Multichannel Video Server Applications in TV Broadcasting and Post-Production

By Mark Ostlund

As we move into the age of digital television signals, new storage technologies will play a dominant role. Video servers based on these new technologies will provide multiple access to a shared pool of television images and sounds, revolutionizing the way television programs are produced and distributed. While providing increased functionality, these servers will greatly simplify operations and allow a reduction in costs. This paper examines the requirements for such a server and presents a vision of the future television production and distribution environment.

As the next-generation digital videotape format wars were being played out in recent years, the continued acceleration of technology has delivered professional recording alternatives that render the outcome of the format wars nearly inconsequential. The advent of multichannel video servers in 1995 will change the nature of the way we store and work with video and audio. Offering improved reliability, instantaneous access, more flexibility, and greater operational efficiencies, a well-designed video server can offer significant advantages over traditional videotape.

Multichannel Video Server Defined

What exactly is a video server? A server is a term that originated with the client/server model of the computer industry. An example would be a data-base server that stores and serves information to many clients or users simultaneously. Thus, a video server provides video and audio to many channels or users simultaneously. The disk resource is shared to allow multiple users to access and update the same data, and also because disk storage is frequently the most expensive part of the system. This shared video storage is analogous to a file server in a networked computing environment. Each client uses only as much space as it requires on the server. Some applications (e.g., data bases) require that the recorded material be available for reading and writing by many clients simultaneously.

Video servers are only now becoming available, due to the high data rates required. A video server does not have the luxury of having one client wait while it services another, as database servers do — each video stream cannot be interrupted. This means that the bandwidth into the storage pool must be a multiple of the bandwidths of all of the input/output (I/O) channels on the system, thus driving the server throughput into the gigabyte-per-second range for very large servers. A computer data-base server has much lower bandwidth requirements, delivering only several hundred or perhaps several thousand bytes of data to each user for every request.

This last difference is an indication as to why it is difficult to turn a general-purpose computer into a video server. Requirements for serving data-base applications are fundamentally different from those for serving video streams. Therefore, general-purpose computers are designed for a completely different problem domain. Video servers must reserve overhead in bus bandwidth, memory caching, and central processing unit (CPU) power to enable them to deliver uninterrupted streams of video data reliably. Computer buses are inherently limited in bandwidth and thus scalability. In streaming applications like video, even a high-speed bus becomes a severe bottleneck when shared with several processors, disk controllers, and I/O controllers. Therefore, a server designed specifically for streams of video and audio would be likely to have significant architectural differences and be much more efficient.

Requirements for a Professional Video Server

To provide a superior recording platform, a professional video server should include:

- Scalability in storage depth (from several minutes to hundreds of hours), storage bandwidth, and number of video and audio I/O channels. The system should be modular, so that it can grow in such increments that channels and bandwidth scale together. Storage time should scale by adding drives to existing channels.
- Random access to all audio and video recorded on the server, which should provide complete interactive response to all users. The server should be able to replace random access still stores and audio recorders.
- Simultaneous access to as many channels as are required in a facility, allowing the storage pool to be a shared resource to all users of the system. This should be possible without the need to replicate clips by copying or dubbing.
- Choice of quality. The server should not be limited to a particular level of compression. It should be capable of full bandwidth when the highest quality is desired, or of using compression to increase storage.
- High reliability. To ensure that operations can continue uninterrupted when revenues depend on it the server should allow for redundancy with no single point of failure. It should be capable of being maintained while it is providing continuous service.

Presented at the 136th SMPTE Technical Conference in Los Angeles (paper no. 136-15) on October 13, 1994. Mark Ostlund is with BTS - Broadcast Television Systems, Inc., Salt Lake City, UT 84130.

