



A Game Developer's Review of

SIGGRAPH 2002: San Antonio

Morgan McGuire <matrix@graphics3d.com>



July 2002 San Antonio, TX. [SIGGRAPH](#) is the [ACM](#)'s annual conference on computer graphics. This year the week-long event brought 20,000 members of the computer science, film, games, web, and art community together with the general public in San Antonio, Texas. The city has a beautiful European style river walk. It winds out of the past by the historic Alamo, through the present downtown Tex-Mex restaurants and hot dance clubs to end right at the convention center. Hosting SIGGRAPH, the convention center is filled with ideas and products straight from science fiction. Several sessions showed that computer generated imagery (CGI) for film and television passed the point where we can tell real from rendered, and papers showed that this trend will continue to dominate with completely synthetic environments and characters. Computer gaming is the common man's virtual reality and it has emerged as a driving force behind graphics technology and research. ATI and NVIDIA both announced new cards and software for making games even more incredible than the current generation, and several game developers were on hand to discuss the state of the industry. The rise of gaming and enthusiasm over the many high quality CGI films this year was offset by an undercurrent of economic instability. Belt tightening across the industry was clearly evident in smaller attendance and a subdued exhibition.



Yoda



It is common knowledge the animation team at [Industrial Light and Magic](#) (ILM) created virtual characters, space ships and



environments for Star Wars Episode II: Attack of the Clones. We know that giant space ships and four armed aliens don't exist and can easily spot them as CGI creations. What is amazing is the number of everyday objects and characters in the film that are also

rendered, but look completely real. The movie contained over 80 CG creatures, including lifelike models of the human actors. When Anakin Skywalker mounts a speeder bike and races across the deserts of Tatooine, the movie transitions seamlessly from the live actor to a 100% CGI one. During Obi-Wan Kenobi and Jango Fett's battle on a Kamino landing pad the actors are replaced with digital stunt doubles several times. One of the hardest parts of making these characters completely convincing is simulating the clothing. Any physical simulation is a hard task—small errors tend to accumulate, causing the system to become unstable and explode. Cloth is particularly hard to simulate because it is so thin that it can easily poke through itself and the continuing deformation makes collision detection challenging. New techniques were developed to simulate the overlapping robes that form a Jedi's costume. These appeared in a [paper](#) by Robert Bridson, Ronald Fedkiw and John Anderson presented at the conference. The purple cloth in the rendered image on the left is an example from their paper. Here, the cloth has dropped onto a rotating sphere and is twirled around it due to static friction, with proper folding and no tearing.

One of the highlights of Star Wars Episode II was the "Yoda fight" where the 800 year old little green Jedi trades in his cane for a light saber and takes on the villainous Count Dooku. A team from ILM: Dawn Yamada, Rob Coleman, Geoff Campbell, Zoran Kacic-Alesic, Sebastian Marino, and James Tooley, presented a behind-the-scenes look at the techniques used to animate and render the Yoda fight. When George Lucas revealed the script to the team just two days before shooting began, they were horrified. At the time, the team was still reviled by Star Wars fans for bringing the unpopular animated character Jar Jar Binks to life in Star Wars Episode I. Now they were expected to take Yoda, the sagacious fan favorite, and turn him into what George Lucas himself described as a combination of "an evil little frog jumping around... the Tasmanian devil... and something we've never seen before." The scene obviously required a computer generated model because no puppet could perform such a sword fight. Their challenge was to somehow make the completely CG Yoda as loveable as the puppet while having his spry fighting style still seem reasonable.

