



A Game Developer's Review of

SIGGRAPH 2003

Morgan McGuire <matrix@graphics3d.com>



Why is the Sky Blue?

July 2003 San Diego, CA. In San Diego you can sun under perfect blue skies and watch gorgeous ocean sunsets, then head out for a night of dining and hot dancing in the [Gas Lamp Quarter](#). You might stop to wonder why those skies are blue and sunsets red. The answer is that Rayleigh scattering favors short wavelengths, so in the afternoon the sky is suffused with blue light. As the sun drops near the horizon we are in the right orientation to observe the red wavelengths. Together, this festive atmosphere down 5th street and nugget of atmospheric wisdom from [Course 1](#) kicked off [SIGGRAPH 2003](#).



Rendered sky from Mark Harris's outdoor illumination talk.

The [ACM](#) special interest group on computer graphics conference draws computer scientists, artists, and film and game professionals. Scientific papers and courses introduce the newest theory and an exhibition shows off the latest practices as NVIDIA, ATI, Pixar, SGI, Sun, and others showcase their products and productions. On the creative side, the Electronic Theater contains two hours of computer generated imagery. It is augmented by a continuously running Animation Festival, art gallery, and emerging technology display.

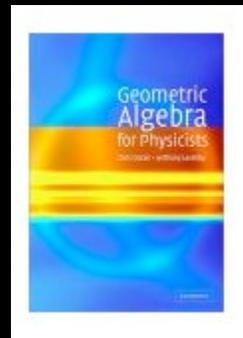
To Infinity and Beyond

Cambridge University astrophysicist Anthony Lasenby delivered the keynote address. He invited an

audience focused on triangles to consider issues of a more cosmic scale. Physicists have long theorized that just as the universe exploded from a tiny point in the Big Bang 14 billion years ago, it will some day collapse back to that point in a Big Crunch. This crunch would be the collapse of reality as we know it. Although it would not happen for an inconceivably long time, the notion of an inevitable death sentence for humanity has hung heavily over science fiction fans and astrophysicists familiar with this theory. The deciding factor for our ultimate fate is whether the universe is currently expanding fast enough. If it is, growth will fight the gravitational forces attracting all matter towards a common center. If not, we had better get started on interdimensional travel right now.

In pursuing the question, a startling fact emerged: the universe is closed in space. This means that if you traveled very fast and long enough in one direction you would come back to where you started. If this sounds strange, think of how the same argument would sound to someone who thought the Earth was flat. They measured the expansion rate as well. Good news: the universe is also expanding sufficiently fast that we are safe from the Big Crunch.

What does this have to do with graphics? The Geometric Algebra (a.k.a. Clifford Algebra) used by astrophysicists to resolve these issues can also be used to simplify movement in 3D. Game programmers typically treat straight-line motion separate from rotating motion. This complicates the process of finding in-between orientations given two positions, since the translation and rotation components must be interpolated separately. Furthermore, rotations by themselves are tricky to interpolate. [Geometric algebra](#) presents a framework where 5-dimensional "rotors" represent both 3D rotation and translation within a unified framework, and also take the pain out of rotation. This leads to smoother camera movement and character animation as well as a much cleaner theoretical framework for motion in computer graphics. How much does this idea appeal to graphics practitioners? Lasenby's \$90 [book](#) on the topic sold out from the bookstore within the next hour and immediately started racking up mail-order requests.



New Tools for Game Development

Games have become pervasive at SIGGRAPH. Game developers are in the audience and behind the podiums and their games are shown in the Electronic Theater and played in the [Vortechs](#) LAN party. Most of the papers presented were relevant to game development, but a handful really stood out as ideas solid enough to begin incorporating into new projects immediately. Surprisingly, these generally were not the hottest new rendering techniques but solid behind-the-scenes algorithms appropriate for game creation tools.

