✓ ²	✓ ³				
nrt-VBR			✓ ³		✓ ⁴	✓ ⁵	
ABR				✓ ⁶	✓ ⁷	✓ ⁸	
UBR							✓

Table 3. Several Possible Association of Service Classes to Request/Assignment Strategies

- Note :
- ¹ RC is set at PCR
 - ² maximum RBDC value is set at PCR
 - ³ RC is set at SCR and maximum RBDC value is set at (PCR - SCR)
 - ⁴ RC is set at SCR
 - ⁵ maximum RBDC value is set at SCR
 - ⁶ VBDC-HP capacity is limited by ACR
 - ⁷ RC is set at MCR, VBDC-HP capacity is limited by (ACR - MCR)
 - ⁸ maximum RBDC value is set at MCR, VBDC-HP capacity is limited by (ACR - MCR)

4.4 Terminal Uplink Access Process

Efficient use of R/A strategies by ground terminals is key to the performance of the ECFDAMA scheme. Each terminal has to avoid wasting any capacity by requesting only the necessary capacity to empty its queues. Each one should efficiently utilize its capacity assignment in order to meet the quality of service requested by its connections. In this section, the uplink access processing in ground terminals will be described. Methods used to calculate requests and utilize capacity assignments will be explained.

4.4.1 Terminal Internal Structure

We will start by describing the internal structure of a terminal. A terminal is responsible for the management of traffic (data and control messages) received from both terrestrial sources and the satellite. Figure 10 illustrates crucial multiple access - related components of a terminal. Solid lines represent the flow of traffic, while dotted lines represent the flow of information (such as queue size and allowable rate for ABR connections).

Traffic received from terrestrial sources is first segmented into 48-byte cells and formatted for uplink transmission. A terminal has five queues (one for each service class) for the formatted cells waiting for transmission. Traffic received from the satellite can be separated into two types, namely data traffic and messages. Data traffic is subject to re-assembly process before it is forwarded to terrestrial destinations. Message traffic is composed of (ABR) rate control and assignment messages and is delivered to the uplink access process. The uplink access process is responsible for the management of uplink transmission and will be described further in the following section. Figure 11 summarizes the tasks each terminal has to perform in each frame.

To compensate for the missing data transfer acknowledgement, a terminal keeps track of requests which have not been granted or expired (in the case of VBDC requests) in the pending request queues. A copy of each transmitted request is inserted into the corresponding pending request queue and is given a time-out counter. The counter value is decremented in every subsequent frame. When the counter value is zero, the request is removed from the pending request queue. This means that the terminal gives up hope for the request and prepares to make another one. The counter is set initially at V for RBDC requests (where V is the sum of round-trip delay and processing delay with the unit of

frame) and at $V + W$ for VBDC requests (where W is the time-out duration of VBDC requests with the unit of frame and is a known value in the whole network).

Figure 10. A Typical Block Diagram of an Earth Terminal

Figure 11. A Summary of Terminal's Tasks Performed at Every Frame

4.4.2 Uplink Access Process

As shown in Figure 11, the uplink access process performs the following tasks :

- monitor traffic queues and pending request queues and take appropriate actions (i.e., request additional capacity requests, release reserved capacity which is no longer needed and remove granted or expired requests from pending request queues)
- maintain and update the burst time plan
- schedule the transmission of cells in the five traffic queues

Figure 12 depicts the algorithm used by the uplink access process. Detailed methods used to calculate capacity requests have been omitted for the purpose of clarity and will be explained more elaborately in the following section.

Figure 12. The Terminal Uplink Access Process Algorithm

4.4.2.1 Request Calculation Method

Before we can describe the method used to compute capacity requests, we need to suggest a mapping of service categories to R/A strategies. The strict delay and delay variation requirements of CBR service can only be satisfied by the RC strategy. Uplink capacity is reserved for CBR traffic at its PCR per frame (i.e., $R = \lceil PCR \cdot T \rceil$ slots are reserved, where T is the frame duration and $\lceil x \rceil$ is the integer ceiling of x). Therefore, conforming CBR cells only need to wait for one frame before being transmitted. This results in a maximum cell transfer delay of close to the round-trip propagation time and in a small delay variation (small delay variation is inherent due to the round-robin scheduling). Next, rt-VBR service is associated with the RBDC strategy (the maximum RBDC value for rt-VBR traffic is set to its PCR per frame) to achieve a higher multiplexing gain of non real-time traffic and still maintain a predictable cell transfer delay and delay variation. Hence, conforming rt-VBR cells will experience a transfer delay of close to two round-trips. Then, nrt-VBR service can be mapped to a combination of RC and RBDC strategies. Uplink capacity is reserved for nrt-VBR traffic at its SCR, while the maximum RBDC value for nrt-VBR traffic is set to the difference between its PCR and SCR. The result is that most nrt-VBR cells will experience a transfer delay of one round-trip (since the average cell rate in the long run should be close to the SCR), while the rest of the cells will experience a transfer delay of two round-trips (since they will be transmitted using time slots requested using the RBDC strategy). ABR service is associated to a combination of RC and high priority VBDC strategies. Uplink capacity is reserved for ABR traffic at its MCR and additional need for uplink capacity above the MCR will be signalled using the high priority VBDC strategy. We will postpone treatment of rate control for ABR service until Chapter 6. In the mean time, a terminal may send as

much ABR traffic as allowed by its capacity assignment. Finally, UBR service is matched with the low priority VBDC strategy.

Considering the above service association, we will now describe how capacity requests for each service category are calculated. CBR traffic has a reserved capacity equivalent its PCR. Therefore, a terminal does not need to request capacity for its CBR traffic. Both rt-VBR traffic and nrt-VBR traffic use the RBDC strategy. Hence, a terminal must calculate a capacity request for these two service classes in every frame. Let us first define the following variables :

Q_{rt-VBR}	:	the current size of the rt-VBR queue
$Q_{nrt-VBR}$	:	the current size of the nrt-VBR queue
$RC_{nrt-VBR}$	:	capacity reserved for nrt-VBR traffic
C_{rt-VBR}	:	current assignment for rt-VBR traffic using the RBDC strategy
$C_{nrt-VBR}$	:	current assignment for nrt-VBR traffic using the RBDC strategy
CR_{rt-VBR}^j	:	the j th RBDC request for rt-VBR traffic in the pending RBDC queue
$CR_{nrt-VBR}^j$	:	the j th RBDC request for nrt-VBR traffic in the pending RBDC queue
TO^j	:	the counter value of the j th RBDC request in the pending RBDC queue
M_{RBDC}	:	the number of RBDC requests in the pending request queue
CR_{RBDC}^{Max}	:	the maximum negotiated RBDC capacity request
$C_{carrier}$	:	the payload capacity in each carrier (frequency)

- CR^{Max} : the maximum capacity request (limited by the size of the capacity field in the request minislot format)
 V : then round-trip delay plus processing delay (with the unit of frame)
 RQ_{rt-VBR} : the calculated request for rt-VBR traffic for the current frame
 $RQ_{nrt-VBR}$: the calculated request for nrt-VBR traffic for the current frame
 RQ_{total} : the calculated total request (using RBDC) for the current frame

Then :

$$\begin{aligned}
 RQ_{rt-VBR} = & Q_{rt-VBR} - C_{rt-VBR} - (V - TO^1) \cdot CR_{rt-VBR}^1 \\
 & - \sum_{j=2}^{M_{RBDC}} (TO^{j-1} - TO^j) \cdot CR_{rt-VBR}^j
 \end{aligned} \tag{1}$$

$$\begin{aligned}
 RQ_{nrt-VBR} = & Q_{nrt-VBR} - (V \cdot RC_{nrt-VBR}) - C_{nrt-VBR} - (V - TO^1) \cdot CR_{nrt-VBR}^1 \\
 & - \sum_{j=2}^{M_{RBDC}} (TO^{j-1} - TO^j) \cdot CR_{nrt-VBR}^j
 \end{aligned} \tag{2}$$

The total amount of capacity request is then calculated as follows :

$$RQ_{total} = (RQ_{rt-VBR} + RQ_{nrt-VBR}, 0)^+ \tag{3}$$

where $(a, b)^+$ is the maximum of a and b .

Next, the total capacity request must not exceed the maximum RBDC value. It must also not exceed the total payload capacity in each carrier (since a terminal is not able to transmit in more than one carrier simultaneously). Finally, the total capacity request cannot exceed the maximum capacity request (limited by the size of the capacity field in

the request minislot format). Therefore, the total capacity request is bounded by the following relation :

$$RQ_{total} = \left(CR_{RBDC}^{Max}, C_{carrier}, CR^{Max}, RQ_{total} \right)^- \quad (4)$$

where $(a, b)^-$ is the minimum of a and b .

Further, due to the bounding operation above, the final total capacity request may be less than the original (calculated) value. Hence, the following operations are necessary to modify the calculated capacity request for each service class. Note the higher priority placed on rt-VBR service the operations imply.

$$RQ_{rt-VBR} = (RQ_{rt-VBR}, RQ_{total})^- \quad (5)$$

$$RQ_{nrt-VBR} = RQ_{total} - RQ_{rt-VBR} \quad (6)$$

Both rt-VBR and nrt-VBR service classes require bandwidth guarantee. Therefore, a terminal must ensure that at least one payload slot is available to send (piggy-back) its request. For the above reason, the following operation is performed :

$$\begin{aligned} & \text{if} \left((RC_{nrt-VBR} = 0) \text{ and } (CR_{RBDC}^{Max} \neq 0) \right) \\ & RQ_{total} = (RQ_{total}, 1)^+ \end{aligned} \quad (7)$$

In order to understand equations (1) and (2), recall that RBDC requests, once granted by the uplink scheduler, is effective for a number of frames (corresponding to the RBDC refresh duration). Therefore, if at frame j a request of x slots is granted by the scheduler and never updated, then up to frame i ($i > j$), the scheduler will have allocated a total of $(i - j + 1) \cdot x$ time slots. For example, the expression :

$$(TO^1 - TO^2) \cdot CR_{rt-VBR}^2$$

gives the total number of time slots which will be assigned to rt-VBR traffic due to the second request in the pending RBDC request queue. Hence, the last two terms of equations (1) and (2) give the total number of slots which will be assigned to rt-VBR and nrt-VBR traffic, respectively, due to requests in the pending RBDC request queue. Furthermore, the term $(V - RC_{nrt-VBR})$ in equation (2) gives the reserved capacity for nrt-VBR traffic in the next V frames. Finally, in whole, equations (1) and (2) give the additional capacity needed to send all rt-VBR and nrt-VBR cells (up to the current frame), respectively, in V frames.

ABR and UBR services are mapped to the high priority and low priority VBDC strategies, respectively (assuming zero MCR). One needs to consider the current queue size and the total capacity that has been requested during the previous frames to calculate additional capacity request for each traffic class. Let us first define the following variables:

Q_{ABR}	:	the current size of the ABR queue
Q_{UBR}	:	the current size of the UBR queue
C_{ABR}	:	current assignment for ABR traffic using the VBDC-HP strategy
C_{UBR}	:	current assignment for UBR traffic using the VBDC-LP strategy
CR_{ABR}^j	:	the j th VBDC-HP request in the pending VBDC-HP queue
CR_{UBR}^j	:	the j th VBDC-LP request in the pending VBDC-LP queue
$M_{VBDC-HP}$	:	the number of VBDC-HP requests in the pending VBDC-HP queue

- $M_{VBDC-LP}$: the number of VBDC-LP requests in the pending VBDC-LP queue
 CR^{Max} : the maximum capacity request (limited by the size of the capacity field in the request minislot format)
 RQ_{ABR} : the calculated request for ABR traffic for the current frame
 RQ_{UBR} : the calculated request for UBR traffic for the current frame

Then :

$$RQ_{ABR} = Q_{ABR} - C_{ABR} - \sum_{j=1}^{M_{VBDC-HP}} CR_{ABR}^j \quad (8)$$

$$RQ_{UBR} = Q_{UBR} - C_{UBR} - \sum_{j=1}^{M_{VBDC-LP}} CR_{UBR}^j \quad (9)$$

The amount of capacity request for each service class is then bounded above and below as follows:

$$RQ_{ABR} = \left(\left(CR^{Max}, RQ_{ABR} \right)^-, 0 \right)^+ \quad (10)$$

$$RQ_{UBR} = \left(\left(CR^{Max}, RQ_{UBR} \right)^-, 0 \right)^+ \quad (11)$$

The third term in equations (8) and (9) gives the total number of time slots which have been requested using the VBDC strategy and have not been granted. Therefore, in whole, equations (8) and (9) give the additional capacity needed to send all ABR and UBR cells (up to the current frame), respectively, in V frames.

4.4.2.2 Transmission Scheduling Strategy

In addition to calculating capacity requests, the uplink access process is also responsible for the scheduling of uplink transmission. This task involves maintaining an updated BTP and scheduling cells from each of the five traffic queues for transmission. In order to perform the transmission scheduling task efficiently, the uplink access process needs to calculate the number of cells it can transmit from each service class. We will now describe the method used to obtain this number.

The uplink access process first subtracts the number of reserved slots for CBR and nrt-VBR traffic from its total capacity assignment (denoted by C_{total}). Next, it determines if it has capacity assignment due to its RBDC requests by examining the oldest (top-most) request in the pending RBDC request queue. If the request's counter value is zero, the process continues to compare its total capacity assignment with the requested capacity (denoted by CR_{RBDC}^1 , capacity requested by the oldest, queued RBDC request). If C_{total} is larger than CR_{RBDC}^1 , the process assumes that the oldest request has been granted (i.e., its capacity assignment due to the RBDC strategy, denoted by C_{RBDC} , is CR_{RBDC}^1), and it can send as many as CR_{rt-VBR}^1 and at least $CR_{nrt-VBR}^1$ of rt-VBR and nrt-VBR cells, respectively. C_{RBDC} has a time-out counter and the counter value is set equal to the RBDC refresh duration every time the value of C_{RBDC} is changed (and is decremented in each subsequent frame). If the request's counter value is larger than zero, the process continues to compare C_{RBDC} from the previous frame with C_{total} . If C_{total} is larger than C_{RBDC} value and if the counter value of C_{RBDC} is larger than zero, the process assumes

that the old C_{RBDC} value is still effective. In any other cases, the process will assume that there is no capacity assignment due to the RBDC strategy.

The process determines capacity assignment due to VBDC and FCA strategies by subtracting C_{RBDC} from C_{total} . Next, the process examines the oldest request queued in the pending VBDC-HP request queue. If the request's counter has been reduced to less than or equal to W (i.e., capacity assignment due to the queued VBDC-HP request should have arrived, if granted), then it assumes that it has capacity assignment for ABR traffic. In this case, the process continues to compare its remaining capacity assignment (C_{total}) with the requested capacity (denoted by CR_{ABR}^1 , capacity requested by the oldest, queued VBDC-HP request), and the smaller value becomes its capacity assignment for ABR traffic (C_{ABR}). Any remaining capacity assignment is due to the VBDC-LP request or to the FCA strategy. The process performs similar operation to determine its capacity assignment for UBR traffic (C_{UBR}).

Any remaining capacity assignment after the above operations is due to the FCA strategy. We restrict the use of FCA to non real-time traffic (i.e., nrt-VBR, ABR, and UBR) for the following reason. CBR traffic reserves uplink capacity equivalent to its PCR, hence, it does not need FCA. Further, the capacity request calculation method for rt-VBR traffic aims at achieving a timely drain of rt-VBR queue so that each rt-VBR cell experiences a transfer delay of $2V$ frames and correspondingly, a small delay variation. If an rt-VBR cell can be transmitted using a slot due to the FCA strategy, the cell's transfer delay will be smaller than $2V$ frames. Furthermore, the effect will spread to cells scheduled for transmission in subsequent frames. As a result, cell transfer delay variation for rt-VBR

traffic increases. Among the three service classes which may use free capacity assignment, nrt-VBR traffic is given the highest priority and UBR traffic the lowest.

We employ a simple priority scheme for transmission scheduling where CBR queue receives the highest priority and UBR queue the lowest. Figure 13 and 14 describe the pseudo-code of the algorithm used to distribute capacity assignment to the five traffic classes. Figure 15 depicts the transmission scheduling algorithm. Note in Figure 15 that any unused reserved capacity may be used by both ABR and UBR traffic (ABR traffic is given a higher priority).

```

Start of Procedure for Time Slot Assignment Distribution
Decrement RBDC counter value
Initialize  $C_{ABR}$  , and  $C_{UBR}$  to 0
Free Capacity Assignment = 0
 $C_{total} = C_{total} - RC_{CBR} - RC_{nrt-VBR}$ 
Find the oldest queued request queued in the pending RBDC queue
if (the queued request counter value = 0)
    Obtain  $CR_{RBDC}^1$  ,  $CR_{rt-VBR}^1$  ,  $CR_{nrt-VBR}^1$ 
    Dequeue and discard request
    if (  $C_{total} \geq CR_{RBDC}^1$  )
         $C_{RBDC} = CR_{RBDC}^1$ 
         $C_{rt-VBR} = CR_{rt-VBR}^1$ 
         $C_{nrt-VBR} = CR_{nrt-VBR}^1$ 
        Reset RBDC counter value to  $V$ 
    endif
else if (RBDC counter value  $\leq 0$ )
     $C_{RBDC} = 0$ 
     $C_{rt-VBR} = 0$ 
     $C_{nrt-VBR} = 0$ 
endif
 $C_{total} = C_{total} - C_{RBDC}$ 
Find the oldest queued request in the pending VBDC-HP queue
if (  $0 \leq$  the queued request counter value  $\leq W$  )
    Obtain  $CR_{ABR}^1$ 
    if (  $C_{total} \geq CR_{ABR}^1$  )
         $C_{ABR} = CR_{ABR}^1$ 
    else
         $C_{ABR} = C_{total}$ 
    endif
     $C_{total} = C_{total} - C_{ABR}$ 
     $CR_{ABR}^1 = CR_{ABR}^1 - C_{ABR}$ 
    if (  $CR_{ABR}^1 = 0$  )
        Remove the queued request from the pending VBDC-HP queue
    endif
else if (the queued request counter value = 0)
    Remove the queued request from the pending VBDC-HP queue
endif

```

Figure 13. The Pseudo-code for Time Slot Assignment Distribution Algorithm

```

Find the oldest queued request in the pending VBDC-LP queue
if (0 ≤ the queued request counter value ≤ W)
    Obtain  $CR_{UBR}^1$ 
    if (  $C_{total} \geq CR_{UBR}^1$  )
         $C_{UBR} = CR_{UBR}^1$ 
    else
         $C_{UBR} = C_{total}$ 
    endif
     $C_{total} = C_{total} - C_{UBR}$ 
     $CR_{UBR}^1 = CR_{UBR}^1 - C_{UBR}$ 
    if (  $CR_{UBR}^1 = 0$  )
        Remove the queued request from the pending VBDC-LP queue
    endif
else if (the queued request counter value = 0)
    Remove the queued request from the pending VBDC-LP queue
endif
if (  $C_{total} > 0$  )
    Free Capacity Assignment = 1
endif
End of Procedure

```

Figure 14. The Pseudo-code for Time Slot Assignment Distribution Algorithm (cont'd)

```

Start of Procedure for the Initialization of Cell Transmission Scheduling
  Set no CBR cells sent, no RT VBR cells sent, no NRT VBR cells sent to 0
  Set no ABR cells sent, no UBR cells sent to 0
End of Procedure

Start of Procedure for Cell Transmission Scheduling
  for every (time slot assignment for the terminal is due)
    if ((CBR queue not empty) and (no CBR cells sent <  $RC_{CBR}$ ))
      send one cell from CBR queue
      no CBR cells sent = no CBR cells sent + 1
    else if ((rt-VBR queue not empty) and (no RT VBR cells sent <  $C_{rt-VBR}$ ))
      send one cell from rt-VBR queue
      no RT VBR cells sent = no RT VBR cells sent + 1
    else if ((nrt-VBR queue not empty) and ((no NRT VBR cells sent <  $RC_{nrt-VBR} + C_{nrt-VBR}$ ) or
      ((no ABR cells sent  $\geq C_{ABR}$ ) and (no UBR cells sent  $\geq C_{UBR}$ )) or
      (Free Capacity Assignment = 1)))
      send one cell from nrt-VBR queue
      if (Free Capacity Assignment = 1)
        Free Capacity Assignment = 0
      else
        no NRT VBR cells sent = no NRT VBR cells sent + 1
      endif
    else if ((ABR queue not empty) and ((no ABR cells sent <  $C_{ABR}$ ) or (no UBR cells sent  $\geq C_{UBR}$ ) or
      (Free Capacity Assignment = 1)))
      send one cell from ABR queue
      if (Free Capacity Assignment = 1)
        Free Capacity Assignment = 0
      else
        no ABR cells sent = no ABR cells sent + 1
      endif
    else if (UBR queue not empty)
      send one cell from UBR queue
      if (Free Capacity Assignment = 1)
        Free Capacity Assignment = 0
      else
        no UBR cells sent = no UBR cells sent + 1
      endif
    endif
  end for
End of Procedure

```

Figure 15. The Pseudocode for Cell Transmission Scheduling Algorithm

4.5 Example of Data Transfer Operations

Figure 16 illustrates data transfer operation using the RBDC strategy. In this illustration, the round-trip propagation delay is assumed to be 11 frames (or 264 ms), which is typical for a geosynchronous satellite. The processing delay is assumed to be 2 frames (one frame in the satellite and in the ground terminal).

Let us assume that terminal X successfully registers itself and negotiates a maximum RBDC value of 8 time slots for its rt-VBR traffic at frame 0. The current RBDC value is initialized to 0 at frame 0. Let us also assume that the RBDC refresh duration is set to 63 frames and that free capacity assignment is granted to X in every frame (due to the low uplink load) so that at frame 1 (and for the next 11 frames) the scheduler assigns 1 time slot to X using the FCA strategy.

