DOCUMENT RESUME

ED 388 212 IR 017 361

AUTHOR Booth, Kellogg S.
TITLE Is There Computer Graphics after Multimedia?
PUB DATE 94
NOTE 7p.; In: Educational Multimedia and Hypermedia, 1994. Proceedings of ED-MEDIA 94--World Conference on Educational Multimedia and Hypermedia (Vancouver, British Columbia, Canada, June 25-30, 1994); see IR 017 359.
PUB TYPE Reports - Evaluative/Feasibility (142) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Computer Graphics; Computer Networks; Foreign Countries; *Futures (of Society); *Multimedia Materials; *Research and Development; Screen Design (Computers); *Technological Advancement

ABSTRACT

Computer graphics has been driven by the desire to generate real-time imagery subject to constraints imposed by the human visual system. The future of computer graphics, when off-the-shelf systems have full multimedia capability and when standard computing engines render imagery faster than real-time, remains to be seen. A dedicated pipeline for graphics will be redundant for all but the most demanding applications; imagery available today only on expensive systems will be supported by standard components. Deficiencies in spatial resolution for current head-mounted displays are one obstacle. However, with predictions of 4K X 4K and even 8K X 8K, it seems that most applications will not be limited by the spatial resolution of the screen once another one or two factors of two are achieved in the number of pixels per inch on a screen. With respect to chromatic resolution (number of bits per pixel used to represent color), research is within one or two factors of two from the ultimate chromatic resolution imagined as being necessary for the human visual system. Temporal resolution comes in two forms: refresh rate and update rate. The future of refresh rate is likely to be tied to consumer video, so the one or two factors of two may be a while in coming. Update rate is mostly a question of memory, transmission bandwidth, and computing power, and will increase almost automatically with the general advance of computing technology. With many of the hardware problems close to being solved and the promise of widespread multimedia applications likely to bring costs down, will there be any difference between computer graphics and multimedia? It is proposed that the difference will be significant, and computer graphics will be relegated to the back seat. The multimedia pipeline of the future will involve computing of the images ahead of time (maybe 1/240 of a second before needed), compression using high-speed circuitry, moving the images to secondary memory, fetching them all back when needed (1/240 of a second later), decompression, then pasting them onto the screen. Networking provides access to moving imagery, and largely eliminates the distinction between real-time and pre-computer imagery. So where will this leave computer graphics? It will still be there, but mostly as a producer of imagery on the network and as a tool for augmenting imagery obtained from other sources. (Contains 17 references.) (MAS)

Is There Computer Graphics After Multimedia?

KELLOGG S. BOOTH
Department of Computer Science
The University of British Columbia, Vancouver, British Columbia, V6T 1Z4, Canada

Abstract: Computer graphics has been driven by the desire to generate real-time imagery subject to constraints imposed by the human visual system. Modern workstations employ special-purpose hardware and software in a "graphics pipeline" to satisfy these constraints. We expect that the next generation of hardware and software may meet the requirements of the human visual system and that subsequent generations will vastly exceed them. Multimedia faces the same challenges, but only for pre-existing imagery because image generation takes place outside the multimedia workstation. It will be interesting to see what becomes of computer graphics when off-the-shelf systems have full multimedia capability and when standard computing engines render imagery faster than real-time. A dedicated pipeline for graphics will be redundant for all but the most demanding applications; imagery available today only on expensive systems will be supported by standard components. This talk will review the technical obstacles to achieving this vision and the impact that we can expect to see in educational uses of computer graphics after the research challenges we face today are met.

Many of the revolutions that have swept the world of computing have promised to usher in a new era in education. None have. There's a reason for this, but it's one educators don't like. The reason is that education is not an end in itself and thus education never leads in the adoption of new technology. Computers are in schools now because the world (and this means the business world) uses computers, so students have to be trained in the use of computers in order to function in the business world. Despite many noble efforts by the education community, realistically it is hard to point to any significant impact of computers in education other than small, isolated successes that are often the result of substantial investments of time, money and good will on the part of educators and the computing community.

A particularly good example of this phenomenon is computer graphics. Who could doubt the potential for enhancing a curriculum by using interactive graphics to not only bring a possibly dull subject to life but to add 3-D realism and maybe even a glimpse of virtual reality thrown in for good measure? But it rarely happens. The reasons, as we all know, are two-fold.

