ATM Traffic Control in Hybrid Fiber-Coax Networks – Problems and Solutions

Nada Golmie[†] *Mark D. Corner*[‡] *Jörg Liebeherr*[‡] *David H. Su*[†]

[†] NIST – National Institute of Standards and Technology
Gaithersburg, MD 20899

[*] Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109

[‡] Department of Computer Science
University of Virginia
Charlottesville, VA 22903

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 higher layer traffic services, namely, ATM Constant Bit Rate (CBR), Variable Bit Rate (VBR) and Available Bit Rate (ABR) traffic classes. In this study, we investigate the inter-operation of the MAC protocol, defined by 802.14, with ABR transmissions. An important finding of our study is that the bandwidth contention 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. Numerous simulation scenarios are presented to illustrate our findings. Parts of the results have been presented to the IEEE 802.14 standardization committee.

Key Words: Hybrid-Fiber Coax, IEEE 802.14, Cable Modems, ATM, Available Bit Rate, ABR, Community Networks.

1 Introduction

The IEEE 802.14 working group is currently standardizing a media access control (MAC) protocol for the emerging Hybrid Fiber Coax (HFC) networks for providing high bandwidth residential networking services. An HFC network (see Figure 1) utilizes the in-place residential broadcast cable system. While downstream communication from the headend to the stations is free of contention, the upstream channel from the stations to the headend is a shared access channel and subject to collisions. The 802.14 working group is currently defining a contention resolution protocol that controls access to the upstream channel and resolves collisions. Crucial for the success of 802.14 will be its ability to support higher layer traffic services, namely, ATM Constant Bit Rate (CBR), Variable Bit Rate (VBR) and Available Bit Rate (ABR) traffic classes.

This study explores the performance and inter-operation of the MAC protocol, as defined by the 802.14 specification, with the 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 [8]. Our current findings are as follows:

- *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.

The contribution of this study is a proposal to solve the unfairness problem of downstream ABR traffic. In our 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 we use the backlog of the so-called grant queue at the headend station as a load indicator. We will show how the proposed solution can be incorporated into the IEEE 802.14 [9] and the ATM Traffic Management [1] specifications, with no modifications to those standards.

The remainder of the paper is organized as follows. In Section 2 we provide an overview of the MAC protocol proposed by the IEEE 802.14 working group. In Section 3 we briefly review the ABR flow control mechanism. In Section 4 we discuss the performance of ABR transmissions over HFC networks. In Section 5 we present our scheme and we offer conclusions in Section 6. In Appendix A we present a more detailed description of the simulation environment used for performance evaluation and to demonstrate its effectiveness.

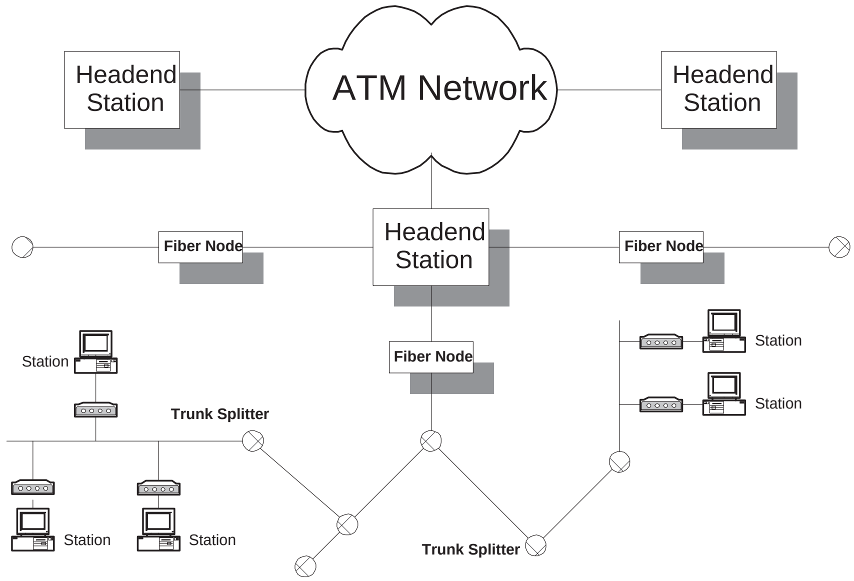

Figure 1: HFC Network Connected to an ATM Network.