They began by creating a super-realistic animated model of Yoda based on footage from Star Wars V: The Empire Strikes Back. This model is so detailed that its eyes dilate when it blinks and it models physical ear-wiggling and face squishing properties of the original Yoda model. Interestingly, puppeteer Frank Oz disliked the foam-rubber character of the original model and wanted the animators to use a more natural skin model, but Lucas insisted on matching the original, endearing flaws and all. To develop a fighting style for master swordsman Yoda, the animators watched Hong Kong martial arts movies both old and new. Clips from these movies were incorporated directly into the *animatics*—early moving storyboards of test renders and footage from other movies used to preview how a scene will look. The SIGGRAPH audience had the rare treat of seeing these animatics, which will never be released to the public. The Yoda fight animatic was a sci-fi fan's ultimate dream: the ILM team took Michelle Yeoh's Yu Shu Lien character from [Crouching Tiger, Hidden Dragon](#) and inserted her into the Yoda fight scene in place of the Jedi. In the animatic, a scaled down Michelle Yeoh crossed swords and battled Count Dooku using the exact moves and motions that Yoda does in the final movie.

The digital wizardry continued with behind-the-scenes shots from other movies. In [Panic Room](#) the entire house was digitally created in order to allow physically impossible camera movement. In [The](#)

One, Jet Li's face is mapped onto stunt doubles to let him fight himself in parallel universes and virtual combatants are moved in slow motion to enhance his incredible martial arts skills. Most of the scenes in [Spider-Man](#) had a stylish rendered look, but amazingly some scenes that seemed like live action were completely rendered. At the end of the movie Spider-Man dodges the Green Goblin's spinning blades. We assume the blades are rendered, but in fact, nothing—the set, the fire, the blades, even Spider-Man himself—is real!

A scientific paper showed the next possible step for virtual actors. In [Trainable Video Realistic Speech Animation](#), Tony Ezzat, Gapi Giger, and Tomaso Poggio of MIT showed how to take video of an actor and synthesize new, artifact free video of that actor reciting whatever lines the editor wants to put into the performer's mouth. The authors acknowledged the inherent danger in a technology that allows manipulation of a video image. It now brings into question how long videos can be used as legal evidence and raises questions about what part of a performance and actor's appearance is owned by the actor and what is owned by the copyright holder of the production.

Although we've seen similar work in the past, nothing has ever been so realistic. Previous work spliced together video or texture mapped 3D models. The author's new approach builds a phonetic representation of speech and performs a multi-way morph for complete realism. This allows synthesis of video matching words, sentences, and even songs. As proof that the system could handle hard tasks like continuous speech and singing, a simulated actor sang along with "Hunter" by Dido and looked completely convincing. The authors followed up with an encore presentation of Marilyn Monroe (from her Some Like It Hot days) singing along to the same song. Both videos were met with thunderous applause. The authors' next goal is adding emotion to the synthesized video and producing output good enough to fool lip readers.

Fast Forward

This year the presentation of scientific papers was improved in two ways: they were published as a journal volume and a brief overview of each was given at a special session.

For the first time, the SIGGRAPH proceedings were published as a volume of the ACM [Transactions on Graphics](#) (TOG) journal. This change is significant for professors because accumulating journal publications is part of the process of qualifying for tenure. The journal review process is also more strict and maintains a high quality standard for the papers. The official proceedings are not yet available on-line, however [Tim Rowley](#) (Imagination Technologies) has compiled a list of recent SIGGRAPH [papers available on the web](#).



The papers were preceded by a special Fast Forward session. Suggested by [David Laidlaw](#) (Brown), this special session gave each author 62 seconds to



present an overview of their paper. This enabled the audience to

preview the entire papers program in one sitting and plan the rest of their week accordingly. The format encouraged salesmanship and many researchers gave hilarious pitches along the lines of “come to my talk to learn how to slice, dice, julienne and perform non-linear deformations on 2-manifolds.” Two short presentations displayed particular creativity. [Barbara Cutler](#) (MIT) recited a sonnet about the [solid modeling paper](#) she co-authored with [Julie Dorsey](#), [Leonard McMillan](#), Matthias Mueller, and [Robert Jagnow](#) – and we thought MIT wasn’t a liberal arts school. [Ken Perlin](#) (NYU), famous for his [Perlin Noise](#) function and work on [Tron](#), donned sunglasses and pulled open his shirt to rap about his [improved noise function](#).

