FAST DIGITAL DATA CHANNELS ENABLE NEW CABLE SERVICES

J. Wendorf and H. Strubbe
Philips Laboratories, Briarcliff Manor, NY

Abstract

With the advent of digital delivery systems in CATV, fast data channels of multiple megabytes per second are feasible. Such channels open new possibilities for multimedia interactive applications. They can be used to provide an enhanced TV guide service that includes still pictures or even video clips. They can provide an electronic advertising magazine with pictures. Under appropriate conditions, they can lead to a variety of CD-I like applications. This paper outlines some of the new cable services made possible with fast, one way data channels, and explains the issues to be addressed in developing those services.

Keywords: television, cable, digital, teletext, multimedia, database, CD-I, TV guide.

INTRODUCTION

For many years broadcasters have included digital data in the video signals they send to TV receivers. Examples are teletext [1,2] and closed caption systems [3]. We consider such data channels "slow", as they carry less than 2 Kbytes/second. For closed captions, the rate is only 60 bytes per second. If in place of video an entire 6 MHz TV channel is used only for data, much higher rates are possible. For example, full channel teletext could carry roughly 0.5 Mbytes/second.

The advent of digital delivery systems in cable TV (CATV) provides the opportunity for even faster data channels of multiple megabytes per second. In the Philips Digital Video system [4], each 6 MHz TV channel can carry roughly 3.5 Mbytes/second. Such fast data channels open new possibilities for multimedia interactive applications.

In the following section we explain the concept of locally interactive applications, based on the repeated, one way transmission of complete databases. This is the primary technique for exploiting fast data channels. We then describe three applications that use this technique: an electronic TV guide, an electronic advertising magazine, and a system help and tutorial facility.

Following these examples we explain how a wide variety of applications, including currently unforeseen applications, can be supported through the transmission and loading of an application's code, along with its database. We then describe a further extension, the use of local storage such as a digital tape or optical disc, to capture all or selected portions of a database for later play back. Finally, we discuss some of the database design issues, including the overall structure of the database and the format of its constituent data items.

FAST DATA CHANNELS

In this paper we assume a one way transmission channel, as supported by most current CATV systems. With a two way interactive channel, many applications would be handled differently. We do not expect two way fast data channels to be widely available at low cost in the near future. However, the addition of a separate, slow back channel to a one way fast data channel is an extension that will be considered in our future work.

We also assume a single, 6 MHz TV channel is being used. In the future, multiple adjacent channels could be combined to provide even faster digital data channels.

With a one way data channel, it is not possible to send a query to a remote database to obtain a specified item of interest. Instead, the entire database can be transmitted from the remote location and scanned sequentially by a receiver to find and extract the required information. This is the basic technique used in teletext systems.

With a fast data channel of 3.5 Mbytes/second, a database of several megabytes can be transmitted within a few seconds. If the database is transmitted repeatedly, a query about the database can be answered within the time it takes to transmit the database again, in this example within a few seconds.

The underlying query handling system within the receiver is conceptually quite simple. It merely waits until the required data are transmitted again and picks them up at that time. In this way an application can interact with a user and provide responses to the user's input, based on the contents of the database. In practice, the high data rate of the transmission channel necessitates the inclusion of special hardware in the receiver, to scan the data stream for the items of interest and extract them for use in the application.

The scanning hardware can be simple, such as comparing a fixed bit field within each block or packet of the database against a given pattern. The scanning hardware could also be quite complex, such as matching each data item against a Boolean combination of several variable length and variable position fields, or against a complex relational database query [5].

The structure and format of the database, and the complexity of the query handling software within the receiver, are strongly affected by the type of scanning hardware used. In this paper we assume simple hardware, like the fixed bit field matcher described above. Some of the associated database design issues are discussed later, in the section on Database Format.

For databases larger than a few megabytes, the query response time can become unacceptably long, because of how long it takes to transmit the database each time. Depending on the application, this problem can be alleviated to some extent by one of two techniques. First, copies of the most frequently accessed data items can be repeated at regular intervals throughout the database. Second, copies of the most likely items to be accessed next can be cached by the receiver, usually in RAM. Both of these techniques are discussed in more detail in the section on Database Format.

To give a more concrete impression of approximately how much data can be carried in a given time on a fast data channel of 3.5 Mbytes/second, the following list provides a number of examples. These examples assume that state of the art digital compression techniques are employed:

- 5000 full TV screen text pages in 1 second.
- 300 fax pages in 1 second.
- 1200 two page letters in 1 second.
- 15 novels in 1 second.
- 25 one minute voice messages in 1 second.
- 15 high quality video stills in 1 second.
- 300 small color still pictures in 1 second.
- A daily newspaper with black and white pictures in 1 second.
- A weekly newsmagazine with color pictures in 3 seconds.
- The contents of four 800K floppy disks in 1 second.
- One full CD-ROM or CD-I disk (600 Mbytes) in 3 minutes.

