Quality of Service Issues in Hybrid Fiber-Coax Networks

A Thesis
Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia

In Partial Fulfillment of the Requirements for the Degree of
Masters of Science
Electrical Engineering

by

Mark D. Corner

May 1998

Approvals

This thesis is submitted in partial fulfillment of the requirements for the degree of
Masters of Science
Electrical Engineering

Mark D. Corner

Approved:

Jörg Liebeherr (Advisor)

Stephen G. Wilson

W. Bernard Carlson

James H. Aylor

Accepted by the School of Engineering and Applied Science:

Richard W. Miksad (Dean)

May 1998

Abstract

The IEEE 802.14 working group is currently standardizing a new media access control (MAC) protocol for the emerging Hybrid Fiber Coax (HFC) networks. Crucial for the success of 802.14 will be its ability to support Quality of Service(QoS) for delay and throughput sensitive applications such as multimedia. Two methods are presented in this thesis to provide QoS, efficient interoperability with Asynchronous Transfer Mode(ATM) services and the design and simulation of a novel priority scheme.

HFC networks must support higher layer traffic services, namely, ATM Constant Bit Rate (CBR), Variable Bit Rate (VBR) and Available Bit Rate (ABR) traffic classes. This first part of this thesis investigates the inter-operation of the MAC protocol, as defined by 802.14, with ABR transmissions. An important finding of this study is that multiple users contending for bandwidth on the upstream channel in the HFC network may interfere with the feedback congestion control mechanisms of ABR traffic control. This interference can result in unfairness between ABR sources, and decreased utilization of the upstream HFC channel. As a solution to the problem we propose a scheme whereby the headend station of the HFC network returns congestion information contained in resource management (RM) cells to the ABR sources. The proposed mechanism can be incorporated into the ABR rate control scheme without modifying the current traffic management specifications. Two sets of simulation scenarios are presented to illustrate this finding. Some of these results have been presented to the IEEE 802.14 standardization committee.

In addition to supporting ATM services, the evolving 802.14 standard must include an effective priority scheme. The second part of the thesis investigates the ability of the current specification to provide priority service and shows that a preemptive scheduler is not a sufficient solution. We propose to augment the scheduler with a novel scheme for implementing priority access in an HFC random access environment. The proposed mechanism integrates a multilevel priority collision resolution system into the proposed IEEE 802.14 MAC. The scheme separates and resolves collisions between stations in a priority order. A set of simulation scenarios is presented that shows that the protocol is robust and efficient, is able to isolate higher priorities from lower ones and provide quick access to high priority requests. We also give analytical results on the space occupied by priority contention slots at any given interval after a collision.

Dedication

Anytime I use the word 'we' in this thesis I refer to a group of people that has made this work possible. I would like to thank my advisor, Jorg Liebeherr, for all of his guidance, hard work and wisdom, and my manager and coworkers at NIST, David Su and Nada Golmie for all of their help and their generous support of my education. In addition, I would also like to thank my friends and fellow students (Kira, Mike, Chris, Dallas, Tyler and Bhupinder) in the Multimedia Networks Group at the University of Virginia. A loving thanks go to my parents who have always, and continue to support and encourage me through my academic pursuits.

Contents

Nomenclature

ABR Available Bit Rate

ACR Allowed Cell Rate

ATM Asynchronous Transfer Mode

$BISDN$ Broadband Integrated Services Digital Network

CBR Constant Bit Rate

CI Congestion Indication

CS Contention Slot

$DQDB$ Distributed Queue Dual Bus

DS Data Slot

DSL Digital Subscriber Line

$EFCI$ Explicit Forward Congestion Indication

ER Explicit Rate

FTR First Transmission Rule

$FTTB$ Fiber to the Building

$FTTC$ Fiber to the Curb

$FTTH$ Fiber to the Home

HFC Hybrid Fiber-Coax

ICR Initial Cell Rate

IEEE Institute of Electrical and Electronics Engineers

ISDN Integrated Services Digital Network

MAC Multiple Access Control

MCR Minimum Cell Rate

NI No Increase

NIST National Institute of Standards and Technology

PNA Priority Newcomer Access

QI Queue Identifier

QoS Quality of Service

RDF Rate Decrease Factor

RIF Rate Increase Factor

RM Resource Management

RQ Request Queue

VBR Variable Bit Rate

WITL Wireless in the Loop

XDQRAP eXtended Distributed Queue Random Access Protocol

List of Figures

List of Tables

1

Introduction

The growth of the Internet and other computer information services has out-paced the expansion of the data rate available to home users using conventional modems. Constrained by the speed of conventional voice-band modems and the high price of other data services like dedicated leased lines and satellite links, users are demanding new means of residential community based network access. To meet this demand, existing community cable television systems are developing into bidirectional Hybrid Fiber Coaxial (HFC) networks [40] [51] that can support bandwidth-intensive interactive applications, including video-on-demand, tele-conferencing, telephony, and Internet access. High-speed data services have, in the past, been expensive, but advances in video-compression, software, and on-line services has lowered the price of residential broadband networking. Bandwidth-intensive applications are part of a variety of services that home users are willing to purchase.

However, the increased use of multimedia applications, such as voice and video, has created a data bottleneck often experienced by many residential users. This rapid expansion in data services has overwhelmed the bandwidth available from the telephone company's basic phone service. [40] Most residential users currently depend on the conventional phone system rather than more current access methods. The phone system is designed for the transmission of the human voice, which is typically limited to frequencies below 4 kHz. To maximize utilization of bandwidth, the phone system filters out frequencies above this threshold. Physical limitations prevent the phone system from providing higher-data rates that consumers demand. Integrated Services Digital Network architecture (ISDN) addresses this problem and gives services to home and business users at a basic rate, called a B-channel, of 64 kbps[45]. Nevertheless the growth of the Internet has created greater demand and in 1988 the networking research community started using the phrase Broadband ISDN (B-ISDN) to refer to any network service at a rate higher than the basic ISDN rate.

The Asynchronous Transfer Mode (ATM) architecture, which is based upon high-speed cell switching, is the standard protocol and service provider for the B-ISDN vision[13]. The difficulty has been delivering higher data rates directly to subscribers. A review of many different possible broad-band access technologies is found in [23]. Fiber to the Home (FTTH) is a possibility, but the high cost of laying fiber is too expensive for residential users, although Fiber to the Building (FTTB) has remained the choice of the commercial sector. To utilize the available in-place wiring, referred to as the last mile of the network, phone companies promote a class of systems called xDSL, and the cable companies have promoted the use of the cable systems using Hybrid Fiber Coax (HFC) networks [51] [38] [17]. Other popular contending technologies include wireless services such as Local Multi-point Distribution Service (LMDS), and Wireless in the Loop (WITL).

The new *Hybrid Fiber Coax* (HFC) system architecture is one of the front-runners in the race to provide residential community network access. An HFC network (see Figure 1.1) is able to utilize the in-place residential broadcast cable system to provide an immediate and cost-effective solution to bandwidth demands. Cable systems in the US are constructed in a hierarchical manner, [51] using a tree-and-branch topology, with as many as 2000 user stations attached at the leaves of the tree. The coaxial wire portion of the network extends from fiber-optic interconnected nodes to the subscribers' homes. Cable systems are currently being upgraded to carry bidirectional traffic from the users to the fiber node, called a headend, and from that node back to the subscribers. Cable operators are dividing the frequency spectrum on these wires into into a downstream region (typically from 50 to 750 MHz) and an upstream region (usually in the 5 to 40 MHz range). [51]

Stations transmit requests and data on the shared upstream channel to the headend, which is located at the root of the cable tree. All users share the upstream channel and collisions occur when more than one station transmits simultaneously. The headend transmits feedback and data to the users on a downstream channel, which is collision-free. Typically, home users receive much more data then they send, and the network bandwidth is allocated for efficient utilization. In order to support larger amounts of traffic in the downstream direction, data rates are approximately 3 Mbps and 30 Mbps in the upstream and downstream directions respectively. Time synchronization occurs at the physical layer, so that each station has a common ref-

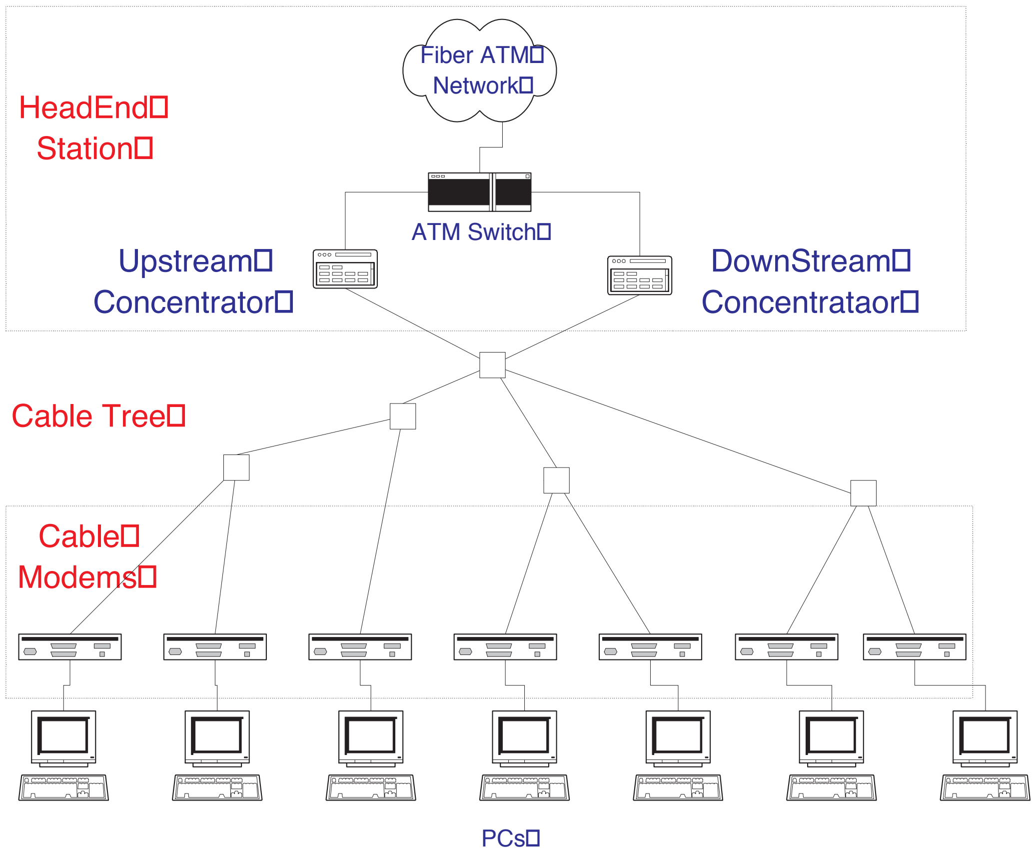

Figure 1.1: Hybrid Fiber Coax Network Connected to an Asynchronous Transfer Mode Network.

erence to base transmissions on. The cable system is often compared to a satellite communications architecture where a large number of ground stations, in this case home subscribers, communicate with a satellite, which is like the headend of the cable system. Due to electrical properties of the cable tree users may not directly communicate with one another and may only communicate with the headend. [43]

All shared network architectures require a Medium Access Control (MAC) protocol. Resource sharing between multiple independent users is arbitrated by the MAC protocol which is a set of rules that controls which users may use the channel at a given time and that settles resource conflicts. If cable operators are to use HFC to offer services to a large numbers of home-users the selection and design of these MAC protocols is essential to providing desirable performance characteristics.[42]

The IEEE 802.14 working group is currently standardizing a media access control (MAC) protocol for these emerging Hybrid Fiber Coax (HFC) networks[4][15]. While downstream communication from the single headend to the stations is free of mutli-user contention, the upstream channel from the stations to the headend is a shared access channel and subject to collisions between transmissions. The 802.14 working group is currently defining a contention resolution protocol for the MAC protocol that controls access to the upstream channel and resolves collisions.

In November of 1995, fourteen companies submitted proposed MAC algorithms to the IEEE working group, including Cabletron[1], Zenith[12], General Instrument[27], Scientific Atlanta[30], Phillips[18], HP[11], Lucent[44], LANcity[48], Panasonic[21] and IBM[53][16]. The IEEE standardization process uses these initial proposals as the foundation for the protocol specification. Each proposal is evaluated by the working group for features that can be included in the standard. Over the last two and a half years, elements of these proposals and countless minor contributions have been combined into a draft standard. As of the time of this writing, the working group is still debating which features and specifications will be included in the final version of the standard.

Most proposals contained collision resolution algorithms, of which S-ALOHA, START-3 and XDQRAP received the most interest from working group members.

The Aloha Network was developed around 1970 for communications access for computer terminals to computing resources at the University of Hawaii. [5] S-ALOHA (Slotted ALOHA) was used for a radio-based network but can be modified for use in any shared medium with a central communications server[2]. The theoretical maximum efficiency of S-ALOHA is 37%. [46] Unfortunately, the Aloha algorithm has been shown to be unstable in the case of infinite user populations and bistable [1] [10] [7] in the case of limited users. One method to stabilize ALOHA is pseudo-Bayesian broadcast whose authors claim to be exceptionally effective and stable in practice. [41]

The current IEEE 802.14 draft is largely based on START-3, an n-ary STAck ResoluTion (START-n) collision resolution algorithm, which has a throughput of 40%.[6] The last protocol, XDQRAP, was created to utilize the research done on Distributed Queue Dual Bus (DQDB) architecture for use in a residential tree-and-branch cable architecture. [52]

It is essential that the 802.14 standard MAC support Quality of Service (QoS) for home users. Quality of Service refers to the measurable characteristics of network delivery such as throughput, delay, delay variation and packet loss. QoS is a metric to evaluate different proposed MACs and judge what additions need to be made to a MAC protocol. Many real-time applications, like video conferencing and interactive game-play, require a specific level of QoS to operate correctly and without loss in quality. For instance, if a user is participating in a video-conference call and he

[1] Having two stable points, one with high throughput and one with low throughput for the same offered load

receives no data from the network for even a short period of time, the video will freeze and the sound will dropout. Thus the user is not receiving the network service quality he needs. The 802.14 MAC must support a high level of QoS to deliver data in a timely manner. This thesis proposes two improvements that can be made to the MAC to improve QoS for community network access.

This thesis is divided into two chapters. Chapter 2 explores the performance and inter-operation of the MAC protocol, as defined by the 802.14 specification, with the ATM ABR service. Using simulation experiments, we evaluate the degree to which contention at the MAC layer of an HFC network interferes with the rate-based flow control mechanisms of ABR traffic. Preliminary results were presented to the 802.14 working group in [32]. In Chapter 3, we evaluate the ability of the 802.14 MAC to provide preemptive priorities in a random access environment. After showing that scheduling is insufficient we present a novel mechanism which combines a priority system with the collision resolution protocol. We evaluate this protocol and show that it effectively implements multiple priorities. Results for this study were presented in [34]. After the submission of this work to InfoCom, this priority system was generalized and incorporated into the 802.14 standard. Both chapters are specifically focused to improve the ability of the MAC to provide QoS to cable modem users. QoS measures are used to evaluate their effectiveness and each are unique contributions to the field of broadband networking.

