THE DIGITAL VIDEO ROUNDTABLE:
SIGGRAPH '95 EDITION

BY CHRIS ALLAIN

Major video industry trade shows are practically the only occasion when large numbers of savvy computer-video users and the software developers and manufacturers that serve them are in the same place at the same time. Seizing upon this fact, VIDEOGRAPHY magazine has, during the past two NAB shows, conducted what it calls its Digital Video Roundtables in conjunction with such co-sponsors as Adobe, Macromedia, and Data Translation. Last August, during SIGGRAPH '95, VIDEOGRAPHY Contributing Editor Craig Birkmaier organized the third in our series of Roundtables. Once again, a select group of users and developers were invited—this time to a conference room in a Los Angeles hotel—for an evening of frank discussion on technical issues impacting the digital domain of computer-video integration.

As reported in our previous coverage of these Roundtables ("A Night at the Digital Video Roundtable," 6/94 and "Another Night at the Digital Video Roundtable," 6/95), these events have tended to be Macintosh-centric meetings. During this latest meeting, however, we sought the participation of developers working with additional operating systems. Co-sponsors Adobe and Macromedia both develop for multiple platforms. No major Silicon Graphics developers showed, but some PC-only developers and Mac developers with cross-platform products represented the "WinTel" (Windows/Intel) camp.

Open DML

The program began with an introduction of a Matrox representative who described his company's Open DML project. The Open DML group includes Matrox, D-Vision and other Matrox OEM partners (see "Promise Versus Reality," 6/95). This group is working to patch Microsoft's AVI standard to make Windows more useful for video, film, and multimedia applications. Their goals include some of the standards already available in QuickTime and others that QuickTime users are waiting on as well. The group has developed a preliminary specification for file format compatibility and will soon work on issues concerning codecs.

One of the goals of the Open DML group is to develop APIs (application programming interfaces) for standardized control of production devices such as switchers and DVEs. The stated goal of the Open DML players is to make WinTel the platform of choice for digital video applications. Before that, how-

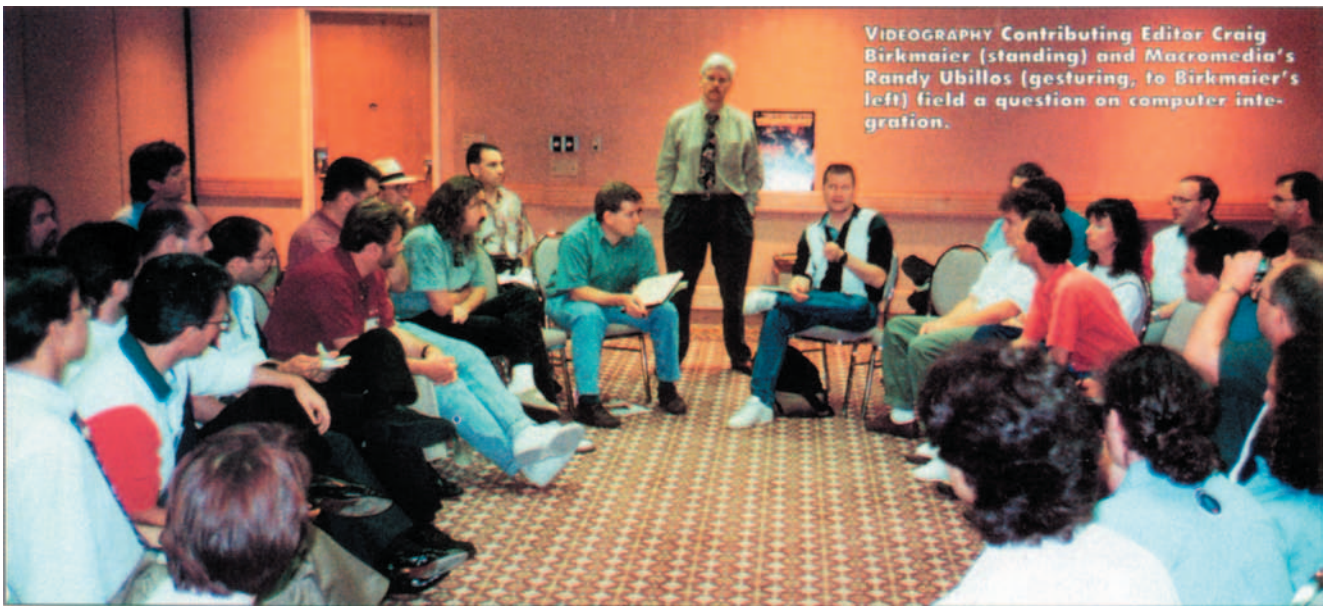

VIDEOGRAPHY Contributing Editor Craig Birkmaier (standing) and Macromedia's Randy Ubillos (gesturing, to Birkmaier's left) field a question on computer integration.

ever, someone pointed out during the Roundtable, they'll have to get Video for Windows to work.

