M.I.T. Laboratory for Computer Science

Request for Comments No. 265
October 10, 1984

A Proposed High Speed Packet Switching Network Built on a CATV System

by David C. Feldmeier

1. Introduction

Computers are increasingly common in the home and small businesses; many applications require access to external computer systems and services. Consequently, the communication capacity of the home or business must also increase. Larger organizations can afford to install local communications with LANs and then connect these LANs to the rest of the world via telephone, microwave, optic and satellite links. For a small business or home, with only a few computers, high speed metropolitan area communications are out of reach. The user is forced to use a phone line at 4800 baud or less. Compounding this problem is the new generation of personal computers arriving on market. These machines have faster processors, more use of graphics and bitmap displays, and larger memory sizes. These increases in data manipulation will increase the need for high speed data communication.

The telephone system is used for communication because it is already installed. Fiber optic cable may eventually replace copper twisted pair phone lines, bringing more communication bandwidth to the home or office; but fiber is currently expensive and widespread use is probably five to ten years away for large business customers and twenty years away for the home. But already a high bandwidth communication link runs to many homes in the form of a coaxial cable for the local CATV system. Unused bandwidth on a CATV system could bring high speed (1-10 MBPS), low cost (less than $1000 per interface), packet switched communication facilities into the home or office.

This paper proposes a preliminary design for a packet switched communication system installable on a CATV network to provide digital data communication over a metropolitan area.

2. System Design

The network design takes advantage of existing commercial CATV systems, so understanding CATV systems helps in understanding the network design.

2.1. Cable Television

A cable TV network is a coaxial cable transmission system that is frequency-split into 6 MHz. bands. In the most common, single-cable design used for residential areas, most of these bands are *downstream* channels occupying the spectrum from 54 Mhz. to 300 MHz., 450 MHz., or higher depending on the particular system and converter technology employed. Downstream channels carry signals, typically entertainment video signals, from the cable TV headend to the subscriber. The *upstream* channels occupy the 5.75 - 29.75 MHz. range and carry signals from the subscriber to the cable headend. Frequencies from 29.75 MHz. to 54 MHz. form a *guard band* and no signals are transmitted at these frequencies.

Sources of noise on a CATV system include ingress of off-the-air signals from short-wave radio stations, insertion noise from equipment connected to the cable, white noise from the amplifiers on the system, loose connections at the drop, signal reflections because of impedance mismatch and beat products from rectification by corroded connectors. [5] The downstream channels of a cable system have low noise because the system is primarily designed for distributing video signals from the headend to the subscriber with good picture quality. The noise upstream is high because a CATV system has a converging tree topology; the noise arriving at the headend is the amplified sum of all the noise over the entire system. In addition, the frequencies used in the upstream direction are routinely used for short-wave communication and the CATV distribution plant make a very effective antenna at these frequencies. Network design must take this high level of noise into account.

2.2. Design Goals

The network consists of stations that transmit data to the headend on an upstream channel. The headend broadcasts the received data from the upstream channel to all stations via the downstream channel. A design goal for data communication over a cable TV network is to change the existing cable system as little as possible. Ideally, the only additions would be a network controller at the headend and modems at the subscriber end. A related goal is complete compatibility with the pre-existing cable equipment, television sets and cable channel allocation. In addition, expensive and complex equipment for data communication should be placed at the headend in exchange for cheaper and simpler equipment at the subscriber nodes. To ensure a high upstream bitrate, the upstream link to headend should be via the same cable as the downstream link; other types of upstream channels such as telephone lines are undesirable. The network should support multipoint communication instead of point-to-point communication, because any station may communicate with any other station without prior arrangement. A system must exist to allocate bandwidth on the network to allow communication and to support addressing of the multiple points on the net.

In summary, the problem of designing a CATV based packet switching network can be broken into four parts:

1. Transmission method from headend to subscriber

2. Transmission method from subscriber to headend

3. Subscriber control logic

4. Access scheme to determine bandwidth utilization

Two of these problems also appear in other data communication systems and only slight redesign is necessary. The two interesting problems are those of transmission from subscriber to headend and the access scheme used for bandwidth allocation.