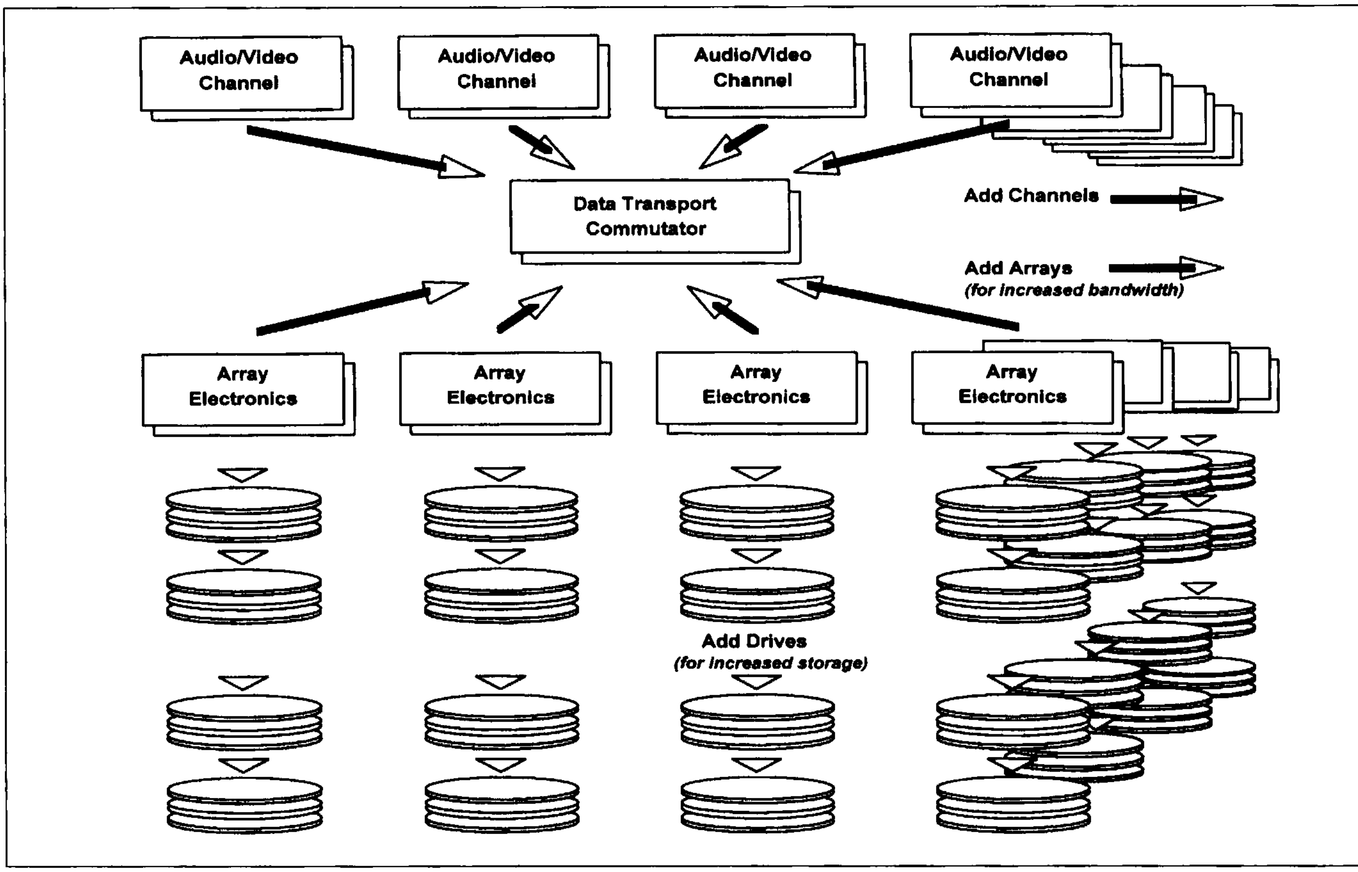

Figure 1. Symmetrical system configuration.

• Compatibility with existing studio equipment. To preserve existing investments in studio equipment a server should provide VTR emulation along with all of the required signal and control interfaces to plug into an existing studio environment. VTR emulation includes audio, insertion edits, and high-quality variable play.

