A Protocol for Efficient Transfer of Data over Fiber/Cable Systems

Dolors Sala
School of Electrical
and Computer Engineering

John O. Limb
College of Computing

Georgia Institute of Technology
Atlanta, GA 30332-0280
E-mail: (dolors,limb)@cc.gatech.edu

Abstract

A revolution is occurring in the scope and range of information, communication and education services that will be made available to schools, libraries, town-halls, clinics and, most importantly, residences. These services will be provided initially, primarily over hybrid fiber-cable systems, either by telephone companies or cable companies. The old cable plant is being upgraded and used in totally new ways.

The topology and physical characteristics of the upstream channel present new challenges for efficient channel access. We present a media access protocol that efficiently transfers data on this channel. A primary goal in the design was to keep the portion of the protocol resident in the station as simple as possible. Thus we use centralized control located in the cable head-end and minimize intelligence in the station. We refer to this protocol as Centralized Priority Reservation or CPR. A station wishing to transmit sends a request to the head-end using a contention channel. The head-end acknowledges the request and then schedules the request, informing the station by means of a grant message when to transmit.

The protocol performs well under heavy load. Performance is affected little by the number of stations, the speed of the system and the physical length of the system.

1 Introduction

1.1 Background

The path to the home will become an important entry and exit point to the National Information Infrastructure. This path should support communications services for telecommuting and access to information sources in addition to many other services. The existing twisted pair has been used to meet this need to date by means of data modems, ISDN and Asynchronous Digital Subscriber Loop [1]. An alternative physical medium to twisted pair is a combination of optical fiber and coaxial cable referred to as Hybrid Fiber Coax (HFC). This medium has caught the attention of both the cable industry and the telephone industry so that today many trials are being conducted using HFC. While HFC is just an evolution of classical cable technology a lot of questions exist as to how best to exploit it for the purpose of supplying information and data services to the home.

This paper addresses the challenge of how best to provide data communication over the HFC plant, irrespective of what the data may be used for. The biggest challenge is to get the data from the home to the Head-end (H/E) referred to as the upstream or reverse channel (from where it might be routed to the office, information supplier or perhaps another home). There are fewer hurdles associated with transmitting the data in the opposite direction from the H/E to the home, referred to as the downstream or forward direction. This asymmetry results from the nature of the HFC plant as described in the next section.

1.2 Hybrid Fiber Coaxial Channel

Fig. 1a represents a piece of a HFC system. Clusters of homes are served by a fiber link that runs from the H/E to an optical fiber node. At the node the optical signal is transformed to an electrical signal and makes the connection to the home via a relatively short section of coaxial cable, perhaps as short as 2km in total length. The number of houses served may vary from 50 to 2000, depending on a number of factors.

HFC is an inherently shared medium. As the signal passes between the optical node and the station it passes through a small number of amplifiers that boost the electrical signal, using separate amplifiers for the upstream and downstream channels (represented by a pair of arrowheads in Fig. 1a). Additionally, directional taps are used to distributed the signal to groups of stations. As a result, while the signal may flow easily from the H/E to the station and from the station to the H/E the signal is *very* highly attenuated between stations due to directional taps and amplifiers which pass the signal in only one direction. This makes it all but impossible to communicate directly between stations. It is this characteristic that distinguishes the HFC channel from many others and rules out the efficient use of a number of common protocols used on shared access media such as CSMA/CD or DQDB. The best analogy to the HFC channel is probably a number of ground stations communicating via a satellite. All signals must pass through the

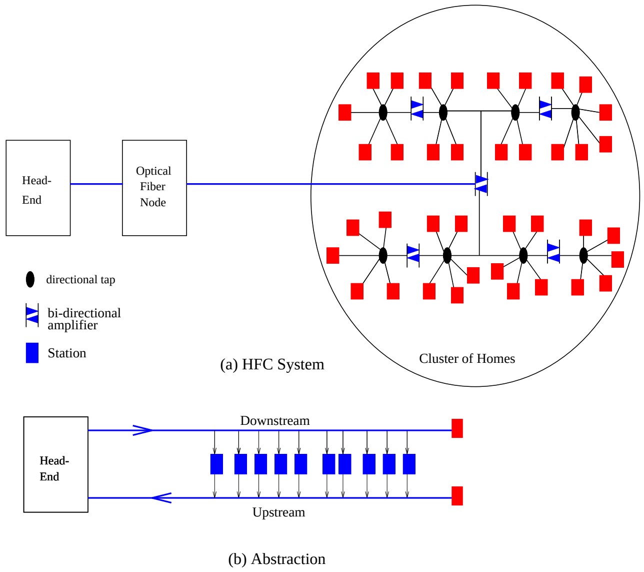

Figure 1: Hybrid Fiber-Coaxial System

satellite; ground stations cannot communicate directly with one another.

