A Game Developer's Perspective of

SIGGRAPH 2001

Morgan McGuire <morgan@graphics3d.com>

Los Angeles, CA. 350 people, 6 terabytes of data, and 5 years of production -- *Shrek* piled up some impressive numbers. More impressive, however, was the end result: a 100% computer graphics (CG) tale of an ogre, a princess, and a donkey with enough personality and realism to show that PDI/Dream Works has what it takes to make an animated hit. Two days before the keynote address, the real kickoff to SIGGRAPH 2001 was the *Shrek* production team's massively attended course about the movie, which doled out these impressive numbers and described the science behind the film.

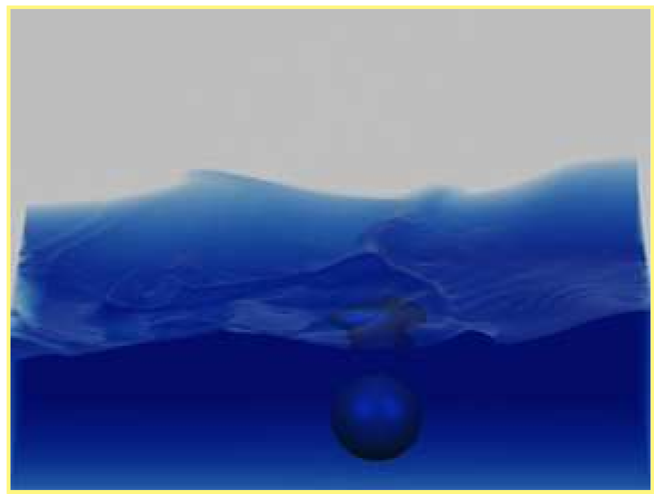

Fluid simulation algorithm used in *Shrek*. Click for animation.

SIGGRAPH is the Association for Computing Machinery's (ACM) special interest group on graphics, and also the name of its biggest conference, held this year at the LA Convention Center. The conference brings together researchers, hardware vendors, game programmers, artists and the motion picture industry to discuss the state of science, technology, and art in computer graphics. The result is matrices, polygons, GeForce3… and Shrek!

Computer animated films like *Shrek* contain a lot of hacks. For example, all lighting is direct rasterization-- no global illumination is used. This allows artists to light one character and not have extra shadows or light splash on to the scene, but it means the results aren't as realistic. To attain realism, effects like the bright focused caustic underneath a crystal sphere or the glint of sun off a sword are faked. Also scenes are built and animated as needed; move the camera a few feet and you'd discover that a character's legs aren't moving, and maybe don't even exist. Increasingly, however, major parts aren't hacks. Fluid simulations of mud, water, and beer in *Shrek* were based on the real 3D Navier-Stokes equations and appeared as a scientific paper by Foster and Fedkiw in the conference proceedings. Instead of moving facial bones and keypoints for animation like in video games or *Toy Story*, the characters of *Shrek* had a full physiological model of facial muscles and skin. Animators actually tensed individual muscles to make an expression, and an in-house program synchronized lip movements against vocal tracks. If only the Final Fantasy movie had used the same system to improve its characters' lackluster lip movements!

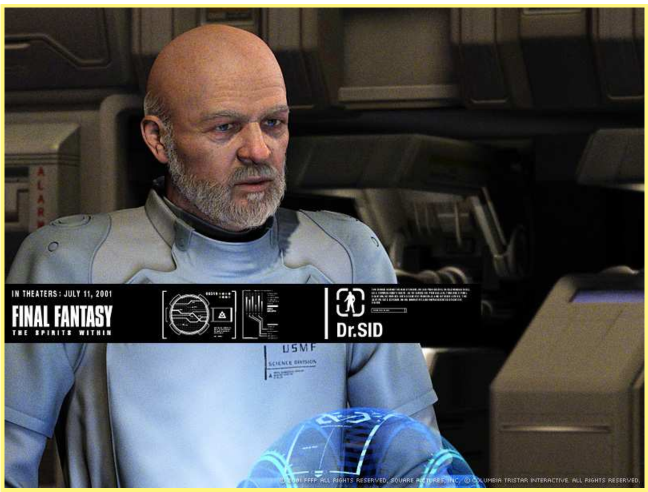

Square USA was rendering Dr Sid and Dr Ross in at 20 fps on an nVidia card. (image from the movie, not the real-time animation)

Despite the disappointing facial animation from their movie, Square Pictures and Square USA swept the other demos away on the trade exhibition floor at the conference. With nVidia's Quadro board in a 2.0GHz Pentium IV machine, they were rendering a scene from *Final Fantasy: the spirits within*... in real time! This incredible display used actual animation and dialog from the movie and simplified models. Using the nVidia card's advanced features, they combined per-pixel lighting, bump mapping, anisotropic shading, skinning, and depth of field effects. The demo was running at approximately 30 fps and was definitely being rendered in real time: you could move the camera and lights and turn on and off rendering features. nVidia presented information about implementing similar functionality on the GeForce3 at a real-time shading panel and released new demo code on their website.