2

Improving the Effectiveness of ATM Traffic Control over HFC Networks

2.1 Introduction

The 802.14 MAC protocol must adequately support Quality of Service(QoS) in the network. Network designers created the mostly complete Asynchronous Transfer Mode (ATM) network architecture [26] [39] to meet the QoS demands of applications such as the transmission of multimedia information. Using ATM's switched network architecture, rather than a shared medium, it is possible to provide hard delay bounds and reliable QoS. Crucial for the success of 802.14 is its ability to support ATM's higher layer traffic services, namely, Constant Bit Rate (CBR), Variable Bit Rate (VBR) and Available Bit Rate (ABR) traffic classes. 802.14 group members have submitted some proposals to integrate ATM services with HFC. [24][29][31][50]

In order to support Quality of Service in HFC Networks it is essential that they support ATM services efficiently. Not only must the MAC protocol and scheduler be capable of supporting Constant Bit Rate (CBR) services through periodic bandwidth allocations to certain users, but it must also support the complex flow control structure contained in the Available Bit Rate (ABR) service. The flow control information is used to ensure a fair division of bandwidth between multiple users, and to prevent congestion in the network. In this chapter we address an interoperability problem that we found in a simulation of ABR over HFC. During periods of high congestion in the HFC network, flow control information being returned by stations attached to the HFC network can be delayed. This may create congestion on other links in the network, leading to lost information cells and long delays.

The ABR service, defined in the ATM Forum Traffic Management Specification [3], allows the available extra bandwidth to be used by several users in a fair and efficient manner. Users send data on the ATM network at an Initial Cell Rate (ICR). The network provides congestion feedback information and the station must adjust its rate to match the congestion conditions in the network. ABR is particularly vulnerable to delays in feedback information and because of this its inter-operation with HFC is of particular interest. The results of the study are summarized here:

- *Upstream ABR Traffic:* ABR transmissions that originate inside the HFC network and have destinations outside the HFC network maintain fairness and quality-of-service (QoS) requirements.

- *Downstream ABR Traffic:* ABR traffic sources that send into an HFC network may experience unfair bandwidth allocation, even if the downstream channel of the HFC network is not congested. The unfairness in bandwidth allocation is due to contention on the upstream HFC channel. Feedback information from the destinations to the sources is delayed, and, as a result, ABR sources transmit at a reduced rate.

This chapter proposes a solution to this unfairness problem in downstream ABR traffic. In this solution, the headend station of the HFC network generates feedback information that is returned to the ABR sources on the ATM network. The feedback is based on the load of the upstream HFC channel, where the network uses the backlog of the so-called grant queue at the headend station as a load indicator. This chapter shows how the proposed solution can be incorporated into the IEEE 802.14 [25] and the ATM Traffic Management [3] specifications, without modification.

The chapter is organized as follows. In Section 2.2 we provide an overview of the MAC protocol proposed by the IEEE 802.14 working group. In Section 2.3 we briefly review the ABR flow control mechanism. In Section 2.4 we discuss the performance of ABR transmissions over HFC networks. In Section 2.5 we present our scheme and offer conclusions in Section 2.6. Appendix A presents a more detailed description of the simulation environment used for performance evaluation and to demonstrate its effectiveness.

2.2 IEEE 802.14 Media Access Control Protocol

In an HFC network up to two thousand stations are connected to a single tree network. All stations transmit to the headend using an upstream communications channel. The transmissions on the upstream channel are divided into fixed-sized time intervals, so-called *minislots*. Stations send transmission requests to the headend in a single minislot; such a slot is then called a *contention slot* (CS). Stations send data in *data slots* (DS), which consist of multiple minislots. At the top of the cable tree, the headend station transmits feedback and data to the stations using a downstream channel. The system of upstream and downstream transmission channels is asymmetrical with typical upstream and downstream channel rates equaling approximately 0.5-10 Mbits/s and 30 Mbits/s, respectively. The IEEE 802.14 MAC is only concerned with the transmission of data on the upstream channel.

Figure 2.1 illustrates the steps taken by the Media Access Control in the HFC network. A station with data to transmit must send a request for bandwidth on the upstream channel to the headend station. Using the downstream channel, the headend acknowledges the request or indicates that a collision between two stations transmissions has occurred. The receipt of a collision initiates a resolution process. Once the collision is resolved, the headend sends a message to the station granting the use of the upstream channel. Since bandwidth is allocated by a reservation process, no collisions will occur during the transmission of data. Only transmission requests, which are transmitted in contention slots, are subject to collisions.

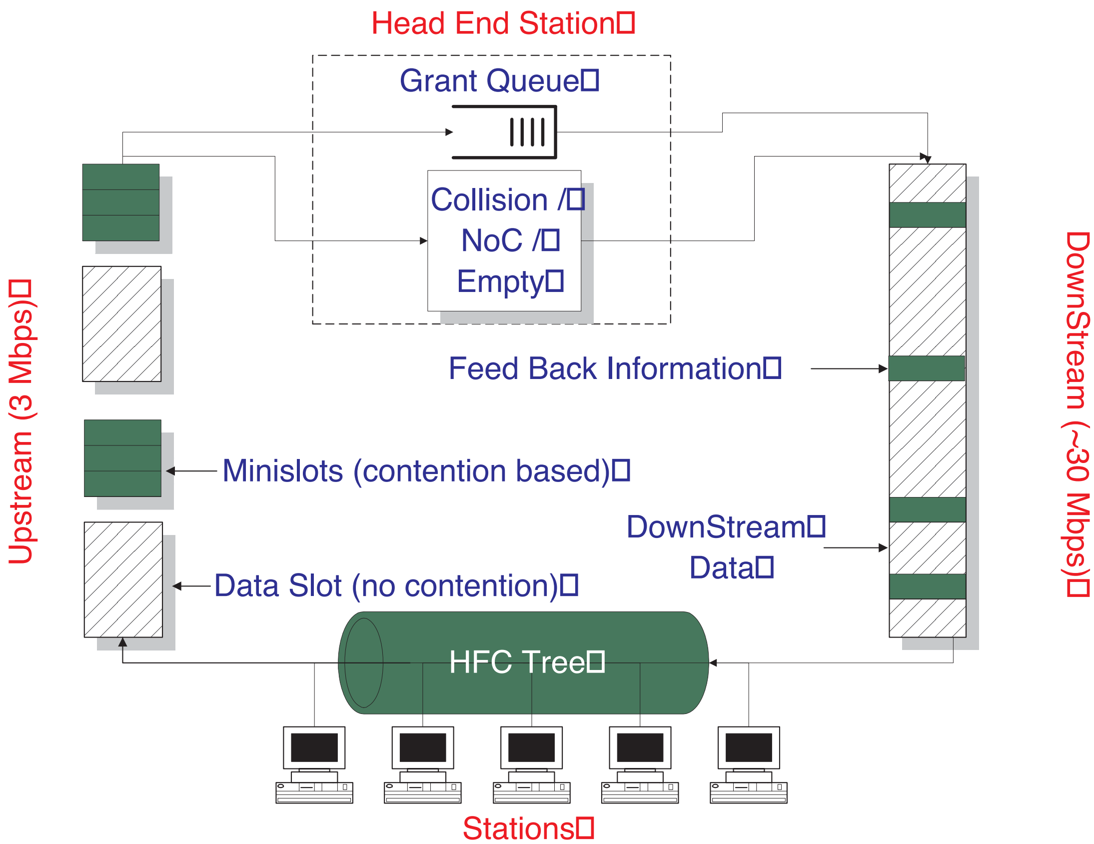

Figure 2.1: Media Access Control in HFC Networks.

The selection of the collision resolution algorithm for the contention slots on the upstream channel has received a lot of interest from the IEEE 802.14 research community. The current version of the MAC protocol is uses a blocking ternary tree splitting algorithm, which is an adaption of the n-ary Stack Random Access Algorithm [37, 7, 8]. Tree splitting algorithms have been used in the past to improve the performance of collision access [5]. In a tree splitting algorithm, all stations that are involved in a collision split into a number of subgroups. After a collision, only the stations in the first subgroup continue the collision resolution step. The stations in the second subgroup resume the collision resolution process after all stations in the first group have successfully transmitted, and so forth. A ternary tree splitting algorithm always divides colliding stations into three subgroups. Some tree-splitting algorithms are *non-blocking*, which allows stations to transmit new requests at any time. *Blocking* tree-splitting algorithms do not allow new stations to transmit during an ongoing collision resolution process [5]. The selection of a blocking algorithm for 802.14 is intended to reduce the MAC access delay variance [19]. The collision resolution scheme used in this paper is based on the status of January 1997 as reflected by [25, 49].

2.3 ABR Service Overview

The Available Bit Rate (ABR) service in ATM networks [3] is intended to carry data traffic, which requires a high degree of data integrity and incurs some transfer delays. An end-system that establishes an ABR connection specifies its maximum required bandwidth, referred to as *peak cell rate* (PCR), and minimum usable bandwidth, referred to as the *minimum cell rate* (MCR). During the lifetime of an ABR connection, the network can set the actual traffic rate, the *allowed cell rate* ACR, of the connection to any value provided $MCR \leq ACR \leq PCR$.

An end-to-end flow control mechanism, known as the rate-based mechanism, controls the ABR source rate as follows. A source starts sending its data at some negotiated Initial Cell Rate (ICR). Periodically, the source sends Resource Manage-

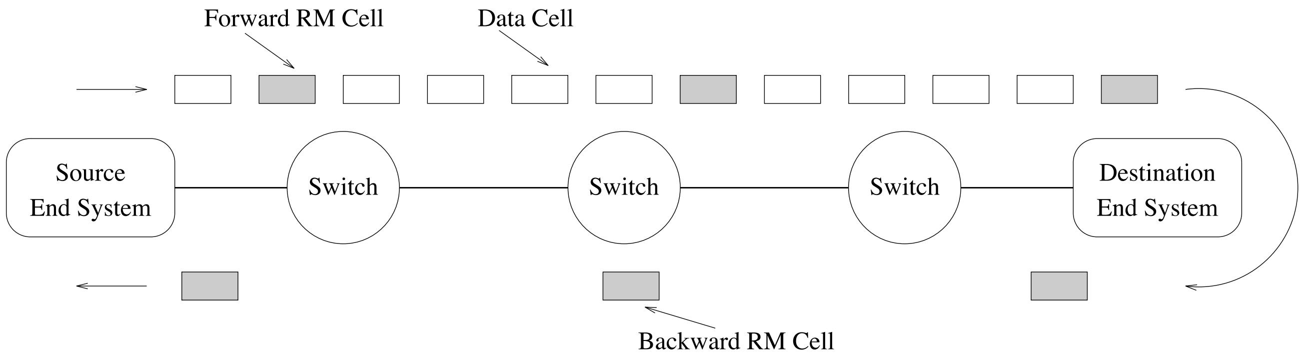

Figure 2.2: Closed-Loop Traffic Control.

ment (RM) cells along with data cells to its destination (see Figure 2.2). When RM cells arrive at the destination, they are returned to the source with some flow control information, such as congestion status and expected cell rate. Any intermediate network switching node can update the feedback information contained in the RM cell on its way back to the source. Based on this feedback information, the source adjusts its transmission rate. If a returning RM cell indicates congestion in the network, the source decreases the Allowed Cell Rate (ACR) multiplicatively by the Rate Decrease Factor (RDF). Otherwise, the source increases its ACR additively by a Rate Increase Factor (RIF).

In [3], two modes of switch behavior can be offered: EFCI (Explicit Forward Congestion Indication) and ER (Explicit Rate). When in a congested state, a switch in EFCI mode (EFCI switch) sets the EFCI bit in the header of all data cells that are forwarded to its destination. The destination conveys the congestion information back to the source by setting the Congestion Indication (CI) field in a returning RM cell. A switch in ER mode (ER switch) is more sophisticated in that it monitors its traffic and calculates an average fair share of its capacity per active connection. This quantity is called ‘explicit rate’ and is given directly to the source. In comparison, an ER switch provides more efficient and fair control of the source rate than an EFCI switch. Due the use of different parameters for the calculation of the explicit rate, there are several variations for ER switch mechanisms. In this paper we use an ER mechanism developed at NIST [22]; this mechanism attempts to achieve maximum network stability in terms of ACR and buffer occupancy oscillations.

2.4 Effectiveness of ABR Flow Control over HFC

The key issue for transmitting ABR traffic over an HFC network is whether the QOS guarantees of ABR connections can be maintained. As far as ABR QOS is concerned, the ABR service category definition in [3] clearly states that no bound is required on the delay or the delay variation experienced by a given connection. There is, however, a requirement to provide a low cell loss ratio for those connections whose end-stations obey a specified reference behavior. Also, it is assumed that all connections experiencing the same congestion conditions should receive an equal (‘fair’) share of the network bandwidth.

In order to evaluate the degree to which contention in an HFC network interferes with the feedback loop of the ABR rate control mechanisms, we have built a simulator of a combined ATM/HFC network. The primary concern in our study is how well the MAC layer in HFC networks supports the ABR service.

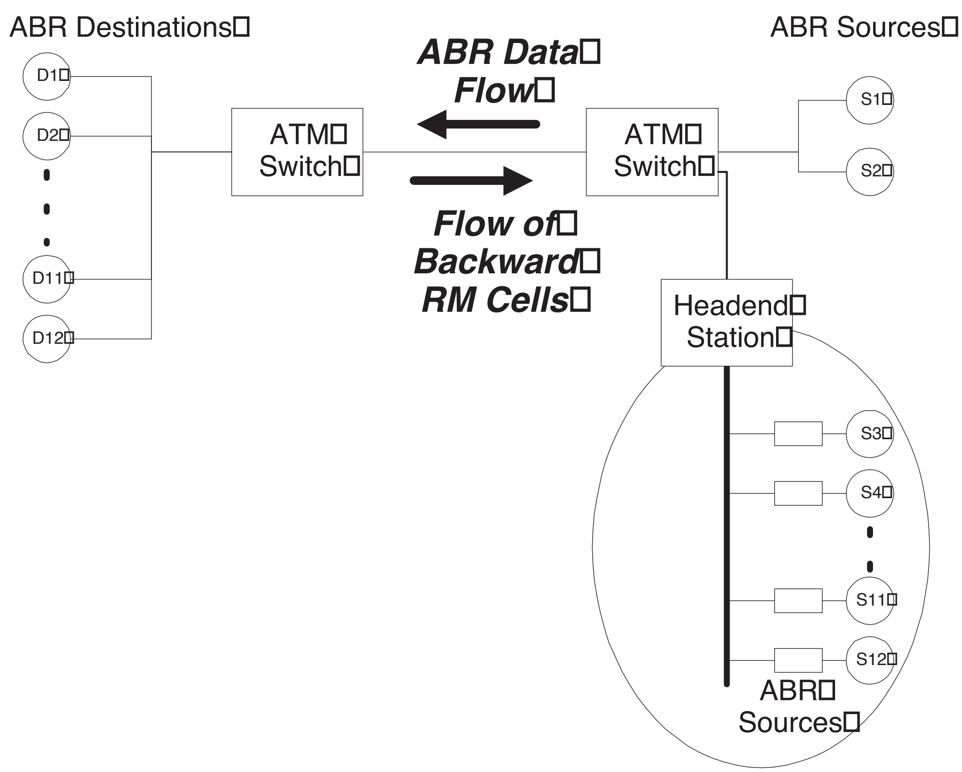

Figure 2.3: Upstream Configuration.

