ANOTHER NIGHT AT THE DIGITAL VIDEO ROUNDTABLE:

BY CHRIS ALLAIN

Across the road from the Las Vegas Hilton and Convention Center, philosophical light-years away from awards to Carol Brunett and other trappings of the NAB's own vision of what the annual conference that bears its name is all about, 70 people gathered in the Rotunda Room of the Las Vegas Country Club. They were there that evening to discuss the real future of "television" and to share their concerns about the technology barriers that still stand in the way of what digital video and media production promises to offer.

One year ago I wrote a report in this magazine on The Digital Video Roundtable held at NAB '94. I described the genesis of the event, an idea that sprang from discussions on America Online (AOL) in the Video.sig (special-interest group). From there it moved to a hotel room during NAB. Then Randy Ubillos and Tim Myers, formerly of Adobe, arranged for a larger room and invited select participants.

For those involved with producing digital video on the Macintosh, The Digital Video Roundtable was immensely useful. It provided users an opportunity to discuss nagging and complex issues with the developers of products they use. It allowed developers to hear what a select group of leading-edge users do with their products, and to find out what new tools they need. Some of the most prolific developers in the industry were listening closely.

Last year's event impressed me enough to ask VIDEOGRAPHY Editor Brian McKernan if this magazine could co-sponsor it. McKernan agreed, and this year we helped organize, fund, and develop an agenda for an even better Digital Video Roundtable. This year, Adobe's Charlie Donaldson and I made the arrangements for the meeting place and took care of invitations. Macromedia, Randy Ubillos' new employer, and Data Translation also helped fund the event.

Attendance was larger than ever, but still kept to a manageable size. Representatives were there from most of the players in the Mac digital video arena. Logistically the event was smooth, that is after we realized that there was enough

Contributing Editor Craig Birkmaier (standing) makes a point as (right to left) Adobe's Charlie Donaldson, Macromedia's Randy Ubillos, and Contributing Editor Chris Allain look on.

cable in the podium PA to pass the mic to each speaker. Unlike some events held during NAB, much of the food remained at the end of the evening. These attendees did not come to eat; when it was over many didn't want to leave. There was still so much more to discuss. So what did the Roundtable accomplish?

We moved through an ambitious agenda briskly, thanks to Ubillos' leadership. The Roundtable helps to identify the problems the industry is trying to deal with and to develop directions for technological advancement. I think it is fair to say that at times we've gone beyond identifying the problems and suggesting solutions. Certainly, there were users who came away with work-arounds that allowed them to temporarily deal with some problems. A week or so after NAB, one AOL user described The Digital Video Roundtable as "the best part of the show." I feel the same way.

> We discussed faster network solutions, since standard Ethernet provides little value.

I've organized this report loosely into a problem-and-discussion format, providing background where I could. So pull up a chair to the 1995 Digital Video Roundtable.

Mass Storage-Scenarios for Working with Data

The problem with storage is simple: You always need more of it than you have, and it's too difficult to move data from place to place. Digital video demands more of data storage and transfer than nearly any other application of computer technology. Large throughput requirements make it necessary to use arrays that are less portable than single fixed hard drives. Ethernet, by far the most common network solution, is sorely inadequate for moving data from system to system. The group clearly recognized the need for much faster networking *now*. Since that will happen in its own sweet time, however, we discussed other solutions currently available.

Physically moving drives from place to place, a common practice referred to as "sneaker net," currently offers the most practical solution. The vast majority of multimedia producers use this approach today. You can't push a 400 or 500 MB animation file across Ethernet in a reasonable time. So, we shutdown the computer, disconnect the drive, move to the workstation of a new user, reconnect the drive, and then restart the system. If done only occasionally, a user can put up with this, but done multiple times a day, it's becomes a problem. The process can damage the drive connectors, and waste time.

The "hot swap" solution moves us a step up from fixed drives. *Hot swap* refers to the ability to remove a drive or add one to a system without shutting down. Typically hot swapping also involves the use of some sort of carriage to cradle a drive or a group of drives. It allows a user to remove the drive by simply ejecting it. This solution is interesting, but adds to the individual drive cost by $200 to $300, or much more for an elaborate system. Using striped drive arrays further complicates the process of hot swapping.