• Nonlinear editing. Clips of video and audio should be playable back to back without regard for internal disk fragmentation, seek times, or whether or not the segment is compressed.

• Archive interfaces. In addition to archiving to standard videotape, a server should be capable of interfacing to computer archive devices, e.g., data tape, robotic tape libraries, and optical disks. If the server is capable of compression, it is desirable to archive in a compressed format to minimize tape usage, download time, and to eliminate generation losses from decompressing and recompressing.

• Cost. The cost of a video and audio server channel, inclusive of disk bandwidth and storage, should be less than the cost of an equivalent VTR. The "cost of ownership" and reduced operational costs should be factored into the cost equation.

• Future proof. Standard interfaces and components should be used that provide compatibility with applications not yet written and storage devices not yet developed. System expansion should allow drives from future generations to work with drives purchased today.

Architecture for a Scalable Multichannel Video Server

The BTS Media Pool multichannel server represents the convergence of high-speed computer technologies. While providing improved reliability, instantaneous access, and more flexibility than traditional tape storage, the server was designed to simplify operations, reduce costs, and provide common access to a shared pool of images and sounds.

Compression is an option. Each video segment can be full-bandwidth, 8-, 10-bit, or any level of compression. This provides a choice of quality for each video segment that is recorded. These segments can then be randomly accessed and played out in any sequence, with Media Pool technology detecting compression levels automatically.

Video distribution and archiving until now has always utilized traditional video equipment. Compression has presented a dilemma. In order to distribute or archive compressed video data using traditional means, it must be decompressed on output and is likely to be recompressed on input, which causes a new form of generation loss. The Media Pool server includes computer network and small computer systems interfaces (SCSI) to distribute and archive compressed video as data, avoiding generation losses and taking advantage of standard computer networks and tape transports.

The Media Pool system combines modules, thus allowing a high degree of scalability. Video I/O modules and disk arrays are added to the system in pairs to expand the number of simultaneous channels in a full bandwidth system. These progress to the right, as shown in Fig. 1. Disk modules are added to arrays to scale the amount of storage from 40 min up to 80 hr. If compression is being used, this could equate to over 1,000 hr.

This description provides a symmetrical system configuration in which all channels can simultaneously access uncompressed video. To save costs in systems where compression is used extensively, the system can also be configured asymmetrically with more I/O modules than arrays (Fig. 2). The I/O modules negotiate for bandwidth, which determines the minimum compression rate that they are required to use. In such asymmetrical configurations some modules may still run at full bandwidth, while others may be limited to a certain minimum compression rate. The general rule requires that the bandwidth of the I/O modules must not exceed the bandwidth of the supporting arrays.