For our purposes the abstract representation of an HFC system shown in Fig. 1b captures the essence of the system. Our aim is to devise an efficient medium access (MAC) protocol. We will assume that a physical signaling system exists that provides digital connectivity between stations and H/E. In practice the analog band from station to H/E is relatively narrow (5 MHz to 40 MHz) and of poor quality. On the downstream, channel bandwidth is more abundant (54 MHz to 550 or 750 MHz) and the quality of the channel is very high (otherwise it would not be suitable for TV transmission). Note also that when the H/E changes from "listening" to one station to listening to another it will take some time to "tune-in" to the new station; it may have to adjust to the amplitude of the new signal, it must then achieve bit timing and finally framing before it can interpret the information. This is not true in the downstream direction since each station can stay locked-in to the H/E data stream.

In addition to a physical layer we assume that the H/E has the means to determine the time it takes for a message to propagate from a station to the H/E and back. To achieve an accurate measure this function will most likely occur at the hardware level and lie between the physical layer and the MAC layer.

1.3 Goals of Centralized Priority Reservation Protocol Design

The most important goal of the design was that it be inexpensive to implement in the station. The cost of the H/E implementation is less important since the cost is amortized over all the stations in the cluster. It is important that the protocol exploit the medium efficiently over a range of transmission speeds (say 1 - 100Mb/s) and over relatively long distances (up to 35 km, [2]). Latency should be relatively low so that voice conversations and interactive games can be supported (less than a few 10s of ms). It is also important that the protocol support fixed-rate streams of information and multiple priorities for data cells, although these aspects will not be addressed in this paper (see [3]).

1.4 Previous Work

As previously mentioned, the HFC channel is similar to a satellite channel and, hence, much of that literature is relevant. In particular we have drawn inspiration from early work by Roberts [4] and on the PODA system [5, 6]. There is little published work directed specifically to the HFC channel. The exceptions are the work on DQRAP and XDQRAP [7, 8], and the paper by Counterman [9]. DQRAP uses contention minislots to make distributed reservations in a DQDB-like manner. Contentions occurring in minislots are resolved through a gating mechanism where new arrivals are blocked until all collisions are resolved. Counterman uses centralized control to send permits to each station. He makes a clear distinction between access control on one hand and bandwidth allocation mechanisms on the other.

By comparison, our proposal, Centralized Priority Reservation, (CPR) uses contention minislots with simple p-persistence to send requests to the H/E [10]. The H/E then makes the reservations according to a centralized policy and informs the stations by means of grant minislots in the downstream channel.

2 Description of CPR Protocol

2.1 Overview

A station transmits its information in fixed length data slots, but controls the flow of information by means of a stream of minislots placed between the data slots. In the upstream direction these minislots carry requests from the station to the H/E. In the downstream direction the minislots are used by the H/E to acknowledge requests from the station and to grant access to the data slots (see Fig. 2). A station, upon receiving a request to transmit information writes in the first arriving minislot in order to inform the H/E station that it has information to transmit. The H/E, upon receiving the request, immediately echoes back an acknowledgment, letting the station know that the request has been received and that no collision has occurred. If an acknowledgment is not received the station assumes that a collision has occurred and competes for a new minislot in which to send the request.

The H/E, upon receiving the request and forwarding the acknowledgment, schedules the request. In scheduling the request the H/E will use information about the overall state of the system, the type of the request, and specific information about the requesting station. The H/E then forwards a message to the sending station granting a number of slots in which to send the message. The granted number of slots may be less than the number requested, in which case the sending station waits for an additional grant of slots to complete transmission of the message. The H/E

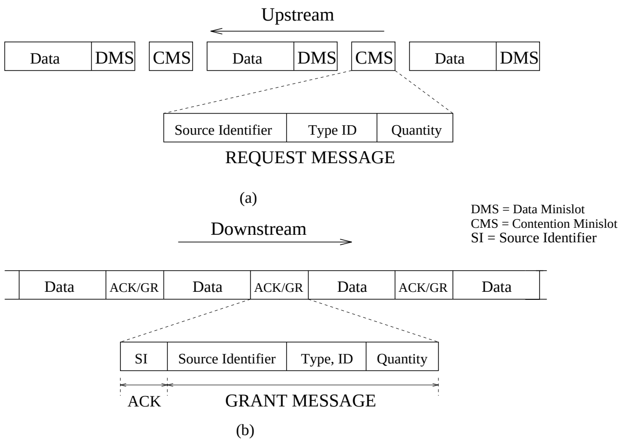

Figure 2: Slot Structure of the Up and Down Channels

times the transmission of the grant so that it arrives at the sending station in the slot prior to the allocated upstream slot.