2.3. Philosophy of Design

Both the headend and the subscriber nodes should be modular in design. In the headend, the control logic can be attached to any transmit and receive modems desirable. The upstream and downstream environments are different, thus different modulation techniques may be appropriate. Modulation and signalling techniques can be experimented with by switching modems.

The subscriber end is also modular, thereby lowering the cost per node of future changes by reusing most of the interface components. The four modules for the subscriber node are the receiving modem, the transmitting modem, the subscriber logic, and the bus interface.

Modularity allows the system to be put together in a building-block manner. The bus interface is easily changed as new computer systems are added and the modems can be changed during the initial testing to determine which transmission techniques work best.

3. Tested Network Components

Transmission from the headend to the subscriber and the subscriber logic are parts of the network design already tested in other data communication systems. Only slight redesign is necessary to use these ideas for the two-way data communications network.

3.1. Transmission from Headend to the Subscriber

The headend transmitter and the subscriber receiver use the same modulation system. Transmission from the headend to the subscriber is the easier of the two transmission paths to engineer because CATV systems are designed for video signal distribution from the headend to the subscriber with a good signal-to-noise ratio. Since the downstream data transmission channel will be a 6MHz. television channel, surrounded by video channels on a system designed for video transmission, it makes sense to have the downstream data method be compatible with video transmission as much as possible. CATV amplifiers expect certain signal levels, amplifier triple-beat characteristics are optimized for video signals, and television sets have filters for eliminating adjacent channel interference that expect carriers to be at certain frequencies. Research at the Sony Corporation has already demonstrated a data communication method compatible with CATV.

The Sony system has the right characteristics needed for the packet switching network:

1. High speed
2. Low error rate
3. Easily introduced into existing CATV networks
4. TV signals not adversely influenced
5. Low receiver cost

The Sony system transmits data at 7.4 MBPS using one channel of 6MHz. Transmission tests on North American and Japanese CATV networks have proven the system to be workable. The system uses VSB (Vestigial Side Band) transmission, the same as that of a television signal, modulated with a bandwidth-restricted, two-level baseband signal. The subscriber receiver is inexpensive because it uses mass-produced TV components. The system has a raw error rate of about 10^{-7} even in the worst carrier-to-noise ratios seen on CATV networks. [3]

A second possibility for the downstream channel is spread-spectrum modulation. In some areas, certain channels on a CATV network cannot be used because radiation of signals from the cable system might interfere with aircraft navigation. [1] Spread-spectrum signals can make the energy across a channel uniformly low so that radiation from the cable might be acceptable. Spread-spectrum receivers are more expensive than VSB receivers, but in an area where cable channels are at a premium, or in the future when demand rises, it may be worth the extra receiver cost to reclaim the unusable channels.

3.2. Subscriber Logic

The subscriber logic should be similar to the logic used in local area network interfaces, such as the Proteon proNET. The network interface should be simple to operate, intelligent enough to do several operations at once, and have multiple transmit and receive buffers to avoid missing packets. The maximum packet size should be about one or two kilobytes to keep transmission delay small.

3.3. Bus Interface

The bus interface will initially be for an IBM PC, since the PC is a popular computer for home and business use. The IBM PC bus is easy to interface to and the bus interface will be connected to the subscriber logic via a standardized interface.

4. New Techniques

Two new pieces of the design of the CATV network are the upstream transmission technique and the access scheme for bandwidth allocation. Most of the research will be focused in these two areas.

4.1. Upstream Transmission

The subscriber transmitter and the headend receiver have the same modulation method. Upstream transmission is harder than downstream transmission because of the high noise level that accumulates on the upstream channels. Low signal-to-noise ratios on the upstream channels suggest three modulation techniques: coherent PSK, phase-continuous FSK and spread spectrum. Coherent PSK and phase-continuous PSK have a low energy-to-noise ratio for a given error rate [4], suiting them for high noise channels. But, the type of noise found on the upstream channel of a CATV system would likely cause these two simple modulation techniques to perform poorly. The noise on a CATV system is probably closer to jamming than to gaussian noise distribution because of strong interference by short wave radio stations entering the upstream channel by ingress.