Hot swapping striped drive arrays could be a bit tricky. One would probably be wise to use the same disk controller and striping software on each of the systems that would share the drives. The cabling could be a nightmare, since some arrays use single-channel Wide SCSI controller cards like those from ATTO and FWB, and others use the dual-internal SCSI 1 controllers built into the Quadra 900/950s and the Power Mac 8100s. Generally, striped arrays are more trouble than stand-alone drives, so if you try this expect to spend some time working out the bugs.

We discussed faster network solutions, since standard ethernet provides little value with these file sizes and data rates.

Although still quite expensive, a few fast networking solutions are beginning to appear. FDDI and 100Mbs Ethernet are network solutions that promise to step up the bandwidth of Ethernet significantly. Apple is working with another called FireWire that is potentially faster still, but none of these are in common use today.

SCSI Net, from Transoft, is a high-speed network solution available today (see Distributed File Sharing," VIDEOGRAPHY 4/95). It uses differential SCSI connections to a central hub. Each of the connected computers—up to six of them—uses an ATTO SCSI card, and may be located as far as 80 feet from the hub for a 160-ft. network diameter. Transoft claims 20MB/s transfers, faster than any of the other high-speed solutions, and fast enough for digital video applications. This system sounds sweet, and Transoft prices it accordingly. A system will cost $15,000 to $25,000 without drives. Transoft says that Avid bought a large quantity of these systems.

More than 70 industry professionals attended the Roundtable, a forum on the future of teleproduction.

Sonic Solutions, known for high-end audio systems, provides a high-speed solution called MediaNet that uses a "distributed" network solution. The product combines the function of an FDDI network card and a SCSI controller card. It connects directly to the hard drive. Its throughput is high, but not high enough for many digital video applications.

Another available solution, CreativePartner from *e*motion, improves throughput on standard Ethernet cabling, and provides an interesting set of tools for media work groups. (See "Hooked on a Feeling," VIDEOGRAPHY 5/95).

Rating Throughput for Hardware Codec's

Rating the throughput of codec systems is an area that many Roundtable participants stated could benefit from improved nomenclature. One consensus clearly arose against comparisons to analog tape formats, such as S-VHS or 3/4-in. U-Matic. The artifacts from compressed digital video and analog tape machines differ to the point where comparisons are meaningless.

Proprietary systems, such as Avid's AVR levels, present problems when comparing them to other systems. The "MBs of QuickTime video" is better because it gives a relative indication of the quality of the video. The megabytes per second (MB/s) rating causes problems though, when confused with a drive's rating in MBs. Due to system overhead, a drive's data-transfer rate is not equivalent to the rate at which it can feed video data to the codec. So, you end up with two sustained data rates described in MB/s.

A consensus eventually developed toward the KB per frame rating used by Data Translation. It is an unambiguous rating with the added benefit of being just as relevant at frame rates other than 30 fps.

Hopefully, other systems will adopt this rating. Although one wouldn't expect Avid to abandon their AVR levels, perhaps they could additionally provide the KB-per-frame ratings of each quality level.

Color Space Conversion and Gamut

The next issue regards color space conversion and color gamut. The science of color is complex and beyond the scope of this writing, but what you must remember is that a color produced in one color space may not be reproducible in another. For instance, If you attempt to reproduce the I & Q components of NTSC color bars with an RGB board, you'll find that you simply can't do it. The I & Q values, from YIQ color space, are not within the gamut for RGB. The I and Q values have saturation and color, but no brightness, so perfect encoding into RGB yields black. At least one of the three colors must include a brightness level greater than zero in RGB color space, otherwise black will result.

Similarly there are colors available in RGB that seemingly can't be recreated with the CMYK color space, and vice versa. CMYK—cyan, magenta, yellow, and black—is the color space used for printing. Have you ever had a client bring you a business card and ask you to use the color on the card in an on-screen graphic, only to find that you can't? You may have also created a great looking image on a monitor that you simply can't print accurately, even with a properly calibrated system.