The first reason that computer graphics has not had a major impact in education is that despite the many advances that have been made in the field of computer graphics, it remains a fact that very few professionals in any field have access to 3-D graphics workstations and, while a great many people use 2-D computer graphics, the computers generally available to schools do not support the functionality necessary to run many of the applications representative of the state-of-the-art in computer graphics. With time, this may change. Costs are coming down, so as computer graphics comes to have a greater impact on everyday life we might expect, just as typewriters found there way into schools when typing became a part of everyday life in the business world, so too will computer graphics be part of the curriculum once it is adopted in the workplace. (Eventually even electric typewriters made it into the classroom.)

The second reason that computer graphics has not had a major impact in education is less likely to be resolved so easily. This reason is related to the difficulties associated with developing and maintaining courseware. To continue the analogy with typewriters, we can observe that by putting a typewriter into a classroom students are instantly enabled to use the typewriter with perhaps only the need of an introductory text on typing skills. No other expenditure is required. An educator trained in typing can probably do a credible job of bringing beginners to an intermediate level of typing skill. There is an immediate payoff. Students are then able to type the papers they write for school (helping both them and their teachers) and they gain a practical skill for their future life in the world of business.

Not so with computers. Courseware is often quite different from the programs used in business (unlike typewriters, which are used identically in schools and in business). And courseware costs a lot to develop, test and maintain. Courseware that utilizes computer graphics is often much harder to write than normal courseware. Ironically, the advent of GUIs (graphical user interfaces), which were invented to make computers more accessible to non-programmers, make the programming of user interfaces much more difficult than it ever was before and thus, to a large extent, make good courseware even more difficult to produce.

Multimedia is the latest darling of the computing community that many predict will change the way that education is done. Why should we believe this will happen any more for multimedia than it did for computer graphics or for the many other areas of computer science that did not, in the end, have an impact on education?

The remainder of this talk will address this question by first discussing the relationship between computer graphics and multimedia and then by pointing out the differences that are likely to be significant factors leading to the ultimate success of multimedia as an educational tool where computer graphics and others have demonstrably failed. Chief among these is the distinction that will be drawn between multimedia as a means for communication and computer graphics as a means for computation.

A (Very) Brief History of Computer Graphics

For three decades research in computer graphics has had as one of its main concerns the development of hardware and software to accelerate the so-called "graphics pipeline" that supports interactive real-time display of computer-generated imagery (Foley, et al., 1990, section 18.3). Many architectural advances were made to facilitate high-performance 2-D and 3-D graphics applications in specialized domains (Akeley & Jermoluk, 1988; Haeberli & Akeley, 1990). About a decade ago this effort merged with mainstream computing when 2-D and 3-D graphics workstations became the platforms of choice even for applications that often were not primarily graphical in nature but which took advantage of features common to graphics workstations to enhance productivity (IEEE, 1985; IEEE, 1988). Today most components of the graphics pipeline are available as standard features or modest upgrades in common computing environments and the distinctions between personal computers, graphics workstations and traditional mainframe computers have become blurred.

Many of the performance limitations that currently exist for high-quality graphics workstations are related to human perception (screen resolution, color resolution, refresh rate, and update rate). These are about to become non-issues during the next decade when the multimedia revolution eliminates most of the remaining distinctions between interactive real-time graphics and "normal" computing. Every computer will be accessed as a multimedia workstation because almost all user interfaces will employ some aspect of multimedia. When this happens, computer graphics research will become a very specialized area concerned with rather arcane aspects of modeling and rendering whose solutions will be achieved primarily with cleverer, more accurate algorithms that require more computing power, rather than with new architectural breakthroughs.

The graphics pipeline will be replaced with a ubiquitous "multimedia pipeline" that will eliminate the distinctions between real-time and pre-computed imagery, between local and remote imagery, and between real and synthetic imagery.

The Perceptual Bottleneck

We are small factors of two away from meeting most of our goals for computer graphics displays if we simply look at the level of performance currently available and compare it to the strictest requirements we expect to encounter in typical applications. For the most part, these requirements are based on properties of the human visual system that are studied in psychology and related fields. Many of the limitations are discussed in detail in a recent survey article on virtual reality (Ellis, 1991). We will quickly review the ones that relate to multimedia.

Deering has commented on the deficiencies in spatial resolution for current head-mounted displays by pointing out that if this were our regular vision, we would be legally blind (Deering, 1992). But for normal display monitors, the situation is much better. Screen resolution ranges from roughly 512 × 512 for anyone who cares to invest in a graphics board for a PC to almost 2K × 2K resolution for high performance workstations. One sees predictions of 4K × 4K and even 8K × 8K, but it seems that most applications will not be limited by the spatial resolution of a display once another one or two factors of two are achieved in the number of pixels per inch on a screen. Dreams of larger screens will always be with us, but there is a practical limit to the physical size of the screens we expect to use on an everyday basis.