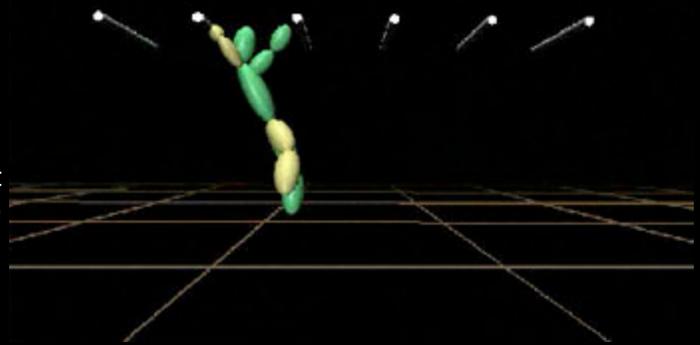
Terrain Texture

[Wang Tiles for Image and Texture Generation](#) by Michael Cohen (Microsoft Research), Jonathan Shade (Wild Tangent), Stefan Hiller, and Oliver Deussen (Dresden University of Technology) is a texture synthesis method that can reduce tiling artifacts on terrains. No matter how seamless the wrapping transition and well constructed the texture, a single repeating grass or rock texture will produce visual tiling artifacts when viewed from far away. To overcome this, games currently use multiple grass tiles and assign them randomly. However, the seams are visible between the

different grass textures. For a texture like rock that has larger features, these seams are completely unacceptable. The Wang Tile paper describes how to fit a set of related textures together to form an arbitrarily large plane that has no repetitive pattern and no seams. Synthesizing these textures requires only an algorithm like [Wei and Levoy's](#) 2000 paper or any of the [several](#) texture synthesis papers from last year.

Easy Motion

Anthony Fang and Nancy Pollard of Brown University presented [Efficient Synthesis of Physically Valid Human Motion](#). Their algorithm runs an optimizer to compute physical motion from an initial set of key poses for a human body. What is amazing is how few key poses it needs. For example, say we want to generate an animation of a person swinging across a set of monkey bars. Fang's algorithm need only be told that the hands touch each rung in sequence. It uses built-in knowledge of the human body to independently deduce how to swing the arms, head, and legs for a realistic animation. This method could be adapted to create a whole sequence of skeletal animations for a game character without requiring an animator to manually guide each limb through its motions. Other examples in the paper include synthesizing walking and a gymnast's somersault dismount from a high bar.



The optimizer creates realistic motion, but games often need stylized motion. Addressing this need, Mira Dontcheva spoke on [Layered Acting For Character Animation](#), a new animation method created with coauthors Gary Yngve and Zoran Popovic at the University of Washington. Their system allows a user to act out motion for a character one body part at a time. Unlike traditional motion capture, the layered approach allows a human being to simulate the motion of a kangaroo, dog, or spider because the acting applies to a single limb and not the whole body. By demonstrating the motion path of the spider and giving eight separate leg performances, a two-legged researcher can become an eight-legged monster. Acting is far easier than using typical animation tools like Character Studio, so now animators can be replaced with actors or even novices to create motion data. They will not be entirely out of a job, however. Someone still has to rig the skeleton within a 3D model and possibly tweak the results.

Easy Texture



Kraevoy et al's Matchmaker makes photo texture mapping easy.

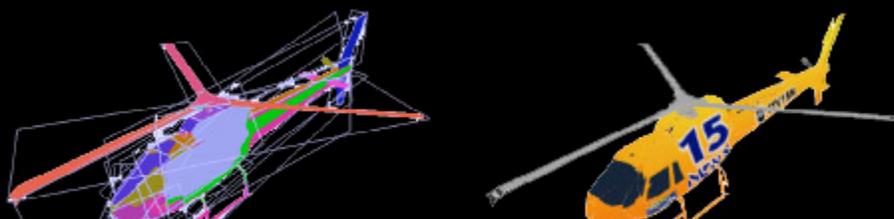
Digital cameras make it easy for artists to capture realistic textures and Photoshop is the tool of choice for editing or painting them from scratch. To map a flat texture onto a 3D model, an artist assigns each 3D vertex a 2D coordinate within the texture map. Even skilled artists have to work hard to make sure this process does not distort the texture and the results often are not as good as they could be. Compare the texture maps from a typical 3D shooter to the images seen in-game. Distortions often ruin beautiful artwork and give faces an inhuman look.

Alla Sheffer presented a paper on this topic, [Matchmaker: Constructing Constrained Texture Maps](#) by Vladislav Kraevoy, herself, and Craig Gotsman from the Israel Institute of Technology. Their program allows a user to make simple constraints like locking eye and mouth position in both the texture and model. It then solves for a consistent way of smoothly mapping the texture across the model. The process is akin to a traditional 2D warp but the output is 3D. Their results (pictured on the left) are incredible. The male head is texture mapped from a single photograph (mirrored to cover the whole face) and the full female body from simple front and back images. Each was produced using a few mouse clicks. Using conventional tools, these results would take hours of manipulation by a skilled artist to create.

