High Speed Data Recording: Digital VTRs Find New Applications

By John Krooss, Philip Livingston, and Stephen Mahrer

Entertainment production, be it cinema, television, or even "game" software, is turning increasingly to computer-generated or manipulated information. Computer processing and digital manipulation has permitted the creation of images otherwise impossible, the production of effects not otherwise practical, and correction of innumerable image defects large and small. Given this combination of voluminous data and expensive hardware, the industry faced a dilemma: how to "move" jobs in and out of a facility quickly so the equipment could be made available to other clients and projects. Existing small format data recorders generally have proven unsatisfactory for these purposes due to the limited data rates and media size. Current digital videotape recorders are now capable of data rates well in excess of 200 Mbits/sec (25 Mbytes/sec). Unfortunately, some recorders incorporate data compression, rendering them suitable only for traditional video. Other video formats have been customized and repurposed for data recording, but these machines are no longer applicable for video recording. As the boundaries between the computer and video worlds blur, users need multipurpose solutions for both fast data storage of nonrasterized image information and high-quality conventional video storage. Full bit-rate uncompressed digital video recorders can provide the desired performance capabilities for both video and high-performance data recording, allowing dual-purpose use of the same transport and cassettes. Panasonic developed the D-5 (uncompressed 10-bit ITU-R 601) videotape recorder (VTR) format for high-end video post-production, and it has met with wide acceptance. The Viewgraphics (Mountain View, Calif.) "Dataview" is an industry-standard computer bus (VME) interface device developed to allow large data files to be recorded on unmodified D-5 VTRs and to allow video to be input as data into platforms like the Silicon Graphics Onyx and Challenge. This paper describes the D-5 format, the Viewgraphics Dataview adaptor, and data recording on videotape.

Film and production companies employ high-end UNIX-based computer systems today to generate and manipulate film and video images. These modern computers, with their phenomenal processing power and ever-increasing speed can release the user from the traditional production constraints and permit the realization of images and special effects that until now were only dreamed of. Such obvious benefits have resulted in a tremendous increase in the use of computer graphics, and their ever more cost-effective applications now extend to the realm of everyday production and post-production.

Fast computers are very efficient at numerically intensive operations such as morphing or multilayer compositing. However, they all suffer from an increasingly important deficiency: bit-rate transfer limitations to peripheral data storage devices. In addition, the file sizes for video, image creation, and manipulation can be voluminous. For example, 10-bit high-definition television (HDTV), 1080 active lines (V) x 1920 pixels (H) requires about 1.25 Gbits/sec (155.2 Mbytes/sec), and film scanned as RGB at a typical 2048 x 1526 sampling requires some 2.25 Gbits/sec (281 Mbytes/sec). Large film clips used for multiple layer compositing thus impose mammoth storage requirements with enormous attendant time delays to transfer data files to an archive tape media.

Users need to get multi-gigabyte files in and out quickly, not only to back up work in progress but to shift jobs in and out of the facility rapidly. Standard data recorders linked to such computers have provided very low bit-rate transfers, typically in the low Mbit/sec range in the early days up to several Mbytes/sec today. Using such low bit-rate data recorders created a problem because the computer could perform the required image processing in mere minutes, but hours were then required to transfer the same images to a data recorder. While the data was safe and secure, the time spent transferring the data to tape resulted in enormous inefficiencies. Expensive equipment was idle, no new project could be undertaken, and the wasted time and money parsing out a few megabits per second to a data recorder negated the benefits of the system's high speed and its attendant high cost.