2.2 Description of basic CPR Protocol

2.2.1 Slot Structure

Fig. 2 shows the slot structure in both the upstream and downstream directions. As we are only interested in the MAC layer protocol the data structure within the data slots will adhere to whatever higher level protocols are deployed. For example, the data structure could be ATM cells or 802.2 frames. The request message consists of three fields. The first is the *source identifier*. The second is the *type* field which indicates the type of communication required. In this paper we are concerned primarily with data traffic (but see discussion in Section 4 and [3]). The third field is a *quantity* field and specifies the number of slots that the station requires for a message or packet. To prevent a single station from hogging the channel a limit may be placed on the maximum number of slots that are asked for in a single request. Alternatively such control may be exercised by the H/E.

A minislot is of two types, a contention minislot (CMS) that is positioned between data slots and a data minislot (DMS) that is appended to each data slot. The field structure of both types of minislots is the same. Any station may attempt to write in a CMS on a contention basis. Only the same station that writes in the data slot may write in the associated DMS.

The downstream channel has a combination acknowledgment (ACK) and grant (GR) minislot sandwiched between data slots. The ACK message requires just one field containing the *station identifier* of the requesting station. The grant message, on the other hand, requires three fields, *source identifier*, *type* (since a station could request to send data of a different type while a data request is outstanding), and *quantity* (indicating the number of consecutive cells granted to the station).

2.2.2 Station Protocol

When a message arrives at a station, it is added to the request queue along with its priority (*type* field) and the number of slots required to transmit the complete message (*quantity* field). While not strictly necessary, for simplicity, we will assume that only one message request of a given type may be outstanding at any time. Upon receiving a message to transmit, a station writes a request in the next minislot arriving at the station, filling in the fields shown in Fig. 2a. A station is aware of the time, in slots, that the acknowledgment (ACK) takes to arrive back at the station. This time is a constant period that includes the time it takes the request to arrive at the H/E, the H/E processing time required to enqueued the request and to generate the ACK and finally the time necessary for the ACK to reach the station. After this period of time, the station checks if the incoming ACK in the downstream slot refers to its own address. If this is the case the request has been successfully transmitted and acknowledged by the H/E. At this point the station continuously monitors the downstream grant channel waiting for grants. When a grant is addressed to the station, the station immediately transmits as many cells as granted. The number of slots used is subtracted from the requested slots in order to keep track of when the requested transmission has been completed. The station waits for more grants until the number of pending requested slots is equal to zero. In other words, the transmission is complete when the total number of granted slots equals the initial number of requested slots. The station can use the DMS of the last granted slot to start the reservation process of the next enqueued transmission (request queue). Since there can be no collisions in the DMS it is not necessary to acknowledge the request (although it may still be done for reliability purposes). If at the completion of the current transmission the request queue is empty the station loses the opportunity to use the DMS so the new incoming request will need to go through the contention channel (CMS) with the possibility of suffering collisions.

If the request is not acknowledged by the H/E in the expected ACK slot then the station repeats the request in the subsequent contention minislots with probability p until the request is transmitted (p-persistence [11]). After transmitting the request the station again waits for an ACK from the H/E. Note that if the H/E is able to immediately grant the request then the grant and the ACK will arrive in the same ACK/GR slot (assuming of course that the H/E is able to schedule the request in time).

Fig. 3 captures the essence of the station protocol as the station moves through four different states.

2.2.3 Head-End Protocol

Upon receiving a request from the station the H/E immediately writes the address of the requesting station in the ACK slot of the next departing downstream

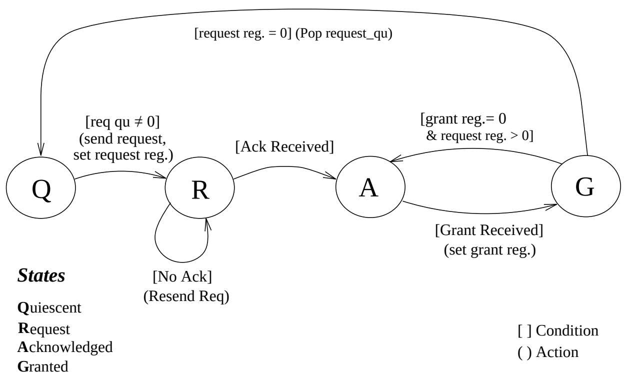

Figure 3: State Transition Diagram of a Station

slot. The H/E then schedules the request. It does this by examining a buffer that contains an image of future slots in which granted slots have been marked. It is not always possible to use the first available free slot because the grant minislot in which the H/E station would need to write may already be taken. Remember that we need to time the grant message to arrive at a station just prior to the arrival of the upstream slot in which the station is scheduled to write. Since the distance of stations from the H/E varies it is possible for another station, writing in a different upstream slot, to have already used the downstream grant minislot.