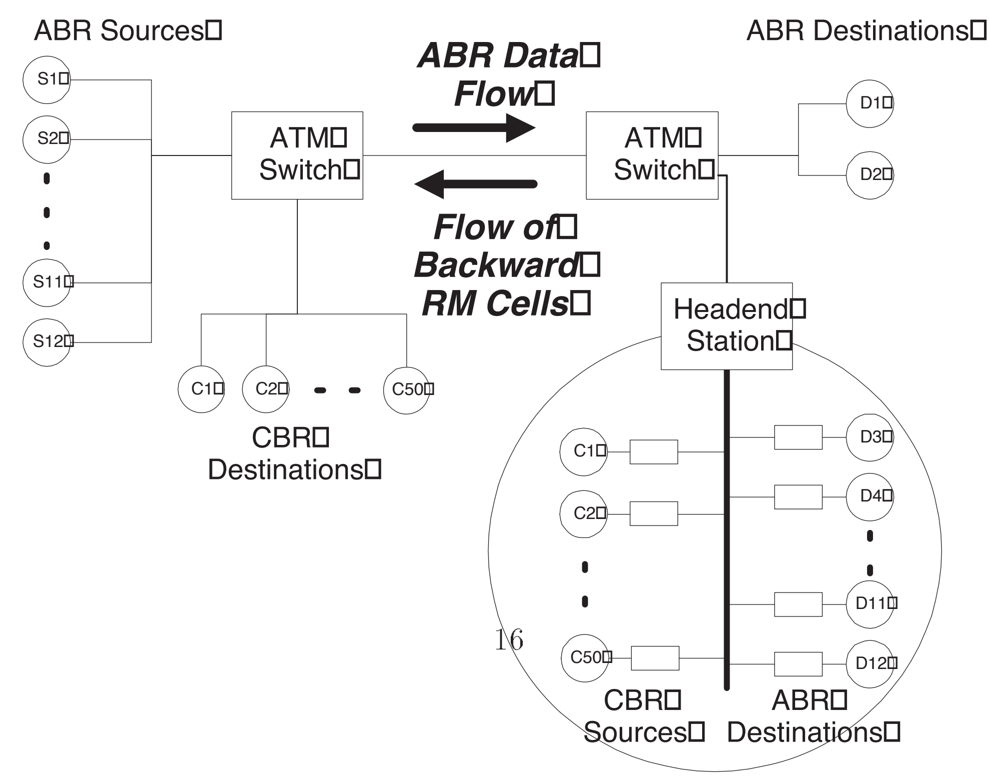

Figure 2.4: Downstream Configuration.

2.4.1 ATM/HFC Simulation

We have built a simulation of an ATM/HFC network to measure the degree to which HFC network will affect ABR rate control in an ATM network. The implementation was done using the NIST ATM Network Simulator [20]. We extended the ATM simulation package with a module for an HFC network with an interface to ATM components.

The simulated network scenarios are depicted in Figures 2.3 and 2.4. In both figures, the network topology is identical. The network consists of two interconnected ATM switches which are connected to sources and destinations of ATM traffic. One of the switches is attached to the headend station of an HFC network; the HFC network itself has stations that are sources and destinations of ATM traffic. Traffic sources send either CBR or ABR traffic. The link bandwidth available to ABR traffic between the ATM switches is set to 6 Mbits/s. By making this link the bottleneck in all simulation scenarios, we enforce that the the ABR rate control algorithms are active throughout all simulations.

In all simulation experiments we assume that cells are generated at a persistent constant rate for both ABR and CBR applications, set to 0.5 Mbits/s for the ABR applications and to 0.13 Mbits/s for CBR applications. The 12 ABR users, each sending at 0.5 Mbits/s create a load of 6Mbits/s which should fill the 6 Mbit/s ATM link. The 50 CBR users at 0.13 Mbits/s create a total load of 6.5 Mbits/s, which is sufficient to congest the 8.2 Mbit/s upstream HFC link. The complete set of parameters for the simulations is given in Appendix A. In the simulations we measure the transient behavior of the system.

Since it is our goal to study (1) the fairness of bandwidth allocation among ABR connections, and (2) the impact of delayed feedback on the ABR sources, we take the following measurements:

- Buffer Occupancy measured at the congested ATM link.
- Allowed Cell Rate (ACR) of the ABR traffic sources.

The two simulation scenarios shown in Figures 2.3 and 2.4 are referred to as Upstream Configuration and Downstream Configuration, respectively. In the Upstream Configuration (Figure 2.3) we evaluate the service given to ABR connections if the sources of the connections are located inside an HFC network. In the Downstream Configuration (Figure 2.4) we evaluate ABR connections that have the destination inside an HFC network. In both configurations, we consider both EFCI and ER switch control.

2.4.2 Upstream Transmissions

In the simulations for the Upstream Configuration, shown in Figure 2.3, we see a total of 12 ABR connections with sources labeled as Si ($i = 1, 2, \ldots, 12$) and destinations labeled as Di ($i = 1, 2, \ldots, 12$). Sources S3, S4, ..., S12 are located inside the HFC network. Since the available ABR capacity between the two ATM switches is given by 6 Mbits/s, we expect that each of the ABR sources obtains a fair share of 0.5 Mbits/sec as end-to-end throughput.

In Figures 2.5 and 2.6 we show the results of the simulations. In Figure 2.5 we depict the results for the ABR feedback mechanism when EFCI switch control is

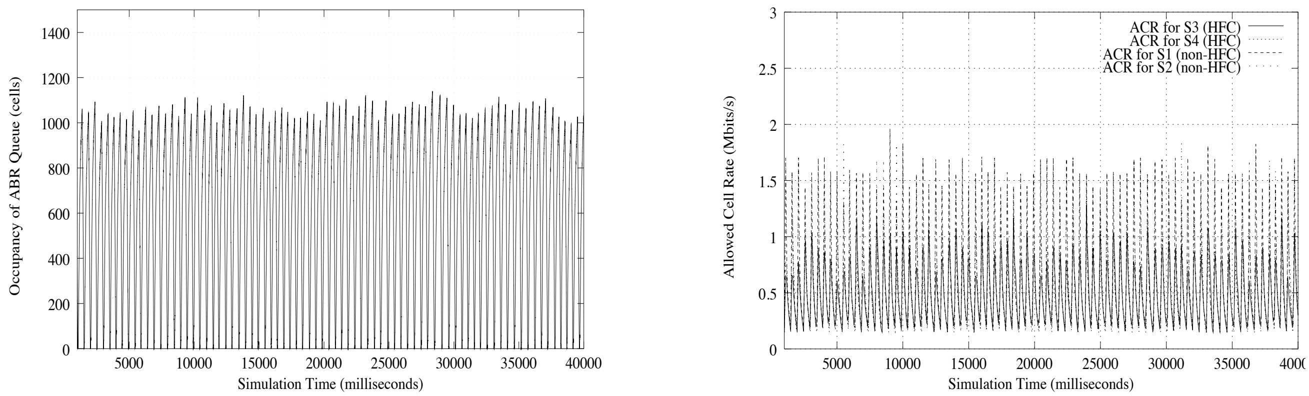

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 2.5: Upstream Transmissions: EFCI Control Mechanism.

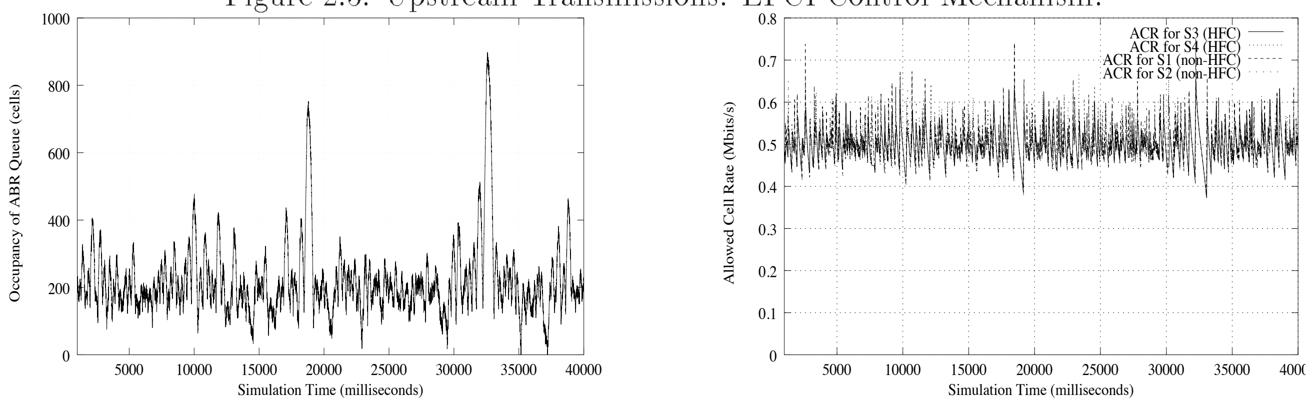

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 2.6: Upstream Transmissions: ER Control Mechanism.

used. Figure 2.5(a) illustrates the buffer occupancy at the congested ATM link, i.e., the output link of the right hand switch in Figure 2.3. Figure 2.5(b) depicts the allowed cell rate (ACR) of four ABR traffic sources: S1, S2, S3, and S4. Note that S1 and S2 are located outside the HFC network, while S3 and S4 originate inside the HFC network.

Foremost, we note in Figure 2.5 the oscillations of the measured values. This outcome is typical for a network with EFCI rate control. The amplitude and frequency of the queue oscillations and the ACR oscillations are due to the binary nature of the ABR feedback mechanism of EFCI. These oscillations, investigated extensively by the ATM Forum, mainly depend on the round trip delay, the parameters of the increase/decrease process (RIF, RDF), the buffer threshold levels, and the Initial Cell Rate (ICR). In Figure 2.5(b) we observe that the throughput oscillations of ABR sources from inside and outside the HFC network are quite similar. However, the peaks of the throughput graphs of sources S3 and S4 are smaller than those of S1 and S2. The smaller peaks are due to the bandwidth limitation of the upstream channel of the HFC network. Since 12 ABR connections with PCR values of 2 Mbits/s are sending on the upstream channel of the HFC network, which has a bandwidth limitation of 8.192 Mbits/s (see Appendix A), the HFC network does not have sufficient bandwidth to support the peak rate of all sources. As a result, the HFC network can become the bottleneck link for the ABR sources.

In Figure 2.6 we show the results for the Upstream Configuration if ER control is used at the switches. As compared to EFCI, the ER algorithms reduce the amplitude and frequency of the buffer and ACR oscillations (Figures 2.6(a) and 2.6(b)). In comparison to EFCI, the maximum buffer occupancy at the congested ATM link (Figures 2.6(a)) is an order of magnitude smaller with ER. Note from Figure 2.6(b) that the ACR rates oscillate around the expected rate of 0.5 Mbits/s.

In summary, we conclude that the Upstream Configuration preserves the properties of both ABR rate control algorithms. Throughput fairness of the ABR sources is maintained with both EFCI and ER.

2.4.3 Downstream Transmissions

Turning to the simulations for the Downstream Configuration shown in Figure 2.4. We observe that the contention on the HFC network results in noticeable unfairness, requiring a change to the ABR feedback mechanism.

In the Downstream Configuration we again have 12 ABR connections transmitting over an ATM link. All sources, labeled S1, S2, ..., S12 are connected to an ATM switch. Destinations D1 and D2 are located outside the HFC network, and D3 – D12 are located inside the HFC network. Note that the downstream bandwidth in the HFC network, set to 30 Mbits/s, is sufficient to support the peak cell rate of all ABR connections that enter the HFC network. In the Downstream Configuration in Figure 2.4 we add fifty CBR connections that transmit from inside the HFC network. The CBR connections, each transmitting at 0.13 Mbits/s, are intended to introduce high traffic load on the upstream HFC channel.

The results of the simulations are summarized in Figures 2.7 and 2.8. For EFCI switch control, we observe in Figure 2.7(a) that the backlog of the ABR queue frequently reaches the maximum buffer size of 10,000 cells, resulting in high cell loss rates due to buffer overflows. Obviously, the EFCI feedback algorithm is not effective in this situation.

An analysis of the situation reveals that the buffer overflows are caused by the CBR connections that are transmitting on the upstream HFC channel. These CBR connections lead to congestion on the upstream HFC channel. As a result, the backward RM cells from the ABR connections that are transmitted on the upstream channel are being delayed at the MAC layer. This increase of the MAC delay results in a rather large cycle time in the ACR oscillations for EFCI control; almost four seconds in Figure 2.7(b). As Figure 2.7(a) demonstrates, the excessive delays of the backward RM cells cause a breakdown of EFCI feedback control.

Figures 2.8(a) and 2.8(b) demonstrate that throughput fairness is maintained under ER switch control, even though the delays of the backwards RM cells are also large if ER switch control is used. However, the ACR values of all sources stay in the expected range of 0.5 Mbits/s for each ABR source.

2.5 Solution to the Downstream EFCI Problem

The problem with downstream transmissions of ABR traffic in an HFC network was not expected since the downstream capacity of the HFC network is rather large. However, as demonstrated by our simulations, the feedback cycle of EFCI switch control can collapse due to congestion on the upstream channel, independent of the bandwidth availability on the downstream channel.

In this section we present a scheme that maintains fairness and prevents a collapse of EFCI switch-control. Our solution has a number of desirable properties:

- Our scheme is implementable within the framework of the ATM Forum Traffic Management 4.0 specification [3]. No modifications to the ATM standard are required.

- The interactions between the MAC and the ATM layers is kept minimal.

- Our scheme does not result in throughput reductions or delay increases for non-ABR traffic.

Our scheme shortens the long feedback loop incurred during periods of upstream congestion of the HFC channel by passing a simple congestion indication signal from the MAC layer to the ATM layer. The solution works within the framework of the ATM Traffic Management 4.0 specification. More precisely, we exploit that the TM standard [3] makes allowances for extra ABR flow control mechanisms, such as the creation of backward RM cells at the switch.

2.5.1 Solution Approach

To prevent a collapse of EFCI rate control during congestion periods in the HFC network, our scheme short-circuits the feedback loop of RM cells in situations of high load on the upstream HFC channel. The solution scheme has three parts. First, there is a method for accurately determining the congestion level on the upstream link. Second, the MAC layer signals to the ATM layer a binary congestion notification (either congestion or no congestion). Third, upon receiving a congestion notification, the ATM switch generates backward RM cells that reduce the feedback cycle time. Next we discuss these steps in more detail.

(1) **Congestion Measurements:** The headend of the HFC network determines the congestion state by taking a measure of the backlog of bandwidth grants that have been distributed to stations. Rather than taking instantaneous mea-

surements of the grant queue size, the headend station tracks a weighted moving average computed as follows:

$$\text{GQ-Length (n)} = \alpha * \text{Current Length} + (1-\alpha) * \text{GQ-Length}(n-1)$$

Here, GQ-Length is the smoothed value of the grant queue size, Current Length is the instantaneous backlog in the grant queue, and α is a design parameter, set to 1/16 (a common parameter for weighted averages) in all our simulations.

(2) Congestion Indication: The headend determines if the upstream link is congested using two thresholds, and the measure of the average queue length. The headend has two design threshold values TH_{high} and TH_{low} which are used in the following manner:

$$Congestion = \begin{cases} TRUE & \text{if GQ-Length} > TH_{high} \\ FALSE & \text{if GQ-Length} < TH_{low} \end{cases}$$

(3) Interfacing with ATM rate control:

We assume that the headend station is directly connected to and integrated with an ATM switch. This allows the MAC layer to signal the ATM switch with the congestion status, as shown in Figure 2.9.

