Holly Rushmeier can be contacted at
IBM Thomas Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598, U.S.A.
e-mail: holly@watson.ibm.com

George Lucas Discovers Computer Graphics
by Alvy Ray Smith

My colleagues and I have been blessed by a sequence of individuals of a unique variety I call *accidental visionaries*—a peculiarly American type, I suspect, and probably more fundamental to U.S. technological leadership than is commonly appreciated. Three of them, Alexander Schure (patron of the New York Institute of Technology's computer graphics lab), George (*Star Wars*) Lucas, and Steven Jobs (cofounder of Apple), have directly influenced my life and helped to make the computer graphics industry robust and pervasive. They all deserve the title "visionary" in the sense that they jumped into the digital fray based almost solely on intuition and paid generously for our years of early experimentation and development. They made it happen. But; I believe they did not recognize the full power of the technology they supported, which is why I call them "accidental visionaries." This story, for example, is about how filmmaker Lucas discovered computer graphics.

By this time, however, it had dawned on me that [George Lucas] did not understand raster graphics.

The time is 1981 in Marin County, California, across the Golden Gate from San Francisco. Ed Catmull is director of the computer division of Lucasfilm, and I am director of its computer graphics branch. Our goal of many years has been to make a completely computer-generated animated film. We believe Lucas hired us to use computer graphics in his films and, hence, help us along the path to our dream. True, he had not explicitly mentioned this as his motivation for engaging us. Instead, his specific request of us was to build three pieces of equipment to modernize the technology of filmmaking:

- a digital film printer (part of which would later be known as a "Pixar");
- a digital audio synthesizer; and
- a digitally controlled video editor.

But surely he must be interested. He had, after all, used computer graphics in his first smash hit *Star Wars*, in the form of Larry Cuba's black-and-white vector graphics—used in training X-wing pilots. And *Star Wars* was known for its use of computers (but no one was quite sure how).

By this time, however, it had dawned on me that he did not understand raster graphics. I had already commented to Catmull that "George doesn't know what he has." Actually, I had said he "doesn't know what he has *either*," since I had made almost the same remark about Schure years earlier at the New York Institute of Technology. Indeed, the use of computers in *Star Wars* turned out to be for control of mechanical devices—repeatable camera dollies for blue-screen shots—not for making pictures directly. Production of *The Empire Strikes Back* was under way, but there was no request that we help with it. In fact, it would have no computer graphics. So, although Lucas was clearly a visionary about digitizing Hollywood's technology, the best I can say about his computer graphics vision at that time was that he allowed me to assemble the best team of computer graphics wizards in the world, probably because he just did not know who they were.

Several of our colleagues from New York Tech, including Tom Duff, Bill Reeves, and David DiFrancesco, had joined us or would shortly. I hired Tom Porter from Ampex, Loren Carpenter from Boeing, and Rob Cook fresh out of Cornell University. Jim Blinn joined us from the Jet Propulsion Laboratory, where I had worked with him briefly between New York Tech and Lucasfilm—although he was to return to the Jet Propulsion Laboratory before the heart of this story transpired. The place was hopping. But there were no requests from Lucas to do what we were *really* good at.

The big break finally came from next door. Industrial Light & Magic (ILM), the Lucasfilm special effects division (at that time, it was working with physical models, not computer graphics), was in the building next to ours in San Rafael, California. ILM produced effects for other companies as well as for Lucasfilm. In this case, the outside company was Paramount Pictures, and the contract work was for *Star Trek II: The Wrath of Khan*. The director expressed his wish for a spectacular effect using this newfangled technology called computer graphics. Luckily for us, although ILM did not know computer graphics, one of its employees used an Apple computer and was, therefore, aware of us computer folks next door. I got the call to go next door for a meeting.

The story of the film hinged on the so-called genesis effect, a phenomenon that instantaneously brought dead matter to life. I quickly concluded that the director and designers did not know what could or could not be done with computer graphics, so I suggested that they let me think about the shot overnight and propose a solution that I knew we could execute. They agreed.

Now, it is important to put this opportunity into perspective. This was to be a major motion picture by a major studio. It would, in fact, become a blockbuster. I had just been given the opportunity to design 60 seconds of this film. Complete novices never get such a chance. But I had it and knew it. I worked sleeplessly all night on a concept employing everything I had at hand—principally an extremely talented team, including Carpenter and his fractal mountains, Reeves and his new particle systems, Porter and his paint program, the 3D rendering skills of Cook and Duff, and the digital filming expertise of DiFrancesco. I was still under the influence of Blinn's *Voyager* flybys of the planets at the Jet Propulsion Laboratory. I threw all this together and sketched a short, six-card storyboard for what would become the "genesis demo" *Star Trek II*. ILM accepted my ideas completely—even giving me another shot to do (the retina ID sequence). This was to be my first directing job for the big screen and our first real movie job, a very lucky break.

At the first production meeting for the genesis demo, I outlined the job to the team and made a statement something like, "We will do a first-rate job of satisfying Paramount and the story requirements. We will captivate the viewing public with our story. But the *real* purpose is to create a 60-second commercial to George Lucas, so that he will know what he has and what we can do for him."