Many developers consider QuickTime for Windows to be a better solution, particularly because it offers cross-platform support. Apple intends to keep QuickTime competitive on both platforms.

QuickTime and Other Issues Mac

The discussion of codec compatibility in open DML eventually lead to talk of the same issue in QuickTime. Currently, M-JPEG systems are virtually incompatible. The Targa 2000 provides the only exception to this, with its ability to read Radius VideoVision material in real time with lower data rate files. Apparently, developers of the various JPEG compression products could make their systems compatible if they could strip off the "wrappers" placed around the data in their proprietary formats. The industry could benefit from some form of software transcoder between the formats. Some called for codec developers and Apple to explore direct reads and writes of files of competing systems. This would reduce the complexity for users trying to share compressed video files in multisystem environments.

Real-world applications of new computer technologies typically fall short of their theoretical limits.

PCI-bus equipped Macs will soon profoundly impact digital video users. Bus speed probably represents the greatest limiting factor for digital video systems on personal computers. Moving data into and out of the CPU—and among various cards in a computer—has restricted the speed of the data-intensive processing required in digital video. Frequently, data movement represents more time than processing the effect in the CPU.

Real-world applications of new computer technologies typically fall short of their theoretical limits, but the PCI bus will bring tremendous data rate benefits. Developers at the Roundtable reported that the current PCI bus Power Macss will load data into memory about three times faster than their previous NuBus counterparts. Eventually the standard should lead to real-time uncompressed systems, but initially it will enable compressed systems to handle higher data rates with greater ease. Users will no longer need to press the envelope just to get adequate video performance. Expect products to work more reliably and for more users to be able to successfully build systems.

Before everyone begins ditching

those NuBus Macintosh systems, however, remember that any new technology must be perfected. Migration to the PCI bus involves major change. A relatively poky NuBus system that works today may allow you to be more productive than a PCI system that you spend the next few months tweaking.

Some Roundtable participants expressed trepidation regarding the entrance into the market spawned by the PCI bus. Will these developers—many of which have never developed for the Macintosh OS—provide the kind of user experience to which Mac users are accustomed? Today, virtually all Mac peripherals install without users setting dip switches or having to endure long, arduous installations. WinTel systems can be brutal to users trying to install peripherals. (See "The Trouble With PCs," in 2/95).

On the positive side, however, competition and economies of scale are leading to prices on PCI products that are a fraction of the cost of their poorer performing NuBus equivalents. Whether the technology will usher in revolutionary change or merely evolutionary change depends on your perspective, but all seem to agree that digital video systems will soon achieve a new performance plateau.

Copland *is the code name for Apple's upcoming operating system release.*

One issue facing users trying to do long-form work on digital video systems regards the volume size limit of 4 GB and the file size limit of 2 GB. A user of the highest data rate systems, can record only 5 to 8 minutes of high-quality video into a 2 GB file, without audio. That's fine for commercial production, but what if you're editing a one hour documentary and you need to digitize a ten-minute interview? When a recording reaches the 2 GB file limit the system reports that the drive is full. Playback, on the other hand, works now. A user can sequence files of various lengths together and play them back without interruption.

One Roundtable participant suggested that developers produce a fix that could link files in digitize mode to use multiple partitions. The question that many developers ask regarding these and other issues, is "Should I develop a fix now or wait for Copland?"

Copland is the code name for Apple's upcoming operating system release, due in the middle of 1996. This new OS will reportedly address the problem by increasing volume size limits to 256 TB (terabytes), which equals 256,000 GBs. It seems safe to assume that Apple will sub-

stantially raise file size limits as well.

The Gauntlet Thrown: Video Streams and Effects Acceleration

The issue of effects acceleration—a hot item of discussion last April during the NAB Roundtable—again generated controversy. Participants were divided into two camps: the real-time, multiple-video-stream people and the linear processing, single-stream people. Both obvisiously want improved performance in the production of effects, but their approaches differ.

The multistream proponents believe that systems should be able to generate two or more video streams and that dedicated special-purpose chip sets could then perform effects in real time on these streams. The disadvantages to this approach include a higher system cost for multiple codecs and arrayed drives; it also limits the effects accelerator to processing video-resolution frames. Additionally, with more complex effects it could require multiple passes through the codec that will increase the likelihood of compression artifacts.

A chip set designed to process a fixed resolution in real time will provide little or no benefit when producing some effects.