EXAMPLE APPLICATIONS

We now briefly describe three example applications that are based on the repeated, one way transmission of multiple megabyte databases. First we consider an enhanced electronic TV guide service that demonstrates some of the features made possible by fast data channels. There already exist some electronic TV guides. Two examples are SuperGuide [6], built into the UNIDEN UST4800 satellite decoder, and InSight [7], soon to be broadcast via the PBS network. However, because these systems work with relatively low data rates, they only present the user with screens and menus consisting of text.

By exploiting a fast data channel, an electronic TV guide can be augmented with still pictures, audio clips, and even video clips. These can be used to attract consumers to the shows, for instance to advertise pay-per-view offerings. A still picture can be used as an identifier for a show, both to make it more attractive and to make it more immediately recognizable. This is particularly valuable for children who cannot comprehend text. Some shows, such as movies, can have their identifying pictures replaced by video clips. This would be like the clips shown on current "preview" channels, or like the "trailers" shown in movie theaters, and hence be much more effective for attracting viewers.

The size of the text portion of a TV guide database, containing schedule information and brief summaries of all the shows over a one week period, is roughly 1 Mbyte. There are approximately 1000 different shows in one week, each of which could have a separate identifying still picture. Each of those pictures requires roughly 5 Kbytes of storage. Thus, the total size of the TV guide database, including still pictures, is roughly 6 Mbytes. It can be transferred on a fast data channel in under two seconds.

Video clips require much larger amounts of information to be transferred. A fast data channel can simultaneously carry over 15 VHS quality, full screen movies, or over 100 small, picture-in-picture (PIP) size movies. By reserving half the bandwidth of the data channel for video clips, over 50 continuously repeating, PIP size clips could be provided. The text and still pictures part of the TV guide database could then be transferred in about four seconds.

Four seconds is too long for a viewer to wait for information about a TV show, so parts of the TV guide database would have to be cached in the receiver. Only information about the shows that are currently playing or coming up in the near future need be cached, since that is what the viewer will most often be interested in. As the viewer looks further into the future, it will take him/her a few seconds to study each screen of information. During that time the data for successively later shows can be cached, so that the information will be immediately available when needed.

Because of the quantities of data involved, video clips are never cached. Instead, they are played in real time, beginning at whatever point they are currently showing. Since each clip is typically quite short (30 to 60 seconds) and repeated continuously, starting in the midst of a clip should be acceptable to most viewers.

Our second example is an advertising magazine that uses a database of classified ads and optional associated pictures. A mechanism to search through the magazine by category or by keyword would make this more powerful than its printed equivalent. Color pictures, audio, and video could be exploited by advertisers to attract and inform a highly targeted audience of viewers. The viewers of an ad are those who, through their interactions and selections, have already indicated an interest in the product or service.

The size of the advertising magazine database, excluding video clips, would be comparable to a weekly newsmagazine with color pictures. Thus, it would take about 3 seconds to transmit on a fast data channel. Data caching techniques similar to those discussed above for the TV guide can be used to ensure that the system response times are acceptable.

A third example application is a system help and tutorial facility that can explain to users how to operate their receivers, and introduce them to the various TV, audio and data services available. By pressing a HELP button on the remote at any time, users can obtain context sensitive help information. Procedures can be explained and demonstrated through tutorial sequences. Pictures, audio and video can significantly enhance the learning experience.

The text part of the help database, along with simple illustrations, would probably be less than 1 Mbyte. With more elaborate pictures for demonstration sequences, the size might grow to around 2 Mbytes. This can still be transmitted in under one second. Hence it may be possible to provide the help facility without any caching of the database.

LOADING APPLICATION CODE

The flexibility of the data receiving device to support a wide variety of applications, including currently unforeseen applications, can be supported through a mechanism for loading application code. One obvious approach is to transmit an application's code along with its database. The receiver would first load the application code and begin executing it. The application code would then interact with the user and access information from the database as required.

If the data processing computer inside the receiver conforms to an open standard, third party software developers and service providers will be encouraged to produce applications for use with the system. A wide variety of applications would then be available. For example, one could imagine a "Video Game of the Week" data service channel that transmitted the code and data of multimedia video games, with the selection of games changing on a weekly basis.

For the standard data processing computer, we advocate a Compact Disc Interactive (CD-I) compatible platform [8]. CD-I was designed especially with consumers and TVs in mind. It has been adopted by a number of consumer electronics manufacturers, and an increasing selection of products and applications based on the standard are now becoming available. Although CD-I was originally developed with compact disc technology in mind, much of the standard and many of the existing software development tools can be readily adapted to the fast data channel environment.

LOCAL STORAGE

A further extension of the data receiving device, to increase its range of capabilities, is the addition of local mass storage facilities. We discussed earlier the use of RAM for caching selected data items from the transmitted database. With local mass storage, such as an optical disc, large portions or the entire contents of a database could be captured. For example, the complete contents of a compact disc would normally take about three minutes to transmit, which is too long to wait for the retrieval of a selected data item. However, if the transmitted data is first stored onto a local disc, it could be "replayed" later with faster interaction.

