

A System for the Delivery of Interactive Television Programming

Roger L. Haskin

IBM Almaden Research Center

650 Harry Road

San Jose, CA. 95120

Frank L. Stein

IBM Networked Application Systems Division

6710 Rockledge Drive

Bethesda, MD. 20817

Abstract

Interactive television is a technology for delivering television programming on demand to households and businesses. Rather than passively choosing from a predetermined set of programs at scheduled times, the viewer can request programming to be delivered to him individually on demand.

This paper describes a system being developed by IBM that is in active use at customer trials. The system supports movies on demand as well as more interactive applications such as home shopping. This system as deployed today on IBM's RS/6000 series of computers supports in the hundreds of viewers, but running on the SP-2 supercomputer and its follow-ons, the system can be scaled to support thousands of interactive viewers.

Introduction

Over the past several years, radical changes have occurred in the underlying technology of computing and communication systems. The plummeting cost of computer hardware and storage, the increasing capacity of wide-area communications, and particularly the advances in data compression technology, have enabled a variety of exciting new applications.

One such application is "interactive television", or ITV, which promises a radical shift in the way entertainment,

education, news, and advertising are delivered to the consumer - away from a one-size-fits-all broadcast paradigm to one in which all information is presented to the consumer under his interactive control.

Much hype has surrounded ITV, and there is still considerable debate about the rate at which this technology will be introduced and what form it will take. While this debate is taking place, the industry is progressing to gain experience with the technology and its applications.

IBM is very active in ITV, participating in trials at Bell Atlantic ([WN93], [PL93]), Cox Cable ([UPI93]), Okazaki City, Japan ([IBM94]), and Hong Kong Telecom ([PAT94]). This paper briefly describes the ITV system deployed in the Hong Kong Telecom trial. The system as described is capable of presenting interactive video applications to several hundred simultaneous viewers. The conclusion of the paper will discuss work that is underway to scale the system up to handle a considerably larger number (thousands) of customers. The following sections will discuss the types of applications the system is designed to handle, will describe the end-to-end architecture, and will discuss some of the major components in more detail.

The message that this paper intends to communicate is that the basic technology required for ITV is largely available now. What is needed next is a better understanding of the applications and services that will be commercially viable. The trial system described here is primarily intended as a

vehicle for content providers to use to gain experience, refine their existing applications and develop new ones, and thereby develop the business of interactive television.

Interactive TV Applications

Most ITV trials include Video on Demand (VOD) as their premier application. In VOD, the viewer selects a particular movie or TV show, which immediately begins playing back on the viewer's TV. VOD is considered to be a well-understood starting point for an ITV system because it is thought to be a simple application, because it is known that the video rental market is about \$12 billion annually in the US ([BS93]), and because customers find many aspects of renting videos inconvenient. Trips to the video store are time consuming, many people do not like going out at night, the videos customers want are frequently not in stock, and the customer has to return the video on time (or incur late charges). While ITV systems may not capture enough of the video rental business to pay for the entire ITV infrastructure, at least this business provides a quantifiable market.

A second application under consideration in many trials is Home Shopping. In this application, the viewer can examine potential items to purchase (using video and/or still pictures), and eventually initiate an electronic transaction to purchase the items, which are then mailed to the person. It is believed that ITV can contribute significantly to the shopping experience compared to catalog shopping, because the items can be demonstrated, modelled, and otherwise promoted using the video and audio capabilities of the system in a manner analogous to home shopping TV channels. However, unlike home shopping channels, the viewer, not the network, gets to select the items for consideration. Catalog shopping and shopping channels combined had revenues of approximately \$55 billion in the U.S. in 1994 ([AW94]), so this market could potentially dwarf the VOD market.

A number of other applications are being considered, including news on demand, electronic banking, education and training, interactive game shows, and action games. Most of these applications, including VOD and shopping, are similar in concept: viewers are primarily making selections that specify the content (movies, news, or product presentations) they wish to view, and then viewing the material.

