TECHNICAL CONSIDERATIONS OF TWO-WAY INTERACTIVE CATV

Niraj Jain, Tarek Saadawi, Mischa Schwartz

PHILIPS LABORATORIES
Division of North American Philips Corporation
Briarcliff Manor, New York 10510

ABSTRACT

This paper focuses on two-way data transmission over CATV. The types of services envisioned, their traffic statistics, as well as the trade-offs between centralized and distributed switching approaches are presented. Some protocol design issues are considered. Single-frame video transfer is discussed as a representative example of wideband traffic.

INTRODUCTION

Several different techniques for providing an alternative to the conventional analog wired local loop are expected to emerge in the next few years. CATV systems have a high potential for local distribution of information. This is understandable in light of the growth in cable penetration (approximately 41% of U.S. house holds, with gross revenues of about four billion dollars for basic cable services). Experiments with two-way interactive CATV have been under way for the last decade. A number of pilot projects are in operation. Examples of such projects are the Qube System by Warner Amex [1], and the INDAX System by Cox Cable [2].

The element common to the different CATV pilot project proposals submitted to the IEEE 802 MAN (Metropolitan Area Networks) committee [3] is that the CATV data system is treated as a Head-End (HE) centralized system. All input traffic (data) travels to the HE on an up-stream frequency, is frequency converted at the HE, and is then transmitted on the down frequency to its destination. Obviously, such an approach emerged from the currently operational entertainment CATV systems, and from the fact that current broadband local area networks, such as Sytek, WangNet and others, utilize similar approaches.

In this paper, we consider an alternate approach for two-way interactive CATV. It is based on a distributed store-and-forward network, with intelligent switches installed at different locations in the CATV plant. These switches (or nodes) must perform the usual packet-switched store-and-forward functions of routing, flow control, error correction and detection, etc. Databases (localized memory) may even be installed near these switches. The users contend in accessing the channel only to the nearest node. This reduces the number of users in contention for the channel, and as a consequence, the channel access problem becomes less critical than in the centralized case. Another obvious advantage is the efficient utilization of the CATV channel. The only drawback to such an approach might be the cost of the switch. Yet with economy of scale, the high rate of data growth, and the gain in channel utilization, this approach might be justifiable.

A set of protocols has to be established for communication over the CATV network. One possibility would be to adapt popular existing protocols such as X.25 for our use. However, it will be seen that several architectural features unique to the CATV network, suggest modifications that will enhance performance.

In CATV systems, a set of 6 Mhz channels are allocated exclusively for two-way activity. These may be further subdivided into broadband and narrowband channels to optimize performance for the various kinds of traffic encountered. Some possible strategies for channel allocation will be considered along with typical examples of traffic.

Single-frame video transfer capability is an interesting option made practical by the high-speed channels that are available on CATV systems. We will look at some typical applications which are expected to be of value to home and business users. Several alternative techniques for single-frame video transfer will be analyzed and comparative performance figures presented.

DESIGN CONSIDERATIONS OF TWO-WAY CATV

Figure 1 shows some design considerations for two-way CATV. Determination of the services envisioned over the system is essential in designing the system. These services could be classified into two generic types of data traffic: narrowband and wideband traffic. The specific bandwidth (or bit rate) requirements of each type would be determined from the kinds of services and applications foreseen. It would appear that at least two broad classes of data transmission must be considered. A narrowband group is represented by interactive bursty traffic, while the second

