The Future of 3D Graphics Technology:

Will the Movies Maintain Their Lead on the Desktop?

George R. Dicker

STS 145

March 18, 2003

Introduction:

Creating cinema-quality graphics on desktop computers has always been a dream, but in the years up to and including 2003, it has become a task within reach. A parallel between the video game industry and Hollywood has always existed, with Hollywood generally having the render-time and money to deliver technologies first. A movie may take years to render, but current graphics cards are getting close to delivering "movie-quality" graphics in real-time. This case study attempts to answer some questions about the history and future of dedicated three-dimensional computer graphics hardware. Will movie technology continue its lead on consumer level 3D, always raising the bar for what a "cinema quality" game would require or will the two become indistinguishable? The paper will focus on game technology but will also talk about movies as points of reference. For a desktop computer game to be interchangeable with a Hollywood creation, what takes hours for movies to render will need to be nearly instantaneous on a PC. The problem with the topic of fast moving technology is that it can easily be out of date in six months. To avoid that limitation, this paper will concentrate on the current state of the art technology, but assume that the next technological advancement will be drastically different and hard to predict. The paper will also include information about future technologies and their possible ramifications.

Early Graphics History:

The history of three-dimensional computer graphics begins before it became available to the public through the movies or later in video games. In the late sixties and

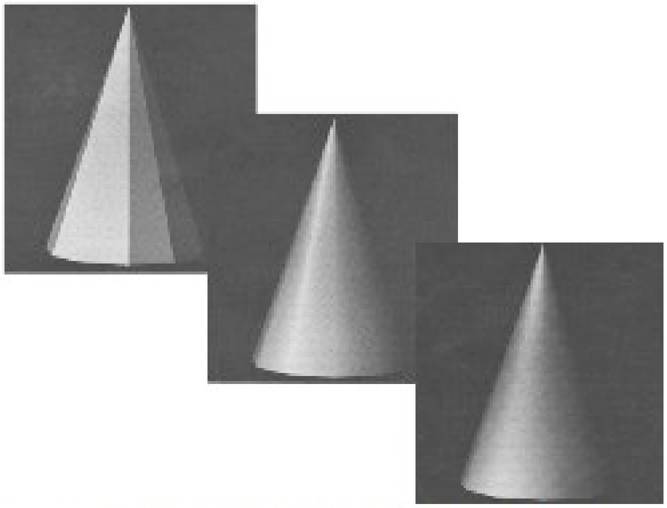
Facet, Phong, and Gouraud Shading (Carlson)

early seventies a number of small research projects began at government funded laboratories and universities. Most notably, the University of Utah founded its Computer Science Department in 1965 and started one of the very first Computer Graphics departments soon after. Over the years a number of key technologies and graphics players came out of this department. While at Utah in 1974, Ed Catmull, who would later go on to work at Lucasfilm and found Pixar, developed z-buffering and rendering of curved surfaces, technologies which were crucial in the quest for photorealism (Carlson). Around the same time and also at the University of Utah, Gouraud (1971) and Phong (1975) were inventing their shading methods that would make rendering of three dimensions possible. Later, Catmull, while still at Utah, would invent texture-mapping, allowing for much more complicated textures and pictures to be mapped and transformed around the surfaces of objects in a scene (Carlson). As a point of reference, this work in 3D computer graphics occurred at the same time Microsoft was

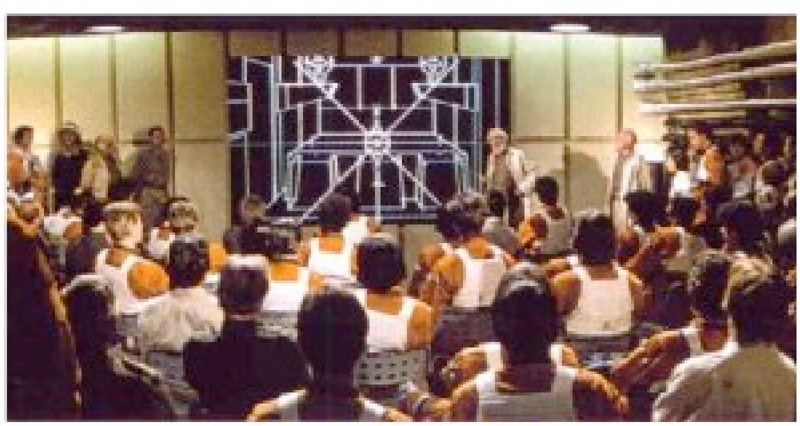
Model of Death Star from *Star Wars* (Carlson)