The required solution is obviously faster storage! However, fast data recorders are frequently built to service a specific application; some are semi-custom and often military. Most are cost prohibitive. Nevertheless, at the same time that advances were being made in image computing, digital VTRs (DVTRs) were improving despite the staggering data rates required to record video and the users' desire for ever-smaller media. For example, conventional ITU-R 601 component 4:2:2 with 10-bit resolu-

Presented at the 138th SMPTE Technical Conference and World Media Expo (paper no. 138-28), Los Angeles, Calif., October 8 to 12, 1996. John Krooss is with Viewgraphics, Inc., Mountain View, CA 94043; and Stephen Mahrer and Philip Livingston (who read the paper) are with Panasonic Broadcast & Television Systems Co., Secaucus, NJ 07094. An unedited version of this paper appears in *Film and Video Synergies: Creation to Delivery*, SMPTE, 1996.

Table 1 — Standard Uncompressed Videotape Formats

Format	Format Bit Rate	"Useful" Payload	BER After ECC	Max Rec Time	Cost $
D-1 (19mm)	220 Mbits/sec	190 Mbits/sec	10^{-6}	90 min	$125k
D-2 (19mm)	125 Mbits/sec	90 Mbits/sec**	10^{-9}	208 min	$75
D-3 in. (1/2)	125 Mbits/sec	90 Mbits/sec**	10^{-11}	245 min	$50
D-5 in. (1/2)	300.6 Mbits/sec	238 Mbits/sec	10^{-11}	124 min	$75

**These figures are for NTSC D-2 and D-3 VTRs. PAL VTRs of both formats have format bit rates of about 20% higher.

tion requires about 210 Mbits/sec (26.2 Mytes/sec). A simple and attractive alternative for data storage would be to use a conventional "broadcast" digital VTR, with its significantly lower purchase cost, and provide it with a direct interface to the computer's memory. That solution has been implemented and will be described.

The Dataview SDA-21 is a very cost-effective solution for backing up and restoring large amounts of digital data using conventional digital VTRs and videotape as the media. Since its availability last year, Dataview has provided digital film producers with a much needed breakthrough throughput volume. Just recently, Dataview made it possible for Pacific Ocean Post in Santa Monica to establish a new mark for sheer volume of digital film used in a single production (*Independence Day*). Over the course of six months Pacific Ocean Post completed more than 200 visual effects that required the scanning and recording of 12.8 Tbytes (1 Tbyte = 1000 Gbytes) of film-resolution data.

Modern Videotape Recorders

Today's digital VTRs can provide payloads in excess of 100 Mbits/sec and feature industry-standardized digital input/output (I/O) and control schemes. These modern videotape recorders are basically divided into two categories: full bit-rate (uncompressed) recorders, and those using some form of video compression. The difference is crucial. Compressed VTRs cannot record external data due to the inherent use of video compression in the format. Uncompressed videotape recorders may be considered rasterized "bit buckets," providing high-speed data recording capability with numeric transparency. Table 1 shows standard uncompressed videotape formats in use for recording data. While most modern videotape formats provide more than adequate error correction for video and audio recording, additional or secondary error correction (ECC) beyond that done within the VTR is usually required for data recording purposes.

Note that the more modern formats have significantly improved post-error correction bit error rate (BER) performance. As witness to almost a decade of progress, the D-5 format often exceeds the D-1 format by a 100,000:1 margin in error rate with a much increased useful payload. The 238 Mbit/sec payload of D-5 is a 25% increase over D-1's 190 Mbits/sec.

The formats shown in Table 1 are conventional VTRs. As such they expect to be fed with very specific, steady and fixed data rate signals. There is no "flow control" in video. It is therefore a relatively simple matter to bit map the computer data into the digital input signal of the VTR. In this manner the data rate is kept constant, and the de-mapping of the recorded data is equally easy to accomplish upon playback. The extra error correction mentioned above is also often added at this step.