Chromatic resolution is, roughly speaking, the number of bits per pixel used to represent color information. On low-end machines, 8 bits per pixel is still not standard, but this is changing. Even a few years ago it was rare on a PC but it is now becoming more common. Frame buffers have had 24 bits per pixel (8 each for red, green and blue) for many years with 10 and even 12 bits per pixel not uncommon. Certainly 64 bits per pixel (16 each for red, green, blue and an alpha or opacity channel) is adequate given the properties of the human visual system (Levinthal & Porter, 1985), and 128 bits (32 bits per channel) seems overly generous. So again we are within one or two factors of two from the ultimate chromatic resolution we might imagine as being necessary.

Temporal resolution comes in two forms: refresh rate, the number of times per second that a CRT screen is painted, and update rate, the number of times per second that the image being painted is modified (Baecker, 1979). A rule of thumb is that refresh must be above 40 times per second (television is either 60 or 50, depending on whether you live in North America or not, and film is 48 except for a few special projection systems). Lighting conditions and other environmental factors may increase the required update rate to perhaps 120 times per second, but even assuming 240 times per second (which might be necessary for stereo displays), the "one or two factors of two" assumption applies here as well.

In the flight simulator community, the required update rate is generally assumed to be 10-20 times per second (television is 30, film is 24, and Saturday morning cartoons are 15 or less). In virtual reality, head-coupled displays may have somewhat stricter requirements (Deering, 1992; Ellis, 1991). But at worst the update rate needs to be the same as the refresh rate, but the two are solving different problems (refresh rate governs persistence of vision, whereas update rate provides an illusion of continuous motion) and hence there is a distinction.

The fact that all of the performance figures for actual displays are not too different from the upper bounds that are cited above is not coincidence. Most of the current numbers exist because they are a good compromise between image realism and economic realism. So which ones are likely to change most rapidly?

HDTV is about to double or triple the spatial resolution of consumer video products. This will no doubt impact workstation displays, at least by reducing costs (due to economies of scale) and perhaps by increasing resolution as well.

Chromatic resolution is really an issue of more bits of storage (and transmission bandwidth) per pixel plus the costs of digital-to-analog converters (DACs). Consumer televisions are not yet digital. So they don't have DACs. When they are digital, we can expect to see prices drop a lot for DACs and no doubt we will see higher resolution ones at affordable prices. Today, 8-bit DACs are standard computer components and 10-bit and 12-bit DACs are used for high-performance systems.

Refresh rate is likely to be tied strongly to consumer video. So the one or two factors of two may be quite a while in coming. But if the often-predicted shift away from CRTs finally does take place, whatever replaces CRTs (liquid crystal displays?) may provide higher refresh rates and more spatial resolution too.

Update rate is mostly a question of memory, transmission bandwidth and computing power. It will increase almost automatically with the general advance of computing technology.

All of these powers of two compound to maybe a factor of 50 to 100 increase in performance. This is a lot, but it is small compared to the changes in the computer industry achieved over a typical decade. Moreover, once the limitations are overcome, they are likely to be entirely over-run (at least those achieved through raw computing power) because the technological advances are being driven by other forces that will not be satisfied.

Something that is important for multimedia, where images are generated separately from when they are displayed, but not for computer graphics, is compression. Computer graphics is often not networked because the real-time requirements for image generation are more easily met when there is a tight coupling, through the graphics pipeline, of the data structures representing the underlying model and the final representation of the image as a raster array of pixels.

Some early line drawing systems did explore the separation of modeling and rendering across a local network, with transmission bandwidth being the limiting factor that determined the distribution of function between the modeler and the renderer (van Dam, Stabler, & Harrington, 1974). The "wheel of re-incarnation" is a phenomenon in computer graphics that happens because we need to generate the images close to where we view them and we need to use special purpose hardware (Myer & Sutherland, 1968): as display processors increase in complexity there is a tendency to split the work and distribute it across some communication channel between a main processor and a satellite processor, then the satellite processor is made more powerful to gain performance until it is itself a critical resource, and so another processor is spawned off (thus completing another circuit around the wheel). Until recently, network transmission bandwidth was inadequate for this division of labor with raster images. Higher-speed networks and better image compression algorithms are starting to tip the balance again in favor of distributed systems.

Where does this leave computer graphics? With many of the hardware problems close to being solved and the promise of widespread multimedia applications likely to bring costs down, will there be any difference between computer graphics and multimedia? I think so. Moreover, I think the difference is significant and one that will in the end relegate computer graphics to the back seat.