When the ATM switch receives a forward RM cell from the ATM link, it forwards the cell to the downstream link. If the switch has received notification of congestion on the upstream link, it generates a new backward RM cell with the

No Increase bit set (NI = 1). This backward RM cell shortens the feedback loop for sources sending to HFC destinations, since it short-circuits the delay that will be incurred on the congested upstream link. The generation of additional backward RM cells works within the framework of the TM 4.0 specification (Section 5.10.6 in [3]); the TM specification permits ATM switches to generate backward RM cells at a limited rate of 10 cells/sec per connection with either the congestion indication (CI) or no increase (NI) bit set.

Next we we demonstrate the impact of our solution method for the Downstream Configuration. We will see that the generation of additional backward RM cells at the headend has a profound effect on the rate oscillations and the buffer occupancy.

2.5.2 Evaluation

To evaluate our scheme, we use the topology and parameters from the Downstream Configuration in Figure 2.4. The network is enhanced by the mechanism described above. The results of the simulation are shown in Figures 2.10 and 2.11 and in Table 2.1. In the simulation, we have used the following threshold values:

$$TH_{low} = 35 \qquad TH_{high} = 40$$

Recall that the headend allocates 20 of the minislots for contention slots and 140 minislots for data slots. There are 4 minislots per data slot so there are 35 data slots in a frame. If there are more than 35 slots backlogged in the grant queue this indicates that cells will be delayed by at least one frame during upstream transmission.

For this reason we chose a lower threshold of 35 cells and an upper threshold of 40 cells. The difference of 5 cells in the thresholds prevents the system from excessive oscillations between a congested and non-congested state. When the grant queue grows larger than 40 cells, the mechanism is triggered, and the ATM switch begins to send backward RM cells.

From Figure 2.10 we observe that the threshold values are effective in preventing buffer overflows at the ATM switch. The congestion indication sent to the ATM switch causes the switch to produce enough backward RM cells to compensate for the delay in the feedback loop. In this particular example, the system is permanently congested by the 50 CBR sources.

In Figure 2.11 we reduce the rate of the CBR sources to 0.065 Mbps per station. This reduces the amount of congestion in the upstream HFC channel. Because the upstream channel is not congested the mechanism is not triggered by this choice of threshold values and the buffer occupancy is nominal.

In Table 2.1 we show the percentage of cells that each of the four ABR destinations should receive during a simulation run. Recall that there are 12 stations sharing a 6 Mbit ATM link between the two switches and thus each station should receive 0.5 Mbps worth of data. The actual amount of data received is dependent on EFCI flow control and the extra backward RM cells that the system provides. A comparison of the two previously mentioned experiments shows that when the HFC network is congested (due to the CBR sources) the ABR destinations on the HFC network receive less data. When congestion occurs the mechanism sends extra RM cells which reduces the overall throughput of sources with destinations on the HFC network. The

Destination	**Percent of Cells Received**	
	High HFC Congestion	*Low HFC Congestion*
D1 (ATM)	231.1%	119.6%
D2 (ATM)	231.1%	120.1%
D3 (HFC)	69.4%	92.9%
D4 (HFC)	73.1%	92.0%

Table 2.1: ABR Cell Success Rates.

mechanism limits the rate of sources when feedback is unavailable and thus the rate is not as high as it normally would be.

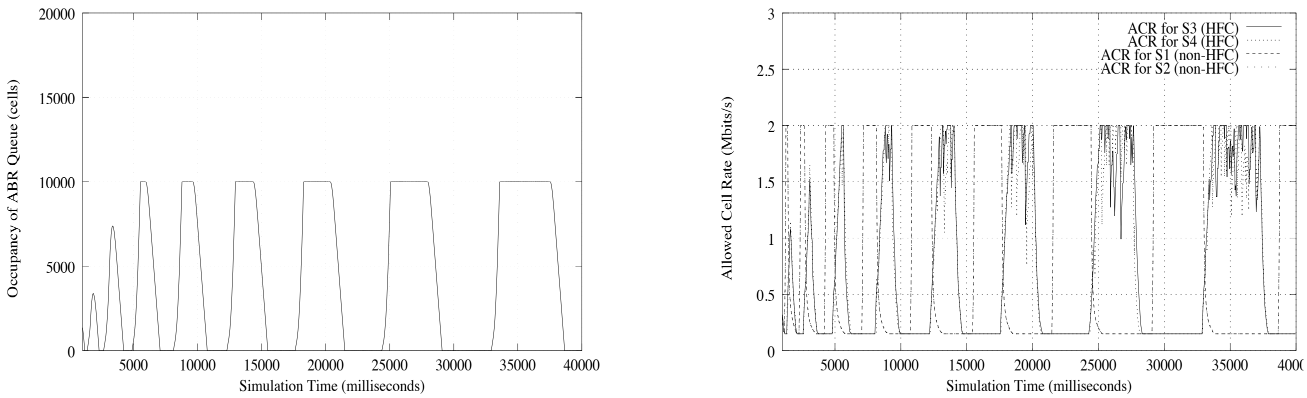

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 2.7: Downstream Transmissions: EFCI Control Mechanism.

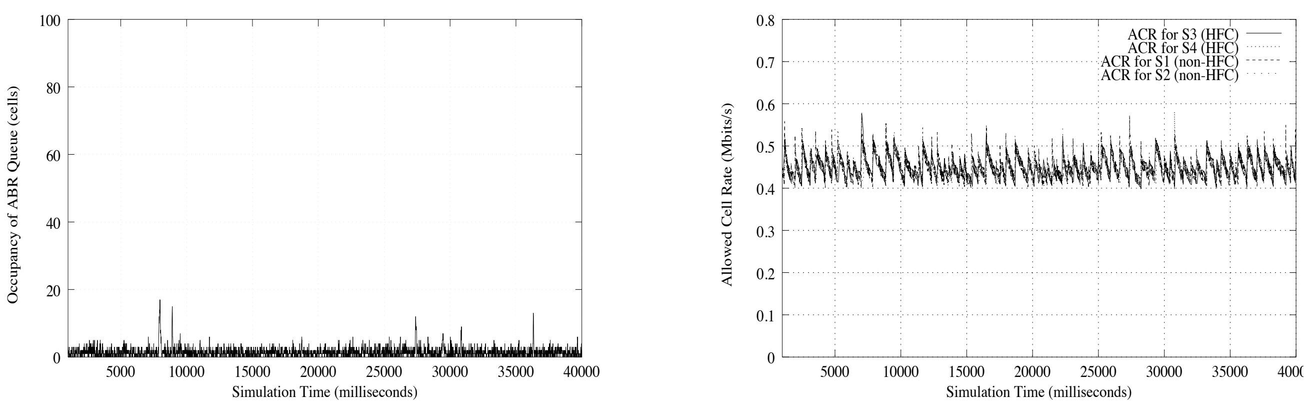

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 2.8: Downstream Transmissions: ER Control Mechanism.

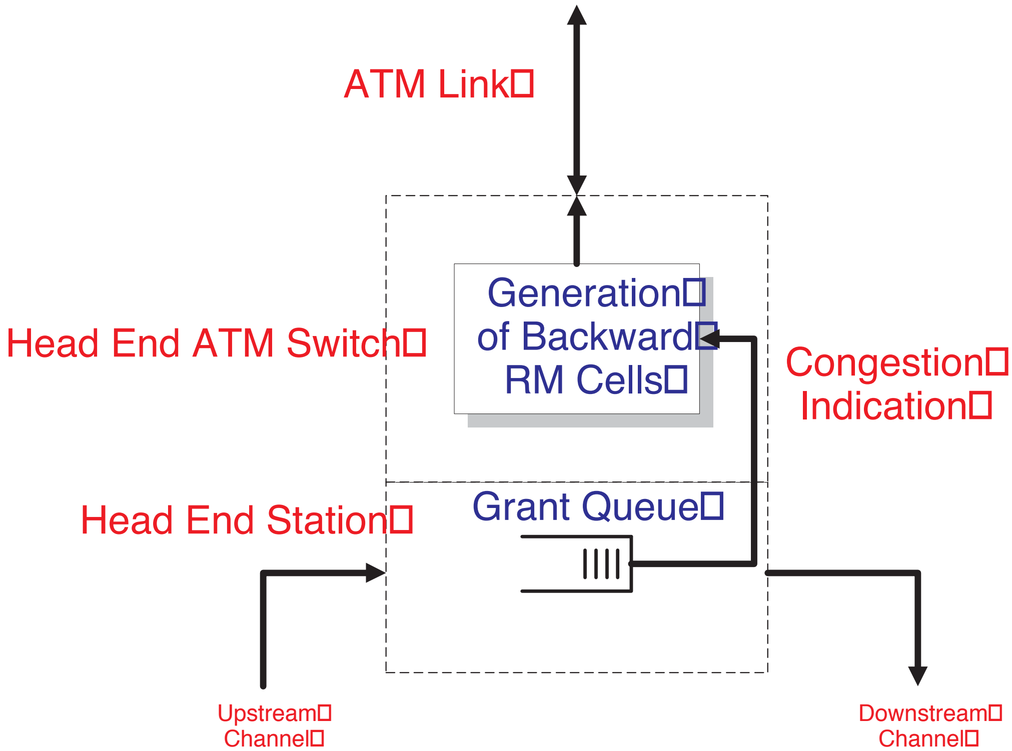

Figure 2.9: Modified Headend Station.

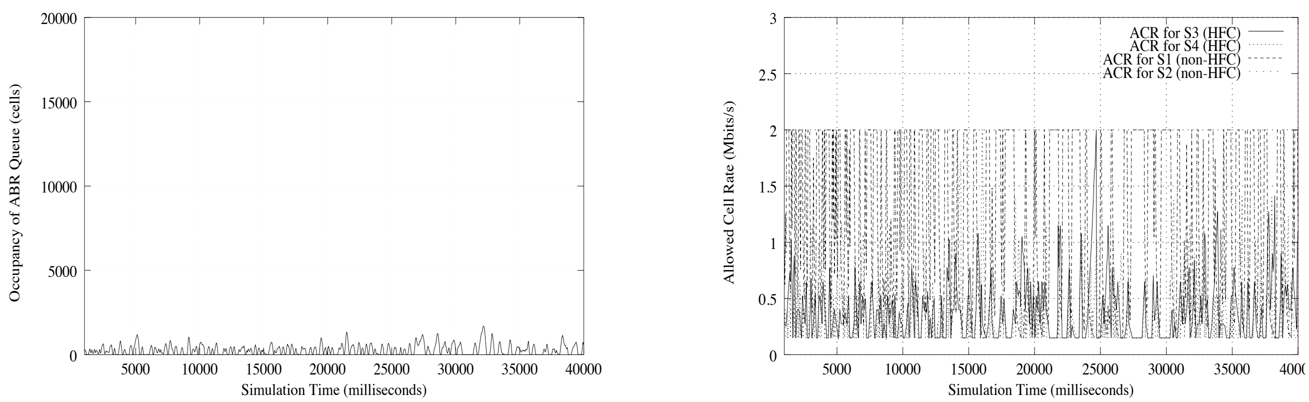

(a) Congested Link Switch Buffer Occupancy.

(b) ABR Allowed Cell Rate.

Figure 2.10: Downstream Transmissions-High Congestion (EFCI Rate Control): $TH_{low} = 35$, $TH_{high} = 40$.

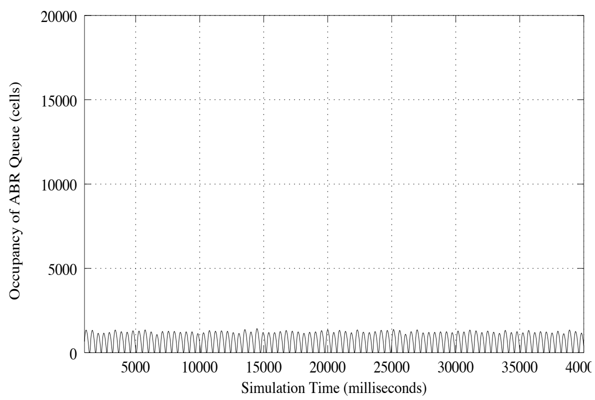

(a) Congested Link Switch Buffer Occupancy.

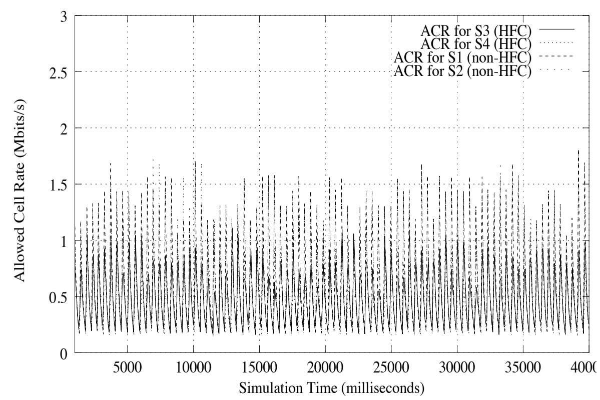

(b) ABR Allowed Cell Rate.

Figure 2.11: Downstream Transmissions-Low Congestion (EFCI Rate Control): $TH_{low} = 35$, $TH_{high} = 40$.

2.6 Concluding Remarks

This chapter have pointed out a possible problem when ABR traffic is transmitted over an HFC network that running the current version of the IEEE 802.14 MAC. We have shown that the fairness requirements of the ABR service may be violated for ABR connections that have destinations inside the HFC network. The problem results from congestion on the upstream HFC channel which may prevent backward RM cells to reach the ABR sources in a timely fashion. I proposed a solution whereby the HFC headend indicates its congestion level to the closest ATM switch, which, in turn, generates additional backward RM cells.

3

A Priority Scheme for the HFC MAC Protocol

3.1 Introduction

In this chapter of the thesis we investigate the ability of the MAC protocol, currently being defined by the IEEE 802.14 Working Group, to provide preemptive priority access to stations using the described MAC procedure. An effective priority system is needed to provide Quality of Service (QoS) in HFC applications and services such as voice, video and ATM[32]. Priority systems have been implemented in recent MAC protocols, such as DQDB[14] and Token Ring[47]. But the priority mechanisms used in those collision-free access protocols cannot be applied to the contention based HFC environment. [33] describes a modification to Extended Distributed Queue Random Access Protocol (XDQRAP)[52][30] that adds an extra slot to each frame to support priorities. However, this only provides access for two priorities with a fixed frame format. Note that the 802.14 MAC should support multiple priority levels and a dynamic frame layout. This is necessary to provide many different levels of service and to intergrate the scheme with the dynamic frame format outlined in the standard. In [12], a priority scheme is implemented with variable probabilities in combination with the p-persistence random access protocol. However, 802.14 cannot use this because the Collision Resolution Protocol(CRP) does not use random p-persistence for collision resolution.

To implement effective priority access, two mechanisms are used in this chapter. First, the headend uses a preemptive scheduler when allocating bandwidth to stations of different priorities. Second, the MAC protocol regulates collisions so that high priority stations are able to transmit requests without interference from lower priorities. We propose a multi-priority mechanism for IEEE 802.14 to implement the latter. This contribution is easily integrated with the standard and we show that it incurs little overhead.