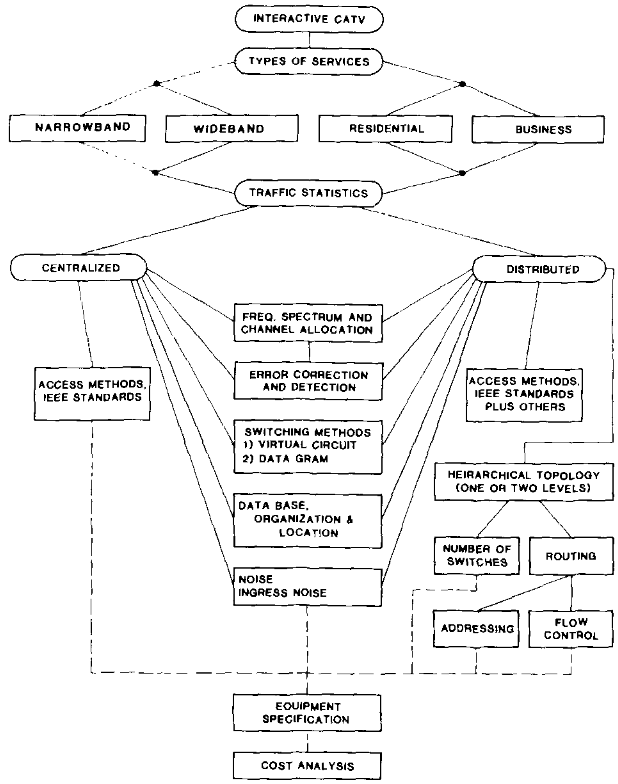

Fig. 1. Design Considerations of Two-Way CATV.

group consists of wideband information such as single-frame video transmission, facsimile communication, videotex, and file transfers. Another possible grouping might be residential and business services.

In many typical applications, users are basically interacting with databases. The databases could be located at the head-end, or in the "outside world" with access through a gateway; they could also be distributed at a number of nodes in the network.

Narrowband and wideband traffic statistics for the network have to be estimated for the services to be provided, and categorized into one of the two classes noted above, the expected length of messages, message inter-arrival time, the number of active subscribers during a busy period, and the mix of services used by the subscribers.

For both approaches, a number of frequency allocation strategies are possible: a single broadband data channel for each direction of traffic with TDMA structure superimposed, multiple FDM data channels (as implemented in the Cox Cable INDAX System [2], the Sytek broadband LAN system and others), or, possibly, two or more frequency subchannels, some for narrowband traffic, others for broadband traffic.

TYPES OF SERVICES AND TRAFFIC STATISTICS

The services foreseen on a two-way interactive CATV system could be classified as follows:

Information. Electronics news, Classified ads, Yellow pages, Employment opportunities, Insurance information, Travelogues/Airline schedules, Education, Income-tax tips, Online library/ encyclopaedia, Community news/Bulletin boards, Weather, "Do-it-Yourself" manuals.

Entertainment. Network programming, Interactive games, Trivia quizzes.

Financial. Banking services, Bill payment, Wall street news, Financial analysis newsletters, Archival database, Credit card verification.

Communications. Videotex, Teletext, Electronic mail, Digital telephone, Single-frame video transmission, Teleconferencing.

Computing. Programming, Data storage, Word Processing.

Miscellaneous. Telemonitoring (Medical alert, Security, Fire, Theft), Telecontrol, Telemetering, Catalog shopping (and shop-at-home).

The IEEE 802 Committee on Metropolitan Area Networks [3] has identified five distinct services to be addressed:

1) Bulk data transfer
2) Digital voice trunks and compressed video
3) Videotex and transaction processing
4) Low-speed terminal traffic
5) Local area network interconnect

Traffic statistics for the different services are generally not available since most of these services are still in the conceptual phase. Nonetheless, one can generate some reasonable numbers. We attempt to model a typical local CATV community and estimate the expected traffic during a day time busy hour. Only interactive (bursty) use is considered. This might represent carrying out "from-the-home" shopping, users accessing bank accounts, branch offices of a business communicating with one another, etc.