Interactive TV Systems

An ITV system is much more than the hardware required to deliver video content to viewers. It includes an end-to-end infrastructure to support the creation of content, delivery of that content to viewers, and the billing systems to collect money to pay for the system. This section will summarize the components of the end-to-end ITV system.

Content and Application Creation

The term "content" refers to the actual video material such as movies, news, and the like, and "application" refers to the dialogs, queries, and other navigational software that presents content to viewers. The IBM ITV system¹ contains several components that support creating ITV content and applications:

- An authoring system for designing and implementing ITV applications,
- Production tools for user interface objects such as screen image displays, and
- A compression facility for preparation of the video content.

The authoring system allows the application writer to define user interface objects (dialog panels, menus, and the like) using pre-defined prototype templates. To generate a new application script, the author customizes the templates most closely matching his designed user interface. Descriptions of the objects and functions referred to in the application are stored in a database. The authoring system is designed to allow a graphical designer to generate an application with only minimal help from a computer programmer.

Targeting the authoring system to graphical designers rather than programmers was an important choice, as television is fundamentally different from the computer as a medium for presenting information. Some of this difference is due to technology (e.g. the much lower screen resolution of TV). A more important difference is that the television viewer community is more diverse - it includes young children, the elderly, and those unfamiliar with, or even hostile to, computers. Graphical designers are experienced in presenting information to ordinary people in an interesting

1. We refer to the system described here as "the IBM ITV system" for convenience, but note that substantial differences exist among the systems IBM has deployed in various trials.

manner, which is crucial to the acceptance of ITV. Even so, unhappy surprises often occur in the design of ITV applications, so formal usability testing is critical to developing successful applications.

Another component of the content creation system is the set of tools that produce the content and other presentation material for the ITV system. The sound and picture quality of both the content (such as movies or a merchandise video) and associated user interface objects (e.g. application dialog screens) are critical to the acceptance of the system.

ITV screen displays, unlike those for kiosks and desktop multimedia, are generally developed off-line, separate from the application development environment, with high-end animation and graphics rendering tools. This is because the audience for ITV is accustomed to the production quality expected of broadcast television. The IBM ITV system can accommodate files created by most high-end tools.

The IBM ITV system, like most such systems, uses MPEG (Motion Picture Experts Group) compressed video. MPEG is a lossy compression technique, and the picture quality strongly depends not only on the playback data rate, but on the compression system used and even the skill of the compression system operators. This is because MPEG is really a standard for *decoding* a video data stream. MPEG encoders are composed of proprietary hardware and/or software that is often optimized for some desired characteristic other than picture quality (i.e. speed, simplicity of operation, or the ability to fit on a chip). We have found in direct comparisons that a high quality system such as IBM's Power Visualization System (PVS), when operated with skilled operators, can produce significantly better results at lower bit rates than other systems.

Content Delivery and Presentation Systems

The content delivery and presentation system consists of a number of components:

- Customer premises equipment, normally a "set-top box" or *STB* that decompresses the video and sends it to the TV, and manages the low-level interaction with the viewer via a remote control.
- The video server, which manages the storage and playback of video content under the interactive control of the viewer.
- The distribution network that carries both the video content and control information between the video server and the customer premises.

In consumer ITV systems, the distribution network is operated by a carrier such as a telephone operating company or a cable company. There are differences between telephony-based and cable-based systems. We will describe the system architecture in terms of the telephony-based system that we have deployed.

A description of a telephony-based distribution network can be found in [BA93a] and [BA93b]. The network consists of a high-speed unidirectional (from server to customer premises) data channel supporting a DS-1 (1.544 Mbit/sec) data rate, and a bidirectional low-speed (16 Kbit/sec.) control channel.

The network comprises an open system designed to allow an individual customer to connect to any of a variety of video servers. To this end, the network contains a component called the Level 1 Gateway, which serves as an intermediary between the customer and a corresponding Level 2 Gateway in each video server. This gateway architecture is the so-called *video dial-tone* system mandated by the FCC for ITV systems that use regulated common carriers. It is analogous to the standard telephony system, which allows a customer to connect to any desired long-distance carrier.

