

A Priority Scheme for the IEEE 802.14 MAC Protocol for Hybrid Fiber-Coax Networks

Mark D. Corner[†]

Nada Golmie[†]

Jörg Liebeherr*

David H. Su[†]

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

* Department of Electrical Engineering
Polytechnic University
Brooklyn, NY 11201

Abstract

In order to provide Quality of Service (QoS) to users with real-time data such as voice, video and interactive services, the evolving IEEE 802.14 standard for Hybrid Fiber Coaxial (HFC) networks must include an effective priority scheme. In this paper we investigate the ability of the current IEEE 802.14 specification to provide priority service and show 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 the robustness and efficiency of the protocol, such as its ability to isolate higher priorities from lower ones and provide quick access to high priority requests.

1 Introduction

Existing community cable television systems are evolving into bidirectional Hybrid Fiber Coaxial (HFC) networks [12][13] that can support interactive broadband applications, including video-on-demand, tele-conferencing, telephony, and Internet access. The current residential network architecture uses a tree-and-branch topology, with as many as 2000 user stations attached at the leaves of the tree. Stations transmit requests and data on an 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 multiple stations transmit simultaneously. The headend transmits feedback and data to the users

on a downstream channel, which is collision-free. 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. Synchronization occurs at the physical layer, so that each station has a common time reference.

A multiple access control (MAC) protocol is used for the upstream communication between stations and the headend, in order to efficiently use the upstream channel. The MAC specifies the rules that stations must employ to request access to the channel. The procedure is as follows: First, a station sends a request for upstream bandwidth to the headend. If more than one user transmits a request at the same time, the requests collide. The headend uses a collision resolution protocol (CRP) to force the stations to transmit at different times. If the stations transmit successful requests, the headend acknowledges their transmission and reserves bandwidth in the upstream channel for the stations. The headend informs the station, using a *grant message*, when to use the channel and the user sends data without contention at the specified time.

In this paper 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. An effective priority system is needed to provide Quality of Service (QoS) in HFC applications and services such as voice, video and ATM [7]. Priority systems have been implemented in recent MAC protocols, such as DQDB [2] and Token Ring [1]. But the priority mechanisms used in those collision-free access protocols cannot be applied to the contention-based HFC environment. [10] describes a modification to

Extended Distributed Queue Random Access Protocol (XDQRAP) [14][9] that adds an extra slot to each frame to support priorities. However, this scheme 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. In [5], a priority scheme is implemented with variable probabilities in combination with the p-persistence random access protocol. However, this mechanism cannot be used in 802.14 because the CRP does not use random p-persistence for collision resolution.

To implement effective priority access we use two mechanisms. 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 bandwidth requests without interference from lower priorities. We propose a multi-priority mechanism for IEEE 802.14 to implement the latter. This scheme can be easily integrated with the standard.

The remainder of the paper 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 offers some conclusions.

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 the 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 [8] is not complete as of the time of writing (December 1997). Our description reflects the state of the draft from October 1997¹.

2.1 MAC Operation

The HFC upstream channel is divided into discrete basic time slots, 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. DS, 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, gains initial access using a so called First Transmission Rule (FTR) [3]. 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 CRP, for each slot that suffered a collision. The CRP specified by the 802.14 MAC is a blocking ternary tree algorithm [4]. 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 this 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 [11].

2.2 Collision Resolution Example

Figure 1 shows an example of the 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

¹In November 1997, after the acceptance of this paper and before the the submission of the final version, a priority scheme was proposed to the IEEE 802.14 group and voted on.

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 1 are initialized to 0, so that they can accept the transmission of new requests. In the first frame, shown in Figure 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 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 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 1(d)), the remaining slots with $RQ = 1$ are allocated and station *I* transmits its request and the system returns to an idle state.




Figure 1: Collision Resolution

3 A Multi-Priority Access Scheme for 802.14

In this section we propose an extension to the IEEE 802.14 MAC to provide priority collision resolution and access to multiple priority traffic. We first motivate 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.

3.1 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 cannot 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 1(c). It shows that nine contention slots are needed for contention resolution, 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.

3.2 Priority Protocol Description

Similar to the priority system suggested in [10], we introduce a scheme which integrates extra priority slots with the 802.14 frame format. The use of an extra slot to indicate high priority traffic was first proposed for XDQRAP [10]. 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 described next.

New Frame Format: In Figure 2 we suggest 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.




Figure 2: Priority Frame Layout

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 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 1(c), where two CS do not fit and must be allocated in the last frame (Figure 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 leftover 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 Example Priority Collision Resolution

In Figure 3 we show an example of the priority collision resolution process. Each frame corresponds to a roundtrip in the system, therefore the example represents a total time of four roundtrip delays. In this case, there are four priority levels, where three is the highest priority and zero is the lowest. We use seven stations labeled A through G , with priority level three for stations A and B , level one for station C , and level zero for stations D , E , F and G . The frame consists of seven CS, which are initialized to an RQ of zero, and two DS. Figure 3(a) illustrates the initialized priority frame as described. Recall that a negative RQ number $-N$ designates the CS as a PNA slot of priority level N . The first three CS, 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.