This chapter is organized as follows. Section 2 explains the relevant details of the MAC protocol. Section 3 describes a new MAC level priority system for use in the HFC network. Section 4 presents several simulation test scenarios that show the performance of the system. Section 5 presents some analysis on the collision resolution space for priority traffic. Section 6 offers some conclusions.

3.2 The 802.14 MAC Protocol

In this section we review the operation of the IEEE 802.14 MAC protocol. Our priority mechanism, to be described in the following section, largely depends on the basic operation and Collision Resolution Protocol(CRP) of the standard, so an understanding of the protocol is essential to the description of the priority system. The 802.14 MAC layer specification [25] is not complete as of the time of writing (May 1998), and this description reflects the most current draft.

3.2.1 MAC Operation

The HFC upstream channel is divided into discrete basic time slots[1], called mini-slots. A variable number of mini-slots are grouped to form a MAC layer frame. The headend determines the frame format by setting the number of contention slots (CS) and data slots (DS) in each frame. CS, which are one mini-slot long, are used by the stations to transmit requests for bandwidth. DSs, which are several mini-slots long, are used by stations to transmit data. Only CS are prone to collisions, which occur when more than one station attempt to transmit a request in the same slot. The DS are explicitly allocated to a specific station by the headend, and are therefore collision-free. The headend controls initial access to the CS slots and resolves collisions by assigning a Request Queue (RQ) number to each CS.

The MAC protocol specifies a multi-step procedure for gaining access to the upstream channel. A station with a new request for bandwidth, or *newcomer* station,

[1]Some of this information is included in Chapter 2 as well and readers of the entire document can skip ahead.

gains initial access using a so called First Transmission Rule (FTR) [7]. The FTR specifies that the station waits for a group of CS with an RQ value of zero, called *newcomer CS*. The station then picks a number, p, between 0 and R (R is designated by the headend). If p is less than the number of CS in the group then the station waits for the p^{th} slot, and transmits the request. Otherwise it waits for the next group of newcomer slots.

After the headend receives a frame, it sends feedback to the stations, on the downstream channel. First, it sends the status of each CS in the frame. This indicates whether the slot was empty, successful, or contained a collision. Then the headend sends an RQ number, determined by the collision resolution protocol (CRP), for each slot that suffered a collision. The CRP specified by the 802.14 MAC is a blocking ternary tree algorithm[9]. The CRP assigns RQ numbers to collisions in descending order, starting with the first collision. The first collision in each frame is assigned the highest RQ number (the actual number depends on collisions that occurred in previous frames) and each subsequent collision in the frame is assigned an RQ number that is one less than the previous one. Then, each station that transmitted in a collided slot, saves that RQ number for future transmissions. The headend allocates three slots in the next frame with the same RQ number. For a station to retransmit its collided request, it must match the saved RQ number to the one found in a group of three CS. The station randomly chooses one of the three CS for retransmission. The assignment of RQ numbers can become complex when collisions have occurred in previous frames. Further details (including state machines and pseudo-code) can be found in [36].

3.2.2 Collision Resolution Example

Figure 3.1 shows an example of a collision resolution process. In this case, the system contains nine users, labeled A through I and each frame contains seven CS and two DS. The frame has the same duration as a round-trip, and 4 round-trips are shown, labeled (a) through (d). All RQ numbers assigned to CS in the initial frame, shown in Figure 3.1 are initialized to 0, so that they can accept the transmission of new requests.

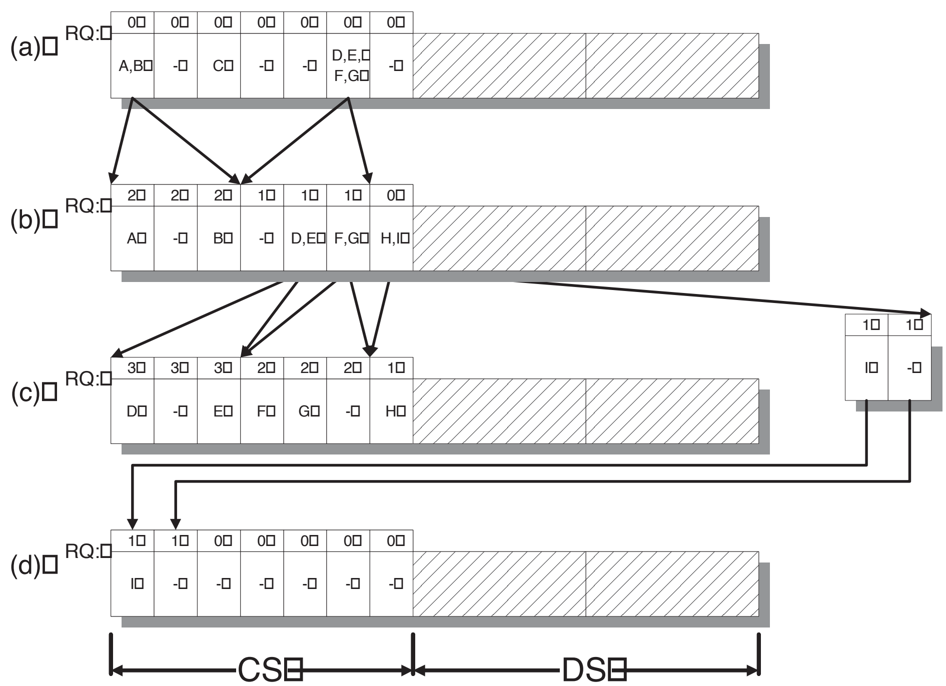

Figure 3.1: Collision Resolution Example

In the first frame, shown in Figure 3.1(a), stations A and B collide in the first slot, station C makes a successful request and stations D, E, F, G collide in the sixth CS. The highest RQ number, in this case, 2, is assigned to the first three slots in Figure

3.1(b) and to the stations, A and B, that collided first in the frame. The next highest RQ number, 1, is assigned to the second collision, involving D, E, F and G, and three CS with an RQ equal to 1 are allocated next. Stations A and B randomly select the first and third slots respectively. Stations D and E collide in the fifth slot, and F and G collide in the sixth slot. The seventh slot is still open for newcomer stations (RQ=0) and new stations H and I transmit in it. The RQ numbers are assigned in the correct order, D and E are assigned RQ=3, F and G are assigned RQ=2 and H and I are assigned RQ=1. In the next frame (Figure 3.1(c))there are not enough CS to accommodate all the slots needed for collision resolution, so station I must wait until the next frame. In the last frame (Figure 3.1(d)), the remaining slots with RQ=1 are allocated and station I transmits its request and the system returns to an idle state.

3.3 A Multi-Priority Access Scheme for 802.14

In this section we contribute an extension to the IEEE 802.14 MAC to provide priority collision resolution and access to multiple priority traffic. We first explain the need for a priority system by showing that headend scheduling is not sufficient and that a system integrated with the CRP is needed to efficiently support QoS.

3.3.1 Motivation for a Priority System

Currently in the IEEE 802.14 draft specification, stations can indicate their traffic type through the use of a Queue Identifier (QI) field in the CS. The exact guidelines for QI use have not been defined yet, but it is expected that this field will be used to indicate a traffic priority level. The headend uses a priority scheduler for stations indicating high priorities in the QI field, therefore a station that transmits a successful request for its priority traffic to the headend will gain immediate access to the channel. The time it takes a station to transmit a successful request to the headend, or request delay, must be kept low for high priority stations, even during periods of high contention.

Two problems exist in the current 802.14 draft. First, during contention all stations are treated equally with disregard for their priority, and newcomers can easily be blocked for extended periods of time, which may result in large delays for high priority stations. If a high priority request is blocked from accessing the channel or suffers a high number of collisions from lower priority traffic, it can not rely on the preemptive scheduler to receive low access delays. Second, the MAC does not provide a mechanism to give higher priority stations immediate access to the channel, nor does it separate and resolve collisions in a priority order.

Both of these problems are depicted in Figure 3.1(c). It shows that nine contention slots are needed to split the three collisions into three CS each, but only seven are available. Therefore, no contention slots with an RQ value of zero are allocated and a newcomer station with high priority data is not able to transmit in this frame. The first problem identified by this situation is that any new high priority requests would be delayed due to blocking. The second problem arises because the high priority stations have to send their request in contention with lower priority stations. When this occurs the CRP is unable to determine the priority level of the stations involved in a collision and thus resolves collisions without taking priorities into account.

3.3.2 Priority Protocol Description

Similar to the priority system suggested in [33], we have developed a scheme which integrates extra priority slots with the 802.14 frame format. However, we use a multiple priority system integrated with the ternary tree resolution protocol. As opposed to the fixed frame format found in XDQRAP, the flexible frame size of the 802.14 standard allows our protocol to allocate more CS to each priority level when needed.

Our scheme addresses both of the problems mentioned in the previous section by allowing higher priority stations to bypass the blocking feature of the CRP and by separating collision resolution for different priorities. In our protocol, areas of contention are defined for each priority level. The mechanism is as follows:

New Frame Format: In Figure 3.2 we illustrate a new frame format for the priority system. Several CS at the beginning of the frame are converted for exclusive use by priority stations. Each of these CS, referred to as a Priority Newcomer Access (PNA) slots, correspond to a single priority level. The headend identifies a PNA slot with a negative RQ number (unused in the current standard), where the RQ value $-N$ is reserved for priority level N. For example, an RQ number of -3 signifies that the slot is reserved for priority level 3. This provides a slot so that priority traffic is not blocked from accessing the channel by a lower priority.

New First Transmission Rule: Priority stations use the PNA slots for initial access. The previously described FTR is only used by stations of the lowest priority to access CS with an RQ value of 0. A new FTR is defined for stations with higher priority requests, which allows the stations to immediately transmit requests in the

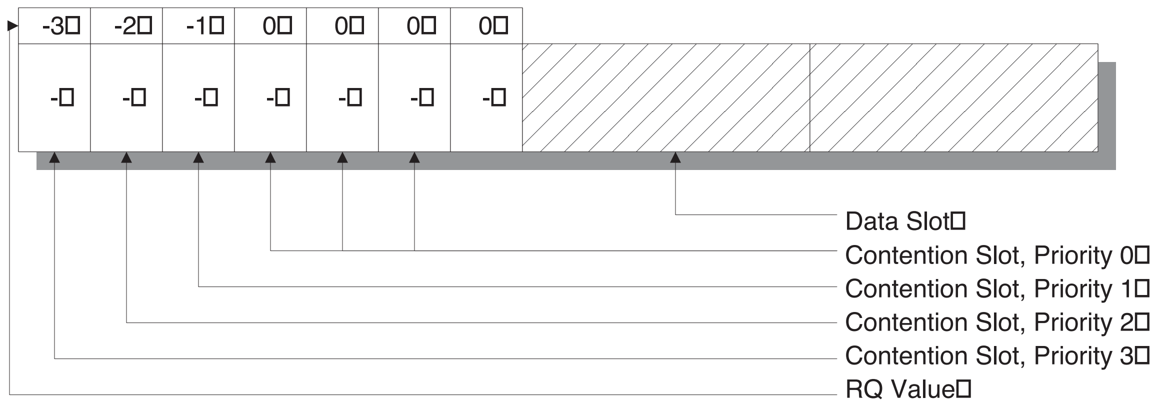

Figure 3.2: Priority Frame Layout

PNA slots. A station with a new request waits for a PNA slot with a priority that matches its own priority, and transmits the request with probability 1. Priority traffic gets immediate access to the channel rather than having to use the range parameter R. Note that this FTR reduces the request delay for stations with priority requests.

Separate Collision Resolution for Each Priority: Collision resolution is performed separately for each priority level. Stations initially transmit in slots exactly matching their priority level, so the headend knows that all stations participating in a particular collision are of the same priority level. The headend allocates three slots in the next frame for each collided slot; and each one of these slots is reserved for requests of the same priority as the first collision. Requests only collide with other requests of the same priority, preventing lower priorities from interfering with them.

Slot Allocation: Since the number of CS available in each frame may not be sufficient to accommodate all the slots needed for ongoing collision resolution and newcomer access, the headend allocates only some of the slots needed. The remaining slots must be allocated in a later frame. An example of this is shown in Figure 3.1(c), where two CS do not fit and must be allocated in the last frame(Figure 3.1(d)). The headend follows a priority order to determine which slots are allocated in the next frame and which are allocated in a later frame when space permits. Given that N is the highest priority, the order is as follows: 1) Collision resolution slots for priority stations at level N, 2) PNA slot for level N, 3) Collision resolution for level $N-1$, 4) PNA for level $N-1$, and so on. Any left over slots are allocated with an RQ equal to zero and used by the lowest priority. The ordering gives the highest priority collision resolution the first allocated slots and if the number of CS is not sufficient, the lowest priority collision resolution slots are allocated in later frames.

3.3.3 Example Priority Collision Resolution

In Figure 3.3 we show an example of the priority collision resolution process. Each frame corresponds to a round-trip in the system, and so the example represents a total of four round-trip delays. In this case, we use seven stations labeled A through G, with priority levels as shown in Table 3.1. There are four priority levels, where three is the highest priority and zero is the lowest. The frame consists of seven CS, which are initialized to an RQ of zero, and two DS. Figure 3.3(a) illustrates the initialized priority frame as described. Recall that a negative RQ number designates the CS as a PNA slot of the $\mid RQ \mid$ priority level. The first three CS, with RQ values -3, -2, and -1, are PNA slots for priority level 3, 2, and 1, respectively. The priority levels assigned to each CS are also shown in the diagram. The PNA slots are assigned the expected priority levels, and the remaining slots are assigned a priority level of 0.

Stations	**Priority**
A, B	3
C	1
D, E, F, G	0

Table 3.1: Station Priority Assignments

In the first frame, shown in Figure 3.3(a), stations A and B have traffic of priority level 3 so they transmit their initial request in the $RQ = -3$ slot. Station C transmits a successful request for priority 1 traffic in the $RQ = -1$ slot. Stations D, E, F, and G choose the same slot that has an RQ value of 0. In the second frame (Figure 3.3(b)) the headend assigns RQ values to split each collision across three slots, each with the same priority as the collided slot they are generated from. Stations A and B use contention slots with $RQ = 2$ and stations C, D, E and F use those with $RQ = 1$. The PNA slots for each priority level are still allocated to provide newcomers of those priorities with immediate access. There is not enough room in the frame to accommodate all the slots needed for collision resolution, and stations F and G are occupying slots of the lowest priority, so they must wait for a later frame. 3.3(c) shows the resolution of stations D, E and F. In the last frame, shown in Figure 3.3(d), all stations complete their requests and the system returns to the idle state.