The Interface Solution

The Viewgraphics Dataview SDA is a single-board VME card designed to be inserted into a graphics computer system to support two primary functions. The SDA allows the computer to archive data to industry-standard digital video recorders, and to connect digital video inputs or outputs directly to the computer system (Fig. 1). In the data mode, Dataview permits users to selectively back up and restore computer data files between their host computer system and the digital video recorder operating as a high-speed "bit bucket." The SDA also includes both ECC and a read-after-write (RAW) circuit to ensure data integrity. In the video mode, Dataview provides real time video connectivity to serial digital video equipment. It provides a full bandwidth path for the video between an external device (usually a DVTR) and the host computer system memory, i.e., the computer system can be used to capture or output component serial digital video. In addition, the SDA can be configured with expanded buffer memory providing up to 20 sec of video storage on the board, and two Dataview adapters can be ganged together to facilitate 4:4:4 interfacing.

The Dataview adapts an industry-standard D-5 VTR for use as a high-speed, high-capacity data tape recorder. With Dataview, users can store up to 125 Gbytes on a single tape cassette, and can transfer data at

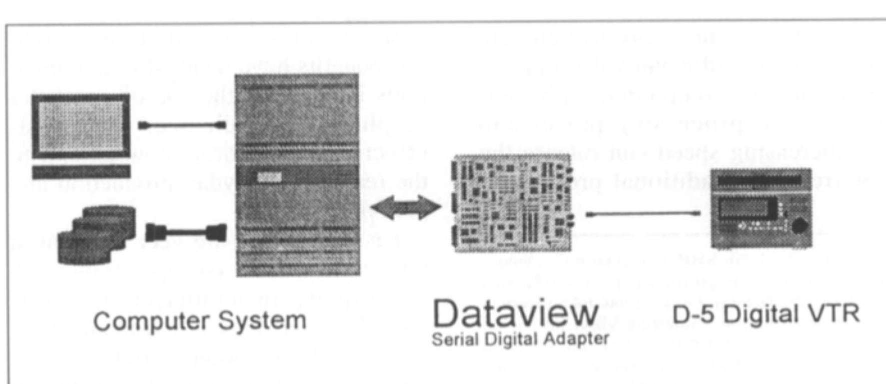

Figure 1. The Dataview SDA allows the computer to archive data to industry-standard digital video recorders, in this case, a D-5 VTR, and to connect digital video inputs or outputs directly to the computer system.