The FCA sent at frame 0 arrives at the terminal during frame 6 and is available for use in frame 7. Let us assume X has only rt-VBR traffic and 8 rt-VBR cells are received by the end of frame 6. At frame 7, X generates and transmits a request of 8 cells. Observe that X does not use the FCA to send its rt-VBR traffic, since this will increase cell transfer delay variation. One new cell arrives during frame 7, hence, X generates a request of 1 time slot in frame 8. Furthermore, 5 new cells are received during frame 8 so that X generates another request of 5 cells in frame 9, since it has requested a total of 9 time slots in the previous two requests.

Assuming an error-free transmission channel, the RBDC request sent in frame 8 arrives at the satellite during frame 12. The scheduler, however, processes the request at the beginning of frame 13 in order to ensure that all in-band requests in frame 12 have arrived. The scheduler grants the requested capacity of 8 slots to X since the requested capacity does not exceed X 's negotiated maximum RBDC value is 8. Capacity assignment

sent at the beginning frame 13 arrives at X during frame 17, hence, X can send 8 rt-VBR cells in frame 18. Observe that all 8 cells received in frame 6 are transmitted during frame 18. Hence they all experience a similar cell transfer delay, and correspondingly, cell transfer delay variation will be small.

Figure 16. An Example of Data Transfer Operation using the RBDC Strategy

4.6 Satellite Model

We will now describe a typical satellite model in order to illustrate multiple access - related processing inside the satellite. The block diagram of an onboard processing satellite is given in Figure 17. A satellite consists of seven components, namely the antennas, the receivers, the transmitters, the multi-carrier demodulators, the input / output ports, the switch, and the onboard network controller (OBNC). The functions of antennas, receivers and transmitters are apparent. The functions of the other four components are described below.

Figure 17. A Typical Block Diagram of an Onboard Processing Satellite

The multi-carrier demodulator (MCD) extracts uplink payload packets, embeds the position information (i.e., the carrier and time slot positions), and forwards them to the

input port. The MCD also extracts OBRs from synchronization slots, encodes them into messages, and forwards these messages to the input port.

The input port receives data traffic and messages forwarded by the MCD. Inside the input port resides the uplink scheduler. The uplink scheduler extracts IBRs from payload packets and OBRs from MCD-generated messages, and inserts them into their owner's request queues. It can determine the owner of each request by cross-referencing the request's position (in the MF-TDMA frame) with the assignment map constructed V frames earlier (V represents the round-trip delay plus some processing delay onboard the satellite and in the ground terminal). In order to do this, it keeps assignment maps built in the previous V frames. Every 24 ms (the duration of a frame) the uplink scheduler performs the assignment process previously described and builds an assignment map. An input port may have more than one uplink scheduler, each is responsible for a group of carriers (or a carrier group).

The output port maintains five priority queues for each downlink channel (one for each of the five ATM service categories), with CBR queue ranks the highest and UBR queue the lowest. Higher priority queues are emptied before cells are taken from the lower priority queues to build a downlink frame. Hence, the scheduling process in each downlink channel is much less complex than the uplink counterpart. An output port may have more than one downlink scheduler, each is responsible for a particular downlink channel.

The switch is responsible for delivering traffic from input ports to output ports and between ports and the OBNC. Finally, the OBNC is in charge of general network management duties, such as handling connections' set-ups and tear-downs and performing congestion control related processing.

Chapter 5

Performance Evaluation of ECFDAMA

The sophisticated scheduling algorithm, the multiple R/A strategies, and the various enhancement incorporated into the ECFDAMA protocol, combined with the variety of service classes supported by ATM, make it extremely difficult, if not impossible, to assess the performance of the protocol through analytical means. Therefore, the performance of ECFDAMA has been evaluated using software simulation. This chapter describes the system's configuration and the miscellaneous test cases employed to assess various performance aspects of the scheme. Simulation results are presented in terms of several important performance measures, namely throughput, cell transfer delay and cell transfer delay variation.

5.1 Simulation Model

The ECFDAMA scheme described in the previous chapter is accurately implemented for discrete event simulation in OPNET environment. The satellite network (shown in Figure 18) is composed of a geosynchronous satellite with onboard processing capabilities and a number of ground terminals. The satellite may support a number of spot beams, each corresponds to a single carrier group. All terminals located in the same spot beam are considered at the same distance from the satellite (35,786 km), resulting in a

round trip propagation delay of approximately 250 ms. Each terminal can support up to five traffic sources (generators), one for each service class defined by the ATM Forum [1]. The model user can associate each traffic generator with one of the available traffic models. For the simulation results reported in this thesis, two traffic models have been used, namely a Markov Modulated Poisson Process (MMPP) model, which can be considered representative for voice applications, and a self-similar model capable of capturing the burstiness of data applications.

Figure 18. Satellite Network Configuration for Simulations

Following the approach in [38], the MMPP model is used to represent N packetized voice sources. Each voice source can be represented as an ON-OFF model with exponentially distributed length of talkspurts and silence periods with average length of α^{-1} and β^{-1} respectively. In the ON state, each voice source generates cells at a constant rate of A cells/sec. Therefore, the peak to average ratio (PAR) of each voice source can be expressed as :

$$PAR = \frac{\alpha + \beta}{\beta} \quad (12)$$

The aggregation of N voice sources is represented by a 2-state MMPP model as described in [38]. The model switches between periods of underload and overload. Traffic arrival rates in underload and overload periods are assumed to be given by Poisson distributions with parameters λ_1 and λ_2 respectively. Similarly, sojourn time distributions in both periods are Poisson with parameters σ_1^{-1} and σ_2^{-1} respectively. Parameter values for the 2-MMPP are found using the parameter matching technique described in [38] and based on [39], which, in addition to the ON-OFF parameters of voice sources, also considers the intended Index of Dispersion for Counts (IDC), which describes traffic burstiness [37]. In this thesis, the model is used to match 6 voice sources ($N = 6$), each with the same parameters ($A = 41.667$ cells/sec, $\alpha = 0.7871$, $PAR = 3.25$), and is set to have an IDC of 10.72. Table 4 summarizes the 2-MMPP parameter values and Figure 19 shows the IDC graph of the model.

$\lambda_1 = 54.0038$ cells/sec
$\lambda_2 = 138.8699$ cells/sec
$\sigma_1^{-1} = 1.0256$ sec ⁻¹
$\sigma_2^{-1} = 2.7721$ sec ⁻¹

Table 4. The 2-MMPP Model Parameter Values

The self-similar model used in this thesis is an empirical model of World Wide Web (WWW) request-arrivals. A detailed description of the model can be found in [40]. This model is selected to represent data traffic because, with the widespread popularity of the Internet, it is likely for data traffic in future satellite networks to have the same characteristics as WWW traffic. The WWW traffic model is an ON-OFF model with sojourn times in ON and OFF periods following heavy-tailed Weibull and Pareto distributions, respectively. The ON state represents the arrivals of WWW requests, while the OFF state represents the “thinking time” after users receive their requested WWW pages. Further, the interarrival of requests in the ON state is also found to be Weibull and heavy-tailed. For the thesis, it is assumed that each WWW request has a fixed size of 48 bytes (the distribution of WWW document sizes is also found to be Pareto [41]).

Figure 19. The IDC of the 2-state MMPP Model

A Weibull distribution has the following probability density function :

$$f(x) = \frac{k}{\theta} \left(\frac{x}{\theta} \right)^{k-1} e^{-(x/\theta)^k} \quad (13)$$

For the ON state, [40] suggests that $0.77 \leq k \leq 0.91$ and that $e^{4.4} \leq \theta \leq e^{4.6}$, while for requests interarrival in the ON state, the suggested values for k and θ are 0.5 and $e^{1.5}$, respectively. Further, a Pareto distribution has the following probability density function :

$$f(x) = \frac{\alpha k^\alpha}{x^{\alpha+1}} \quad (14)$$

Here, k represents the smallest length of OFF periods. The paper does not suggest a range of values for both α and k (except that $\alpha \leq 1$), but it uses the values 0.5 and 60 for α and k , respectively, as an example of the empirical model for the OFF periods.

A Pareto distribution has infinite mean and variance when $\alpha \leq 1$. Consequently, a source may be idle for a long time. In order to limit the required simulation time, the OFF periods are bounded by a maximum value, the cut-off time. In this case, if a value (OFF period) drawn from the distribution is larger than the cut-off time, a new value is drawn until one that is smaller or equal to the cut-off time is picked. In order to get a reasonably accurate model, however, the cut-off time should be set sufficiently large. For the thesis, a cut-off time of 2000 seconds is considered adequately large, since it is safe to assume that users will be disconnected after being inactive for more than 2000 seconds.

Table 5 summarizes the chosen parameter values for the WWW traffic model, while Figure 20 shows the IDC graph of the model. Higher IDC values (compared to the MMPP model) are indicative of greater variability of the traffic generated by the model.

$k_{ON} = 0.75$
$\theta_{ON} = e^{4.61} = 100.484$
$k_{IA} = 0.5$
$\theta_{IA} = e^{1.6} = 4.953$
$k_{OFF} = 96.0$
$\alpha_{OFF} = 0.5$
<i>cut - off time</i> = 2000 sec

Table 5. The WWW Model Parameter Values

Simulations have been performed on a scaled-down model, in which the satellite supports only one carrier group. Each uplink MF-TDMA frame has a duration of 24 ms and consists of 4 non-overlapping carriers. Each carrier is composed of 4 synchronization slots and 128 payload slots, resulting in a carrier burst rate of 2048 Kbps and an uplink capacity of 8192 Kbps. Downlink capacity is set to be slightly larger than uplink capacity to offset the signaling overhead (assignment messages). With the above assumptions, each cell arriving to the satellite is almost immediately transmitted downlink, i.e., there is no possibility of onboard congestion. Log slots are omitted since the login/logoff process of ground terminals does not need to be simulated to evaluate throughput, delay and delay variation. Correspondingly, the number of ground terminals can be fixed at the start of each simulation. Acquisition slots are also omitted since terminals are assumed to be synchronized at all times.

Figure 20. The IDC of the Self-Similar (WWW) Model

5.2 Service Association and Traffic Mixing in Ground Terminals

As indicated in Chapter 4, the ECFDAMA scheme allows for flexible association of ATM service classes to R/A strategies. Clearly, the scheme's performance under mixed traffic environment depends upon the particular service association employed by the network. For the purpose of illustrations, the mapping of service classes to R/A strategies described in Section 4.4.2.1 is selected, since it emphasizes the strengths of each strategy (or a combination of strategies). Continuous traffic (CBR, rt-VBR, and nrt-VBR) is associated to either one or a combination of RC and RBDC strategies. First, CBR service class is mapped to the RC strategy. A CBR source reserves $R = \lceil PCR \cdot T \rceil$ slots in each frame. Next, rt-VBR service class is mapped to the RBDC strategy, with the maximum RBDC value set at $R = \lceil PCR \cdot T \rceil$ slots. Then, nrt-VBR service class is mapped to a combination of RC and RBDC strategies, in which capacity is reserved at $\lceil SCR \cdot T \rceil$ slots

and the maximum RBDC value is set at $\lceil (PCR - SCR) \cdot T \rceil$ slots. Data (ABR and UBR) traffic, however, relies on the VBDC strategy, since the strategy is aimed toward maximizing uplink utilization. Further, not being a rate-based request strategy, this strategy is suitable for bursty sources. In this case, ABR service class is mapped to a combination of RC and high priority VBDC strategies. An ABR source reserves $R = \lceil MCR \cdot T \rceil$ slots in each frame and requests any additional capacity using the VBDC-HP strategy. Further, UBR service class is mapped to the low priority VBDC strategy. Finally, non real-time traffic (nrt-VBR, ABR and UBR) is eligible to use free capacity assigned slots. Table 6 summarizes the selected association of service classes to R/A strategies.

Service Class	RC	Max RBDC	VBDC-HP	VBDC-LP	FCA
CBR	$\lceil PCR \cdot T \rceil$				
rt-VBR		$\lceil PCR \cdot T \rceil$			
nrt-VBR	$\lceil SCR \cdot T \rceil$	$\lceil (PCR - SCR) \cdot T \rceil$			✓
ABR	$\lceil MCR \cdot T \rceil$		✓		✓
UBR				✓	✓

Table 6. The Mapping of ATM Service Classes to Request Strategies

CBR traffic is represented by a constant rate traffic generator, hence, CBR cells will experience a transfer delay of approximately 250 ms. Both rt-VBR and nrt-VBR traffic are represented by the 2-MMPP traffic model. Cell interarrival times for each rt-VBR or nrt-VBR source is scaled to give an average cell rate of 4 cells/frame (64 kbps) and a peak cell rate of 16 cells/frame (256 kbps). ABR and UBR traffic are represented by the WWW traffic model. Similarly, cell interarrival times for each ABR or UBR source is scaled to give an average cell rate of 1 cell/frame (16 kbps).

In order to allow meaningful study of protocol behaviour and performance assessment, traffic is mixed in ground terminals. Ground terminals are grouped according to the combination of traffic mixes and the traffic loading strategy. Table 7 summarizes the 10 types of ground terminals used in simulations. Type 1 terminals, for example, only generate rt-VBR traffic, while type 7 terminals generate a mixed (rt-VBR and ABR) traffic. Further, traffic load in type 1 terminals is fixed at an average rate of 64 kbps and a peak rate of 256 kbps, while that in type 6 terminals is variable, starting at an average rate of 16 kbps.

5.3 Simulation Test Cases

Several simulation test cases have been defined to assess throughput and delay performance of the ECFDAMA scheme under a variety of traffic mixes and terminal populations. Table 8 summarizes the 15 test cases considered for performance evaluation. In each case network load is gradually increased by increasing the number of terminals in the network, by increasing the load in each terminal, or a combination of both. Further, in order to achieve a low cell transfer delay for jitter-tolerant traffic (nrt-VBR, ABR and UBR), the selective free capacity assignment strategy is only used at an extremely high load (above 95 %). At lower loads, the uplink scheduler considers all terminals to be equally eligible for free capacity assignment.

Test cases 1 to 5 are useful in evaluating the scheme's performance under homogeneous traffic. Test case 1 measures the throughput and delay performance of the network as the load is increased by increasing the number of rt-VBR terminals (sources). Similarly, test cases 2, 3, and 4 evaluate the scheme's performance under an increasing number of nrt-VBR, ABR, and UBR terminals, respectively. In test case 5, on the other

hand, network load is increased by increasing the load in each ABR terminal. This test is useful to evaluate the scheme's performance when each (ABR) source is actively competing for uplink capacity.

Terminal Type	Traffic Type(s)	Traffic Load	Average Cell Rate	Peak Cell Rate	Traffic Generator
1	rt-VBR	fixed	64 kbps	256 kbps	MMPP
2	nrt-VBR	fixed	64 kbps	256 kbps	MMPP
3	ABR	fixed	16 kbps	bursty	WWW
4	UBR	fixed	16 kbps	bursty	WWW
5	ABR	fixed	8 kbps	bursty	WWW
	UBR	fixed	8 kbps	bursty	WWW
6	ABR	variable	≥ 16 kbps	bursty	WWW
7	rt-VBR	fixed	64 kbps	256 kbps	MMPP
	ABR	variable	≥ 16 kbps	bursty	WWW
8	nrt-VBR	fixed	64 kbps	256 kbps	MMPP
	ABR	variable	≥ 16 kbps	bursty	WWW
9	rt-VBR	fixed	64 kbps	256 kbps	MMPP
	ABR	variable	≥ 8 kbps	bursty	WWW
	UBR	variable	≥ 8 kbps	bursty	WWW
10	nrt-VBR	fixed	64 kbps	256 kbps	MMPP
	ABR	variable	≥ 8 kbps	bursty	WWW
	UBR	variable	≥ 8 kbps	bursty	WWW

Table 7. Various Terminal Types and Configurations for ECFDAMA Performance Evaluation

Test cases 6 and 7 should give the first insight into the scheme's performance in mixed traffic environment. In test case 6, a combination of rt-VBR and ABR traffic is considered, while in test case 7, a combination of nrt-VBR and ABR traffic is considered. In both cases, each ground terminal only generates one type of traffic. The number of VBR terminals is fixed, while network load is increased by increasing the number of ABR

terminals. These two tests are useful for evaluating the effectiveness of the VBDC strategy in utilizing uplink capacity unused by continuous (VBR) traffic for data (ABR or UBR) traffic. In test case 6, continuous traffic terminals only request capacity needed to empty their queues, by using the RBDC strategy. In test case 7, however, these terminals always have capacity reserved in every frame, even though they do not have traffic to transmit. Therefore, delay performance for data traffic should be better in test case 6 than in test case 7. Due to the reserved capacity, however, nrt-VBR traffic should experience a lower average transfer delay than rt-VBR traffic.

Test cases 8 and 9 are similar to the previous two cases except that network load is increased by increasing the load in each ABR terminals. This should give insights into the expected delay performance of data traffic when data sources are actively competing for uplink capacity and when continuous, variable rate traffic is present in the background.

Test case 10 and 11 are similar to test cases 6 and 7, except that data terminals are generating traffic of different priorities, namely the higher priority ABR traffic and the lower priority UBR traffic. These cases will provide insights into the expected performance of lower priority data traffic with a background of continuous and higher priority data traffic.

In test cases 6 to 11, continuous terminals are separated from data terminals. These cases are useful to evaluate the inter-terminal multiplexing aspect of the ECFDAMA, that is, how effectively are data terminals in utilizing capacity not requested by continuous traffic terminals, by using the VBDC strategy. In test cases 12 to 15, however, all terminals have both continuous and data traffic components. The assessment focus for these cases is on the intra-terminal multiplexing aspect of the scheme, that is, how effectively are data sources in utilizing capacity not used by continuous sources.

Test cases 12 and 13 consider a combination of rt-VBR and ABR traffic and a combination of nrt-VBR and ABR traffic, respectively. Since capacity is not reserved for rt-VBR traffic, ABR (data) terminals have to rely on VBDC-HP and FCA strategies to transmit their traffic. On the contrary, since some capacity is reserved for nrt-VBR traffic, ABR (data) terminals may transmit their traffic using unused reserved capacity. Therefore, delay performance for data traffic in the latter case should be better than in the former. Finally, test cases 14 and 15 are similar to the previous two cases, except that data traffic is composed of an equal proportion of ABR (high priority) and UBR (low priority) traffic.

Test Case	Terminal Type	Quantity
1	1	variable (4 to 32)
2	2	variable (4 to 32)
3	3	variable (64 to 512)
4	4	variable (64 to 512)
5	6	fixed (64)
6	1 3	fixed (32) variable (64 to 384)
7	2 3	fixed (32) variable (64 to 384)
8	1 6	fixed (32) fixed (64)
9	2 6	fixed (32) fixed (64)
10	1 5	fixed (32) variable (64 to 384)
11	2 5	fixed (32) variable (64 to 384)
12	7	fixed (32)
13	8	fixed (32)
14	9	fixed (32)
15	10	fixed (32)

Table 8. Simulation Test Cases for The ECFDAMA Performance Evaluation

5.4 Simulation Results

Figure 21 (a) - (d) show the simulation results of the first and second test cases. Figure 21 (a) and (b) show the load applied to each rt-VBR or nrt-VBR terminal and the uplink load (throughput). Figure 21 (b) shows a linear relationship between throughput and the applied (generated) load. Figure 21 (c) confirms an average transfer delay of close to two round-trips (approximately 22.5 frames or 540 ms) experienced by rt-VBR cells (each rt-VBR cell has to spend one round-trip delay waiting for its due-requested slot and one round-trip delay in transmission). This substantial delay may prevent many real-time applications from using the RBDC strategy. Nevertheless, the small and predictable mean delay variation experienced by rt-VBR cells (shown in Figure 21 (d)) indicates that the RBDC strategy is suitable for connections which can tolerate delay, but not variation in delay (such as video distribution). Furthermore, by requesting only the necessary capacity to empty their queues, rt-VBR connections contribute to efficient utilization of network resources. The average transfer delay experienced by nrt-VBR cells is much smaller due to the reserved capacity. The mean delay variation for nrt-VBR cells, however, is much larger, since nrt-VBR cells might escape their due-requested slots and are transmitted using undue or free assigned slots. As a result, escaping nrt-VBR cells will experience a shorter cell transfer delay than the rest.

Figure 22 (a) - (d) present the simulation results of test cases 3 and 4. The results of both cases are identical since both ABR and UBR services are associated to the VBDC strategy (only of different priority). Figure 22 (b) shows a linear relationship of throughput and applied load up to an extremely high network load (97 %). A stable mean delay of approximately 17 frames (408 ms) is experienced up to a load of 75 % (384 data terminals). As the load is further increased, the mean cell delay is rapidly increasing.

Nevertheless, even at a load of 97 % (512 data terminals), a modest mean delay of 27 frames (648 ms) is experienced. This excellent delay performance is due to fairness features incorporated in the uplink scheduler's assignment process. By queueing (instead of accumulating) requests per terminal and allocating capacity in a round-robin fashion, considering only one request per request type per terminal in each cycle, the scheduler is able to treat each data terminal fairly.

Figure 21. Test Cases 1 and 2 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 22. Test Cases 3 and 4 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 23 (a) - (d) show the results of test case 5, where network load is increased by increasing the load applied to each data (ABR) terminal. As network load increases, each terminal will be generating more traffic and correspondingly, larger and more frequent request for capacity. Hence, data terminals will be actively competing for uplink capacity. Figure 23 (b) shows that the increase in throughput is still linear. Further, the average cell transfer delay is better than in test case 3, which indicates that the scheduler is able to grant requests from most data terminals in each assignment cycle, due to fairness features of the assignment process. The average cell delay of less than 24 frames (576 ms) at extremely high load (97 %) is only slightly larger than that experienced by rt-VBR cells in test case 1. This indicates that the average size of terminal request queue (onboard) is 2, so that each ABR cell has to wait one additional frame before it can be transmitted. The average cell delay variation of 1.5 frames (i.e., one frame larger than the average delay variation of rt-VBR cells in test case 1) at the intended 100 % load confirms the previous statement.