Spread-spectrum could be the modulation technique needed to overcome noise on the upstream channel. The military uses spread-spectrum techniques to produce systems that work even in the presence of intentional signal jamming.

Upstream spread-spectrum can be implemented in two ways. The first is to allocate a 6 MHz. slot in the upstream band exclusively for the network data. An advantage of this method is that it fits with the modularity of the CATV system - to add more networks, allocate more upstream channels.

The second method is to use spread-spectrum across the entire upstream band (24 MHz.) at a low enough signal level so that it can sit "underneath" conventional services occupying the same bandwidth. The advantage is that the network upstream channel can exist independently of other services using the upstream channels. The disadvantage is the possible effect on signal-to-noise ratio on the upstream channels by the spread-spectrum modulation and vice-versa.

Digital regenerators could replace some of the upstream linear amplifiers on the appropriate channels. Not all of the linear amplifiers would be replaced - only enough to reduce the noise to an acceptable level. The advantages are that noise is reduced and a simpler modulation system can be used. The disadvantage is the need to change some upstream amplifiers, making the system more complex and expensive.

4.2. Access Scheme

The access scheme for a MAN is a very important and difficult problem. Because the network serves the home or small business, the overall packet distribution will be bursty. Bursty traffic is difficult to schedule on a network with a large propagation delay and static allocation of time or frequency slots would make network utilization unacceptably low

The downstream channel will probably run at a higher data rate than the upstream channel because of the difference in signal-to-noise ratio between them. Therefore, the downstream channel can be broken into two parts: downstream data repeated from the upstream channel and access control information. The headend then gives permission to send on the access control part of the downstream channel.

The upstream channel could be split into two channels, a data communication channel and an access channel for queueing. The separate access channel allows stations to enqueue while data is transmitted by some other station. When the transmitting station is finished, the headend notifies the next station in the queue to start transmission. The downstream channel could be split into data communication channel and a enqueue acknowledge channel. The headend sends acknowledgments as the requests for transmission are received so the next station can begin data transmission without delay.

The stations must somehow make the desire to transmit known. CSMA/CD may work efficiently on a small network, but as size and speed increase, efficiency decreases. On a geographically large system, the efficiency would degrade to that of an Aloha system. RF collision detection is very hard to do; because the system is broadband, detection of conflicting signals on the upstream channel is uncertain, so the only way to know that the network is in use is to watch the downstream channel. Collision detection on the downstream channel only increases the delay of collision detection. Uniform station distribution implies that the number of stations at a distance r from the headend is proportional to r^2 because a CATV network covers area and is not linear as is an Ethernet. The average station distance is $R^2/2$, where R is the radius of the system. The large average distance between colliding stations also slows down collision detection. [2] Inefficient collision detection increases the loss of bandwidth because of collisions.

Polling is an access scheme frequently used on cable systems. The headend can poll every station to allow those with data to send to transmit. But, since most stations would not be ready to send traffic, polling is slow and thus not suited to the interactive environment that this cable network would support. Polling a large network (100,000 customers) can take six seconds, causing low throughput. [1] Adaptive polling polls active stations more often than inactive ones, so the polling is more efficient, but still not fast enough.

Multiple polling systems poll all stations at once by giving each station an orthogonal channel for reply to the headend. Frequency division or code division (for spread-spectrum) orthogonal channels allow stations to reply within a short amount of time. The headend determines the stations with traffic and assigns a transmission order. The headend polls again when the queue is empty.

For a network with R meter spokes (2R meters in diameter) and a propagation velocity of V meters/second, the one-way propagation time is R/V seconds (amplifier delay is negligible). [2] A worst case polling cycle includes a request out (R/V seconds), delay time at modem (D_m), arrival time of the last possible carrier (R/V seconds), carrier frequency determination (F), permission to send packet out (R/V seconds), delay at modem (D_m), and time until the first message arrives (R/V seconds). The minimum polling time is $4R/V + 2D_m + F$ seconds. If there are N stations on the network, each assigned a unique poll-response frequency in a B hertz bandwidth slot, then each has a slot bandwidth of B/N. A rough approximation of bitrate is B/3 bits/second for a bandwidth B, so each of the N slots would have a bitrate B/3N bits per second. Only one bit per station (request to send) is needed, so the time per station is 3N/B seconds per bit. Since all stations transmit carriers simultaneously, the total time to receive the one bit for each of the N stations is 3N/B seconds, so F = 3N/B. The maximum polling time is $(4R/V + 2D_m + 3N/B)$. If R = 30km, V = .87c, $D_m = 1\mu s.$, B = 6 MHz., and N =1024, then the maximum polling time is 974 microseconds. Compare this with $N(2R/V + D_m + 3/B)$ for conventional polling, which is 28.7 milliseconds.