Having found the first usable free slot the H/E then reserves subsequent slots until the number requested has been reached or a slot that has already been granted to a previous request is encountered. If the later case, the scheduled slots are recorded together with other grant message information and transmitted to the sending station in the appropriate downstream minislot. The remaining slots are scheduled in a second (or more) round. Thus a single request might result in one or more grants being transmitted to a station.

In this manner the H/E can schedule upstream data slots very densely; the only limitation, which occurs very occasionally, is that an upstream data slot may not be used because the corresponding downstream grant minislot may have already been taken. This limitation could be even further minimized by providing more than one grant minislot per data slot however since this effect occurs very rarely it does not warrant the added complexity (see Section 4).

3 Performance

3.1 Reference Parameters

The performance of CPR is described using a reference set of parameters. After we explore the reference condition we vary each of the important parameters, in turn, in order to test the performance of the protocol over a range of conditions. The simulations to characterize the performance were written in C and were typically run for 550,000 slots with statistics being gathered after the first 50,000 slots. The input

Table I: Values of the input parameters for the reference case.

Parameter	Reference
Number of Stations	50
Length of Fiber Trunk	2 slots[1]
Length of Coaxial Distribution	1 slot [1]
Head-End Processing Time	1.4 slots[1]
Number of Collision Minislots, CMS	1
Number of Data Minislots, DMS	1
P-persistent probability, p	0.2
Burst Poisson Traffic Model, (i/j)	(9/3)

parameters of the simulations and the set of reference values are shown in Table I.

We have assumed a simple data traffic model, the burst Poisson model. The length of the burst may be adjusted to approximate LAN traffic, an important initial application of data over HFC. Our burst Poisson traffic model assumes that messages arrive with a burst length of either 1 cell or i cells and the ratio between messages of length 1 and length i is j. Thus (9/3) denotes Poisson burst arrivals with sizes of 1 cell and 9 cells and a ratio of 3 to 1 between messages of size 1 and 9. In this instance the average message size is of length 3 slots.

The measurements that were collected on the performance of the system were queuing delay of cells (not messages), the number of cells transmitted by a station and the probability that a collision occurs in a collision minislot. (A data minislot cannot have a collision). The delay measurement could be used to calculate the average delay at a station or the average delay averaged over all stations.

All results are normalized to the cell transmission time. A utilization or load of 100% would correspond to every upstream slot being utilized. The total channel utilization would be less than 100% since the overhead of the minislots and the necessary guard bands between slots are not taken into account. Since these durations depend strongly on the transmission rate and the implementation details we feel it is more broadly useful to present the utilizations in this way. The results can then be used to calculate the expected throughput in bits/s for a specific implementation.

3.2 Performance of Reference Case

The queuing delay averaged over all stations is shown in Fig. 3.2. The results of a number of simulations (each of 500,000 slots) at each value of nominal load is shown as a small circle. It can be seen that the statistical fluctuation is very small, particularly at small loads. The delay expected at very low loads

[1]Distances are measured in units that correspond to the distance the signal on the medium travels during the time it takes to transmit one cell, i.e. a slot duration. While this normalized unit of distance may appear unusual it simplifies the comparison of systems operating at different speeds over different distances.

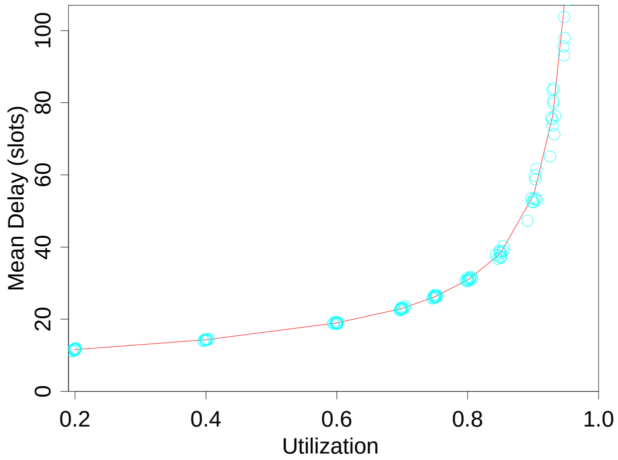

Figure 4: Mean queuing delay as a function of Utilization for the Reference Configuration: Stations = 50, Fiber = 2, Coax = 1, p = 0.2, Traffic (9/3), 1 collision minislot and with data minislot. Each nominal value of utilization has been repeated 10 times.