Figure 24 (a) - (d) give the first insight into the scheme's performance in mixed traffic environment. In this test, a steady load of background rt-VBR traffic (25 %) is injected into the network, while data traffic load is increased by increasing the number of data terminals. Figure 24 (b) shows that rt-VBR throughput is steady, while ABR throughput increases linearly with the applied load. Figure 24 (c) and (d) show a similar cell delay performance of rt-VBR traffic as those of test case 1 (Figure 21 (c) and (d)). This indicates that continuous traffic can be isolated from bursty data traffic by using two different request strategies. With isolation, QoS guarantee for continuous traffic can be honoured at any loading situation. The cell delay performance of ABR traffic deteriorates faster at high loads compared to that of test case 3 (Figure 22 (c)), due to the higher

priority rt-VBR traffic. Nevertheless, a modest mean delay of 28 frames (672 ms) is experienced at an extremely high load (97 %), which is only slightly larger than that of test case 3. The cell delay variation performance of ABR traffic remains similar to that in test case 3 (Figure 22 (d)).

The results of test case 7 shown in Figure 25 (a) - (d) also indicate an isolation of continuous (nrt-VBR) traffic from data (ABR) traffic. Figure 25 (c) and (d) show a similar cell delay performance of nrt-VBR traffic as those of test case 2 (Figure 21 (c) and (d)), except at high loads. This is because at high loads there are fewer slots available for free assignment. As a result, most nrt-VBR cells are transmitted using reserved and due (RBDC) requested slots. Figure 25 (c) shows a slightly lower cell delay performance of ABR traffic compared to that in test case 6 (Figure 24 (c)), since a portion of uplink capacity is reserved for continuous traffic. This reserved capacity will be wasted if not used since data terminals, which are physically separated from continuous traffic terminals, cannot utilize these unused slots.

Figure 23. Test Case 5 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 24. Test Case 6 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 25. Test Case 7 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 26 (a) - (d) show the simulation results of test case 8. Test case 8 is similar to test case 6, except that the total load from ABR traffic is increased by increasing the load in each ABR terminal. Similar to previous results, throughput increases linearly with applied load. The isolation of continuous traffic from data traffic is again observed. Further, the cell delay performance of ABR traffic is almost identical to that in test case 5 (Figure 23). The slightly longer delay it encounters at extremely high loads is due to the variability of available capacity due to capacity requests from rt-VBR terminals. However, compared to test case 6 results (Figure 24 (c) and (d)), ABR cells enjoy a much better delay performance, since they receive free capacity allocation more frequently due to a low terminal population.

Test case 9 results shown in Figure 27 (a) - (d) indicate the same isolation effect up to a load of 90 %. The cell delay performance of nrt-VBR traffic is identical to that of test case 7 (Figure 25 (c)) up to a fairly high load (70 %), and is marginally better at high loads, since free capacity assignment is distributed more frequently due to a low terminal population. Correspondingly, ABR cells experience a much better delay performance than that of test case 7 (Figure 25 (c)) at high loads.

Figure 26. Test Case 8 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 27. Test Case 9 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

In test cases 10 and 11, data terminals generate an equal proportion of traffic with different priorities, namely the high priority ABR traffic and the low priority UBR traffic. In both cases, continuous traffic load (rt-VBR and nrt-VBR, respectively) is maintained at 25 %, while data traffic load is increased by increasing the number of data terminals. Continuous traffic terminals are still physically separated from data terminals so that the multiplexing of traffic occurs onboard the satellite through the scheduling process.

Figure 28 (a) - (d) present the simulation results of test case 10. As shown, rt-VBR throughput remains steady, while data throughput increases linearly with applied load. Continuous traffic is still isolated from data traffic, as indicated by the stable delay and delay variation experienced by rt-VBR traffic at all loads. Further, data traffic enjoys an excellent cell delay performance (less than 24 frames and 26 frames for ABR and UBR traffic, respectively, at a load of 97 %). UBR traffic experiences a slightly worse cell transfer delay than ABR traffic, since ABR traffic is given a higher request priority (by using the VBDC-HP strategy) and a higher priority in using free capacity allocation.

Figure 29 (a) - (d) show the simulation results of test case 11, where continuous traffic is of nrt-VBR type. The cell delay performance of nrt-VBR traffic is almost identical to that of test cases 7 and 9 (Figure 25 and Figure 27, respectively), while the cell delay performance for ABR traffic is slightly worse than that exhibited in test case 8 (Figure 28), due to a waste of uplink capacity (unused reserved capacity for continuous traffic). Finally, the cell delay performance of UBR traffic is marginally worse than that of ABR traffic due to its lower priority.

Figure 28. Test Case 10 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 29. Test Case 11 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

In test cases 12 to 15, each ground terminal generates both continuous and bursty data traffic. Similar to previous cases of traffic mixing, continuous traffic load is held at 25 %, while data traffic load is gradually increased by increasing the load in each ground terminal (instead of adding more data terminals). Therefore, the multiplexing of traffic mainly occurs within each ground terminal (intra-terminal multiplexing). In this case, data traffic can utilize time slots unused by continuous traffic. In test cases 12 and 13, a combination of rt-VBR and ABR traffic and a combination of nrt-VBR and ABR traffic, respectively, are considered. In test cases 14 and 15, data traffic is composed of an equal proportion of high priority ABR traffic and low priority UBR traffic.

Figure 30 (a) - (d) show the simulation results of test case 12. It is clearly seen that the delay performance of rt-VBR traffic is identical to that achieved for previous test cases, indicating the isolation effect. Further, since rt-VBR traffic is transmitted only using RBDC requested time slots, no bandwidth is wasted, resulting in a similar cell delay performance of data traffic to that exhibited in test case 8 (Figure 26).

Figure 31 (a) - (d) depict the simulation results of test case 13. For this case, the cell delay performance of nrt-VBR traffic is excellent, with a stable average transfer delay of approximately 13 frames (182 ms) up to a total network load of approximately 87.5 %, and decreases slightly at higher loads, due to the more infrequent allocation of free capacity. The cell delay performance of ABR traffic is also excellent, since it is now able to use any reserved capacity unused by nrt-VBR traffic.

Finally, as already seen in several previous test cases, the average ABR cell transfer delay variation decreases as the load is increased. This is due to the more infrequent assignment of free capacity, so that most ABR cells have to be transmitted using their due-requested time slots. The benefit of cell transfer delay reduction brought

about by the FCA strategy has a cost of high delay variation. This confirms the decision (see Chapter 4) to prevent CBR and rt-VBR traffic from utilizing free capacity assigned slots.

Figure 32 (a) - (d) present the simulation results of test case 14. The isolation of continuous (rt-VBR) traffic from data traffic at all loads is again exhibited. Further, as compared to the results of test case 10 (Figure 28), it is clearly visible that the cell delay performance of ABR traffic is higher. This is due to a low number of terminal population (32 ground terminals), compared to a terminal population ranging from 96 to 416 ground terminals in test case 10. As a result, each terminal receives free capacity assignment more frequently. Furthermore, since rt-VBR traffic cannot use free assigned time slots, FCA slots are mostly used to send ABR traffic. Consequently, the average cell transfer delay of ABR traffic is reduced. This is confirmed by a slow decrease in mean delay variation for ABR traffic, indicating that it often uses FCA slots up to a fairly high network load.

Figure 33 (a) - (d) depict the simulation results of test case 15. Notice that the cell delay performance of UBR traffic is almost identical to that of ABR traffic. This is because nrt-VBR traffic is eligible and has the highest priority to use free assigned time slots. Therefore, most FCA slots will be used by nrt-VBR traffic, and consequently, the average nrt-VBR cell transfer delay is maintained at a low value of approximately 13 frames (182 ms). Due the low number of terminal population and the benefit of higher priority for free capacity assignment than data traffic, the cell delay performance of nrt-VBR traffic is maintained at a low value of approximately 13 frames (182 ms). Only at very high loads, where FCA slots are more rarely distributed, does it reduce slightly.

Figure 30. Test Case 12 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 31. Test Case 13 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 32. Test Case 14 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

Figure 33. Test Case 15 Simulation Results : Packet Generation Load vs. (a) Generated Load; (b) Uplink Load (Throughput); (c) End-to-End Mean Delay; (d) End-to-End Std. Dev. Delay

5.5 Summary

In summary, all simulation results have confirmed the excellent performance and features of the ECFDAMA scheme. All results indicate that stable performance is maintained up to the maximum load. Next, continuous traffic experiences the benefit of isolation from data traffic, allowing the network to maintain the QoS guarantee promised to each continuous traffic connection (CBR, rt-VBR, or nrt-VBR connection). Further, the cell delay performance of data traffic is excellent at low to medium loads, while a graceful degradation in performance is experienced at high to extremely high loads due to several fairness features incorporated into the scheduler's assignment process. Finally, the FCA strategy is shown to be useful in maintaining high uplink utilization and reducing the average transfer delay of jitter-tolerant traffic (nrt-VBR, ABR, and UBR traffic). However, the strategy has a negative effect of increasing the mean delay variation, since ground terminals do not have control in obtaining FCA slots. Therefore, the decision to prevent CBR and rt-VBR traffic from using free capacity assigned slots is justified.

Chapter 6

Review of Congestion Control Schemes

Onboard the satellite, congestion on a certain downlink channel occurs when packet arrival rate to the downlink channel exceeds the channel bandwidth or transmission capacity. Considering the typical internal architecture of an onboard processing satellite shown previously in Figure 17, congestion situation can arise when more than one uplink channel simultaneously transfers packets to the same downlink channel. In this case, the downlink queue will quickly fill up and finally packets will be discarded. These discarded packets may have to be re-transmitted, which would further worsen the situation. The conclusion is that congestion control schemes employed in a satellite network have a significant role in determining the network's performance.

There are several reasons why congestion control is a very difficult problem for satellite networks. First, the phenomenon of self-similarity in traffic arrival pattern has been identified in various networks, including local area networks and the Internet. It is then reasonable to assume that the traffic arrival pattern in satellite networks is also self-similar. This will bring about stochastic and unpredictable bandwidth demands to downlink channels. Second, due to the high delay-bandwidth product, in the case of congestion a large number of packets will be in transmission before the satellite is able to notify the earth terminals about the congestion situation. Finally, in destination directed packet-

switched satellites, the destination of each packet is unknown until it arrives to the satellite [10].

In spite of these difficulties, a large number of congestion control schemes have been developed in the past. In [43] R. Jain stated several criteria which should be taken into consideration when one intends to design and implement a suitable congestion control scheme, namely scalability, optimality, fairness, robustness and implementability. A good scheme should be scalable in the sense that it is able to work in a wide range of network configurations and sizes, and optimal in the sense that it allows for optimum use of network resources (bandwidth and buffer size). The scheme should be fair to all users or terminals and robust, that is, it should not be sensitive to minor network imperfections, such as lost control messages and occasional measurement inaccuracies. Finally, a good scheme should not be constrained to certain hardware architecture. With respect to satellite networks, one may need to include yet another design criteria - the complexity. The current lag in hardware technology for space applications and the mass/power constraints onboard the satellite limit the processing power onboard, and correspondingly, the amount of processing associated with the congestion control scheme.

This chapter reviews some of the schemes which have been developed for both ATM and non-ATM networks. Congestion control schemes can be grouped into three layers depending on the time scale of the congestion, namely call-layer, burst-layer, and packet-layer. Call-layer control schemes are specifically designed for congestion which lasts longer than a session duration. Burst-layer control schemes are meant for congestion which lasts for a duration comparable to a burst (active period). Finally, packet-layer control schemes are targeted for congestion which lasts shorter than the end-to-end round-trip delay. It has been argued and accepted that all control layers should be present

in a network, since congestion may manifest over different duration [10, 42]. Nevertheless, call-layer and packet-layer controls do not have the capability to react to the changing network load, so that in order to be effective they require a conservative use of network resources [10]. Burst-layer control schemes, with their reactive or even predictive components, do have this capability. Hence, the focus of this chapter is on burst-layer congestion control schemes, which may have potential for implementation in satellite environment.

6.1 Congestion Control Schemes for Non-ATM Networks

Several selected congestion control schemes for non-ATM networks will be described shortly. These schemes are selected due to their suitability, at least with respect to some aspects, for implementation in satellite-based systems. Selection criteria include overhead, implementation complexity, capability to cope with the long delay inherent in satellite-based systems, and performance.

In [58] R. Jain proposed a delay-based approach for congestion avoidance. In his scheme, Jain is more concerned with congestion avoidance rather than congestion control in the sense that the scheme was designed to bring the network to the operating point giving low delay and high throughput, which he called the knee. This is the point where, with any further increase in network load, throughput will be stagnant or even deteriorate, while the queue will start to increase quickly. In his delay-based approach, each terminal attempts to deduce network load by measuring packet transfer delay. Further, the scheme assumes that a window-based flow control mechanism is employed, so that throughput is directly proportional to the size of the window. It then attempts to determine the optimum window size by computing the power, which is defined as the ratio of throughput and

delay. The delay-based approach only requires that the network provides some sort of acknowledgement to packet transfer operation, so that a terminal is able to estimate packet transfer delay. Since the network does not have to provide explicit feedback to terminals, the scheme is categorized as implicit feedback scheme. Some features of this scheme include zero overhead, distributed control, robustness, and scalability. Further, the scheme can be employed as a complement to other (explicit) feedback schemes, to increase system performance by the increase in the available information on congestion situation. The only drawback of delay-based approach is that it has not been explored and tested thoroughly.

A feedback congestion control scheme tailored for networks with long propagation delay was proposed in [59]. The scheme is based on the use of “start” and “stop” signals to bring the queue size within a desired value. Virtues of the scheme include a simple algorithm and capability to achieve zero cell loss ratio (CLR) while avoiding starvation. The scheme has several drawbacks, the most important being the connection between the maximum throughput and the targeted overhead. Low overhead can be achieved at the price of low maximum throughput for lossless operation. This is because the lower the overhead, the more infrequent is the generation of “start” and “stop” signals. Combined with the goal of zero CLR, this results in a higher ratio of unavailable periods to available periods, and correspondingly, a lower maximum throughput. The scheme also assumes that each source is persistent, i.e., each source always has packets to send. It is expected that the scheme’s performance under bursty traffic degrades considerably.

A predictive control scheme for packet switching satellites was presented in [52]. The author argued that reactive schemes will not be efficient in handling congestion situations due to the long propagation delay between the satellite and earth terminals. A

reactive scheme reacts on the onset of congestion. Due to the high bandwidth-delay product, however, a large number of packets will be in transmission before the first control message (notifying the onset of congestion) is effective. As a result, in order to obtain a reasonably small CLR, the onboard switch must have a very large output buffer. Unfortunately, buffer size is an expensive resource of a satellite network due to the mass/power and technological limitations. The paper, therefore, claimed that a predictive scheme would be necessary for efficient operation. The paper suggested two predictive methods, one based on the use of packet arrival rate statistics, the other based on the use of a neural network whose input data includes packet arrival rate, packet size distribution, source address, and destination address. In the scheme, when near-term congestion is predicted, the satellite will broadcast a rate control ("rate reduce") signal to the earth terminals, governing the maximum rate at which a terminal may transmit in its uplink channel to each of the satellite's downlink channels. Even though the paper did not present any results, the scheme's merits can be easily recognized. First, due to its predictive features, the required output buffer size for efficient operation should be reduced considerably. Second, by taking advantage of the inherent broadcast capability of a satellite, fast delivery of rate control signals is achieved under any loading situation, thereby increasing the scheme's effectiveness. Finally, low overhead is attained since a control signal for one uplink channel only consists of "rate reduce" message for each of the satellite's downlink channels, and is effective for all terminals sharing the uplink channel. Nevertheless, predictive schemes do have several disadvantages. In order to be effective, the scheme must be able to estimate, with sufficient accuracy, the probability of congestion for one round-trip duration, which is close to 250 ms for geosynchronous satellite-based systems. This requires a large amount of data for a reasonably accurate

estimation, which require a large storage memory and a fast onboard processor. Further, with bursty and self-similar traffic, it is highly probable that previous data will not reflect future one. As a result, accurate estimation may not be obtained, and correspondingly, the system's performance will suffer.

Whereas the previous scheme delegates the control responsibility to the satellite, the reactive scheme described in [53] attributes most of this responsibility to ground terminals. Similar to the previous one, this scheme also takes advantage of the capability of the satellite to broadcast global network status to all ground terminals simultaneously. In this scheme, the satellite will continuously broadcast the status of its output buffer and the ground terminals will adjust their rates based on this information and according to some pre-defined thresholds. When the satellite's output buffer size is below the first threshold, each terminal is allowed to transmit at its original rate. Then, as the buffer size increases and exceeds the first threshold, each terminal is required to reduce its original rate by a pre-defined fraction. Each time the buffer size exceeds a higher threshold, the terminal is required to reduce its original rate by a larger fraction. The most important feature of the scheme is the low complexity required onboard the satellite. The satellite is only required to monitor and broadcast its buffer size. The computation burden is moved to the ground terminals, which have better information regarding their incoming traffic and can have more sophisticated hardware. The scheme has several drawbacks, however. First, the scheme uses buffer size rather than packet arrival rate to detect congestion. Buffer size is an inaccurate measure of congestion. A system with a large but decreasing buffer size is in a more favourable situation than one with a small but rapidly increasing buffer size. Next, the scheme can be unfair to low-rate terminals. In the event of congestion, each terminal is required to decrease its rate, even though the rate is already below the fair

value. Further, as the queue size decreases below the first threshold, each terminal is allowed to send again at its original rate, which would bring back the congestion situation. The network will then oscillate between periods of congestion and non-congestion and never converge to the ideal operating point.

Another reactive congestion control scheme for packet-switched satellite networks was described in [54]. The scheme correctly uses the rate of arrival and departure of jitter-tolerant packets as a measure of congestion. In this scheme, the satellite continuously broadcasts two congestion signalling bits for each of its downlink channels (or buffers), namely a Congestion Indication (CI) bit and a No Increase (NI) bit. Defining $\lambda(k)$ as the estimated (weighted moving averaged) packet arrival rate at frame k , $\mu(k)$ as the estimated (weighted moving averaged) packet departure rate at frame k , and β as the overload protection factor, the satellite sets the CI bit of a given downlink channel when the condition $\lambda(k) > \mu(k)$ holds, or sets the NI bit of a given downlink channel when the condition $\beta\mu(k) < \lambda(k) \leq \mu(k)$ holds. Hence, at any time, at most one signaling bit may be set. Ground terminals adapt their rates to each of the satellite's downlink channels in response to these status bits. A ground terminal may increase its rate if neither the CI bit nor the NI bit is set, should maintain its rate if the NI bit is set, and must decrease its rate if the CI bit is set. Further, the scheme dictates that successive rate increases should not be too close (to prevent congestion), while successive rate decreases can be close to each other to quickly avoid network congestion. Ground terminals are responsible for determining the rate of increase or decrease. Defining γ_{INC} , γ_{DEC} , and $R(k)$ as the calculated rate of increase, the calculated rate of decrease, and the rate at frame k , respectively, the rate increment equation is given by

$$R(k+1) = \gamma_{INC} \cdot R(k) + (1 - \gamma_{INC}) \cdot PCR \quad (15)$$

while the rate decrement equation can be expressed as

$$R(k+1) = \gamma_{DEC} \cdot R(k) + (1 - \gamma_{DEC}) \cdot MCR \quad (16)$$

As a result, the closer the current rate to the PCR the slower the rate of increase is, while the closer the current rate to the MCR the slower the rate of decrease will be. Next, γ_{INC} is computed such that repetitive execution of (15) yields $R(N_{INC}) = f_{INC} \cdot PCR$ starting with $R(0) = MCR$ and f_{INC} close to but smaller than one. Similarly, γ_{DEC} is calculated such that repetitive execution of (16) yields $R(N_{DEC}) = f_{DEC} \cdot MCR$ for $MCR > 0$, or $R(N_{DEC}) = \varepsilon \cdot PCR$ for $MCR = 0$, starting with $R(0) = PCR$, f_{DEC} close to but smaller than one and ε close to but larger than zero. The use of ICR, ACR, MCR, and PCR parameters in the scheme indicates that the scheme may be intended for use in satellite ATM networks. However, the scheme is not compatible with the rate-based congestion control framework set by the ATM Forum [1], since it does not use RM cells for the transport of control signalling and it employs a different rate adaptation algorithm in terminals. Merits of the scheme include low overhead, simple algorithm, correct usage of congestion measure and capability to maintain a negligible CLR. Drawbacks of the scheme include a moderate maximum achievable throughput (approximately 77 %) and an unfavourable trade-off between the system's responsiveness and the desired CLR. This is so as N_{INC} should be large in order to maintain a small CLR, and as a result, it would take a long time for a ground terminal to reach its desired transmission rate (PCR), even though the network is underloaded.

6.2 Congestion Control Schemes for ATM Networks

Among the five ATM service classes, ABR is the only service class required by the ATM Forum to use some form of flow control [1]. This requirement is a consequence of the fact that ABR service is targeted to use as much of the available bandwidth as possible, while maintaining a low CLR. There has been a considerable debate on the type of congestion control scheme appropriate for ATM networks. Most of the debate, however, centers around two flow control approaches, namely the rate-based approach [44] and the credit-based approach [45].

The rate-based approach controls the rate at which a source may transmit its traffic. Generally, a source is allowed to increase its rate when the network is underloaded, and vice versa. It is the responsibility of the switch and the destination to detect and notify the source of a potential congestion situation. A source periodically transmits resource management (RM) cells to its destination, containing information on its current (allowed) cell rate and the desired (peak) cell rate. Each upstream switch continuously monitors its status and uses some fields in RM cells to inform the source about the congestion situation. When these RM cells reach the destination, they will be returned back to the source. Finally, based on the information contained within the RM cell, the source can determine how it should adjust its rate. The strengths of the rate-based approach include inexpensive hardware requirement, scalability (i.e., it can work in a wide range of network speed and configurations) and implementability (i.e., it does not require a specific hardware implementation). The weaknesses of the approach include the non-zero probability, though small, of cell loss (due to congestion) and the inability to handle bursty traffic efficiently [45].