The STB contains a control port, used to communicate with the video server and the Level 1 Gateway over the control channel using X.25, and a data port, over which the video server sends video and graphical images and downloads code to the STB at DS-1 data rate.

The video server contains an X.25 control port, used to communicate with the Level 1 Gateway and STB control ports. The video server's control interface to the Level 1 Gateway is called the Level 2 Gateway. This "gateway" architecture is what allows video servers from multiple service providers to connect to the system in an open manner. The video server also contains a number of DS-1 rate data ports sufficient to support the maximum expected number of active viewers.

Operational Support Systems

Even in early trial systems, the requirements for reliability, availability, and serviceability are stringent. First of all, customer acceptance depends on the perceived reliability of the system - unreliability leads to frustration, which in turn leads to disuse. Furthermore, the service provider usually desires to preserve his investment by placing the server into production at the end of the trial. The IBM ITV system supports multiple, redundant server components, and al-

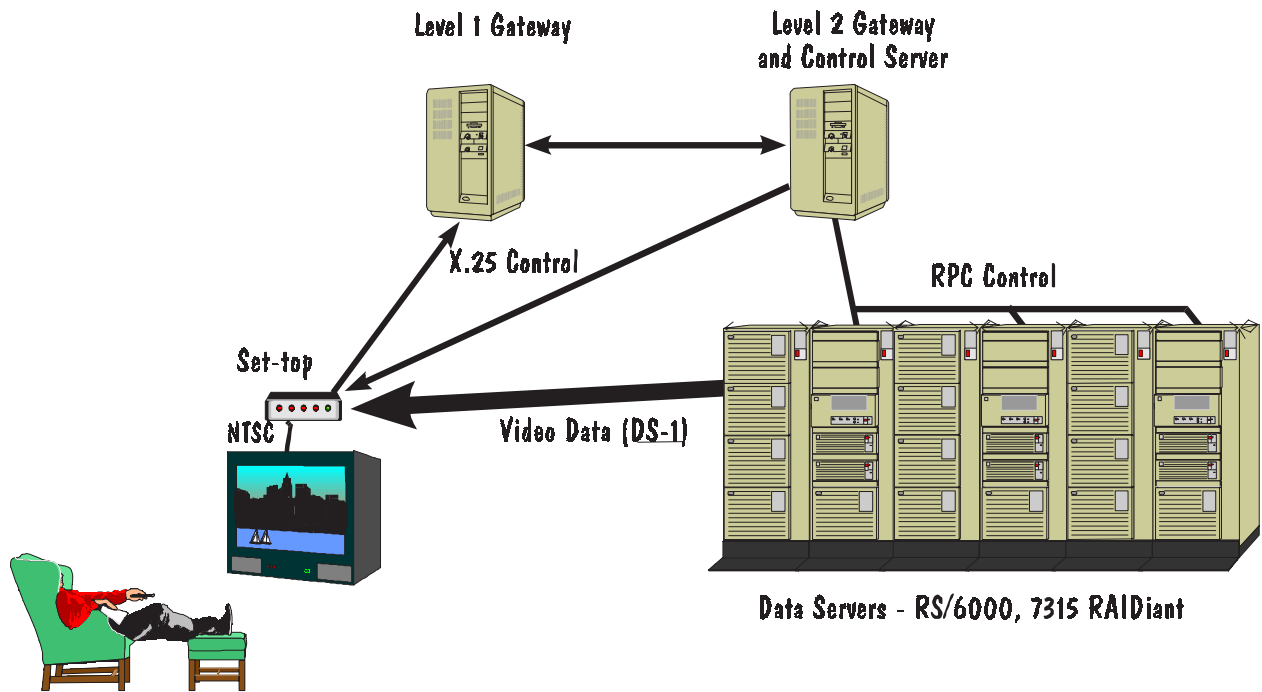


Figure 1 - Interactive Television System Architecture

allows the entire server complex to be managed by IBM's Netview/6000 Network Management product. Netview/6000 allows the operator to manage not only the server, but also the associated control and broadband network equipment.