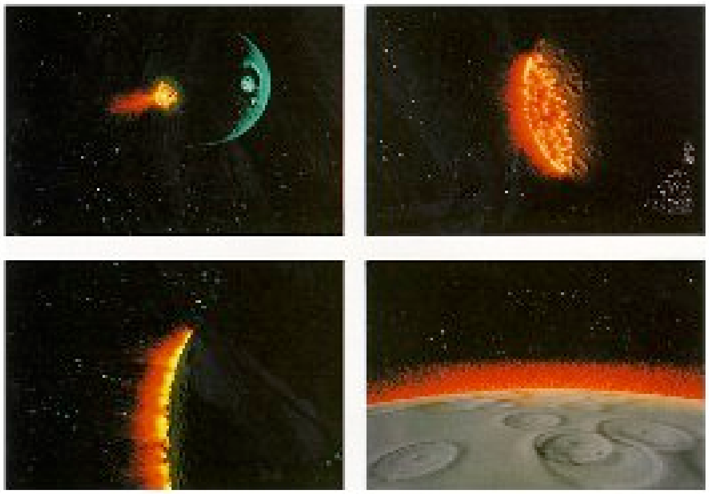

The Genesis Effect from *Star Trek II* (Carlson)

founded on the business of text-based computing in 1975. In 1977, the general public saw the first 3D generated scene in a movie with the wire-frame (meaning only outlines, and no shading) flight sequence in *Star Wars: A New Hope*. Lucas requested a full 3D test sequence of five X-Wing fighters flying in formation, but chose to use traditional models instead when he was not impressed with the rendering technology of the time (Yaeger). At the dawn of the eighties, people saw more complicated 3D graphics in *Star Trek II: The Wrath of Kahn* (1982), where Pixar used the first "particle rendering system to achieve its fiery effects" in the *Genesis* sequence (Yaeger). Eighties band Dire Straights soon opened the door for (cable subscribing) pop-culture with the 1984 video *Money for Nothing,* featuring "a low-detail, but very engaging animated CGI character" (Yaeger). Each of these rendering examples took days to complete and were done on huge computers with processors that were not specialized to handle the kinds of matrix and floating-point mathematics that graphics requires. In 1986, however, Intel came out with the *82786* graphics coprocessor chip and Texas Instruments introduced *TMS34010*

1984: *Money for Nothing* (mtv.com)

Graphics System Processor, beginning the revolution of dedicated 3D graphics hardware (Carlson). But 3D games on the home computer or video game console were still years away and required very expensive workstations from SUN Microsystems or Silicon Graphics (founded by James Clark, another University of Utah graduate) (Shoaff). In 1989 an alien with a water-like appearance that mimicked people was created for *The Abyss,* a long way from the effects used in *Star Trek II* or *Money for Nothing*. Both leading 3D design firms of the time bid on the project, and Industrial Light and Magic (ILM) beat out Pixar, but still used Pixar Software to create the alien (Farley, et al). The decade closed with 3D graphics beginning to bloom on the big screen but still quite far from the home computer.

The Nineties and the Dawn of Desktop 3D:

Starfox: the SuperFX chip (nintendo.com)

