A Cast of Thousands

ILM Reveals Six Techniques for Creating Artificial Crowds.

by Chris Tome

On the day *Star Wars Episode One: The Phantom Menace* opened at movie theaters, many companies in California's Silicon Valley gave their employees the day off. Even banks and insurance companies, it seems, fell under George Lucas' spell.

One organization that might not have so been happy was the Screen Actors Guild. After all, it's not easy to get digital actors to pay their yearly dues. And make no mistake about it, *The Phantom Menace* is chock full of non-union thespians. The movie swarms with them—from humanoids to droids to alien swamp creatures.

Notwithstanding prominent CG life forms such as Jar Jar Binks and Boss Nass (the subject of last month's in-depth character study), most of the film's artificial actors appear in crowd scenes. Concocted by the ingenious team at Industrial Light & Magic (ILM), these sequences present an utterly convincing cast of thousands, even when the actual number of flesh-and-blood actors onscreen is no more than a dozen, and sometimes none at all.

How did they pull it off? In more ways than you can shake a light saber at. Some didn't even involve computers, such as the attentive audience of painted Q-Tips that fill the stadium bleachers during the pod race on the planet Tattooine. Digital technology did fill the stadium grounds and aisles, however, where a throng of animated human and alien figures was emitted by a Maya-based particle system that calculated their paths on the fly using artificial intelligence. Meanwhile, a herd of otherworldly animals stampedes through the Naboo jungle along fixed spline paths that were laid out and tweaked by hand. In the enormous spherical Senate chamber on the planet Coruscant, where the Trade Federation debates the fate of Naboo, preshot and prerendered delegates were projected onto an array of 3D cards. The climactic battle between the amphibious Gungans and the mechanical droid army involved an immense field of motion derived from a small number of mocap files. And, in what might be the simplest method conceptually—but a technical challenge nonetheless—the adoring crowd that lines the streets cheering the triumphant Gungan victory parade actually consisted of many small groups of live actors filmed and composited into positions determined using Softimage3D.

We'll look at each of these techniques in turn, delving into the thoughts that led to it, the tools used, and the obstacles overcome. The very scope of solutions ILM's artists and animators found to the common problem of how to make artificial crowds illustrates the no-holds-barred creativity they bring to their work. Their efforts were directed toward only one purpose: to make the shot work.

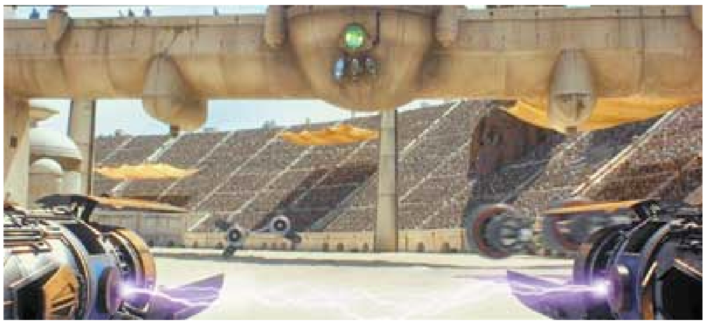

Practical Populations

Sometimes all it takes to make the shot work is a bunch of bobbing Q-Tips. That was the conclusion reached by Michael Lynch, who has served as chief model maker at ILM for the past nine years, when he set out to populate the stadium where Anakin Skywalker bests the alien Sebulba in a hair-raising pod race. Lynch comes from the old school of practical effects. He might even be considered a dinosaur by today's high-tech standards, but one that, like the alligator, is far from extinct.

"As much as you can get in camera initially, the better," he points out while making the case for practical effects over CG. "We built the entire stadium practically, and if you can get at least some of the crowd in there, you get interactive light, shadows, and so on." Even if the practical crowd becomes replaced by CG imagery later—which, in some shots, it was—the CG artists benefit from having a mock-up of the entire setting, shot in natural light, available for reference.