Billing information is collected in a database on the server. In this early phase of ITV, service providers are trying to understand how viewers are using the system, so much more data is collected than is actually needed for determining a customer's bill. Almost any of the actions that the viewer takes can be collected for later off-line analysis by market researchers. For example, it's possible to collect the titles of all movies a viewer considered watching in addition to the movies that were actually watched.

ITV System Architecture

Figure 1 illustrates the overall architecture of the IBM ITV system. The function of the video server is broken down into three components: the Level 2 Gateway, which interfaces to the network, the Data Servers, which store and play back video and other data via the high-speed data ports, and the Control Server, which acts as an intermediary between the Level 2 Gateway and the Data Servers.

These components communicate via RPC, and as such can reside on a single machine or several machines depending on the size of the system. Figure 1 shows the Level 2 Gateway and the Control Server on one machine, with three Data Servers on separate machines.

Set-Top Box

The Set-Top Box (STB) manages interaction with the viewer and produces the video and audio signals for the TV. It contains the following major components:

- A control microprocessor with boot ROM and DRAM for program and data storage.
- A low-speed, bidirectional control port and an associated X.25 interface connected to the network's control channel.
- A high-speed (DS-1) data input port, connected to the network's data channel.
- An MPEG decoder.
- Hardware to generate the TV video and audio signals from the MPEG decoder's output.
- A detector for an infrared remote control.

The STB's boot ROM contains the code necessary to allow the STB to connect to the Level 1 Gateway and choose a video server.

When the server connects to the STB, the server downloads "set-top enabling code" (or *STEC*) that enables the STB to interact with the server.

The ITV application is executed cooperatively by the STB and the video server. Since the server downloads the application code executed by the STB, the partitioning of function between the server and the STB is up to the server. The IBM ITV System partitions the function such that low-level input processing (displaying menus and dialog panels, tracking the selection cursor, and handling input from the remote control) is handled by the STB, and all higher-level application functions are handled by the server. For example, to generate a menu, the server sends the STB instructions to receive video or image data for the menu background over the data port, and sends a definition of the menu (which screen areas represent which choices, along with overlay text and graphics) via the control port. Cursor tracking and menu highlighting are handled locally in the STB. When the viewer picks a menu item, the STB sends the choice (e.g. "item 3 selected") to the server. A description of the server side of applications is presented below.

It should be noted at this point that the STB described here is only one design point. Many design alternatives are possible, for example:

- The STB could be connected to a cable TV coax. In this case, data could be sent as MPEG-compressed video at a higher data rate (e.g. 6 Mbits/sec) or could even be sent as analog video. The control port could use any of several existing techniques for implementing a LAN on a broadband cable.
- The STB could be directly connected to an ATM network. In this case, a separate control port might not be needed, and the "upstream" (STB to server) connection could employ a higher data rate.
- The application could be partitioned so more, or perhaps all, of its logic resided in the STB. The video server could in this case be used strictly as a file and database server rather than as an application engine.

All of these design points are reasonable, and are being explored by various ITV systems. There are several advantages to partitioning application function between the server and STB as described here. First, customer premises equipment is subject to tampering, and as such is a

hostile environment with respect to both reliability and security. This argues for minimizing the function running in the STB. Second, the partitioning we have described allows a simpler, cheaper STB, which is important since unlike other system components, the cost of a STB is not related to the amount of use it gets.

Level 1 Gateway

As discussed above, the Level 1 Gateway (GWL1) provides the Video Dial Tone for the set tops. The GWL1 processes requests for service from the set tops. If the set top is a valid subscriber, the GWL1 responds with a list of available video servers. Alternatively, and appropriate for many trials, the set top can be programmed to automatically select a default server in the initial request message.²

The GWL1 is also responsible for establishing the high-speed data connection between the server and the STB. This is done in response to a request by the GWL2, specifying the server data port address and the STB identifier. To this end, the GWL1 maintains tables containing the data port address for each STB, allowing it to issue the appropriate commands to the data network switching equipment to establish the connection between the server port and the STB. The GWL1 supports several switch control interfaces including Bellcore TL1.