Games Panels

Games have steadily increased in visibility at SIGGRAPH. This year, several panels and talks directly concerned games. Jennifer Olsen, senior editor at Game Developer Magazine led a panel in the exhibitor forum on the need for more game play. Jet Set Radio, Tony Hawk Pro Skater 2 (Game Boy Advance), and SSX were cited as examples of new games with great game play. The challenge is incorporating storyline, multiplayer, and solid arcade game play. Right now these elements are usually in opposition. The extreme sport games mentioned above all have game play but laughable storylines. Games like Deus Ex target an immersive experience but lack the simple drop-in fun of an arcade game, and don't work well as multiplayer games. Ion Storm promised that in Deus Ex 2 gamers will see more of the freedom and immersion that made the original so popular.

George Suhayda (Sony Picture Imageworks), Tom Hershey, Dominic Mallinson, Janet Murray (Georgia Institute of Technology) and Bill Swartout participated in another panel on game play entitled "Video Game Play and Design: Procedural Directions". Their focus was on leveraging the power of new graphics cards to improve dynamic storytelling.

A panel called "Masters of the Game" featured presentations by Academy of the Interactive Arts and Sciences awards winners. These presentations consisted of award show style speeches and lacked useful content, but the high level of attendance showed that SIGGRAPH attendees are interested in games.

Hecker vs. Pfister

The "Computer Games and Viz [Scientific Visualization]: If you can't beat them, join them" panel was intended to focus on ways scientists can take advantages of game technology. It turned into a series of antagonistic barbs as misunderstandings between the scientific and gaming communities bubbled to the forefront. Though the debate was kept good natured and everyone laughed a lot, the undercurrents of animosity were real.

Hanspeter Pfister, a MERL researcher well known for his work on volume and surfel rendering, led the charge against game developers. He characterized developers as Neanderthals eager for whiz-bang features with no long term vision, and living only to hack assembly and consume more fill rate. Chris Hecker, editor at Game Developer Magazine and the Journal of Graphics Tools, was witty and gracious as he defended game developers from the stereotype while acknowledging the contributions of scientists to games. Granted, wearing shoes and getting a haircut might have made Hecker's case for the maturity of game developers more convincing.

SSX on the PlayStation 2 by Electronic Arts. The lead render programmer Mike Rayner shared his experience with rendering and game play on the title.

It was clear that many scientists don't understand what game developers do and are upset that consumer grade graphics are oriented towards games and not scientific applications. They are in a tough place: cheap PC hardware has the power of a workstation, but lacks the finer appointments. Missing are high (greater than 8 bits per channel) bit depths, hardware accumulation buffers, image based rendering primitives, and high texture bandwidth. PC's are displacing workstations in many institutions, driving up workstation prices, yet the PC graphics cards don't really offer the feature sets that researchers need.

Scientific visualization is behind MRI's, computer assisted surgery, military network simulations, and earthquake research. In each of these situations, a misplaced pixel could be the difference between life and death. Results need to be correct, fast, and computed without crashing the machine. When researchers see "features" like embossed bump mapping come and go and a new API with each version of DirectX, they think this is what game developers want and that the technology used for gaming is too unstable for research.

Hecker and Nate Robins (Avalanche Software) explained that game developers want many of the same features that researchers want -- just check John Carmack's .plan. Higher internal precision, better filtering, better 3D texture support are at the top of the wish list. As to game developers using hacks, Hecker replied, "Where do you think we get our algorithms? We read SIGGRAPH." Bump mapping, Phong illumination, texture mapping, environment mapping, radiosity and other common game algorithms all appeared as scientific publications long before they found their way into Quake. Sophisticated game developers want a stable, platform independent graphics library. They have been at the mercy of nVidia, ATI, and Microsoft's feature wars as much as scientists.

Panel organizer Theresa-Marie Rhyne pointed out that games like SimCity have actually affected the way urban planning simulation is done. Other scientists on the panel, Peter Doenges (Evans & Sutherland) and William Hibbard (University of Wisconsin-Madison) were respectful of the graphics progress by games developers. Rhyne is no stranger to the convergence between games and science; she's written several papers on the topic and is experimenting with game hardware in her own lab.

In the end, few panel members seemed to change their viewpoint. An exasperated Chris Hecker was finally driven to object, "scientists need to stop bitching and moaning." Organizer Rhyne intended to stir up controversy, promoting these issues to the forefront, and got it. Hopefully she'll bring us a re-match next year.