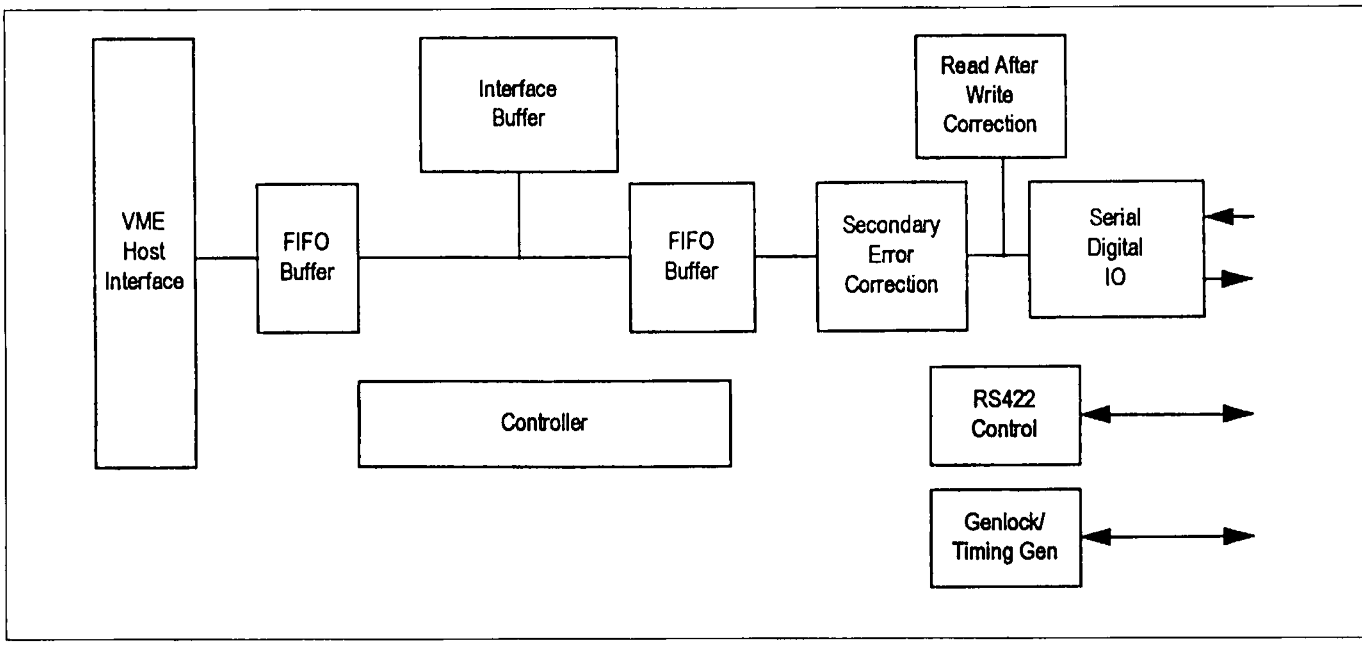

Figure 2. Detailed block diagram of the Dataview SDA.

speeds above 17 Mbytes/sec. This phenomenal capacity and speed allows backup or restoration of the entire contents of a large disk array in minutes instead of hours or even days. By striping together multiple D-5 tape units, data rates in excess of 100 Mbytes/sec can be achieved, and, by using the Panasonic M.A.R.C. robotic system, the multiple machines can have access to hundreds of tapes.

The Dataview adapter was designed for production environments, and a key objective was to ensure reliable operation independent of host computer configuration and I/O performance. Because the host computer is not normally a real-time system, there may be occasions when it is difficult if not impossible for the computer, no matter how powerful the configuration, to continually input or output video data streams. By utilizing an on-board interface buffer, the Dataview SDA continues to receive data during these brief periods and thereby relieves the host computer system of the demanding requirement of a real-time response. As long as the host computer can sustain the full-bandwidth transfer rate on average over time, the Dataview SDA can provide real-time video input and output as well as data transfers at the limit of the component serial digital interface. A detailed block diagram of the Dataview SDA is shown in Fig. 2.

The Dataview SDA can sustain a data transfer rate of 17.5 Mbytes/sec to a D-5 VTR directly from host computer memory. This can be maintained without interruption for the full duration of a D-5 cassette (124 min maximum). Similarly, in the 4:2:2 video mode, rates above 20 Mbytes/sec can also be sustained (full uncompressed digital video). To maintain this high data rate to and from the storage device connected to the host computer, the device must be able to transfer data to and from host memory at a similarly high sustained data transfer rate.

The file systems of most servers are not able to sustain the necessary data transfer rate. Care must be taken by the system integrator to ensure that not only is the raw hardware performance available, but that the file system of the server is sufficient for the task, for example, Silicon Graphics' Extended File System™ (XFS). For those file systems that can transfer data at an average of 20 Mbytes/sec, the ability of the adapter to buffer data on board has the effect of allowing the host computer to keep up the transfer at a high rate even if the connected storage device can only sustain high data transfer rates for short periods of time.

Table 2 shows the data transfer rates of the Dataview SDA in different modes of operation. In the data modes, data is written to tape with added ECC data into the active range of a 4:2:2 video frame.

The Dataview SDA can handle files up to 4 Gbytes in size when used to back up a normal UNIX file system, the file size limit imposed by UNIX. However, when used with an XFS file system from Silicon Graphics, much larger files up to the capacity of a tape, approximately 127 Gbytes for a 124-min D-5 cassette, can be stored. Table 3 shows the capacity of different sizes of D-5 cassettes in each data mode.

The Dataview employs a proprietary Data On D5 (DOD5) encoding/decoding scheme for storage of data on tape. Additionally, it utilizes its capabilities for ECC, data block checksum (DBC), and RAW functions to guarantee data integrity. The data mode capabilities of the Dataview allow it to provide high-speed, resolution-independent, digital video, or image file backup and retrieval.