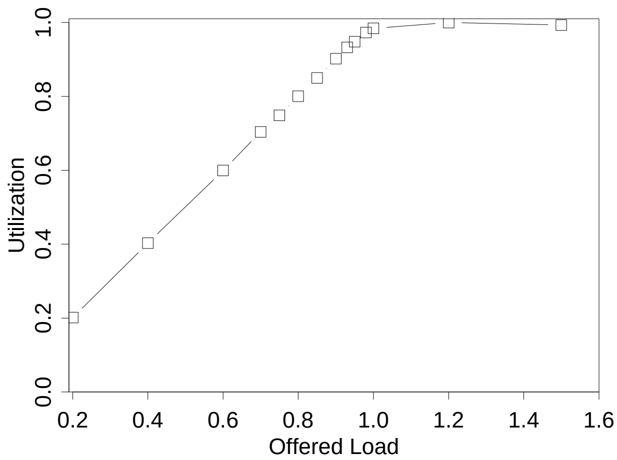

Figure 5: Throughput versus Load in the data channel. The length of the simulation is 550,000 slots up to 0.95 load, for higher loads, the simulation length has been reduced to 110,000 (gathering statistics after 10,000 slots) due to queue build up.

would be the average round-trip delay (time required to make the reservation) plus the average time a cell waits to be transmitted (0 for a message of 1 cell duration and 4 for a message of 9 cells duration). Since stations are randomly distributed over the length of the coax cable the average round-trip delay is 7 slots, the total delay is 10 slots. It can be seen that the delay at low loads asymptotically approaches to this value. The delay increases with load to a delay of approximately 40 cells at a 90% utilization. Much of this delay is due to the inherent burstiness of the traffic (see Section 4.2).

The throughput increases linearly with load up to a maximum throughput of almost 1.0 in overload (Fig. 5).

The probability of a minislot collision as a function of load is shown in Fig. 6. Since the average length of a message is equal to the duration of 3 slots, and since the reservation channel is a "classical" Slotted Aloha channel we should expect a throughput that would support a data load in excess of 100%. In fact collisions are even less frequent than such an analysis might suggest because at high loads queues start to build up and this enables the data minislot to be used for making reservations. At very high loads this leads to a decrease in the number of minislot collisions. With such a small probability of minislot collision we can assume that the reservation channel is not a bottleneck for this traffic. For shorter average message lengths this will not be the case (see Section 3.3).

It is important that stations continue to obtain fair access to the channel during overload since, in practice, virtually all data networks experience loads in excess of their capacity. Fig. 7 shows the number of cells sent by each station when the applied load is 1.5. Stations are numbered according to their distance from the far end of the coaxial distribution plant. There appears to be no positional bias and the access granted each station is very similar.

3.3 Variation of Parameters

In Fig. 8 we see that performance of the CPR changes little as the number of stations changes. For a large number of stations, 100 or greater, it is possible to get excessive congestion in the reservation channel if p is too small (curve c). This is because there is a small probability of a large number of stations being simultaneously busy and overloading the reservation channel, a situation from which it takes a long time to recover. This does not occur with a smaller value of p of 0.05 (curve d). For a small number of stations the total throughput will decrease for long systems (see Section 4.2). In Fig. 9 the mean delay increases as the length of the line increases. The total round-trip-delay (including the H/E processing delay of 1.4 slots) is 8, 22 and 16 for curves a, b and c respectively.

Changes in the traffic model are shown in Fig. 10. In comparing burst traffic (9/3) with (16/4) we see, as expected that the more bursty traffic experiences greater delay, and this difference increases with load. At low loads the delay is due just to the increase in time that it takes to transmit the longer message. One would expect that delay for the (1/1) traffic (single cells with Poisson arrival) would be even less than for the burst traffic. In fact it is greater, and the reason, again, is contention in the reservation channel. This can be seen in Fig. 11. For the (1/1) traffic the value of p was 0.05.

In most applications an average message length less than 2 would be unusual. A preponderance of voice cells could produce this situation, but a better way to handle voice traffic would be to allocate a stream

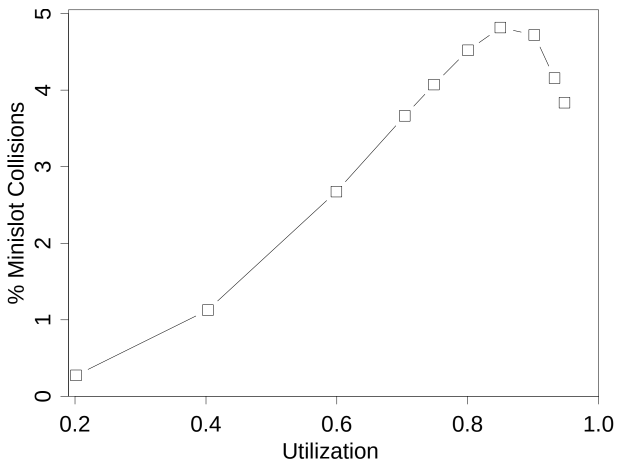

Figure 6: Percentage of collisions in the reservation channel as a function of Utilization. At high loads the usage of the data minislot increases, decreasing the collisions in the collision minislots.

