DIGITAL Technical Journal

Updated: 11 April 1996

FOREWORD

Robert Ulichney
Senior Consulting Engineer
Research and Advanced Development,
Cambridge Research Lab

"Can you dig it ... New York State Throughway's closed, Man. Far out, Man," announced a young Arlo Guthrie in the vernacular on the stage at Woodstock in 1969. Reading these words may evoke a mental picture of the event, but it sure is a lot more fun to hear and see Arlo deliver this message. Audio and video technology is the featured theme of this issue of the *Digital Technical Journal*.

Four years *before* Arlo's traffic report, in the year that a young Digital Equipment Corporation introduced the PDP-8, an interesting forecast was made. Gordon Moore, who was yet to co-found Intel, asserted in a little-noticed paper that the power and complexity of the silicon chip would double every year (later revised to every 18 months). This prediction has been generally accurate for 30 years and is today one of the most celebrated and remarkable "laws" of the computer industry.

While we enjoyed this exponential hardware ride, there was always some question about the ability of applications and software to keep up. If anything, the opposite is true. Software has been described as a gas that immediately fills the expanding envelope of hardware. Ever since the hardware envelope became large enough to begin to accommodate crude forms of audio and video, the pressure of the software gas has been great indeed. Digitized audio and video represent enormous amounts of data and stress the capacities of real-time processing and transmission systems.

Digital has participated in expanding the envelope and in filling it; its hardware performance is record-breaking and its audio and video technologies are state-of-the-art. Looking specifically at the four categories into which computer companies segment audio and video technologies, Digital is making contributions in each of these: analysis, synthesis, compression, and input/output.

MIT's Nicholas Negroponte believes that practical analysis, or interpretation, of digitized audio and video will be the next big advance in the computer industry, where nothing has changed in human input (keyboard and pointing device) since, well, the Woodstock era. Digital is actively investigating methods for speaker-independent speech recognition and, in the area of video analysis, means to automatically detect, track, and recognize people.

The synthesis of still and motion video, more commonly referred to as computer graphics, has traditionally been a much larger area of focus than the handling of sampled video. Synthesis of

audio, or text-to-speech conversion, is the topic of one of the papers in this issue; DECtalk is largely considered to be the best such synthesis mechanism available.

When audio or video data are represented symbolically, as is the case after analysis, or prior to synthesis, a most efficient form of compression is implicitly employed. However, the task of storing or transmitting the raw digitized signal can be overwhelming, especially at high sampling rates. Compression techniques are relied upon to ease the volume of this data in two ways: (1) reducing statistical redundancy, and (2) pruning data that will not be noticed by exploiting what is known about human perceptual systems. In this climate of interoperability and open systems, Digital recognizes the importance of adhering to accepted standards for audio and video compression versus the promotion of some proprietary representation.

The last category is that of I/O. Audio and video input require a means for signal acquisition and analog-to-digital conversion. The focus here is on preserving the integrity of the signal as opposed to interpreting the data. Proper rendering is needed for good-quality output, along with digital-to-analog conversion. For both audio and video, trade-offs must be made to accommodate the highest degree of sampling resolution in time and amplitude.

Digital is a leader in the area of video rendering with our AccuVideo technology, aspects of which are described in part in three papers in this issue. Video rendering incorporates all processing that is required to tailor video to a particular target display. This includes scaling and filtering, color adjustment, dithering, and color-space conversion from video's luminance-chrominance representation to RGB. In its most general form, Digital's rendering technology will optimize display quality given *any* number of available colors.

The earliest form of AccuVideo appeared in a 1989 testbed, known internally as Pictor. This led to the widely distributed research prototype called Jvideo in 1991. Jvideo was a TURBOchannel bus option with JPEG compression and decompression and was the first prototype to combine dithering with color-space conversion. Jvideo was the basis for design of the Sound & Motion J300 product, which included a remarkably improved dither method. A follow-on to J300 is a PCI-bus version called FullVideo Supreme.

In products that render RGB data instead of video, Digital's rendering technology is referred to as AccuLook; except for this one difference, the rest of the rendering pipeline is identical to AccuVideo. AccuLook products include graphics options for workstations: ZLX-E (SFB+) designed for the TURBOchannel and ZLXp-E (TGA) designed as an entry-level product for the PCI bus.

AccuVideo rendering is a key feature in the DECchip 21130 PC graphics chip and in the TGA2 high-end workstation graphics chip. While noted for its high image quality, AccuVideo is also efficiently implemented in software; it is available as part of a tool kit with every Digital UNIX, OpenVMS, and Windows NT platform.

With Moore's law on the loose, it can be argued that hardware implementations of video rendering are not justified as software-only versions grow in speed. Although today's processors can indeed handle the playback of video by both decompressing and rendering at a quarter of full size, little is left for doing anything else. Moreover, users will want to scale up the display sizes, and perhaps add multiple video streams -- and still be able to use their processors to do other things. For the near term, hardware video rendering is justified.

The five papers that make up the audio and video technology theme of this issue are but a small sampling of the work under way in this area at Digital; look for more papers to follow in subsequent issues of this *Journal*. As the audio and video gas continues to fill the ever-expanding hardware envelope, we look forward to an enriched and more natural experience with computing devices. Arlo's Woodstock pals would likely agree that this sounds like more fun.