Another key feature of the Dataview Adapter critical to the post-production process is direct file access. In contrast to other data tape devices, Dataview permits users to access individual files stored on the tape directly by subdirectory and file name. This feature greatly speeds production throughput and simplifies systems integration.

A comprehensive application programmers interface (API) is provided

Table 2 — Transfer Rate of the Dataview SDA in Different Operational Modes

	DOD5 525 VTR	DOD5 625 VTR	Video 10 bit (4:2:2) Y, R–Y, B–Y
Number of bytes/line	1,260	1,260	1,800
Number of lines/frame	486	576	486
Tape rate frames/sec	30	25	30
Transfer rate Mbytes/sec	17.52	17.30	25.03

Table 3 — Data Capacity for Different D-5 Cassette Sizes

Tape Capacity for Different Sizes of D-5 Cassettes (Gbytes)	Data on D-5 525 VTR	Data on D-5 625 VTR
Medium cassette (12 min)	12.32 Gbytes	12.17 Gbytes
Small cassette (23 min)	23.61 Gbytes	23.32 Gbytes
Medium cassette (33 min)	33.88 Gbytes	33.46 Gbytes
Medium cassette (48 min)	49.27 Gbytes	48.67 Gbytes
Medium cassette (63 min)	64.67 Gbytes	63.87 Gbytes
Large cassette (94 min)	96.50 Gbytes	95.30 Gbytes
Large cassette (124 min)	127.29 Gbytes	125.72 Gbytes

to allow custom development. Viewgraphics has over one dozen developers, including industry leaders Kodak and Discreet Logic, who have easily integrated their programs with Dataview.

The Data Recording Solution: The D-5 Component Digital VTR

The first digital VTR format, the 19 mm D-1, made its debut in 1987, followed by D-2 in 1989, finally making digital recording technology a viable alternative to that of analog. Rapid developments in technology prompted development of a 1/2-in. digital composite VTR format that could be implemented in a camera recorder similar to 1/2-in. analog component, and D-3 made its debut in 1991. However, the world fully expected to enter the era of high-definition television in the late 1990s, and this has created a demand for component signal-based systems that facilitate editing and flexible format conversion with high resolution.

Since this transition to component recording was envisioned by the developers, notably Matsushita and NHK, the D-3 composite format was created not only to satisfy the present demand for a cost-effective composite format, but also to form a base or platform for a component format. D-5 realizes full-bit component recording with full production features and is based on the D-3 format and technology.

While D-5 has been created on the D-3 platform, there are significant differences. In the D-5 format, the linear tape speed is twice that of the D-3 format in order to maintain the same track pitch as D-3 while still doubling the number of recording channels from two to four. The number of digital signal processors, recording and playback amplifiers, rotary transformers, and heads are also approximately doubled — a D-5 VTR has 18 heads while a D-3 VTR has 10 heads. It is thus possible to record digital component signals whose video data rate is about twice that of digital composite signals. The maximum recording time of the D-5 format is 2 hr (half the 4 hr of D-3).

As was previously described, two different methods have been conceived for new digital component VTRs. One is full-bit recording, and the other is bit-rate reduction (BRR) recording. Since D-5 is a full bit-rate (uncompressed) format, it is not only appropriate for the 525/60 and 625/50 standard-definition television (SDTV) systems, but can be extended to HDTV recording using a 4:1 compression adapter, and to true digital data recording as well.

Data Capacity

The D-5 format consumes 21.2 cm^2/sec of recording media with input data rates of 270 Mbits/sec (10 bit) or 288 Mbits/sec (8 bit), as compared with 25.0 cm^2/sec for D-2 (143 Mbits/sec 8 bit) or 54.3 cm^2/sec for D-1 (216 Mbits/sec 8 bit). This allows up to 2 hr of program material in a single large cassette (1237 m of tape). This 2-hr program length equals approximately 127 Gbytes of data storage, or about a 100-Gbyte payload per cassette.