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 rates equaling approximately 0.5-10 Mbps and 30 Mbps, respectively. The IEEE 802.14 MAC is only concerned with the transmission of data on the upstream channel.

Figure 2 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 has occurred. The latter initiates the collision 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.

The selection of the collision resolution algorithm for the contention slots on the upstream channel has received a lot of interest from the research community. The collision resolution scheme adopted by the IEEE 802.14 group is based on a blocking ternary tree splitting algorithm [3]. Tree splitting algorithms have been used in the past to improve the performance of collision access [2]. In a tree splitting algorithm, all stations that are involved in a collision split into a number of subgroups. After

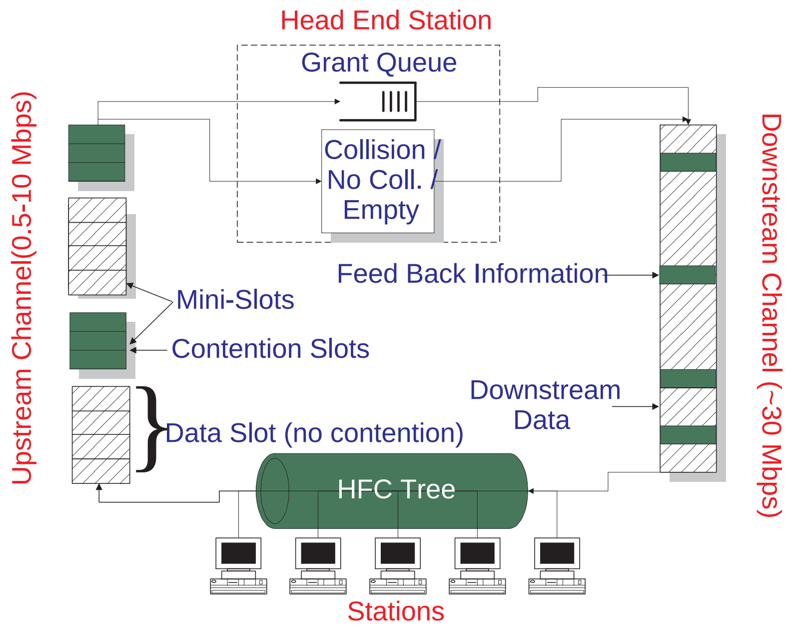

Figure 2: Media Access Control in HFC Networks.

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 [2]. The selection of a blocking algorithm for 802.14 is intended to reduce the MAC access delay variance [5].

The current version of the MAC protocol is heavily influenced by an adaption of the tree splitting algorithm, the n-ary Stack Random Access Algorithm [11, 3, 4]. The collision resolution scheme used in this paper is based on the status of January 1997 as reflected by [9, 12].

3 ABR Service Overview

The Available Bit Rate (ABR) service in ATM networks [1] 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* or ACR, of the connection to any value which satisfies $MCR \leq ACR \leq PCR$.

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 Management (RM) cells along with data cells to its destination

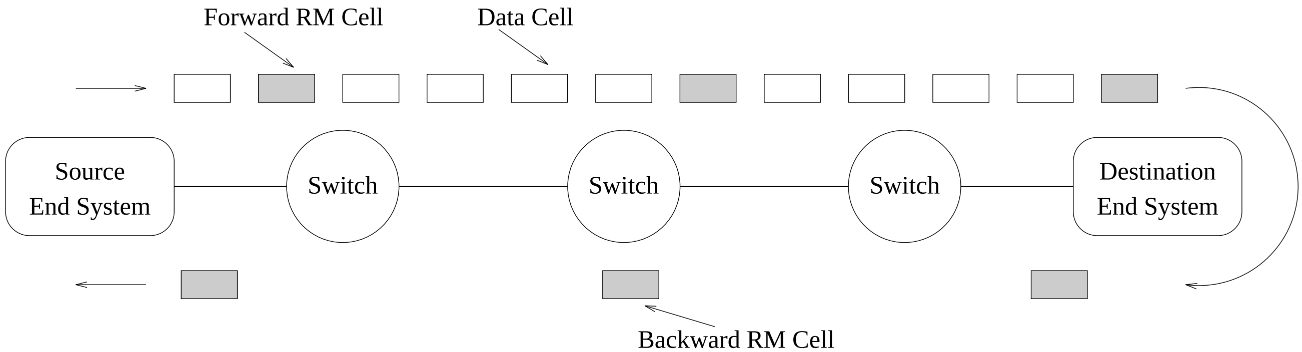

Figure 3: Closed-Loop Traffic Control.