While SUN and Silicon graphics battled in the workstation market, very little outside of the high-end occurred in 3D graphics. Around 1994 very early 3D graphics started to make it to consumer-level video games. A game called *Starfox* for the Super Nintendo put a RISC number cruncher developed by Argonaut into the cartridge that normally only contained ROM to create a very low polygon environment for the game (pcparadox.com). After Sony set out alone to make PlayStation and the PC manufacturers caught site of it, they realized they were behind in the 3D realm and began to move fast. Early players included nVIDIA and ATI, but also a number of other companies including Matrox and 3Dfx (with their popular VooDoo line). The early nineties also saw an amazing in-flux of 3D graphics on screen. As the technology trickled down to the consumer level, it was being displaced at the top by more and more powerful (albeit expensive) hardware and software. In 1994 ILM brought dinosaurs to life for *Jurassic Park,* melding live-action backgrounds with animated herds of dinosaurs. Only a year later, in 1995, Pixar upped the ante with a fully 3D animated movie, *Toy Story*. *Toy Story* was 4 years in the making and was the first 3D animated feature length film (Farley, et al). 1994 also brought the first real first-person-shooter, John Carmack's *Doom*, which despite using the CPU instead of dedicated graphics chips to do the rendering rendering, was still the beginning of 3D gaming on the PC. As the nineties continued there was a clear focus on games and hardware moving into the third dimension. The technology was advancing but some people felt that it was actually a step back in game quality. Critics complained that "rendered 3D graphics will never be able to create the look some designers may want for the games" and challenged the industries devotion to 3D (Rouse). But 3D would not be a passing fad, and as quality soared critics gave in.


DOOM screenshot (game-research.com)

RADEON 9800 Pro demo movie (ati.com)

The Current Players (2003):

As the nineties went on, graphics cards improved and newer games similar to or based on Quake III allowed for fast paced 3D games with amazing graphics, but still fell short of even some early movies. At the beginning of 2003, however, graphics cards seem to be more powerful than any of the games that people play on them because "technology is outpacing demand" (Liaw). The only software to come close to using the full potential of the graphics card is the demo software from the vendors themselves. Since the nineties, the 3D graphics industry has almost completely been narrowed down to two major players: ATI and nVIDIA. With their top cards both vendors show they are striving to match Hollywood, and they each use the word *cinematic* to describe the cards, alluding to the ultimate goal. ATI calls the just-released Radeon 9800 Pro, "stable, reliable, highly optimized" and say it "represents the industry's only cinematic Visual Processing Unit (VPU) in its 2nd generation," while nVIDIA counters that "the GeForce FX delivers cinematic effects beyond imagination [with] its CineFX engine," which produces "gaming effects on par with the hottest motion pictures" (ati.com, nvidia.com). Each claim sounds amazing and the chips have the power to back the words up. Both of these companies have made great gains in technology and performance in the past couple of years through the introduction of vertex and pixel shaders, "pixel-by-pixel graphics programming that can dramatically increase the realism of animated images" (Becker). Starting around the time of the GeForce 3, in 2000, graphics chips gained the ability to "be programmed to do different graphics tasks, the same way a CPU can be programmed to run an application" (Patrizio). In an interview for this paper, Joey Liaw, formerly of Ion Storm, said, "Corrinne Yu once joked to me that, 3d cards are getting so ridiculously powerful that we're going back to 2d … we're back to the level of operating on the pixel level rather than the triangle level" (Liaw). The cards of 2003 are so fast, developers don't quite seem to know what to do with them.

GeForce FX Demo picture (nvidia.com)

The Console Market:

Currently, the console market is divided between the Sony PlayStation2, the Microsoft XBOX, and the Nintendo GameCube. Each of these consoles is about the same generation, and when games are made for more than one of them they look similar. The current generations of consoles are now behind both the desktop and Hollywood, judging by *The Lord of the Rings: The Two Towers* for PlayStation 2. There are three levels of decreasing quality of graphics. The game contains full-motion MPEG2 video from the movie, live-rendered, but pre-scripted cut-scenes, and live-rendered game play. Since consoles have to be able to stay on the market longer than a PC graphics card they have to be able to last a couple of generations. For this reason, the next generations of game consoles from Sony, Nintendo, and Microsoft are already in progress and talked about as the next big thing. According to industry experts, the PlayStation 3 will "feature a revolutionary architecture that will allow it to pack the processing power of a hundred of today's personal computers on a single chip" (Takahashi). The same chip that might also make its way into IBM servers "will be able to process a trillion math operations per second -- the equivalent of 100 Intel Pentium 4 chips and 1,000 times faster than processing power of the PS 2" and will presumably only cost about $199 when it comes out in 2005 (Takahashi).