D-5 Specifications

- Scanner rotation: 90/1.001 Hz
- Tracks per rotation: 8
- Track azimuth angle: ±20°
- Drum diameter: 76 mm
- Helix angle: 4.9°
- Wrap angle: 178.11°
- Total average data rate: 300.6 Mbits/sec
- Channel bit clock: 133 MHz (7.5 nsec)
- Shortest wavelength: 0.64 mm
- Tape consumption: 7.6 m^2/hr

Figure 3 shows the D-5 tape footprint for one field of recording (525).

8-Bit/10-Bit Signal Processing

The D-5 VTR accepts two types of digital component signals: 18-MHz sampling with 8-bit quantization, and 13.5-MHz sampling with 10-bit quantization. In the 13.5-MHz mode, the 10-bit samples are reformatted to 8-bit samples, creating new 8-bit words from the aggregated 2-bit "difference" in word length. At the same time, the bit rate is changed to 18 MHz in the 10-bit to 8-bit conversion process in order to achieve efficient and identical processing regardless of input signal. Figure 4 illustrates the 10-bit/8-bit data mapping used by a D-5 VTR.

At 13.5 MHz, each active line contains 720 samples of Y and C to be recorded, but 48 samples are added to each, making them equal 768 samples. These 768 samples undergo the same channel and segment distribution as that for 18-MHz signals. The four components Y_E, Y_O, P_B, and P_R are distributed evenly to all channels and

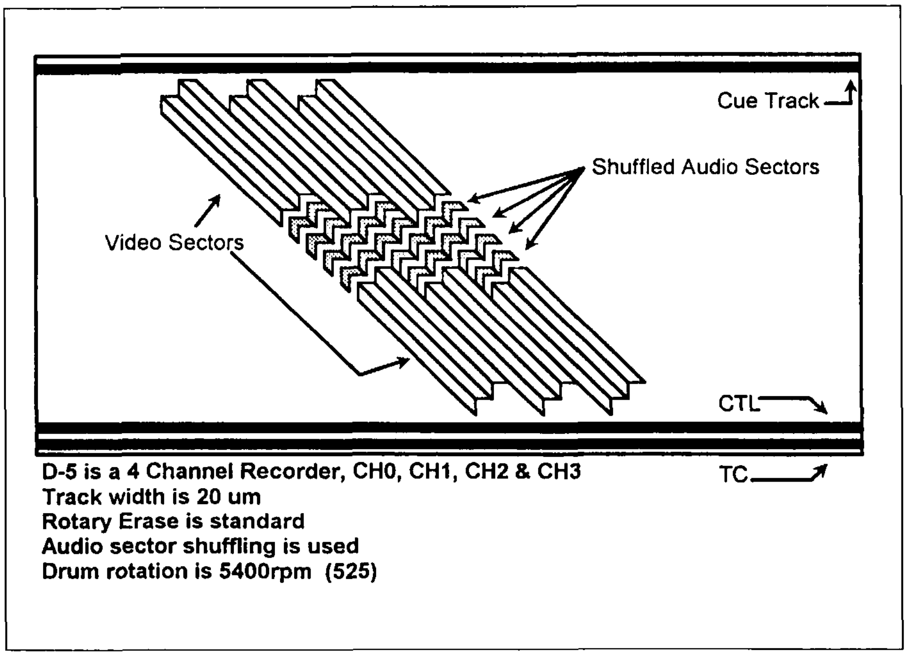

Figure 3. D-5 tape footprint for one field of recording (525).