Local mass storage opens additional application possibilities. For example, music libraries or video game libraries could be made available on data service channels. Subscribers could then create personal collections of their favorite music or games by capturing selected items from the transmitted libraries and storing them on digital tapes or optical discs. The various copyright issues to be addressed by such applications, while very important, will not be discussed in this paper.

Another way to exploit local mass storage is to combine a local database with a transmitted database. For example, the contents of an encyclopedia could be stored on a CD-I disc. The latest updates to the encyclopedia could be transmitted on a fast data channel. By combining the two databases, the user would have interactive access to the latest information in a very large database. Updates to the CD-I disc could then be made at relatively infrequent intervals.

DATABASE FORMAT

Having looked at some of the potential applications of fast digital data channels, we can now summarize some of the associated database design issues. First it should be noted that although the above multimedia databases are transmitted over a TV channel, the format of their pictures, video and sound need not be identical with the compressed video and audio formats of a digital cable TV system. Sometimes the TV formats are the most convenient to use, such as for the encoding of video clips. But in other cases different formats might be preferred.

The type of data processing computer used, in our case a CD-I compatible platform, has a strong influence on the format of data items. To maintain CD-I compatibility, most items, such as still pictures, graphical objects, and even application code, will follow CD-I conventions. However, some database providers may prefer to supply their information in other formats. This can be accommodated by converting the data into a CD-I compatible form, either before transmission or after it is captured in the receiver. The needs of the database scanning hardware and software could also force additional restructuring of the data items to facilitate retrieval.

One major concern when designing the overall structure of the database is to minimize the number of times the user might have to wait for long periods to receive requested information. Two main techniques are available for dealing with this problem: repetition and caching. Repetition involves copying the most frequently accessed data items and transmitting them more frequently, thereby reducing the expected latency. The down side is that the overall size of the database increases, and thus the time to receive the less frequently accessed data items increases.

The second technique is caching. This involves storing in the receiver the most likely information to be requested next. The main problem is determining what information should be cached. This also has to be traded off against the amount of memory, usually RAM, to be used for the cache. Caching can be especially effective for hypertext structured applications where the user is presented with screens of information that are linked together in a predetermined fashion. While the user is viewing one screen, all of the screens directly linked to it can be retrieved. One of those screens will be what the user wants to see next.

For accessing the database information in different ways, index files can be transmitted along with the database and cached in the receiver. For example, the TV guide database might have one index file containing pointers to all of the shows for each channel, and a second index file with pointers to all of the shows for each time slot. By carefully designing the index files according to the application's needs, it should be possible to answer most queries in a single scan of the database, looking for an item identified by a particular pointer. Simple scanning hardware that matches a fixed bit field in each block or packet of the database against a given pointer/identifier is all that is needed.

CONCLUSION

In this paper we have explained how the fast digital data channels, made possible by the advent of digital delivery systems in CATV, can enable new cable services. We described a variety of locally interactive, multimedia applications, based on the repeated, one way transmission of complete databases. We also discussed some of the issues involved in the design of those applications.

Work is currently proceeding on the development of an experimental system that can support the types of applications discussed above. For demonstration purposes, an experimental, enhanced, electronic TV guide is being developed as an initial application. Studies by ourselves and others indicate that users like the concept of an electronic TV guide, and are even willing to pay for it. We believe that the TV guide may be the driving application that finally leads to the broad acceptance of interactive services by American consumers.

Throughout the paper, a single 6 MHz TV channel was assumed as the basis for a fast digital data channel. However, nothing prevents multiple adjacent channels from being combined into a single, even faster data channel in the future. This will enable even more elaborate interactive, multimedia applications to be developed.

ACKNOWLEDGMENTS

The ideas and work presented in this paper are the result of comments and contributions from many people. We especially want to thank Paul Rutter, Jhumkee Iyengar, Bill Lord, Brian Johnson, and everyone associated with the Briarcliff TV Data Services project.

REFERENCES

[1] Electronic Industries Association, *Joint EIA/CVCC Recommended Practice for Teletext: North American Basic Teletext Specification (NABTS)*, EIA-516, May 1988.

[2] UK Department of Trade and Industry, *World System Teletext and Data Broadcasting System Technical Specification*, 1986.

[3] J. Lentz, *et al.*, *Television Captioning for the Deaf: Signal and Display Specifications*, Rep. No. E-7709-C, PBS, May 1980.

[4] M. Balakrishnan and W. Mao, "An MPEG Standard Based Video Compression System and Applications", May 1992, in these Proceedings.

[5] T. Bowen, *et al.*, "A Scale Database Architecture for Network Services", *IEEE Communications*, Jan. 1991, pp. 52-59.

[6] P. Hallenbeck and J. Hallenbeck, "Personal Home TV Programming Guide", in *Proc. of ICCE 1990*, pp. 310-311.

[7] Chicago Tribune, "One-Button Recording Latest VCR Technology", Sept. 12, 1990, p. 3.

[8] Philips and Sony, *Compact Disc Interactive Full Functional Specification*, Sept. 1990.