

Architectures for Video Servers

Manu Thapar and Bill Koerner

Hewlett Packard

1501 Page Mill Road, Palo Alto, CA 94304

thapar@hpl.hp.com

Introduction

In the past few years a number of important activities and technology advancements that will influence the future of the industry have emerged. Cable companies have been laying out fiber cable providing enormous amounts of bandwidth capacity to residential neighborhoods. Memory storage is quickly becoming cheaper. Compression and decompression techniques are becoming more well understood and standards such as MPEG are gaining significance. The above factors now enable the possibility of various new services that can be provided to residences. One such service is video on demand (VOD). VOD will allow customers to choose the movies they want to view. Instead of driving to a video store to rent a tape, they will be able to choose a movie that is delivered to their home over a high speed digital network. VOD has the potential of being a multi-billion dollar business in a few years.

The essence of video on demand is that a number of digitized, compressed movies are stored in a video server and transmitted over a distribution network to viewers in their homes. The overall architecture is shown in figure 1. There are a number of possibilities for the distribution network [1]. The bandwidth into the residence is limited by the final interface at the home. Due to cost constraints, the equipment in the home is assumed to be minimal, comprising little more than decompression, decoding, demodulation, and display on a TV. The movie is played back in real time over the network and is not downloaded into storage in the home. The system cost is dominated by the cost of storage, and it is less expensive to concentrate the storage in the server where only enough storage for the number of simultaneous viewers is

needed, rather than having idle storage in the homes of all non-viewing subscribers.

This paper describes the key issues in the design of video servers. We provide an overview of some of the video server architectures currently being explored. We also briefly overview the potential services, service requirements, features, and communication facilities required for the whole infrastructure. We highlight the design issues that to us are the interesting research problems and outline some of our views on their solutions.

Video Server Requirements

There are a number of services in the area of entertainment and information/education that can be conceived. These include video on demand, news services, shopping services and educational information, to name a few. Of all these services motivating the development of a video server, the most immediate is the delivery of movies to home viewers. Cable companies want to expand their revenues by providing new services as well as providing telecommunication services. Already there are cable companies that are utilizing the large bandwidth of cable to provide internet services to the home. Not to be left out, and experiencing a threat to their revenues, the telephone companies are planning new integrated services. The telephone companies, cable companies, and third parties see an opportunity simply to displace a major part of the revenue collected by video stores for movie rentals. There are a number of other service opportunities. Hotels will offer movies on demand to their guests plus on-demand videos of hotel facilities, tourist attractions, and other information. Companies large and small will offer training programs and company news

Architectures for Video Servers

Manu Thapar and Bill Koerner

Hewlett Packard

1501 Page Mill Road, Palo Alto, CA 94304

thapar@hpl.hp.com

Introduction

In the past few years a number of important activities and technology advancements that will influence the future of the industry have emerged. Cable companies have been laying out fiber cable providing enormous amounts of bandwidth capacity to residential neighborhoods. Memory storage is quickly becoming cheaper. Compression and decompression techniques are becoming more well understood and standards such as MPEG are gaining significance. The above factors now enable the possibility of various new services that can be provided to residences. One such service is video on demand (VOD). VOD will allow customers to choose the movies they want to view. Instead of driving to a video store to rent a tape, they will be able to choose a movie that is delivered to their home over a high speed digital network. VOD has the potential of being a multi-billion dollar business in a few years.

The essence of video on demand is that a number of digitized, compressed movies are stored in a video server and transmitted over a distribution network to viewers in their homes. The overall architecture is shown in figure 1. There are a number of possibilities for the distribution network [1]. The bandwidth into the residence is limited by the final interface at the home. Due to cost constraints, the equipment in the home is assumed to be minimal, comprising little more than decompression, decoding, demodulation, and display on a TV. The movie is played back in real time over the network and is not downloaded into storage in the home. The system cost is dominated by the cost of storage, and it is less expensive to concentrate the storage in the server where only enough storage for the number of simultaneous viewers is

needed, rather than having idle storage in the homes of all non-viewing subscribers.

This paper describes the key issues in the design of video servers. We provide an overview of some of the video server architectures currently being explored. We also briefly overview the potential services, service requirements, features, and communication facilities required for the whole infrastructure. We highlight the design issues that to us are the interesting research problems and outline some of our views on their solutions.