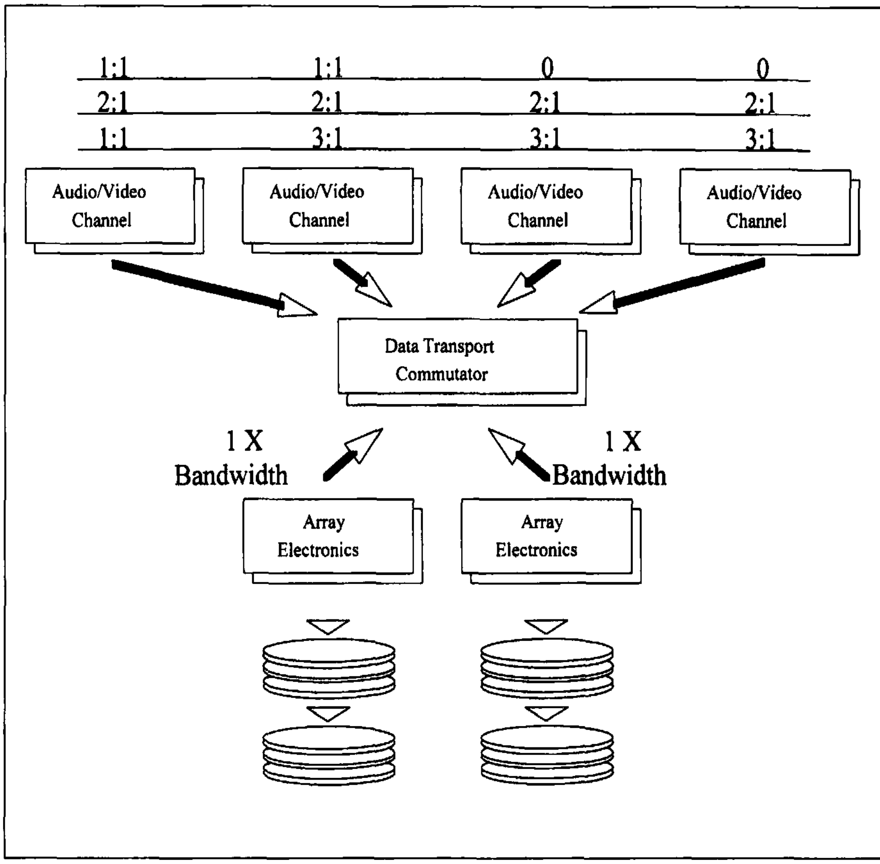

Figure 2. Asymmetrical system configuration.

The Key to the Media Pool Server — The Data Transport Commutator

The data transport commutator (DTC), shown in Fig. 2, is a scalable interconnect that makes all of the arrays in a system appear as a single pool of video and audio to each channel. Only a single copy of any video or audio segment is needed, and all channels can have simultaneous access to it without restriction. This interconnect is a scalable, high-speed switching fabric that makes it possible to share data and to balance the bandwidth load of the system. The DTC is the fundamental system component that allows scaling to dimensions not previously achieved in a server architecture. With no inherent bandwidth limitations, it is capable of cost-effective, modular expansion to aggregate bandwidths of many gigabytes per second.

RAID Technology

Fast SCSI-2 disk drives typically contain internal error detection and correction systems that reduce the "unrecoverable data error rate" to one error per 10^{13} bits transferred, or better. The major concern in a disk array is therefore not so much the random error rate, but is instead the possibility of a total failure of one of the drives in the array.

To correct for drive failures, schemes involving redundant arrays of independent drives (RAID) have evolved. Using a data redundancy scheme similar to what is known in the computer industry as "RAID Level 3," BTS servers are able to continue operation after a failure of any single drive with no effect on system performance. The failed drive is instantly identified, and its bad data is "corrected" on the fly using redundant data stored on a parity drive. At the same time, a background process is begun to reconstruct the data that was on the failed drive onto a spare drive that is part of the array. Once the spare drive has its data rebuilt, the system is once again prepared to withstand another single-drive failure with no loss of data or operability. The failed drive is then marked for replacement.

In large disk arrays the drives are divided into groups, each group having its own private parity drive and optional spare drive. Using this technique, it is possible to have a very large mean time between failure (MTBF) for the disk system as a whole. A system failure only occurs if there is a double-drive failure within the same group of drives, because a single-drive failure is error-corrected as described above.

For example, consider a system containing 40 drives, each with an MTBF of 500,000 hr. Assuming a parity group size of 10 drives, and that a spare drive is rebuilt in less than 1 hr, the probability of a system failure due to double-drive failure over a five-year operating period is about 0.0008%. Properly implemented RAID technology renders system failures due to drive failures extremely unlikely.