The Multimedia Pipeline

The much touted Information Superhighway is a mechanism for transmitting multimedia data via universally accessible networks. There are few applications that will require the bandwidth being planned for unless they are using accessing multimedia documents. So we can expect to see a situation that already exists for many high-end multimedia users. Typical workstations will both send and receive highly compressed imagery over a network utilizing specialized components designed to support multimedia applications.

If the workstation is generating imagery, there will need to be a data path that starts where the traditional graphics pipeline started (with a data structure describing a scene to be rendered, usually a dynamic scene whose content and viewing parameters change with time) and ending up on the network feeding the compressed representation of the rendered scene to other workstations or archival storage servers. Similarly there will need to be a data path that accepts compressed imagery from the network and puts it onto the screen for when the workstation is displaying imagery.

Why go to the trouble of compressing, transmitting, and decompressing the image? Why not send the description of the scene and let the remote workstation generate the image? After all, we argued above that future workstations will be able to generate imagery faster than real-time for most applications. The answer is that the scene description will be bigger than the image! We noted above that current image sizes of $1K \times 1K \times 24$ bits are within striking distance of what we may need in the future. Yet scene descriptions grow without bound and, much worse, they are idiosyncratic in terms of the primitives they support and the auxiliary information they require (textures, bi-directional reflection functions, etc.). Standardization is hopeless at this time and for the foreseeable future. But raster images, even with arcane compression schemes, are easily described and, at least in principle, easily translated into various formats.

Example: PostScript is the standard way to ship text and images, regardless of how they were produced. Multimedia formats can be expected to play this same role, hiding many of the details behind how an image was produced. A related point is that PostScript delays the final binding of an image to the physical device on which it will be displayed (or printed), but much of the "value" of the program that prepared the image is added before it is translated into PostScript. Future multimedia servers may in fact be selling not the images (or streams of images) they produce, but the technology that produced them, much like Hollywood sells the movies that go on to video casettes for home viewing.

What will we do with all of the multimedia images that arrive over the network? We'll paste them into windows. Which means we will need to integrate hardware and software support for common windowing operations into the multimedia pipeline. Compositing functions will be common place; every workstation will have the equivalent of a video special effects box.

So why will we need to have a graphics pipeline any more? Just compute the images ahead of time (like maybe 1/240 of a second before you need them), compress them using high-speed circuitry designed for the task, move them to secondary memory (to save precious main memory), then fetch them all back when they are needed (1/240 of a second later), decompress them, and finally paste them onto the screen blended with whatever else is being looked at. Sound wasteful? Sure. But the hardware will be (almost) free and otherwise idle, so why not use it. And it will be cheaper than specialized graphics hardware because that will have a very limited market compared to multimedia hardware.

So we again ask, where is computer graphics in all of this? There will certainly remain problems at the frontier of computer graphics. There will always be imagery that is difficult to produce and research will be needed to determine the best techniques. But this will be of little concern to the average user. A second comparison will illustrate this point.

Example: Numerical analysis is not concerned at all with the design of floating point processors, just the standards employed in the representation and the conventions for exception handling and rounding. Instead numerical analysis concentrates on the fundamental algorithms. It has largely returned to its roots in mathematics, leaving the hardware details to computer architects whose concern is mostly speed and not accuracy. Computer graphics will be the same, concentrating on algorithmic aspects of imagery and ignoring many of the hardware details that have previously dominated many of the discussions of new techniques.

Whither Computer Graphics?

Hermann Maurer has argued convincingly that multimedia fills a void resulting from what he calls "the missing organ" (Maurer, 1992). Humans have a highly developed vision system, but little or no facility for creating imagery without technical aids. This is in significant contrast to sound, where humans have relatively equal facility for producing and hearing sounds as a means of communication among themselves.

Some of you probably grew up in the 50s and the 60s as did I. In the U.S., at least, Life Magazine played an interesting role in education. Whatever you thought of it, Life was a ubiquitous source of imagery. The pictures that Life provided were cut-and-pasted onto the classroom bulletin board of every classroom I was in until high school. The pictures arrived in the mail every week. Multiple copies, as many as you wanted. If you needed more, you could just ask a neighbor or a friend; lots of people had subscriptions and garages and attics provided a treasure trove of back issues. The pictures were almost literally free, so there was never any problem cutting them up and using them whenever imagery was needed. And there was almost always a large collection of imagery available, so long as you had access to the "archive."

