Cover Story

VOLUME RENDERING

Volumetric rendering speaks volumes for data 20 orders of magnitude apart—from human anatomy to neuroanatomy, and from electrostatic charges of macromolecules to failure analysis of manufactured parts.

KAREN A. FRENKEL

A young man in his late twenties suffered a crushed pelvis in an auto accident. His orthopedists said that the fracture was too complicated to operate on and elected to treat him conservatively; he would be in traction for a few months. The doctors were certain that the young man would be permanently crippled.

Luckily, the man's father, also a physician, knew of research in 3-D rendering of computed tomography (CT) scan data. He sent his son's CT scan studies to the researchers, a radiologist, an orthopedic surgeon, and a computer graphics expert, who studied the volumetric rendering of the pelvis that was created with specially designed hardware and software. Able to see it from all angles, they determined the extent of the fracture and the locations of several key fragments. The pelvis was operable and the next day, the surgeon set the fragments. Three months later the patient returned for a check up and demonstrated full-range hip motion.

This case coupled great medicine and great computer science. The technique of volume rendering changed the course of treatment by providing the physicians with more data. This data ultimately gave them the confidence to operate and thereby improved the patient's quality of life. While volume rendering helped manage the medical complexities, this case also represents departures from tradition for both disciplines.

In the medical realm, radiologist Elliot K. Fishman, Director of Computed Body Tomography at Johns Hopkins University, has been pioneering volume renderings of CT (also known as computerized axial tomography) data for three years. But not all radiologists, whose fundamental training involves interpreting two-dimensional data into three dimensions, have embraced this new technique. Many are concerned that computer-generated artifacts and pseudo-color can lead to misdiagnoses. Widespread use and acceptance of volume rendering has also been hindered by competing scanner vendors who have kept their data formats proprietary. This has forced those on the computer graphics end to reverse engineer tapes of data to uncover the formats. On the other hand, surgeons can benefit from the precision that volume rendered CT and magnetic resonance imaging (MRI) can yield.

On the computer graphics side, Bob Drebin and his colleagues at Pixar have departed from tradition by abandoning *surface* rendering, which has its foundations in geometry based modeling. They maintain that volumetric data should not be skimmed to yield only surface renderings. They have developed new algorithms that take full advantage of 3-D arrays of data, rather than just using the surface data found in such arrays. Their approach also reflects the early computer graphics dilemma over slowly generated, photorealistic graphics versus fast, less detailed image processing.

With the imaging and graphics application market expected to reach $1.6 billion by 1990, according to Dataquest, several graphics and medical imaging vendors are merging or teaming up on large projects. In 1987, Sun Microsystems, Inc. bought Trancept Systems, Inc. to form a graphics accelerator division in Research Triangle Park, North Carolina, where the TAAC-1 add-on for Suns was developed. Also that year, Philips Medical Systems, Inc., in Shelton, Connecticut, formed a partnership with Cemax, Inc. in Santa Clara, California, and Island Graphics and Pixar, both in San Rafael, California, to launch Project Pegasus. Fifteen medical centers (including Johns Hopkins) are exploring different challenges in medical imaging, evaluating equipment, and helping Philips develop hardware and software. Another project, named for the Renaissance artist Leonardo da Vinci, is a database of the entire human body on a Cyber 910. A joint effort of Control Data and the University of Illinois, Chicago, the project integrates the activities of medical illustrators, anatomists, radiologists, and biomedical engineers. Some researchers, like Craig Upson of Stellar Computer, Inc., in Newton, Massachussetts, say that volume rendering is doing for

Volume renderings of the broken pelvis using CT scan data by Elliot Fishman of Johns Hopkins. The extent of the fracture and location of the fragments is clearly visible. Although radiologists have been using CT data for almost 20 years, volume renderings of CT offer a new way of interpreting such data.

computer graphics in the late 1980s what the introduction of perspective did for drawing and painting during the Renaissance.

Why is volume rendering more feasible now then ever before? In volume rendering, computer graphics has found a general means of visualization that is effective for *two* types of data—real (measured) and numerical (calculated). Huge amounts of data, collected a variety of ways, can be processed by miniaturized chips with immense, and often parallel processing, power. Besides CT and MRI data, 3-D measurements are now taken with the help of solid state cameras and lasers, and improved microscopes. One new device, the very high resolution, confocal microscope, is expected to vastly expand cell biology experimentation.

Neurobiologist Vincent Argiro of Maharishi International University in Fairfield, Iowa, who has used confocal data to volume render neurons, says, "Volume rendering is one of the most direct, straightforward ways of doing three-dimensional image processing. It's computer graphics on the one hand, but it's image processing on the other. You're dealing with real data about real structures. These are not synthesized pictures that just come out of a mathematical equation." Further, he sees volume rendering of any kind of data as a catholic approach to understanding a huge range of phenomena in nature. Far from belittling numerically based volumetric simulations generated on supercomputers, he adds, "It's intriguing that you have one method that can be used to display a three-dimensional relationship among either real data from the real world—from a huge variety of measurement methods that span a range of 20 orders of magnitude—that can also just as easily be applied to the simulation of phenomena generated on supercomputers by purely numerical methods. So it's a unifying approach to looking at the world."

In fact, researchers are volume rendering right up the scale, from molecules and their electron clouds measured in nanometers, to the distribution of gasses throughout the galaxy, measured in light years. Other applications include nondestructive testing and failure analysis of manufactured goods.