Although modern SCSI drives are rated with an MTBF of 500,000 hr or more, the above discussion only applies during the "life" of the drives, which is likely to be between seven and ten years. Normal drive life ends when the spindle or actuator bearings fail. This implies that all of the drives in a system should be replaced after

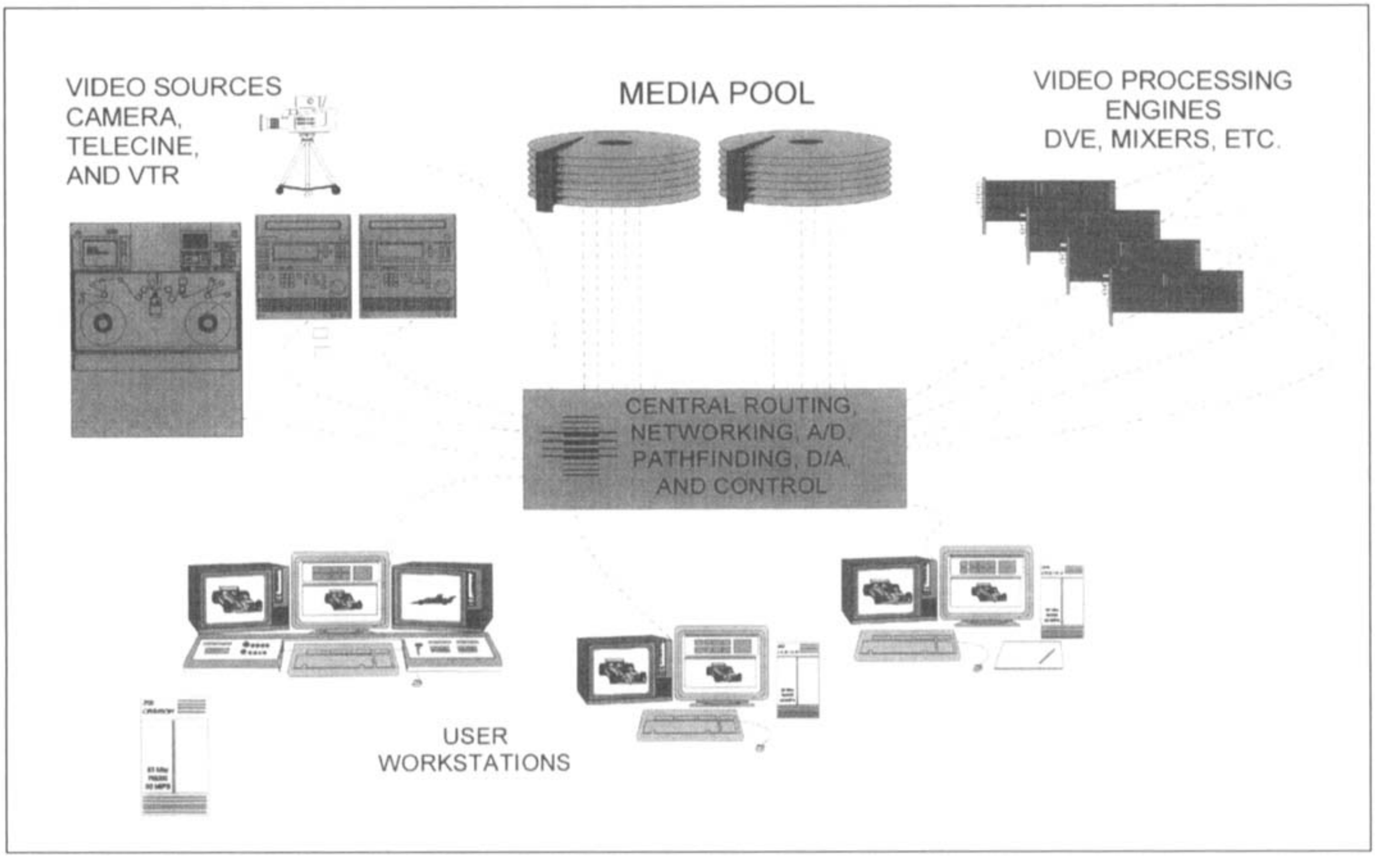

Figure 3. Post-production diagram.

seven years of use to maintain a high degree of reliability. The current price/performance curve of drives yields a doubling of storage capacity at a constant price every 14 months. If this curve holds for the next seven years, this would mean that the drives in a system could be replaced in seven years for 2.5% of the original purchase price, or that 40 times the storage could be purchased at the same price.