The credit-based approach, on the other hand, controls the amount of traffic a source can transmit. In this approach, a source is given a fixed (static) or variable (adaptive) value of buffer allocation, or credit. The source maintains the balanced of unused credit, which is the difference between the value of buffer allocation and the number of packets which have been transmitted but not yet forwarded by the upstream switch. This balance is the number of packets which the source is allowed to send. The advantages of this approach include a simple algorithm, the capability to guarantee that no cell is lost due to congestion, and its suitability for bursty traffic. Its weaknesses include expensive hardware requirement (due to the need of per-VC queueing) and the rather rigid need for each switch to use per-VC queueing with round-robin service [43].

There was an attempt to integrate the two approaches in order to combine their merits [46]. In [46] the authors argued that rate-based approach is more suitable for wide area network (WAN), since the use of credit-based approach will require a large amount of buffering in order to achieve high throughput. On the other hand, the credit-based scheme is more suitable for local area network (LAN), due to its ability to guarantee zero cell loss due to congestion. Further, the number of connections in a LAN will be much less than in a WAN, so that each switch does not have to maintain a large number of queues. In the end, however, the ATM Forum decided to adopt the rate-based approach, mainly because of its simple hardware requirement.

6.2.1 The Rate-Based Congestion Control Framework for ATM Networks

The complete description of the rate-based congestion control framework for ATM networks is given in [1]. In this section, only a concise description is provided. The rate-based control scheme is of closed loop type and is performed per connection. It is based on the use of RM cells to carry upstream congestion information to each ABR

source. The source periodically generates and transmits forward RM cells. The contents of these RM cells can be examined and modified by downstream switches. The destination turns around RM cells and sends them back to the source. The source is then able to adjust its transmission rate by examining the congestion information within these backward RM cells.

An ABR source has four rate parameters, namely the Allowed Cell Rate (ACR), Initial Cell Rate (ICR), Minimum Cell Rate (MCR) and Peak Cell Rate (PCR). The ACR is the rate at which the source is allowed to send (i.e., the allocated rate) and is a variable quantity subject to congestion situation and the source's MCR and PCR values. The source's ACR is initially set equal to its ICR. Then, in response to congestion information carried by the RM cells, the source's allocated rate may increase or decrease. The value of the source's ACR, however, is lower bounded by its MCR and upper bounded by its PCR.

There is a set of rules governing the frequency of RM cell generation. In a normal situation, RM cells are generated in proportion to the transmission rate (i.e., the ACR). The NRM parameter specifies the maximum number of cells the source may transmit for each forward RM cell (i.e., the source may transmit up to $\text{NRM} - 1$ data cells before it generates a new RM cell). Nevertheless, when the rate is very small, a source may not generate a new RM cell for a long time. As a result, the congestion-related information will not be up-to-date. An upper bound is therefore established on the time between forward RM cells, specified in the TRM parameter. Further, in order to control the amount of overhead due to RM cell generation, there is a fixed minimum number of data cells between subsequent forward RM cells specified by the MRM parameter. Finally, a source is allowed to send out-of-rate RM cells when it deems necessary for, as an example, obtaining an up-to-date congestion status. These RM cells will not be included in

the count of cells transmitted in-rate, hence, they are called out-of-rate cells. However, since they do contribute to the overall control overhead, the frequency of out-of-rate RM cell generation is limited to the Tagged Cell Rate (TCR) parameter, which is a constant value of 10.

Another set of rules are established to control how each ABR source should adjust its rate on receiving backward RM cells or on encountering lost or delayed RM cells. In the event a source is not able to send a new RM cell due to a low transmission rate or capacity unavailability, the source has to reduce its rate to its initial cell rate (ICR). The ACR Decrease Time Factor (ADTF) parameter specifies the maximum time permitted between subsequent sending of RM cells before the source's rate is decreased to its ICR. Further, RM cells may be dropped in the network due to congestion. The source maintains an account of forward and backward RM cells. If the difference between forward and backward RM cells exceeds the CRM parameter, the source is required to decrease its ACR by a factor specified by the Cutoff Decrease Factor (CDF) parameter.

In order to describe the source's rate adjustment algorithm on receiving a backward RM cell, the functions of several bits and fields in each RM cell will first be explained. The MCR field stores the source's MCR parameter, while the Current Cell Rate (CCR) field stores the source's ACR parameter. Next, the Explicit Rate (ER) field is initially set to the source's PCR, but may be modified by an explicit rate switch to store the maximum rate the switch can allocate for the source. An explicit rate switch is able to calculate the explicit rate it can allocate for each connection. Such a switch usually requires a sophisticated hardware, and as a result, its use may not be widespread. Therefore, each RM cell has two congestion signaling bits which can be used by less advanced relative rate switches to indicate congestion situation, namely the Congestion

Indication (CI) bit and the No Increase (NI) bit. In the event of congestion, a relative rate switch can request a source to reduce its rate by setting the CI bit in the source's RM cells. The switch may also prevent a source from increasing its rate by setting the NI bit, usually in the situation where congestion is imminent. Nevertheless, due to the high-speed operation, some switches may not be able to support even relative rate marking. Fortunately, the ATM Forum establishes a minimum requirement for an ATM switch, in that it must support Explicit Forward Congestion Indication (EFCI) marking. In the event of congestion, such a switch will set the EFCI bit in the data cell header. The destination terminal is then required to update the CI bit based on the EFCI bit status of the previous data cell received from the source.

When the source receives a backward RM cell, it inspects the RM cell's CI bit, NI bit, and ER field. If the CI bit is set, indicating congestion, the source reduces its rate by a factor specified in the Rate Decrease Factor (RDF) parameter according to the relation :

$$ACR = ACR - ACR \cdot RDF \quad (17)$$

Next, if the NI bit is not set, indicating no imminent congestion, the source may increase its rate by a factor specified in the Rate Increase Factor (RIF) parameter according to the relation :

$$ACR = ACR + RIF \cdot PCR \quad (18)$$

The source's final ACR value is then lower bounded by its MCR and upper bounded by the minimum of its PCR and ER field value. It may be useful to note that (17) specifies that the rate of decrease is exponential, while (18) states that the rate of increase is linear. This combination of exponential rate of decrease and linear rate of increase has been shown to lead to a fair control scheme [57]. The concept of fairness, therefore, is incorporated in the rate-based control framework.

Finally, the rate-based control framework also provides a provision for segmenting the network into several individually controlled ABR segments. This would expedite the delivery of RM cells to the originating sources, in order to increase the overall system's responsiveness. A switch may assume the behaviour of a destination end system (by becoming a virtual destination), and hence it is able to turn around RM cells it receives. The same switch must also assume the behaviour of a source end system (by becoming a virtual source), and generate its own RM cells. This virtual source / virtual destination (VS/VD) configuration may be beneficial in the satellite networks. The hardware complexity required for implementation onboard the satellite, however, may be a major obstacle.

6.2.2 Examples of Rate-Based Schemes

Several rate-based switch algorithms will be described shortly. These algorithms strive to achieve fast convergence of fair rate allocation for each connection, and in many cases, have successfully accomplished the task.

6.2.2.1 Congestion Control with Explicit Rate Indication (CCERI)

The CCERI algorithm [47] was originally developed for a general packet network. At the time of its development, the congestion control framework for ABR service has not yet been standardized. In fact, the algorithm's distributed feature and the approach to the calculation of optimum fair rate allocation had significant influence on the development of newer switch algorithms.

In CCERI algorithm, each source periodically sends a control packet containing its current cell rate (stamped rate). When the control packet arrives at the switch, the switch compares this stamped rate to the switch's advertised rate. When the stamped rate

is above the advertised rate, the switch reduces the stamped rate to the advertised rate and sets the underloading bit (“u-bit”). Otherwise, the switch leaves these fields unmodified. Finally, the switch forwards this control packet to the next upstream switch or to the destination. When the control packet reaches the destination, it will be turned around and sent back to the source. The source will use the stamped rate in the returned control packet as its current (allowed) cell rate. Hence, the control packet has the same function as the RM cell, although its format is still elementary.

The switch maintains a database of all connections along with their last stamped rates. The switch also calculates a quantity called the advertised rate. Connections whose recorded rates are below this advertised rate are marked as restricted or bottlenecked, while those whose rates are above the advertised rate are marked unrestricted. The switch calculates the advertised rate, denoted by μ , as follows :

$$\mu = \frac{C - C_R}{n - n_R} \quad (19)$$

where C, C_R, n , and n_R are the total link capacity, the total capacity used by the bottlenecked connections, the total number of connections and the total number of bottlenecked connections, respectively.

The algorithm is proven to converge to the max-min fairness and the convergence time is upper-bounded to $4 \cdot M \cdot D$, where M is the number of bottlenecked links and D is the round-trip delay. The most serious drawback of the CCERI algorithm is the $O(n)$ computational complexity, since the calculation of new advertised rate requires information from all active connections.

6.2.2.2 Explicit Rate Indication for Congestion Avoidance (ERICA)

Many schemes developed after the CCERI aim to achieve $O(1)$ operation. The ERICA scheme [48] accomplishes this by the periodic calculation of the load factor z as follows :

$$z = \frac{ABR \text{ Input Rate}}{ABR \text{ Capacity}} \quad (20)$$

The switch uses this load factor to compute each connection's VC share :

$$VCShare = \frac{CCR}{z} \quad (21)$$

The switch also computes the fair share by using the relation :

$$FairShare = \frac{ABR \text{ Target Rate}}{N} \quad (22)$$

where N is the number of active connections and *ABR Target Rate* is the target rate of arrival of ABR traffic. When the switch receives a backward RM cell, its ER field will be updated with a new explicit rate value calculated as follows :

$$ER \text{ Calculated} = \left(ER, (FairShare, VCShare)^+ \right) \quad (23)$$

The algorithm has been shown to bring the rates to converge quickly under various situations and is insensitive to minor measurement inaccuracies. Further, even though the basic algorithm was found not to reach the max-min fairness for connections which started late, modifications to the basic algorithm are supposed to fix the problem at the cost of increased complexity. In this case, the switch stores the highest allocation made during one averaging interval and ensures that all sources can also get this allocation. The switch maintains two variables, *MaxAllocPrevious* and *MaxAllocCurrent*, the maximum allocation given in the previous and current measurement interval, respectively. When the

load factor z is bigger than $1 + \delta$, where δ is a small fraction, the ER value is still computed as in (23). Otherwise, the new ER value is computed as follows :

$$ER_{Calculated} = \left(ER, (FairShare, VCShare, MaxAllocPrevious)^+ \right) \quad (24)$$

Further, during transient periods, the queue size may increase quickly, resulting in high cell loss, if the target utilization is set close to one. In order to avoid these transient overloads, the basic ERICA algorithm is improved by incorporating yet another modification. If the calculated ER value is greater than the source's fair share and the source's CCR is below its fair share, then its ER value will be limited to its fair share, that is :

$$\begin{aligned} & \text{If } ((CCR < Fair\ Share) \text{ and } (ER_{Calculated} \geq Fair\ Share)) \\ & ER_{Calculated} = Fair\ Share \end{aligned} \quad (25)$$

A further enhancement to the basic ERICA algorithm is the incorporation of buffer size monitoring as the secondary control mechanism, in order to achieve 100 % link utilization. The resulting algorithm is called ERICA+. The ERICA+ algorithm attempts to maintain a non-empty ABR queue, so that an ABR cell is always available when higher priority traffic is temporarily absent. This control algorithm uses a target queueing delay T_0 as its operating point. When the estimated ABR queueing delay is less than T_0 , ERICA+ increases the amount of available ABR capacity in order to increase each VC's rate allocation and reach the target queueing delay. Vice versa, if the estimated queueing delay is larger than T_0 , ERICA+ decreases the amount of available ABR capacity to let the queue to drain, thus bringing the queueing delay to target.

6.2.2.3 Max-Min Rate Control Algorithm (MMRCA)

The MMRCA algorithm proposed in [49] uses two queue size thresholds to indicate congestion, namely QT and DQT . When the queue size is above QT , the switch is

considered to be congested. Further, when the queue size is above DQT , the switch is considered to be heavily congested. The switch also maintains the values of maximum (MAX) and the minimum (MIN) rates of all connections sharing the switch along with the corresponding VC number (MAX_VC and MIN_VC). When the queue size is below QT , all connections except MAX_VC are allowed to increase their rates. When the queue size exceeds QT but is still below DQT , the switch selectively indicates congestion to connections whose rates are above MIN . In this case, all connections whose rates are higher than $MIN \cdot IPF$ (IPF being the Increase Pressure Factor) are required to reduce their rates. However, if the difference between the maximum and the minimum rate is smaller than MAX_MIN_DIFF , all connections will be required to reduce their rates. Finally, if the queue size is above DQT , then all connections are instructed to reduce their rates.

Further enhancements to the MMRCA algorithm include the use of either the rate of queue size increase or the total rate of all connections, as an indicator for potential congestion. In the case of the former enhancement, the switch is considered to be in potential congestion if the queue size increases by an amount larger than RQL (RQL being the rate of queue length increase threshold) within a given measurement interval. In the case of the latter enhancement, the switch is considered to be in potential congestion if the total packet arrival rate of all connections is greater than the outgoing link rate. In the case when potential congestion is detected, the switch instructs all connections whose rates are higher than $MAX \cdot DPF$ (DPF being the Downlink Pressure Factor) to decrease their rates.

The MMRCA algorithm has been shown to achieve fast convergence of rate allocation and high link utilization, while requiring a small buffer size. Its enhanced

versions are found to perform even better. The algorithm has some shortcomings, however [50]. In networks with multiple switches, where some connections are bottlenecked, the MMRCA is found to lead to unfairness. This is because the rates of bottlenecked connections may prevent the minimum rate of all connections to converge to the fair share, yet this minimum rate value will still be used in the selective marking leading to the beat-down problem.

6.2.2.4 Dynamic Max Rate Control Algorithm (DMRCA)

The DMRCA algorithm [50] solves the shortcoming of the MMRCA algorithm by using the adjusted maximum rate of all connections, scaled by a function of the queue size at the switch, as the threshold to perform selective marking of connections to be notified of congestion. The reason behind the use of the maximum rate as the threshold for selective marking is that unlike the minimum rate, the maximum rate varies more rapidly and stays above the fair share before convergence. The minimum rate, however, may never converge to the fair share if the terminal is bottlenecked.

Similar to the MMRCA, the DMRCA also uses two queue size thresholds, QT and DQT , to detect congestion. Further, the switch maintains the value of the maximum rate of all connections (MAX) along with the corresponding VC number (MAX_VC). In order to filter out variations, the adjusted maximum rate, denoted by A_MAX , is calculated as the exponential average of the maximum rate, as shown by the following relation :

$$A_MAX = (1 - \alpha) \cdot A_MAX + \alpha \cdot MAX \quad (26)$$

This adjusted maximum rate will then be scaled by a function of the queue size to give the threshold for selective marking (the estimated fair share).

The switch performs selective marking only when the queue size is between QT and DQT , i.e., when the switch is congested. In this case, all connections whose rates are

above the estimated fair share are instructed to reduce their rates. The recommended means to indicate congestion is by setting the CI bit. Any other connections are allowed to increase their rates. When the queue size exceeds DQT , the switch is considered heavily congested. In this case, all connections will be required to reduce their rates and the switch sets the ER field in all RM cells equal to $A_MAX \cdot MRF$, where MRF is the Major Reduction Function. Finally, when the switch is not congested, all connections are allowed to increase their rates.

The DMRCA is shown to achieve fairness in various situations and network configurations and to be insensitive to the mistuning of parameters [50]. Compared to the ERICA algorithm, however, the DMRCA algorithm appears to require a larger buffer size and does not provide a mechanism to maintain a stable operation under full link utilization.

6.2.2.5 Fast Max-Min Rate Allocation (FMMRA)

The FMMRA algorithm [51] achieves the computational complexity of $O(1)$ by performing incremental update of the switch's advertised rate. The advertised rate in link l , denoted as γ_l can be expressed simply as follows :

$$\gamma_l = \frac{C_l^A - \bar{C}_l}{N_l - \bar{N}_l} \quad (27)$$

where C_l^A , $\bar{C}_l$, N_l , and $\bar{N}_l$ are the total capacity available for ABR traffic, the total capacity consumed by bottlenecked connections, the total number of connections and the total number of bottlenecked connections, respectively. Besides the above variables, the switch also maintains two variables for each connection i , namely its bottleneck indicator, β_l^i , and its bandwidth λ_l^i .

When the switch receives a backward RM cell, it uses the RM cell's ER value, denoted as λ_i^{ER} , to determine the change in the bottleneck status, $\Delta\beta$, and the bottleneck bandwidth, $\Delta\lambda$, as follows :

$$\Delta\lambda = \begin{cases} \lambda_i^{ER} - \lambda_l^i \beta_l^i, & \lambda_i^{ER} < \gamma_l \\ -\lambda_l^i \beta_l^i, & \lambda_i^{ER} \geq \gamma_l \end{cases} \quad (28)$$

$$\Delta\beta = \begin{cases} 1 - \beta_l^i, & \lambda_i^{ER} < \gamma_l \\ -\beta_l^i, & \lambda_i^{ER} \geq \gamma_l \end{cases} \quad (29)$$

Using these changes in bottleneck status and bandwidth, the new advertised rate is then calculated as follows :

$$\gamma_l = \begin{cases} C_l^A, & N_l = 0 \\ \gamma_l + \frac{\gamma_l \Delta\beta - \Delta\lambda}{N_l - (\bar{N}_l + \Delta\beta)}, & N_l > \bar{N}_l \\ \gamma_l, & N_l = \bar{N}_l \end{cases} \quad (30)$$

Finally, the connection's bottleneck status and bottleneck bandwidth are updated according to the following relations :

$$\lambda_l^i = \begin{cases} \lambda_i^{ER}, & \lambda_i^{ER} < \gamma_l \\ 0, & \lambda_i^{ER} \geq \gamma_l \end{cases} \quad (31)$$

$$\beta_l^i = \begin{cases} 1, & \lambda_i^{ER} < \gamma_l \\ 0, & \lambda_i^{ER} \geq \gamma_l \end{cases} \quad (32)$$

A unique feature of the FMMRA algorithm is its bi-directional ER marking, that is, ER marking is performed on both forward and backward RM cells. This feature significantly improves the speed of convergence to the max-min fair allocation. The reason is that by bi-directional marking, upstream switches are able to learn the bottleneck status of downstream switches through forward RM cells, and vice versa. Hence, a switch is able

to quickly gather information about other switches. The ER field itself is updated as follows :

$$\lambda_i^{ER} = \left(\gamma_l, \lambda_i^{ER} \right)^- \quad (33)$$

A further enhancement to the basic FMMRA algorithm can be achieved by computing the ABR load factor ρ_l , which is the ratio of the arrival rate of ABR traffic and the available bandwidth for ABR traffic, and by using two thresholds for the queue size, QT and DQT . Here, when the queue size is below QT , i.e., the switch is not congested, the ER field is updated as follows :

$$\lambda_i^{ER} = \left(\left(\frac{\hat{\lambda}_l^i (1 - \beta_l^i)}{\rho_l} \right)^+, \gamma_l \right)^-, \lambda_i^{ER} \quad (34)$$

When the queue size is between QT and DQT , i.e., the switch is congested, the ER field is updated according to (33). Finally, when the queue size exceeds DQT , the switch is considered to be heavily congested and the available bandwidth for ABR traffic is reduced according to a function of queue length, to allow the queue to drain.

The performance of the FMMRA algorithm is shown to be better than the performance of ERICA in terms of the speed of convergence [51]. The FMMRA is also able to avoid the fairness problem exhibited in the basic ERICA algorithm. Finally, the maximum link utilization permitted by the FMMRA is comparable to that by the ERICA+.

6.3 Congestion Control Scheme for Satellite ATM Networks

The long propagation delay inherent in satellite networks limits the effectiveness of any congestion control schemes. Any adopted scheme should therefore be very responsive to changing network conditions. This requirement of fast response has been early indicated

in [4] and it essentially means that the appropriate control scheme should be rate-based. Further, due to the tremendous improvement in system's responsiveness and link utilization permitted by explicit rate flow control for wide area network (such as the Internet) [60], the preferred congestion control scheme should be based on explicit rate notification.

With respect to satellite ATM networks, some questions remain regarding the need for any modifications to the adopted rate-based scheme, especially considering the existence of a rigid frame structure and the significant influence of multiple access protocols to the performance of such networks. Several papers have indicated the required modifications with respect to some ABR parameters [56] or the need for an integrated MAC and congestion control scheme [55], in order to achieve efficient utilization of uplink bandwidth. In the next chapter, the topic of an integrated MAC and ABR rate control scheme will be explored further. Several modifications to the rate-based congestion control approach will be described. The adaptations of the popular ERICA switch algorithm to MAC-related information stored in the network controller will be presented. The objective is to assess the performance of an integrated MAC and rate control scheme for ABR service. Finally, since terminal population in a satellite network is potentially very large, some satellite networks may not be based on ATM. A novel rate-based congestion control scheme for implementation in non-ATM satellite networks will therefore be described. Nevertheless, the new scheme can be implemented in satellite ATM networks by segmenting the networks and applying the scheme in the access (uplink) channel. Performance improvement offered by this novel scheme compared to the integrated MAC-ERICA scheme will also be presented.

Chapter 7

Integration of Multiple Access and Rate Control Schemes for ABR Service

Due to the long propagation delay, the multiple access scheme employed in a satellite network has a significant influence to the network's overall performance. Therefore, in developing a congestion control scheme for such a network, one needs to consider the access scheme implemented, or intended to be implemented, in the network. This chapter describes some modifications to the ABR rate control framework and to the switch algorithm. The ERICA scheme is selected from the many well-developed switch algorithms due to its popularity and modest complexity. Further, in order to show the potential performance improvement offered by a scheme that takes advantage of the broadcast capability of satellites, a new rate control scheme will be described.

The proposed modifications to the ABR rate-based control framework and the switch algorithm are based on the following suggestions :

- Establish control of transmission (sending) rate by controlling the number of time slots to request per frame.
- Enforce fairness in resource allocation for both uplink and downlink channels.