Rendering Trees

It is common in games to simulate complex 3D objects with a single square facing the viewer. This square is called a billboard or sprite, and is textured with a picture of the 3D object to make it look convincing. Older games like Wing Commander and Doom used this technique exclusively. In new games characters and vehicles tend to be true 3D objects but explosions and trees are still rendered with billboards. One problem with billboards is that they appear to rotate as the viewer moves past since they always have the same picture. Using two or more crossing billboards with fixed orientations creates an object that appears to stay fixed in 3D, but up close the individual billboards are visible.

[Billboard Clouds](#) is a more sophisticated technique proposed by Xavier Decoret, Frédo Durand, François X. Sillion, Julie Dorsey of MIT, INRIA, and Yale. Their algorithm takes a 3D model as input and produces a set (or confusingly named "cloud") of billboards on various angles that collectively approximate the model. For distant objects the billboards appear to





capture all of the complexity of the 3D model using far fewer polygons.

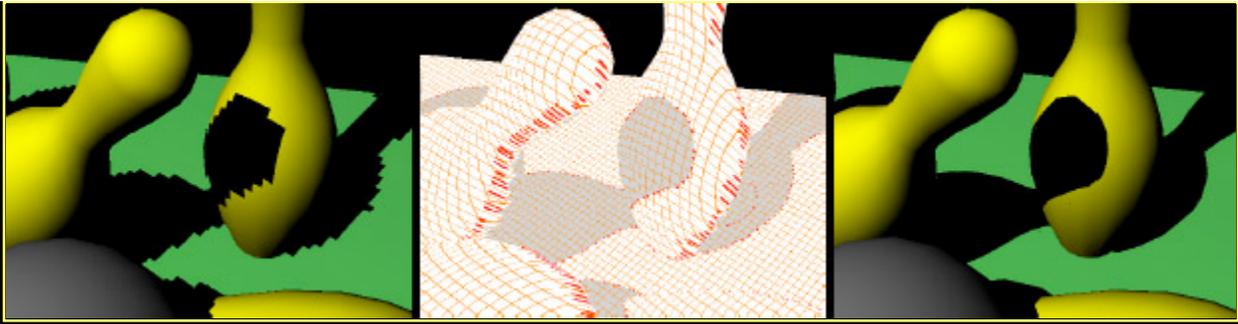
In addition, the natural MIP-map texture filtering performed by a graphics card for small objects causes the edges of a billboard cloud to blend smoothly into the background. A true polygon object would present jagged edges that shimmer as it moves. However, it is sometimes possible to see cracks and distortions up close on a billboard cloud. While the algorithm is advertised for rendering distant objects, its greatest strength may be for rendering objects with many thin parts, like trees and scaffolding. The paper contains images of a tree and the Eiffel tower, both of which appear far nicer than their geometry counterparts typically would in a game.

Watch this Space

The shadows in many games are blurry and blocky because they are rendered with the Shadow Map technique. This algorithm captures a depth image of the scene relative to the light and compares it to the depth from the viewer. A point seen by the viewer and not by the light is in shadow and is rendered without extra illumination. This creates shadows whose blurry jagged edges reveal their underlying square depth pixels. While it is easy to interpret the blurring as a soft shadow, it is really an error created by bilinear interpolation on the depth pixels. To make the shadow boundaries crisp requires more memory than most games can afford to throw at the problem.

In 2001, Randima Fernando, Sebastian Fernandez, Kavita Bala, and Donald Greenberg from Cornell suggested [Adaptive Shadow Maps](#) for dividing up the screen according to the resolution needed. Strikes against their approach are the limited performance that comes from by GPU/CPU cooperation and viewer dependence in the remaining artifacts. In the end, it gets the job done and handles some difficult cases while cutting down on the ultimate memory requirements. Last year Marc Stamminger and George Drettakis of INRIA presented [Perspective Shadow Maps](#) to address the resolution problem. Their method is mathematically elegant but complicated to actually implement, and still fails when the viewer is looking towards the light or shadows fall nearly straight along a wall.

This year's approach is [Shadow Silhouette Maps](#) by Pradeep Sen, Michael Cammarano and Pat Hanrahan at Stanford. They use a conventional shadow map but augment it with geometric information about the silhouette edges of shadows (borrowing a page from the competing shadow volume algorithms). The edges are tracked in a separate buffer from the image and are used to control the sampling process when comparing light depth to image depth. They inherit the other flaws of shadow maps-- low resolution when the camera faces the light and 6x the cost when rendering point (instead of spot) lights-- but the blurry and blocky edges are replaced with crisp shadow edges.