(see Figure 3). 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 its 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 [1], 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 [7]; this mechanism attempts to achieve maximum network stability in terms of ACR and buffer occupancy oscillations. The ER mechanism in [7] was selected since it is, in comparison to other proposed solutions [10], relatively simple to implement. The results in this paper will not change significantly if a different ER algorithms is selected.

4 Effectiveness of ABR Flow Control Traffic 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 [1] 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

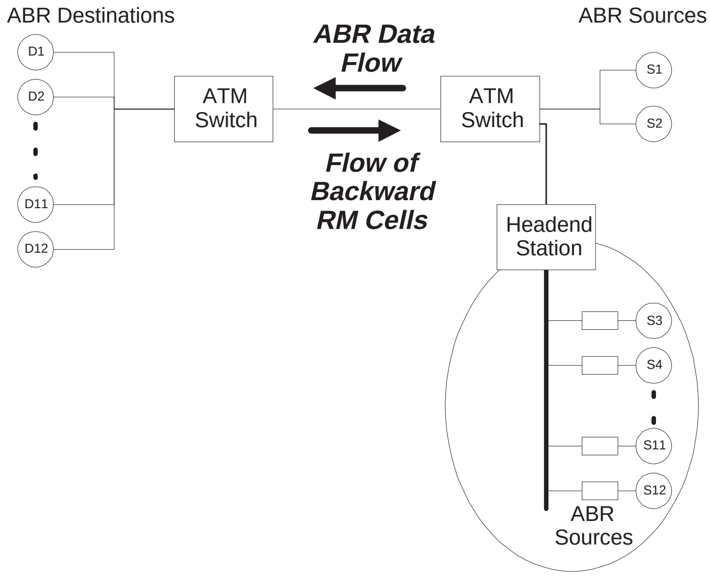

Figure 4: Upstream Configuration.

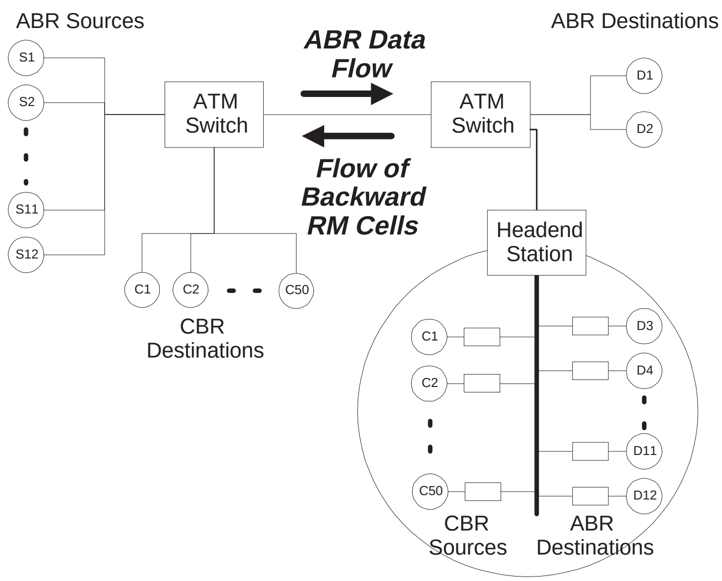

Figure 5: Downstream Configuration.

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.

4.1 ATM/HFC Simulation

We have built a simulation of an ATM/HFC network to measure the degree to which an HFC network can impact the effectiveness of ABR rate control in an ATM network. The implementation was done using the NIST ATM Network Simulator [6]. The ATM simulation package was extended by a module for an HFC network with an interface to ATM components.

The simulated network scenarios are depicted in Figures 4 and 5. 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 Mbps. 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 Mbps for the ABR applications and to 0.13 Mbps for CBR applications. 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 4 and 5 are referred to as Upstream Configuration and Downstream Configuration, respectively. In the Upstream Configuration (Figure 4) 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 5) we evaluate ABR connections that have the destination inside an HFC network. In both configurations, we consider both EFCI and ER switch control.

4.2 Upstream Transmissions

We first discuss the outcome of the simulations for the Upstream Configuration shown in Figure 4. 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 Mbps, we expect that each of the ABR sources obtains a fair share of 0.5 Mbps as end-to-end throughput.

