Media Pool — Flexible Video Server Design for Television Broadcasting

By Charlie Bernstein

It appears to be an accepted fact that disk-based video storage systems are viable solutions for broadcasters to store and play video in their facilities. One of the more accepted disk-based video products in the broadcast facility is the video server. There are various approaches and technologies that manufacturers can use to build servers. All of these design decisions and trade-offs contribute to various video server architectures to address the task of delivering program content and commercials. The Media Pool video server offers a unique architecture and flexibility that leverages disk bandwidth and system modularity to provide a system capable of multiple full-bandwidth ITU-R 601 streams, system scaleability, and standard industry interfaces — protocols to integrate the system into a facility.

A video server is made up of various subsystems to allow the storage and retrieval of video and audio to and from hard disk in real time. The architecture of each of those subsystems and the interrelationship between them have a major impact on the performance of the system in terms of system responsiveness, bandwidth capabilities, and scaleability. In addition, disk array configurations, system timing, internal data transport, data buffering schemes, and disk access algorithms will affect operation. Providing separate subsystems that focus on specific tasks allows for increased performance and more flexible operation.

The Media Pool video server from Philips Broadcast Television Systems (BTS) Co. includes subsystems for storage, input/output (I/O), video transfer, and control. The storage subsystem is responsible for reading and writing data to and from hard disk drives in real time, disk synchronization, and data buffering. The I/O subsystem provides all of the connections to a broadcast facility, similar to the connections found on a videotape recorder. The video transfer subsystem takes care of bandwidth management from one or more storage arrays to one or more I/Os. The control subsystem allows for system monitoring, database administration, simultaneous multiuser access, and support for integrated applications.

Other variables affect the performance of the system, including system compression techniques, data striping, system bandwidth, and reliability features. All of these topics and other information are presented in this paper, in detail, describing a very effective approach to streaming real-time video from disk-based storage to multiple channels in a broadcast facility.

Video Server Architecture

Most broadcast video servers today have many things in common in their architecture. All servers have major subsystems that handle storage, video transfer, input/output, and control. Low-end systems have taken a more off-the-shelf, computer-oriented approach. These systems contain many limitations due to the nature of the components that are used. In general, they were built for computing and not for addressing the problem of streaming real-time video data from hard disks. Other vendors have taken a more dedicated hardware approach to solving the problem. Dedicated hardware offers specific solutions that are built to solve streaming video issues and provides greater flexibility and higher performance.

When evaluating the requirements of broadcasters and production houses, careful attention must be made to performance factors of the system. How will the system be used? What bandwidth requirements do the facilities have? How many channels will the facility need? All of these factors impact system architecture, design, and performance.

The Philips BTS Media Pool video server combines both off-the-shelf components and dedicated hardware that makes it a very effective solution for broadcasters. It is a scaleable, multichannel system that provides the highest quality video and audio in a facility. Each channel in a Media Pool system is capable of supporting a full-bandwidth ITU-R 601 uncompressed video stream along with up to two AES/EBU 48-kHz pairs of audio. In addition, the server can compress the video stream to leverage economies of scale in storage and bandwidth. The Media Pool is a modular system that can be built in various configurations to meet specific broadcasters' needs. The modular architecture offers increased flexibility in video channels, storage capacity, system bandwidth, integrated applications, simultaneous multiuser access, and control alternatives.

The server architecture uses multiple arrays of hard disk technology to provide high bandwidth and high storage capacity. Each array provides a full bandwidth stream of video. The system supports from one to eight storage arrays per system. The Media Pool also provides modular scaleable I/O channels. A Media Pool system can be configured starting with a single I/O module up to 12 I/Os in a single system. A system purchased today with fewer than 12 channels can be upgraded with additional channels in the future; additional drives and arrays can be added to the system to increase system storage, and additional workstations can be added to run Media Pool's integrated applications in a real multiuser environment.

Media Pool Basics

The Media Pool is divided into four basic subsystems (Fig. 1):

Presented at the 138th SMPTE Technical Conference and World Media Expo (paper no. 138-11), Los Angeles, Calif., October 8 to 12, 1996. Charlie Bernstein is with Philips Broadcast Television Systems (BTS) Co., Salt Lake City, UT 84119. An unedited version of this paper appears in *Film and Video Synergies: Creation to Delivery*, SMPTE, 1996.