The first suggestion is intended to avoid waste of capacity without increasing the complexity of the uplink scheduler's assignment process. To explain the potential waste of

bandwidth when control is exercised on the sending rate, consider the situation where an earth terminal's request of x slots is granted by the scheduler, the corresponding assignment message has arrived to the earth terminal and the terminal is ready to use its allocated time slots in the current frame. If the terminal's allocated rate is low (due to congestion), it may not be able to use one or more of its allocated time slots. Since other terminals cannot use these slots for their traffic, even though it is destined to a non-congested downlink channel, transmission capacity is wasted. In order to avoid this potential problem, control should be enforced with respect to the amount of time slots a terminal can request in a given frame. The terminal should then be able to utilize all time slots assigned to it, as a result of its requests. This implies that an ABR source should not be eligible for free capacity assignment, since such assignment is not based on the terminal's requests.

The second suggestion is the consequence of our understanding of a communication satellite as a queueing system with two servers [3]. The first server is the uplink scheduler, with service rate equal to the uplink capacity. Traffic buffering for this server is performed in ground terminals. Since memory is relatively inexpensive, the terminal's buffer can be assumed to have infinite capacity. The second server is the downlink scheduler. A downlink scheduler services several downlink queues (including the ABR queue) of an output port at a rate equal to the downlink channel capacity. Since any additional hardware occupies space in and adds weight to the satellite, the downlink queue has a limited capacity, though a possibly large one, to compensate for the long propagation delay. Based on this two-server view, it is proposed to enforce fairness in rate allocation in both uplink and downlink channels. Additional control related to the queue size, such as that employed in ERICA+, may be enforced in the downlink channel.

The first suggestion leads to modifications to the scheduling of ABR traffic in ground terminals, while the second suggestion has as a result more processing in the switch. The next section describes modifications to the transmission scheduling in more detail. Modifications to the ERICA switch algorithm and the description of the broadcast rate control scheme will be presented in later sections.

7.1 Transmission Scheduling of ABR Traffic in Ground Terminals

Earth terminals receive traffic from various ATM traffic classes and from several users. Following the association of service classes to R/A strategies described in Chapter 5, traffic arriving from other than CBR sources is queued in the terminal until capacity is available (CBR traffic is not queued since capacity is reserved at the source's PCR). A scheduler (the uplink access process) located in each terminal is responsible for managing various queues. The management responsibility includes the calculation (and scheduling) of requests to empty each of the queues, the decomposition (breakdown) of capacity assignment (the assignment message only specifies the total number of slots a terminal may utilize in the next frame, not the number of slots for each of its service classes), and the final transmission of traffic. See Chapter 4 for a detailed description of the terminal processing related to the access scheme. Further, in order to support ABR service, the scheduler is also responsible for the calculation (or adaptation) of allowed cell transmission rates.

Before the scheduling of ABR traffic transmission can be described, it is necessary to explain the followed approach to congestion control problem (for ABR traffic) onboard a satellite, since it affects the processing within the ground terminal. The satellite's downlink scheduler places a higher priority for satellite's control messages and high

priority traffic (which include CBR, rt-VBR and nrt-VBR traffic). Hence, ABR traffic needs to be queued until the downlink scheduler has finished servicing these two types of traffic. Ideally, ABR traffic should be queued per-VC onboard the satellite and correspondingly, the downlink transmission should be scheduled per-VC. Per-VC queueing, however, requires a more sophisticated onboard hardware, and in addition, adds mass to the satellite. Due to constraints imposed on the satellite hardware, it is assumed that ABR traffic is queued per downlink destination (downlink channel). Consequently, performing per-VC queueing in terminals, although feasible, may not bring much benefit. Congestion control in ground terminals is therefore performed per downlink channel. Based on the above assumption, the queueing structure in each earth terminal is illustrated in Figure 34.

Figure 34. Queueing Structure in Ground Terminals

Since traffic from more than one VC may share the same queue, the sum of each VC's MCR, ICR, ACR and PCR is used as MCR, ICR, ACR and PCR, respectively, of

the aggregate traffic in the queue. The scheduler uses the aggregate ACR to request uplink capacity to a certain downlink. Note that each request is made on behalf of all VCs sharing the queue. Therefore, in order to meet traffic contracts for all of the terminal's VCs, it is necessary that each VC traffic be policed to meet its individual traffic contract (by using the leaky-bucket algorithm for example). Otherwise, non-conforming traffic from one VC will affect that from other VCs. For the sake of simplicity, it is assumed that each ABR queue contains traffic from only one VC, so that an ABR source is equivalent to a terminal's ABR traffic generator to a certain destination.

As indicated earlier, the task of transmission scheduling includes the calculation and scheduling of capacity requests, the decomposition (breakdown) of capacity assignment, the final transmission of traffic and the calculation (adaptation) of ACR values. The first three tasks are described more elaborately below. The algorithm for ACR adaptation (in ground terminals) for the ERICA scheme is described in [1], while that for the broadcast scheme will be described in a later section.

7.1.1 Capacity Request Calculation Algorithm for ABR Traffic

The capacity request calculation algorithm for ABR traffic assumes that ABR service is associated with a combination of RC and VBDC-HP strategies. Uplink capacity is reserved for an ABR source at its MCR per frame, i.e., $R = \lceil MCR \cdot T \rceil$ slots per frame are reserved using the RC strategy, where T is the duration of a frame and $\lceil x \rceil$ is the integer ceiling of x . Any other request for capacity above the source's MCR will be signaled using the VBDC-HP strategy.

The scheduler uses the allocated rate to request capacity for its ABR traffic. The maximum number of time slots, n_{Max} , it may request in each frame for a given destination

depends on its allocated rate for the corresponding destination and can be calculated as follows :

$$n_{Max} = \lfloor \eta + ACR \cdot T \rfloor \quad (35)$$

where $\lfloor x \rfloor$ is the integer floor of x . In the above expression η is the fractional part of n_{Max} obtained in the previous frame by using the following relation :

$$\eta = (\eta + ACR \cdot T) - \lfloor \eta + ACR \cdot T \rfloor \quad (36)$$

The actual amount of uplink capacity to be requested is determined by the current queue size and the amount of capacity that has been requested in previous frames. The calculation is performed such that capacity is not wasted by requesting more than what is needed. The total number of time slots which have been requested for a given downlink channel but not yet granted, *TotalRequestedCapacity*, is increased every time a request is made or capacity is reserved, and decreased every time an ABR cell is sent to the destination. If the MCR for a given downlink channel is greater than zero, the scheduler determines how many of the reserved slots it may utilize by the following calculation :

$$ReserveCapacityToUtilize = (CurrentQueueSize - TotalRequestedCapacity, R, n_{Max})^- \quad (37)$$

TotalRequestedCapacity and n_{Max} are then updated as follows :

$$TotalRequestedCapacity = TotalRequestedCapacity + ReserveCapacityToUtilize \quad (38)$$

$$n_{Max} = n_{Max} - ReserveCapacityToUtilize \quad (39)$$

Next, the scheduler calculates how many slots it should request for a given downlink channel using the following expression :

$$CapacityToRequest = (CurrentQueueSize - TotalRequestedCapacity, n_{Max})^- \quad (40)$$

and similarly, *TotalRequestedCapacity* will be updated as follows :

$$TotalRequestedCapacity = TotalRequestedCapacity + CapacityToRequest \quad (41)$$

At this point, the unused capacity (n_{Max}) will be lost (i.e., a source cannot accumulate capacity which it does not use to request uplink time slots). Finally, the total number of slots to be requested using the VBDC-HP strategy is the sum of *CapacityToRequest* for all downlink channels the terminal has traffic to send to. Figure 35 summarizes the terminal algorithm for capacity request calculation in the form of flow charts.

Figure 35. Flow Charts of Terminal's Capacity Request Calculation Algorithm

7.1.2 Capacity Assignment Decomposition

The assignment message contains the number of time slots a terminal may utilize for the current frame and the starting slot location in the MF-TDMA frame (see Chapter 4). It does not specify the composition of the assigned time slots (i.e., the number of slots

for each service class). It is the scheduler responsibility to determine the number of time slots for ABR traffic.

The scheduler uses the same method as that described in Section 4.4.2.2 to determine how many time slots in a given assignment message belong to ABR traffic. The scheduler first sets aside capacity reserved or requested for CBR, rt-VBR and nrt-VBR traffic. The remaining capacity is therefore due to VBDC and FCA strategies.

Recall that a copy of each transmitted VBDC-HP request is inserted into the pending VBDC-HP request queue and is given a time-out counter. The counter value is initialized to $V + W$ (where V is the sum of round-trip delay and processing delay and W is the time-out duration of VBDC requests) and is decremented in every subsequent frame. When the counter value is zero, the request is removed from the pending request queue. By inspecting the counter value of the oldest (top-most) request, the scheduler is able determine the number of time slots for ABR traffic. If the request's counter value has been reduced to less than or equal to W (i.e., capacity assignment due to the queued request should have arrived, if granted), the scheduler assumes that it has capacity assignment for ABR traffic. The number of time slots for ABR traffic is the smaller of the requested capacity and the remaining capacity assignment.

7.1.3 Transmission of ABR Traffic

The scheduler distributes its ABR capacity assignment equally among all downlink channels (destinations) it has ABR traffic to send. In order to do this, it chooses a destination in a rotating manner (by arranging its traffic destinations in a circular linked-list and maintaining a record of the last chosen destination). Each time a slot for ABR traffic is due, the scheduler chooses the next destination in the linked-list whose *TotalRequestedCapacity* is greater than zero (i.e., it chooses a destination for which the

terminal has requested uplink capacity) as the destination to which it will send traffic. Figure 36 summarizes the terminal algorithm for the transmission of ABR traffic in the form of flow charts.

Figure 36. Flow Charts of Terminal Algorithm for Transmission of ABR Traffic

7.2 The Modified ERICA Switch Algorithm

From an implementation point of view, the ERICA switch algorithm is separated into two components : the first is performed in the satellite's input and output ports and the other is performed in the network controller. The reason behind this separation is that a satellite switch, unlike terrestrial switches, may not be able to perform all the measurements, not to mention the computation tasks associated with the ERICA algorithm. Due to a greater aggregation of traffic in the switch and a greater resource sharing in a satellite network, the switch may have to handle millions of VCs. However, a much smaller number of VCs shares each uplink carrier group, hence making the task of per-VC accounting in each uplink carrier group feasible. Further, being a component with possibly the most computing power, the network controller is the best place to conduct computation tasks. Hence, the port algorithm part of ERICA is mainly performing various congestion control related measurements, while the network controller algorithm part of ERICA (or ERICA+) is responsible for the task of computing the fair share (based on port measurement results) and updating the ER field in each RM cell. Note that the task of updating the ER field in RM cells is handled by the network controller. This implies that each RM cell has to be redirected to the network controller, which will then forward it to its destination or return it back to its source.

Due to this separation, many enhancements incorporated in ERICA [48], such as the exponential averaging of active VCs and load factor, have to be performed differently if the network controller is not located onboard the satellite. Otherwise, the amount of signaling between the satellite and the network controller will be tremendous. It is therefore assumed in the thesis that the network controller is located onboard the satellite.

As a result, many of the ERICA's enhancements can be incorporated without any modification.

The only required modification with respect to the original ERICA algorithm is in the calculation of ABR input rate in input ports. Originally, the ABR input rate was calculated as the total number of ABR cells seen in a measurement interval. However, this calculation will always result in a load factor less than 1.0, since the maximum number of ABR cells which can arrive in each input port in a frame is equal to the uplink capacity per frame. The modified method of calculating the ABR input rate in an input port is as follows :

$$ABR\ Input\ Rate = \frac{TotalInPortHPVBDCCount}{m \cdot T} \quad (42)$$

where *TotalInPortHPVBDCCount* is the number of time slots requested using the VBDC-HP strategy, *m* is the measurement interval expressed as an integer number of frames and *T* is as previously defined. This modified calculation method has another benefit : by using the total number of capacity requested for ABR traffic instead of the rate of arrival, unfairness situation in an uplink channel can be detected earlier and more accurately.

7.3 Broadcast Rate Control Allocation for Onboard Processing

Satellites

The Broadcast Rate Control Allocation (BRCA) scheme strives to reach fairness situation in the uplink channels, and fairness and optimum loading situation in the downlink channels. The optimum loading situation is the "knee" in the throughput-delay curve as indicated in [58], i.e., the point at which queueing delay starts to increase rapidly. In this scheme a satellite broadcasts the fair rates at which a terminal (the scheduler) can

request capacity, and eventually can transmit traffic, for each of its downlink destinations. In addition to the fair rates, the satellite also broadcasts the loading condition of its input and output ports, so that each terminal can increase or decrease its rates accordingly.

The BRCA algorithm uses the arrival rate of high priority traffic, rate reserved for ABR traffic and rate requested for ABR traffic to detect congestion in the uplink channels (input ports), while in the downlink channels, it uses the arrival rate of both high priority and ABR traffic. The algorithm sets a target load for the input and output ports. It then controls the arrival rate of ABR traffic so that the target load in the input ports is not exceeded and the target load in the output ports is achieved. In addition, as in ERICA+, BRCA employs downlink queue size as a secondary mechanism to avoid congestion. Unlike ERICA, however, BRCA does not attempt to maintain a stable downlink queue size. Rather, it only attempts to maintain a queue size below a certain value. This is performed for the following reasons. In a satellite, uplink capacity is usually matched with downlink capacity. In fact, downlink capacity will be set slightly larger, to account for assignment and other control messages. Therefore, an attempt to bring the size of a downlink queue close to the target value may introduce unfairness during the process. Fair rates for some downlink channels will be set larger than their fair values, to allow their queues to increase. Earth terminals will respond by requesting more time slots for traffic destined to these downlink channels. Correspondingly, the uplink scheduler will assign more time slots for traffic destined to these downlink channels and suppress or delay the assignment for traffic destined to other channels. Hence, an unfair situation is created.

Due to hardware and software limitation onboard the satellite, the complete BRCA algorithm is divided into three sub algorithms : the port, the onboard network controller (OBNC) and the terminal algorithms. The port algorithm is mainly composed of the

various congestion control-related measurements performed by the input or output ports. The OBNC algorithm involves the calculations of fair rates based on port measurement results, as the OBNC is the component with the most processing power. The terminal algorithm controls how each terminal should react to congestion control messages broadcasted by a satellite. Each of these algorithms will be explained in detail below.

7.3.1 The Port Algorithm

The port is responsible for measuring bandwidth requests for ABR traffic and cell arrival rates, and for monitoring downlink queue size. Most of the measurement tasks are handled by the input port, since performing the tasks in the output port requires more intensive hardware and software processing, due to the greater aggregation of traffic. Furthermore, one can take advantage of some multiple access-related processing which have already been performed, such as extraction of requests from OBR or IBR slots and identification of slot owner terminal ID. Measurements are performed periodically, with each measurement interval corresponding to a fixed integer number of frames. At the end of every measurement interval, the port sends the measurement results to the network controller.

Measurements are performed per uplink carrier group. Each earth terminal belongs to one carrier group and the terminal may transmit its traffic on any of its carrier group's uplink carriers. More than one uplink carrier group may exist in each input port. However, for the sake of simplicity, for the rest of this description it is assumed that each input port contains only one carrier group. Hence, in effect, an input port is equivalent to an uplink carrier group, and each reference to an input port can be substituted by a reference to an uplink carrier group, or vice versa. Similarly, an output port may contain more than one downlink carrier (or channel or destination), but, it is assumed that each output port

consists of only one downlink carrier, so that an output port is equivalent to a downlink carrier. Each reference to output port can therefore be substituted by a reference to a downlink carrier, and vice versa.

The port algorithm assumes that ABR traffic class is mapped to a combination of RC and VBDC-HP strategies. An ABR source with MCR greater than zero reserves capacity on the uplink corresponding to its MCR. Any other request for capacity above the source's MCR will be signaled using the high priority VBDC strategy.

In each frame the input port measures the total capacity requested for ABR traffic and the high priority (CBR, rt-VBR, and nrt-VBR) and ABR cell rates to each of the satellite's output ports. Further, in order to handle ABR sources with non-zero MCR, the input port also measures the total number of reserved capacity to each output port. Finally, the input port measures the total number of ABR sources (or terminal weights) sending traffic to each of the output ports. Each terminal is given a weight for each of its ABR traffic (downlink) destinations which relates to the number of VCs established in the terminal for a given ABR traffic destination. In counting the number of ABR sources, the input port also takes into account sources which have not sent any traffic but have sent their requests through OBR or IBR subslots. This approach gives a better estimate of the number of active sources sharing an uplink channel, since it accounts for the contribution of terminals which use their free capacity assignments to send their requests. The tasks of an output port is much simpler than the input port. The output port is in charge of monitoring its queue size. The complete port algorithm is described in the flowcharts given in Figure 37 and 38. Definitions of the variables used in these flow charts are self-explanatory.

Figure 37. Flow Charts for BRCA's Port Algorithm

Figure 38. Flow Charts for BRCA's Port Algorithm (cont'd)

7.3.2 The Onboard Network Controller (OBNC) Algorithm

The OBNC receives and queues congestion control information sent by all ports. Once information from all ports has been received, the OBNC starts processing this information and finally broadcasts congestion control messages to earth terminals. A set of congestion control messages will be sent to each carrier group in every measurement interval. Congestion control messages sent to a carrier group contain the fair rate per frame, i.e., the rate at which every terminal in the carrier group may request capacity for each of the satellite's downlink destinations, the loading condition in each downlink destination and the corresponding amount of overload (during overload condition) or underload (during underload condition). Earth terminals listen to their downlink channels, extract control messages as they arrive and act accordingly.

7.3.2.1 The Consolidation Process

The information sent by all ports needs to be consolidated to obtain the total number of high priority and ABR cells, the total number of slots reserved by ABR sources with non-zero MCR, and the total number of ABR sources for each input and output ports. The resulting totals may have high variances, therefore they are averaged using the well-known exponential averaging as the following :

$$\bar{X} = (1 - \alpha) \cdot \bar{X} + \alpha \cdot X$$

where $\bar{X}$ and X are the average and current values respectively, and α determines the contribution of the current value to the average value. The consolidation process algorithm is described in Figure 39.

```

For each input port (src) do
  Initialize TotalInPortHPCellCount (src)
  Initialize TotalInPortMCRCount (src)
  Initialize TotalInPortABRSources (src)
End For

For each output port (dest) do
  Initialize TotalOutPortHPCellCount (dest)
  Initialize TotalOutPortABRCellCount (dest)
  Initialize TotalOutPortMCRCount (dest)
  Initialize TotalOutPortABRSources (dest)
End For

For each input port (src) do
  For each output port (dest) do
    TotalInPortHPCellCount (src) = TotalInPortHPCellCount (src) + TotalHPCellCount (src, dest)
    TotalInPortMCRCount (src) = TotalInPortMCRCount (src) + TotalMCRCount (src, dest)
    TotalInPortABRSources (src) = TotalInPortABRSources (src) + TotalABRSources (src, dest)

    TotalOutPortHPCellCount (dest) = TotalOutPortHPCellCount (dest) + TotalHPCellCount (src, dest)
    TotalOutPortABRCellCount (dest) = TotalOutPortABRCellCount (dest) + TotalABRCellCount (src,
dest)
    TotalOutPortMCRCount (dest) = TotalOutPortMCRCount (dest) + TotalMCRCount (src, dest)
    TotalOutPortABRSources (dest) = TotalOutPortABRSources (dest) + TotalABRSources (src, dest)
  End For
End For

```

Figure 39. The Consolidation Process Algorithm

7.3.2.2 The OBNC Algorithm for Input Ports

As previously stated, the BRCA algorithm sets a target load for each input port, so that the target input port capacity can be expressed as follows :

$$InPortCapacity = \rho_u \cdot m \cdot UplinkCapacity \quad (43)$$

where ρ_u is the target input port (uplink) load and m is as previously defined.

The fair rate for requests per frame for each input port can be calculated as follows:

$$CurrentInPortFairRate = \left(\frac{InPortCapacity - AvgInPortHPCellCount - AvgInPortMCRCount}{AvgInPortABRSources}, 0 \right)^+ \quad (44)$$

where $(a,b)^+$ means the maximum of a and b . Notice the inclusion of the average number of MCRs in the calculation. This means that the broadcast fair rate is the rate at which a terminal may request capacity using the VBDC-HP strategy, on top of the terminal's reserved capacity. That is, the actual fair rate at which a terminal may eventually send its traffic is the sum of the terminal's MCR and the broadcast fair rate.

The current load in each input port can be calculated as follows :

$$CurrentInPortLoad = \frac{AvgInPortHPCellCount + AvgInPortMCRCount + AvgInPortHPVBDCCount}{InPortCapacity} \quad (45)$$

An input port is in overload condition if $CurrentInPortLoad > 1.0$ and in underload condition if $CurrentInPortLoad < 1.0$. Note that the average number of reserved capacity and the average number of slots requested using the VBDC-HP strategy is used instead of the average number of ABR cells. This is performed for the same two reasons as in the ERICA case. First, if the average number of ABR cells is used, then $CurrentInPortLoad$ will always be less than 1.0, since the maximum number of ABR cells which can arrive in each input port per frame is equal to the input port capacity (per frame). Second, by using the number of reserved and requested time slots instead of the rate of arrival, unfairness situation in an input port (uplink channel) can be detected earlier and more accurately.

When the load of an input port exceeds its target load, the algorithm recognizes this condition as leading to a potential unfairness situation and instructs earth terminals to decrease their rates. The amount of decrease requested by the network controller is determined by the method of rate decrease employed in earth terminals. In BRCA algorithm, exponential rate decrease in earth terminals is used since [57] suggested that such decrease will bring the system closer to fairness condition in every step. Therefore,

the amount of decrease is calculated such that

$$CurrentInPortLoad \cdot (1 - \lambda_{Up})^j = 1 \quad (46)$$

where λ_{Up} is the amount of rate decrease requested in the input port (uplink channel) and j is $\lceil r/m \rceil$, where r is the round trip propagation delay expressed as an integer number of frames and m is as previously defined. Hence, the amount of decrease is set such that the load in the input port is equal to the target input port load in one round trip delay, which is when the effect of the current rate decrease message will be experienced by the satellite's uplink schedulers (it takes one one-way propagation delay for the message to reach earth terminals, and another one one-way propagation delay for their corresponding requests to arrive at the satellite). The amount of decrease is bounded by a maximum rate decrease λ_{dMax} , so that the system does not become very sensitive to variation in traffic.