Ground Battle Sequence Stats

A few statistics from the production of the ground battle between the Gungun army and the FederationÕs battle droids:

Number of cycles:

- Gungan cavalry/kaadus 27
- Gungan infantry 55
- Bdroids 48
- Ddroid, faamba, fulump, catapult, aat 13

Total number of cycles/ animations checked-in and baked: 143

Total number of frames for all cycles: 27,534

Average number of frames per

Lynch began with sketches created by design director Doug Chiang and architectural plans by architect Bill Beck. "The finished product looks very much like the original drawings," Lynch says. "We were well provided for in terms of information. A lot of times we don't get that much, but in this case, Doug Chiang really knew what he wanted."

Under the direction of model supervisor Steve Gawley, he built the 32 x 42-foot stadium model. Then he set out to fill the seats with hundreds of thousands of spectators. First, Lynch cut head-and-shoulder silhouettes out of cardboard and attached them to cams, which he stacked to make rows. The cams allowed him to make the silhouettes undulate for a more realistic effect. It worked well, he found, in shots where the camera was low and looking up into the stadium. However, higher camera angles destroyed the effect, so he discarded this approach.

Instead, he decided to populate the seats with cycle: 192.5

Total number of .rib files baked (includes multiple resolutions and cycles with multiple creatures): 87,490

A typical shot (code name ggbb038):

Amount	Description	Number of cycles used
3	Faamba	2
248	Infantry throwing balls	1 (almost hidden)
114	Infantry with shield in front	3
2,048	Walking and dying droids	13
2	Balls thrown from catapult	1
395	Balls thrown by infantry	1

thousands upon thousands of Q-Tips—335,000 at final count, which he dipped into different colors of paint a fistful at a time. Stapling lengths of hardware cloth to the stadium, he dropped Q-Tips of various colors into the cloth mesh, which held up their heads while leaving their four-inch dowels hanging below. An electric fan positioned behind and below the bleachers made the dowels sway, giving the Q-Tip crowd a lifelike motion. "The dowels acted like pendulums," he says. "If you wanted to, you could make them do a wave."

Before launching into the project full-on, Lynch built and populated a small section of the stadium to see if this technique would be cost-effective. He installed roughly 1,500 cotton-headed spectators in a space roughly 1.5 feet square, which gave him a sense of how long it would take to fill the stadium. This section took nearly an hour and a half, including painting the heads.

"Eventually, we built an entire stage for this shot outside," he explains, "with an airplane hangar on a track to cover it up at night. It was placed on a rotating base so we could position the set properly in natural sunlight."

Despite his ingenuity and hard work, Lynch was aware that his creation might be replaced by CG effects. "This was really just gravy," he says. "They could use as much of it as held up and fix whatever didn't work on camera. It was just something they didn't have to do."

While Lynch was positioning Q-Tips, a film crew documenting the making of *The Phantom Menace* moved in for a close-up. He began to count, "Three hundred twenty thousand and two, three hundred twenty thousand and threeÉ" Then he reached for a paintbrush with one tiny hair on it, and, heaving a deep sigh, he said, "Now it's time to paint the faces!"

Particle People

Q-Tips may have worked in the bleachers, where audience members were seated and visible only from the neck up. Populating the stadium floor and aisles with a milling crowd, on the other hand, required methods at the opposite end of the technology scale. In these shots, the team composited into the

shots 3D imagery generated by a custom plug-in to Alias|Wavefront Maya and driven by MEL scripts. The plug-in made it possible to substitute geometry for particles, while the MEL scripts controlled their paths intelligently according to complex rules that mimicked the behavior of people in a crowd.

Mocap Techniques for a Disposable Army

Motion capture is a volatile subject among 3D artists. Many see it as a poor substitute for keyframing at best. Others envision a dark future in which mocap will have replaced their artistry altogether. While mocap can be a panacea in some cases, it's a nightmare in others.

Nobody understands this better than technical animation and droid supervisor James Tooley, who oversaw the capture and application of motion files for the droid army in the ground battle between the droids and the Gungan cavalry.

"The droids used a lot of motion capture, and they also had to be sliced and diced by lasers, so I dealt with the dynamic simulations for that too." Tooley explains. "We wanted to be able to produce as much motion as we could. Once you have control over the technical aspects of getting the motion files onto the models, the rest is just a matter of how fast can you capture motions, clean them up, and get them into the animation."