video segments, and each video segment of each channel is comprised of 32 samples in 525/60 mode, or 24 samples for 625/50 mode. The 10-bit words are divided into the eight most significant bits and two least significant bits; the two least significant bits from four samples are collected together to make one 8-bit pseudosample. To assure that potential errors are largely contained within a single pixel, the two least significant bits of the Y_E, P_B, and P_R samples of a single pixel are placed into the same 8-bit pseudosample. (This philosophy is repeated in channel and segment distribution.) The two least significant bits of the Y_O sample are placed into pseudosamples, separated by 16 samples (for 525/60) or 12 samples (for 625/50) in each video segment. Through this process, 32 bytes (525) or 24 bytes (625) of 8-bit data are assembled from the two least significant bits for each video segment in each channel. Distributing this data evenly into four outer codes for each video segment results in 8 bytes (525) or 6 bytes (625) each, and adding these to the "eight most significant bit" group of 32 bytes (525) or 24 bytes (625) data yields 40 bytes (525) or 30 bytes (625) of 8-bit data, which is the same as that for 18 MHz. After intra-subsegment shuffling, processing is performed in the same way as for 18-MHz signals, eliminating the need to greatly increase the hardware for field memory and error correction to 10-bit and yet maintaining full-bit recording.

Clearly, this formatting allows the remainder of the processing to function on 8-bit data at 18 MHz (one of the two signal formats accepted) and minimizes the difficulty of handling two different input signals while passing the full 10-bit 13.5-MHz signal. What may not be so clear is the contribution that this makes to using D-5 as a data storage device and the robustness of the internal error correction scheme due to the cosited data.

D-5 Video Data Arrangement

In D-5, the three cosited pixel samples (Y, P_B, and P_R) are handled as a single group and are written to the tape in the same sync block if at all possible. This is preferable because any erroneous sample(s) of the three cosited Y, P_B, and P_R samples result in only a single pixel error in the image, whereas three erroneous samples of distributed Y, P_B, and P_R would result in three pixel errors. Therefore, as much as possible, each group of cosited samples is placed into one sync block in D-5.

The outer Reed-Solomon codeword of 120 data bytes and 8 check bytes can correct up to 8 bytes of error. The inner codeword is either 85 data bytes (525/60) or 76 data bytes (625/50), with 8 check bytes attached, and can correct up to 4 bytes of error.

D-5 Tape/Format Interchange

D-5 is typical of most modern tape formats in that the tape path is carefully designed to be stable and repeatable. In this way tapes from other VTRs may be played with confidence, knowing that there will be no marginality of data recovery. Head-to-tape interfacing is very stable due to D-5's low tape tension (< 30 G). Head wear and tape life are also greatly improved over earlier formats.

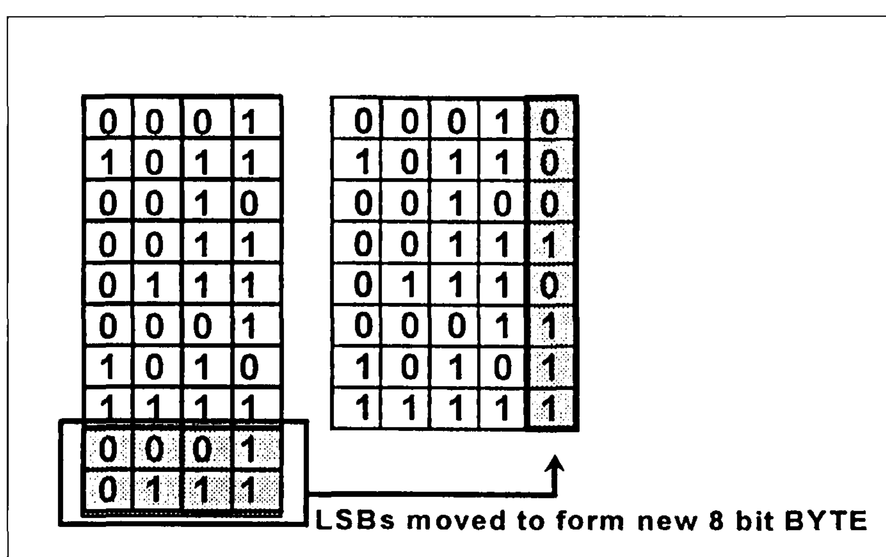

Figure 4. 10- and 8-bit data mapping used in D-5.