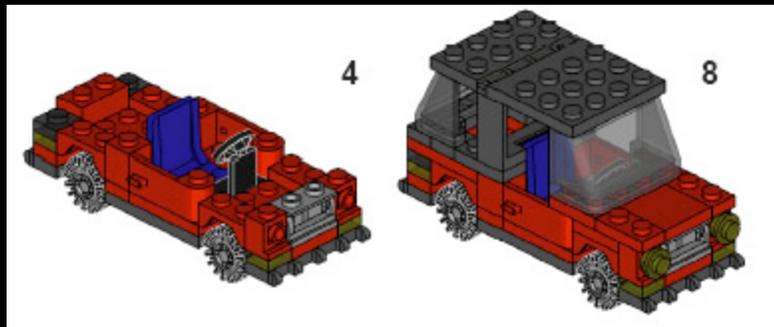
Traditional shadow map (left) vs. Shadow Silhouette Map by Sen et al.

Compared to previous shadow algorithms, theirs is well suited to the overall architecture of modern graphics cards and is straightforward to implement. With direct hardware support it will no doubt replace conventional shadow maps. Without hardware support a graphics card will be bogged down by long pixel shading instructions and be unable to render many other special effects beyond shadows. Branching statements like IF...THEN in programmable hardware, built-in silhouette map support for controlling interpolation, and hardware support for the covering rasterization needed to find all pixels touching a line are possible new directions for graphics cards.

Tomas Akenine-Möller and Ulf Assarsson have been working on [real-time soft shadow rendering](#) for several years at the Chalmers University of Technology. This year they presented a paper in SIGGRAPH and [one](#) in the companion hardware graphics conference on making their SIGGRAPH technique fast. Their algorithm finds the geometry bounding the soft area of the shadow called a penumbra and explicitly shades it using programmable hardware. The method allows volumetric and area light sources, even video streams, to be used and produces beautiful results. Like other shadow volume inspired techniques the implementation is a bit complex to optimize but their paper provides the necessary details.

Blown Away

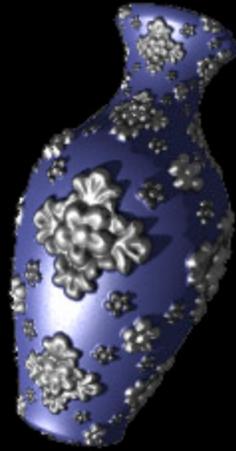
Some papers had particularly awe-inspiring results. Maneesh Agrawala presented a paper he wrote with the Stanford graphics group while at Microsoft Research. Their method for [Designing Effective Step-By-Step Assembly Instructions](#) takes a 3D model of a piece of furniture or machine and produces a sheet of diagrams explaining how to construct it. They even created a page of classic Lego™ instructions using their program.



Kavita Bala presented a ray tracer that uses a method for [Combining Edges and Points for Interactive High-Quality Rendering](#). Created with coauthors Bruce Walter and Donald Greenberg at Cornell University, this ray tracer gets its speed by only sampling the illumination at a scattered set of points in an image. The areas in between are colored by fading near-by intensities. This trick has been employed for years on the Demo Scene. The new work is impressive because it features sharp silhouettes and shadows. The ray tracer finds the silhouette and shadow edges and uses them to

cut cleanly between samples. The cut is at a sub-pixel level so colors do not fade across the edges and edge pixels are cleanly anti-aliased. This approach is strikingly similar to the shadow silhouette map proposed this year for real-time shadow rendering, clearly an idea whose time has come.

SIGGRAPH isn't all 3D. [Adaptive Grid-Based Document Layout](#) by Charles Jacobs, Wilmot Li, Evan Schrier, David Barger, and David Salesin of Microsoft Research and the University of Washington solves the 2D web layout problem for online magazines. Their system takes an article and fits it to one of several prototype layouts then fixes pagination and scaling issues in real-time. Now the New Yorker can retain its distinctive style when viewed under any window size (even on a PDA) and readers can actually read the content without scrolling and hunting ill-wrapped text.



3D models don't have enough detail to simulate the real world so developers use texture and surface shape maps to create complex shading on otherwise flat surfaces. This is convincing within the center of a polygon but the illusion fails at the silhouette.