Mocap, which is designed to capture the subtlety of actual muscular motion, isn't the most obvious choice for animating mechanical robots. However, Tooley found that it worked beautifully. "The droids only have two bones in their spines," he observes. "If you apply a curvy motion from a human onto that, you lose resolution in the spine, so it ends up looking mechanical, which is fine for the droids."

The optical mocap process they used

Initially, visual effects supervisor John Knoll enlisted Habib Zargarpour to help with the pod race—more precisely, the pod crashes. Knoll explained how he wanted the high-speed vehicles to crash, and Zargarpour created physically accurate rigid-body simulations.

Zargarpour had been visualizing instanced objects that appeared as tiny squares in Alias|Wavefront Dynamation or dots in Maya. Then Hiromi Ono of ILM's R&D department wrote a Maya plug-in that drew the objects themselves directly in the Maya GUI. This made it easy to add thousands of pieces of crash debris and preview their scale, rotation, and so forth, making the creative process far more interactive.

"One day I thought, ÔHey, wait a minute! We could stick people on these things,'" he recalls. "That thought segued perfectly into doing crowd simulations." Ono's plug-in was fast enough to render over 4,000 relatively small people in Maya's hardware rendering mode, including textures and shadows, making it possible to preview complex particle-based crowds in real time and even avoid exporting to Pixar Renderman for test rendering. Each crowd member was not quite an instance, but one of a selection of animated geometries. Moreover, their motions could be controlled by logic composed in a combination of Maya's dynamic expressions and custom code.

For his first simulations, Zargarpour used battle droid models, as they already had a library of mocap animation cycles. In one early test translated the positional data generated by optical sensors attached to the mocap performer's joints and triangulated their positions in 3D space using specialized cameras. "We chose optical over magnetic because there's no umbilical cord and no backpack involved," Tooley explains. In addition, he points out, optical is more precise and less prone to magnetic distortions.

Tooley and his crew had to plan their work carefully to capture the motions they needed. In the one sequence in the ground battle, the droids push through the Gungan energy shields, straining on one side, then nearly falling through to the other. Mocap supervisor Jeff Light came up with a clever approach to creating the right kind of motion. "We had a rope harness on the performer, and as he moved forward, somebody held the rope, which gave a leaning-over-forward look. Then they just let go of the rope, and the actor had to take a couple of steps to regain his balance. Of course, because it's mocap, you don't see the rope or the person holding it."

For motions that were intended to be applied to Gungans, a special technique was used to deal with the difference between human and Gungan anatomy. "The Gungans had really tall necks," he points out. "To extrapolate from the human head to the Gungan head, we had a specialized rig that would offset the motion and expand the rotations, or otherwise amplify the motion, to make the head work correctly."

After the data was captured, it was fed through a proprietary conversion program called MC2ANIM. The software automatically scales the data to an IK-constrained character model and applies it in a way that, according to Tooley, makes it look less mechanized and simplifies tweaking the data. "A lot of the battle droids and Gungan warriors had such good conversions, that they were almost plug-and-play," he relates. "The resulting animations didn't have any noise problems or glitches."

"For the hand to hand combat," Tooley continues, "they were keyframing a lot of complex motion. Although we did use some of the keyframed data, we went back to the he showed me, several droids walk over hilly terrain. "If the hill's not too steep and the characters are walking fast enough, they can get over it," he points out, describing the logic that governs their behavior. "But if it's too steep, they'll go around it. You can imagine what it would take to animate this by hand."

Eventually, he replaced the droids with human models. Mocap cycles were attached to the particle-based people, which were exported as prebaked Renderman .rib files. Renderman can combine .rib files under scripted control at render time. Using separate .rib files made it possible for Renderman to load character models only as needed, which saved time and processing resources and also simplified workflow should any changes need to be made later.

"Let's say you have an object, and you want to include 1,000 copies of it in a scene," Zargarpour explains. "One option is to instance the object, but that requires a lot of rendering overhead. Another option is to encode it into a single .rib file. Renderman can merge chunks of data from various .rib files into the main .rib, and you can set it up so it loads the object only when it comes into view. The person that made that workable was Christophe Hery," CG sequence supervisor for the ground battle between the battle droids and the Gungans, in which the immense number of characters made rendering a special challenge.