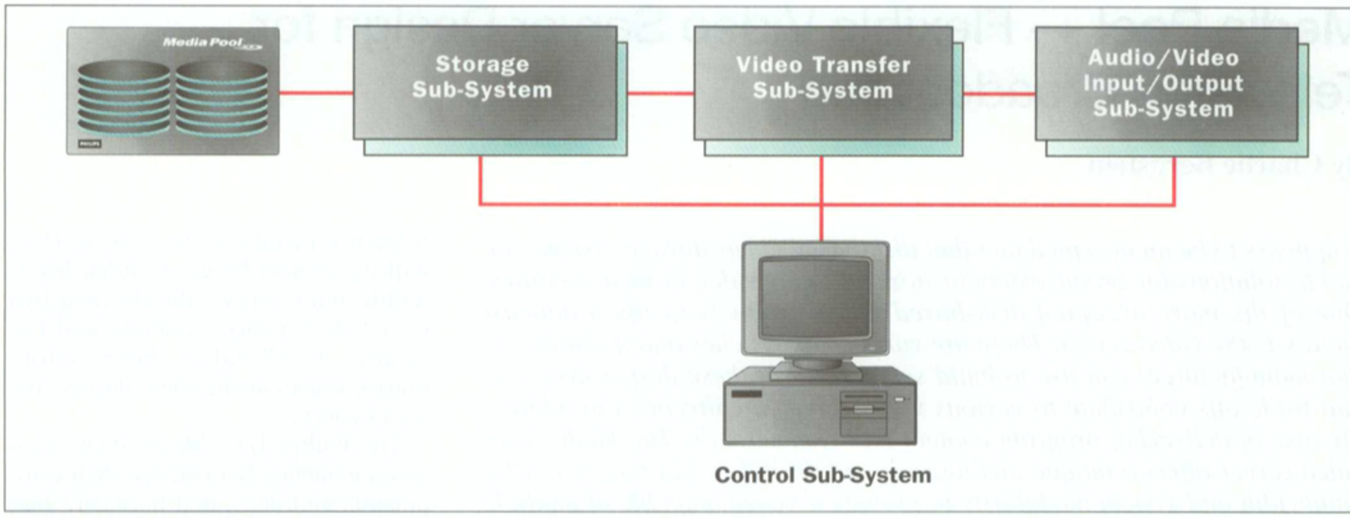

Figure 1. Components of the Media Pool server.

• Storage subsystem — a shared pool of video, audio, and time code information accessible to all channels in the system. Multiple access of shared files is available with disk-based storage systems creating new applications for video storage in the broadcast environment (i.e., multicast feeds, time shifting).

• Input/output subsystem — provides individual access to video files as well as control of streaming video data. It contains interfaces to the broadcast facility similar to a videotape recorder (VTR).

• Video transfer subsystem — provides high-speed fiber channel switching between various storage arrays and I/O modules enabling the system to provide up to full bandwidth video on multiple channels simultaneously.

• Control subsystem — allows for system configuration, media file database tracking, support for integrated software applications, and multiuser capabilities.

Storage Subsystem

Hard disk storage systems have evolved quickly over the past few years. The ability to connect multiple drives as a single addressable storage system has greatly increased the total capacity and sustainable bandwidth of storage via disk. These advances, along with cost reductions in hard disks and random access memory (RAM), have made video server technology viable today.

Multiple drives connected together as a single storage unit is called an array. The Media Pool can be configured with one to eight storage arrays per system. Each array can contain from 10 to 42 drives. This allows for terabytes of online data storage, which translates into hundreds of hours of compressed video on a single server.

RAID Arrays

There are various options to the interconnection of drives. Some vendors use proprietary algorithms and software to connect drives, others use standards-based solutions. One of the most popular and widely used standards-based approaches to connect drives is the use of redundant array of independent drives (RAIDs, also known as redundant

Table 1 — RAID Levels Basic Description

RAID Level	Short Reads	Long Reads	Short Writes	Long Writes	Redundancy	Notes
LEVEL 0	Normal	Normal	Normal	Normal	None	No protection from any drive failures
LEVEL 1	Faster	Faster	Slower	Slower	100%	Very expensive due to doubled disk cost
LEVEL 3	Normal	Normal	Slower	Slightly Slower	Excellent	Cost slightly more than no redundancy. Separate drive used for redundancy
LEVEL 4	Normal	Normal	Slower	Slightly Slower	Excellent	Similar to level 3 with block striping instead of byte
LEVEL 5	Normal	Faster*	Slower	Slower	Excellent	Striping plus data protection is spread among all drives

*Faster when no disk errors occur.