On the other hand, since the controller is only concerned with fairness situation in the uplink channels, the controller does not instruct earth terminals to increase their ACR values when the current load of an input port is below its target load. Instead, the decision to increase ACR values will be based on the current load in the output ports, where the algorithm is also concerned with reaching the optimum loading situation, the "knee". Figure 40 summarizes the network controller congestion control algorithm for input ports.

Figure 40. Flow Charts of OBNC Algorithm for Input Ports

7.3.2.3 The OBNC Algorithm for Output Ports

The BRCA algorithm also sets a target load for each output port or downlink channel. Similar to the input port, the target output port capacity can be expressed as follows :

$$OutPortCapacity = \rho_d \cdot m \cdot DownlinkCapacity \quad (47)$$

where ρ_d is the target downlink load and m is as previously used.

Therefore, the current load in each output port is calculated as follows :

$$CurrentOutPortLoad = \frac{AvgOutPortHPCellCount + AvgOutPortABRCellCount}{OutPortCapacity} \quad (48)$$

Further, similar to the ERICA+, BRCA also uses downlink ABR queue size as one of its mechanism to avoid congestion. It sets a maximum delay D_{Max} , for the output port to empty their ABR queues. Available output port capacity for ABR traffic is a function of the output port's service rate and the arrival rate of higher priority traffic (i.e., CBR, rt-VBR and nrt-VBR traffic and messages), and can be expressed by the following relation :

$$ABRCapacity = OutPortCapacity - AvgOutPortHPCellCount \quad (49)$$

Hence, the time required to empty the output port's ABR queue is :

$$D = \frac{ABRQueueSize}{ABRCapacity} \quad (50)$$

When D is larger than the maximum delay D_{Max} , a portion of output port capacity is set aside to allow the queue to drain and D to be reduced to below D_{Max} . BRCA uses a negative exponential function for this intentional introduction of draining capacity. Finally, transmission capacity set aside to drain ABR queue is calculated as follows. First, the current offset of ABR traffic queueing delay relative to the target delay is simply the following :

$$offset = \frac{D - D_{Max}}{D_{Max}} \quad (51)$$

Next, the fraction of output port capacity set aside for draining is computed as follows :

$$DrainFrac = \beta \cdot (1 - e^{-v \cdot offset}) \quad (52)$$

where β is the maximum fraction of output port capacity that can be set aside to drain the

queue and v is the negative exponential function rate parameter which determines how fast *DrainFrac* reaches its maximum value. Hence, the fraction of output port capacity that is left for ABR traffic is clearly :

$$OutPortFrac = 1.0 - DrainFrac \quad (53)$$

The final output port capacity for the calculation of fair rate is then :

$$OutPortCapacity = \begin{cases} OutPortCapacity & , D \leq D_{Max} \\ OutPortFrac \cdot OutPortCapacity & , D > D_{Max} \end{cases} \quad (54)$$

The fair rate for requests per frame for each output port can be calculated as follows :

$$CurrentOutPortFairRate = \left(\frac{OutPortCapacity - AvgOutPortHPCellCount}{AvgOutPortABRSources}, 0 \right)^+ \quad (55)$$

Next, the OBNC determines whether earth terminals' ACR values need to be adapted in response to the current load and ABR queue size of the output port. In the case when ABR queue size of the output port is large and consequently capacity is set aside for draining, the OBNC will instruct terminals to decrease their ACR values. The amount of decrease is calculated such that the output port load reaches *OutPortFrac* within two round-trip delays, which is when the effect of the current control message will be experienced in the satellite's output ports (it takes one round-trip delay for requests to arrive in the satellite, and another one round-trip delay for the traffic corresponding to those requests to arrive in the output port). Hence, when *CurrentOutPortLoad* is larger than *OutPortFrac*, *TargetRateDecrease* (the amount of rate decrease in two round-trip delays) is set to $\frac{CurrentOutPortLoad - OutPortFrac}{CurrentOutPortLoad}$. When *CurrentOutPortLoad* is smaller than or equal to *OutPortFrac*, *TargetRateDecrease* is set to *DrainFrac*. The network controller then calculates the immediate amount of rate decrease required in

earth terminals such that :

$$(1 - \lambda_{Dn})^k = 1 - TargetRateDecrease \quad (56)$$

where λ_{Dn} is the amount of rate change (rate decrease) requested in the output port (downlink channel) and k is $\lceil 2r/m \rceil$, with r and m as previously used. Similar to the algorithm in the input ports, the amount of rate decrease is bounded by λ_{dMax} .

When $D \leq D_{Max}$, the network controller may request earth terminals to increase or decrease their ACR values depending on the current load in the output port. During an overload situation, the controller instructs earth terminals to decrease their ACR values. The amount of rate decrease is computed as for the input port, that is λ_{Dn} is calculated such that

$$CurrentOutPortLoad \cdot (1 - \lambda_{Dn})^k = 1 \quad (57)$$

and is bounded by λ_{dMax} .

Vice versa, during an underload situation, the controller will instruct earth terminals to increase their ACR values. An exception is made when the output port is intentionally brought into an underload situation in order to drain its buffer. In this case, the network controller will not instruct earth terminals to increase their ACR values even though the port is lightly loaded. The network controller will prevent rate increase in an output port for two round-trip delays after the last rate decrease message intended to drain the port's buffer is sent. In any other situation, an output port is considered to be underloaded when its current load is slightly below 1.0, i.e., $CurrentOutPortLoad < 1 - \varepsilon$, where ε , the minimum load offset, is a small number. Following the rate-based algorithm set by the ATM Forum, the increase in terminal's allocated rate will be linear. However, instead of using the terminal's negotiated PCR as the base of increase, BRCA uses the

broadcast fair rate as the base of increase. The amount of rate increase is calculated such that

$$CurrentOutPortLoad \cdot (1 + \lambda_{Dn} \cdot k) = 1 \quad (58)$$

where λ_{Dn} is the amount of rate change (rate increase) requested in the output port (downlink channel). The amount of rate increase is bounded by a maximum rate increase λ_{iMax} for the same reason as bounding the amount of rate decrease, that is, to prevent the system from being overly sensitive to variation in traffic. Figure 41 and 42 summarize the OBNC congestion control algorithm for output ports.

7.3.2.4 The Final Decision and Broadcast Process

In the next step, the OBNC determines whether rates should be increased or decreased for each input port - output port (uplink channel - downlink channel) combination. The OBNC final decision process is as follows. If earth terminal's ACR values need to be reduced, due to uplink or downlink congestion, then the amount of rate decrease is set to be the maximum of the calculated rate decrease in the input port and in the output port. Rate increase is only allowed if both the input port and the output port are lightly loaded. In this case, the amount of increase is equal to the amount of rate increase calculated in the output port. Figure 43 summarizes this decision process.

The final step in OBNC congestion control algorithm is the broadcast of control messages to earth terminals. Control messages sent to a given carrier group contain the fair rate and rate increase or decrease information for terminals within the carrier group, to adjust their ACR values to each of the satellite's downlink channels or output ports.

7.3.2.5 Summary of the OBNC Congestion Control Algorithm

The complete OBNC congestion control algorithm can be summarized as follows. The algorithm is invoked once information from all ports is received. This means that it is executed once every measurement interval. First, information from all input ports is consolidated. Then, the OBNC calculates the fair rate and the required amount of rate adaptation in each input port and each output port. Next, it determines the final fair rate and amount of rate adaptation for each input port – output port combination. The OBNC will then prepare and finally broadcast a set of control messages for each carrier group. Figure 44 depicts the flow diagram of the OBNC congestion control algorithm outlined above. Note that although the network controller is assumed to be onboard the satellite in the above description, no modification is needed if the network controller is located on the ground. Performing these operations on the ground only affects the signalling overhead and the system's response time, since one round-trip delay is needed for port information to arrive at the network controller and its corresponding control messages to arrive at the satellite's output ports.

Figure 41. Flow Charts of OBNC Algorithm for Output Ports

Figure 42. Flow Charts of OBNC Algorithm for Output Ports (cont'd)

Figure 43. Flow Charts of OBNC Final Decision Algorithm

Figure 44. Summary of the OBNC Rate Control Algorithm

7.3.3 The Terminal Algorithm for Rate Adjustment

The BRCA algorithm does not use RM cells to convey information related to congestion situation in the satellite. Instead, it relies on the periodic broadcast of congestion control messages. Since all terminals in a given carrier group share the same broadcast messages, BRCA has to employ a different ACR adjustment algorithm than that specified by the ATM Forum [1] in order to have selective control over each terminal's behaviour. The algorithm for rate adjustment in ground terminals is explained below.

At the beginning of its life, an ABR source should initialize its ACR equal to its ICR. The source's ICR, however, may be larger than its actual fair rate (the sum of its

MCR and the broadcasted fair rate). *An ABR source should initialize its ACR equal to the minimum of its ICR and its actual fair rate.* The current actual fair rate value can be obtained from the next (new) congestion control message broadcasted by the satellite. As a result, *a terminal has to wait for the immediate forthcoming congestion control message before it can make the first request on behalf of its new ABR source.*

The OBNC can start estimating traffic demands of ABR connections (sources) with non-zero MCR immediately after they are established. The OBNC, however, can only estimate traffic demands of ABR sources with zero MCR through their requests. The first congestion control message which represents the influence of a source's reserved capacity to the broadcasted congestion control information will arrive shortly after the arrival of capacity assignment due to the reserved capacity. Similarly, the first congestion control message which exhibits the influence of a source's first VBDC-HP request will arrive shortly after the arrival of capacity assignment due to the request, if the request is immediately granted by the uplink scheduler (if it is not immediately granted, the congestion control message may arrive sooner than the capacity assignment). Since it is very important to respond to the correct congestion control information, the terminal should not start adjusting its source's allocated rate until it has received capacity assignment for the source (due to reserved and/or requested capacity) and a congestion control message which represents the source's traffic demands. In other words, *a source's ACR is maintained at the minimum of its ICR and its actual fair rate until the terminal has sent traffic from the source.*

The situation where VBDC-HP requests cannot be immediately granted by the scheduler may be viewed as uplink congestion, although no cells are lost in the satellite due to this situation. The situation is certainly unfair for terminals whose requests have to

be delayed or even denied (due to time out). In order to avoid uplink congestion, extra care should be taken in reacting to a rate-increase message, since such an increase may aggravate the congestion situation. A terminal should only increase its source's allocated rate in response to a correct rate-increase message, that is, it should not increase its source's allocated rate until it has received capacity assignment for the source (due to reserved and/or requested capacity) and a rate-increase message. In other words, *a source's ACR can be increased after the terminal has sent traffic from the source*. Rate increases are additive and are functions of the source's MCR, its ACR, the broadcasted fair rate, the RIF and the fraction of rate increase indicated in the message. The new ACR value due to a rate increase message is calculated as follows :

$$\begin{aligned}
 & \text{if } (ACR_{current} \leq MCR + FairRate) \text{ then} \\
 & \quad ACR_{new} = (ACR_{current} + RIF \cdot PCR, MCR + FairRate \cdot (1 + \lambda))^- \\
 & \text{else} \\
 & \quad ACR_{new} = ACR_{current} + \lambda \cdot FairRate
 \end{aligned}$$

where λ is the amount of rate increase calculated (Figure 41) and determined (Figure 43) by the OBNC. Notice from the above expression that an ABR source is allowed to quickly reach the actual fair rate in increments proportional to its PCR and further tune its rate based on the loading condition in its traffic destination.

On receiving a message that requests a rate decrease (i.e., the message indicates an overload situation in the source's traffic destination), the terminal responds by decreasing its source's allocated rate if the rate is larger than its actual fair rate. Rate decreases are exponential and are functions of the source's MCR, the broadcasted fair rate and the fraction of rate decrease indicated in the message. The new ACR value due

to a rate decrease message is calculated as follows :

if ($ACR_{current} \leq MCR + FairRate$) *then*
 $ACR_{new} = ACR_{current}$
else
 $ACR_{new} = (ACR_{current} \cdot (1 - \lambda), MCR + FairRate)^+$

where λ is the amount of rate decrease calculated (figures 41 and 42) and determined (Figure 43) by the OBNC. Notice from the above expression that once an ABR source reaches its actual fair rate, it is always entitled to transmit at the actual fair rate.

Figure 45 summarizes the terminal algorithm for ACR adjustment.

Figure 45. Flow Charts of Terminal Algorithm for Rate Adjustment

Chapter 8

Performance Evaluation of An Integrated MAC and ABR Rate Control Scheme

In this chapter, the performance of the integrated multiple access scheme and rate-based congestion control scheme will be evaluated. An important performance measure of the effectiveness of a congestion control scheme is the buffer occupancy, or equivalently, the probability of cell loss due to congestion. One should be able to estimate this cell loss probability through analytical means when the control scheme is only modestly complex. However, with the sophisticated rate-based control described in the previous chapter, in addition to the complexity involved in the multiple access protocol and the interaction between the two, it is impossible to obtain this performance measure through analytical means. Issues such as optimality, fairness and robustness are difficult to evaluate analytically as well. Therefore, we will again resort to software simulations.

The two rate control schemes (ERICA and BRCA) described in the previous chapter are accurately implemented for discrete event simulation in OPNET environment. The base network configuration for performance evaluation is similar to that used to evaluate the performance of the multiple access scheme in Chapter 5 (Figure 18). The satellite network is composed of a geosynchronous satellite with onboard processing capabilities and a number of ground terminals. The satellite may support a number of spot

beams, each corresponds to a single carrier group. All terminals located in the same spot beam are considered at the same distance from the satellite (35,786 km), resulting in a round trip propagation delay of approximately 250 ms. Each terminal can support up to five traffic sources (generators), one for each service class defined by the ATM Forum [1]. The model user can associate each traffic generator with one of two available traffic models, the MMPP model and the self-similar WWW model.

Simulations are performed on a scaled-down model, to which each uplink MF-TDMA frame (of 24 ms duration) consists of 2 non-overlapping carriers. Each carrier is made up of 4 synchronization slots and 128 payload slots, resulting in a carrier burst rate of 2048 Kbps and an uplink capacity of 4096 Kbps. Log and acquisition slots are omitted since the login/logoff process of ground terminals will not be simulated and since terminals are assumed to be synchronized at all times. Downlink capacity is set to be slightly larger than uplink capacity, to offset the signalling overhead (assignment and congestion control messages). Each downlink frame consists of 264 payload slots resulting in a total downlink capacity of 4224 Kbps.

The association of service classes to R/A strategies is similar to that described in Chapter 5 (Table 6). Continuous variable-rate traffic (rt-VBR and nrt-VBR) is associated with the MMPP traffic model, whose parameters are given in Table 4. Data traffic (ABR and UBR) is represented by the self-similar WWW model whose parameters are given in Table 5.

Several test cases are used to assess various aspects of the integrated scheme. These test cases can generally be grouped according to the traffic pattern employed. There are three traffic patterns commonly used to assess the performance of a congestion control scheme, namely persistent sources, staggered sources and bursty sources [43]. Persistent

sources always have traffic to transmit. These sources will utilize all their allowed cell rates to request for uplink capacity, so that fairness and optimality become two important issues. Staggered sources start at different times and are useful to study the ramp-up and ramp-down times of the control scheme. Bursty (non-persistent) sources are always in one of two states : active state, when they actively generate cells, and idle state. In simulations involving bursty sources, if the average total network load is less than 100 %, fairness is not an issue anymore. Throughput, however, is still important when the total network load is close to 100 %. The most important issue for bursty sources is the burst response time [43], which is defined as the time from the arrival of the first cell of a particular burst to the departure of the last cell of the corresponding burst. When queueing delay is negligible in the onboard switch, the burst response time can be equivalently represented as the end-to-end transfer delay (a short response time translates to a short end-to-end cell transfer delay). In the thesis, the simulation test cases are grouped according to two traffic patterns, namely the persistent source pattern and the non-persistent (bursty) source pattern. The case of staggered sources will be incorporated in the persistent source simulation group. Test cases for each of these two simulation groups are described in detail below.

8.1 Persistent Source Simulations

There are 4 network configurations and 7 test cases defined for persistent source simulations. Performance measures of interest in these simulations are fairness (indicated by allowed cell rates granted to each ground terminal), optimality (expressed in the form of uplink and downlink utilization) and downlink buffer occupancy. Each test case was run for a simulated time of 25 seconds. This does not appear to be sufficiently long for

congestion control study. However, as can be seen later in the results, allowed cell rates of ground terminals stabilize (converge) within 5 seconds after turning active. Consequently, due to the persistent nature of each of the ABR sources, the dynamics of throughput and buffer occupancy can be captured within that 5-second period. Hence, there is no incentive for longer simulations.

8.1.1 Network Loading Configurations

As stated earlier, 4 network configurations are used for simulations with persistent sources. The first is the balanced network configuration. Here, the satellite supports two beams (two carrier groups) and ground terminals within each carrier group generate traffic to each of the two downlink channels in equal proportion. There are 8 ground terminals within each carrier group (a total of 16 terminals in the network), each with the same peak and initial cell rates to each of the downlink channels and zero MCR. Figure 46 depicts this balanced configuration in detail. Notice that the total load of each downlink channel is less than 1.0. This is due to the slightly bigger transmission capacity in the downlink channel, compared to the uplink channel capacity.

The second configuration is almost similar to the first. In this configuration, however, each carrier group has a terminal population of 6, which gives a total of 12 ground terminals in the network. Furthermore, not all terminals have the same peak and initial cell rates. In fact, in each carrier group, the total PCR and ICR of two terminals are larger than the total PCR and ICR of the other four terminals. Traffic, on the other hand, is still distributed equally to each of the two downlink channels. Figure 47 clarifies this second network configuration.

Figure 46. Configuration for Balanced Traffic Loading

Figure 47. Configuration for Balanced Traffic Loading with Different ABR Source Rates

The third network configuration is given in Figure 48. Here, the satellite supports three carrier groups, each with 16 terminals (a total network population of 48 terminals). This configuration is called non-balanced since the distribution of traffic in each uplink and downlink channel is not equal. For example, approximately 67 % ($0.4/0.6$) of traffic from carrier group 0 goes to carrier group 1, while the rest is transmitted to carrier group 2. Next, 67 % of the traffic from carrier group 1 and 33 % of the traffic from carrier group 2 is destined to the carrier group 0. Furthermore, all traffic from carrier group 0 corresponds

to 60 % uplink utilization, while both carrier group 1 and 2 uplink channels should be utilized at their maximum. Downlink congestion is expected to occur in downlink channel 1 (output port 1), due to the intentional overload. Finally, downlink channel 2 utilization is expected to be approximately 52 % considering the amount of traffic send to that downlink from carrier group 0 and 1. Notice that some terminals are not sending any traffic. These terminals are idle terminals and will not be considered in the scheduler's assignment process.

Figure 48. Configuration for Non-Balanced Traffic Loading

Finally, in the fourth network configuration, most of the traffic is destined to one carrier group, carrier group 0. The satellite still supports three carrier groups and 16 terminals within each carrier group (a total of 48 terminals in the network). Uplink utilization in carrier group 1 and 2 is expected to be 50 % and permanent congestion is set

to occur in downlink channel 0 (output port 0). Figure 49 shows this configuration in more detail.

Figure 49. Configuration for Concentration Traffic Loading

8.1.2 Simulation Test Cases

Based on the 4 network configurations described earlier, 7 test cases have been defined for persistent source simulations. The first test case uses the first network configuration (Figure 46). This test case is intended to examine the capability of both ERICA and BRCA schemes to allocate a fair rate to earth terminals in the simplest case (i.e., homogeneous terminals and homogeneous, equally distributed ABR traffic). Each terminal starts simultaneously at the beginning of the simulation and has a PCR ratio (relative to the carrier group uplink capacity) of 0.25. This means that each source in each ground terminal (one ABR source is associated with terminal's traffic to one downlink

channel) has a PCR of 32 packets/frame (512 kbps). The expected allocated rate (ACR) of each source is 16 packets/frame or approximately 666 packets/sec.

The second test case uses the second network configuration (Figure 47) and is intended to check whether each of the schemes is capable of allocating a fair rate to terminals with different ABR cell generation rates. Each source in terminals 0 to 3 in each carrier group (low-rate terminals) has a PCR of 16 packets/frame (256 kbps), while the rest of the sources (high-rate terminals) have a PCR of 64 packets/frame (1024 kbps). The starting time of each terminal is still the same. Each of the low rate sources is expected to have an ACR of 16 packets/frame (666 packets/sec), while each of the high rate sources should have an ACR of 32 packets/frame (1333 packets/sec).

The third test case is useful to assess the performance of both schemes under mixed traffic (rt-VBR and ABR traffic) environment. The speed of response of both ERICA and BRCA schemes in utilizing the transmission capacity unused by the higher priority rt-VBR traffic is assessed. The test case uses the first network configuration as depicted in Figure 46. 25 % of the generated traffic is of rt-VBR type, which has a higher priority than ABR traffic. The PCR of each ABR source is 32 packets/frame (512 kbps) and the starting time of all terminals is the same. Due to the higher priority given to the background VBR traffic, however, the expected allocated rate of each ABR source is 12 packets/frame (500 packets/sec).