The aliens in the sequence were keyframed by hand, match-moved, and composited later. To add the final degree of realism, live-action footage was composited in, seamlessly mingling real and CG characters. Sometimes Zargarpour himself had difficulty differentiating human actors from their digital equivalents. "We made takes in which the CG people would flicker," he says, "so we could tell what was real and what wasn't."

It took some experimentation to finalize the logic behind the crowd intelligence. In one test, as the crowd meandered toward the stadium doors, characters bumped into walls repeatedly before finding the entrance. "That was when we didn't have much intelligence built stage and reshot two people actually struggling hand-to-hand. Then the animators came in and helped clean it up. Maybe a hand wasn't completely connected to an arm because of differences in the physiology of the mocap actor and the model, so we had to do minor reconstraining to replot the droid's hand back to gripping the Gungan's arm, for example."

An extensive array of proprietary tools made the mocap process more efficient by giving Tooley and his team a great deal of flexibility in how they applied captured motions. "Using a proprietary tool we call MoJo," he reveals, "we could take multiple pieces of mocap, align or offset them in time, and blend from one to another. For more complicated blends, we used another in-house tool to change the path of the motions. That is, the actor may have been mocapped straight down a path, but we may apply the motion to another path. MoJo helped us take care of the leg placement and so on."

Yet another in-house tool helped blend the start and end of the mocap cycles. "A person who does a walk cycle won't have the start and end motions exactly the same," he explains, "so we used a shearing tool that compared the beginning and end frames and automatically blended between them to create a seamless loop."

In other sequences, re-using motion threatened to impart a uniformity that worked against the impression of a believable crowd, but the droids turned it into a virtue. "They were what George called a disposable army," Tooley says. "They had to look like they came from the same mold and had the same programming. I decided that if we were going to use mocap, we had to use one person to do all of it. If you get somebody who's really good, you can do a piece of motion and then have them do it again ten days later, and it will be almost dead on.

"It's not as easy as it looks," he continues. "It's like acting in a way; you have to hit certain beats at certain times. But it's also making your body move just right, like pantomime. I regard it as an extension of traditional animation, but instead of using a mouse to get my input, I can use my whole body."

in," Zargarpour says. "They try one side, then the other side, and finally they make it in. Of course once they make it in, they die, because we don't need them anymore," he concludes cheerfully.

Little by little, Zargarpour gave crowd members the intelligence to find their way through the stadium doors, through corridors, along railings, and the like. To make them even more lifelike, he added behaviors that enable them to make friends. When two characters come into close proximity with one another, they can become companions and walk in tandem. As they approach other characters, they can leave their original partners and join new ones. This lent the crowd's motion complexity while avoiding the impression of randomness.

One of Zargarpour's most ingenious techniques involves "carrots and skunks." A carrot is a null object that tends to attract individual audience members. A skunk tends to repel them. This made it possible to control the crowd's overall motion without determining it absolutely.

How does Zargarpour approach the task of programming rule-driven automated animation? "I imagine what the rules are," he replies simply. "Whether walking can go to running, or dying, or whatever. Then I ask, what conditions cause a transition to happen? Then you start getting into the trench of actually doing it, and that's where you're going to hit stumbling blocks. The computer is only able to do what you tell it to, nothing more."

More>>

It's surprising to find a technician with such a fine appreciation of performance, but Tooley knows what he's talking about. Asked who performed the droid motions, he replies, "You're looking at him."

A Cast of Thousands

(Continued)

Cloned Creatures

Compared to the artificially intelligent digital humans in the stadium, the creatures that inhabit the swamps of Naboo are brain-dead. When Federation ships land on Naboo, sending ground troops and battle droids barreling through the forest, a stampede ensues in which all manner of alien beasts run along fixed paths, every frame of which was analyzed and tweaked.

The heaviest shot in the sequence includes roughly 35 creatures—an assortment of peko pekos, pikobis, motts, ikopis, nunas, and fulumpasets, all of which were derived from a wire model of a dog in the ILM database. "They're not in the hundreds like in other sequences," remarks CG supervisor Tom Hutchinson, "but it was pretty significant."