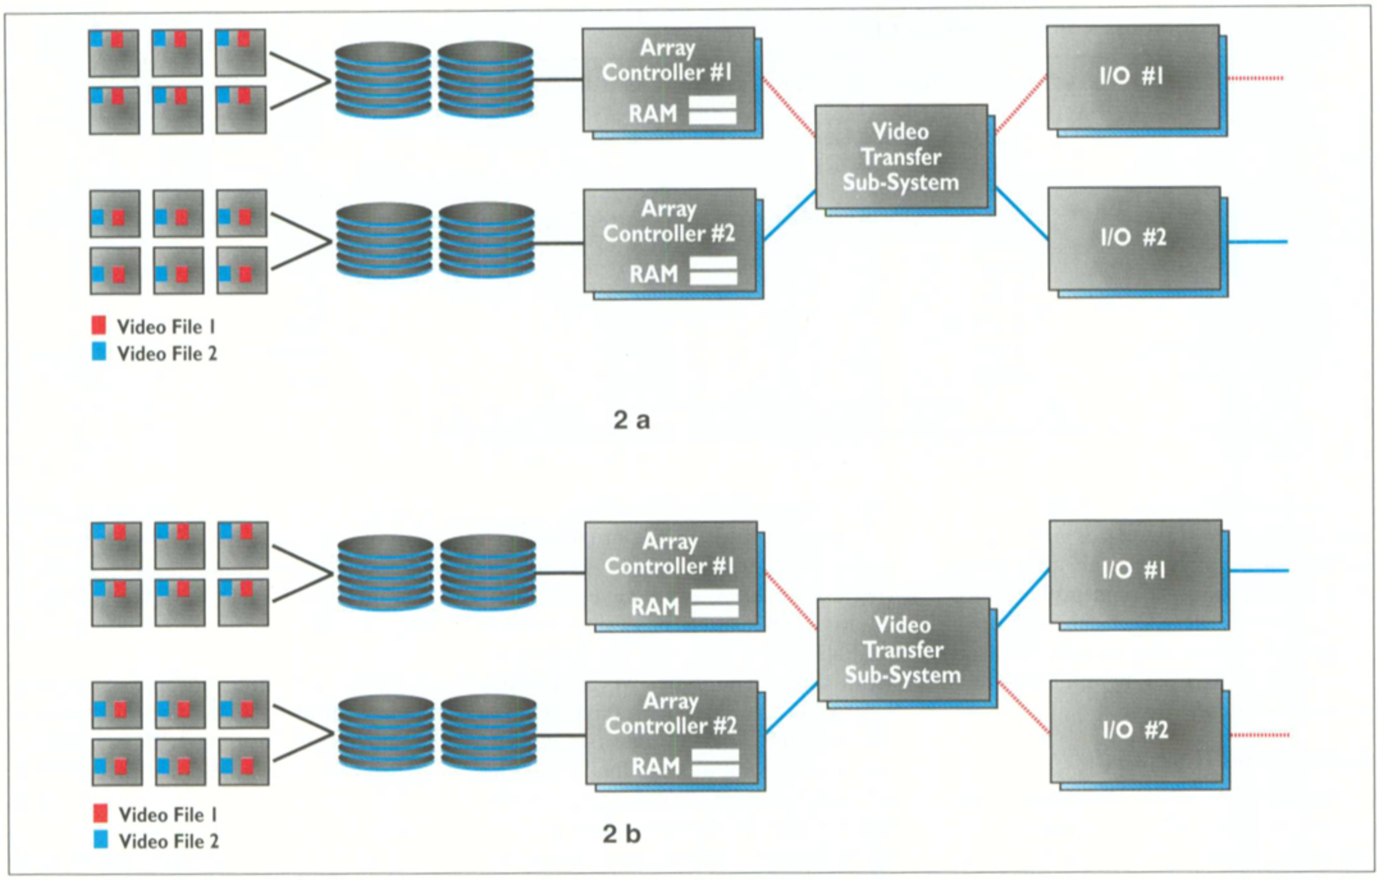

Figure 2. Media Pool data striping.

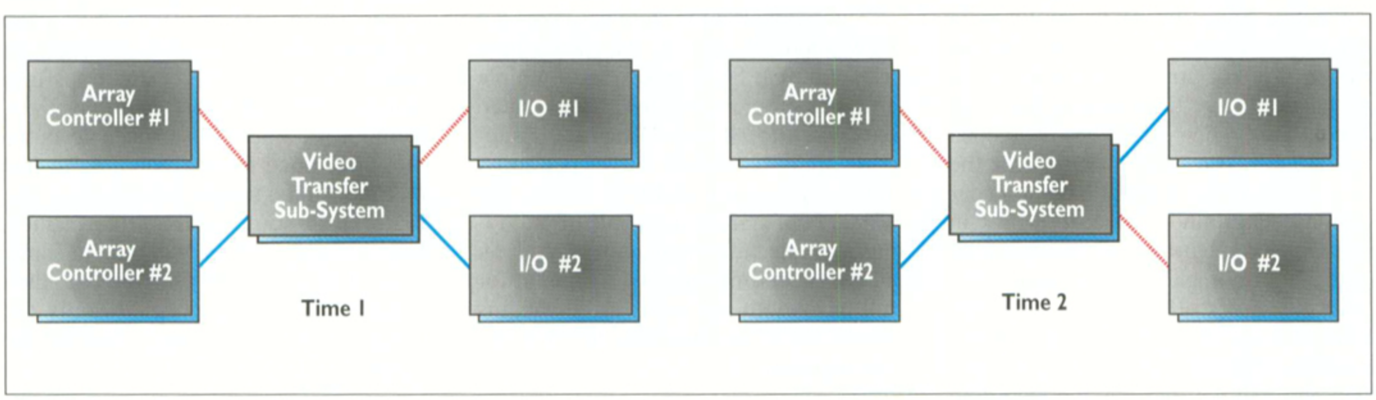

Figure 3. Media Pool connection diagram.

array of inexpensive drives). RAID is the computer industry standard for the connection of multiple drives.

RAID offers various levels of data bandwidth and data protection; the characteristics of each of these levels are summarized in Table 1. RAID level 0 provides for multiple read and writes simultaneously, as all data are spread to all disks. Unfortunately, if any single drive in the system fails, all data on the system are permanently lost. Some manufacturers use RAID 0 in their systems today, but this is not appropriate for real-time, on-air operations. RAID level 1 provides the highest protection, but at a great cost. RAID level 1 mirrors all data of all drives. Since it requires many disks to store a meaningful amount of video, RAID level 1 becomes prohibitively expensive for video server applications.