If you overlaid accurate color space charts, which you can't since CMYK printing would accurately represent only the CMYK space, you could demonstrate which colors can't be reproduced. Those would be the colors where the charts did not overlap.

It would be useful to know whether a specific product was scaling the color values into the gamut of a new color space, or if they were merely truncating the excessive values, which might lead to loss of detail.

A basic comprehension of color space helps in understanding where conversion is taking place, and in dealing with the color gamut of 8-bit CCIR 601 video. The eight bits in digital video refer to each of the 4:2:2 components of YUV or Y, R-Y, B-Y. It is similar to the 24-bit color standard in computer RGB. It differs in that it only consumes two-thirds the data storage because it samples only half as often along the horizontal line for the color-difference signals, R-Y and B-Y, as are made for Y, the luminance signal.

The 601 specification calls for even fewer values than you might think. The range of values in an 8-bit signal is from 0 to 255, but the specification reserves several values for other purposes, so that the range for Y, or luminance, for each pixel may vary from 16 to 236. The range of values for the color-difference signals, runs from 16 to 240. To further complicate matters, digital video devices can also use an extended value range of from 1 to 254. This is a set-up option with Intelligent Resources' Video Explorer boards equipped with digital output, and with Digital Betacam VTRs.

Problems from improper setup can arise when the extended-range signal feeds a device expecting the more limited range. This results in a loss of detail in the dark areas, as all values from 1 to 15 become black.

You can also experience problems in normal use when you create a color ramp in RGB space that you must convert into the limited color gamut of 8-bit YUV space. Occasionally the range of colors used in the RGB ramp are not available in 8-bit YUV. Banding can result from the conversion, when 8-bit YUV offers too few values to express the range. The fix for this is to add a small amount of noise to the image before the conversion. This randomizes the variation and reduces or eliminates the banding. The dithering that results effectively simulates a greater range of values.

These problems illustrate the reasons that vendors are trending toward 10-bit systems. The additional color values available in a 10-bit system eliminate the primary shortcomings of 8-bit component digital video. Of course, at next year's Roundtable we'll be asking the software vendors to add support for 10-bit rendering and output.

Field Labeling

Field labeling is another topic returning from last year's roundtable and another where the industry would benefit from better nomenclature. It has to do with the labeling and dominance of the interlaced fields composing a frame. It seems that every manufacturer has come up with their own method of labeling the fields. Since dealing with interlacing is already complex, it shouldn't be necessary for the user to have to figure which field is which based on non-standard terms.

Typical names include: odd and even, A and B, Upper and lower, field 1 and 2, and others simply refer to a frame as being field-1 dominate or field-2 dominant. Dominance refers to the chronological priority of the fields. Which ever name vendors select, we simply wish they could be consistent.

Just for comic relief, here's a top-ten list of names that are only slightly more obtuse than some of those used now:

10. North and South
9. Beavis and Butthead
8. Boutrous and Boutrous
7. Ying and Yang
6. Leaded and Unleaded
5. Inee and Outy
4. Abbott and Costello
3. Dominate and Submissive
2. Plain and Peanuts
1. Darell and Darell

QuickTime Support for Alpha Channels

Several animators and editors voiced the need for QuickTime to support motion-alpha channels, sometimes referred to as a flying key. A single compressed movie file cannot currently contain both an animation's color frames and the alpha channel frames. This requires two files that you must separate and later reattach during composition. If a user wanted to employ the codec's compression with key-channel frames, it would require multiple steps to create the second file containing the alpha-channel frames.

We discussed two methods of dealing with the key-channel data stream: non real-time, and real-time using two chip sets. We briefly covered the dual-codec approach.

A monochrome key channel, or flying key, would have a data rate of about half that of the "4:2:2" full-color stream, since the codec would compress only the luminance, the "4" or Y signal. At the moment, it's impractical, with most compression products to provide a second real-time compression stream anyway, since the current bus architecture restricts the data rate for a single video stream.

The nature of the content of the key channel might also change our approach to its compression. The vast majority of pixels in a typical key channel frame are pure black or white, but the border areas between these extremes are critical to quality keys. Therefore, a lossless RLE (run-length encoding)-style compression would better suit the key-channel stream. Because of a high propensity for a few values, it would usually compress very efficiently without loss.