In the end, it was 2002 Papers Chair [John “Spike” Hughes](#) (Brown) who stole the show with an incredible intellectual stunt. Several paper authors failed to appear to give overviews of their work. In the madcap parade of 62 second speeches, these sudden breaks left the stage strangely quiet and empty. Rather than let the session stall on these occasions, Hughes jumped to the podium and gave detailed introductions... to other researcher’s papers. The audience responded with laughter and applause as he explained the importance and mechanism of various techniques, demonstrating both his broad knowledge of graphics and familiarity with all of the papers at the conference.

Several papers were directly interesting to game developers. Shadow Maps, [originally](#) proposed by Lance Williams in 1978, are now a hot topic because recent hardware advances have made them possible to apply in real-time in games like Halo. Williams’ Shadow Maps suffer from poor resolution when the camera and light are pointing in different directions, however. Through application of projective geometry, Marc Stamminger and George Drettakis’ (REVES/INRIA) new [Perspective Shadow Maps](#) resolve this situation without affecting performance.



In recent years, [Hugues Hoppe](#) (MSR) has created several new methods applicable to gaming. His [Lapped textures](#), [Texture mapping progressive meshes](#), [Real-time hatching](#), and [Real-time fur over arbitrary surfaces](#) were important work for real-time rendering and texture mapping. His [Displaced subdivision surfaces](#) and [Silhouette clipping](#) proposed new hardware modifications for letting low-polygon objects appear

to have fine detail, speeding up rendering. This year he’s back, presenting [Geometry images](#) with coauthors [Xianfeng Gu](#) and [Steven Gortler](#) (Harvard). This new method encodes a detailed 3D model directly into a texture and normal map, allowing the use of standard texture compression, MIP mapping, and manipulation algorithms on geometric data.

Most game 3D animation is currently performed by fitting a skeleton inside a 3D mesh and interpolating the mesh position to fit the skeleton. This approach lacks a solid feel (it’s more like chicken wire over the skeleton than a body filled with fat and muscle) and can’t handle certain movements like twists and interaction with the environment very well. The University of Washington group: Steve Capell, Seth Green, Brian Curless, Tom Duchamp, and Zoran Popovic presented a paper on [Interactive Skeleton-Driven Dynamic Deformations](#) that resolves these problems and still allows real-time rendering and animation. Their technique fits a spring lattice over the body that moves with the skeleton. Physical simulation is run on the model to achieve lifelike results. By modeling bones with actual thickness instead of an infinitely thin skeleton, they handle twists of skin and muscle mass about bone much more realistically than traditional models. Demos of a walking cow and one subjected to having its ears and udders tugged looked very

organic compared to what we're used to seeing in games. Except for some hand tuning needed in the lattice creation, the results look immediately applicable to game development.

ATI

[ATI](#) is one of the leading consumer and workstation graphics card developers. This year they announced and demonstrated the Radeon 9700 graphics card. Showing just how incredible hardware has become, the engineers from ATI built a [demo](#) that renders [Paul Debevec's](#) *Rendering With Natural Light* film *in real-time*. This is amazing because the original film took hours per frame to render when presented four years ago at SIGGRAPH 98. Another demo showed a photorealistic race car with two tone paint, light bloom, normal mapped hood scoops, reflection, and sparkle paint flecks. On the organic side, ATI's toothy [bear](#) looks so good you want to run before it bites you.



To be released next month, the Radeon 9700 supports the expected 2x bump in most specs with AGP 8x and 128 MB of video RAM. When available on shelves it will be the card of choice for serious gamers as it is the only production card supporting DirectX 9.0 and 2.0 pixel & vertex shaders and the current OpenGL 2.0 proposal, as well as outperforming the fastest card being sold today: GeForce4.