Along with RAID level 0, RAID level 3 and RAID level 5 are also used by server manufacturers today. Both offer economical storage and protection for video data. RAID level 5 provides a higher level of security at the cost of system bandwidth. RAID level 5 spreads redundancy data over all drives. If a single drive fails, the system will continue to run. Unfortunately, because all data are spread out, it takes longer to access data, thus losing valuable bandwidth for video. RAID level 3 offers better bandwidth capabilities. RAID level 3 can lose a drive and continue operation without bandwidth loss. This ensures that video is uninterrupted.

Table 2 — Compression Comparison Chart

Feature	MPEG-2 MP@ML	MPEG 4:2:2	Motion JPEG
Band width	1.2 to 15 Mbits/sec	up to 50 Mbits/sec	3.5 Mbits/sec to Full Bandwidth
Trim Video	No	Yes	Yes
Compression Type	Discrete Cosine Transform	Discrete Cosine Transform	Discrete Cosine Transform
MPEG Encoders*	Yes	Yes	Yes
Spacial Compression	Yes	Yes	Yes
Temporal Compression	Yes	Yes or No	No
Video Quality	Lowest	Higher	Highest
Storage Capacity	Low	Higher	Highest
Encoder Cost	High	Highest	Low

*This signal can be decoded and passed through an MPEG encoder for uplinks to satellite.

The Media Pool uses RAID level 3, which provides support for parity protection of data. It uses an additional hard disk to store parity information for the disk array. RAID 3 can support a single disk failure with no data loss and provide high data bandwidth even after a disk has failed.

Background Rebuilds

In addition to RAID protection, video servers can also offer automated recovery from failed drives. Background rebuilding of a failed drive ensures that the system will maintain a high level of reliability. Since the server must remain on air to generate income for a broadcaster, it is not practical for a server to be taken off-line to replace a disk. The Media Pool video server has developed system software that provides for the rebuilding of failed disk drives to occur in the background. Since each disk in the server can be removed independently from any other, it is not required to take the server off-line, or unplug any other drives when one fails. The background rebuild is integrated into the system and will not interrupt the recording or playback of video from the server, thus increasing system reliability.

Disk Bandwidth

All disk drives used today are simple computer system interface (SCSI)-based drives. SCSI, which is a standard for the connection of data devices to personal computers (PCs). A single disk drive today can sustain about 3 to 5 Mbytes/sec. A SCSI bus has a maximum data rate per bus of 20 Mbytes/sec. A full 10-bit digital serial component video signal with audio and time code takes 270 Mbits/sec or 33.75 Mbytes/sec.

It takes multiple SCSI hard disks on a SCSI bus to provide up to 20 Mbytes/sec. As you can see, even multiple drives on a single bus cannot sustain the bandwidth required for full bandwidth video. There are two approaches that manufacturers have taken to solve this problem. The first one is compression. Using compression, the 34 Mbyte/sec video data can be reduced to fit into the less than 20 Mbyte/sec data space of the SCSI bus. Second, some manufacturers have used multiple SCSI buses to provide the bandwidth required.

The Media Pool System has taken the latter approach. In fact, the Media Pool uses multiple SCSI buses and multiple arrays to provide and sustain high video bandwidth to all ports. The server supports up to 42 9-Gbyte drives per array. With support for multiple arrays, the Media Pool system is capable of providing multiple full bandwidth streams on all channels simultaneously. Support for full bandwidth and compressed streams allows a user the flexibility to select the video quality they require.

Compression

All video servers use some type of compression to provide the required bandwidth on output or provide more economical use of disk storage space. Lower end systems offer a single fixed compression level or a selection of a few fixed compression levels. Offering a single compression level is very easy for the manufacturer. The system will always use the same amount of disk space per second and have a fixed disk bandwidth requirement on each channel. This may make the system easier to build but does not provide as much flexibility to the user. Low-end variable compression schemes allow for a range of compression usually starting at a high compression ratio and going higher (e.g., MPEG-2 has a maximum data rate of 15 Mbits/sec).

There are three compression standards being used for server development today: Motion JPEG, MPEG-2 MP@ML, and the emerging MPEG 4:2:2 standard. They are described in Table 2.

There are two methods used to compress a video signal: spatial compression and temporal compression. Spatial compression is used to compress each individual frame. All three compression standards use discrete cosine transform (DCT) for spatial compression. Temporal compression, used to compress from one frame to another to further reduce the image, throws away like material in some consecutive frames. For example, if the sky is blue in frame 1 and the sky is blue in frame 2, then using temporal compression, only the pixels in the sky that have changed are stored.

MPEG-2 MP@ML

MPEG compression uses spatial and temporal compression to reduce storage requirements. Spatial compression is accomplished with DCT, as previously mentioned. Temporal compression is accomplished using three types of frames: I frames, P frames, and B frames. I frames are index frames and are similar to those stored in Motion JPEG. They contain the entire frame (compressed). P and B frames use temporal compression and motion estimation to predict what the next frame will be.

Although MPEG requires less space, it creates some limitations. First, the cost of an MPEG encoder today is greater than that of a Motion JPEG encoder. Even though there is cost savings in disk space, additional cost is added to the encoder.