Black Oil

At last year's art and emerging technology exhibition, a mechanical display device called the Wooden Mirror stole the show. This year's coolest techno-art was Minako Takeno and Sachiko Kodama's *<Protude, Flow>*. To see the piece, viewers lined up to enter a crowded black room tucked into the back of the exhibition. Once inside, people first stared at the writhing spiky shapes back-projected onto a wall. The shapes moved in response to a re-echoing sountrack of screams and clicks. Then it became clear that the image was a real-time video feed from a table in the room. In the middle of the table was a creation that was pure science fiction. A pool of black oil moved under its own control, forming independently roving lumps that looked like a koosh-ball's evil twin. Then they launched upward to a cylinder suspended over the table, a flow of oil that ran *up* to pool on the bottom of the cylinder. Without warning, the oil suddenly released itself and splashed down to pool and ripple like normal, inanimate fluid. Mechanical trick? Alien virus? Nope, the black oil is a ferrofluid, a suspension of regular oil and magnetic micropowder. A microphone picked up sounds from the audience and the soundtrack and used them to drive electromagnets in the table and the floating cylinder. The ferrofluid moved in response to this field, like the iron filings in a high school science class. Want one in your living room? Write to info@campuscreate.com… and don't leave your credit cards on the magnetic table.

Hardware Rendering

Behind the trade show, panels, art exhibition, and courses, comes the original purpose of SIGGRAPH: scientific papers. Many of this year's papers are directly relevant to game development. Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen's paper "Visual Simulation of Smoke" used the inviscid Euler equations to simulate realistic smoke in 2D and 3D. These equations are similar to the Navier-Stokes equations used for both real fluid flow simulation and demo effects of rain in a shallow pool. Most physical dynamics approaches suffer from numerical error that either adds or removes energy from the system after each simulation step. Adding energy causes simulations to explode and removing it causes them to deflate. Both kinds of instability lead to poor results. Fedkiw, Stam, and Jensen's approach avoids this by allowing the simulation to lose energy, keeping it stable, then selectively injecting energy back into the vortices. Because vortices are the most visually important part, this keeps fine-level detail in the smoke and makes it look realistic. Although the full 3D version of the algorithm takes 1 second to simulate movement within a 40x40x40 voxel grid (and render it with self shadowing!), 2D versions can run fast enough for real-time game effects. At the end of his talk, Stam held his PocketPC in the air and rendered 2D smoke reacting to real-time disturbance by his stylus and received thunderous applause.

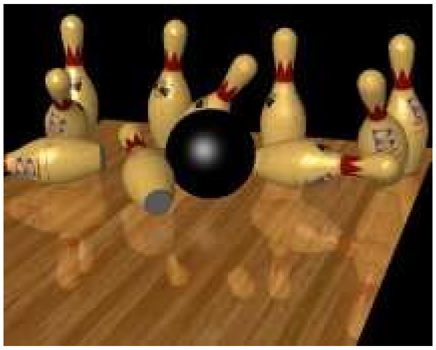

Images from A Real-Time Procedural Shading System for Programmable Graphics Hardware

Seven papers on hardware rendering showed that not all scientists are upset with their GeForce3 cards. The WireGL and Lightning-2 papers described systems for leveraging armies of PCs to render immense scenes at interactive rates. Kekoa Proudfoot, William R. Mark, Svetoslav Tzvetkov, and Pat Hanrahan presented "A Real-Time Procedural Shading System for Programmable Graphics Hardware". Their images, shown on the left, look like what we used to expect to see from ray-tracing, and are rendered in real-time.

Erik Lindholm, Mark J. Kilgard, Henry Moreton's "A User-Programmable Vertex Engine" paper describes the GeForce3 architecture. Although the GeForce3 is actually clocked slower than the last-generation GeForce2 Ultra, it is substantially more powerful because of the programmable pipeline. With a GeForce3, the programmer can download code onto the graphics processor for custom transformations of individual points and per-pixel shading and blending. This allows effects like per-pixel Phong shading, realistic hair and fur, and reflective bump mapping...although it does require programmers to work in a painful custom assembly language. Extensive examples are available in nVidia's newly released NVSDK 5.1.

Michael D. McCool, Jason Ang, and Anis Ahmad's "Homomorphic Factorization of BRDFs for High-Performance Rendering" shows that high quality shading and real-time can coexist. The Bi-Directional Reflectance Function, or BRDF, is a four dimensional function that describes how a surface reflects light based on varying viewer and light angles. The BRDF has long been considered the most scientifically accurate measurement of material lighting properties. It was traditionally used in high end ray-tracers, but McCool et al. show that clever programming of a GeForce3 can produce sophisticated light reflection in real-time.