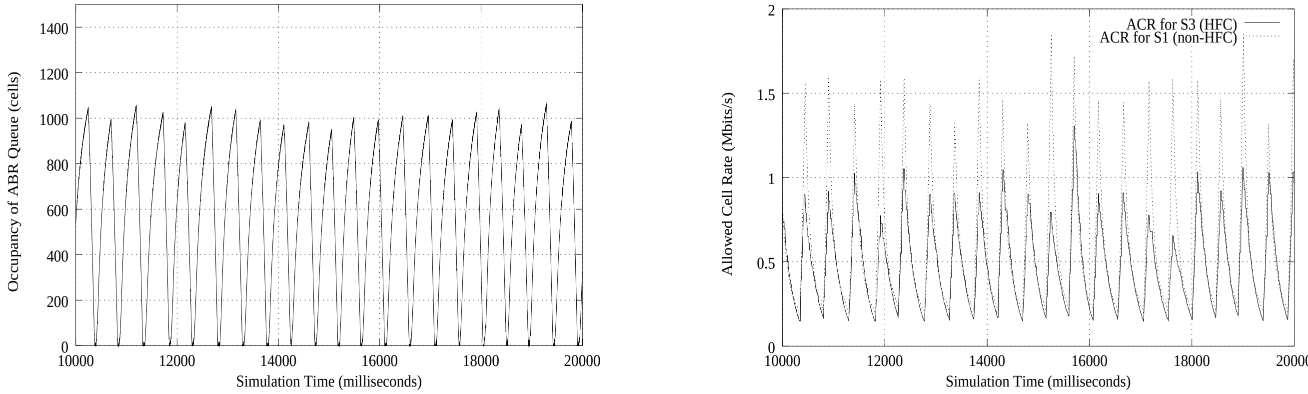

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 6: Upstream Transmissions: EFCI Control Mechanism.

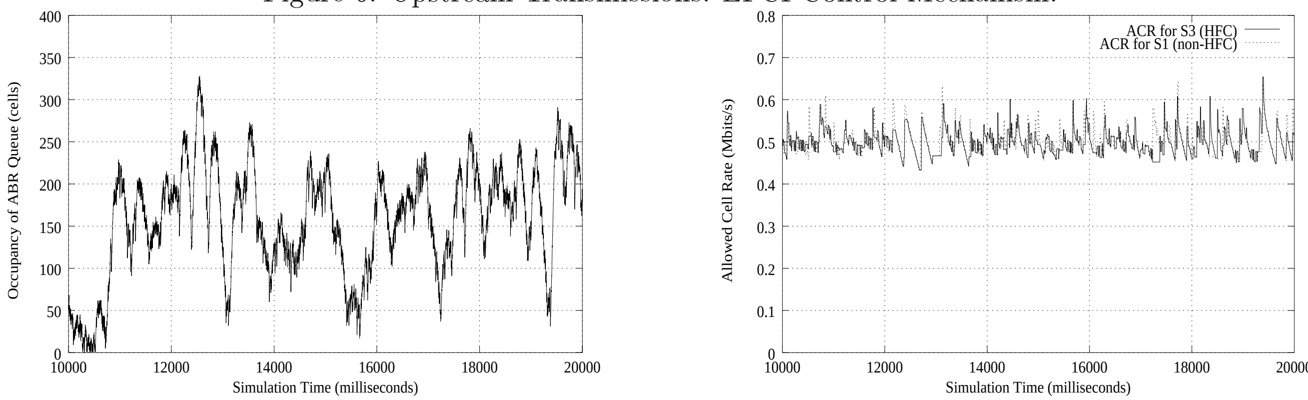

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 7: Upstream Transmissions: ER Control Mechanism.

In Figures 6 and 7 we show the results of the simulations. In Figure 6 we depict the results for the ABR feedback mechanism when EFCI switch control is used. Figure 6(a) illustrates the buffer occupancy at the congested ATM link, i.e., the output link of the right hand switch in Figure 4. Figure 6(b) depicts the allowed cell rate (ACR) of two ABR traffic sources: S1, which is located outside the HFC network, and S3, which is located inside the HFC network.