Of course with the industry on the cusp of true cinematic graphics, everyone is trying to outdo each other, and Microsoft and Nintendo both plan to deliver their next generation systems in 2005. It is hard to compare things that have not come out yet, but most likely between now and 2005, both ATI and nVIDIA will keep up their pace and stay on par with SONY's prediction. Patrick Hanrahan, of Stanford University, claims that this trend will continue. Within recent years nVidia's speed has been increasing by 225% every year or "Moore's Law Cubed" (Hanrahan). Hanrahan goes on to claim that the gap in speed and efficiency between the central and graphics-processing unit will continue to increase. Not only are the hardware companies moving as fast as they can to stay on top of the other but it is rumored that "Microsoft will make improving the PC gaming experience one of the central appeals of Longhorn, the next major version of

Windows that's expected in about two years" (Parker). Microsoft is also taking advantage of the fact that "ATI, Nvidia, and other hardware companies have spelled out their road maps in enough detail that [DirectX 9.0] has previously unannounced support for all the next-generation graphics features" (Parker). The desktop technology of tomorrow is going to be "ridiculously powerful," but will still face challenges in catching up to Hollywood (Liaw).

Hollywood Effects in 2003:

The level of detail that current movies are able to attain comes from the amount of time and care put into each frame. Pixar claims that, "Each frame represents 1/24 of a second of screen time and takes about six hours to render, though some frames have taken as many as ninety hours" (pixar.com). In 2001, SquareSoft, a game company of Japan, finally came out with a movie version of their extremely popular game series, *The Final Fantasy*. Each three to five minute scene took one month of production time, and delivered more believable living creatures than any other movies before it (Farley, et al). Not surprisingly, SquareSoft actually used Pixar's software package *RenderMan*. Along with pure graphics power, for games to equal Hollywood standards they are going to have to keep up with the AI and mass-simulation. In *Massive*, a software package created for *The Lord of the Rings*, "every agent has its own choices and a complete brain" (Macavinta). Similarly large pieces of software have been written for other movies such as *Fitz*, created to simulate 3 millions hairs for the lead character of Pixar's 2001 movie, *Monsters Inc* (Clewley). The pure processing power to run these simulations is still years away from real-time on the desktop and will play out to Hollywood's advantage. In just one year, the difference between the character of Gollum in the first and second *Lord of the Rings* installments shows how quickly they are moving from a distant cameo to a fully digital and believable animated character.

Graphics Processing Beyond Games:

With the amazing power of today's graphics processors being unused while people are not playing games, and the gap getting larger all the time, many companies are striving to use the power of the GPU for other tasks. Apple Computer, who has been putting fast chips from ATI and nVIDIA in their machines for three years, has decided to help out the whole operating system with them. They have a technology called *Quartz Extreme* which they claim "uses the integrated OpenGL technology to convert each window into a texture, then sends it to the graphics card to render on screen" and "while other operating systems hope to introduce comparable technology in late 2004, [MacOS X 10.2] Jaguar has it now" (apple.com). Microsoft is also interested in harnessing the cycles of graphics cards not being used and along with technology similar to Apple's, the next major version of Windows, called Longhorn, "will likely include a