Video Server Requirements

There are a number of services in the area of entertainment and information/education that can be conceived. These include video on demand, news services, shopping services and educational information, to name a few. Of all these services motivating the development of a video server, the most immediate is the delivery of movies to home viewers. Cable companies want to expand their revenues by providing new services as well as providing telecommunication services. Already there are cable companies that are utilizing the large bandwidth of cable to provide internet services to the home. Not to be left out, and experiencing a threat to their revenues, the telephone companies are planning new integrated services. The telephone companies, cable companies, and third parties see an opportunity simply to displace a major part of the revenue collected by video stores for movie rentals. There are a number of other service opportunities. Hotels will offer movies on demand to their guests plus on-demand videos of hotel facilities, tourist attractions, and other information. Companies large and small will offer training programs and company news

to employees at their individual convenience. Real estate agents will allow customers to view potential acquisitions and retail merchants will be able to provide information about products more effectively.

The specific requirements for a video server are not totally defined, since commercial trials are only about to begin. We expect to see small trial systems with 50-100 titles serving 10-100 subscribers. Pilot roll outs will probably have 100-1000 titles and 200-3000 subscribers. And real services, when fully configured, will likely require 10000 titles per server (with subscriber access to a total of over 100000 titles on a distributed array of servers) and 10000 subscribers per server. These numbers are also likely to vary with the area served and the service provider. Because of the wide range of these numbers, we feel it important to make the server scalable in at least the dimensions of number of viewers and number of titles, so that a common design and set of components can be the basis of a whole family of servers.

Another factor that influences the design of the server is the set of features that it must provide. Reliability, quality and privacy (security) are some of the key features. Since the video server provides a service that is can be accessed by

thousands of viewers and each call last around 2 hours, reliability of the server is of key importance. We use redundancy to increase the reliability of the architecture. The quality of the video stream, which translates into jitter free bitrate, is another important feature. We assume that movies are encoded digitally using MPEG (the standard from the Moving Picture Experts Group) at rates that begin at around 1.5 Mb/s but could also include rates such as 3, 4.5, 6 or 8 Mb/s. Ideally, a server would allow subscribers to select their quality level, so it would have to play movies back at a variety of bitrates. In some cases, the distribution network bandwidth and settop specifications will determine an upper bound to the bitrate that can be delivered. The storage requirement for a movie is proportional to the bitrate of its encoding; a minimal bitrate of 1.5 Mb/s results in a 2-hour movie taking up 1.3 giga bytes of storage. Enough buffering is included in the system to allow a continuous stream of data to be delivered to the network. Privacy and security are provided by encryption of the digital data stream. The encryption key is changed periodically to eliminate piracy of the data, so that no one views something they are not paying for. Another feature is the degree of interactivity. To explain this, we describe two varieties of services.