The second issue is video quality. In MPEG, the higher the compression ratio, the farther the distance between I frames. If a scene change occurs between I frames, the quality of the image will degrade until the next I frame appears. This is because you do not have all the information for the whole frame when the scene changes until the next I frame is played.

The third issue affects cutting material in MPEG. Since you may not have all the information at the edit point in the video (not on I frame boundary), it takes more processing and intelligence for a video server to cut two streams together. It is possible to edit two MPEG transport streams together on frame boundaries on output, but additional processing and bandwidth is required to read the previous I frame and all the frames between it and the current frame. This is not a simple feat for some video servers where bandwidth is a premium. The maximum data rate of MPEG-2 MP@ML is 15 Mbits/sec.

MPEG 4:2:2

MPEG 4:2:2 is similar to MPEG-2 MP@ML in that it supports an MPEG transport stream and uses I, B, and P frames. MPEG 4:2:2 differs in that it allows for higher bit rates and supports an I-frame-only mode (no temporal compression), where media can be cut on frame boundaries.

Motion JPEG

Motion JPEG supports full variable compression. Compressing video on a frame-by-frame basis, it allows the system to support limited editing capabilities not currently supported with MPEG-2 MP@ML compression. Video recorded onto the server can now be trimmed and manipulated on a frame basis without concern for loss of video quality or the requirement of additional bandwidth or processing.

The Media Pool video server uses Motion JPEG compression and supports variable compression from full bandwidth video up to 50:1 compression selectable by the user on a frame-by-frame basis. It uses a separate mezzanine board to support compression requirements of the system. The compression mezzanine card can be replaced in the future to support other compression algorithms, like the emerging MPEG 4:2:2 standard being established.

Motion JPEG and MPEG Compatibility

Since all of these compression algorithms use DCT for spatial compression of each image, interchanging between MPEG and JPEG is not a concern. A Motion JPEG server can be used to feed an MPEG encoder for uplinks to satellites without concern for artifacts. The two compression algorithms are compatible.

Today, Motion JPEG offers the most flexibility for video server users. MPEG-based solutions are viable for video servers offering higher compression with some limitations.

Scaleability

The key to the Media Pool's storage subsystem is its scaleability. Multiple drives and multiple arrays allow for storage and bandwidth management. Using compression, storage capacity can be increased and bandwidth shared. Flexible storage capacity and flexible system bandwidth are the result of the scaleable architecture of the storage subsystem.

Input/Output Subsystem

There are a range of choices developers can make to interface to a server. Using an off-the-shelf architecture, video server manufacturers require PC-based video cards. These cards are limited in physical space and data bandwidth due to the limitations of PC industry standards and the low data rates the buses support (ISA about 8 Mbytes/sec). Higher end systems have developed system specific I/O subsystems that perform the task of playing and recording video into the server. The Media Pool video server has dedicated hardware developed for this task. The server's I/O subsystem provides serial digital component video output on all channels. Since all video information is stored in the server digitally, and the industry is migrating to the digital domain, it makes sense to provide a digital signal to a broadcast facility from a server. It ensures that the best quality video and audio will be provided from the server, and it provides support for the digital future.

The Media Pool I/O, known as the VR-7000, contains all control and I/Os required for a broadcast facility. The VR-7000 supports ITU-R 601 serial digital video input and output. There are two selectable video inputs and a single video output with dual connections for support of electronic-to-electronic (E-E) mode operations during recording. The VR-7000 contains two AES/EBU (48-kHz, 24-bit) digital audio stereo pair inputs and outputs per I/O. The I/O supports longitudinal time code (LTC), and general purpose inputs and outputs for interfacing to a master control switcher or other trigger devices. The VR-7000 contains a fast/wide SCSI interface for data transfer to storage libraries. The board also supports RS-422 ports for control via the beta control protocol and the video disk control protocol for automation system support. Local area network (LAN)-based control is also provided.

Each I/O houses a separate audio mezzanine and compression mezzanine. Either board can be replaced to support other audio or compression standards. A Motion JPEG-based I/O could be converted to an MPEG-based output with a board swap. The Media Pool can support from 1 to 12 I/O subsystems, or channels, in a single server.

The modular architecture of the I/O subsystem allows a system installed today to be upgraded in the future with additional channels or new compression algorithms. This provides customers with flexibility and investment protection not offered in other systems.

Video Transfer Subsystem

The video transfer subsystem can take on various forms. At the low end, standard PC-based buses can be used to distribute video signals to multiple I/O ports. At the high end, high-speed links can be switched to create multiple channels of I/O. The use of separate high-speed links to transfer data from storage arrays to the I/O modules is preferred. It provides the system with the horsepower needed to serve high-bandwidth applications or can be used to cut the bandwidth into smaller units for distribution to many I/Os.

In the Media Pool, this subsystem is comprised of a gigabit matrix switch, called the data transport commutator. Signals between the array subsystem and I/O subsystem are switched very rapidly over a matrix fiber channel switch to allow the system to provide up to full bandwidth digital serial component video for each I/O simultaneously. The timing of the switch is regulated by system software in the I/O subsystem and the control subsystem, allowing for flexible bandwidth allocation and support for multiple data formats. The ability to provide this bandwidth allows the Media Pool various architectural options. These include full bandwidth output for high-quality video applications, system modularity, bandwidth allocation management, and variable compression.