Consider a community of 1000 subscribers, with 900 homes and 100 stores/businesses. Let 400 homes and all 100 businesses access the system during the daytime busy hour. (Home users may be communicating with one another, with stores/ businesses, or elsewhere in the larger network). This model would thus provide statistics for a portion of a typical CATV system. Businesses also communicate with homes, with one another, or elsewhere in the network.

a. Home User Traffic

Let each home use the network for a typical 15 minute session, with 5 interactive transactions carried out during this time. A transaction consists of 60 characters from the home terminals, with 400 characters being received in reply. One half of the stores/businesses are involved in these transactions. The average interactive information rate during this busy hour is thus:

400 homes x 5 transactions per hour x $\frac{460}{3600}$ char./sec ≈ 240 char. per sec.

b. Stores/Businesses

The remaining stores and businesses in a typical scenario might be grouped as follows:

a. One large business, 1000 transactions/ hr., 80 characters outbound, 400 characters inbound, producing about 400,000 characters/ hour or about 100 characters/second.

b. 50 small businesses, 1 terminal each, one 400-character transaction/2 min. or about 160 character/sec., total.

The total community traffic, averaged over the hour, is thus about 500 characters/sec. Other examples will be considered as well in the course of discussion.

CENTRALIZED APPROACH

Most of the proposed two-way interactive systems employ the centralized approach in which all users on the system share the same channel. (In the INDAX system, communities of users are assigned one narrowband channel each.) Once the user accesses the channel, he transmits his message to the Head-End on the Up-stream frequency. The message is then frequency converted to the down frequency and transmitted to its destination.

Such an approach has a number of drawbacks. First, all users in the system are accessing the same channel, and hence the access scheme becomes very critical and may result in a long delay for the user until he is successful in securing the channel. Second, it requires more channel capacity than a distributed scheme. This is due to the fact that all traffic must travel to the Head-end, even if the two communicating users are neighbors of each other. Lastly, a problem often encountered in a centralized scheme is the ingress noise that funnels back towards the Head-end from over 100,000 taps. This can be limited somewhat in non-random-access systems by "Bridger Switching" as employed in QUBE [1]. The Warner-Amex QUBE and Cox-Cable's INDAX are two examples of commercial systems employing centralized control. Other schemes proposed are essentially TDMA-based schemes, such as Reservation - ALOHA, a discussion on which follows later in this section. The Cox Cable INDAX [2] uses Carrier Sense Multiple Access scheme with Collision Detection (CSMA/CD). Because of the nature of the CATV environment, this access scheme significantly limits the bit rate as will be shown below.

CSMA/CD normally operates for wideband Local Area Networks (LAN) that cover a distance of 1 Km or less. A critical parameter in the design of CSMA/CD networks is the ratio, "a", of the round-trip propagation delay τ (in seconds) to the packet length, T (packet transmission time in seconds), $a = \tau/T$. This ratio, "a", must be much less than unity to ensure that all users may sense the channel correctly for prior transmissions before they begin to transmit. Otherwise, a large number of collisions will occur and this access scheme will degenerate to a random one. For local area networks, at distances close to 1 Km, the propagation delay is small compared to the packet length T and hence a is maintained very small. But in a CATV environment, where the geographical coverage may be 30-50 Km, the "a" parameter becomes too large for a wideband channel. Systems that propose using CSMA/CD overcome this problem by using narrowband (e.g., 28 Kbps in INDAX [2]) channels. They keep the packet length T, large, and hence "a", small.

The use of narrowband channels precludes high data rate traffic on the system. This defeats the very purpose of wideband CATV and works against the competitive edge which CATV potentially has over the data services provided by the telephone company.

We now examine the centralized approach further using a Reservation-ALOHA access scheme. The channel time is divided into fixed length slots. The length of the slots is determined by the length of data packets. Typically, users make requests for reserving a slot before transmitting. The slots are grouped into a frame of fixed length in which the first Y slots are subdivided into mini-slots during which reservation request packets may be transmitted. As shown in Figure 2, each slot is divided into V mini-slots, i.e., there are VY request slots in a frame. The length of the mini-slot is determined by the length of the reservation packet. The remaining slots in the frame (M slots) are reserved for data packets. The protocol behaves like normal slotted ALOHA at low channel utilization and moves gradually over to Time Division Multiplexing (TDM) as the channel load grows. This allows much better utilization of the channel and hence better throughput in heavy load conditions. A detailed description can be found in [4].