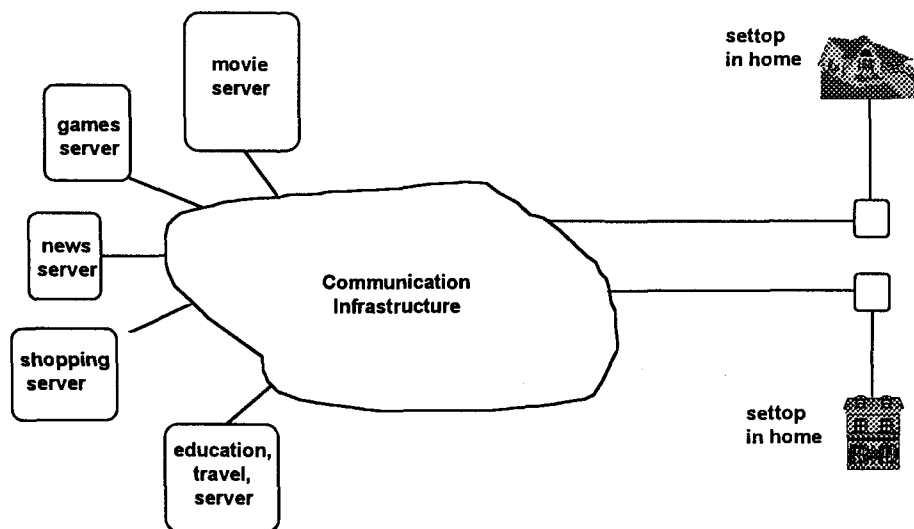


Figure 1: VOD Infrastructure

One potential service is near video on demand (NVOD). In NVOD, a popular movie is repeated at intervals of, say, 5 minutes, and a viewer can simply choose to watch continuously or jump between streams. The movie is broadcast to all the viewers, hence a large number of subscribers can be served by the same stream. This delivery mechanism is good for very popular movies (hot movies). The server needs to transmit only a single stream to satisfy multiple requests. For example, a two hour movie transmitted at an interval of 5 minutes requires 24 streams to serve an unlimited amount of viewers. At the high end, we have fully interactive video on demand (IVOD), sometimes called true video on demand, in which the viewer has complete virtual VCR functions of pause, fast forward, and rewind. Whereas NVOD allows a considerable degree of broadcasting and therefore a very large number of viewers per copy, IVOD is restricted to one viewer per stream. It is therefore the most expensive in terms of storage and network bandwidth per viewer.

Network Delivery of Video

There are a number of delivery mechanisms possible. We provide a brief overview here. Further details are available in [1]. Currently there are two main physical communication channels entering homes. These are the twisted pair wires used for telephones and the coax cable used for the cable services. Asymmetric digital subscriber line (ADSL) can be used to transmit digital data over twisted pair. ADSL I has an upper bound of 1.5 Mb/s for transmission of data. ADSL II improves the upper bound to 6.4 Mb/s. Delivery of information over twisted pair is thus limited by the data rate and cannot support very large data rates. Coax cable has a bandwidth of around 1GHz. It is split up into 6MHz analog channels. Using 64 QAM (quadrature amplitude modulation), 25Mb/s data rates can be transmitted in each of the 6MHz channels. Typically VOD related information will be carried in the frequency spectrum above 450 MHz, with most of the lower frequency spectrum being used for traditional cable delivery that is being done today. A limited

number of channels are reserved to carry reverse, or upstream, information to the server. However, upstream amplifiers do not exist in all locations and need to be included. A hybrid combination, using the twisted pair for upstream communication, and the cable for downstream data delivery is also possible.

We assume that many servers will reside in telco central offices or cable company head ends. Customers will directly access movies on their local server. If the movie they wish to view is not stored there, their server will establish a link to other local servers or regional libraries. We believe that by the time video on demand is widely deployed this communication infrastructure will be based on SONET/SDH and ATM (asynchronous transfer mode). Already the long-distance trunks are predominantly fiber and SONET based.

Video Server Architectures

The main physical components of a video on demand service are a video server, data delivery network and set-top converter box. Figure 1 shows these main components. The video server consists of the storage and control required to store movies in compressed format and play them back on request. It differs from a traditional database server in various ways. It has to perform a number of functions such as admission control, request handling, data retrieval, guaranteed stream transmission, stream encryption and support of VCR functions. Admission control is done for each request by determining if the request can be serviced by the available resources in the system. Because the transmission of video data is stream oriented, it needs to be delivered to the end viewer without any glitches. The system can service the request only if continuous delivery of the video stream can be guaranteed, once the stream has been started. Because of the non-deterministic nature of disk accesses, intermediate buffer memory is used to transform the bursty disk accesses into a continuous stream that is guaranteed to be glitch free. The communication paradigm is also different from that used in a conventional system that is part of

an enterprise network. In the case of VOD, the data is continuously sent from the server to the settop. There is no feedback from the settop asking for more data when the settop buffer is about to underflow. In an enterprise network the settop acts as a client and is continuously pulling the data from the server. In the VOD case, the data is "pushed" to the settop instead of being "pulled" by the settop. The buffer requirements at the settop need to be sufficient to prevent any underflow or overflow to occur.

The digital data is taken in compressed format from the disks and delivered to the network interface, where it is modulated for transmission over the distribution network to the home receivers, or set-tops. The set-top box, next to the TV, decompresses the data in real time and displays it on the TV screen. Decompression chips conforming to the MPEG standard are available from a number of vendors. The set-top box will provide the user interface that will allow the user to choose movies from a menu and also have a VCR-like interface for play, pause, rewind, and fast forward. The set-top box has to be very inexpensive in order to compete with existing CATV converters and

VCRs. Therefore it should contain only enough storage to compensate for network jitter, not enough to download the whole movie.

There are a number of important issues that need careful consideration in the design of a video server. Current computer systems are mainly designed for good computations. For the VOD application, the compute power of the server is irrelevant. Good input/output bandwidth, or efficient data movement, is significantly more important. Multiprocessors may be used as video servers, but that would not be a very cost effective choice since a lot of hardware is not utilized.