Reliability

For those applications where redundant systems are deployed to guard against going off-air or losing revenue, the Media Pool server provides for complete redundancy without requiring a complete backup system. Given the reliability of the drive system using RAID technology, it is not desirable to have a redundant mirrored set of drives, as the drives can be the most expensive part of the server. Instead, the video server augments RAID by providing redundant and hot-swappable electronics, power supplies, and fans, so that if any server component fails, it can continue to operate and even be repaired while remaining in operation.

An Open Server Platform

Industry-standard protocols allow the Media Pool server to work with existing systems. In addition, the full capabilities of the server will be opened up to developers by providing protocols that export its complete functionality. This server protocol is a superset of the SMPTE VTR Dialect, which will provide a standard way for applications to control multichannel video servers.

Archival Storage

Those applications needing large amounts of storage will best be served by a combination of disk and tape. The main question to ask is whether the tape portion of the storage should be videotape or data tape. To avoid generation losses from moving video to and from the server in systems where compression is being used, data tape is the better choice. Compressed images are best archived in the compressed data format to avoid multiple decode and encode stages. Since compression reduces both data and the required bandwidth, it allows transfer rates many times faster than real time.

Post-Production Applications

Nonlinear Editing

Full-bandwidth multichannel video servers represent the enabling technology to bring nonlinear editing into the on-line suite, as shown in Fig. 3. As many channels as are required for layering or effects can be allocated to an edit suite, and all operate in real time and at full bandwidth. Multiple-source reels can be brought into the server simultaneously in the background without interrupting edits in progress. When an edit is finished, it can be dubbed to tape while editing begins on the next job. Background loading and dubbing maximizes editing time. By providing a pool of channels and storage media that can be dynamically assigned to edits as needed, the server