Data Striping

To gain the benefits of multiple disk drives and multiple arrays to provide maximum bandwidth to all I/Os in the system, data striping is used. Striping of data is performed by the data transport commutator and array electronics and allows video from a single file to be placed on all disks of the system (Fig. 2). This is an important feature of the server, because it allows each I/O to be connected to a different array at any time and still get simultaneous access to the file it is playing.

In Fig. 2a, I/O 1 is reading video file 1 from the drives in array 1 while I/O 2 is reading video file 2 from array 2. In Fig. 2b, the video transfer subsystem switches the connection and I/O 1 reads the next block of video file 1 from array 2 and I/O 2 reads the next block of video file 2 from array 1. All video is read into a RAM buffer so that when the I/O is ready to get video from the array, the video is ready to be transferred. Figure 3 depicts the connections between I/O subsystems and array subsystems and the switching concept.

Simultaneous Access of Media Files

The Media Pool's RAM buffer allows for simultaneous access of video from the same file at the same time. Blocks can be read into the buffer separately in the background, and once the buffer is full, you have multiple virtual copies of the file. By adding more arrays and more I/Os and allowing switching to occur very rapidly (gigabit rates), you can see that each I/O can get access to any file from any and all arrays simultaneously.

Full Bandwidth Video and Bandwidth Allocation

The Media Pool arrays are capable of providing full bandwidth digital serial component video with two AES/EBU stereo pairs of audio, LTC time code track, and full preservation of the vertical interval on each I/O. This bandwidth can be broken up into smaller segments to allow for increased system modularity. By providing bandwidth allocation management software, and variable compression, the storage system can now be leveraged by multiple I/O channels.