[View-Dependent Displacement Mapping](#) is a technique for true per-pixel displacement of surfaces by Lifeng Wang, Xi Wang, Xin Tong, Stephen Lin, Shimin Hu, Baining Guo, and Heung-Yeung Shum of Microsoft Research Asia and Tsinghua University. Unlike other methods that introduce new surface geometry around the silhouette, their approach runs a tiny ray tracer (it could also be considered a light field renderer) for every pixel. To make this fast, they store an enormous amount of precomputed visibility data. The results look amazing, but a 128 x 128 pixel detail map requires 64 MB of texture memory. The authors claim it is good for games but the memory requirements are ridiculous for modern games. Some day 64 MB may be a drop in the bucket, but the geometry approach is likely to still be preferred for its simplicity then.

To learn about the other SIGGRAPH papers, visit Tim Rowley's [index](#) of SIGGRAPH papers available online.

Haptics Gone Wild

Haptic input devices provide feedback by rumbling, pushing, or even shocking the user. Gamers are familiar with the simple vibrators built into console controllers. Researchers envision a future where full body suits provide an immersive experience akin to the Star Trek holodeck or Lawnmower Man virtual environment. Two unorthodox inventions in the Emerging Technology Exhibit advance this technology, albeit in an amusing and perplexing fashion.



The Food Simulator by Hiroo Iwata, Hiroaki Yano, Takahiro Uemura, and Tetsuro Moriya from the [VR lab](#) at the University of Tsukuba looks like an opened curling iron with a condom on the end. When the user sticks it in his or her mouth and bites down the device resists, simulating the sensation of biting into an apple or carrot. To complete the experience a small tube ejects artificial flavoring into the user's mouth. Unfortunately, the overall effect is exactly what you would expect. It feels like having a

haptic device wearing a condom in your mouth.

Attendees who were left unsatisfied by food simulation could cross the aisle and strap themselves into [SmartTouch](#). This device is like a tiny optical mouse that reads the color of paper under it. White paper is harmless, but black lines deliver an electric shock. The power of the shock is adjustable by a little knob. Just make sure it is set all the way to the left when you start, and don't forget to sign the provided liability waiver before using the device.

Hail to the Chief

Conference chair Alyn Rockwood from the Colorado School of Mines did an excellent job of hosting a blockbuster conference on a fraction of the budget to which SIGGRAPH had become accustomed. The crash of tech stocks at the end of the '90s and spending cutbacks across the board have shrunk SIGGRAPH. A full day was cut from the schedule to cut costs, resulting in marathon 8am to 8pm days with parallel sessions. This was like a three ring circus where watching one act meant missing another, but the arrangement generally fell into separate conference tracks for artists, developers, and designers. A few cutbacks stood out. The official reception party was held in a concrete gymnasium under fluorescent lights while the convention center empty ballroom mysteriously sat empty just next door. Poor A/V support meant several broken presentations, burned out projectors and audio glitches. Yet even with these problems the quality was far beyond most scientific conferences. Overall, one could have expected to be underwhelmed from looking at the budget, but SIGGRAPH 2003 decidedly exceeded any such low expectations.

One other source to thank for the quality is Doug Zongker and his [Slithy](#) presentation library. Presentations created with this library blew away their Power-Point counterparts by smoothly integrating video, animations, and 3D programs directly into slides. Of course, Zongker shouldn't get all the credit. These presentations are crafted with code, not pointing and clicking, and their authors took the time to do so. Perhaps next year we'll see a paper describing a visual programming tool for making this process easier.

SIGGRAPH returns to its frequent home in LA for the next two years. What is likely to be there? Based on current trends we're likely to see more real-time lighting and shadow techniques, high dynamic range imagery, geometry image techniques, and animation algorithms, all targeted at games. Hopefully there will be a new course on geometric algebra. With a new Pixar film under production and the final Lord of the Rings, Star Wars, and Matrix movies all being released, 2004 should be a banner year for computer generated imagery in films. Either in a special session on these movies or in the Electronic Theater the attendees are likely to see some behind the scenes footage and review the amazing effects. The January 24th paper submission deadline for the next conference is already approaching and poster, sketch, and film deadlines will follow throughout the spring.



This year was packed with content in a short week. If you missed the conference you can catch up with it all on the proceedings DVD and course notes, both available from the ACM. But do it quickly. By the time you've watched all the great graphics it will be time to book your hotel room and head out to [SIGGRAPH 2004!](#)

About the author. [Morgan McGuire](#) is a PhD student at [Brown University](#) performing [Games](#)

[Research](#) supported by a NVIDIA fellowship. He is a regular contributor to flipcode.com.