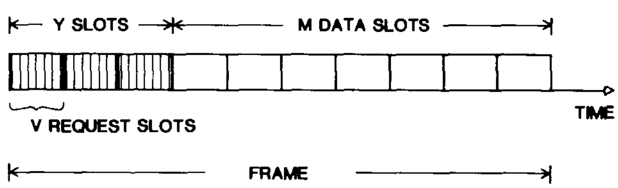

Fig. 2. Reservation ALOHA scheme.

A typical calculation has been performed for a 1 Mbps Reservation ALOHA channel. Assuming 10 char. request packets, 100 char. data packets, and an average message size of 6 packets, the optimal frame size is determined to be 53 slots, of which 3 slots are reserved for request packets and the rest for data. Each user will be transmitting data at a rate of 18.87 Kbps. In such a situation, the 1 Mbps channel can support a total traffic of 200 messages/sec without becoming unstable. The total number of users may be 2000 with each user initiating a traffic at a rate of 0.1 message/sec., or 20,000 users with each user initiating traffic at a rate of 0.01 message/sec.

We have looked at the performance of Reservation ALOHA as a typical centralized access scheme. This will be compared with a distributed control strategy in the next section.

DISTRIBUTED APPROACH

We consider two general cases of distributed control. In the first case, switches are installed only on a single main trunk [Fig. 3]. An appropriate location is at the bridger amplifiers on the trunk. In the second case, switches are installed on both the trunk and the branches [Fig. 4]. Locations of these switches could be both at the bridgers and at the splitters. We have carried out a detailed comparative analysis of both cases for an interactive, bursty traffic model. Details of the analysis appear in [5].

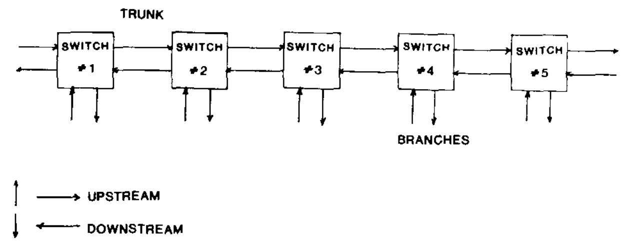

Fig. 3. Case 1; Switches only on the Trunk.

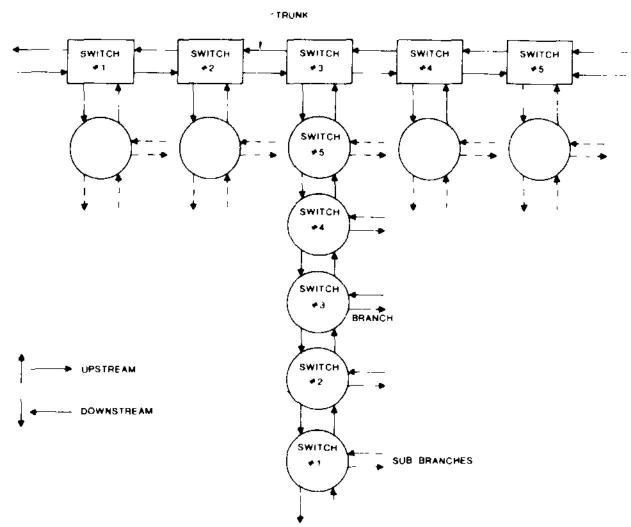

Fig. 4. Case 2; Switches on Both the Trunk and Branches.

In the analysis mentioned above, the network topology assumes one main trunk with 5 switches, 5 branches on the trunk (one at each switch location), with 5 sub-branches per branch. In case 1, we have a one-level hierarchy with switches only on the trunk for a total of 5 switches. Case 2 considers a two-level hierarchy with switches placed on the branches as well (at each of the sub-branch locations) for a total of 30 switches. In the traffic model assumed, each sub-branch inputs 0.1 Mbps of traffic for a total aggregate bit arrival rate or network load, over the 25 sub-branches, of 2.5 Mbps.