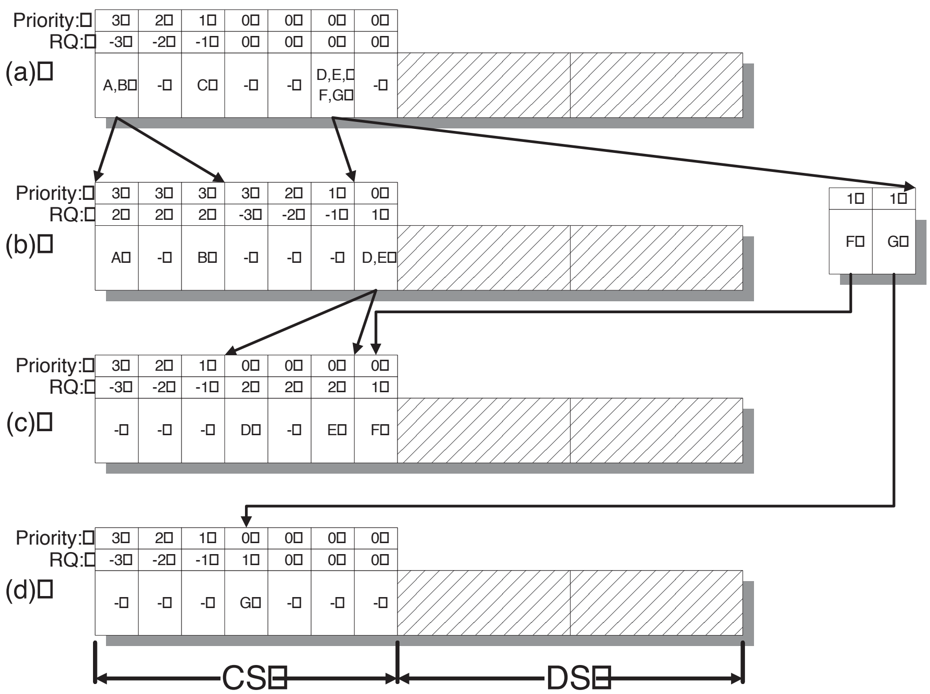

Figure 3.3: Priority Collision Resolution

3.4 Performance Evaluation

We built a simulation to evaluate the performance of the priority system. The implementation was created as part of an HFC module for the NIST ATM simulator[20]. We used the configuration and system parameters for the HFC network shown in Table 3.2. All simulations, with the exception of Experiment 5, were run for 10 seconds of simulated time and the first 10% of the data was discarded. We present the results from five different simulation experiments that measure the effectiveness of the priority system using mean request delay, request delay variation and transient throughput. A summary of the experiments is shown in Table 3.3. In all simulations the maximum number of priority levels is set to three.

- In Experiment 1 we quantify the overhead caused by the allocation of PNA slots in each frame in a lightly loaded network. We compare the request delay for low priority traffic in a system with PNA slots to a system without PNA slots. Since no higher priority traffic is present, this experiment evaluates the amount of overhead due to just the priority system.

- In Experiments 2 and 3 we show the impact of increasing the load of one priority on the request delays of the other priorities. Experiment 2 varies medium priority load and Experiment 3 varies the high priority load. As the traffic from a particular priority is increased, traffic from lower priorities is expected to be preempted. At the same time, high priority traffic should not be effected.

- In Experiment 4 we evaluate the bandwidth that should be reserved for priority newcomer stations. We verify that our selection of one PNA slot per frame is sufficient. Priority stations are given only one newcomer slot, while low priority stations are given the remaining CS in a frame. Typically, this is more than one slot and this experiment verifies that priority traffic still receives lower request delays.

- In Experiment 5 we evaluate how fast our priority scheme can preempt lower priority traffic if higher priority traffic becomes active. We also verify that the priority system is fair within a priority level.

Simulation Parameter	Values
Distance from nearest/furthest station to headend	25/80 km
Downstream data transmission rate	Not considered limiting
Upstream data transmission rates	3 Mbits/sec
Propagation delay	5 μs/km for coax and fiber
Length of simulation run	10 sec
Length of run prior to gathering statistics	10% of simulated time
Guard-band and pre-amble between transmissions	Duration of 5 bytes
Data slot size	64 bytes
CS size	16 bytes
DS/CS size ratio	4:1
Frame size	52 slots
CS	Fixed 18 slots
Round-trip	1 Frame
Maximum request size	32 data slots
Headend processing delay	1 ms

Table 3.2: Simulation Parameters

	Low Priority		**Medium Priority**		**High Priority**	
Experiment	Stations	Load	Stations	Load	Stations	Load
Protocol Scheme Overhead	80	[5%,45%]	0	0	0	0
Varying Medium Priority	100	20%	80	[10%,45%]	20	5%
Varying High Priority	50	12.5%	50	12.5%	100	[10%,45%]
Low Load Performance	50	[2.5%,35%]	50	[2.5%,35%]	0	0
Transient Throughput	50	100%†	100‡	100%†	50	100%†

Table 3.3: Simulation Scenarios
† Continuous backlog. ‡ 50 in Group 1 and 50 in Group 2

Each simulation scenario and its results are presented here:

3.4.1 Experiment 1: Overhead of the Priority Scheme

In Experiment 1 we quantify the system overhead for a system that sends all traffic at the same (lowest) priority level. We compare two cases. In the first case, the PNA slots are not present and the low priority stations can use the entire range of CS, which corresponds to the current 802.14 MAC draft[25]. In the second case, three contention slots are marked as PNA slots for higher priorities, therefore stations can only use part of the CS in the frame. We plot the average request delay and coefficient of variation versus traffic load in Figure 3.4(a) and 3.4(b) respectively. Figure 3.4(a) shows that the reserved PNA slots cause only a slight increase in request delay while Figure 3.4(b) shows a similar increase in the coefficient of delay variation (note that the coefficient of request delay variation is the ratio of the the request delay standard deviation to the mean of the request delay[28]). This quantifies the minimum increase in delay for lower priority traffic. Note that larger delays may be incurred in the presence of other priority traffic because higher priority traffic will resolve its collisions first. The small increase in average request delay and delay variation are limited and have to be weighed against the benefits of a structured priority system.

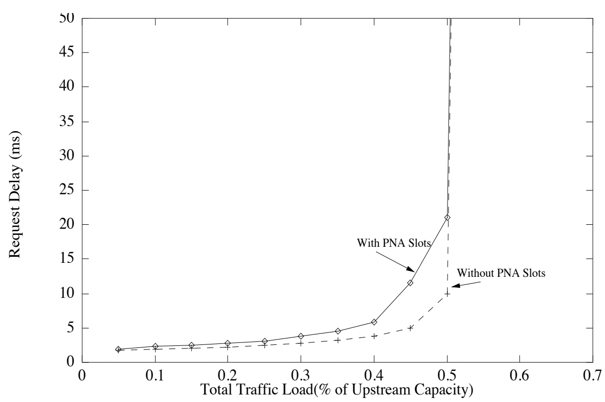

(a) Average Request Delay

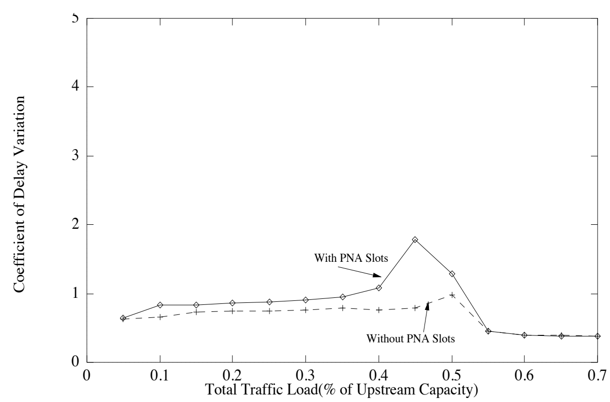

(b) Coefficient of Request Delay Variation

Figure 3.4: PNA Overhead

3.4.2 Experiment 2: Varying Medium Priority Load

Here we show the effect of the load from a particular priority level on other priority levels. In the experiment a total of three priority levels are used. There are 20 high priority stations which contribute 5% of the channel capacity to the load and 100 low priority stations which transmit at a total load of 20%. 80 medium priority stations are introduced to the system that generate a load that is varied from 10% to 45%. In Figure 3.5 we plot the request delay versus load for each priority level. We observe that as the medium priority traffic increases, the headend allocates more CS for the medium priority contention and less for the low priority stations. This causes the delay for low priority traffic to increase. This results in a relatively flat request delay for the medium priority stations. The high priority stations retain the same average request delay at any medium priority load. This reflects the robust operation of the priority scheme.

3.4.3 Experiment 3: Varying High Priority Load

Experiment 3 shows the effect of varying the load of the high priority stations. There are three groups of stations. Two of the groups consist of low and medium priority stations and each group consists of 50 stations at 12.5%. The third group consists of a large number of high priority stations (100) which transmit at loads that is varied from 10% to 45%. Figure 3.6 shows that as the high priority stations' load increases, the low priority stations are delayed. Then, as the load increases further request delays for medium priority stations increase as well which results in low request delays for

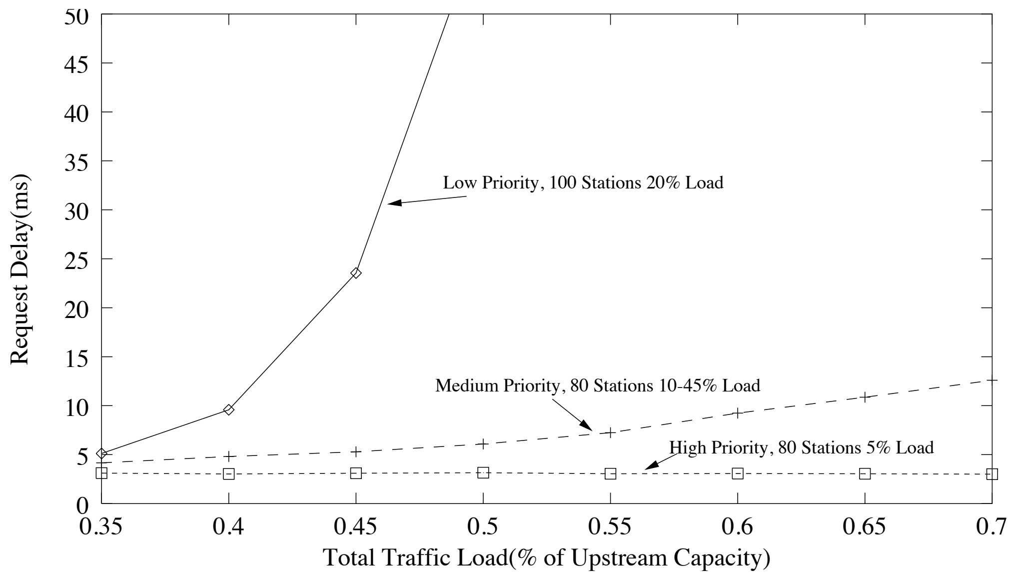

Figure 3.5: Varying Medium Priority Load

the high priority stations. Although it is unlikely that a system would be operated with such a large amount of high priority traffic, the high priority stations still receive a flat request delay.

3.4.4 Experiment 4: Low Load Performance

Experiment 4 compares the performance of low priority traffic, which is given multiple newcomer slots, to higher priority traffic, which is given one PNA slot per priority. Two sets of stations transmit in the system, one group with 50 stations at low priority and the other with 50 stations at medium priority. In Figure 3.7a one PNA is allocated per frame for medium priority requests. We observe that the medium priority traffic has slightly higher request delay than the low priority traffic (about 1 ms between 5% and 45% load). This can be attributed to the fact that medium priority newcomer

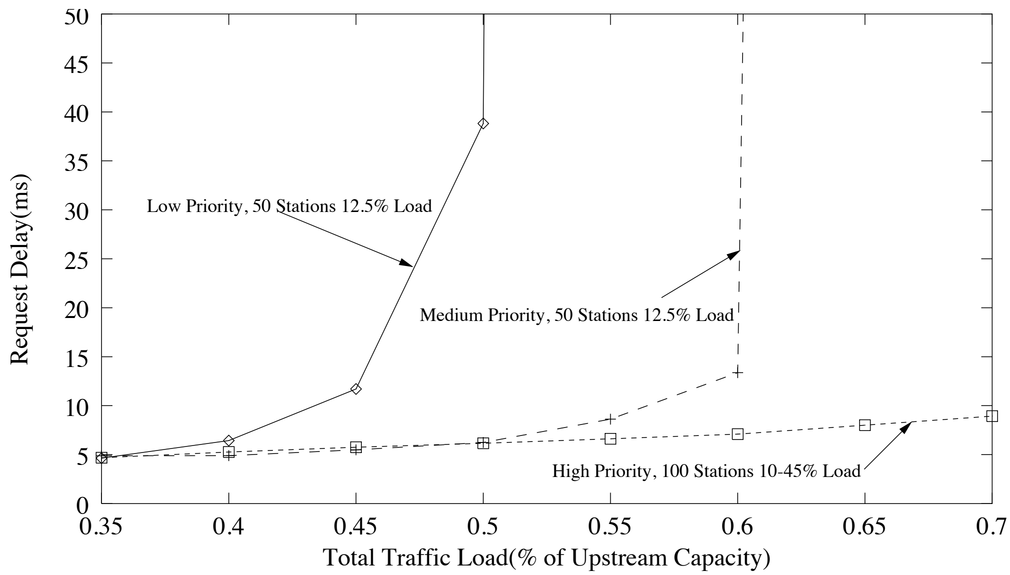

Figure 3.6: Varying High Priority Load

stations confined to only one PNA slot, while the remaining CS are used by the low priority traffic. At low loads, since collisions are infrequent, the request delay is mostly comprised of the time to transmit the first request. At higher loads, (above 45%) the request delay is mostly attibuted to collision resolution. This shows that one PNA for priority traffic newcomers may not be sufficient at low loads. This can easily be corrected, as shown in Figure 3.7(b), where 5 PNA slots are allocated to medium priority traffic. Note that the protocol is flexible to accommodate different priority traffic mixes and is not limited to one PNA slot per priority. If the headend controller knows that a large amount of high priority traffic will be sent, then the number of high priority newcomer slots can easily be increased.

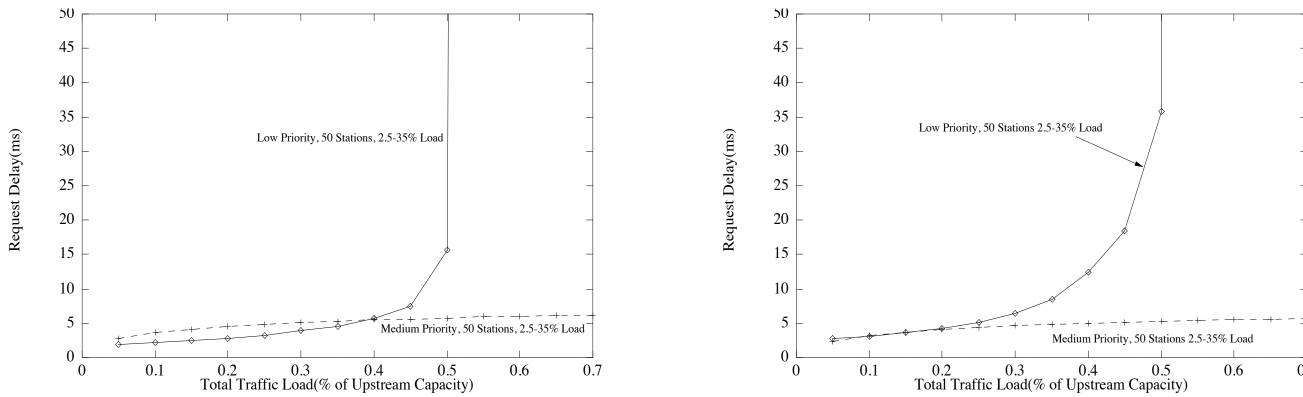

(a) 1 PNA for Medium Priority (b) 5 PNAs for Medium Priority

Figure 3.7: Low Load Performace

3.4.5 Experiment 5: Transient Throughput

In Experiment 5 we show the transient performance of the protocol. The experiment measures the throughput attained by the stations of a priority class per roundtrip delay. The entire experiment measures the throughput values over a total length of 350 round trips. At the beginning of the simulation the entire upstream bandwidth is occupied by users of the lowest priority. After 150 roundtrip delays, a group of medium priority traffic stations begins to transmit. A second group of medium priority stations and a group of high priority stations begin to transmit at 175 and 225 roundtrip delays respectively. Figure 3.8 shows throughput measurements, taken at one frame, or one roundtrip, intervals. A comparison of Figures 3.8(a) and 3.8(b) shows that the medium priority stations can preempt the low priority traffic within one or two roundtrip delays. Figure 3.8(c) shows that when the second group of medium priority stations is added, both groups share the bandwidth equally, which shows that the system is fair within a priority level. High priority stations can preempt all lower priorities immediately, as shown in Figure 3.8(d).