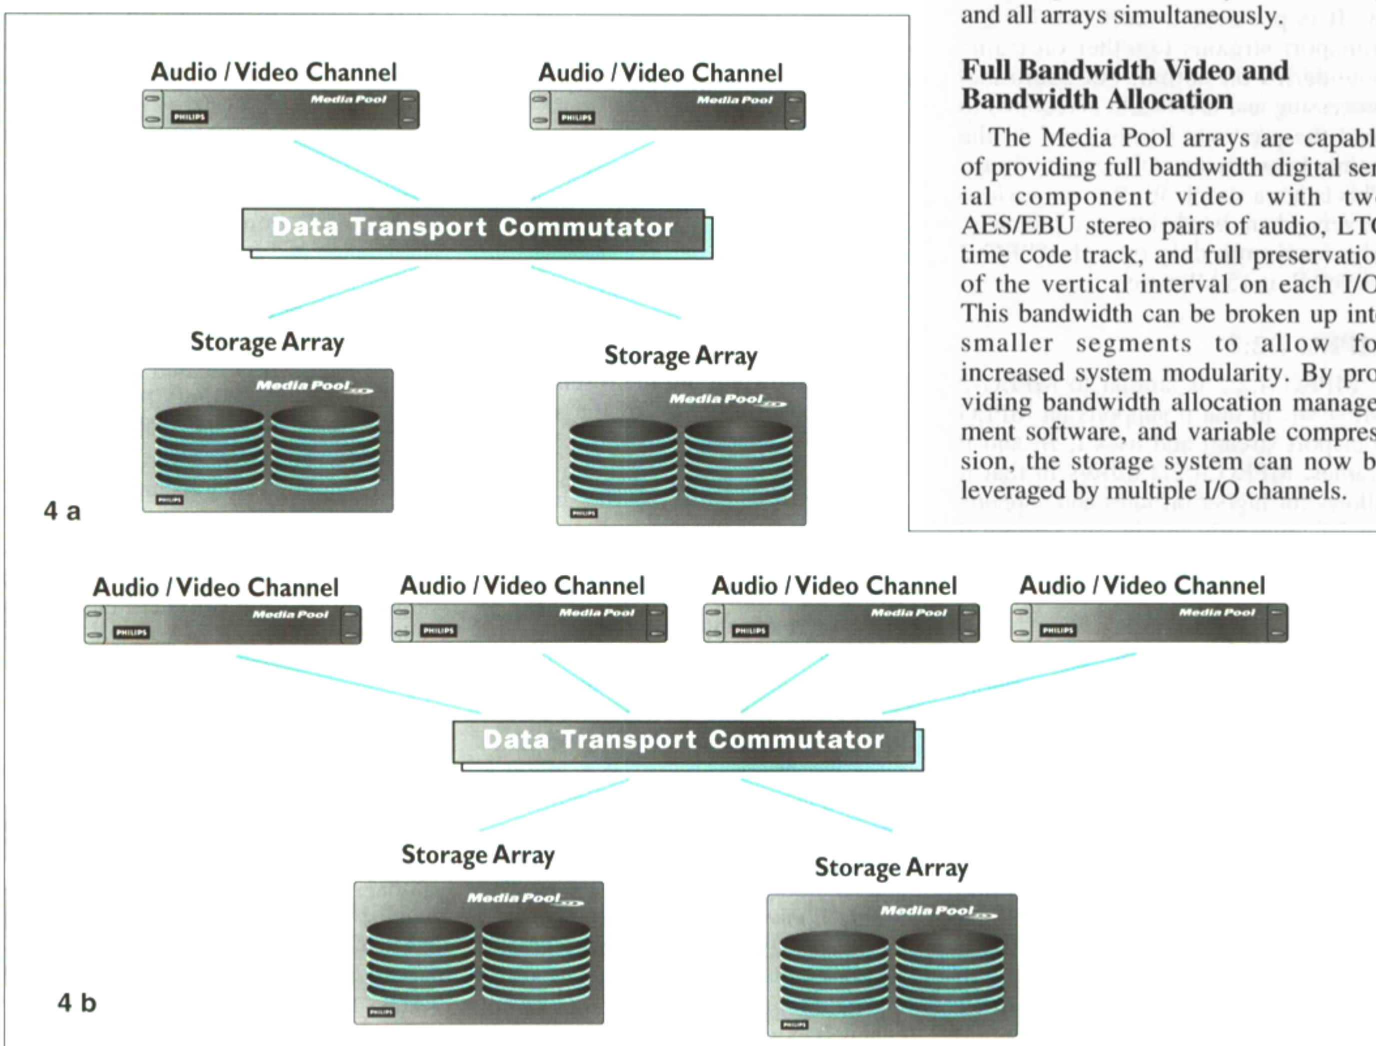

Figure 4. Media Pool configuration and bandwidth management.

Table 3 — MPS Series Media Pool Configurations

Series	Model No.	Channels*	Arrays*	Workstations	Drives	Min Storage	Storage @ 8:1
MPS-2100	2110, 2112	2	1	1	10 to 42	57 min	28 hr
MPS-3100	3110, 3112	3	1	2	10 to 42	57 min	28 hr
MPS-4100	4110, 4112	4	1	2	10 to 42	57 min	28 hr
MPS-4200	4220, 4224	4	2	2	20 to 84†	114 min	56 hr†
MPS-5200	5220, 5224	5	2	3	20 to 84†	114 min	56 hr†
MPS-6200	6220, 6224	6	2	3	20 to 84†	114 min	56 hr†
MPS-6300	6330, 6336	6	3	3	30 to 126	172 min	84 hr
MPS-7300	7330, 7336	7	3	4	30 to 126	172 min	84 hr
MPS-8300	8330, 8336	8	3	4	30 to 126	172 min	84 hr
MPS-8400	8440, 8448	8	4	4	40 to 168	229 min	112 hr

*Note: Larger Media Pool systems are available. †Note: Some drive configurations may not be supported by all systems.

For example, if a user needs two full bandwidth channels, the Media Pool server can be configured as shown in Fig. 4a. The two arrays can support two full bandwidth channels as configured. This is referred to as a symmetrical system. In Fig. 4b, you can see that the system can also be configured with more I/Os than arrays; this is known as an asymmetrical configuration. In this case, the total bandwidth of the combined arrays is divided among the four channels. The full bandwidth signal is compressed (in this case 2:1) and striped across both arrays. The data transport commutator provides the appropriate bandwidth allocation and switching for all four channels.

The Media Pool system is very modular and can be configured in many combinations of channels and arrays. This modular approach allows broadcasters to build a system that meets their needs and also allows for expansion of the system in the future.

Media Pool Configurations

The Media Pool is capable of 74 different configurations (up to 12 channels and 8 arrays per system). Almost any combination of 1 to 8 arrays can be combined with 1 to 12 channels with some minor limitations. Additional configurations will be added in the future. Each array can be connected to up to four I/O modules in a system. This is referred to as a four-to-one fan out. This allows a facility to leverage a few storage arrays with multiple channels when compression is used.

MPS Series Media Pool Systems

Although the Media Pool video server can be configured in any of 74 configurations, there are a smaller number of configurations that most facilities require today. As a result, Philips has created the MPS Series Media Pool Systems. Each model in the MPS series is preconfigured and includes all I/Os, arrays, drives, cables, and RAM buffers a customer needs to integrate a video server into their facility. There are 10 MPS Series families and 20 models available. Table 3 depicts the configurations available in the MPS series.

Control Subsystem

The control subsystem in the Media Pool consists of one or more workstations running X Windows. UNIX was chosen because it is one of the most reliable operating systems available today. The workstation was chosen so that the controller would have enough horsepower to drive multiple channels concurrently. The Media Pool system controller is responsible for initiating system booting, system configurations, compression level setting, system monitoring, media file database administration, and for running the server's integrated applications. The system controller, which communicates to the other subsystems in the server via an internal communications network, contains software applications for system administration, commercial insertion, and record/playout. Additional workstations can be added to allow for simultaneous multiuser access.

Multiuser Capabilities

The Media Pool supports multiple workstations in a single system. Each workstation can run one or more applications or protocols allowing multiple users to operate and use the Media Pool system simultaneously. This is a very important distinction, because the operations within a broadcast facility, although fairly basic, are very repetitive and time consuming. By allowing for multiple simultaneous access to each channel of the system, broadcasters can increase operational efficiencies and reduce operational costs; now there is less idle time as conflict for resources is reduced.