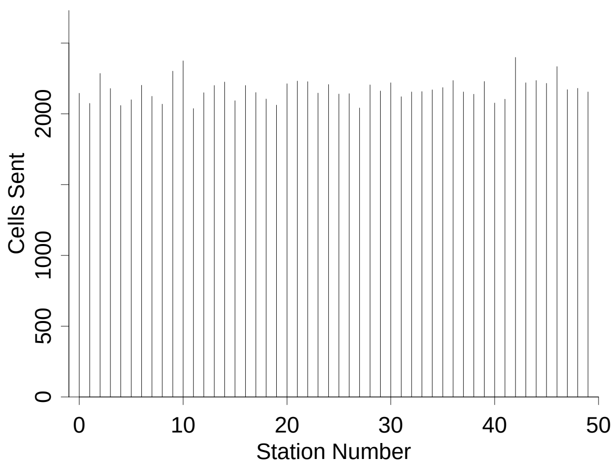

Figure 7: Cells sent by each individual station in overloaded condition (load = 1.5). There is no positional advantage in accessing the medium.

to each voice circuit (not described in detail in this paper, but see Section 4.3 and [3]). The presence of streams reduces the load on collision minislots since they are set up with a single request. If a network did have a traffic load with an average message length approaching 1 then one should consider using two collision minislots per slot so as to keep minislot collisions low and cell delay low.

Delay versus utilization curves are shown in Fig. 12 for (1/1) traffic comparing 1 collision minislot and one data minislot with 2 collision minislots, with and without a data minislot. First we observe that 1 collision minislot with 1 data minislot performs better than 2 collision minislots at high load where the data minislot takes over. At low loads 2 collision minislots are clearly better as would be expected. Low delay is obtained with 2 collision minislots and a data minislot. It appears that contention in the reservation channel is not a bottleneck at any load for this condition.

4 Discussion

4.1 Centralized vs Distributed

The scheduling adopted here could have been implemented in a distributed manner to achieve identical performance, so performance was not an issue in deciding whether to centralize or not; it was the overall cost of system implementation. As mentioned in Section 1.3, our primary goal was to keep station costs to a minimum. However, additional system advantages derive from a centralized implementation. In particular, management administration and control are simplified through retaining all intelligence in the H/E. System parameters are easily changed at the H/E in response to changing load or changing mixture of traffic. Fairness (or favoritism) can be exercised simply through the H/E scheduling process. Policy is executed in the H/E and therefore can be changed more easily than if new software has to be downloaded to each station.

4.2 Performance

There are three factors that will limit the throughput of the CPR protocol. The first is the maximum throughput limitation of an individual station; it must wait one round-trip delay between each message that is sent. If we assume a maximum length system of 70km then the round-trip delay is approximately 400μs. If we assume a station sends messages that are three ATM cells long, on average, the duration of the message is 848μs at 1.5Mb/s and so a station could seize about 70% of the total capacity. If on the other hand the transmission speed is 20Mb/s, the message is only 64 μs long and the station seizes only 14% of the capacity (still a rather large 2.8Mb/s for an individual station). Note that if the station were making a large file transfer the data could proceed at close to the full speed of the channel, in the absence of competing traffic and assuming that the H/E policy permitted it.

The second factor that can limit throughput is contention for the minislots. As previously mentioned, our aim has been to design the system so that minislot collisions are very unlikely, less than a few percent. By attaching a minislot to the data slot we keep collision probability very small, so that for realistic traffic it is unlikely that more than one contention minislot would be needed.

The protocol as evaluated herein makes one request to the H/E for each message. It would be possible to group all messages waiting to be transmitted together as one "super message" and send just one request to the head end. This would further reduce the load on the minislots. However, we suspect that real traffic will have a large enough average burst length that this approach would not warrant the additional complexity.

The third factor that could limit performance is

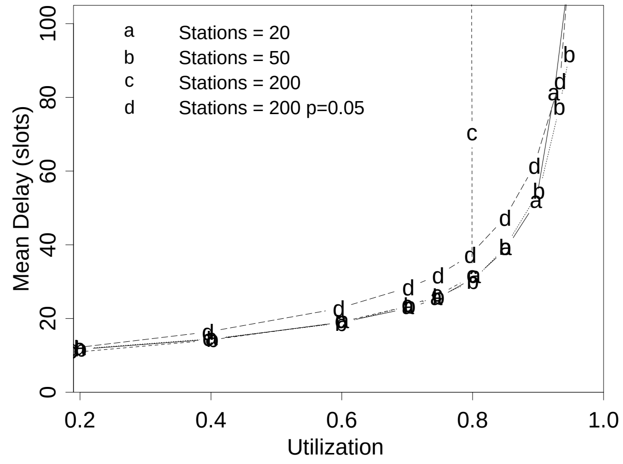

Figure 8: Mean queuing delay versus Utilization for different numbers of stations.