Hutchinson's background prepared him well for animating large photorealistic creatures. His accomplishments at ILM include work on Jumanji as well as the sequences in *Jurassic Park* in which two raptors tear through the kitchen and the T. Rex chases the lawyer to the bathroom. For the stampede sequence, his team included five animators and six technical directors (TDs). "It fluctuated a lot, though, because at the same time we were working on the underwater chase sequence," he says. "We pulled TDs back and forth, and our R&D people were working on the ground battle sequence at the same time."

Using a proprietary application called Choreography, Hutchinson and his team attached animation cycles to the models and set them running along spline-based paths. "We could get lots of creatures going and manipulate their paths without worrying about the geometry at the same time," Hutchinson says. "The schedule on this was a fraction of the time we had on *Jumanji*." Whereas Maya was used extensively in other crowd animation shots, Choreography and Softimage|3D were used to visualize the stampede.

Choreography is fairly intuitive, Hutchinson explains. "You load up the 3D environments, basically the match moves from a Softimage file, then lay out a 3D spline path where you want the characters to move. Then you assign animation cycles to each of the various paths. From that point on, you can load snippets of animations or cut them out in a very interactive way."

The majority of the animation time was spent making the basic animation cycles as realistic as possible. These cycles were copied among animals of the same type and then enhanced with head turns, jumps, moving around obstacles, and so on. The animators worked together, with some of the team focused on blocking out the animation and others whose job was to tweak secondary motions and other seemingly minor yet critical details.

Since the cycles were freely duplicated, it was necessary to vary them to keep the crowd from looking mechanical. "That was ongoing, and we really had to discover what worked, depending on the shot," Hutchinson recalls. "Once we got one cycle of animation, between 60 and 120 frames, [visual effects supervisor Dennis] Muren would say, ÔOK, let's pick a couple of others, and we'd vary them to make a library of three or four cycles for each character. Let's give this guy a limp or whatever.'"

The skeletal setups posed their own challenges, especially that of the lumbering fulumpasets. "They've got no real mass over their hips," Hutchinson observes. "They're designed for looks rather than for proper locomotion. They really work best coming head-on."

Much of the technology involved in making these animals move realistically evolved from research and development done for Jurassic Park. "The software has been progressing at a pretty quick clip," Hutchinson says, "so more of the power comes for free, depending on how the bone structure is animated and enveloped."

Apart from the animation, creating the environment itself was a major effortÑeverything in the sequence was 2D except the animals. "There's a whole background plate that was synthesized mostly from still photography that we cut and pasted together," Hutchinson says, "along with a lot of work from the matte and art departments." To make trees fall and break apart as the animals ran by, he sought out tree-cutting crews and filmed them at work. "We had a few windy storms in the area," he recalls. "We got in touch with power company ground crews and shot them cutting the trees down. Then we comped them into the sequence."

More surprising still, the space ships were practical models. This posed a problem as the work progressed and it became clear that the scale in which they were filmed wasn't appropriate. "The scale depends on the camera's lens and angle, as well as how distant it is from the object," Hutchinson explains. "If you're assuming a ship is going to be a certain distance away, it may be moving too fast for the scale it was shot in. As we developed the sense of scale and progression, we had to go back and reshoot them."

Committee of Cards

If the way in which Naboo's 3D swamp creatures were composited against 2D background elements has become conventional in cinematic effects, the melding of 3D and 2D in the Senate chamber turns convention inside out. The chamber, its cavernous interior lined with innumerable pods that hold delegations from across the galaxy, is literally a house of cards. Each pod holds cards onto which the ILM team projected precomposited imagery.

As the sequence progresses, it becomes evident that the pods can detach from the wall and fly into the center of the room, where the representatives they carry can address the Chancellor or others face to face. The flying pods were made up of a combination of 3D models and practical models.

Development of the chamber began late in production, in November 1998, when CG sequence supervisor Steve Molin started making previsualizations based on Doug Chiang's sketches. Molin moved to ILM eight years ago from a defense contracting company, where he got his first taste of 3D on Silicon Graphics Irix boxes. Previsualization was critical to the scene—indeed, to the entire film—because the animatics provided a basis for the 3D work.

"Based on the animatics," Molin explains, "we did camera match moves for the whole sequence, so we could get a handle on which camera angles we were going to need. Before we started, we had a pretty good idea what needed to happen."