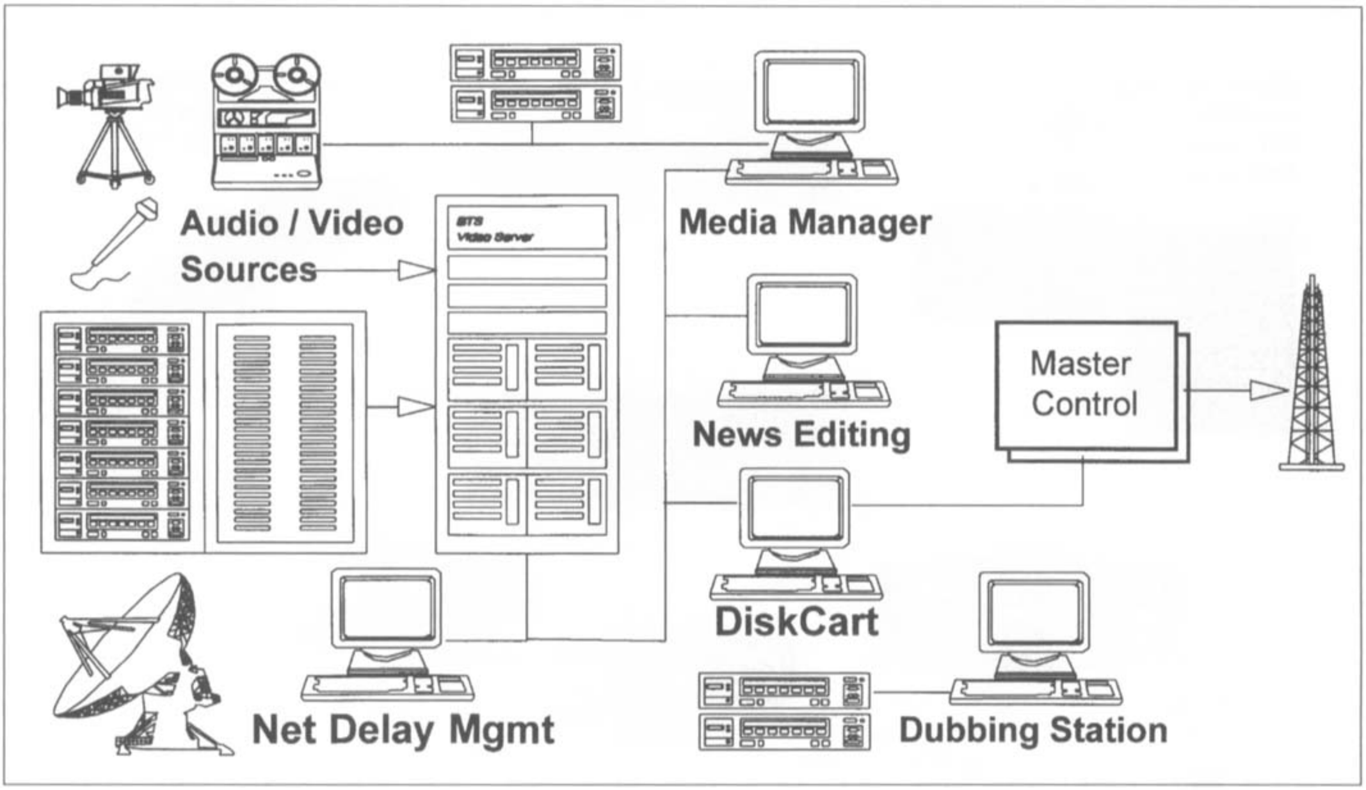

Figure 4. Broadcast applications diagram.

resource becomes much more cost-effective.

This is not the only benefit servers offer for editing. With most edit decisions still being done off-line, there remains a time-consuming and expensive edit decision list (EDL) conformance session in an on-line suite. As a multichannel video server, the Media Pool will be able to dramatically speed up this process by loading multiple source reels simultaneously. Conformance edits will be faster than real time and provide the ability for changes and alternate cuts without doing separate linear edits.

VTR Emulation

This is an application that provides full emulation of industry-standard VTR control protocols and a graphical user interface (GUI) to match the type and functionality of a normal VTR, allowing the server to drop into virtually any type of post-production environment where a VTR would be used.

Disk Mastering

Disk mastering is a use of the server that simply uses VTR emulation just described. It allows editors to master to disk to facilitate making changes later. Rather than repeating a linear process they can make changes (e.g., trims or insertions) instantaneously.

Broadcast Applications

Commercial Insertion

DiskCart is an application that allows the Media Pool server to complement or replace a cart machine (Fig. 4). In an entry-level configuration DiskCart buffers a cart machine, saving wear and tear on VTRs and robotics, eliminating scheduling bottlenecks, and simplifying the generation of play lists and spot reels. In a larger configuration the server will be capable of replacing a cart machine. It has the advantages of reducing labor, more reliability, and more functionality and flexibility. In this configuration it may still interface to a cart machine or robotic data tape system for archival purposes. DiskCart provides standard interfaces to allow it to interface to traffic, automation, and master control systems. As video servers become ubiquitous and automation becomes capable of remote management, it may become possible to deliver a "broadcast station in a box" completely unattended, FM-radio style.

As many broadcasters move toward multiple channels for cable or regional commercial insertion, a multichannel video server provides an elegant solution for inserting on a number of channels simultaneously.

Program and Network Delay

The Media Pool server can effectively delay any source. Incoming video, such as network feeds from satellite, can be shifted to later times easily, as media is buffered on disk until needed. Currently this task is accomplished by a bank of three or more VTRs, with a full-time operator.

Conclusion

A well-designed multichannel video server can provide a powerful alternative to tape. Providing unprecedented reliability, simultaneous multiple channels, random access, choice of quality, scalable storage, and compatibility, video servers can dramatically improve the way we store and work with video and audio media.