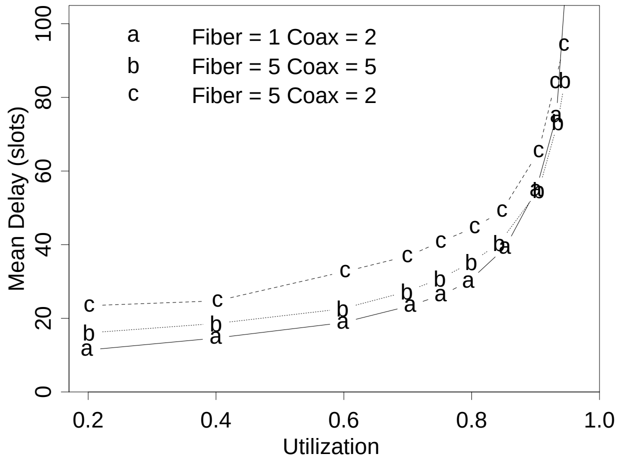

Figure 9: Mean delay versus Utilization for different cable lengths.

that the optimum grant minislot may be busy, requiring that the grant message be delayed. The probably of this happening will increase with the physical spread of the variation in distance between the closest and furthest stations from the fiber node. In our most extreme cases we found this factor to have a negligible effect.

We have defined throughput as the fraction of dataslots that are filled. However, there is overhead in the minislots and guard bands between slots that must be accounted for. In an implementation the overhead will depend on the details of exactly how (and how well) the physical layer is implemented as well as system design specifications. Take, for example the preamble required in front of each slot/minicell. The physical receiver at the H/E must adapt to the change in amplitude from one station to another. This variation will depend on the system design and how well the system is maintained. The preamble on each slot is used to obtain timing and framing. The speed with which this is done will depend on the quality of the phase-locked loop in the receiver, and the synchronization scheme that is used. A practical system could easily have from 5 to 30% overhead. Short slots, noisy signals and high transmission rates will all result in increased overhead for a given level of performance.

4.3 Priority and Stream Traffic

In this paper we described the performance of CPR for single-priority data traffic. The framework permits the addition of priorities and stream type traffic. The format of the request slot is quite general and data can be marked with a priority or a request can be made for a constant bit-rate (CBR) stream. In the case of priorities the scheduler at the H/E maintains a separate queue for each priority and serves the traffic according to its policy. In the case of a CBR stream the H/E overlays a frame structure on top of the slots and reserves periodic slots for the requested streams. A station with only data to transmit need not be aware of the frame structure. The performance of CPR for prioritized traffic and stream traffic is described in [3].

5 Conclusions

We have presented a protocol called Centralized Priority Reservation (CPR) for transmission of data over a hybrid Fiber/Cable TV System. The system uses contention slots to forward requests via the upstream channel to the head-end. Here the requests are acknowledged and scheduled for transmission. A "grant" message from the head-end then informs the station when to transmit. Simulations using Poisson and batch traffic show that the protocol performs well under heavy load. By attaching a request slot to each upstream data slot we reduce the number of collisions on the minislots to the point that only one contention minislot per data slot is necessary for the typical range of loads. Centralized control permits a low-cost, flexible system and simplifies the handling of multiple levels of priority and constant bit-rate streams.

References

[1] D. L. Waring, "The Asymmetrical Digital Subscriber Line (ADSL): A New Transport Technology for Delivering Wideband Capabilities to the Residence", *IEEE Proc. Globecom'91*, pp. 1979-1986, December 2-5, 1991.

[2] IEEE 802.14, "Cable TV MAC/PHY Protocol Working Group Functional Requirements", October 19, 1994.

[3] J. O. Limb and D. Sala, "An Access Protocol to Support Multimedia Traffic over Hybrid Fiber/Coax Systems", *Proc. 2nd International Workshop on Community Networking*, pp. 35-40, June 20-22, 1995.