A chip set designed to process a fixed resolution in real time will provide little or no benefit when producing some effects. Adobe After Effects users will often design a composition with high-resolution images. For example, consider a project that called for an effect where the viewer's point of view zoomed from a wide shot of a note on a table, into a full-screen close up of the signature on the note. The After Effects artist could create an oversized note in which the signature fills the screen at the zoomed position. The note image might require a 2,000 by 3,000 pixel resolution to make that move and appear sharp at the final framing. Set resolution chips would offer little or no benefit accelerating this high-res rendering, and would similarly limit feature film and high definition media production.

After Effects users also routinely compose with several layers of video. Many report using ten to 20 layers on a regular basis. Even if resolution was not a problem, only

being able to add one or two effects in a pass would lead to artifacts as the user repeatedly ran material through the compression system. -- I've seen problems with Digital Betacam material routed through a component analog switcher with as few as eigth generations. Image content significantly influences the likelihood of these problems, but, even if it occurs only occasionally, for many it is unacceptable.

Clearly, not all video producers create the way After Effects users do, but industry trends favor trhs style. For other users, however, the multistream approach would provide large efficiency gains over today's more common "non-real-time rendering" approach. The news or documentary producer that requires only simple transitions and text overlays might be able to avoid effect rendering altogether. This producer might rarely use multiple layers, so compression problems would be less common. This style of production does represent the vast majority of program material. Therefore, within its limits, a multistream system can provide efficiency and speed.

The alternate camp—the linear processing, single-stream proponents—believe that flexibility and long-term vision should prevail over immediate speed gains. Speed stands as the one obstacle to the single-stream, linear processing approach, but none of the other limitations apply. This limitation diminishes as hardware gets faster and incorporates multiple processors.

The linear processing model offers only one video stream in and out, but includes more flexible effect acceleration hardware. The After Effects style of production allows the layering of a nearly unlimited number of images at well beyond TV resolution. This approach is unlikely to produce compression problems because the system loads each of the layers required for a frame into memory once, it composites all layers, and then writes the rendered image to disk. No effect should require multiple passes through the codec.

Today's personal computer systems can not load five or ten image layers into memory and render them in under one thirtieth of a second. You can't build a real-time linear processing system affordably now, but we are approaching that day rapidly. The linear processing proponents would rather see manufacturers focus on this model and make do with linear processing systems that are as fast as possible today, systems that build on and perfect the linear processing model. Designers could achieve big speed gains through this approach, that would be adequate for many users. Some linear processing systems today can quickly produce many of the effects required for mainstream video programming, such as dissolves and text overlays. Obviously most producers can't wait ten minutes to preview a two-second dissolve, but can they wait 20 seconds? Is three seconds too long? Three seconds isn't real time, but is it fast enough?

Perhaps tech support people should get more credit for their difficult jobs.

Of course, both approaches have a place in today's market and despite what anyone says, development will continue on both. Any user would welcome a tool that allows more productivity regardless of the design model used by its developer.

An intermediate approach or a hybrid system, while possible, complicates a programer's job as he or she tries to program for both models at the same time. As a developer, the issue you face is whether to take the long or short view. A multistream system design loses merit as processor speeds increase. To paraphase Field of Dreams, if you build an affordable multistream system today, they will come—but for how long? The window of opportunity could be narrow for those beginning today to develop a multistream system, especially if it takes one or two years to bring to market.

Stability VS Progress OR Till Copland Comes

The mere mention of system stability opened up wounds, as Roundtable developers and users alike provided emotional testimonials of their personal system horror stories. Many in the group felt that systems were becoming too complicated and that Apple had to take steps to improve things. Tim Schaaff, a member of the QuickTime team and the company's representative at the Roundtable, admirably endured the complaints of many participants. Schaaff listened with interest to the criticism and agreed that others at Apple need to hear those same comments. He pointed out that Apple was addressing many of the stability and performance issues with Copland. He also suggested that the press was one of the best vehicles for getting such a message to Apple.

Concerning such third-party products such as video boards, complaints were registered back and forth. Users complained that technical support was often inadequate and that vendors could do a better job of documenting known problems. Technical people representing manufacturers described problems caused by some users who, for instance, sometimes report "bugs" that were actually caused by operator error. These users often discover their mistakes later on but fail to notify the technical support staff. Meanwhile the techs are working to reproduce the problem. Then there's the user that calls repeatedly for help locating a problem caused by some buggy extension that keeps "creeping" back into his or her system. Perhaps tech support staffs should get more credit for their difficult jobs.

Finally, all Roundtable participants agreed that the most useful tool of all is effective communications. Many sited valuable online resources, such as digital video sections on American Online. Others called for more and better press reviews of video products, and that is, of course, one of the reasons why VIDEOGRAPHY exist.

Special credit for this year's Roundtable goes not only to all the participants that showed up to make it an important event, but also to Craig Brikmaier (who helped me recall and respond to the issues raised that evening) and to Macromedia's Randy Ubillos, who once again did a fabulous job as moderator. □