Consider the significance of the traffic assumed. For sub-branch traffic of 0.1 Mbps, average, and 450 character messages, the average message arrival rate per sub-branch is about 30 messages/sec. This corresponds to 300 subscribers/sub-branch, each inputting one 450 character message (5.5 lines of a terminal screen) every 10 seconds, or 1800 subscribers per sub-branch, each inputting a message every 60 seconds. The former case corresponds to a network with 7500 subscribers per trunk, the latter to 45,000 subscribers. The activity per subscriber is rather high, even in a busy-hour period, particularly if most of the subscribers represent single-family homes.

Consider an alternate example using the same numbers. Assume 900 home subscribers/sub-branch each inputting one 450-character message every 3 minutes. This provides an aggregate average sub-branch traffic of 5 messages/sec. Let an additional 25 messages/sec come from 10 large offices or businesses, each inputting an average of 2.5 messages/sec. Some of these "messages" could in reality be 450 character packets representing portions of a much larger file to be transmitted over the same channel as a non-interactive application. Other mixes of home and business traffic also result in the aggregate sub-branch traffic of 0.1 Mbps traffic assumed here.

It has been assumed that an originating data message is destined to any point on the network with uniform probability. Homogeneous message traffic statistics are assumed, with all users transmitting messages at the same average rate; messages are exponentially distributed in length, with the same average value at an originating point. Three performance measures are of interest. These are the: 1) trunk and branch flows throughout the network, 2) the throughput of each switch, and 3) the average message delay as a function of the total load on the network.

The results of the above mentioned analysis are summarized in Table 1 below. Note that the numbers given for average message delay implicitly include an access delay of 25 ms for case 1 and 13 ms for case 2. These numbers are based on the assumption that a Reservation ALOHA scheme is used for access.

TABLE 1. Comparison Between Different Approaches to Interactive CATV.

		DISTRIBUTED	
	CENTRALIZED	One Level Hierarchy (Case 1)	Two Level Hierarchy (Case 2)
Max Trunk Flow (Mbps)	2.5	0.6	0.6
Max Branch Flow (Mbps)	0.5	0.5	0.4
No. of Switches	1	5	30
Max Switch Throughput (Mbps)	2.5	1.7	1.6
Avg Message Delay (msec)	Not Possible	45	55

Note that the maximum trunk flow for the distributed approach (case 1 and case 2) is 0.6 Mbps. In the centralized case, the trunk would carry the full 2.5 Mbps which represents a significant increase of the data flow on the trunk. As a consequence, less channel capacity would be required for the distributed approach.

The centralized approach requires a trunk capacity of more than 2.5 Mbps. As an example, for a trunk utilization of 0.6 or less, at least 4 Mbps capacity would be required. The distributed approach, on the other hand, would require a trunk capacity of 1 Mbps or more, to keep the trunk utilization to 0.6 or less.

In case 2, each sub-branch inputs a traffic of 0.1 Mbps for a total network traffic of 2.5 Mbps. The resultant traffic flow on a branch varies from approximately 0.1 Mbps to 0.4 Mbps as opposed to 0.5 Mbps for the centralized approach or the distributed approach of case 1. On the trunk, the maximum traffic flow is 0.6 Mbps (similar to case 1) as opposed to 2.5 Mbps for the centralized case.

The maximum switch throughput for case 2 is approximately 1.6 Mbps in contrast with 1.7 Mbps for case 1 and 2.5 Mbps in the centralized case. Since the sub-branch channel capacity is 1 Mbps, its utilization is 0.1. This implies that the sub-branches are so lightly loaded that the simplest access strategies, such as random ALOHA would work well with relatively low access delay. For the range of traffic considered, the two-level hierarchy gives rise to a higher message delay than the one-level hierarchy. The store-and-forward process requires retransmission of each packet at each switch. There are now 30 switches, compared to 5 previously. Ultimately, the message delay for the one-level scheme will exceed that for the two-level hierarchy since the latter allows for more traffic to be input in the network.