The fourth test case is associated with the third network configuration (the non-balanced traffic configuration, Figure 48). This test case is the most complex test case with persistent sources, because of the unbalanced distribution of traffic and the use of terminals with different ABR cell generation rate. Terminals' starting time is still the same. Each ABR source in the active terminals in carrier group 0 has a PCR of 12.8

packets/frame (204.8 kbps), while those in carrier group 1 and 2 have an identical PCR of 64 packets/frame (1024 kbps). The expected allocated rate of each of the ground terminals is more difficult to assess. Notice that due to the low PCR of each carrier group 0 terminal, each ABR source in this carrier group should be granted its PCR (533 packets/sec). Consequently, the expected allocated rate of terminals 12 to 15 of carrier group 2 is $(1 - 0.4 \cdot (256/264))/4 \cdot 264 = 40.4$ packets/frame (1683 packets/sec). This rate accounts for 63.125 % of carrier group 2 uplink traffic destined to downlink channel 1. Hence, the expected allocated rate of terminals 8 to 11 of carrier group 1 is $(1 - 0.63125)/4 \cdot 256 = 23.6$ packets/frame (983 packets/sec). Carrier group 2 uplink traffic will then place a load of $23.6 \cdot 4/264 = 0.357$ in downlink channel 0. Consequently, the rest of the traffic to downlink channel 0 (64.3 %) will come from carrier group 1, so that terminals 0 to 7 in carrier group 1 will each have an ACR of $0.643 \cdot 264/8 = 21.22$ packets/frame (884 packets/sec). This rate accounts for 66.3 % of carrier group 1 uplink traffic load destined to downlink channel 1. Traffic from terminals 12 to 15 in carrier group 1 will occupy the rest of carrier group 1 uplink capacity, resulting in an expected allocated rate of $(1 - 0.663) \cdot 256/4 = 21.57$ packets/frame (900 packets/sec) for each of these terminals.

The fifth test case is the test of staggered sources. This test is used to examine the dynamics of uplink and downlink channel utilization and ABR buffer occupancy level as more terminals are added to the system. The network configuration is depicted in Figure 46 (the first network configuration). Each ABR source has a PCR of 64 kbps (1024 kbps). The starting time of each terminal is different, however. In this test case, one terminal within each carrier group becomes active every 80 frames (slightly less than 2 seconds). For example, terminal 0 becomes active at frame 0, terminal 1 becomes active at frame 80,

and so on. The final ACR value of each ABR source is similar to that expected in test case 1, i.e., 16 packets/frame (666 packets/sec).

The sixth test case uses the fourth network configuration (Figure 49) - the concentration of traffic configuration. This test is useful to assess the speed of response of each scheme to congestion situation. In this test case, each ABR source in carrier group 0 has a PCR of 6.4 packets/frame (102.4 kbps), while those in carrier group 1 and 2 have a similar PCR of 32 packets/frame (512 kbps). Terminals' starting times are the same. The expected allocated rate of each active terminal in carrier group 0 is equal to its PCR (533 packets/sec), while that in carrier group 1 and 2 is $264/16 = 16.5$ packets/frame (687 packets/sec).

The seventh test case, the last test case for persistent sources, also uses the same network configuration (Figure 49), only the starting time of each terminal within a carrier group is different. Therefore, in addition to assessing the speed of response to congestion situation, this test is also useful for studying the dynamics of uplink and downlink channel utilization and ABR buffer occupancy level as the terminal population grows. Similar to the fifth test case, one terminal within each carrier group becomes active every 80 frames, with the last terminal turning active at frame 560 (13.44 seconds simulated time). The expected allocated rate of each of the ground terminals is similar to that in the sixth test case (533 packets/sec for each terminal in carrier group 0 and 687 packets/sec for each active terminal in carrier group 1 and 2).

Each of these 7 test cases is summarized in Table 9.

Test Case	Terminal Traffic	Configuration	ABR Source's Peak Cell Rate Ratio	Terminal's Start Time
1	100 % ABR	figure 46	0.25	Simultaneous
2	100 % ABR	figure 47	Term 0 .. 3 : 0.125 Term 4 .. 5 : 0.5	Simultaneous
3	25 % rt-VBR 75 % ABR	figure 46	0.25	Simultaneous
4	100 % ABR	figure 48	CG 0 : 0.05 CG 1, 2 : 0.25	Simultaneous
5	100 % ABR	figure 46	0.5	Variable
6	100 % ABR	figure 49	CG 0 : 0.05 CG 1, 2 : 0.25	Simultaneous
7	100 % ABR	figure 49	CG 0 : 0.05 CG 1, 2 : 0.25	Variable

Table 9. Persistent Source Test Cases for ERICA and BRCA Performance Evaluations

The measurement interval for each simulation is two frames with either ERICA or BRCA scheme. Many of the ERICA enhancements described in [48] are implemented and utilized, including those to achieve max-min fairness and avoid transient overloads, the exponential averaging of active VC and load factor (to filter out variations in the number of VCs and the load factor in each uplink and downlink channel) and the ERICA+ scheme in the downlink channel. Similar parameters are used for each persistent source simulation involving the ERICA scheme, these parameters are given in Table 10. Similarly, each persistent source simulation involving the BRCA scheme uses the same parameters, as given in Table 11.

$\rho_u = 1.0$ (target uplink utilization)
$\rho_d = 1.0$ (target downlink utilization)
$\delta = 0.05$ (achieving max-min fairness)
$\alpha = 0.125$ (exponential averaging of active VC and load factor)
$T0 = 0.048 \text{ sec} / 2 \text{ frames}$ (target queueing delay)
$a = 1.15$
$b = 1.05$
$QDLF = 0.25$ (queueing delay limit factor)

Table 10. ERICA Parameters for Persistent Source Simulations

$\alpha = 0.125$ (exponential averaging of ABR sources and channel load)
$\rho_u = 1.0$ (target uplink utilization)
$\rho_d = 1.0$ (target downlink utilization)
$\lambda_{iMax} = 0.5$ (maximum rate increase)
$\varepsilon = 0.05$ (minimum load offset)
$\lambda_{dMax} = 0.5$ (maximum rate decrease)
$D_{Max} = 0.048 \text{ second} / 2 \text{ frames}$ (maximum queueing delay)
$\nu = 2.5$ (rate of exponential drain function)
$\beta = 0.25$ (maximum drain capacity)

Table 11. BRCA Parameters for Persistent Source Simulations

The ICR of each ABR source is $1/8$ of its PCR and each ABR source shares several common rate parameters, as given in Table 12. Notice that RIF, NRM, MRM, TRM, and CDF are set to their default value [1]. Furthermore, since the only switch in the network is an explicit rate switch, the RDF parameter is not used. The TBE is set to 4096, which, given the maximum ABR source's PCR of 64 packets/frame for all simulations,

should allow each source to use all of its ACR during the start-up period until its first RM cell returns. CRM is calculated by using the equation given in [1], i.e., $CRM = \left\lceil \frac{TBE}{NRM} \right\rceil$.

Finally, the ADTF is set to 2.0 seconds, to allow for sufficient time between subsequent RM cells before the source's ACR value is reduced to its ICR.

RIF = 0.0625
NRM = 32
MRM = 2
RDF = 0.125
CRM = 128
ADTF = 2.0 seconds
TRM = 0.1 seconds
TBE = 4096
CDF = 0.0625

Table 12. Common ABR Rate Parameters for Persistent Source Simulations

8.1.3 Simulation Results

This section discusses the results for each of the simulation test cases, for both ERICA and BRCA schemes. For each case, graphs of terminal's ACR values, uplink or downlink load (throughput) and downlink ABR buffer occupancy are presented.

The graph of terminal's ACR values is used to evaluate the fairness issue. A terminal is referenced with the form Term (*carrier group, ID, downlink channel destination*), based on the carrier group in which it is registered, its identification number within the carrier group and its traffic destination. For example, Term (0,1,2) refers to the ABR source within terminal 1 which is registered in carrier group 0 and is sending traffic to downlink channel 2. For clarity, not all ACR values will be shown, but at least one from

each group will be presented. It should be clear that fairness is not an issue in the BRCA scheme, since all terminals are receiving the same control messages. This will be proven later with results. The throughput is used to assess the optimality issue. In the case of ERICA, RM cells are not counted for in the calculation of the throughput. Based on the NRM value of 32, the expected load from RM cells is 6 % (3 % from each forward and backward RM cells). Hence, the maximum throughput achievable with the ERICA scheme is approximately 94 %. Finally, in order to maintain a manageable size document, only the maximum downlink buffer size seen in each frame is reported.

Figure 50 (a) - (d) and Figure 51 (a) - (d) depict the test case 1 simulation results for both BRCA and ERICA schemes, respectively. Figure 50 (a) and (b) and Figure 51 (a) and (b) show that both ERICA and BRCA schemes achieve fast convergence of rate allocation and are fair. The maximum uplink utilization is also attained, that is, 100 % for the BRCA scheme and 94 % for the ERICA scheme. Finally, downlink buffer occupancy is kept at a small value in every frame. Notice in Figure 51 (d) that the downlink buffer occupancy does not reach the target level set by the ERICA+ scheme. This occurs because further increase in terminals' ACR values cannot be granted, since uplink utilization is already at its maximum. The ERICA+ attempts to reach the target queueing delay by increasing terminals' ACR values. Since in most frames the network controller detects maximum uplink utilization, however, the increase cannot be maintained sufficiently long to allow downlink queues to build. This explains the coarse fluctuation in terminals' ACR values.

Figure 52 (a) - (d) and Figure 53 (a) - (d) present the test case 2 simulation results for both BRCA and ERICA schemes, respectively. Figure 52 (a) shows that the BRCA scheme achieves fast convergence of fair rate allocation to low-rate and high rate

terminals. Uplink utilization is close to its maximum and the variation in uplink load is caused by the rigid uplink scheduler's assignment process, which only considers one VBDC-HP request per terminal in the terminal table at every frame. Downlink utilization is always less than 100 % due to the slightly bigger capacity of each downlink frame. Finally, low buffer size is achieved at all times.

Figure 53 (b) shows that uplink utilization is at its maximum with the ERICA scheme. However, ERICA fails to allocate the expected fair rate to high-rate terminals, as shown in Figure 53 (a). The ACR for each of the high rate terminals is 50 % higher than the expected fair rate (i.e., high-rate terminals dominate the use of uplink capacity). This result confirms the fairness problem which exists, in spite of several modifications and enhancements, in the ERICA scheme, as noted in [51]. The high rate allocated to high-rate terminals explains why the scheme can accomplish maximum uplink utilization. ERICA is able to maintain a low buffer size at all times.

The test case 3 simulation results for both BRCA and ERICA schemes are shown in Figure 54 (a) - (d) and Figure 55 (a) - (d), respectively. Both schemes achieve quick convergence of fair rate allocation. Variations in terminals' fair rates are caused by the variability of rt-VBR traffic which in turn affects the available capacity for ABR traffic. Both schemes are also able to keep uplink utilization close to its maximum, as can be seen in Figure 54 (b) and (c) and Figure 55 (b) and (c). Nevertheless, BRCA responds faster than ERICA in utilizing the available capacity, as indicated by a less fluctuation in uplink load.

(a)

(b)

(c)

(d)

Figure 50. Persistent Source Test Case 1 Results (BRCA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size

Figure 51. Persistent Source Test Case 1 Results (ERICA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size

Figure 52. Persistent Source Test Case 2 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size

(a)

(b)

(c)

(d)

Figure 53. Persistent Source Test Case 2 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size

Figure 54. Persistent Source Test Case 3 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Uplink Load (per traffic classes); (d) Downlink ABR Queue Size

(a)

(b)

(c)

(d)

Figure 55. Persistent Source Test Case 3 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Uplink Load (per traffic classes); (d) Downlink ABR Queue Size

Figure 56 (a) - (d) and Figure 57 (a) - (d) present the results of test case 4 with both BRCA and ERICA schemes, respectively. Figure 56 (a) shows that BRCA scheme successfully achieves fast convergence of rate allocation for all terminals, with each terminal being allocated a fair rate equal or close to the expected rate discussed earlier (see Section 8.1.2). The scheme is also able to accomplish optimum uplink utilization in carrier group 1 and 2 and downlink utilization in carrier group 0 and 1. Further, Figure 56 (d) shows that port 0 ABR buffer size is quite large at the beginning of the simulation. The reason is that as the network controller finds out that downlink channel 1 is continuously congested, it attempts to maximize carrier group 2 uplink utilization, by allowing more traffic from this carrier group to be sent to downlink channel 0. This action affects the allocated rate of carrier group 1 terminals sending traffic to downlink channel 0. During the process of finding a fair rate for terminals sharing downlink channel 0, the ABR buffer size of port 0 increases. The ABR buffer size of both ports 0 and 1 stabilizes to a rather low number relatively quickly.

The ERICA scheme, on the other hand, is not able to bring rates into convergence (see Figure 57 (a)). This is caused by the ERICA+ implemented on the downlink, which attempts to build queues by initially allocating larger rates to terminals. Due to the slightly bigger downlink capacity and the rigid assignment process, however, uplink utilization cannot be maximized, and this would in turn cause the OBNC to allocate larger rates in order to reach the target uplink utilization. As port 0 ABR buffer size reaches a large value, downlink capacity is set aside to drain the queue, and consequently, each terminal sharing the downlink channel 0 is allocated a smaller rate and uplink utilization is reduced. The reduction in uplink utilization later leads to another overload situation after the buffer is drained, as the controller is again attempting to reach the target utilization (in both

uplink and downlink). This overload and underload cycle occurs continuously and rates of carrier group 1 and 2 terminals cannot converge. Figure 57(d) shows the periodic increase and decrease in port 0 buffer occupancy level, confirming the cycle of overload and underload. The BRCA scheme successfully avoids this cycle by not attempting to reach a target queueing delay, but by setting a target maximum delay above which downlink capacity will be set aside to drain the queue and bring the queueing delay below this maximum delay.

Figure 58 (a) - (d) and Figure 59 (a) - (d) show the test case 5 simulation results for both BRCA and ERICA schemes, respectively. In this test case, one terminal in each carrier group is becoming active every 80 frames (approximately 2 seconds). Figure 58 (a) and (b) and Figure 59 (a) and (b) show that both schemes are able to achieve quick convergence of fair rate allocation. Notice, however, the smooth rate allocation to each terminal as more terminals are becoming active with the BRCA scheme, compared to that obtained with the ERICA scheme. Both schemes are able to fully utilize uplink capacity and maintain a low downlink ABR buffer size in both ports 0 and 1.

Figure 60 (a) - (d) and Figure 61 (a) - (d) depict the test case 6 simulation results for both BRCA and ERICA schemes, respectively. In test case 6, all traffic from carrier group 1 and 2 is destined to downlink channel 0, creating a permanent congestion in this downlink channel. Figure 60 (a) and (b) prove that BRCA is able to allocate a fair rate to all terminals quickly. Figure 60 (c) shows the stable load in downlink channel 0, the congested channel. Finally, port 0 ABR buffer occupancy level only rises at the beginning of the simulation before rates converge. A low buffer size is maintained once the convergence state is reached. The ERICA scheme, however, is not able to bring the system into convergence (see Figure 61 (a) and (b)). This is again due to the ERICA+ scheme implemented in each downlink, which has been shown previously to bring the system into a cycle of overload and underload situation.

(a)

(b)

(c)

(d)

Figure 56. Persistent Source Test Case 4 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size

(a)

(b)

(c)

(d)

Figure 57. Persistent Source Test Case 4 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size

Figure 58. Persistent Source Test Case 5 Results (BRCA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size

Figure 59. Persistent Source Test Case 5 Results (ERICA) : (a) ACRs for Downlink 0; (b) ACRs for Downlink 1; (c) Uplink Load; (d) Downlink ABR Queue Size

(a)

(b)

(c)

(d)

Figure 60. Persistent Source Test Case 6 Results (BRCA) : (a), (b) ACRs for Downlink 0;
(c) Downlink Load; (d) Downlink ABR Queue Size

(a)

(b)

(c)

(d)

Figure 61. Persistent Source Test Case 6 Results (ERICA) : (a), (b) ACRs for Downlink 0; (c) Downlink Load; (d) Downlink ABR Queue Size

Figure 62 (a) - (d) and Figure 63 (a) - (d) present the results of the last test case for persistent sources, test case 7, for both BRCA and ERICA schemes, respectively. This test case is similar to test case 6, except that terminals' starting times are not the same. Instead, following test case 5, one terminal in each carrier group is becoming active every 80 frames (≈ 2 seconds). Figure 62 (a) shows the fast convergence of fair rate allocation as more terminals are becoming active with the BRCA scheme. The port 0 downlink ABR buffer size increases every time a terminal is becoming active, but quickly decreases once the queueing delay exceeds D_{Max} . Overload situation occurs in downlink channel 0 every time a terminal is becoming active, which will be followed by an underload situation, since capacity is set aside to drain the downlink ABR buffer. However, since the OBNC does not attempt to reach the target queueing delay, the cycle of overload and underload situation ends as the last terminal is becoming active. The ERICA scheme, however, fails to bring rates into convergence (see Figure 63 (a)), due to the reason previously explained.

Figure 62. Persistent Source Test Case 7 Results (BRCA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size

(a)

(b)

(c)

(d)

Figure 63. Persistent Source Test Case 7 Results (ERICA) : (a) ACRs; (b) Uplink Load; (c) Downlink Load; (d) Downlink ABR Queue Size

8.2 Bursty Source Simulations

A bursty source alternates between active and idle states. For simulations with bursty sources, the average total network load is set below 1.0. These simulations are not intended to assess fairness, but rather the burst response time as defined earlier.

8.2.1 Simulation Test Cases

There are 2 test cases used in simulations with bursty sources, each test employs the balanced load network configuration shown in Figure 46. The first test case is the simpler of the two and is used to assess the performance of both ERICA and BRCA schemes (in terms of the burst response time) under simple homogeneous traffic environment. In this test case, all traffic is of ABR type. The total load in each uplink is set to be 95 % of the total capacity, distributed equally between the 8 ground terminals, that is, each terminal has an average cell generation rate ratio (with respect to uplink capacity) of 0.11875 (or 1267 packets/sec). All terminals start at the beginning of the simulation.

The second test case is useful to assess the performance of both schemes under mixed traffic (rt-VBR and ABR traffic) environment. In the second test case, the traffic is mixed and 25 % of the average total uplink load is due to the higher priority continuous (rt-VBR) traffic, while 70 % is due to ABR traffic. Therefore, the total load in each uplink is maintained at 95 % of the maximum. The generation rate of ABR traffic from each terminal is equal at a ratio of 0.0875 with respect to uplink capacity (933 packets/sec).

Table 13 summarizes each of these 2 test cases. For simulations involving the ERICA scheme, it is assumed that 6 % of the uplink capacity will be utilized for the transfer of RM cells. Therefore, the average ABR cell generation rate ratio of each

terminal is adjusted accordingly in order to obtain the intended load. The simulated time for each simulation is 40 seconds.

Test Case	Terminal Traffic	Configuration	ABR Source's Average Cell Rate Ratio	Terminal's Start Time
1	95 % ABR	figure 46	0.11875	Simultaneous
2	25 % rt-VBR 70 % ABR	figure 46	0.0875	Simultaneous

Table 13. Bursty Source Test Cases for ERICA and BRCA Performance Evaluations

Similar to simulations involving persistent sources, the measurement interval for each of these 2 test cases is two frames. The ERICA parameters are mainly the same as those used for persistent source simulations, with some differences in the target uplink and downlink utilization (0.95 instead of 1.0). In addition, ERICA+ on downlink channels is replaced by the less complex ERICA algorithm, since downlink load will be, on average, less than 0.95 (due to the slightly larger capacity of downlink channels). The BRCA parameters used in these simulations are also mostly the same, with some differences in the target uplink and downlink utilization (0.95 instead of 1.0) and in the minimum load offset ($\varepsilon = 0.10$). This minimum load offset is set larger than in the persistent source case so that the scheme does not become overly sensitive to fluctuations in traffic (recall that in the BRCA scheme, the OBNC will request a rate increase to a certain downlink channel when the current downlink channel load is below the target value by at least ε). Finally, the PCR of each ABR source is set to be ten times its average cell generation rate, while the RIF is set to 0.0125 (instead of 0.0625) to avoid sudden load increases.

8.2.2 Simulation Results

As indicated earlier, performance measures of interest in simulations with bursty sources are no longer fairness and throughput, but rather, the burst response time. This is the time between the arrival of the first cell to the departure of the last cell of a burst. As also mentioned earlier, this response time can be equivalently represented by the end-to-end delay, provided that queueing delay onboard the satellite is negligible. As it will be seen later, the queueing delay is indeed very small. Therefore, one of the assessed parameter is the end-to-end cell transfer delay. The second assessed parameter is the cumulative distribution function (CDF) of the downlink ABR buffer size, which is useful to estimate the probability of cell loss given a certain buffer size, or to dimension the buffer to meet a certain CLR requirement.

It is useful to compare the results of these test cases with those obtained when rate-based congestion control is not implemented. This comparison will show the buffer saving achieved with the implementation of rate-based control and the optimality of the control scheme in handling traffic burstiness. Figures 64 (a) - (b) and 65 (a) - (b) present the results of both test cases when no rate-based control is applied. Figures 64 (a) and 65 (a) all indicate that with a buffer size greater than 40, the CLR due to congestion becomes negligible. Figures 64 (b) and 65 (b) show that with no rate control applied, the end-to-end cell transfer delay is approximately 34 frames.

Figure 64. Bursty Source Test Case 1 Results (No Rate-Based Control) : (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay

Figure 65. Bursty Source Test Case 2 Results (No Rate-Based Control) : (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay

Figure 66 (a) - (b) and Figure 66 (c) - (d) present the test case 1 simulation results for both BRCA and ERICA schemes, respectively. It is obviously seen that with the rate-based control applied, the buffer saving is significant. From the two CDF graphs in Figure 66 (a) and (c), it can be implied that the CLR becomes negligible when downlink port ABR buffer size is greater than 10. Between the two schemes, BRCA gives superior performance in terms of buffer requirement. BRCA also shows superior performance in handling traffic burstiness, as is apparent from the achievable end-to-end delay of approximately 36 frames (only two frames higher than in the case with no rate control is applied), compared to the end-to-end delay of approximately 42 frames obtained with ERICA.