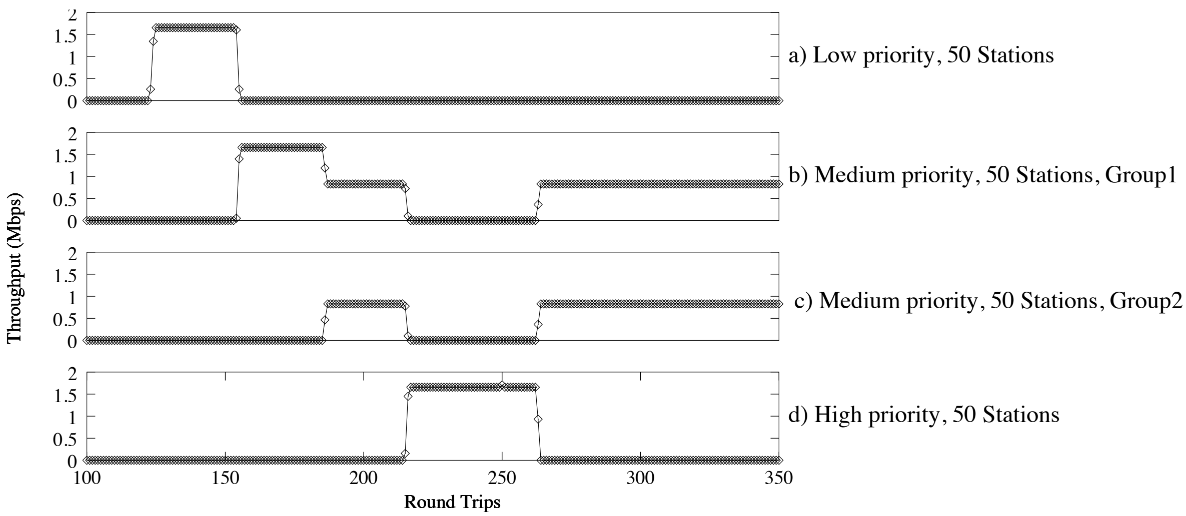

Figure 3.8: Transient Throughput

3.5 Analytical Results

In this section we analyze the collision resolution process of our priority scheme. Note that our scheme resolves collisions in strict priority order, that is, all collisions from a high priority level are resolved before any low priority collisions. Since the collision resolution is performed independently at each level, we can investigate the length of the resolution process for each priority level in isolation. The following analysis is done for an arbitrary, but fixed, priority level.

Assume that a collision occurs in a PNA slot of a priority level. (The use of a PNA slot precludes the analysis to all but the lowest priority level). Then, the ternary tree CRP allocates three CS in the next frame to resolve this collision. The allocation of additional CS can be thought of "splitting" the slot which contains a collision into three new slots in the next frame. If another collision occurs in one of these three CS, another three new slots will be generated in the next frame, and so on. Therefore, the entire collision resolution process can be represented as a tree: an internal node represents a slot which contains a collision, and a leaf node is a slot with a successful transmission. The height of the tree represents the duration of the collision resolution, and each level of the tree corresponds to one frame time. Next, analogous to [35][42][37], we derive expression for the width of the collision resolution tree as a function of the frame times after a collision.

We make the following assumptions:

- We assume that a collision occurs in a single PNA slot at some priority level.
- After the collision, no other station transmits requests until the collision is resolved. Subsequent collisions are resolved afterwards and do not affect the current resolution.
- The frame always contain a sufficient number of CS for the collision resolution. In other words, the width of the tree is smaller than the number of available CS. The results of our analysis verify that this assumption is valid; typically, a collisions does not require a large number of CS.

We denote the width of the tree k frame times after a collision has occurred between n users (of a fixed priority level) to be $W_n(k)$. We use $k = 0$ to denote the frame in which the collision occurs. If zero or one users transmit a request at the same time, the CRP is not started, and we obtain a tree with width equal to zero:

$$W_0(k) = W_1(k) \quad = 0 \quad \text{(for all k)} \tag{3.1}$$

Since, per assumption, the first collision ($k = 0$) occurs in a single slot and since, per definition of the ternary tree protocol, the second frame ($k = 1$) contains 3 slots for that collision, we obtain:

$$W_n(0) \quad = 1 \quad (n > 1) \tag{3.2}$$

$$W_n(1) \quad = 3 \quad (n > 1) \tag{3.3}$$

Given n users and m slots, the probability that i of them pick a particular slot is given by:

$$Q_i(n, m) = \binom{n}{i}\left(\frac{1}{m}\right)^i\left(1 - \frac{1}{m}\right)^{n-i} \tag{3.4}$$

Then, for each iteration, we can compute the expected width of the contention resolution tree, given the number of collided slots in the previous step, via the following recursive formula.

$$W_n(k) = \sum_{i=0}^{n}(Q_i(n,3)W_i(k-1) + \sum_{j=0}^{n-i}(Q_i(n,3)Q_j(n-i,2)(W_j(k-1) + W_{n-i-j}(k-1)))) \tag{3.5}$$

Equation (3.5) together with the base cases from Equation (3.1) and (3.3) allows us to compute the width of the collision resolution tree. In Figure 3.9(a) we plot the results obtained with Equation (3.5). The figure shows the expected width of the collision resolution tree in consecutive frame times if a collision occurs at frame time '0'. In Figure 3.9(a) we plot the graphs when $n = 2, 3, 4, 10$ stations collide at frame time '0'. We see that the expected width of the collision tree is limited even in the unlikely event that a large (n=10) number of stations are involved in a collision. Since the number of CS in a frame is large (we used a 'typical' value of 18 in our simulations), the width of the tree does not exceed the number of available CS in a frame. Also note that the number of frame times for which the tree width is nonzero can be interpreted as the length of the collision resolution process. Figure 3.9(a) shows that collisions are resolved quickly. Even for $n = 10$ the expected tree width in the 7-th frame after a collision is almost zero.

For verification purposes, we show results from our HFC simulator for the average number of CS used in each frame after a fixed collision size in Figure 3.9(b). Since the analysis is exact and does not depend on any stochastic assumptions, the match of simulation and analysis is not unexpected.

From the results shown in Figure 3.9 we note that, on average, the maximum number of CS slots used in any one frame to resolve a priority collision of ten stations,

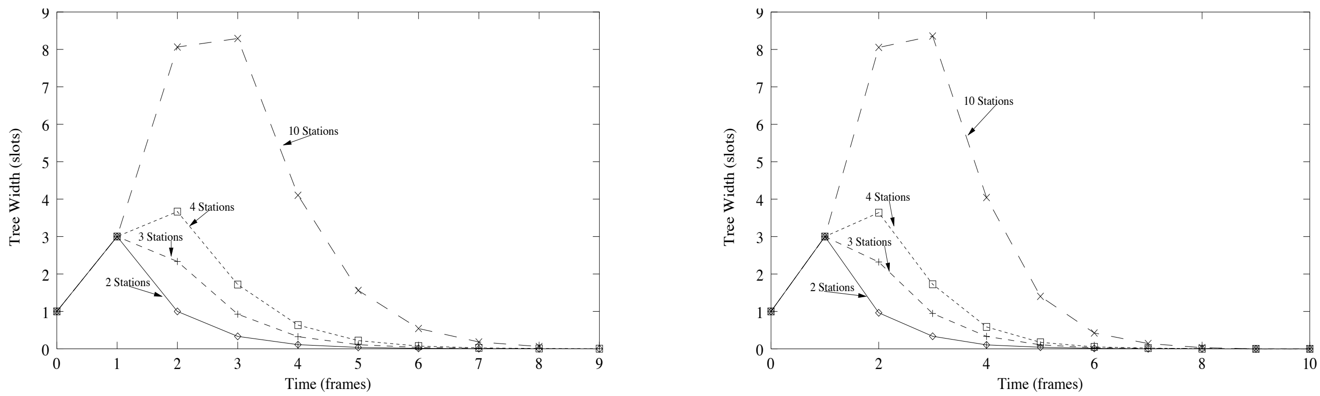

(a) Analytical Result (b) Simulation

Figure 3.9: Tree Collision Resolution "Width"

is nine. If we assume that more than 10 CS are allocated in a frame, collisions in PNA slots involving up to ten users will not block stations of lower priorities.

3.6 Concluding Remarks

In this chapter we have shown the need for a priority system for the IEEE 802.14 MAC protocol. We have contributed the design of a multilevel priority system that can easily be integrated into the current specification using the existing RQ numbering system. The protocol gives immediate access to stations of high priorities and separates and prioritizes collision resolution for different priority levels. The protocol has low overhead and we have shown its robustness and fairness with a wide variety of traffic mixes and priority levels. We have also presented analytical findings that describe the contention space needed by priority based collisions. The scheme can be easily incorporated into the 802.14 MAC standard and enhances the network's ability to provide users with Quality of Service.

4

Conclusions

In this thesis we have presented two methods to improve the Quality of Service offered by the IEEE 802.14 MAC protocol for Hybrid Fiber-Coax networks. In the second chapter we presented the results of an interoperability study between the ATM Available Bit Rate service and the 802.14 MAC. It was shown that congestion on the HFC upstream channel may interfere with the ABR's flow control mechanism. The chapter includes a mechanism for the HFC headend to provide extra congestion feedback information to ensure fairness and congestion free operation. The third chapter shows that the use of a priority scheduler is an insufficient solution for providing preemptive priorities to time sensitive data. We show that a new priority system integrated with the collision resolution algorithm can provide priorities to improve Quality of Service. Both of these improvements to the emerging HFC standard will enable home users to effectively use time sensitive services such as video teleconferencing, telephone service and interactive applications.

In addition to the work done in this thesis, many areas are left to be investigated. The Multimedia Cable Network System (MCNS) Corporation is promoting an industry standard called the Data-Over-Cable Service Interface Specifications (DOCSIS). They have been successful in getting it standardized by the Society of Cable Telecommunications Engineers (SCTE) and by the Internet Engineering Task Force (IETF). This competing standard uses much of the same physical layer but employs a different upstream messages and uses a binary exponential back off to resolve collisions. We have completed a comparison of the two standards that will be presented at the LANMAN workshop this year in Banff. Unfortunately this work was not complete in time for inclusion in the thesis. The DOCSIS standard may be used more in the future because of its support from major cable operators such as Time Warner.

Although the DOCSIS standard may be used by cable operators and cable equipment manufacturers, some additions will be needed to provide Quality of Service. Future work in this area may involve formulating new allocation algorithms for CS and DS to provide low access delays and high throughput. Dynamic allocation of bandwidth to contention and data may also improve delay variation. Also, future work may investigate methods to incorporate the structured collision resolution algorithm found in 802.14 into the DOCSIS standard. 802.14 can provide the basis for a second version of the standard to better support ATM traffic and QoS requirements.

A

Appendix

A.1 Simulation Environment used in Chapter 2

Here we present to some greater level of detail, the simulation parameters of our simulator for ATM networks and HFC networks.

The MAC simulation parameters are set according to Table 1. Table 2 describes the simulation parameters used for ABR sources, and Table 3 describes the parameters for ABR switches. The buffers sizes of all ATM switches are limited to 10,000 cells.

Simulation Parameter	**Values**
Distance from nearest/furthest station to headend	25/80 km
Downstream data transmission rate	Not considered limiting
Upstream data transmission rates (aggregate for all channels	8.192 Mbits/sec
Propagation delay	5 μs/km for coax and fiber
Length of simulation run	10 sec
Length of run prior to gathering statistics	10% of simulated time
Guard-band and pre-amble between transmissions from different stations	Duration of 5 bytes
Data slot size	64 bytes
CS size	16 bytes
DS/CS size ratio	4:1
Frame size	2.27 ms (Max 160 CSs)
Maximum request size	16 data slots
Number of Contention Slots	20
Number of Data Slots	40
Headend processing delay	1 ms

Table A.1: MAC Parameters.

Simulation Parameter	**Values**
Number of ABR sources	12 (Upstream Configuration)
	12 (Downstream Configuration)
Number of CBR sources	None (Upstream Configuration)
	50 (Downstream Configuration)
CBR Parameters	
Peak Cell Rate (PCR)	0.13 Mbits/s
ABR Parameters	
Nrm (Number of RM cells)	16
Available ABR Bandwidth on Congested Link	6 Mbits/s
Link Cell Rate	149.76 Mbits/s
Allowed Cell Rate (ACR)	Dynamically adjusted
Initial Cell Rate (ICR)	0.5 Mbits/s
Peak Cell Rate (PCR)	2 Mbits/s
Minimum Cell Rate (MCR)	0.149 Mbits/s
Rate Increase Factor (RIF)	0.063
Rate Decrease Factor (RDF)	1/16

Table A.2: ATM End System Parameters.

Simulation Parameter	**Values**
Maximum Buffer Size	10,000 cells
Explicit Forward Congestion Indication Switch	
High Threshold	225 cells
Low Threshold	200 cells
Explicit Rate Switch	
High Threshold	15 cells
Low Threshold	10 cells
Target Rate (TR)	10 Mbits/s
Average Factor (AVF)	1/16
Mean ACR Additive Increase Rate (MAIR)	0.015 Mbits/s
Mean ACR Reduction Factor	0.95
Measurement Interval (N)	100 cells

Table A.3: ABR Switch Parameters.

A.2 Brief notes on the use of the HFC module

Please note that the HFC module is highly experimental and complex. Any user who wishes to use this module should at be aquatinted with Hybrid Fiber Coax Networks and MAC protocols. The user should also be willing to examine the code to familiarize ones self with the details of the simulation.

To add an HFC module to the simulation the user should select shift $->$ left mousebutton $->$ BTE $->$ HFC Module. The user selects a placement for the HFC module and connects it to other components in the same manner as any other new component. This section explains the general operation and architecture of the HFC module, the parameters used in the module and which algorithms the code contains.

The HFC module is a special instance of the BTE module. Although the HFC module is based on the code for the BTE, it differs significantly in its usage. The HFC module contains code that simulates the activity of an entire HFC cable tree . For this reason the module will be referred to in the manual as the HFC network. The network represents the combination of the cable modem access units, the bi-directional cable tree, the media access control protocol (MAC) and the head-end. Some of these elements are actual physical devices, while some are protocols and processes. These are further explained here:

Cable Tree. A residential HFC network is a tree structure with users connected at varying distances from a headend, at the top of the cable tree. The system is bi-directional with an upstream and downstream channel. Each user connects to the HFC tree through the use of a cable modem. This allows the user to communicate with the headend, which receives messages from all users. Cable Modem. Each user is represented by a cable modem inside the HFC network. Each time a connection (ABR, TCP, VBR, etc.) is attached to the network, the HFC module creates an internal queue, called the middle queue. This middle queue represents each cable modem, or station, on the network. Each new connection creates a new queue, so multiple connections cannot be attached to the same station. MAC Protocol. The upstream channel is divided into discrete time slots called minislots. The headend designates each one of these slots as either a contention slot (CS) or as a data slot (DS). The stations use the CS to transmit requests for the reservation of a data slot. The CS are subject to collisions when more than one user transmits at a time. DS are explicitly allocated to stations by the headend and are therefore collision free. A frame is a collection of CS and DS. The headend designates the type of each slot in the frame. Access to the CS and the resolution of collisions is handled by the collision resolution protocol (CRP).

