M.I.T. Laboratory for Computer Science May 28, 1981

Local Network Note No. 30

SUBURBAN DATA COMMUNICATION VIA CATV--PRELIMINARY THOUGHTS

by J. H. Saltzer

Personal computers and intelligent terminals are beginning to show up in private homes, with several motivations:

- professionals who use these tools extensively in their office find with home facilities that they can continue work on evenings, weekends, and holidays; some are also working at home one or more days per week to avoid commuting and interuptions.
- entrepreneurs are discovering that desktop personal computers are powerful tools for starting and managing small or large businesses.
- hobbyists have started acquiring personal computers, for education of their children, experimenting with new technology, playing games, etc.

These three users of personal computers and terminals all have a common need--data communications, for movement for both data and programs between their own computers, data bank services, company computers, software suppliers, and future computer-oriented service establishments, such as neighborhood printers or data storage services.

A second data communication requirement is emerging in the area of energy and resource management.

- labor-intensive meter-reading (for water, gas, and electricity) is becoming less attractive as falling prices for electronics make meter-reading sensors feasible.
- centrally initiated shutdown of water heaters and similar energy management communications are becoming economically feasible as electronics prices drop and energy prices rise.
- security monitors (fire and water alarms, etc.) are being more often connected to a central monitoring site (or a neighbor's house).

Data communication is suburban areas has been an expensive proposition because although the distances are not especially great the only in-place

---

communication facilities are those of the telephone company. These are expensive for several reasons:

- The central switching equipment is both configured and priced according to their value for voice communication.
- The equipment is depreciated with 20-40 year lifetimes, so the average economic age is 10-20 years, thus the recent revolution in prices of electronics has not had much effect.
- Because the plant is designed for voice, data must be converted to resemble voice. This conversion is expensive and limiting in potential bandwidth.
- The telephone facilities are designed for continuous connection between pairs of parties, rather than short high-intensity bursts of data. Telephone connections, when used for data, tend to be alternately overloaded and idle, an inefficient method of use.

Recently, a new kind of communication plant has begun to appear in suburban areas: the 2-way cable television system. This plant, being new, takes full advantage of modern low electronics costs. In addition, the plant is designed as a high-bandwidth broadcast plant, which means that certain very high-efficiency data communication strategies that involve broadcast of low-average-rate data at a very high peak rate can be directly applied. The remainder of this note discusses some of the techical issues involved in using a suburban CATV system for data communications.

This discussion is specifically oriented to a proposed M.I.T. experiment in transmitting data over a two-way cable system just being installed in Newton, Massachusetts, by Continental Cablevision of Massachusetts, Inc. This 400 Mhz cable system will carry 52 television channels past 22,000 homes. The cable plant is organized as a set of spokes radiating from a head-end, each with a maximum length of about 8 Km.

Basic mechanics

The simplest approach to data communications on CATV seems to be the following: allocate one full television channel inbound toward the CATV head-end (Transmit) and one full television channel outbound from the CATV head-end (Receive). At the head-end, translate everything received on the Transmit channel to the Receive channel frequency, and broadcast it to the entire system. With this approach, a data signal applied to the Transmit channel at any point in the system will travel in to the head-end, and then out to all possible receiving points via the Receive channel. The idea is to use this pair of channels, which can get a data signal from any point in the system to any other point, as a shared "party line", in which data is sent in short, rapid bursts, just as in the Xerox Ethernet[1], which operates at a baseband rather than with TV frequencies.

The party line approach involves the following steps:

1. Divide data to be transmitted into packets of, say, 8000 bits maximum length.
2. Assign each data station a unique address.
3. To transmit a packet, wait until the Receive channel is quiet, then transmit the address of the target station followed by the data packet, on the Transmit channel.
4. All stations monitor the Receive channel, and make copies of any packets that are preceded by their own address.
5. The transmitter should also listen on the Receive channel to verify that the packet was transmitted properly. If, for example, two stations that are far apart both are waiting to transmit when the channel goes quiet they will probably "collide". If a transmitter always listens to its own signals on the Receive channel, it will immediately notice the colliding data and can stop transmitting to clear the channel pair. If, following a collision, a station waits a random interval (or an interval proportional to its station address) before trying again, the probability of repeated collisions is minimal.

This technique, known as Carrier-Sense Multi-Access with Collision Detection (CSMA/CD), has been analyzed extensively in the technical literature[2], and is used at baseband in several high-performance local area network products.* It has also been applied, experimentally, on CCTV systems at the MITRE Corporation, and at least three companies now offer CSMA/CD products that are intended for use on cable TV systems.** A preliminary mathematical analysis suggests that with a 2 Mbit/sec. data rate, 8000-bit packets, and a 19 Km system diameter, a channel utilization of some 80% of capacity can be achieved without excessive delays[3].

---

* E.g., The Ungermann and Bass Network/ONE, the Digital Equipment Corp/Xerox/Intel Ethernet, and the Xilog Z-net.

** M/A Com, Sytek, and Ungermann and Bass.

More sophisticated ideas such as using channel time allocation strategies developed for satellites (which require a "smarter" head-end installation) are worth exploring. However, the particular cable plant configuration planned for the Newton system appears to allow the simple CSMA/CD approach, so that approach is proposed as the initial access control strategy.