Branch and trunk capacities in the distributed case are all 1 Mbps as noted earlier. In this case, the centralized system could not be used. Its apparent from Table 1 that the two-level hierarchy affords a slight improvement in branch flow and switch throughput. The average message delay has been increased, however, as well as the number of switches required. It would be our conclusion therefore, that for this traffic load, the one-level hierarchy is to be preferred.

ROUTING AND PROTOCOL ISSUES

So far we have only looked at performance issues. Ongoing work is addressing implementation issues as well. Paramount is the problem of designing a data-communication protocol appropriate for the distributed CATV environment.

Routing is a major issue in long-haul computer networks. At every node there are multiple routes for every source-destination traffic. In a CATV environment using a distributed approach, the routing problem is not as critical as in long haul networks, because of certain restrictions inherent to the tree topology of the CATV plant. At every node there is only one possible path for every source-destination pair.

The X.25 interface recommendation has been widely accepted worldwide. Ongoing work is considering the applicability of the X.25 protocol in the CATV environment. X.25 has originally been designed as an interface between a network and customer data equipment. We are looking at ways to adapt an extension of X.25 to network wide capability. The large number of users (20000-50000) motivates the need for a fast switching protocol. With the CATV tree topology and cascaded switches on the main trunk, a large number of virtual circuits must be supported on every switch. To reduce the processing overhead, we suggest a scheme [6] in which packet level acknowledgement and error correction is handled strictly on an end-to-end basis with minimal processing at intermediate nodes. This is a practical approach with the low error rate ($< 1 \times 10^{-9}$) and high speed 1 Mbps on the CATV data channel.

CHANNEL ALLOCATION STRATEGIES

As previously mentioned, we consider two types of traffic, narrowband and wideband, for two-way CATV applications. Suppose that two TV channels, one for the upstream and one for the downstream traffic (6 MHz each) are allocated for two-way services. With a proper modulation scheme, we can achieve a bit rate of at least 6 Mbps over the 6 MHz TV channel. Then a possible strategy would be to divide the 6 Mbps into 2 sub-channels: a 1 Mbps sub-channel for the narrowband traffic and 5 Mbps for the wideband traffic. We use a 1 Mbps sub-channel to conform to the previous analysis. Further study would be required to make proper choices for the channel capacities.

a) The one Mbps sub-channel would be allocated to low-volume traffic where a user (at an interactive terminal) transmits or receives messages of 1 screen page length or less. From the foregoing analysis of the distributed approach with switches installed only on the trunk, the 1 Mbps channel can support a maximum aggregate traffic of 4 Mbps with average message delay of approximately 45 msec.

b) The 5 Mbps sub-channel would be allocated on a dedicated basis to bulk data transfer such as digitized single-frame video and file transfer (a discussion of single-frame video follows in the next section). A user, wishing to transfer a single-frame or a file, would send a reservation message (call set-up message) over the 1 Mbps narrowband channel, indicating the origin and destination of the single-frame or the file. This reservation message is sent to all switches along the origin-destination path reserving the 5 Mbps channel along this path. Note that with the distributed approach, there could be more than one origin-destination path set up at the same time as long as there is no overlapping between the paths. An example for the single-level hierarchy distributed case is shown in Fig. 5.

Consider a data file of 10 pages, 24 lines/page, 80 characters/line. The data file has 19,200 characters or 154 Kbits. The transmission time over the 5 Mbps channel is 30.7 msecs.

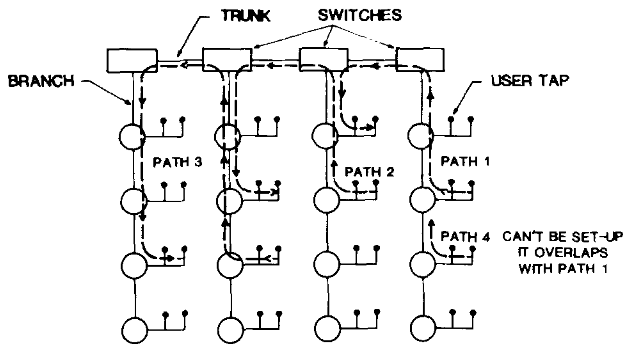

Fig. 5. Three Simultaneous Single-Frame Transmissions over the Mbps Sub-channel.