Foremost, we note in Figure 6 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 6(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 for source S3 are below those of S1. The smaller peak is due to the bandwidth limitation of the upstream channel of the HFC network. Since 12 ABR connections with PCR values of 2 Mbps are sending on the upstream channel of the HFC network, which has a bandwidth limitation of 8.192 Mbps (see Appendix A), the HFC network does not have sufficient bandwidth to support the peak rate of all sources. As a result, the

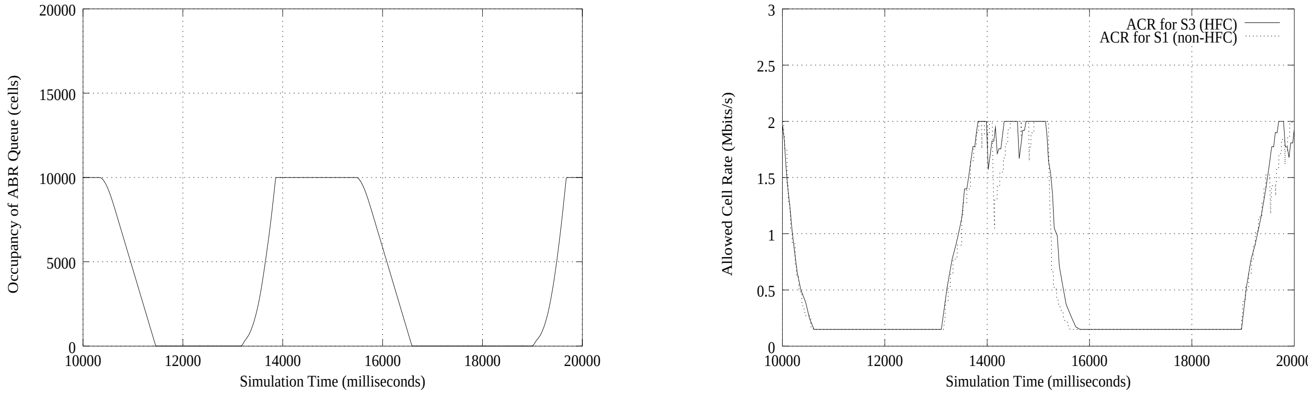

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 8: Downstream Transmissions: EFCI Control Mechanism.

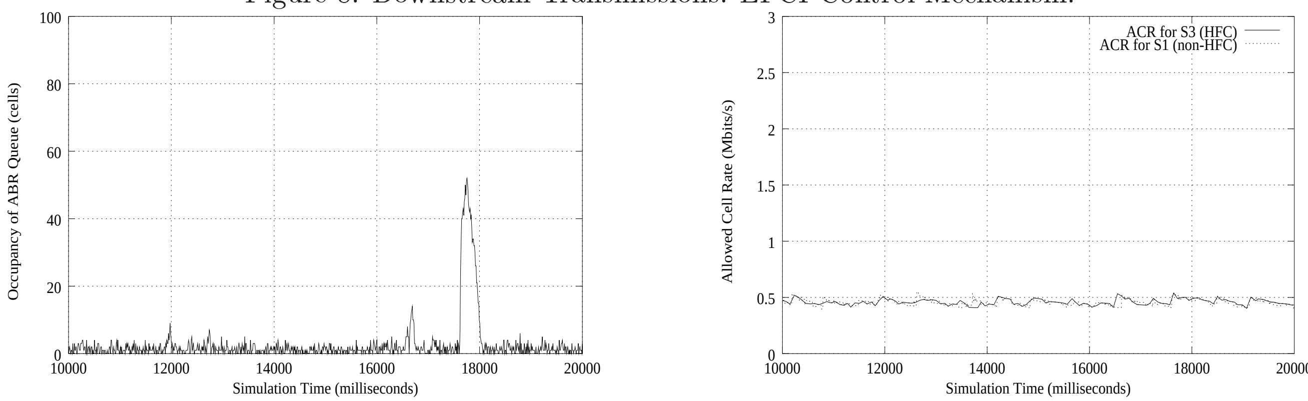

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 9: Downstream Transmissions: ER Control Mechanism.

HFC network can become the bottleneck link for the ABR sources.

In Figure 7 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 7(a) and 7(b)). In comparison to EFCI, the maximum buffer occupancy at the congested ATM link (Figures 7(a)) is an order of magnitude smaller with ER. Note from Figure 7(b) that the ACR rates oscillate around the expected rate of 0.5 Mbps.

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.

4.3 Downstream Transmissions

Next we present the outcome of the simulations for the Downstream Configuration shown in Figure 5. We will 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 Mbps, is sufficient to support the peak cell rate of all ABR connections that enter the HFC network. In the Downstream Configuration in Figure 5 we add fifty CBR connections that transmit from inside the HFC network. The CBR connections, each transmitting at 0.13 Mbps, are intended to introduce high traffic load on the upstream HFC channel.