Collision Resolution. Several different CRP can be simulated using the HFC simulator. The protocols can be broken into two classes, continuous mode and cluster mode. Contributions on the differences between continuous and cluster mode are available. During cluster mode operation, all contention slots (CS) are allocated at the beginning of a frame, and the rest of the frame is allocated to data slots (DS). In continuous mode operation the headend allocates two CS for every DS in a continuous stream. This is also called frameless operation. Further explanation of each MAC protocol can be found in the parameter section below.

Head-End. A controller at the head of the cable tree, called the headend, coordinates the MAC protocol and the transmission of data. Like the BTE, the headend MUST be connected to an ATM switch to function. The headend uses the ATM switch to implement rate control for ABR sources connected to the network. The headend receives each CS and DS, processes the information contained in them and provides feedback information to the stations.

Adding an HFC Node. When adding a BTE station to the simulation, the simulation gives the user a choice between a traditional BTE and a Hybrid Fiber Coaxial (HFC) module. The HFC module models the entire HFC network rather than an access point for a single user (such as the BTE). Each connection attached to the HFC module represents one user on the cable network. The HFC module uses a series of queues to accurately simulate the operation of a real network.

Connecting Sources. The same connections that can be used with the BTE can be used with the HFC network, including rate controlled ABR sources, TCP sources and self similar traffic. Remember that the HFC node must be connected to an ATM switch to operate correctly.

All traffic sources that are connected to the HFC network are classified into two classes of traffic, ABR traffic and non-ABR traffic. Three queues are used for each ABR connection, an input queue, a middle queue and an output queue. Only the middle queue and the output queue are used for non-ABR traffic. When data arrives from an ABR connection it is put into the input queue for that connection. Cells are removed from the input queue and put into the middle queue at the rate 1/ACR for that ABR connection. Cells that arrive from a non-ABR connection are immediately placed in the middle queue.

The MAC mechanism only operates on the cells found in the middle queue. When data arrives in the middle queue for a connection, a request is transmitted. When the request is acknowledged and a data slot is allocated to the station, the cell is transmitted to the headend. The headend takes the cell and places it in the output queue. The output queues from thew HFC network connect to a link to an ATM switch. There are two output queues, one for ABR and one for VBR(or other non-ABR traffic)). These two output queues are shared between all stations on the network.

When setting routes for connections attached to the HFC network, the user MUST link an ATM switch to the HFC module and all routes must follow a path similar to this: Connection − > HFC − > Link − > ATM switch − > (any Link).

The parameters that differ from the BTE are explained here.

Max. Queue Size. There are three queues. The Input queue resides between the source and the HFC network. The middle queue is only used for ABR connections and cells move from the input queue to the middle queue at the Allowed Cell Rate (ACR). The output queue buffers cells between the HFC headend and the ATM link.

Limit Requests to x. When stations request data slots in the upstream channel they may only have one outstanding request at a time. The stations are limited to requesting this many data slots at a time.

Collision Multiplicity. This is the number of users that transmitted in the most recent contention slot. 0=empty, 1= sucessful request, ¿2 is a collision.

Access Delay. This is the delay from the time that the cell is received at the input queue to the time it is transmitted on the upstream channel.

Head-End Processing Time. This is the time that the head-end uses to process requests and collisions.

Cluster Size. The number of minislots in a frame (or cluster). Cluster mode protocols allocate all the contention slots at the beginning of the cluster. Continuous mode protocols distribute them over the cluster.

Congestion Thresholds. When the average grant queue size exceeds the high threshold a message is sent to the ATM switch that congestion is present in the HFC network. When the average drops below the low threshold a message is sent to indicate no congestion.

Grant Queue Length. This is the number of data slots in the grant queue before grants are issued by the head-end.

Average Grant Queue Length. This is a Weighted Moving Average of the Grant Queue size. It uses a constant of 1/16 to compute the average. Interelaving Factor. Separate colllision resolution engines can be used in the same system. This can increase the number of contention slots that can be used in a frame during cluster mode operation. The normal value of this paramter is 1.

MAC Protocol. This is the MAC and collision resolution protocol used by the system. Only some of the paramters apply to each protocol. The protocols are as follows:

0- Ternary Tree Blocking (Cluster Mode) 1- Ternary Tree Free Access (Cluster Mode) 2- Ternary Tree Unblocking (Cluster Mode) (IEEE 802.14 Standard) 3- One

Stage P-Persistance (Cluster Mode) 4- Two Stage P-Persistance (Cluster Mode) 5- Bitmap (Cluster Mode) 6- Ternary Tree Unblocking (Continuous Mode) 7- One Stage P-Persistance (Continuous Mode)

Num CS (if ¡= 0 variable allocation). This is the numebr of contention slots per cluster. If it is set to ¡=0, a variable allocation is used. (See code for exact algorithm)

Maximum Number of CS (if ¡= 0 computed). Maximum number of contention slots per frame. This is ignored if the number of CS is fixed. If it is less than 0, it is computed to be the max it can be without violating the rule that the feedback ust arrive before the next cluster starts.

Newcomer Range (R) (if ¡= 0 variable). The Range paramter R can be fixed to a value or it can be optimized by an algorithm (see code).

Retry Range (P) (if ¡= 0 variable). For P-persistance, retrys use this number in place of R. Can be computed in the same manner as R.

Noise factor (fractional). This fraction of succesful requests will be interpreted as collisions by the headend.

Bitmap Density (Stations/MS). Only used by the bitmap MAC protocol. The density is the number of stations than can transmit requests per contention slot.

Upstream/Downstream Trans. Time (usecs). The time it takes to transmit one minislot on the channels.

CS per DS. The size of the data slots in terms of contention slots. Typical value would be 4.

Bibliography

[1] R. Agarwal. A Simple and Efficient Multiple-Access Protocol. Contribution No. IEEE 802.14/95-136, IEEE 802.14 Working Group, November 1995.

[2] ALOHA. Packet system with and without slots and capture. *Computer Communications Review*, April 1975.

[3] ATM Forum, ATM Forum Traffic Management Specification Version 4.0, April 1996.

[4] A. Azzam. *High-Speed Cable Modems.* McGraw-Hill, 1997.

[5] D. Bertsekas and R. Gallager. *Data Networks, 2nd Ed.* Prentice Hall, 1992.

[6] C. Bisdikian. MLAP: A MAC Level Access Protocol for the HFC 802.14 Network. Technical Report RC 20284, IBM Research Division, December 1995. T.J. Watson Research Center.

[7] C. Bisdikian. A Review of Random Access Algorithms. Technical Report RC 20348, IBM Research Division, T.J. Watson Research Center, January 1996.

[8] C. Bisdikian. msSTART: A Random Access Algorithm for the IEEE 802.14 HFC Network. Technical Report RC 20466, IBM Research Division, T.J. Watson Research Center, June 1996.

[9] C. Bisdikian, B. McNeil, and R. Zeisz. MLAP: A MAC Level Access Protocol for the HFC 802.14 Network. *IEEE Communications Magazine*, 34(3):114–121, March 1996.

[10] A. Carleial and M. Hellman. Bistable Behavior of ALOHA-Type Systems. *IEEE Transactions on Communications*, COM-23(4):401–409, April 1975.

[11] R. Chiu and R. Hutchinson. Multimedia Transmission Link Protocol: A Proposal for the 802.14 MAC Lsyer. Contribution No. IEEE 802.14/95-160, IEEE 802.14 Working Group, November 1995.

[12] R. Citta and D. Lin. Formal MAC Layer Protocol Proposal: Adaptive Random Access Protocol for CATV Networks. Contribution No. IEEE 802.14/95-144, IEEE 802.14 Working Group, October 1995.

[13] M. DePrycker. *Asychronous Transfer Mode: Solution for B-ISDN.* 1993.

[14] IEEE Std 802.6-1990, IEEE Standards for Local and Metropolitan Area Networks: Distributed Queque Dual Bus (DQDB) of a Metropolitan Area Network, July 1991.

[15] J. Eng and J. Mollenauer. IEEE Project 802.14: Standards for Digital Convergence. *IEEE Communications Magazine*, 33(4):20–23, April 1995.

[16] B. McNeil et. al. Formal Proposal for 802.14 MAC Protocol (part 2 of 2) MLMP (MAC Layer Management Protocol). Contribution No. IEEE 802.14/95-158, IEEE 802.14 Working Group, November 1995.

[17] C. Bisdikian et. al. Cable access beyond the hype: On residentail broadband data services over hfc networks. *IEEE Communications Magazine*, 34(11):128–135, November 1996.

[18] N. Achilleoudis et. al. A Multi-Service Interactive System for Hybrid Fiber-Coaxial Networks. Contribution No. IEEE 802.14/95-155, IEEE 802.14 Working Group, November 1995.

[19] N. Golmie et. al. Performance Evaluation of Contention Resolution Algorithms: Ternary Tree vs p-Persistance. Contribution No. IEEE 802.14/96-241, IEEE 802.14 Working Group, November 1996.

[20] N. Golmie et. al. The NIST ATM Network Simulator: Operation and Programming. National Institute of Standards and Technology, Internal Report 5703, August 1995.

[21] S. Lam et. al. A MAC Layer Proposal for IEEE 802.14 WG. Contribution No. IEEE 802.14/95-164, IEEE 802.14 Working Group, November 1995.

[22] Y. Chang et. al. A Rate Based Flow Control Switch Design for ABR Service in an ATM Network. In *Proceedings of the Twelfth International Conference on Computer Communications*, Seoul, Korea, August 1995.

[23] M. Gagnaire. An Overview of Broad-Band Access Technologies. *Proceedings of the IEEE*, 85(12):1958–1972, December 1997.

[24] D. Grossman. Signaling Requirements for ATM over HFC. Contribution No. IEEE 802.14/97-008, IEEE 802.14 Working Group, January 1997.

[25] IEEE 802.14 Working Group. Media Access and Control. Draft Supplement to IEEE Std 802.14, IEEE 802.14 MAC/V1.1, IEEE 802.14 Working Group, December 1996.

[26] R. Handel and M. Huber. *Integrated Broadband Networks*. Addison-Wesley, 1991.

[27] V. Hou. 802.14 MAC Layer Proposal. Contribution No. IEEE 802.14/95-129, IEEE 802.14 Working Group, November 1995.

[28] L. Kleinrock. *Queueing System. Volume 1:Theory*. John Wiley and Sons, 1975.

[29] F. Koperda. Providing CBR Services over XDQRAP. Contribution No. IEEE 802.14/95-147, IEEE 802.14 Working Group, November 1995.

[30] F. Koperda and B. Lin. A Proposal to Use XDQRAP for the IEEE 802.14. Contribution No. IEEE 802.14/95-068, IEEE 802.14 Working Group, July 1995.

[31] M. Laubach. ATM HFC Overview. Contribution No. IEEE 802.14/95-022, IEEE 802.14 Working Group, March 1995.

[32] N. Golmie M. Corner J. Liebeherr and D. Su. Improving the Effectiveness of ATM Traffic Control Over Hybrid Fiber-Coax Networks . 1997. Proceedings of Globecom 1997.

[33] H. Lin and G. Campbell. PDQRAP - Prioritized Distributed Queueing Random Access Protocol. Technical Report 93-2, Illinois Institute of Technology, Computer Science Department, March 1994.

[34] M. Corner, N. Golmie, J. Liebeherr, D. Su. A priority scheme for the ieee 802.14 mac protocol for hybrid fiber-coax networks. In *Proceedings of InfoCom '98*, March 1998.

[35] P. Mathys and P. Flajolet. Q-ary Collision-Resolution Algortithms in Random-Access Systems with Free or Blocked Channel Access. *IEEE Transactions on Information Theory*, 31(2):217–243, March 1985.

[36] B. McNeil. Implementation Overview of Tree-Based Algorithm. Contribution No. IEEE 802.14/96-246, IEEE 802.14 Working Group, November 1996.

[37] L. Merakos and C. Bisdikian. Delay Analysis of the n-Ary Stack Random-Access Algorithm. *IEEE Transactions on Information Theory*, 34(5):931–940, September 1988.

[38] A. Paff. Hybrid Fiber/Coax in the Public Telecommunications Infastructure. *IEEE Communications Magazine*, 33(4):40–45, April 1995.

[39] M. Prycker. *Asynchronous Transfer Mode: Solution for Broadband ISDN*. Ellis-Horwood, 1991.

[40] S. Ramanathan and R. Gusella. Toward Management Systems for Emerging Hybrid Fiber-Coax Access Networks. *IEEE Network*, 9(5):58–68, September/October 1995.

[41] R. Rivest. Network Control by Bayesian Broadcast. *IEEE Transactions on Information Theory*, IT-33(3):323–328, May 1987.

[42] R. Rom and M. Sidi. *Multiple Access Protocols: Performance and Analysis*. Springer Verlag, 1990.

[43] D. Sala, A. Mukherjee, and J. Limb. Protocol For Efficient Transfer Of Data Over Fiber/Cable Systems, June 1995. Unpublished, available from authors (currently at Georgia Insititue of Technology).

[44] P. Sriram. Adaptive Digitial Access Protocol (ADAPt): A MAC Protocol for Multiservices Broadband Acccess Networks. Contribution No. IEEE 802.14/95-168, IEEE 802.14 Working Group, November 1995.

[45] W. Stallings. *ISDN and Broadband ISDN*. Ney York: Macmillan, second edition, 1990.

[46] W. Stallings. *Data and Computer Communications, 4th Ed.* Macmillan Publishing Company, 1994.

[47] IEEE Std 802.5-1985, Token Ring Access Method, 1985.

[48] J. Ulm. A MAC Protocol for 802.14. Contribution No. IEEE 802.14/95-134, IEEE 802.14 Working Group, November 1995.

[49] P. van Grinsven et. al. An Example of a MAC based on the Convergenece Algorithm. Contribution No. IEEE 802.14/96-217, IEEE 802.14 Working Group, September 1996.

[50] J. Vandenameele. How to Upgrade CATV Networks to Provide Interactive ATM-Based Services. In *Proceedings of GlobeCom 95*, Singapore, November 1995.

[51] M. Vecchi. Broadband Networks and Services: Architecture and Control. *IEEE Communications Magazine*, 33(8):24–32, August 1995.

[52] C. Wu and G. Campbell. Extended DQRAP (XDQRAP): A Cable TV Protocol Functioning as a Distributed Switch. Technical Report 94-2, Illinois Institute of Technology, Computer Science Department, June 1994.

[53] R. Zeisz. Formal Proposal for 802.14 MAC Protocol (part 1 of 2) MLMP (MAC Layer Management Protocol). Contribution No. IEEE 802.14/95-157, IEEE 802.14 Working Group, November 1995.