SINGLE FRAME VIDEO TRANSMISSION

We consider single-frame video transfer as a representative example of wideband traffic. The use of wideband CATV makes transmission of high quality single-frame video possible. Existing teletext and videotext systems, using conventional dial-up networks, of necessity are limited to simple graphics and mosaics. The use of single-frame video may be classified into two major categories:

(i) User-to-database applications: In this category, we would typically have uniformly distributed users accessing centralized databases. Typical applications incude illustrated 'How-To' manuals, art collections, numismatic and philatelic catalogs, car/appliance repair manuals, and travel brochures.

(ii) User-to-user applications: Here the traffic is of the user-to-user variety rather than from centralized databases. This could include multiple locations of a business exchanging video facsimile or teleconferencing, users sharing a picture of their favorite projects with other members of the cable community, signing up on a video dating service(!), or moving patient records between various departments of a hospital.

There are many ways in which single-frame video can be 'shipped' over the cable network. The choices are restricted somewhat by the limited number (typically 4 for sub-split entertainment systems) of return or upstream channels. We propose different strategies for handling each of the two classes mentioned above. The video information can be transmitted either digitally or in analog form.

Digital video has the advantage of allowing flexibility in transmission. It may be combined with other digital data, including voice, and sent out over dedicated circuit-switched channels. This requires more bandwidth than analog transmission, however, and additional hardware is required at the transmitters and receivers to handle the analog-to-digital and digital-to-analog conversions. Digital video may also be packetized and sent over store-and-forward networks.

Analog transmission for single-frame video in a CATV system would be carried out using a full TV channel reserved for this purpose. Reservation for the single-frame TV channel can be done through the 1 Mbps narrowband channel used for the interactive data traffic discussed earlier. As in the case of a single digitized frame, a number of simultaneous single-frame transmissions could occur as long as there is no overlapping between the routes (Fig. 5).

The first category, user-to-database, probably represents the bulk of the traffic. The source(s) may be located at the headend with 'n' broadcast channels allocated for this purpose. Consider the activity on one of these channels. A user interacts with a menu-structured database over the data channel to identify the sequence of video frames that he wishes to view on his home terminal. The frames would probably be stored on one or more video-disc player(s) under control of the database. Each frame, when available, is tagged with the user's address and transmitted over a shared broadcast (analog) video channel. All users expecting a frame, monitor this channel and when they see a frame tagged with their address, they acquire it (in a local frame storage device) for display. Assuming that frames are available to be transmitted at all times, and that NTSC video format is employed, 30 frames/sec may be transmitted, or stated equivalently, the system is capable of servicing 30 users every second on each channel. Each subsequent request for a new frame by a user is serviced similarly. The average delay (response time) is determined in large part by the traffic (request rate for frames) and the number of channels employed. The assumption of zero delay at the database can be satisfied in several ways: (a) Use a large number of source device working in parallel so that a frame is available on one of them at any given time. (b) Use a multiple frame storage device (high-speed) which queues up the frames to be transmitted, guessing the next frame in sequence if necessary.

The following example illustrates the performance of a typical channel:

Assuming a frame transmission time (including overhead) T, of 50 ms, and users requesting frames once every 30 sec., a peak load of 30/T = 600 users may be supported. If a typical user draws upon system resources for an aggregate of 30 min. over a 6 hour daily usage period, the average load that may be supported is 600 x 6 hrs/30 min = 7200 users.

In the example described above, the "users supported" are active users per channel per trunk. If the numbers are deemed acceptable, we see that a system with 4 trunks and 2 channels will support a peak load of 4 x 2 x 600 = 4800 users, and an average load of 57600 users.