Level 2 Gateway

The Level 2 Gateway (GWL2) in the video server provides a standardized, open interface between the server and the network. As specified in [BA93b], the GWL2 performs the following functions:

- *Connection to GWL1.* Upon start-up, the GWL2 initiates an X.25 connection to the GWL1 for the exchange of control information.
- *Receiving STB connection requests from GWL1.* When the GWL1 receives a request from an STB to connect to a server, it sends a connection request to that server's GWL2. This request contains the STB's unique ID, its brand and version, and its network address.
- *Establishing control and data connections to the STB.* If the GWL2 determines that the STB is a valid user of

2. Using the earlier telephone system analogy, this is like pre-subscribing to a particular long-distance carrier.

the server, it allocates an unused data output port, and sends a request to the GWL1 to establish a connection between the data output port and the data input port of the STB. The GWL2 then establishes an X.25 control connection to the STB.

The GWL2 then downloads the version of Set Top Enabling Code appropriate to that brand and version of STB. The STEC software provides the set top with the means to process the subsequent commands it will receive.

The GWL2 also acts as the application engine in the IBM ITV server. For each connected STB, the application engine maintains the state of the application script. For each input event (e.g. a menu selection), the application engine acts on the event as directed by the script, usually by sending data or content to the STB over the control or data channels, and saves the new state of the script.

Applications employ various lower-level services. In addition to the services that support communication with the STB, these include a validation service used to authenticate purchase transactions, a database service, and the port and storage management (*PSM*) services offered by the Data Servers.

Data Servers

The Data Server accepts commands to play files over the high-speed data ports. The command set includes functions to start, stop, and change position in the file, allowing the flow of data to be controlled in much the same manner as a VCR or CD player.

The Data Server is a derivative of the Shark continuous media server described in [HASK93]. Shark is capable of storing the large amount of data required for ITV, and plays back large numbers of video streams simultaneously without interruptions or *glitches*. Shark also acts as a conventional file system, storing and retrieving conventional data. In the ITV server, Shark is used to store video, graphics (e.g. screen images), STEC software, and other application data.

Shark has been extended over the version described in [HASK93] to improve its ability to support ITV. Its new features include:

- *Wide striping*. Shark allows file systems to be striped across an arbitrary number of disks.

- *Low latency*. Shark employs a variety of scheduling and caching techniques to minimize video stream start-up latency.
- *Play lists*. Shark accepts requests to play multiple files sequentially with no intervening delay. A related feature is *looping*, which plays a video clip repeatedly with no delay between the end of one play and the start of the next.
- *Real-time file copy*. In systems with multiple Data Servers, files can be copied between servers at a guaranteed data rate. The use of this is discussed in the Control Server section below.

Wide striping balances load across the disks. The popularity of different pieces of video content will vary widely. Consider the VOD application. A blockbuster movie on its premiere night may attract virtually all the prime-time viewers. This is more than a single disk or even a single RAID subsystem can handle. Wide striping allows the load to be spread out across multiple disk drives or even multiple RAIDs. Popularity of content will also vary with time - children's movies might be popular in the early evenings, with first-run movies predominating during prime time, and more eclectic titles seeing activity as the night progresses. Striping all these movies across the same disks maintains balanced load even as activity against individual titles changes. Naturally, wide striping increases the system's vulnerability to disk failures. The Shark file system stripes data across multiple ranks of IBM 7135 RAID; the parity drives guard against disk failures.

Application menus and dialogs include images and video stored in the video server. This puts the server in the viewer interaction loop, making low latency important. Shark allows the start of selected files to be cached in memory, making start-up latency almost instantaneous.

Looping is used to generate video backgrounds for menus and in similar situations. The looping video is usually a computer-generated animation whose start and end frames are identical, for example a wheel turning through a revolution. Looping video introduces a sense of activity into menus and other dialog. Play lists are used to generate effects similar to the coming attractions at the start of a movie. Play lists eliminate the delays that would occur between videos if they were played separately, and avoids the alternative of concatenating the list into a single file. This makes it easy to change the lead-in material for a movie over time, or to vary it for particular viewers.