The second method discussed for treating the key channel would use RLE compression for a non real-time solution. RLE is a lossless compression scheme.

It was suggested that QuickTime could handle this by alternating the storage of frames between the full-color frame and the encoded alpha-channel frame, that is, color frame-key channel frame-color frame-key channel frame, and so on. Although this might not currently offer real-time playback, hopefully the standard might support that when hardware allows.

Since the Roundtable, work—or at least serious talk—has begun on the subject of an RLE-encoded alpha channel for motion JPEG-streams. Eventually, though, Apple needs to do it in a QuickTime standard. Who knows why this wasn't done from the beginning of QuickTime? Considering that Apple built 32-bit Quickdraw into the operating system when most paint applications were 8- and 16-bit, how could they have missed this one? True, today's compression boards can't play it back, but the file format should have supported it. This would convenience animators and editors trying to deal with files that use a separate key channel. At the rate this technology is advancing, it won't be long before playback is reasonable as well.

Since the Roundtable, work has begun on the subject of an RLE-encoded alpha channel for motion JPEG streams.

Square Pixel vs. Rectangular Pixels (Using 720 by 486 Resolution)

This issue returned from last year, although it has advanced somewhat. Now, at least, we are identifying solutions along with providing work-arounds.

The resolution problem faced by system designers and users results from the fact that video standards—such as CCIR 601 digital video—call for non-square pixels, and computer systems typically deal with square pixels. The problem presents itself in various ways. At the most basic level, images created for one pixel ratio are distorted when displayed at another ratio. When designing images in a paint environment, tools don't work as they should, for example, round paint brushes will paint with a different weight when drawing vertically and horizontally. Circles and squares become ovals and rectangles.

When manipulating images for two-dimensional effects, software distorts rectangular pixel images when it handles them as square-pixel images. For instance, if a rectangular pixel image is rotated about the Z axis in After Effects, as the image rotates it becomes warped. Rotating the image 90 degrees squeezes it vertically. A square becomes a rectangle. Also the rectangle appears skewed as it is rotated.

To work around the problem, produce images at an oversized vertical resolution of 720 by 540, and scale down by .9 vertically as the last step to generate the required 720 by 486.

The solution requires the software to handle the non-square pixels. A user could instruct a program to deal with an incoming image as a non-square pixel image, or to render the image for output onto a non-square display. Digital video effects devices and dedicated broadcast paint systems deal with similar problems, so we know they can be solved. Hopefully programs such as Photoshop and After Effects will soon begin to provide solutions for this problem.

Software Codecs for Hardware Systems

Another topic returning from last year is software codecs. Users need a software codec to accompany the hardware codec so that they can view images on another system without having the hardware codec present. Even though it would not be able to access the frames in real time or play back the video at 30 fps, it would be useful. A software-only codec would allow a user to convert a video file to another codec, in non-real time, just to get some form of playback. This would allow sharing of video clips where a sender was not certain of the hardware available on the receiving end.

Another important reason for the software codec is to allow users who own a hardware-codec system to render on another computer that does not have the hardware codec present. For instance, if you have your Radius VideoVision system installed in a Quadra 840AV to get top throughput, but you want to render After Effects compositions on a PowerMac 8100/110 since it is much quicker, you cannot now do without moving the boards to the 8100. If you have two copies of After Effects, and you'd like to render on two systems simultaneously, you'd have to convert the file to an uncompressed format, render, and then reassemble the frames. You'd save no time. Few users would buy a second expensive board set just to be able to render on a second computer, so it is not really an issue of a vendor costing itself sales by releasing the software codec.

Avid announced its software codec at NAB, as a part of its OMF standard. At least one other vendor (Radius), although unable to announce, privately hinted that they would release one as well.

Most agreed that free distribution of a software codec would not cannibalize sales of hardware products, and that it would be necessary for a vendor's product to become dominant in the market.