In the second category, frames are transmitted from one user to another. Now the broadcast technique cannot work so a dedicated path must be set up between the source and destination. This may be in the form of a dedicated circuit (Circuit-Switching) or a virtual circuit (Packet-Switching) for the duration of a frame transfer. We look at each of these cases in some more detail.

In the circuit-switching scenario, a possible system would be one which uses circuit switches at each of the network nodes to handle the video traffic (fig. 5). The interactive data network can be used to establish and tear down circuits. A recent analysis of the blocking probability, under various load conditions, for this network-architecture, can be found in [7] which yields the following typical results.

Ex 1. #Active users = 8000
Switches = 5
#(Fdx)Channels= 3

With a traffic of	1 frame/100s	block.prob. =	6%
" "	1 frame/200s	block.prob. <	1%
" "	1 frame/1000s	block.prob. <	.1%

Ex 2. If the statistics are changed slightly to:

8000 home users transmitting 1 frame/1000s incl.
100 business users " 100 frames/1000s bp<.5%

If packet switching is used for the video transfers, digitized video is necessary. We assume here that the 1 Mbps data channel discussed earlier, will be used. A digitized video frame is typically 256 kbytes long. Sending this as one packet is not feasible. To break the frame up into smaller packets, the actual packet size would depend on the access scheme chosen. If 1000 bit packets were used, one frame would require about 2000 packet transmissions. On the 1 Mbps data channel, this corresponds to a total time of 90 seconds (2000 x 45 ms) to transfer the frame. Assuming that frame transfers are relatively infrequent, there may be little impact on network statistics. Even a single-frame transfer will nonetheless impact heavily on competing bursty traffic at the time of transmission unless it is spread out thinly over the packet stream. Larger packet sizes would help reduce delay but severely impact the buffer size requirements at the packet switch.

DISCUSSION

In this paper, we have covered some of the issues relating to the design of interactive CATV. The following are closing remarks:

- Distributed switching for data transmission over two-way CATV has a number of advantages over the centralized approach:

a) it can carry more traffic over a channel with a given capacity

b) it offers less message delay

c) it appears to offer reduced ingress noise

- The single-level hierarchy (switches installed only on the trunk) seems to be appropriate for the time being. It requires fewer switches than the two-level hierarchy (switches on both the trunk and the branches) and provides a smaller delay at lower traffic levels. (The two-level hierarchy can accomodate potentially more traffic and can be attained in an evolutionary add-on manner).

- An efficient network protocol is needed to utilize the inherent advantages of the CATV tree topology with the cascaded switches.

- Single-frame video must be considered an important application which offers an advantage over existing systems. Both circuit-switched and packet-switched techniques may be used. If a packet-switching system completely distinct from the data network were to be used, it must be compared in cost and performance against an equivalent circuit-switched system.

REFERENCES

1. T.P. McGarty, G.J. Clancy, Jr.; "Cable-Based Metro Area Networks", IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 5, Nov. 1983, pp 816-831.

2. M.L. Ellis, G.W. Gates, J.M. Smith, G.L. Peckham, H.P. Gray; " INDAX: An Operational Interactive Cabletext System", IEEE Journal on Selected Areas in Communications, Vol. SAC-1, No. 2, Feb. 1983, pp 285-294.

3. J.F. Mollenauer, "Data Communication Standards for Cable Systems", Conference Record IEEE International Conference on Communications, Boston, June 1983, pp 1527-1528.

4. Mischa Schwartz; "Computer-Communication Network Design and Analysis", Prentice-Hall, 1977.

5. Tarek N. Saadawi, Mischa Schwartz, "Distributed Switching for Data Transmission Over Two-Way CATV", IEEE Journal on Selected Areas in Communications, Vol. SAC-3, No. 2, March 1985.

6. T. Saadawi, N. Jain, M. Schwartz, "Protocols for a Two-Way CATV Network", Under preparation.

7. A. A. Jafari, T. Saadawi, M. Schwartz, "Blocking Probability in Two-Way Distributed Circuit Switched CATV", Under preparation.