One possible architecture for a video server is to use an existing shared memory multiprocessor. Figure 2 shows a shared memory multiprocessor. The data is taken from the storage subsystem, transferred over the bus to main memory, and then transferred again over the bus to the network interface. The caches and processors are not utilized and the bus is also not effectively utilized due to the dual transmission of the data over it.

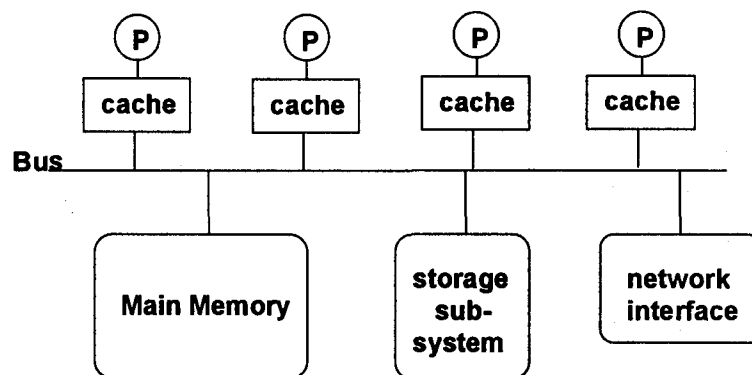


Figure 2: Shared Memory Multiprocessor

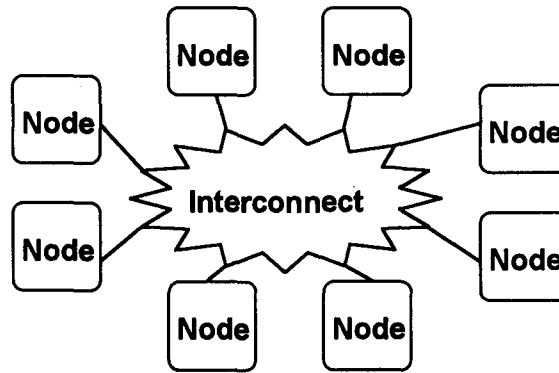


Figure 3: Distributed Memory Multiprocessor

Another possible video server architecture is a distributed memory multiprocessor. Figure 3 shows a distributed memory multiprocessor. Each node consists of a processor, cache, memory and interface to the interconnect. These architectures are designed for large scale computations using the message passing programming paradigm. Each node can be adapted to include a network and or storage subsystem interface. The computation related resources again are underutilized.

Scalable Architecture

As had been mentioned earlier, it is important for the design to be scalable, both in terms of the number of titles and in terms of the number of viewers. The VOD application involves large amounts of data transfers and comparatively small computations. We now describe an

architecture that can be used effectively for the large data transfers that are typical for video servers. It is tailored for the *data movement* problem instead of a general purpose *data computation* computer. Figure 4 shows a generic architecture for a scalable video server, consisting of a main controller and a number of data servers. Each data server contains its own storage devices. Scalability in the number of titles is achieved by increasing the amount of storage per data server. Scalability in the number of viewers is achieved by increasing the number of data servers. This makes the design modular in structure and scalable to larger systems. We follow an open systems design approach by using standard interfaces and protocols. These interfaces range from SCSI for the storage devices to ATM (asynchronous transfer mode) packets that are transported over a SONET interface over the distribution network.