This isn't just an overclocked Radeon 8500 with more bandwidth, however. The design is truly revolutionary with the use of a 128-bit floating point frame buffer and a completely floating point pipeline. This solves color quantization errors and will allow games to have smoother textures and higher dynamic range—a truly blazing sun won't get cut off at saturated yellow just because there is some detail on screen in the shadows as well. Game developers are now used to encoding 3D coordinates as colors in a texture map in order provide massive amounts of geometry data for per-pixel programs. This practice allows effects like per-pixel lighting, reflections, and bump mapping but was limited by the 8-bits per channel allowed on previous cards. The 9700's floating point pipeline allows 24-bits (floating point) per channel throughout the pipeline, making these per-pixel effects more realistic for sharper reflections and smoother shading.

The programmable pipeline has been extended from the Radeon 8500. The per-pixel shaders used primarily for lighting effects can now be up to 64 instructions long and per-vertex shaders for geometry deformation can be 256 instructions, including limited branching and looping statements. These numbers are all hard to evaluate because in many cases the card is limited by the DirectX API but can issue more instructions due to loop unrolling in the driver—the rough order of magnitude is the important thing to take away.

With great power... come great debugging headaches. ATI recognizes that the current system for programming graphics cards is a complete mess. Different APIs are in competition (including NVIDIA's [Cg](#), 3D Labs' OpenGL 2.0 shading language proposal, Stanford's [RTSL](#), and SGI's ISL) to be the shading language of choice, and all of them require a great deal of hardware knowledge on the programmer's part. Moving shaders between art programs like Maya and game engines is a nightmare, as is asking 3D artists to poke around assembly code to tweak art parameters.

ATI's [RenderMonkey](#) tool suite is intended as a solution to these problems. It allows programmers to create shaders within a development environment that looks like Microsoft Visual Studio, previewing and debugging them in real-time. These shaders can be written in any shading language. ATI has even provided a translator for Pixar's RenderMan shaders so they can run in real-time on consumer graphics hardware. The tool suite allows programmers to export shaders that work in Maya, 3DS Max, and SoftImage. The file format is human readable ASCII and is compatible with DirectX 9.0 .FX files—important for revision control. OpenGL support will be provided when the Architectural Review Board (ARB) specifies an official OpenGL shading language.

ATI held the graphics hardware pole position in the 1990's with massive market share, but competitor NVIDIA's GeForce cards led consumers to associate NVIDIA with innovation and top performance. Part of the reason that NVIDIA has been beating ATI in recent years was NVIDIA's incredible developer support program and regular release of reliable high performance software drivers. ATI has clearly recognized their past problems in this area. ATI isn't just releasing an amazing consumer card next month—they are also releasing tons of technical information through their [developer site](#) and hiring more staff for the [developer relations](#) group. A new emphasis on quality over performance is aimed at fixing their reputation for bad drivers. This is supported by a corporate reorganization and new hardware architectures for which it is easier to write reliable, high performance drivers. Combined with ATI's support for open standards and development of the device-independent, API independent RenderMonkey tool suite, the new emphasis on quality makes ATI a clear corporate leader that is responsive to the industry, developers, and customers.

What's next? ATI is focusing on providing vertex and pixel shaders on lower end and mobile hardware. They will also ship more Render Monkey functionality, with support for 3DS Max, Maya, and SoftImage XSI. While the Radeon 9700 is an awesome card, there are some features that would really round out the feature set if added. Currently, there is no alpha blending when using a 128-bit frame buffer. This requires developers to rewrite their rendering pipelines in order to make effects like shadows, translucency, and particle systems work with high precision color. Adding high precision alpha blending, even longer vertex and pixel programs, and support for the NV_DEPTH_CLAMP OpenGL extension (important for [shadow volume](#) rendering) would really make this a next generation card with features to match its raw performance.