What needed to happen was not simple: First, the team had to rustle up enough characters to fill 1,024 pods, each containing between one and five humans and/or creatures. Then the delegation occupying each pod needed to be shot or rendered from a "comprehensive" variety of camera angles, sufficient to cover all of the camera motion planned for the sequence. The footage was composited onto cards, and finally the cards were switched from footage shot at one angle to footage shot from another to match the camera moves.

"We didn't have that many unique creature elements," Molin recalls, "so we had to do some repetition to fill the whole scene. Many of the creatures from the film were reused to populate the Senate chamber, such as the creatures in the pod race.

"If you look at our database of image elements," he continues, "we have the creatures shot from different elevations and azimuths, and then we distributed them based on the camera angles using a MEL script in Maya. The images that corresponded to the camera angles were mapped to flat planes and composited into the scene using alpha channels."

The camera angles used in the footage didn't always align with the match moves exactly, and subtleties such as shadows and lighting weren't precise either. But Molin didn't consider this a problem. "This stuff was far enough away from camera that we didn't really worry," he admits.

"When we were starting work on the sequence, we looked at the plates they had shot for us, mostly for the hero characters"—that is, those who are relatively close in frame—"and tried to select a generic lighting environment for them. That way, as we changed the angle of the people to match the camera motion, they had a repetitive lighting model. As we got into the closer stuff, we stayed with the lighting arrangement that matched the lighting of the main characters."

Ultimately, careful planning was the factor that enabled the team to complete such a complex task in a very short time—but it didn't hurt that the setup was the part Molin found most rewarding. "I got a kick out of being able to set it up so that when the TDs started running shots, it was easy for them to generate imagery in a fairly short time span. For some of the easier shots, we had a three- or four-day turnaround, which is just amazing from my point of view. Anticipating problems and helping people solve them was probably my favorite thing."

Duplicate Droids

By far the most complex sequence in *The Phantom Menace* is the final ground battle between the Gungans and the battle droids. As the Federation prepares to wipe out the brave but seriously outgunned Gungans, 7,000 characters pour across the screen in a single shot. The fact that they were all derived from only one battle droid and one Gungan didn't make matters any simpler: the combined weight of all the models, motions, and elements amounted to a staggering 70 gigabytes.

With thousands of characters driven by hundreds of animation cycles and distributed at varying distances from the camera, previsualization—in stark contrast with the Senate chamber sequence—was no help at all. "We started with an animatic," explains CG sequence supervisor Christophe Hery, "but that's better for hero shots than it is for large crowd scenes. A lot of research was necessary for each shot, and at the same time we were doing the research and design for the tools we were building. It was kind of scary to see everything, not knowing how it would come together."

Hery is a Renderman master, and Dennis Muren approached him in June 1998 to develop strategies for rendering the battle sequence's vast amount of data. From there, he went to work increasing the efficiency of the crowd animations. It's not exaggeration to say that many of these shots would not have been possible without his optimizations.

The battle was laid out in Maya and driven by MEL scripts that allocated animation cycles to groups of characters while conforming their animations to the terrain and avoiding collisions. The basic animation library was imported from Softimage|3D with details such as flapping Gungan ears and cloth simulations added in Maya. All the Gungan warriors were derived from one model and likewise for the battle droids, with seven variations of coloring and texture and variations in scaling and proportion added randomly via MEL scripts.

The Gungan warriors' animation cycles were mostly keyframed. The battle droids, on the other hand, were motion-captured by technical animation supervisor James Tooley. "I deal with animating things that tend to be mathematically or technically difficult," he says. Tooley, who came to ILM five years ago by way of Disney, also served as droid supervisor at large.

As in the shots of the stadium floor and aisles, plug-ins incorporating technology developed by Habib Zargarpour provided very fast animation previews. However, unlike the characters on the stadium floor, an enormous number of characters in the battle sequence needed to be seen relatively close in frame, so optimizing rendering time was critical. "We could see a hardware render of the crowd in Maya almost in real time," Hery reports, "so we were able to refine it as much as possible before we hit the render button."