Bandwidth

The long-term needs for data communication bandwidth are very difficult to assess, but for the short term we can make a variety of observations that allow one to choose sensible parameters for an experiment. The experience of the Xerox Palo Alto Research Center in providing a 3 Megabit/second shared party line among 150 personal computers in a business environment suggests that individual computers tend to be limited in their average appetite for data communication; Shoch and Hupp report that in the busiest minute of the day only 17% of their 3 Megabit/second capacity is used; in the busiest second under 40% of the capacity is used[4].

From the other side, we can observe that of the possible data rates one might propose for a 6 Mhz wide television channel, a single frequency shift keying system with two or four frequencies can operate at 2 Megabits/second with off-the-shelf chips intended for TV use. Lower speeds do not offer much extra economy (the same number of components must be used) while higher speeds would require a substantially more sophisticated modulation scheme.

Thus a data rate of 1 to 2 Megabits/second appears to be both

- achievable at moderate cost
- adequate for an experimental system with up to a few hundred stations.

Thus we propose to use a data rate in this range.

Privacy

Any use of a CATV system for point-to-point data communications must introduce a problem of communication privacy, since a signal targeted to any one party must go down a wire that is shared with many other parties. Further, in a suburban CATV system there is no control over who the other parties might be or now they might behave. Data encryption seems to be the best approach to deal with this problem. We propose to use the National Bureau of Standards Data Encryption Standard (DES)[5] on an end-to-end basis, using key distribution protocols of Needham and Schroeder[6]. (For the first round of experiments, privacy is not important, so encryption will be omitted.)

A second security problem is "denial of service". In a CATV system, there seems to be no way to prevent a hostile intruder from disrupting data communications by jamming one of the pair of data channels. This problem is intrinsic to use of a shared transmission system, and experience will be required to learn how serious it is. For most of the applications suggested above, there is no real motivation for an intruder to disrupt service; disruption may occur more in the vein of vandalism. It seems appropriate to avoid or discourage using a suburban CATV system to communicate data that might be the target of disruption, e.g., sensitive corporate data, police communications, burglar alarms, bank transfers, etc.

Traffic Monitoring

Because the system is inherently broadcast, one can place at any point a traffic monitor to gather statistics on actual use. Such a traffic monitor is essential for three purposes:

- to compare the presented load with the channel capacity, so as to obtain planning information that warns that more capacity will be needed (or new data customers should be discouraged.)
- to help isolate troubles such as a berserk transmitter.
- to gather statistics on who is sending messages to whom, both for capacity planning purposes and probably for billing.

This traffic monitor is probably a slightly modified, receive-only station that makes a copy of the target address, source address, and length of every data packet that comes down the Receive channel; it could either record those three items on tape or disk or else simply update a matrix of data volume indexed by source and target address.

Interconnection

An obvious extension of a suburban CATV-based data communication system is to interconnect it with other, similar data communication systems to allow communications to flow across city boundaries; many interesting problems, both technical and policy, are raised by such internetwork connections. To experiment with the issues involved, we propose connecting the Newton CATV data network with the M.I.T. data network, using a gateway computer attached on the one hand to the Newton cable and on the other to a microwave (or, if necessary, telephone) link to the M.I.T. campus in Cambridge. The gateway computer will forward packets between the two networks using the Internet Protocol (IP) already in use at M.I.T.[7].

An interesting property of the broadcast strategy for data communications within the Newton cable is that the gateway can be placed anywhere, not necessarily at the site of the cable head-end. This flexibility may be especially useful in providing interconnection between CATV data services of adjacent cities, since a gateway can be conveniently placed at the city boundary; if properly sited a single gateway computer could connect three cities and, perhaps, a microwave link to other tri-city gateways or a satellite earth station. Thus one can envision the CATV plant becoming a local data distribution system that connects both the metropolitan and national data communication systems.

References

1. Boggs, D.R., and Metcalfe, R.M., "Ethernet: Distributed packet switching for local computer networks," Comm. ACM 19, 7 (July, 1976) pp. 395-404.
2. Tobagi, F.A., and Hunt V.B., "Performance Analysis of Carrier Sense Multiple Access with Collision Detection," Proc. Local Area Comm. Network Symposium, May, 1979, pp. 217-245.
3. Estrin, D., "Some Technical Considerations in Using a 400-Mhz CATV System for Data," M.I.T. Laboratory for Computer Science Local Network Note No. 29, May 14, 1981.
4. Shoch, J.E., and Hupp, J.A., "Measured Performance of an Ethernet Local Network," Comm. ACM 23, 12 (December, 1980) pp. 711-720.
5. "Data Encryption Standard," Federal Information Processing Standards Publication 46, National Bureau of Standards, January, 1977.
6. Needham, R.M., and Schroeder, M.D., "Using Encryption for Authentication in Large Networks of Computers," Comm. ACM 21, 12 (December, 1978) pp. 993-999.
7. Postel, J., "DOD Standard Internet Protocol," Information Sciences Institute, Internet Experiment Note #128, January, 1980.