Media Pool Integrated Applications

There are currently four integrated applications and three integrated protocols available for the server. The applications provide integrated functionality that takes advantage of tightly coupled control and high-end features. The protocols provide the server with an open-system approach, so the system can be integrated into the facility with other devices and also to allow for custom application development by broadcasters and other vendors. The applications are software package for logical administration of system hardware (SPLASH), system administration, stream, VTR emulation, DiskCart cart emulation, and DiskCache broadcast facility. The protocols supported are BVW-75 serial control, videodisk communications, and PoolNet-native Media Pool.

SPLASH

The main interface to system administration and system configuration is a software program called SPLASH. SPLASH allows a systems administrator to monitor system activity, initialize and boot the system, and configure the system and its integrated applications. SPLASH allows a user to create Media Pool "instance" files. An instance file in the Media Pool server is a configuration file for a specific application or protocol and can be assigned to a specific operator. Multiple configuration files or instances can be created for each I/O, but of course only one instance can be run per I/O at a time. For example, if a broadcaster has a four-channel Media Pool with two arrays, two of the channels can be configured as instances for dubbing spots and another two for commercial spot insertion. The instance scheme described as follows allows for very flexible configuration management and usage of the pool.

STREAM Record/Playback Software Application

The STREAM record/playback application emulates the transport functions of a VTR. Stream runs on the native Media Pool workstation and provides the ability to record, playback, stop, pause, cue, shuttle, set variable compression, and create spot files on a Media Pool system. E-E mode is supported as well as external general purpose interface (GPI) triggers. STREAM supports the ability of clicking in the timeline to go to specific frames in a video file and also supports setting cue points for easy navigation through the media file or spot. STREAM has the ability to remotely control a tape machine for easy loading of material.

DiskCart Cart Emulation Software Application

The DiskCart spot playback application provides a complete integrated cart emulation application for the video server. DiskCart runs on a native Media Pool workstation and provides the capability to create or import playlists, control event playout, and monitor the system.

DiskCart can import a playlist from external traffic systems or create one in the playlist editor. The editor supports drag and drop capabilities to add or delete events. The application supports multiple event types including absolute timed events, externally triggered events, or follow events. External events can be controlled via a GPI from master control or an automation system. DiskCart supports real-time, on-air playlist editing and the ability to nest playlists, insert break lines, comment lines, or delete events. All events are moved to the as-played log for reconciliation later.

DiskCache Broadcast Facility Application

The DiskCache broadcast facility application offers seamless integration of the Media Pool video server into an existing broadcast facility. DiskCache provides the Media Pool with cart machine functionality along with integrated control of other devices. DiskCache offers automatic control of dubbing material into the Media Pool server from multiple devices and can control the Media Pool and other devices for recording and playback for air. DiskCache can also control routers and switchers or master control.

External Protocol Support

The Media Pool also supports a number of control protocols for external control of the system. BetaCP is an RS-422 Beta protocol emulation. Using BetaCP, a channel of Media Pool can be controlled like a VTR. The videodisk communications protocol (VDCP) was developed for interfacing to an automation system. Broadcast facilities can control a Media Pool via the Louth or Alamar automation systems by using the VDCP protocol. Finally, PoolNet is the native control protocol of the Media Pool. PoolNet can be used for custom applications and integration of third-party devices. PoolNet has been used by Philips for development of integrated applications. DiskCache uses PoolNet to control the Media Pool.

Time Delay

The Media Pool is very capable of supporting time delay or time shifting applications. Currently using an automation system, the Media Pool can record satellite feeds into a broadcast facility and, via automation, can be scheduled for playback at a later time. After the video has played, the automation system can delete the files from the pool. In the future, additional applications and functionality will allow the creation of third-party applications and integrated applications to allow for new time delay features, such as destructive playout.

Archiving Solution

Using the Media Pool's I/O subsystem SCSI data port, new applications have been developed to connect medium and long-term storage devices to the server. Otherwise known as near-line storage, files can be moved to an archive device from the server and brought back when needed. Video is transferred as compressed data, which offers a number of advantages. First, since the data are compressed going out and coming back into the server, there is zero data loss (also known as generation loss). Also, when data are compressed, it is possible to transfer information at greater than real-time speeds. Since the Media Pool can begin playing while a file is still being recorded onto the pool or transferred from an archive, it is possible to start playing a file stored on the archive before it is completely transferred. This allows for flexible on-air operation, when a file is needed for a last minute schedule change.

Conclusion

The Media Pool video server architecture uses independent subsystems, variable compression, and integrated applications to provide a flexible solution for storing and serving professional video. The system offers flexible configurations, system scaleability, and support for multiple full bandwidth video streams. A Media Pool system configured today can be upgraded in the future with additional storage, bandwidth, channels, and users.

THE AUTHOR

Charlie Bernstein is currently the director of product management for the video server line of business for Philips Broadcast Television Systems Co. in Salt Lake City, responsible for the Media Pool video server product line. His background is in computer science, business, and marketing. He has been involved in product development and marketing for the past 14 years, and for the past 10 years he has been dealing with video products.