One such refinement involved distributing droid animation cycles judiciously, so the smallest amount of data could generate the greatest range of motion at render time. Many droids shared exactly the same motion file, which was appropriate for a phalanx of mass-produced fighting machines designed to serve as a disposable army. However, Tooley made sure to introduce variations. He set up various prebaked .rib files, each containing the same motion file augmented by a head turn to the right to survey the landscape, a head turn to the left, or another subtle difference.

Having optimized the crowd animations, Hery took a careful look at the lighting setups. "We came up with a compromise that would give it a good daytime look," he says, without unduly contributing to the time required to render. "We used between five and eight lights, depending on the shot, not all the lights cast shadows."

To create interactions between the characters and the grassy field in which the battle takes place, the team began with photographs of matte paintings depicting the field. "All we had to do was make it look like the foot was inside the grass," Hery explains. "It was just a simple shader with some noise to break up the edge of the bottom of the feet." No collision detection was involved—merely a visual trick, and a cheap one at that, at least in terms of render time.

Rendering took place in three passes: a beauty pass, a cast-shadow (or depth) pass, and a foot-in-the-grass pass. For the beauty pass, Hery set up Renderman to load geometry on an as-needed basis. "We baked the .rib files with animation and motion blur," he explains, "but instead of baking in the geometry, we used a feature called delayed read archive, which is also known as procedural primitive. When you start the .rib, there's no geometry in it. Renderman checks for bounding boxes that we provide, and if a bounding box intersects the pixels currently being rendered, it loads the geometry."

To further conserve rendering resources, Hery took advantage of Renderman's Level Of Detail (LOD) capability, which automatically selects among low-, medium-, and high-resolution versions of a model depending on the amount of screen space it occupies in a given frame. "The idea is that the low-res version will appear when the model is far away, while the high-res version will appear when it's close in frame," Hery explains. "In some cases, we pushed the medium-res version as close [in frame] as we could to conserve memory." It was necessary to make sure various LOD versions of a character had the same skeletal setup so the animations would transition properly from one to the next.

Not all the battle's close-in-frame characters were rendered in this manner. Some were treated as hero characters: keyframed in Softimage|3D and Cary (ILM's proprietary facial animation tool), rendered separately, and composited in post. However, many of the characters that appear close in frame actually were treated just like characters with less prominence. This made Hery's job much simpler—otherwise, there would have been too many hero shots to manage.

Composite Crowds

Every crowd-scene technique we've examined so far required that the synthetic extras be three-dimensional. For the crowd that lines the street during Gungan victory parade that closes the film, the ILM team decided to use 2D extras. Shooting several small groups of people against a bluescreen (for film) or greenscreen (for digital video), the shots were composited to create two dense rows of adoring onlookers. The tricky part, though, was to position these shots coherently into the 3D scene—a task that fell to CG sequence supervisor David Meny and compositing supervisor Greg Malony.

"The crowd extras were shot in blocks of six by four or six, so we had more than 100 squares of 30 people in different variations," explains Meny, who began his career at ILM four years ago after graduating from California Polytechnic University with a summer intership and parlayed that into a job as animator. After working on the ground battle sequence, he was charged with helping to animate the parading Gungans as well as the CG confetti.

"We could have eyeballed them together," he continues, "but that would be painstaking with more than 100 elements just for one shot on the street. Then I got the idea: We had the CG set in Softimage|3D to animate the Gungans parading down the center of the street, so why don't we just lay a grid on the plaza floor to determine where the 2D crowds should go? As long as the grid had enough resolution, we could figure out where each crowd element should be positioned. That made it easier for the compositors, since they could tweak two numbers instead of having to deal with exact pixel locations. The grid also helped them adjust the scale of the crowd elements toward the horizon. It made things much more manageable."

Maloney, who worked as a photographer before joining ILM's optical department in 1989, agrees. "By tracking it this way," he says, "we were able to take a radical match-move with a moving camera and lock the greenscreen elements to those 3D points in space. A lot of stuff we did depended on having a really high quality match move." The artists brought the plates into Softimage|3D. Then they got the floor-plane coordinates to overlay the images on the plates and locked down the 2D elements. "After we got our first row of people on," Maloney says, "it became a lot of noodling and hand work."