The results of the simulations are summarized in Figures 8 and 9. For EFCI switch control, we observe in Figure 8(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 8(b). As Figure 8(a) demonstrates, the excessive delays of the backward RM cells cause a breakdown of EFCI feedback control.

Figures 9(a) and 9(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 Mbps for each ABR source.

In the next section we propose a solution to the delayed feedback with EFCI switch control when the upstream HFC channel is congested.

5 Solution Approach to the Downstream EFCI Problem

The problem with downstream transmissions of ABR traffic in an HFC network that we observed in the previous section is somewhat counterintuitive, as 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 [1]. 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 [1] makes allowances for extra ABR flow control mechanisms, such as the creation of backward RM cells at the switch.

5.1 Solution Approach

Our scheme to prevent a collapse of EFCI rate control during congestion periods in the HFC network is based on short-circuiting 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, i.e., 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 the steps of our scheme in more detail.

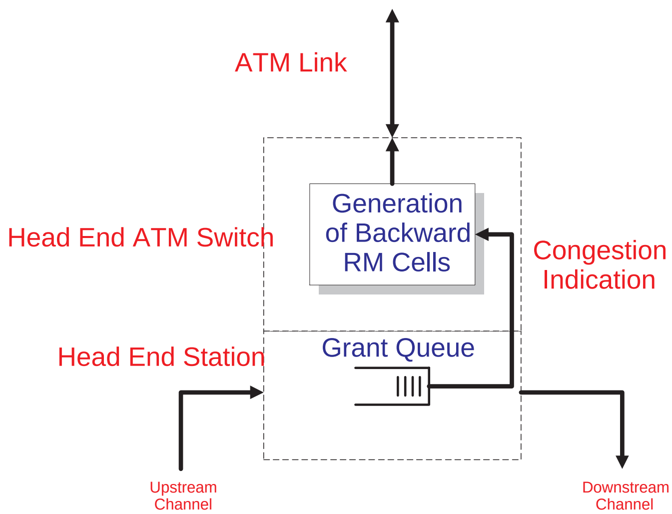

Figure 10: Modified Headend Station.

(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 measurements of the grant queue size, the headend station tracks an exponentially weighted moving average (EWMA) 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 in all our simulations. EWMA is a technique frequently applied in computer networking to smooth noisy measurements. In particular, several ABR flow control Explicit Rate (ER) algorithms employ EWMA [1].

(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 10.

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 [1]); the TM specifications 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.

5.2 Evaluation

For evaluation, we use the topology and parameters from the Downstream Configuration in Figure 5. The network is enhanced by the mechanism described above. The results of the simulation are shown in Figures 11 and 12 and in Table 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 11 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.

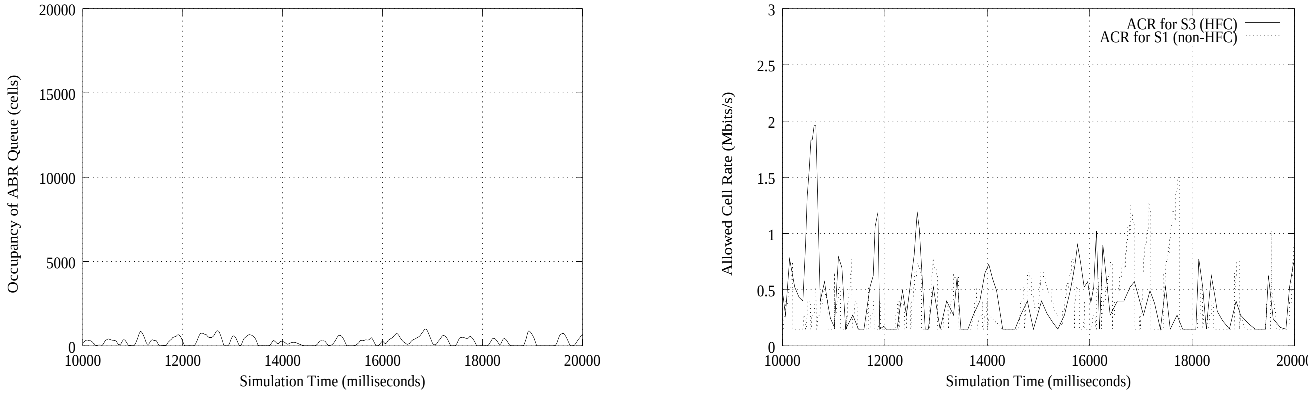

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 11: Downstream Transmissions-High Congestion (EFCI Rate Control): $TH_{low} = 35$, $TH_{high} = 40$, CBR rate = 0.13 Mbps.