New BSSRDF Technique

Just in time to raise the bar, Henrik Wann Jensen, Steve Marschner, Marc Levoy, and Pat Hanrahan introduced a new variation on the BRDF that accounts for subsurface effects. It can't be rendered in real-time with current techniques; to illuminate a single point on an object, all other points must be considered, making the rendering a time-consuming process. The effects look incredibly realistic, producing marble and skin that are almost identical to photographs.

Emil Praun, Hughes Hoppe, Matthew Webb, and Adam Finkelstein "Real-Time Hatching" used hardware accelerated textures and shading to produce hatched line art images of 15k polygon models in real-time. Randima Fernando, Sebastian Fernandez, Kavita Bala, Donald P. Greenberg's Adaptive Shadow Maps fix the biggest problem with the hardware shadowing support: aliasing. SGI's OpenGL shadow map extension is now supported by GeForce3 hardware, but unless the camera and light are pointed in the same direction, ugly rectangular blocks appear in shadows, making them look like they were rendered on an old Atari. The new adaptive shadow map fixes this problem by selectively allocating more resolution to the areas where the blocks would appear. When implemented directly in hardware, this new method may be the generic shadow algorithm game developers have longed for.

The full proceedings aren't yet on-line at the ACM (they will appear in the Digital Library soon), but some are freely available from the original authors.

The War in Graphics

The introduction of nVidia's GeForce3 has put the last nail in SGI hardware's coffin. Scientists are dropping expensive workstations in favor of (relatively) cheap, programmable GeForce cards. Nearly every demo at SIGGRAPH was run on a GeForce card, and nVidia was all the buzz despite their subdued showing at the trade show. With the GeForce hardware in the X-Box, the GeForce2Go laptop graphics card, and the nForce chipset for cheap GeForce2 graphics directly on a motherboard, nVidia is poised to seize to take low-end graphics away from ATI this year. Chris Hecker pointed out that for years, the PC graphics industry just copied old SGI designs. Recently, they passed SGI and started innovating on their own. Although early designs littered APIs with useless special-purpose effects, new PC graphics cards are showing generality and strong long-term architectural trends.

Jet Set Radio Future on the X-Box.

The industry is about to experience an all out war between ATI and nVidia that may clutter DirectX and OpenGL with more useless marketing-driving features. But if the new programmable pipelines are any indication, we'll also see real innovation, and possibly a single graphics chip manufacturer that will pick up where SGI left off.

The war won't only be on the PC's -- although all sides were silent at SIGGRAPH, this Christmas Microsoft X-Box, Sony PlayStation2, and Nintendo GameCube will face off. Sony has the existing market share, but GameCube's bringing world-renowned game designer Miyamoto the computing power he needs to make

Mario and Metroid the way gamers have always dreamed. With a modified GeForce3 under the hood, Microsoft hopes to take nVidia's winning streak into the console arena. The market may only be big enough for one to survive.

Looking Forward

In a few months, Pixar will release the most ambitious CG movie ever made. "Monsters, Inc." has billions of polygons in every scene and features dynamic simulation of the 2.3 *million hairs* on its furry hero,Sullivan. The scenes contain so much detail that Pixar cannot use regular 32-bit integers to index their polygons, and is rapidly approaching the point where they will spend more time transferring geometry over their gigabit network to the render-farm than producing the frames when they get there.

Most of this geometry is well below the pixel level, and a simplification would look exactly the same at the 1024x768 resolution at which many games run. If nVidia can stay on their incredible performance curve for another year, will we see Monsters, Inc. with fewer polygons rendered in real-time on a Geforce4? How will the new X-Box Oddworld game look next to the upcoming Oddworld movie? Even if real-time Monsters, Inc. is just a fantasy, the buzz about Doom III and Quake 4 says id Software is going to raise the standard for 3D games once again. Doom brought us first person shooters and network gaming, and turned id from a shareware company into a multi-million dollar games powerhouse. Quake made true 3D the standard and established John Carmack as a programming super-hero. Taking full advantage of nVidia's new features, the Doom III engine should finally cross the photo-realism barrier for games and reestablish id software as king of the industry.

With the upcoming battles of the graphics hardware and console titans and game developers and scientists discovering how to really leverage the awesome programmable GeForce3, this is sure to be a stellar year for computer graphics.

High noon in San Antonio, July 23, 2002: showdowns between Chris Hecker and Hans-peter Pfister; nVidia and ATI; Microsoft, Nintendo, and Sony; id Software and Epic MegaGames; PDI/Dreamworks, Pixar, and Square... *expect to be blown away at SIGGRAPH 2002*!

[*Updated 8/23/01. Thanks to Theresa-Marie Rhyne (NCSU), Max McGuire (Iron Lore), and Matt Kimball (Avalanche) for errata submissions. -Morgan*]

return to flipCode