Standards for Hardware Acceleration of Effects

This was also a return from last year. Clearly Apples must lead here. It seems we need a standard architecture where QuickTime "brokers" the acceleration of effects. The software and hardware would each check in with QuickTime, allowing one to take advantage of the other without direct support. This is essentially the way compression hardware works with QuickTime software. One development that could perhaps come sooner is a unity in plug-in architecture between Premier and After Effects. Many users do most of their work in After Effects and can't benefit from accelerators and third-party effects that work as Premier Plug-Ins.

The Roundtable discussion came to focus on data rates. You can't make effects happen much faster if you can't deliver the data to the processor more quickly. Of course talk then turned to the PCI bus, which hopefully will deliver the kind of bus speed needed to move digital video data among processing cards more quickly.

Another facet of this issue came up. It regards the use of captive VS common media pools. Devices such as the Digital Magic from ADI (Advanced Digital Images) have SCSI controllers on the codec board, and so can deliver high-data-rate video even with slow CPUs. Although this approach has several advantages

Apple should have sent a representative; no one came.

there are potentially tradeoffs. For instance, users can't easily make the video data available to others on a network.

In a related issue, it matters where the processing takes place. Hold the data for more than one frame in memory to process it on the codec board, and you can avoid a tremendous amount of bus traffic. The Targa 2000, from Truevision/RasterOps, and the Video Explorer, from Intelligent Resources, have this capability. Another approach might be to connect a processing board directly to the video card's dedicated bus, like the VideoBahn on the Video Explorer. These approaches can keep the data rate down on the main CPU's bus. Perhaps, following on that thought, board vendors would be wise to build even greater effects processing power on the video/codec board, avoiding as much bus traffic as possible. We should all be aware of, and moving toward, the ultimate solution, which would move and process multiple video streams in real time.

Chroma Keying and Compression

Roundtable participants report that chroma keying using blue-screen matting does not yield satisfactory results when using JPEG compression on the source material. Apparently the JPEG artifacts reduce the sharpness of the key. According to ADI's Scott Auchmoody, who has performed tests in the area, you can get successful results with compression levels in the range of 3:1 and below. Auchmoody reports that this is partially due to JPEG compressing blue more heavily than other colors. He says that you can perform green-screen chroma keying successfully at higher compression levels.

The Best Computer

A number of users and developers expressed their preference for Apple's Quadra 840AV. The current top-of-the-line machines —PowerMac 8100s—don't allow for as high a data throughput as the 840AV, even though the processors offer much more speed. The newest 8100s are better than the earlier ones, but when you're pushing the envelope in performance the way digital video applications do, you need all the speed you can get. The 840AV barely delivered the data rate for very good quality video, so for many the step down in throughput is unacceptable. Unfortunately, high-level users also want the benefit of the faster processor for rendering. It is really disappointing that Apple would release a top-of-the-line machine with less throughput than its predecessor. I suppose Apple saw it as a temporary problem, since the PCI bus machines due out this summer are expected to provide significantly better bus performance.

It was also pointed out that this bus speed only presents a problem on video compression systems that route all video streams over the primary computer bus, as most do. The Digital Magic uses an onboard SCSI controller with dedicated hard drives. It can achieve top video data rates on virtually any NuBus Mac computer.

In Closing

That's that way it was at the 1995 Digital Video Roundtable. Thanks to: Dave Herbstman and Dan Wilk of Adobe, Hage' VanDijk, of Radius; and Scott Auchmoody of ADI.

They helped me to reconstruct the discussions and develop the background material I've written on some of these issues. Thanks also to Charlie Donaldson, my partner in staging the event, and to Randy Ubillos, who managed the discussions.

One final issue that came up repeatedly at the meeting regarded a no-show. Apple should have sent a representative; but no one came. Last year Peter Hoddie, of the QuickTime, team attended, and I believe he would have come again if he could have. I suppose we should have worked harder to contact more Apple people earlier on, but why weren't members of the QuickTime team all over NAB? You have to wonder sometimes, after years of unpaid evangelism, why Apple doesn't better recognize the needs of the entertainment-production market. You can bet that Silicon Graphics does! Regardless, we thank the folks at Apple for the brilliant products they've delivered and look forward to working with them in the future. ☐