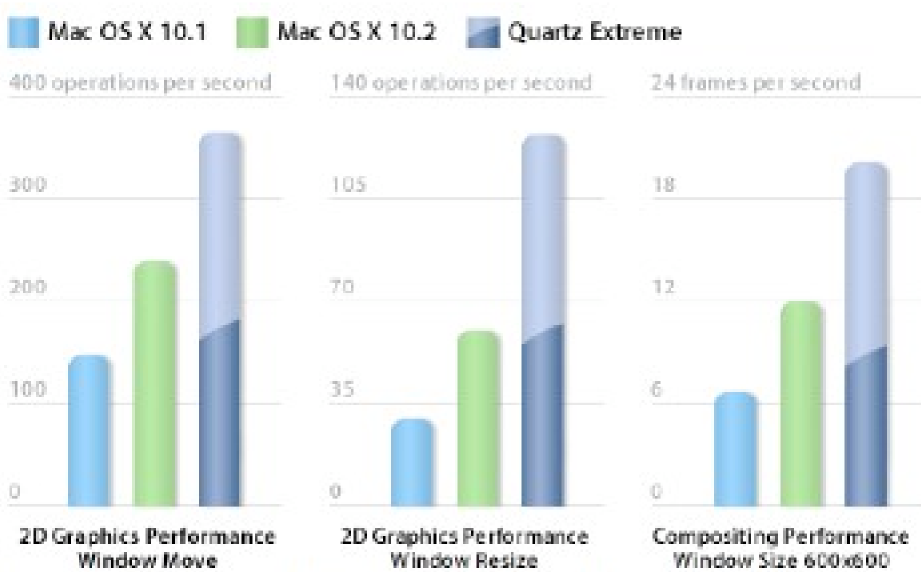

Speed improvements by offloading to GPU (apple.com) major overhaul in Window's visual presentation [and] may include 3D interface elements" (Parker). Some people have even turned to the power of 3D graphics coupled with game-engines to produce their own movies. A genre of movies has sprung up, not filmed in real life, but played out through the engine of a computer game. Customizations to the engine and graphics allow for diversity that Hollywood directors have trouble attaining (King). Machinima, as the genre of movies is called, came about because "video-game developers have long tried to integrate the sights and sounds of cinema into their games, with mixed results" (King). Even Machinima directors are not satisfied with the current crop of tools and "dozens of Machinima filmmakers anxiously await the release of the Doom III game, which will include the most advanced game engine yet," coupled with the next generation of hardware (King). Other researchers are using GPU's to do graphics never before possible in hardware. In 2002 researchers from Stanford University published a paper saying they had tricked graphics hardware into actually doing ray-tracing through the custom vertex and pixel-shading capabilities of current generations of hardware" (Purcell, et al). In their abstract they say "in the near future, the graphics pipeline is likely to evolve into a general programmable stream processor capable of more than

DOOM III screenshot (wired.com)

simply feed-forward triangle rendering" (Purcell, et al). They go onto "demonstrate that ray tracing on graphics hardware could prove to be faster than CPU based implementations as well as competitive with traditional hardware accelerated feed-forward triangle rendering" (Purcell, et al). If successful, this technology could shortcut the difference between a home PC and the movies.

Conclusion:

Before consumer-level machines had the power to render scenes in real-time, Hollywood was able to flex its muscles by modeling complicated scenes and then spending hours rendering them. By the time that a home-pc will be able to render a cinema-quality model in real-time, which both main vendors claim to be able to do in 2003, Hollywood pulls ahead with its more complicated rules and simulations. The question of whether or not a home computer will ever be able to achieve the exact same quality as the movies is difficult. It is the same as asking whether computers will ever be powerful enough to handle anything instantaneously. So far in the history of computational power, the more powerful the hardware gets, the software just gets more demanding and is able to achieve more things. The desktop computer user will certainly never be able to match the budget of a Hollywood movie, but with advances in 3D graphics hardware, the maturation or programmable hardware vertex and pixel-shaders, and the expansion of Machinima, the next big thing to come to the theatres will have to be revolutionary to stay ahead. But consumers have nothing to worry about, it will only be a short while before that new level of complexity, realism, or effects arrive on their desktop as well.

Bibliography:

Articles:

"Nvidia, ATI rev graphics engines "
David Becker
CNET News.com
http://news.com.com/2100-1043-991299.html

"A Critical History of Computer Graphics and Animation"
Wayne E. Carlson
Ohio State Univeristy
http://www.accad.ohio-state.edu/~waynec/history/timeline.html

"Monsters, Inc. Used Monster Tools"
Robin Clewley
Wired News
http://www.wired.com/news/digiwood/0,1412,48052,00.html

"Computer Graphics in Film"
Carolyn Farley, Joel Vogt, Michelle Wilson
The University of Newcastle, Australia
http://www.cs.newcastle.edu.au/~comp3320/history/animation/index.html

"Why is Graphics Hardware So Fast"
Pat Hanrahan
Stanford University
http://graphics.stanford.edu/~hanrahan/talks/why/