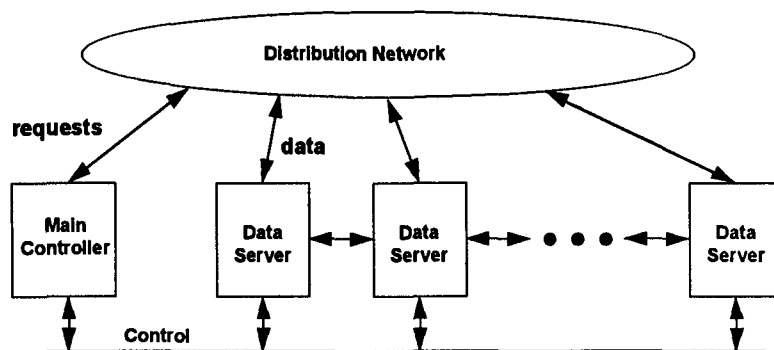


Figure 4: Scalable Video Server Architecture

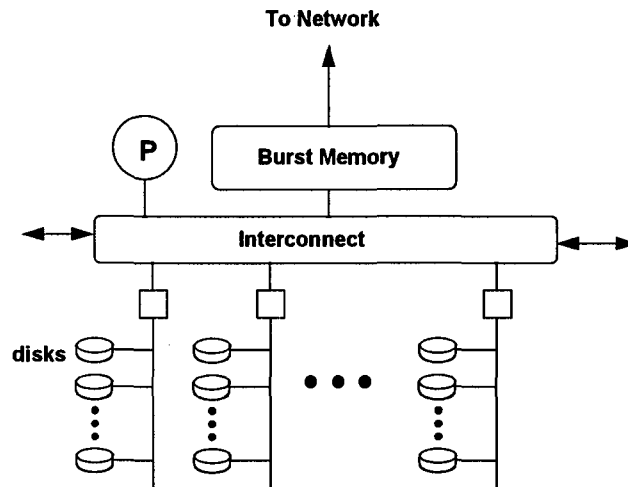


Figure 5: Data Server

The main controller can be a general purpose workstation, or an embedded device. It receives requests from the network and performs the admission control functions. It has the mapping of movies on I/O devices and controls the data servers. Control information consists of commands such as "play stream A", or "rewind/fastforward stream A". A data server is responsible for a particular video stream. Once it receives a command from the name server, it determines locally how to manage the available resources most efficiently. Each data server is connected to the storage devices. A data server is the replication unit and additional data servers may be added to the system to scale the bandwidth of the system. The main controller can run Unix. We follow the open systems approach throughout the architecture. Besides controlling the data servers, general services such as billing are also performed by the main controller. Viewer requests are received at the main controller via the back channel. The main controller has the necessary software and database to perform admission control and control the data servers. Each data server is connected to SCSI subsystems and to an ATM switch. For small configurations requiring only one data server, the ATM switch is not required, resulting in lower overall cost.

Figure 5 shows the structure of a data server. Each data server has information about the

movies that are stored on the subsystems it controls. The movie can be striped across multiple disks. By striping the data fully across all the disks in a data server, load balancing can be achieved more effectively and hot spots in the system can be prevented. A data server performs its own resource management for the movies that are stored locally. It can operate as a NVD server or a IVOD server. A standard control protocol between the main controller and the data servers allows more than one implementation of the data server to be used in the same overall architecture. The data server controls the transfer of data from the disks to the burst memory. The burst memory is used to buffer the disk blocks and remove the non-determinism in disk access times. Data is taken from the burst memory and transmitted in a continuous stream over the distribution network. One likely network interface is a SONET OC3 interface that is capable of transmitting data at a rate of 155Mbps. Other interface conforming to the open systems approach are also possible. Each data server module has the potential of serving a few hundred viewers.

Storage

Traditionally, database servers have been designed for a large number of small transactions. The throughput for these

transactions is optimized (transactions per second), but there is no notion of guaranteed, or real time delivery. In a video server, once a transaction (movie) is started, the complete movie file needs to be delivered in a guaranteed, stream oriented manner. A video server also needs to access a large amount of data of the order of a terabyte. This makes the storage cost a very important design criteria. The data layout on the disks needs to be optimized in order to obtain maximum utilization of the system.

Compression is required to reduce the storage requirements for the video database. A 2 hour NTSC format movie requires approximately 100 GBytes of storage. At \$0.80/MByte for disk storage, this translates to \$800,000 per movie, which would make the system uneconomical. Compression reduces the storage requirement to approximately 1.35 GBytes per 2 hour movie using a 1.5 Mb/s MPEG compression scheme.

From the experiments in our lab, this compression rate produces good visual results on the monitor. But this has not been tested commercially, and if 1.5 Mb/s yields inadequate quality, 3 Mb/s or more will be needed, resulting in close to 3 GBytes per movie. To cope with the vast amounts of storage needed even for compressed movies, we expect to use a combination of magnetic tape and disk storage to reduce the total system cost, which is dominated by the storage cost.

Figure 6 shows the various technologies in the storage hierarchy. RAM is the most expensive and has the highest bandwidth. At the low end we have jukeboxes that can be used to store large amounts of data, but they have large access times. A proper intermix of the various technologies can result in a cost-effective storage system for the video server.