Control Server

The Control Server acts as the intermediary between the Level 2 Gateway and the Data Servers. It receives commands from the GWL2 and translates them into commands to the appropriate Data Server.

The Control Server maintains a catalog of all resources (files and data ports) on each Data Server. The catalog contains the state of each data port (i.e. in use, idle, or out of service) and information about the network topology (in some networks, not all ports can be used to communicate with all set tops). The GWL2 uses this information to choose a data port when connecting to the STB.

The file content directory is used to map the external names of content to their file names on the data server, and to manage hierarchical storage. This last is important because in many installations there will not be enough disk storage on each Data Server to hold the entire content library, so it may be necessary to move infrequently accessed files from hierarchical storage on demand. The Control Server implements a special form of hierarchy among the Data Servers. Each file has a home location on one of the Data Servers. When it is played by an STB attached to another Data Server, the Control Server initiates a real-time copy from the home server to the playback server over an FDDI ring. Since the copy rate is higher than the playback rate, playback can start after only a small amount of data is received by the playback server. Space on each Data Server not devoted to storing home copies of files is managed as a cache by the Control Server in order to keep file copying at a minimum.

Discussion

The IBM Interactive Television system described here has an impressive level of power and scalability. Nevertheless, like the ITV field itself, the IBM ITV server is in an early stage of development, and many changes will occur as ITV is deployed on a larger scale and extended to support new applications.

To enable the server to scale to larger numbers of viewers, new Data Server software has been developed. This new software, called Tiger Shark, allows the server to run on the IBM SP-2 parallel computer, rather than on independent RS/6000 machines. Tiger Shark implements a parallel file system for the SP-2, allowing any data port to access all data on all disks. This eliminates the need to cache separate copies of files on independent Data Servers. Furthermore,

the SP-2 can be managed essentially as a single machine rather than as a group of independent machines, simplifying operation of the server.

The authoring environment and application engine in the GWL2 supports applications of considerable sophistication, considering its simplicity. Nevertheless, as more experience is gained with existing applications, the authoring environment is evolving. A variety of changes are under evaluation, including integration with external databases and transaction processing systems to simplify order fulfillment for shopping applications, and a more general, GUI-based authoring environment.

Finally, use of the server in cable, ATM, and LAN environments is prompting us to generalize the functions of the GWL2, Control Server, and the stream management component of the Data Server. We are confident that a single, scalable architecture for all of these environments is attainable in the near future.

References

[AW94] AdWeek Magazine, Sept. 5, 1994, p.27.

[BA93a] Bell Atlantic Inc., *Signalling Specification for Video Dial Tone*, TR-72540, Dec. 1993.

[BA93b] Bell Atlantic, Inc., *ADSL Video Dial Tone Network Interface Requirements*, TR-72535, Dec. 1993.

[BS93] Bear-Stearns Research Report, *New Age Media*, May, 1993.

[HASK93] R. Haskin, *The Shark Continuous Media File Server*, Proceedings of IEEE 1993 Spring COMPCON, San Francisco, CA, Feb. 1993, pp. 12-17.

[IBM94] *IBM Japan Accepts Contract to Construct a Video-on Demand System in Okazaki City*, IBM Japan Press Release, Sep. 8, 1994.

[PAT94] Alan Patterson, *Hong Kong Telecom, IBM Map Video Effort*, Electronic Engineering Times, Aug. 1, 1994, p. 20.

[PL93] Larry Plumb, *Bell Atlantic Demonstrates Video on Demand over Existing Telephone Network*, Bell Atlantic press release, Jun. 14, 1993.

[UPI93] *Cox to Launch Interactive Tests in Omaha*, United Press International, Dec. 1, 1993.

[WN93] *Bell Atlantic, IBM Announce Agreement for Video on Demand Server*, World News Today, Jan 8, 1993.