NVIDIA



ATI's main rival, NVIDIA, announced their [CineFX NV3x](#) architecture at SIGGRAPH (see a [comparison](#) to NV2x on SharkyExtreme.) NVIDIA is a company associated with innovation and developer support. Their GeForce card was the first consumer hardware texture and lighting card, a major technical leap forward when it debuted. The GeForce3 card introduced programmable consumer hardware, another first. The CineFX architecture is again going to take us forward when it appears on shelves around Christmas 2002.

Like the ATI Radeon 9700, the NVIDIA CineFX cards offer a 128-bit frame buffer and floating point pipeline. NVIDIA raises the bar with 32-bit floats (vs. 24-bit) in the pipeline and blows away ATI's vertex and pixel shaders. Instead of offering a few hundred instructions, they went for 1k (that's a thousand!) instructions in pixel shaders and 64k instructions in vertex shaders. These include branches and loop instructions. As has been the case in the past, the NVIDIA card promises higher raw performance and longer pipelines but gives less texture support. These vertex and pixel shaders may eclipse the ATI Radeon 9700's limits when the CineFX cards ship, but the ATI card has twice the texture access rate. However, the cards are so difficult to program that it may be the tools and not the hardware that makes the difference in the end.

CgFX is NVIDIA's tool suite, directly opposite ATI's RenderMonkey. CgFX is being developed closely with Microsoft's DirectX 9.0 FX API and at this early stage is nearly indistinguishable. The DirectX FX system allows bundling of different effects like bump mapping and reflections into a single program without requiring a programmer to rewrite the code. This addresses one of the major problems with shader development. Developers can apply to join the DirectX 9.0 beta program at [Microsoft's beta place](#) using user ID "DirectX9" and password "DXBeta" and check out NVIDIA's programmable architecture [documents](#) online. CgFX is also closely integrated with NVIDIA's [Cg](#) shader API. Although this API lacks the multipass abstraction of SGI and Stanford's APIs and the multi-vendor support of OpenGL 2.0, it boasts an efficient, working implementation and complete DirectX 9.0 and OpenGL 1.4 support. The Microsoft Xbox runs on NVIDIA audio and video processors and the CgFX Xbox port is 95% complete. CgFX will also ship with plugins for popular 3D modeling packages like Maya, Max, SoftImage/XSI 3.0.

While waiting for NVIDIA's CgFX release, developers have a number of other tools at their disposal. NVIDIA has a dedicated team of 20 engineers led by Director of Developer Technology John Spitzer that is continually

Improving Developer Relations

NVIDIA's weaknesses have never been technical. Their fantastic cards, reliable drivers, and regular innovation have made them a favorite of gamers and game creators. The company has had trouble disseminating clear information to both developers and users, however.

For years, the [developer website](#) has been rich in content but poorly organized so that content is impossible to find. NVIDIA recognizes this problem and has launched into a redesign of their website. This new site will include a better keyword system for searching and a knowledge base section. The registered developer site that contains the NVTune DirectX profiling tool and unreleased drivers crashed previously, deleting user accounts, but is now getting more stable.

Unfortunately, their end-user information problem doesn't look like it will go away soon. NVIDIA has repeatedly shipped

producing and improving the NVIDIA toolset. [NVMeshMender](#) adds tangent and binormal data to 3D models to help transition new shader features into existing art pipelines. The [Cg Browser](#) replaces the NVEffectsBrowser 4.0 for previewing effects. [NVTriStrip](#) is a library for packing triangle geometry into more data efficient formats. The original [NVParse](#) shader tool is still needed to work with Cg under OpenGL because it can convert DirectX shaders for use with OpenGL. DirectX texture formats can be accessed under both OpenGL and DirectX using the open source [DXT](#) compression library.

products with [misleading names](#). For example, a GeForce4 MX card lacks pixel and vertex shaders—it is effectively *less* powerful than a GeForce3 card! This makes it impossible for game developers to give clear requirements to consumers and creates a mess out of what should be a clear purchase decision for users. Will CineFX ship as GeForce5 while the GeForce5 MX is really just a GeForce3, or will their marketing department finally fix this embarrassment?