"Machinima: Games Act Like Films"
Brad King
Wired News
http://www.wired.com/news/games/0,2101,53929,00.html

History of Computer Graphics
Marc Levoy
Stanford University
http://graphics.stanford.edu/courses/cs248-99/History-of-graphics/sld001.htm

"Digital Actors in Rings Can Think"
Courtney Macavinta
Wired News
http://www.wired.com/news/digiwood/0,1412,56778,00.html

"Microsoft's big PC plans"
Sam Parker
GameSpot
http://gamespot.com/gamespot/stories/news/0,10870,2912498,00.html

"NVidia's 3-D Vision Is Clear"
Andy Patrizio
Wired News
http://www.wired.com/news/technology/0,1282,42027,00.html

"Ray Tracing on Programmable Graphics Hardware"
Timothy J. Purcell ,Ian Buck ,William R. Mark , and Pat Hanrahan

Stanford University
http://graphics.stanford.edu/papers/rtongfx/

"Do Computer Games Need to be 3D?"
Richard Rouse III
SIGGRAPH Computer Graphics Newsletter
Vol.32 No.2 May 1998
http://www.siggraph.org/publications/newsletter/v32n2/contributions/rouse.html

"A Short History of Computer Graphics "
William Shoaff
Florida Institute of Technology
http://www.cs.fit.edu/~wds/classes/graphics/History/history/history.html

"Sony chip could transform video-game industry"
Dean Takahashi
San Jose Mercury News
http://www.siliconvalley.com/mld/siliconvalley/5310853.htm

"How Video Game Systems Work"
Jeff Tyson
How Stuff Works.com
http://www.howstuffworks.com/video-game.htm

"A (Spotty) History and Who's Who of Computer Graphics"
Matthew Ward
Worcester Polytechnic Institute
http://www.cs.wpi.edu/~matt/courses/cs563/talks/history.html

"A Brief, Early History of Computer Graphics in Film"
Larry Yaeger
Apple Computer
http://www.beanblossom.in.us/larryy/cgi.html

Websites:

"Quartz Extreme: Faster Graphics"
Apple Computer
http://www.apple.com/macosx/jaguar/quartzextreme.html

"Radeon 9800 Pro"
ATI
http://www.ati.com/products/pc/radeon9800pro/

"PressPass Microsoft Researcher Takes Computer Graphics Into the Future While Preserving the Past"
Microsoft
http://www.microsoft.com/presspass/features/2000/jul00/07-26salesin.asp

"GameCube"
Nintendo
http://www.nintendo.com/systems/gcn/specifications.jsp

"GeForce FX"
nVIDIA
http://www.nvidia.com/view.asp?PAGE=geforcefx

"How We Do it"
Pixar.com
http://www.pixar.com/howwedoit/index.html

"A Brief History of Computer Graphics"
PC Paradox.com
http://www.pcparadox.com/Editorials/History/1994.shtml

"Video Game Console Specs "
pcvsconsole.com
http://www.pcvsconsole.com/features/consoles/

http://www.xbreporter.com/xbox_vs_ps2_vs_gamecube.php

Interview Transcript with Joey Liaw, 3/16/2003, conducted online:

George: You may remember these questions from our [discussion during] nights of coding [cs140].

How powerful will GPU's get in the next couple of years?

Joey: Ridiculously powerful. Right now it's getting to the point where technology is outpacing demand, because there's healthy competition between ATI and nvidia. PC cards are so powerful, they've got so many multitexturing passes and effects, that it's becoming more and more difficult to make all the content for it.

For example, mid 90's... pc games are running at 320x200, 8-bit. You wanna make a texture for doom, ok, you make a 64x64 flat texture from a fixed palette of 256 colors in a simple pixel pusher painting program. You want to make a character sprite, you draw a low-res cartoon character from multiple angles, and you're done.

Fast forward to November 1999, Quake 3 Arena (q3a) goes gold. Cards have higher bandwidth and more memory, your game needs 256x256 16-million color textures, or you get criticized as looking "dated" or "blocky". Now we've got multitexture hardware, to take advantage of that, for each texture on a wall or character you have the option of several passes, so you're drawing two or three hi-res, hi-color textures, in addition the artists have to learn a scripting language to render the textures with the proper effects. Your character models are hitting a thousand polygons, each with their own sets of textures, plus you have to model them in a 3d package and do real 3d animation.