5. Other Considerations

5.1. Error Correction Codes

Packet error checking should be done end-to-end, but a lower-level check might be used to determine the location of defective cable, amplifiers or sources of jamming. If the raw bit error rate is low enough, then a parity bit should work; but a CRC will probably be necessary. CRC calculation on a chip, as on Ethernet cards, is a simple way to add this.

5.2. Security and Availability

For data security, users can implement end-to-end encryption. Using end-to-end encryption allows the user to use any type of data encryption desired (including none at all) and thus set his own level of security. Encryption will not stop traffic analysis, but for the type of traffic this network will carry, traffic analysis should be of minimal concern.

Interference is a more difficult issue to deal with. Interference on the downstream channel would cause local problems for other users on the same cable segment. If the upstream channel is still open, the interfaces being interfered with could complain to the headend. Knowing the network topology and the station numbers of the stations complaining should make isolation of the interfering station fast.

Interference with the upstream channel is more difficult to handle. One way to find the source of the interference would be to use a binary search with intelligent bridger amplifiers. The headend can remotely switch the amplifiers on and off to find and isolate the offending unit, leaving most of the system intact. [1]

6. Test Bed

Newton, Massachusetts has a cable system that is an ideal test site - a 400 MHz. hub-shaped system with upstream amplifiers. The Newton cable system is near MIT and the city of Newton has many potential test subjects - MIT faculty member that already have IBM PCs.

Bell Communications Research has a CATV system built inside a lab for experimental use. This system is smaller, has less noise and is not for commercial use. Running preliminary tests at Bell would determine whether the network equipment works and insure that there is no interference with other channels on the cable before testing on an operational commercial system.

Testing on campus can be done on the MIT cable system because the system is small and near by for easy access. Probably most of the system testing will be done on the MIT cable system for convenience.

7. Conclusion

The need for inexpensive, high speed data communication within a community is growing. Because the physical plant of CATV is already installed, a MAN built on a CATV system allows a communication system to be installed quickly and economically.

Initially, downstream transmissions will use the Sony VSB system and the upstream transmissions will use spread-spectrum. The subscriber control logic will be similar to the proNET interfaces, and the access scheme will be multiple polling. During system testing, other alternative systems will also be tried.

The CATV MAN project explores three important areas: transmission over a broadband cable, spread-spectrum communication and access schemes for MANs. All these areas are important to data communication in the future.

Acknowledgments

My thanks to Jerry Saltzer and Deborah Estrin for their ideas and comments on this paper. Their groundwork on CATV data networks really helped to get this project going.

References

[1]
Estrin, D. L.
Data Communications Via Cable Television Networks: Technical and Policy Considerations.
Technical Report MIT/LCS/TR-273, MIT Laboratory for Computer Science, May, 1982.

[2]
Deborah Estrin.
Some Technical Considerations in Using a 400-MHZ catv System for Data.
Local Network Note No. 29

[3]
Yuichi Kojima, Etsumi Fujita, Masakatsu Toyoshima and Yasuhiro Hideshima.
A New Digital Audio and Data Transmission System Using the CATV Network.
IEEE Transactions on Consumer Electronics CE-30(3):462-7, August, 1984.

[4]

John D. Oetting.
A Comparison of Modulation Techniques for Digital Radio.
IEEE Transactions on Communications COM-27(12):1752-7, December, 1979.

[5]

Semir Sirazi.
Data Communications on Coaxial-Based CATV Facilities.
Presented at an 802.6 meeting