The full specs for the NVIDIA CineFX cards have not yet been publicly released. Ideally, the card will ship with as many texture units as they can cram onto it, support for two-sided stencil testing and GL_DEPTH_CLAMP (proposed in Everitt and Kilgard's [shadow paper](#)), and nView multi-monitor support without needing a dongle like the GeForce4 cards. NVIDIA should also get those extensions into drivers for older hardware like the GeForce3 so game developers can start using them right away. Losing alpha blending on the Radeon 9700 when operating in 128-bit mode is really unfortunate. Hopefully NVIDIA's card will make up for trailing ATI's next generation card by four months by having full alpha blending support or a hardware accumulation buffer for the same purpose.

Gaming Panels

There were three gaming sessions at SIGGRAPH this year: panels "[Games: the Dominant Medium of the Future](#)" and "[The Demo Scene](#)" and a forum "[Game Development, Design, and Analysis Curriculum](#)." As moderator of the games panel, Ken Perlin let the audience submit questions for the panel on index cards, then selected questions to maximize interesting discussion. This technique kept the session topical and interesting. It also maintained a brisk pace for two hours and brought in much more audience participation than we usually see in a SIGGRAPH panel.

The demo scene panel was primarily targeted at educating outsiders about the [scene](#), which consists of amateur programmers using game techniques to create abstract real-time art. Vincent Scheib (UNC) and Saku Lehtinen aka Owl/Aggression (Remedy) gave a solid introduction and showed several screen shots and live demos including [Edge of Forever](#) by [Andromeda Software Development](#). Another scene member was on hand to share his experiences: Theo Engell-Nielsen, aka hybris/NEMESIS.. For games programmers and scene members, one of the highlights was seeing a photograph of [Future Crew](#) in 1992 at Assembly92, a scene conference/party. Future Crew was one of the premiere demo groups, many of whose members later joined [Remedy Entertainment](#), developer of [Max Payne](#). In 1992, they were a bunch of geeky looking students bumming around a classroom with "Assembly92" written on the blackboard. In contrast, Remedy is now a leading game company and Assembly2000 filled a hockey arena with computers... both have traveled far from their humbler beginnings.

Industry veteran game designer Warren Spector (ION Storm) shared his thoughts on the industry. He believes that games are a serious art form, capable of political, ethical, and social commentary, and are receiving recognition as such. He points to NY Times game reviews and Newsweek's game columnist as games going mainstream in the media, and joked that there was no greater recognition than Congress trying to put game developers out of business. At the same time, he doesn't think the growth of games is a foregone conclusion. Publishers push for very unhealthy practices in order to "make a quick buck." These practices include emphasis on violent content, sequels, and content licensing. Spector fears these practices will marginalize games, like what happened to American comic books post-1950. He encourages game developers to make their games playable as soon as possible in the development cycle in order to focus on game play. It is easy to improve geometry resolution or pixel resolution but the key is "behavioral and emotional resolution"—how immersive is the game, how emotionally does the player react? He recognizes that brand name designers like Sid Meier (...and Warren Spector) have the luxury of creating games the right way, with emphasis on design, while most developers are at the mercy of their publishers. Reiterating the publisher problem, he directly said that the publisher-developer battle is getting out of hand. Currently, there are few game publishers and each is very powerful. The situation that results is much like that of the music industry. Publishers force abusive contracts on developers, taking all of the intellectual property rights while forcing developers to operate on compressed time scales and relatively small budgets.