[4] L. G. Roberts, "Dynamic Allocation of Satellite Capacity Through Packet Reservation", *AFIPS*

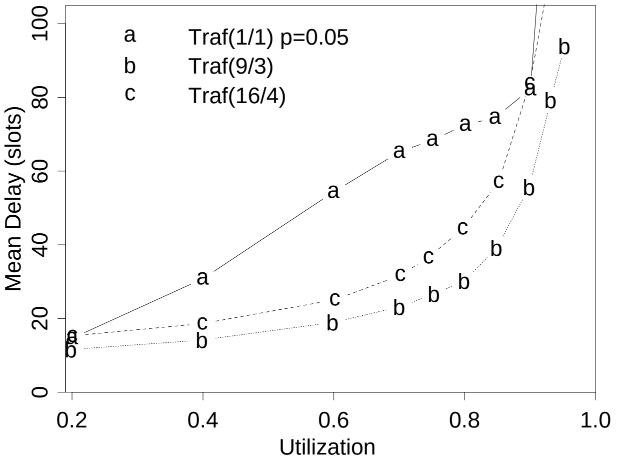

Figure 10: Delay versus Utilization for different traffic models using the Reference Configuration except for the p-persistence probability for the Traffic (1/1) where the p has been adjusted to 0.05.

Conference Proc., National Computer Conference, 42, pp. 711-716, June 1973.

[5] R. Binder, "A Dynamic Packet Switching System for Satellite Broadcast Channels" *Proc. ICC'75*, pp. 41-1 41-5, San Francisco, CA, June 1975.

[6] I. M. Jacobs, R. Binder and E. V. Hoversten, "General Purpose Packet Satellite Networks", *Proc. IEEE*, Vol. 66, No 11, pp. 1448-1467, November 1978.

[7] W. Xu and G. Campbell, "A Distributed Queuing Random Access Protocol for a Broadcast Channel" *ACM SIGCOMM'93*, pp. 270-278.

[8] C-T. Wu and G. Campbell, "Extended DQRAP (XDQRAP) A Cable TV Protocol Functioning as a Distributed Switch", *Proc. 1994 1st International Workshop on Community Networking*, pp. 191-198, July 13-14, 1994.

[9] R. Counterman, "Shared-Channel ATM-based Access for Hybrid Fiber and Coaxial Architectures", *Proc. 1994 1st International Workshop on Community Networking*, pp. 183-189, July 13-14, 1994.

[10] D. Sala, A. Mukherjee and J. O. Limb, "A Protocol for Efficient Transfer of Data over Fiber/Cable Systems", *Georgia Tech Technical report GIT-CC-95/18*, June 1995.

[11] L. Kleinrock and F. A. Tobagi, "Packet Switching in Radio Channels: Part I - Carrier Sense Multiple Access Models and their Throughput Delay Characteristics", *IEEE Trans. Commun.* , Vol. COM-25, pp. 1400-1416, December 1975.

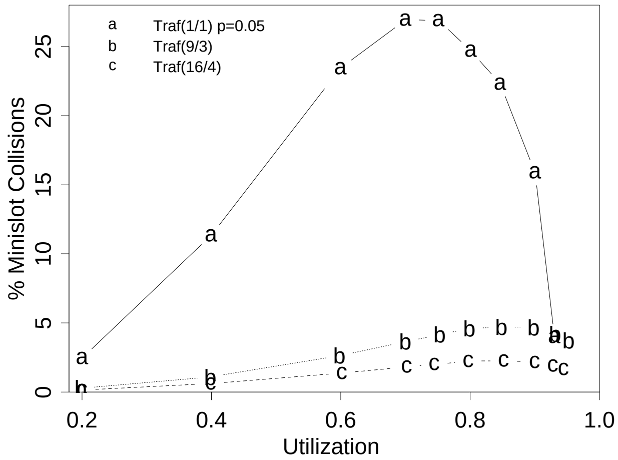

Figure 11: Percentage of minislots collisions for different traffic models. Data minislot usage is much greater for traffic models with a small burst length, reducing significantly the contention in the reservation channel at higher loads.

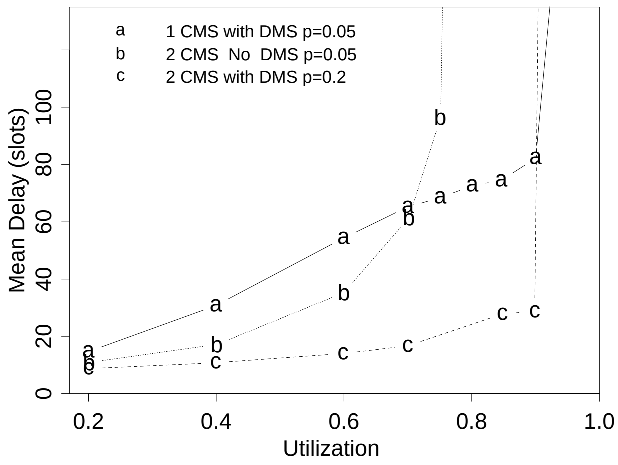

Figure 12: Mean queuing delay versus Utilization for three different reservation channels: (a) 1 collision minislot (CMS) with a data minislot (DMS), (b) 2 CMS with no DMS, (c) 2 CMS with a DMS. Even with 2 collision minislots the data minislot plays an important role at high loads.