In the first frame, shown in Figure 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(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. Figure 3(c) shows the resolution of stations D , E and F . In the last frame, shown in Figure 3(d), all stations complete their requests and the system returns to the idle state.




Figure 3: Priority Collision Resolution

4 Performance Evaluation

We have built a simulation program to evaluate the performance of the priority system. The implementation was created as part of an HFC module for the NIST ATM simulator [6]. We used the configuration and system parameters for the HFC network shown in Table 1. 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. Note that the request delay is the time it takes a packet request to reach the headend from the time the packet arrives at the station. A

summary of the experiments is shown in Table 2. 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 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 affected.
- In Experiment 4 we evaluate the bandwidth that should be reserved for priority newcomer stations. We verify that our selection of one PNS slots 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.

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 [8]. 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 Figures 4 and 5, respectively. Figure 4 shows that the reserved PNA slots cause only a slight

Experiment	Low Priority		Medium Priority		High Priority	
	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 2: Simulation Scenarios

[†] Continuous backlog. [‡] 50 in Group 1 and 50 in Group 2

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)	3 Mbits/sec
Propagation delay	5 μ s/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	52 mini-slots
CS	Fixed 18 slots
Roundtrip	1 Frame
Maximum request size	32 data slots
Headend processing delay	1 ms

Table 1: Simulation Parameters

increase in request delay while Figure 5 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). 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.

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 6 we plot the request delay versus load for




Figure 4: PNA Overhead-Average, Request Delay




Figure 5: PNA Overhead-Coefficient of Request Delay Variation

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 and 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 which reflects the robust operation of the priority scheme.

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% of the capacity. The third group consists of a large number of high priority stations (100) which




Figure 6: Varying Medium Priority Load

transmit at loads that is varied from 10% to 45%. Figure 7 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 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.




Figure 7: Varying High Priority Load

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 8 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 stations are 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 attributed to collision resolution. Since one PNA for priority traffic newcomers may not be sufficient at low loads, in Figure 9 the problem is easily fixed by allocating 5 PNA 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.




Figure 8: Low Load Performance-PNA for Medium Priority




Figure 9: Low Load Performance-5 PNAs for Medium Priority

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 roundtrips. 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 10 shows throughput measurements, taken for one roundtrip (frame) intervals. A comparison of Figures 10(a) and 10(b) shows that the medium priority stations can preempt the low priority traffic within one or two roundtrip delays. Figure 10(c) shows that when the second 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 10(d).

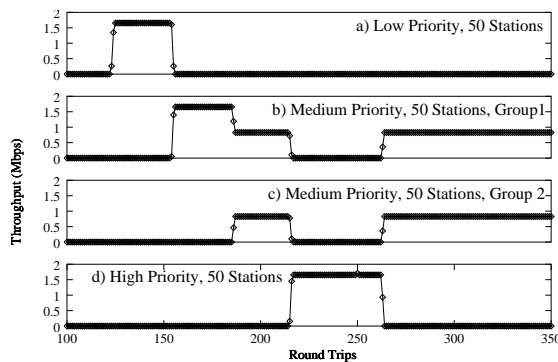


Figure 10: Transient Throughput

5 Concluding Remarks

In this paper 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 with the current specification. 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. 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.

References

- [1] IEEE Std 802.5-1985, Token Ring Access Method, 1985.
- [2] IEEE Std 802.6-1990, IEEE Standards for Local and Metropolitan Area Networks: Distributed Queue Dual Bus (DQDB) of a Metropolitan Area Network, July 1991.
- [3] C. Bisdikian. A Review of Random Access Algorithms. Technical Report RC 20348, IBM Re-

- search Division, T.J. Watson Research Center, January 1996.
- [4] 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.
- [5] 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.
- [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] N. Golmie, M. Corner, J. Liebeherr, and D. Su. Improving the Effectiveness of ATM Traffic Control Over Hybrid Fiber-Coax Networks. Proceedings of IEEE Globecom '97, November 1997.
- [8] IEEE 802.14 Working Group. Media Access and Control. IEEE Std 802.14, Draft 2 Revision 2, IEEE 802.14 Working Group, October 1997.
- [9] 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.
- [10] 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.
- [11] B. McNeil. Implementation Overview of Tree-Based Algorithm. Contribution No. IEEE 802.14/96-246, IEEE 802.14 Working Group, November 1996.
- [12] S. Ramanathan and R. Gusella. Toward Management Systems for Emerging Hybrid Fiber-Coax Access Networks. *IEEE Network*, 9(5):58–68, September/October 1995.
- [13] M. Vecchi. Broadband Networks and Services: Architecture and Control. *IEEE Communications Magazine*, 33(8):24–32, August 1995.
- [14] 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.