Loren Lanning (Oddworld) picked up this thread, observing that with more money and more risk going into each title, publishers are also more cautious. They want to rely on licensed content (e.g. [Spider Man: The Movie](#) game, Star Wars games) instead of supporting novel gaming ideas and really improving game development. He also recognizes that there is little that the developers can do in this battle but recommends savvy lawyers and negotiators because contract negotiation is operating at the Hollywood, not

Spotlight on America's Army

Michael Zyda, Director of the Naval [MOVES Institute](#), which created [America's Army](#), was one of many SIGGRAPH attendees with a background in games. In addition to attending the sessions, he was promoting his game and scouting for a new programmer to join the team.

America's Army became incredibly popular within days of its July 4 release. There are 2,000 servers with 3,000 players on each of them running 24/7—that's 600,000 players on-line! The game was prototyped in 12 months, then scrapped and rebuilt again in 12 months by an 18-person team, including art, managed by lead level designer Alex Mayberry and producer Mike Capps. The entire game was produced on time \$6.3M, plus \$1.5M for marketing. Crucially, it was ready for release on schedule and stable enough to support so many players. The Army team's development strategy involved a talented team, close management, and use of the unreleased Unreal Tournament 2 (UT2) engine under the hood. By leveraging the Army's massive marketing system, the game was primarily promoted at no cost to the developer. Of course, having the game funded, staffed, and specified from the start was a major help—most game companies have to scrap along until they get picked up by a publisher and hence can't build as solidly as they would like.

Traditional problems with online games were resolved through new ideas and new technology. To create a socially accountable gaming community, the Army's strict rules of engagement are automatically enforced by the game system. A player who injures a


amateur software level these days. [Valve](#), the developer of [Half-Life](#) has created a new platform called [Steam](#) that is a potential way out of this mess—it promises to help developers to deliver their games on-line without publisher involvement. Xbox creator Seamus Blackley's [Capital Entertainment Group](#) (ceg) is another option. Ceg acts like a venture capitalist for innovative gaming companies; people making exactly the kind of games that Spector and Lanning applaud but publishers fear.

Will Wright (Maxis) revealed some of the techniques he uses while working on hit games like [Sim City](#) and [The Sims](#). He thinks of game design as programming two platforms: a technology platform using computer code and a psychology platform inside the player's head. The game begins when the player picks up the box in a store and imagines what playing it will be like. During actual game play the symbols from the game form a virtual world inside the player's head, which is how early games could be so compelling even without sophisticated graphics. In this view, there is the game itself and the meta-game: off-line social interaction, websites, the box design, and action figures. All combine to create the compelling experience for users. Success in a game isn't just accumulating items or points. Often a player's idea of success is different from what the game specifically rewards. For example, players in Half-Life will sometimes stack avatars on top of one another or create silly maps rather than playing the game the way it was originally designed. To understand the meta-game and what players enjoy about his games, Wright collects and analyzes player information. He tracks how beta testers (or actual players for an on-line game) move through the state space formed by individual statistics like wealth and multiplayer interaction patterns. These can be viewed as three-dimensional graphs that show an "enjoyment" landscape setting one dimension against another. Understanding this landscape enables him to design for what players actually enjoy, rather than impose artificial goals. He also encourages player content creation ("mods"). Creating the game and modifying the world in which it takes place are part of the meta-game and he sees himself as orchestrating and producing that, not just the bits in the box.

teammate with friendly fire or engages a target that has not been identified as a threat will find himself or herself booted from the online game and locked in jail. Multiple violations result in a permanent ban from the game. Statistics are tracked for players and stored on central servers. These actually affect game play, not just bragging rights—you had better get a good marksmanship score if you want to be a sniper!

Because statistics and player histories are so significant to the game, strong encryption is used to protect the integrity of the entire gaming network. Thus far, there have been no successful attacks or cheats, a feat no other online game has achieved.

The Moves Institute is still hard at work. In the coming months, expect to see add-ons for America's Army (such as the highly anticipated Sniper School). Also in the works is an entirely new game: Soldiers, a boot camp role playing game that uses video sprites and pre-rendered environments to provide complete photorealism.