Nowadays, we've got cards that are capable of another order of magnitude in both total texture memory and texture passes. Who's got the money to have their artists make a hi-res bump maps, reflection maps, texture maps, animations, write shader scripts, etc. etc. New technologies like pixel shaders, 3d textures, we're giving artists for real-time 3d games more options than most know what to do with, these are things whose potentials I think we haven't fully explored, and it's hard for everybody, programmers and artists alike, to keep up.

Corrinne Yu once joked to me that, 3d cards are getting so ridiculously powerful that we're going back to 2d. Basically saying that since we're back to the level of operating on the pixel level rather than the triangle level. Coding those pixel shaders sort of feels like those days of tweaking your faux-3d special effects assembly routines operating directly in your 2d framebuffer.

George: Is there a limit to how powerful they need to be?

Joey: Well, you mean like audio technology? I guess so. Once your eye can't tell the difference between what's real and what's not, we're done. But only in terms of technology, not in terms of art. 44.1 16-bit stereo is good enough for 99% of the population (as is audio compression technology like mp3 and aac),

and samplers sound so good that it's hard to tell from the real thing, unless you're an expert.□But to make visuals on that level.... there's no way to predict when graphics technology will be powerful enough to do such a thing, but I think we're really, really□close.□t's just that nobody's made the real-looking game yet.
□
I think a good example is looking at particle systems.□You can make a cool looking particle system that takes 100,000 triangles to draw, but only exists inside your 3d off-line rendering package, but chances are you can make a particle system that looks 90% as good as that, but uses only 1000 triangles, with good textures and particle dynamics, and runs in real time.

George: Will other things like cpu speed or screen-resolution soon become the
bottle-necks?
□
Joey: No, I'm optimistic, I think the technology is very close to what we need for "realism."□t's a question of the massive art investment necessary to make it look real.□Or somebody clever will come up with a way to do it cheap and procedurally (like, effectly construct a 3d scene and textures out of photos or paintings).

George: What will happen to game technology when increasing GPU performance is no longer the state of the art?
□
Joey: Maybe then they'll start concentrating on the content. :P
There will be some point in the future where a large part of the population will be spending a lot of their time in massive online 3d systems, I think, so people are going to start concentrating on network software technology for games.□There's always the issue of how to optimize your art path, something like that might even go the p2p route where people can modify or create art for an online game.□There are already games that do distributed networking (massively online games with no central servers).

George: How are people using 3D technology outside of the realm of 3D games and Modeling?

Joey: Well... not so much 3d, but I think the Massive AI system used for Lord of the Rings battle sequences is interesting, because they've basically taken game AI and put it into an offline 3d package, so you're rendering a "game" that's not running in real-time.□ think I'd kind of want to play that game, actually.
□
3d graphics hardware becoming so powerful and so cheap, I'm sure people are using it for non-gaming applications.□ understand there's a push in financial markets to research how to visualize market flows in (using that old 80's term) VR, where they render the markets in realtime and you can manipulate data views with a dataglove.□Of course, we see it in advertisements all the time, but I guess that's not realtime.□Hmm..
□
I'd like to see a nice 3d GUI, nobody's done it right yet.□ think there's some usefulness to be had there, it's easier to see the relationships between data, even windows, when you've got them relationally navigable and interactive in 3d.□Oh yeah, I think the medical folks are using real time 3d applications for a lot of cancer diagnosis and treatment these days, they use a big x-ray laser mounted on a car assembly line mechanical arm and shoots beams into your head without opening you up, and they position the targets with an interactive 3d application (is my understanding, you should verify).□
□
Oh yeah, a really sad one.□America's Army, the recruitment game.□War is a video game.□t's still a video game, but it's also outside of standard video games, because you play in it, you get good, and then you can use those points for something to get into the army.□The military and NASA do a lot of 3d simulation for combat to train pilots and tank people and what not, I guess those are sort of games too, though =)