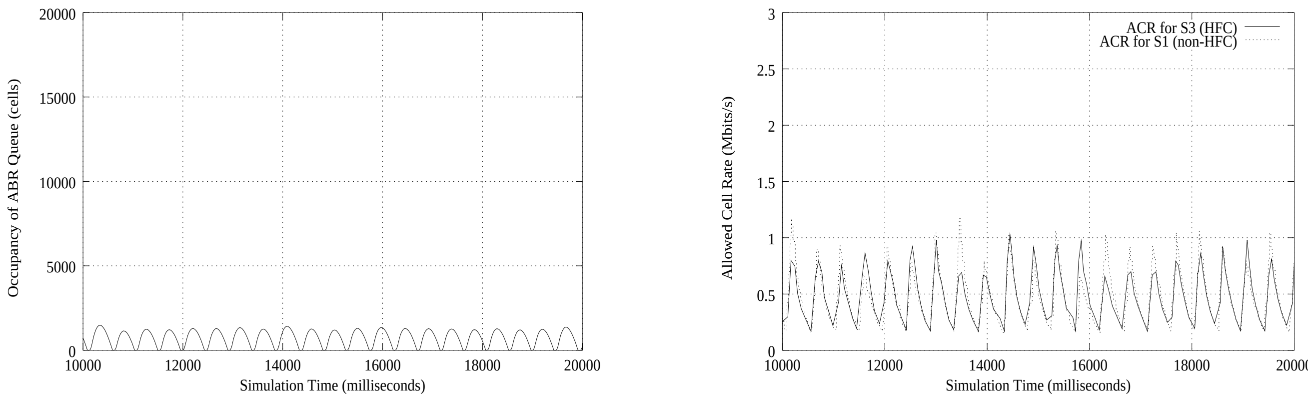

(a) Congested Link Switch Buffer Occupancy. (b) ABR Allowed Cell Rate.

Figure 12: Downstream Transmissions-Low Congestion (EFCI Rate Control): $TH_{low} = 35$, $TH_{high} = 40$, CBR rate = 0.06 Mbps.

In Figure 12 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 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 Mbps ATM link between the two switched 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 mechanism limits the rate of sources when feedback is unavailable and thus the rate is not as high as it normally would be.

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 1: ABR Cell Success Rates.

6 Concluding Remarks

The results presented in this paper have pointed to a possible problem when ABR traffic is transmitted over an HFC network that runs 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. We proposed a solution whereby the HFC headend indicates its congestion level to the closest ATM switch, which, in turn, generates additional backward RM cells.

References

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

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

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

[4] 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.

[5] 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.

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

[7] 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.

[8] N. Golmie, M. Corner, J. Liebeherr, and D. Su. Simulation Study of ABR Service over IEEE 802.14 MAC. Contribution No. IEEE 802.14/97-011, IEEE 802.14 Working Group, January 1997.

[9] 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.

[10] Y.-C. Lai and Y.-Dar Lin. "Interoperability of EFCI and ER Switches for ABR Services in ATM Networks". *IEEE Network Magazine*, 12(1):34–42, January 1998.

[11] 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.

[12] 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.

A Details on the Simulation Environment

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 Mbps
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 preamble 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 2: 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 Mbps
ABR Parameters	
Nrm (Number of RM cells)	16
Available ABR Bandwidth on Congested Link	6 Mbps
Link Cell Rate	149.76 Mbps
Allowed Cell Rate (ACR)	Dynamically adjusted
Initial Cell Rate (ICR)	0.5 Mbps
Peak Cell Rate (PCR)	2 Mbps
Minimum Cell Rate (MCR)	0.149 Mbps
Rate Increase Factor (RIF)	0.063
Rate Decrease Factor (RDF)	1/16

Table 3: 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 Mbps
Average Factor (AVF)	1/16
Mean ACR Additive Increase Rate (MAIR)	0.015 Mbps
Mean ACR Reduction Factor	0.95
Measurement Interval (N)	100 cells

Table 4: ABR Switch Parameters.