Art

The art and emerging technology exhibit wasn't as large or creative



as in previous years. Two research projects did shine above the crowd as “toys I must have, please, Mom, please!” Sony’s Block Jam is a midi sequencer that feels more like a combination of dominoes and the old demoscene [ScreamTracker](#) than a piano keyboard—or maybe it’s LEGO meets [Frequency](#). The individual white cubes have gloss black tops and connect magnetically to each other horizontally and vertically. Tapping the cubes or circling a finger over them causes color LEDs to shine through and makes the system play music. Different digital sound samples, beats, and mix levels can be accessed by changing the connection pattern and value displayed on each block. On a monitor, the block configuration is rendered in 3D with each block animated as it plays its sample. The project was conceived as a collaborative space for kids but ends up as a terrific DJ tool for remixing techno and hip-hop on the fly. Unfortunately, the project is pure research and isn’t scheduled to be released as a product, so you can’t put it on your Christmas list.

Daijiro Koga’s Virtual Chanbara project from The University of Tokyo uses a head mounted display and real sword hilt to simulate Japanese sword fights. Players don the display and hold the sword hilt, which uses a gyroscope to provide force feedback. They see a low polygon 3D world with a virtual opponent and battle to the (virtual) death, ala Hiro Protagonist in [Neal Stephenson’s Snow Crash](#). The system is inherently cool, especially for a student staffed project. Its most impressive feature was (mostly) holding together under the enthusiastic swordplay of attendees. Next year, Koga’s group hopes to be back with a better system and multi-player support.



Only Forward

Despite being one of the best funded scientific conferences in the world, SIGGRAPH has been steadily feeling the pain of a falling economy. Attendance and exhibition vendor registrations fell steadily from [2000](#) to [2001](#) and the sinking trend continued this year. In the exhibition, the vendor booths were noticeably scaled back from the extravagance typical for a trade show. The art and emerging technology exhibits lacked the diversity and inspiration we’ve come to expect (we aren’t seeing innovation like the [wooden mirror](#) or [black oil](#) of past years). Research papers are increasingly focused on incremental work and industry applicable solutions rather than pushing the state of the art and our concept of computer graphics years into the future. Clearly larger research budgets and less emphasis on near-term applications are needed to let scientists explore radical new ideas. With its own budget declining, SIGGRAPH is also looking at scaling back the conference. Proposals have been made for eliminating the panels sessions altogether and shortening the entire conference by one day. Eliminating panels would be especially damaging for the games hardware and software industry and research communities. Presentations on current graphics trends like programmable hardware, OpenGL 2.0, and computer games tends to appear in the courses and panels sessions.

Graphics professionals should [e-mail SIGGRAPH 2003 Conference Chair Alyn P. Rockwood](#) now to share what they’d like to see at [SIGGRAPH 2003](#). Bringing game development more centrally into conference, the Fast Forward papers review and increasing the credibility of SIGGRAPH papers by

publishing them in ACM TOG were great improvements made in 2002. In 2003 these should continue. The game panels should continue as well. Increasing the registration fee or reducing cost by scaling back facilities is preferable to seeing SIGGRAPH move away from games at a time when game development is driving so much of the graphics industry and research.

Check out the [Official SIGGRAPH 2002 Reports](#). Thanks to the SIGGRAPH student volunteers for photographs of the convention.

[Morgan McGuire](#) is Co-Founder of Blue Axion and a PhD student at [Brown University](#) performing [games research](#). Morgan and flipcode thank Bob Drebin (ATI Director, Engineering), Callan McInally (ATI Manager, 3D Application Research Group), Chris Seitz (NVIDIA Manager of Developer Tools), and Michael Zyda (Moves Institute Director) for granting interviews for this article, and the [Brown University Graphics Group](#) for funding the author's trips to SIGGRAPH for the past two years.