D-5 Bit Error Rate

D-5's powerful ECC combined with its stable and robust tape path help provide a very low post-ECC bit error rate. D-5 BER with the ECC off can be compared with D-1 with its ECC active. When both formats are working as designed, D-5 produces a post-ECC BER of about 10^{-11}, about a 10^5 improvement over D-1. D-5's proven track record of superb reliability, low BER, and very useful 238 Mbit/sec recordable payload provide a very useful format not only for conventional 4:2:2 component graphics, but also in conjunction with the Viewgraphics Dataview as a high-speed data recorder to back up today's high-speed graphics computers.

Conclusion

Archiving of computer data on digital videotape allows graphic and image manipulation users to realize the full potential of their investment in computer technology. With Dataview's diverse set of operational modes and D-5's cost-effectiveness, speed, and reliability, the computer graphics and film production process can retain the high throughput methods and rapid job shifting once reserved only for video production. At the same time, the system benefits from the dual nature of the Dataview adapter, acting as both a high-speed data and video storage/retrieval interface. By combining speed, capacity, and economy, these two new technologies will forever break the bottleneck that now exists in computer-based production

THE AUTHORS

Philip Livingston

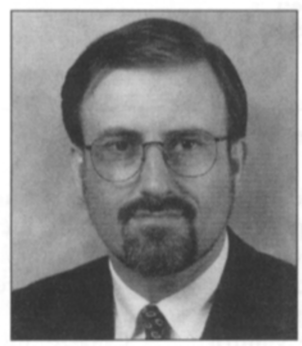

Stephen Mahrer

John Krooss is the founder of Viewgraphics, Inc., and has served as president since 1985. Previously he was an engineer with Ford Aerospace's Western Development Labs, specializing in image processing and remote sensing systems. He received a B.S. degree in electrical engineering from the Massachusetts Institute of Technology in 1982. Based in Mountain View, Calif., Viewgraphics manufactures imaging systems for television research as well as a wide range of digital video interface products.

Philip Livingston is director and general manager, Digital Systems Group, Panasonic Broadcast and Digital Systems Co., a division of Matsushita Electric Corp. of America. During almost 20 years with Panasonic, he has held the positions of chief engineer, video; manager, high tech center; and assistant general manager, product management and engineering. Prior to Panasonic, he served at the State University of New York at New Paltz for ten years in Instructional Resources and as engineering vice-president of TeleCommunications, Inc., in Nashua, N.H. While his engineering career began in radio, he participated in the construction of WOKR-TV in Rochester, N.Y., and later worked as Chief Television Engineer for the City School District of Rochester. Livingston attended Union College in Schenectady and the Millbrook School. A member of SMPTE, he currently chairs the Television Production Technology Committee (P.18).

Stephen Mahrer, whose background represents over 22 years of design and engineering on both camera and VTR products, systems, and product support, is currently senior technologist for the strategic accounts group of Panasonic Broadcast and Digital Systems Co. He had previously held such positions within Panasonic as digital VTR engineering manager and Olympic Project manager. He has been involved in digital VTR engineering for Panasonic's D-3 and D-5 1/2-in. formats, which often required the presentation of technical papers to such groups as the SMPTE and NAB.

Prior to joining Panasonic, Mahrer was senior staff engineer, Technical Development Laboratory for NBC, Inc., in New York City. His projects there included work on advanced television, the conversion of NBC operations to MII, engineering support for the 1988 Seoul Olympics, and the custom design of a digital video data signaling system (for which he was later awarded a patent).

Mahrer came to NBC from RCA Broadcast Systems, after being transferred to the U.S. from RCA's European manufacturing base, RCA (Jersey) Ltd. His work at RCA concentrated on CCD camera design and product support for RCA's existing PAL/SECAM equipment, much of which was extensively customized by RCA (Jersey) Ltd. for the European market.