"For the big shots where you see all the way down the street," Meny adds, "we were fortunate that most of the camera angles were lockoffs. Only one shot was a pretty crazy move, where it started high, tilted down, and followed the parade to the right. It makes life easier if you have a lockoff shot. Then you can weigh the benefits of using 2D, 3D, or a mixture."

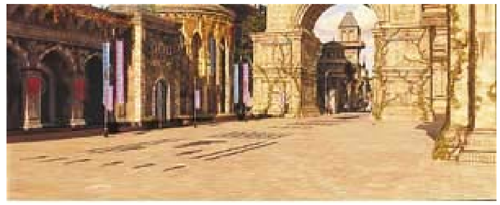

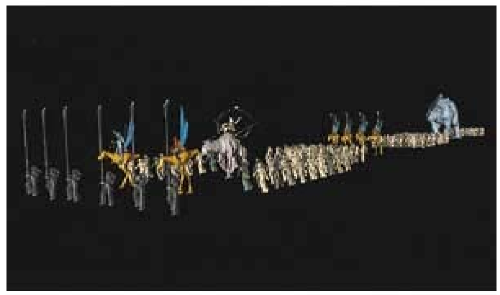

The compositing challenges of the parade sequence didn't stop there. "In the film we got from Pinewood Studios in England," Maloney recalls, "most of the parade shots included less than half the set. Usually, it was just the base of the buildings, up to eight or ten feet high. A lot of stuff just wasn't there."

To make matters more challenging, the original plates were shot outdoors over several hours. "If you look at the end parade with a critical eye," he points out, "you'll see the shadows going all kinds of ways. Some shots were shot on overcast days, some on sunny days, and so on. We had to bring all those plates into into some kind of visually congruent space." Maloney got the matte paintings, the model shots and the CG to match as closely as possible, working with color, matting, and isolating different areas to correct for the inconsistencies in the lighting.

To correct the lighting, Maloney sat down with visual effects supervisor Scott Squires and decided when it looked right. "It's a sort of a cheat," he comments, "color correcting, giving it the right contrast, and making the black levels match. It's all those fine little points that sell it."

Crowds In Perspective

And sell it they did, though not always to Maloney's satisfaction. "You can get 80 percent of the way on a shot in a few days. The next 10 percent take another four, and the final 10 percent a few days more. By the time it gets on the screen, you may not see half the things you worked for weeks on. Keeping your

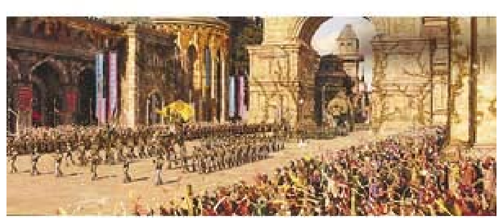

standards very high gets you to a point you can be happy with. We feel like we've pushed it as far as we need to. But it's my job to push and push and push until the supervisor tells me to stop."

He pauses for a moment. "It really is an evolution," he reflects. "Every show builds on a show before it. Two years ago I don't think we could have done this movie."

The range of solutions devised to achieve these awesome scenes was as diverse as the people who thought them up. George Lucas waited many years until he thought graphics technology had reached the level sufficient to realize his cinematic visions, but without the people to make it happen, no amount of computing horsepower would matter. *Star Wars Episode I: The Phantom Menace* serves as a wellspring for CG artists seeking to define how good is good enough, and how good is great. The team at ILM knows the meaning of greatness. May the Force be with them!

Cast of Thousands

The parade sequence, in which the Gungan army parades to Queen Amidala's palace to celebrate their victory over the Trade Federation's battle droids, was assembled from a variety of practical, video, and CG elements. The buildings on the live-action set were no more than eight feet tall; they were augmented by CG buildings in the background plate [A] with a cast shadow pass on the ground. A test using low-res meshes of the Gungan army was rendered [B] followed by a test using hi-res models [C-D]. A crowd of spectators [E-F] was shot in small groups of roughly 30 people and composited into the scene using Softimage|3D to position of the groups. No parade would be complete without confetti [G], most of which was created in 3D and comped in later. The final comp shows all the elements [H].

~~**Chris Tome is the editor in chief of DesignFreak.com.**~~ ^

(This article originally appeared in the September, 1999 issue of *3D* magazine.)