I told them how we would accomplish this latter goal. I was acquainted with Lucas well enough by this time to know that when he watches a movie, he is completely aware of the camera—the choice of angle, distance, movement, etc. This is much easier to say than to do. A director's purpose, after all, is to engage the audience so thoroughly in the emotion of his story that they pay absolutely no attention to the technology creating the illusion. Lucas could resist the director's emotional pull while watching a movie and concentrate instead on the cameraman's choices.

So my instructions to the team were that we would create a camera move that would "knock George's socks off." It would not be a gratuitous 3D camera move that beginners in computer graphics sometimes implement on their first attempt. It would at all times be relevant to the story and help build its emotional force. And, Lucas would know that it could not possibly be executed by a physical camera.

We proceeded to design a move based on the idea of a spacecraft flying by a dead, moon-like planet with a camera attached to the craft. Carpenter was the main contributor to this move, designing a 6D spline (for the six degrees of freedom of a rigid body in 3D space) through about 150 points. It is a twisting, spiraling, accelerating, decelerating, sweeping, reversing, minute-long, continuous camera move that follows the sudden appearance of flame over the limb of the planet and tracks the wall of fire to build tension. The fire overtakes the camera and melts the planet. As the spacecraft pulls away on its outward-bound trajectory, the camera—now looking back—reveals that the molten sphere has been reborn as an earth-like, green and blue, alive, planet.

I will not go into further details of the production and its credits, because I have covered them carefully elsewhere ("The Genesis Demo, Instant Evolution With Computer Graphics," *American Cinematographer*, Oct. 1982). The important point to my story is that the day after the premiere of *Star Trek II*—the special first screening to its creators—the quiet, some say shy, Lucas placed one foot into my office and said, "Great camera shot!" and then was suddenly gone.

We were in Lucas's next film, *Return of the Jedi*, in a 3D holographic shot that Reeves and Duff executed. And Lucas's close friend, Steven Spielberg, used the team shortly thereafter in *The Young Sherlock Holmes*. The team had now been augmented with a spectacular new animator, John Lasseter. From there, digital filmmaking continued to evolve, reaching a crescendo with *Toy Story*, directed artistically by Lasseter and technically by Reeves. Our goal to produce a completely computer-generated film was achieved, after 20 years, in 1995.

Alvy Ray Smith helped begin such famed graphics facilities as Pixar, Lucasfilm, and the New York Institute of Technology. He has received a technical Academy Award for the alpha channel and recently a second one for digital paint systems and has received the Siggraph achievement award. Presently, he is a Graphics Fellow at Microsoft Corp.

Alvy Ray Smith can be contacted at
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399, U.S.A.
e-mail: alvys@microsoft.com

My Inspiration for Dithered (Nonuniform) Sampling
by Rodney Stock

Back in the spring or summer of 1979, I had the good fortune to be dining with a longtime friend, Lynn Conway. I had not seen her in a while—in fact, not since she had gone to work for the Xerox Palo Alto Research Center. At dinner, she excitedly told me about work she was doing on VLSI design and about a book she was coauthoring with Carver Mead of the California Institute of Technology, called *An Introduction to VLSI Systems*. The basic concepts boiled down to the ideas that transistors would be very cheap, wires would be expensive, and tools could be sufficiently automated to allow for design of large custom systems. At the time, I was doing hardware design at Ampex on a frame buffer for a video paint system, but I was still quite interested in 3D hardware design from my days in the early 1970s at Evans & Sutherland, where I worked on the design of ship and flight simulators (alongside two colleagues who also wound up at Ampex during my interval there). Although I was quite busy with video design, in my limited free time I could not help but think about how one could apply Conway's and Mead's work to design high-performance 3D graphics engines.

If transistors were cheap, one obvious idea would be to have multiple processors working on rendering one image. Interleaving processors in image space looked like a decent approach to partitioning the workload, and for point-sampled (i.e., unfiltered or aliased) rendering, this looked like a good way to go. On the other hand, people were just beginning to get fussy about jagged edges and "crawling ants" on images, and the papers coming out of Siggraph were making it clear that spatial filtering had to be done in some fashion to antialias synthetic images.

Filtering required that the interleaved processors communicate with each other if they were to avoid recomputing their neighbors' pixel values. Unfortunately, using lots of communications (i.e., wires) was anathema to VLSI design. This led me to think more about supersampling, where multiple subpixels are computed and averaged together to get a pixel value. This approach had been tried and usually worked fairly well—better than just increasing the pixel resolution—but too often would encounter pathological cases. Ed Catmull had pointed out at Siggraph that even with subpixel averaging, a receding picket fence would always alias at some distance away unless prefiltering were done. I realized that with high-resolution subpixel averaging, the error from one sample could usually be made quite small (with reasonable restrictions on the data input), but that the problem arose when the small errors added up to large ones—I thought of it as the errors being correlated, because the sampling pattern had a high correlation with the image.

That led to the *flash*—the insight that the correlation of the errors was the problem with the pathological cases—if the errors could be uncorrelated (i.e., if the sampling pattern did not