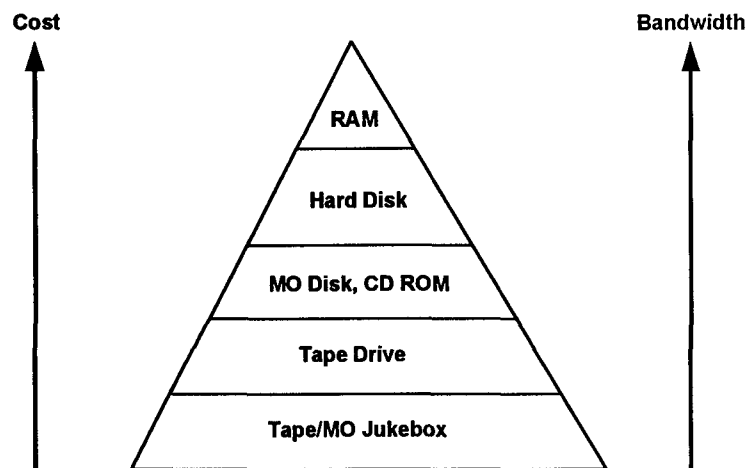


Figure 6: Storage Hierarchy

Storage Technology	Bandwidth (MBytes/sec)	Cost (\$/MByte)
RAM	> 80	60
Disk	4	0.80
Tape	2	.025
CD-ROM	.3	.05

Table 1: Storage bandwidth and cost

To have a cost effective design, the storage hierarchy described above needs to be exploited by using forms of caching. We expect SCSI devices to be used widely, with fiber channel devices to follow. The small computer system interface (SCSI) is a parallel, multimaster I/O bus that provides a standard interface between computers and peripheral devices. SCSI is a widely used protocol for connecting disks, tape drives, CD-ROMs and magneto-optic (MO) drives. Fast and wide SCSI systems can provide a peak bandwidth of 20 MBytes per second.

Disk densities are increasing at a rate of 60% each year, and prices are falling at a rate of 12% per quarter. The video server is a disk based cache of data that is retrieved either locally or remotely from archival storage. Corruption of data on disk is therefore relatively unimportant as it can easily be retrieved from tape or other media. However the failure of an entire disk is important as it will disrupt the service of paying viewers. High availability is therefore more important than data integrity. Specifically the storage system must have excess bandwidth available to cope with component, usually disk, failures.

To provide the bandwidth necessary to supply many streams from the same material the movies will be striped across a number of disks. Redundant Arrays of Inexpensive Disks (RAID) techniques that add additional disks can be used. There are a number of raid levels that have been defined ranging from RAID 0 to RAID 1. For VOD, data retrieval needs to satisfy some hard real time constraints. The array needs to be able to supply the data at the correct rate even when one of the disks malfunctioning. Due to this hard real time constraint, the RAID levels that are most useful for video are RAID 1 and RAID 3. In RAID 1 data is mirrored, or replicated.. The cost is thus double than that of a simple array of disks. In RAID 3 one additional disk is required for each array to store parity. If a disk fails, the parity data is used to generate the correct data. However the sophisticated rebuild algorithms required to reconstruct the data on a replacement drive while maintaining bandwidth to the application may not be necessary as the movie can potentially be replaced from tape.

If a movie is not available at the server, the data can be obtained from an archival storage. It is likely that this transfer will be non real-time initially, with real time capability in the future.

Summary and Conclusions

We have described the various issues related with the design of video servers and described some architectures that can be used for this VOD application. The scalable architecture that we have described allows a very cost-effective solution for the VOD application and other multimedia applications. It follows our commitment to open standards and reflects our belief that the server, while computer controlled, is not a computer. The VOD application is primarily a data movement applications as compared with a data computation applications, hence neither a scaled down mainframe nor a scaled up workstation provides an optimal cost effective solution.

Acknowledgments

We are thankful to our colleagues at Hewlett Packard Laboratories and the Video Communications Division for the many stimulating discussions. The discussions with Dan Pitt, Vivian Shen, Nalini Venkatasubramanian, Al Kovalick, John Youden, and Paolo Siccado were very useful.

References

- [1] Yee-Hsiang Chang, David Coggins, Daniel Pitt, David Skellern, Manu Thapar, and Chandra Venkatraman, An Open-Systems Approach to Video on Demand, *IEEE Communications Magazine*, May 1994.