Figure 67 (a) - (b) and Figure 67 (c) - (d) present the simulation results of test case 2 for both BRCA and ERICA schemes, respectively. Here the saving in buffer requirement is smaller than in the first test case, although it is still quite considerable. Similar to the previous case, BRCA yields superior performance (when compared to ERICA) in reducing buffer requirement. BRCA also achieves a considerably lower end-to-end delay than ERICA. Figure 67 (b) show that the mean end-to-end delay decreases gradually and at the end appears to be stabilizing at around 35 frames. With ERICA, however, the mean end-to-end delay quickly decreases and stabilizes at approximately 40 frames.

(a)

(b)

(c)

(d)

Figure 66. Bursty Source Test Case 1 Results : (BRCA) (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay; (ERICA) (c) CDFs of Downlink ABR Queue Size; (d) End-to-End Mean Delay

(a)

(b)

(c)

(d)

Figure 67. Bursty Source Test Case 2 Results : (BRCA) (a) CDFs of Downlink ABR Queue Size; (b) End-to-End Mean Delay; (ERICA) (c) CDFs of Downlink ABR Queue Size; (d) End-to-End Mean Delay

8.3 Summary

Performance evaluation of the integrated MAC and rate-based control scheme for ABR service has been presented in this chapter. In many test cases involving the ERICA (ERICA+) scheme, the scheme's performance is not always satisfactory. The ERICA+ scheme can bring the system to a cyclic transition between overload and underload conditions, which will prevent the convergence of rate allocations and consequently will yield suboptimum channel (resource) utilization. This finding indicates that an appropriate ABR rate control scheme for a satellite network has to consider conditions specific to satellite environment (such as the rigid frame structure and the significant influence of the uplink scheduler). A control scheme which has been found to perform well in terrestrial ATM networks may not be suitable for satellite networks.

The BRCA scheme was proposed as an alternative to the ERICA scheme. It does not rely on RM cells, but rather takes advantage of the broadcast capability of a satellite to convey congestion control information to ground terminals. Consequently, it is not compatible to the rate-based control framework defined by the ATM forum. Nevertheless, it can still be implemented in satellite ATM networks with rate control segmentation. In this case, the BRCA scheme is used in the access (uplink) channel, where the access scheme significantly influences network performance, while anywhere else, the existing rate-based control framework and any compatible switch algorithm can be applied. The BRCA scheme has been shown to possess all the desired qualities of a congestion control scheme for satellite networks, including fairness, optimality and robustness. Furthermore, it has been shown to be superior to ERICA in many aspects.

Chapter 9

Conclusions and Future Research

9.1 Summary and Conclusions

The Enhanced Combined Free Demand Multiple Access (ECFDAMA) scheme has been proposed as a demand-based multiple access scheme for general multimedia satellite networks. The scheme contains various improvements over the general CFDAMA scheme described in [17], among the most important being the introduction of three new request/assignment strategies and the incorporation of several fairness features in the scheduler's assignment process. The four R/A strategies, namely reserved capacity (RC), rate-based dynamic capacity (RBDC), volume-based dynamic capacity (VBDC) and free capacity assignment (FCA), are intended for different types of applications. The RC strategy is directed toward real-time applications which cannot tolerate delay or delay variations and is appropriate for those applications whose bandwidth requirements are generally constant and known in advance. The RBDC strategy is directed toward real-time applications with considerably bursty traffic characteristics and a more relaxed constraint on cell transfer delay. The VBDC strategy is intended for jitter-tolerant applications, i.e., those that have no delay and delay variation constraints. Finally, the FCA strategy is used to maximize uplink utilization by distributing unused bandwidth to earth terminals and is geared toward non real-time applications. These four R/A strategies can be flexibly

combined to suit the individual requirements of each application. In fact, through an appropriate association of service classes to R/A strategies, the ECFDAMA can be tailored to suit the ATM technology, making it an ideal scheme for satellite ATM networks. One example of a possible mapping from service classes to R/A strategies has been presented in Chapter 5. Fairness features built into the scheduler include the round-robin scheduling, the queueing (instead of accumulating) of VBDC requests per terminal, and the priority scheme employed in the distribution of free capacity.

Being a demand-based multiple access scheme, the effectiveness of the ECFDAMA also depends heavily on the efficient use of capacity assignment by earth terminals. Each terminal has an important role in achieving efficient bandwidth utilization. Each terminal should only request the capacity needed to satisfy its traffic and exercise care in utilizing the assigned capacity, so that the QoS requirements for each of its connections can be met.

The performance of the ECFDAMA scheme has been evaluated extensively through discrete-event simulations. Simulation results confirm the excellent performance of the scheme. A stable performance even at the maximum load can be attained, due to the fair assignment process and the efficient use of requests and assignments by earth terminals. Through the use of two different request strategies (one based on rate and the other on volume), continuous traffic (CBR, rt-VBR, and nrt-VBR) can be isolated from the highly bursty data traffic (ABR and UBR). Therefore, the more stringent requirements of continuous traffic can be satisfied irrespective of the load from data traffic. Furthermore, it has been shown in all test cases that data traffic delay performance is excellent at low to medium loads and that a graceful deterioration in performance is seen at high to very high loads.

Due to the significant influence the access scheme has on the performance of the whole network, it is highly desirable to integrate the rate control scheme for ABR service to the multiple access scheme. This may require modifications to the currently adopted ABR rate-based control framework and the intended switch algorithm. Two suggested modifications are given in this thesis, namely exercising control over the rate of request (rather than the rate of transmission) and enforcement of fair resource allocation both in uplink and downlink channels. In this respect, the maximum number of time slots an ABR source is allowed to request is equivalent to its ACR per frame. In addition, the unrequested capacity (other than the fractional part) cannot be accumulated. Enforcing fairness in uplink channels requires a new method of detecting overload situations. The conventional method of using the number of arrivals is not accurate, since the maximum number of cell arrivals per frame is limited to the uplink capacity. As an alternative, the number of requested time slots is suggested. This has the benefit of earlier and more accurate detection of overload situation.

A new rate control scheme based on the broadcast of congestion information, the Broadcast Rate Control Allocation (BRCA), has been proposed for onboard processing satellites. The scheme was developed taking into account the strengths and limitations of a satellite with onboard processing. The overall BRCA algorithm can be subdivided into three sub-algorithms, namely the port, the onboard network controller, and the terminal algorithms. The port algorithm mainly consists of the measurement of capacity requests for ABR traffic, the measurement of traffic arrival rates and the monitoring of downlink ABR buffer size. The network controller, being the component with the most processing power onboard the satellite, is in charge of the calculation of the fair rate (for requests) for each uplink channel - downlink channel combination. Finally, the terminal algorithm

controls how each terminal should adjust its rate in response to the broadcasted control (rate) information and gives the scheme an implicit, selective control over the behaviour of each ground terminal. The BRCA scheme is not compatible with the current rate-based control framework for ATM networks. Nevertheless, it can be used complementary to any compatible switch algorithm in the network segment where the influence of the access scheme is substantial, such as in the uplink access segment.

The performance of the integrated MAC and ABR rate control scheme has also been evaluated through simulations. The BRCA scheme has been shown to possess all the desired qualities of a congestion control scheme for satellite environment, including fairness, optimality and robustness. In many aspects, the BRCA scheme has been shown to be superior to the ERICA scheme. Since BRCA does not use RM cells, a typical 6 % overhead due to the exchange of RM cells can be avoided. Further, by recognizing the usual capacity matching in the uplink and downlink channels and consequently by not setting a target downlink buffer size, the BRCA scheme is able to avoid the cyclic transition between overload and underload conditions exhibited by the ERICA scheme. This cycling has been shown to prevent convergence of fair rate allocation and consequently lead to suboptimum resource utilization. Finally, the superior performance of the BRCA scheme indicates that a suitable rate-based congestion control scheme for a satellite network has to consider conditions specific to satellite environment.

In summary, a multiple access scheme for satellite ATM networks has been presented. The access scheme has been shown to give excellent performance and robustness independent of the network load and provide a measure of performance guarantee to continuous, delay sensitive traffic. A integrated multiple access and rate control scheme for ABR service has also been described. Performance evaluation of the

integrated scheme has indicated that efficient use of network resources can be attained. Nevertheless, further improvement can still be achieved by taking into consideration conditions specific to satellite environment. The superior performance of the BRCA scheme over the modified ERICA scheme confirms this remark.

9.2 Future Research Directions

Several research directions related to multiple access scheme and ABR rate control for satellite ATM networks for further investigation are presented below.

9.2.1 Call Admission Control

Recall in Chapter 4 the restriction where the total maximum RBDC value of all terminals sharing an uplink channel cannot exceed the uplink capacity. This restriction is necessary in order to guarantee assignment to all valid RBDC requests. The restriction may be removed by incorporating call admission control. This may require modifications to the RBDC strategy and will introduce a probability of cell loss to continuous traffic associated with the strategy (rt-VBR and nrt-VBR). It would be interesting to investigate the required modifications to the strategy and the resulting cell loss probability to continuous traffic.

9.2.2 Scheduling

One disadvantage of the combination of the round-robin scheduling and the absolute assignment strategy is the inherent delay variability. The average delay variation of approximately 0.5 frame (12 ms) experienced by rt-VBR traffic, as shown in Chapter 5, confirms this remark. Further reduction in delay variation can be achieved by reducing the duration of a frame or by using other assignment strategy (such as that used in

MASCARA/PRADOS). It would be useful to see the trade-off between the attainable improvement in delay variation performance and the reduction in uplink utilization from using different assignment strategies.

9.2.3 Traffic Modeling

On one end of the spectrum, measurement-based congestion control techniques are relatively simple to implement but the response time may be too slow to handle short term congestion. On the other end of the spectrum, neural network-based techniques have the capability to learn traffic behaviour and anticipate future congestion but may be impractical for onboard implementation. In the middle of the spectrum are techniques based on statistical prediction. These techniques are still feasible for onboard implementation and, if developed properly, are capable to forecast future congestion with good accuracy. A worthy statistical-prediction technique should take into account the effect of the employed multiple access scheme in transforming traffic arrival characteristics in each terminal to traffic arrival characteristics in downlink channels. It would be useful then to study this transforming effect of the ECFDAMA.

9.2.4 Rate-Based Control

The modified ERICA and the BRCA schemes described in Chapter 7 assume that the network controller is located onboard. As indicated earlier, more extensive modifications to the ERICA scheme are necessary if the network controller is located on the ground (i.e., far from the onboard scheduler), in order to reduce the required signalling between the switch (ports) and the controller. With respect to the BRCA scheme, the algorithm for rate adjustment in ground terminals must be modified if the network controller is not located close to the scheduler. Further investigations are necessary to

adapt existing explicit rate-marking switch algorithms and the BRCA scheme to a general satellite network topology. It would also be interesting to investigate if, by delegating a more active role to uplink schedulers in controlling the rate of arrival, system's performance (network utilization and cell loss ratio) could be improved.

References

1. The ATM Forum Technical Committee. Traffic Management Specification Version 4.0.
2. Ian F. Akyldiz and Seong-Ho Jeong. Satellite ATM Networks : A Survey. IEEE Communications Magazine, July, 1997.
3. Anthony Hung, Marie-Jose Montpetit, and George Kesidis. ATM Via Satellite : A Framework and Implementation. Baltzer Journals, 1997.
4. M. Chitre, et. Al. Asynchronous Transfer Mode (ATM) Operation Via Satellite : Issues, Problems, and Resolutions. Milcom, Communications on the Move, Conference Record, 1993.
5. Nikos Passas, et. Al. Quality-of-Service-Oriented Medium Access Control for Wireless ATM Networks. IEEE Communications Magazine, November, 1997.
6. Osama Kubbar and Husein T. Mouftah. Multiple Access Control Protocols for Wireless ATM : Problems Definition and Design Objectives. IEEE Communications Magazine, November, 1997.
7. Norihiko, N. Masao, K. Ryuji. New Concepts and Technologies for Achieving Highly Reliable and High Capacity Multimedia Wireless Communications Systems. IEEE Communications Magazine, January, 1997.
8. Jim Gilderson and Jafaar Cherkaoui. Onboard Switching for ATM via Satellite. IEEE Communications Magazine, July, 1997.

9. Sun, F. P. Coakley, and B. G. Evans. Satellite ATM for Broadband ISDN. *Telecommunication Systems* 4, pg. 119 - 131, 1995.
10. D. Ivancic, Pong Chu, and Dong-Jye Shyy. Technical Issues Regarding Satellite Packet Switching. *International Journal of Satellite Communications*, Vol. 12, pg. 257 - 267, 1994.
11. Marie-Jose Montpetit and Anthony Hung. Design Issues for a Satellite ATM Network. 18th Biennial Symposium on Communications, June 2, 1996.
12. Andrew J. Viterbi. A Perspective on the Evolution of Multiple Access Satellite Communication. *IEEE Journal on Selected Areas in Communications*, Vol. 10, No. 6, August, 1992.
13. Shuzo Kato, et. Al. A TDMA Satellite Communication System for ISDN Services. *IEEE Journal on Selected Areas in Communications*, Vol. 10, No. 2, February, 1992.
14. Uyles Black. *Mobile and Wireless Networks*. Prentice Hall, Inc. Upper Saddle River, 1996.
15. L. Schilling, et. Al. Broadband CDMA for Personal Communication Systems. *IEEE Communications Magazine*, November, 1991.
16. Marie-Jose Montpetit and Martin Cote. LAN Interconnection by Onboard Switching Satellites. *Proceedings of the Canadian Electrical Engineering Conference*, Montreal, September, 1995.
17. Jahangir I. Mohamed. Combined Free/Demand Assignment Multiple Access Protocols for Packet Satellite Communications. Master Thesis. Concordia University, June 1993.

18. A. Hung and G. Kesidis. Performance Evaluation of Hierarchical Round-Robin Bandwidth Scheduling for ATM. Technical Report 96-07, E&CE Dept, University of Waterloo, July 1996.
19. N. Abramson. The Aloha System : Another Alternative for Computer Communications. Proc. AFIPS Conf., Vol. 37, 1970, pp. 281-285.
20. M. H. Hadjitheodosiou, F. P. Coakley, and B. G. Evans. Multi-Access Protocols for a Multiservice VSAT Network. Conference Record ICC'94, 1994, pp. 864 - 868.
21. Shuji Tasaka, Keisaku Hayashi, and Yutaka Ishibashi. Integrated Video and Data Transmission in the TDD Aloha-Reservation Wireless LAN. Conference Record ICC'95, 1995, pp. 1387-1393.
22. Izhak Rubin, Kirk K. Chang, Gianpiero Forcina. Performance Analysis of Load-Adaptive / TDMA Systems with Applications to Integrated-Services Satellite Networks. Conference Record ICC'94, 1994, pp. 848-852.
23. D. Raychaudhuri and Newman D. Wilson. ATM-Based Transport Architecture for Multiservices Wireless Personal Communication Networks. IEEE Journal on Selected Areas in Communications, Vol. 12, No. 8, October 1992.
24. B. Jafarian and A. H. Aghvami. Performance of Multiservice Dynamic Reservation (MDR) TDMA for Multimedia Wireless Networks. Conference Record ICUPC'96, 1996, pp. 291-295.
25. M. M. Balagangadhar and R. L. Pickholtz. Analysis of a Reservation Multiple Access Techniques for Data Transmission Via Satellite. IEEE Transaction on Communications, October 1979.

26. W. Crowther, et. al. A system for Broadcast Communication : Reservation Aloha. Proceeding 6th Hawaii International Conference System Science, January 1973, pp. 596-603.
27. D. J. Goodman, et. al. Packet Reservation Multiple Access for Local Wireless Communications. IEEE Transaction on Communications, August, 1989.
28. Giuseppe Bianchi. C-PRMA: The Centralized Packet Reservation Multiple Access for Local Wireless Communications. Conference Record Globecom'94, 1995, pp. 1340 - 1345.
29. Gang Wu, Kaiji Mukumoto, and Akira Fukuda. A PRMA Integrated Voice and Data System with Data Steal into Voice (DSV) Technique. Conference Record ICC'94, 1994, pp. 580-586.
30. Farooq Khan and Djamal Zeghlache. Analysis of Aggressive Reservation Multiple Access Schemes for Wireless PCS. Conference Record ICC'96, 1996, pp. 1750-1755.
31. Jeong Geun Kim and Indra Widjaja. PRMA/DA : A New Media Access Control Protocol for Wireless ATM. Conference Record ICC'96, 1996, pp. 240-244.
32. Yu-Dong Yao and Tho Le-Ngoc. CREIR, A Multiaccess Protocol for Onboard Processing Satellite Systems. Canadian Conference on Electrical and Computer Engineering, pp. 26.1.1-26.1.4, September 1991.
33. F. Bauchot. MASCARA: A Wireless ATM MAC Protocol. Proc. Wireless ATM Workshop, Helsinki, Finland, September 1996.
34. Jahangir I. Mohammed and Tho Le-Ngoc. Performance Analysis of Combined Free/Demand Assignment Multiple Access (CFDAMA) Protocol for Packet Satellite Communications. Conference Record ICC'94, 1994, pp. 869-873.

35. S. V. Krishnamurthy and Tho Le-Ngoc. Performance of CF-DAMA Protocol with Pre-Assigned Request Slots in Integrated Voice/Data Satellite Communications. Conference Record ICC'95, 1995, pp. 1572-1576.
36. Eko A. Wibowo, A. Iuoras, P. Takats, J. Lambadaris, and M. Devetsikiotis. Guaranteeing Quality of Service in Packet-Switched Satellites by Medium Access Control. To be presented in the CCBR'98 Conference, Ottawa, June 21st - 24th, 1998.
37. Victor S. Frost and Benjamin Melamed. Traffic Modeling for Telecommunications Networks. IEEE Communications Magazine, March, 1994.
38. Jun Huang, Tho Le-Ngoc, and J. F. Hayes. Broadband Satcom System for Multimedia Services. Conference Record ICC'96, 1996, pp. 906-910.
39. H. Heffes, D. M. Lucantoni. A Markov Modulated Characterization of packetized Voice and Data Traffic and Related Statistical Multiplexer Performance. IEEE Journal on Selected Applications in Communications, September, 1986, pp. 856-868.
40. Shuang Deng. Empirical Model of WWW Document Arrivals at Access Link. Conference Record ICC'96, 1996, pp. 1797-1802.
41. M. E. Crovella and A. Bestavros. Explaining World Wide Web Traffic Self-Similarity. Computer Science Dept., Boston University, Technical Report TR-95-015, August 29, 1995.
42. R. Jain. Myths about Congestion management in High-Speed Networks. Internetworking : Research and Experience, Vol. 3, 1995, pp. 101-113.

43. R. Jain. Congestion Control and Traffic Management in ATM Networks : Recent Advances and a Survey. *Computer Networks and ISDN Systems* 28, 1996, pp. 1723-1738.
44. Flavio Bonomi and Kerry W. Fendick. The Rate-Based Flow Control Framework for the Available Bit Rate ATM Service. *IEEE Network*, March/April, 1995, pp. 25-39.
45. H. T. Kung and Robert Morris. Credit-Based Flow Control for ATM Networks. *IEEE Network*, March/April, 1995, pp. 40-48.
46. K. K. Ramakrishnan and Peter Newman. Integration of Rate and Credit Schemes for ATM Flow Control. *IEEE Network*, March/April, 1995, pp. 49-56.
47. A. Charny, D. D. Clark, and R. Jain. Congestion Control with Explicit Rate Indication. *Conference Record ICC'95*, 1995, pp. 1954-1963.
48. R. Jain, S. Kalyanaraman, R. Goyal, S. Fahmy, and R. Viswanathan. ERICA Switch Algorithm: A Complete Description. *ATM Forum Contribution No. 96-1172*, August, 1996.
49. S. Muddu, F. M. Chiussi, C. Tryfonas, and V. P. Kumar. Max-Min Rate Control Algorithm for Available Bit Rate Service in ATM Networks. *Conference Record ICC'96*, 1996, pp. 412-418.
50. F. M. Chiussi, Ye Xia, and V. P. Kumar. Dynamic Max Rate Control Algorithm for Available Bit Rate Service in ATM Network. *Conference Record Globecom'96*, 1996, pp. 2108-2117.
51. A. Arulambalam, X. Chen, and N. Ansari. An Intelligent Explicit Rate Control Algorithm for ABR Service in ATM Networks. *Conference Record ICC'97*, 1997, pp. 200-204.

52. E. A. Bobinsky. Predictive Onboard Flow Control for Packet Switching Satellites. 14th International Communications Satellite Systems Conference, March 22, 1992.
53. P. P. Chu, W. D. Ivancic, and H. Kim. A Rate-Adjusting Congestion Control Scheme for Satellite Based Packet Switches. AIAA 94-1062, 15th International Communication Satellite Systems Conference, February 1994, pp. 1042-1052.
54. A. Baiocchi, N. B.-Melazzi, M. Listanti, C. Soprano. Definition and Analysis of a Reactive Congestion Control Scheme for Application to Satellite Switched Networks. Conference Record Globecom'96, 1996, pp. 2138-2142.
55. Hakan Mitts and Simo Veikkolainen. Use of ABR Flow Control in Wireless ATM Systems. Conference Record ICUPC'96, 1996, pp. 702-706.
56. S. Fahmy, R. Jain, S. Kalyanaraman, R. Goyal, and Fang Lu. On Source Rules for ABR Service on ATM Networks with Satellite Links. Proceedings of the First International Workshop on Satellite-based Information Services, November 1996.
57. Dah-Ming Chiu and Raj Jain. Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks. Computer Networks and ISDN Systems 17, 1989, pp. 1-14.
58. Raj Jain. A Delay-Based Approach for Congestion Avoidance in Interconnected Heterogeneous Computer Networks. Computer Communications Review Vol. 19, No. 5, October, 1989, pp. 56-70.
59. Ilias Iliadis. A New Feedback Congestion Control Policy for Long Propagation Delays. IEEE Journal on Selected Areas in Communications, Vol. 13, No. 7, September 1995, pp. 1284-1295.

60. Lawrence G. Roberts. Explicit Rate Flow Control - A 100 Fold Improvement over TCP. Online Publication (<http://www.ziplink.net/~lroberts/Explicit-Rate/Explicit-Rate-Flow-Control.html>), April 1997.

IMAGE EVALUATION TEST TARGET (QA-3)

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

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