The promise of multimedia is a return to that Golden Age. We have had electronic cut-and-paste of still images for quite some time now, but not a lot of access to the images. Networking promises to solve this (for a price – but the price will be small once it works) and to provide access to moving imagery. This will largely eliminate the distinction between real-time imagery and pre-computer imagery; all images will be pre-computed even if only just a little before they are viewed. The distinction between local and remote imagery will be lost; all images will go through the same multimedia pipeline just as if they came across a network. And compression techniques will wipe out the distinction between real and synthetic imagery; the two will be indistinguishable and many images will in fact be a blend of the two.

So where will this leave computer graphics?

It will still be there, but mostly as a producer of imagery on the network and as a tool for augmenting imagery obtained from other sources, both to customize it (to enhance its relevance) and to re-style it (to improve its fit with other images). Indeed, future multimedia systems, like existing text processing systems, will concentrate as much on style as on content (Beach & Stone, 1985).

The irony may well be that the very phenomenon that knocks computer graphics from its throne as the hottest thing in computing, that relegates computer graphics to a relatively isolated area of computer science research, will also be the force that finally brings the long-awaited promise of interactive graphics to education.

The power of multimedia lies in its ability to provide a powerful authoring tool that lets us appropriate imagery for the purpose of communication. Unlike computer graphics, which lets us create imagery, multimedia allows us to manipulate imagery. There is a world of difference. While creating imagery sounds like a good idea at first, anyone who actually does it realizes that it is very time consuming even with the best tools. If our goal is communication, we are better off using existing imagery to communicate our ideas or, if that fails, modifying existing imagery to our needs.

At roughly the same time there were two important advances made in the field of computing: Ivan Sutherland's PhD thesis *Sketchpad*, which set much of the the research agenda for computer graphics for a number of decades (Sutherland, 1965), and Doug Englebart's *NLS* (oN-Line System), which set much of the research agenda for multimedia and hypermedia for a number of decades (Englebart & English, 1968). Of the two, I think that Englebart, and the earlier vision of Vannevar Bush (1945), will be the more important for education and humanity in general simply because computer graphics is in the end just a set of technical tools for producing imagery whereas multimedia is a set of tools for using imagery.

References

Akeley, K., & Jermoluk, T. (1988). High performance polygon rendering. *Computer Graphics*, 22(4), 239-246.

Baecker, R.M. (1979). Digital video display systems and dynamic graphics. *Computer Graphics*, 13(2), 48-56.

Beach, R.J., & Stone, M. (1984). Graphical Style – Towards high quality illustrations. *Computer Graphics*, 17(3), 127-135.

Bush, V. (1945). As we may think. *Atlantic Monthly*.

Deering, M. (1992). High resolution virtual reality. *Computer Graphics*, 26(2), 195-202.

Ellis, S.R. (1991). Nature and origins of virtual environments – A bibliographical essay. *Computing Systems in Engineering*, 2(4), 321-347.

Englebart, D.C., & English, W.K. (1968). A research center for augmenting human intellect. *Proceedings FJCC*, Washington, D.C.: Thompson Books.

Foley, J.D., van Dam, A., Feiner, S.K., & Hughes, J.F. (1990). *Computer Graphics: Principles and Practice (2nd edition)*. Reading, MA: Addison-Wesley.

Haeberli, P., & Akeley, K. (1990). The accumulation buffer: Hardware support for high-quality rendering. *Computer Graphics*, 24(4), 309-318.

IEEE (1985). *Proceedings of the 1st IEEE International Conference on Computer Workstations*. San Jose, California. Washington, D.C.: IEEE Computer Society Press.

IEEE (1988). *Proceedings of the 2nd IEEE Conference on Computer Workstations*. Santa Clara, California. Washington, D.C.: IEEE Computer Society Press.

Levinthal, A., & Porter, P. (1985). Chap – A SIMD graphics processor. *Computer Graphics*, 18(3), 77-82.

Maurer, H. (1992). Seminar in the Department of Computer Science at The University of British Columbia.

Myer, T.H., & Sutherland, I.E. (1968). On the design of display processors. *CACM* 11(6), 410-414.

Sutherland, I.E. (1963). Sketchpad: A man-machine graphical communication system. *Proceedings SJCC*, Baltimore, MD: Spartan Books.

van Dam, A., Stabler, G.M., & Harrington, R.J. (1974). *Proceedings of the IEEE*, 62(4), 483-492.

Watson, R. (1970). *Timesharing System Design Concepts*. New York, NY: McGraw-Hill

Acknowledgements

The preparation of this paper was supported by the Natural Sciences and Engineering Research Council of Canada and by the Media and Graphics Interdisciplinary Centre at the University of British Columbia.