Interest in volume rendering among broad groups of researchers is growing. The volume imaging group that met at SIGGRAPH '88 had quadrupled to about 100 from the year before. Upon opening the meeting, Chairman Nick England of Sun Microsystems declared, "I'm going to teach radiologists to love voxels," and then polled the group to find out who they were. The group was equally divided between those studying real world and simulated data, as it was for dynamic and static, monochrome and color, and continuous and discontinuous. Between 10 and 20 were comparing multiple volumes and using stereo to help see an image more clearly. The smallest data set being massaged was $10 \times 10 \times 30$ and the largest was $10{,}000 \times 10{,}000 \times 4{,}000$ voxels. *Voxel* comes from *volumetric element*.

THE MATERIAL MIXTURE MODEL

The geometry-based model that Drebin discarded calls for describing an object in terms of lines, polygons, patches, and other abstract geometrical concepts that are then converted into pixels. But few objects in nature lend themselves to such descriptors. Alvy Ray Smith, Pixar's Executive Vice President, points out that "Ninety percent of the data [in scientific visualization] does not come from geometric sources—geometrically described objects—in the first place." Instead, Drebin, Loren Carpenter, and Pat Hanrahan developed the "material mixture" model, whereby materials are represented as adjacent to each other and/or mixing with each other.

The quest for the algorithm that finally helped Fishman visualize the fractured pelvis actually began with another set of broken bones. In 1984, when Drebin came as a programmer to what is now Pixar, it was still a division of Lucasfilm Ltd., handling special effects and the Pixar processor was still in the design phase. Just as the company was beginning to market the Pixar Image Computer outside the film industry, Drebin broke his arm. With his arm in a sling, he was unable to program, so he joined the team that was identifying potential markets for the machine. They landed at Siemens, where the medical imaging group was dissatisfied with its 3-D imaging results. Siemens' researchers gave Pixar a tape that originated with Fishman. The images were unsatisfactory, Drebin felt, because the users were trying to fit the data into available systems. "The problem seemed to be that people were trying to extract a geometric model from their data so that they could use a traditional renderer," he says. But the geometric model "didn't do justice to the subtlety of the image data." For example, the resolution of 48 CT slices was too low to reveal details of porous bone. "The way it was sampled, the porous bone appeared fuzzy," Drebin explains, "It was not a very distinct, sharp region. So it's difficult to extract geometry and it's not always possible."

I'm going to teach radiologists to love voxels.—Nick England at SIGGRAPH '88

Until now, most techniques for visualizing volumes have relied on displaying surfaces by reducing volume arrays to the *boundaries* between materials. Instead of working with pixels, arrays of voxels are created. Some researchers have manually traced two-dimensional contours from individual slices and then used graphics systems to connect them and form triangle strips or surface patches, but problems arise if the distance between the sections is large relative to the size of the voxels. Other surface techniques output polygons at every voxel; each voxel might be treated as a cube whose faces are output as square polygons, or values at each vertex are used and estimates made of where a surface cuts through the cube.

Because these techniques extract surface information from 3-D arrays, they are an indirect way of visualizing volumes. They assume that a thin surface, suspended in air, accurately represents the original volume, but often data are taken from volumes containing fluids and tissues that interface and form local mixes. They absorb and emit light differently, information which is lost if the data are reduced to shell-like surfaces. "We *have* a three-dimensional data set and we'd like to see *into* that three-dimensional data set," says Drebin, "So let's treat it as a three-dimensional image." To preserve the continuity of the 3-D data and see the exterior, interior, and local variations in that interior, Drebin developed a display process consisting of three steps: classification, modeling, and rendering. Classification assigns percentages of materials contained in each voxel. The result is not an all-or-none decision, but a best estimate of how much of a material is present in a voxel. The information from this probabilistic classifier is stored fractionally, not in binary, because binary representation disrupts the continuity of the data.

In the modeling and rendering steps, the material mixture model is represented by assigned colors and opacity. Boundaries occur when there is a local change between densities, so in a sense, surfaces can be extracted. The change in densities, or gradient, is used to estimate the surface strength, or amount of surface present. Later, that information is used for shading. Perspective on an image is provided by lighting a model in such a way that light can be reflected, absorbed, and emitted.

Because volumetric renderings are done for color images, red, green, and blue data can be computed by three channels in parallel. When rotating a volume, for example, operations for each set are done simultaneously. A fourth parallel channel, the alpha channel, handles opacity, compositing, overlays, anti-aliasing, and matting. Originally developed for creating special effects, matting is the technique responsible for space ships zooming past celestial bodies, as seen in *Star Wars*. In this application, it enables a viewer to peel away layers by performing cut-aways, or to remove regions if the data is unreliable or uninteresting.

A DIAGNOSTIC TOOL?

In case of artifacts, Drebin says working with percentages of materials in a mixture allows noise to appear *as noise* in an image, whereas geometric modeling does not allow noise to be directly modeled. Since it is identifiable, he says it is a big step toward getting rid of most artifacts. Others are not as optimistic about false details. Henry Rusinik, who works on the Pegasus Project site at New York University Medical Center, says, "We don't claim to be able to do a diagnosis yet. False positives are very important if you want to claim whether or not a person is sick. I don't know of anybody who can claim that 3-D is better than 2-D for diagnosis. . . . Eventually it will be, but not in the next 10 years."

Looking back on the successful treatment of the man with a crushed pelvis, Drebin remembers a videotape made two months after surgery. "It's amazing to see this man, with all the pins they put in, walking on a platform without a limp and doing deep knee bends." He says orthopedic surgeons at Johns Hopkins now regularly ask for "a Pixar" to be done. Over 2,000 studies have been performed so far. Also interested in chemotherapy for cancer patients, particularly radiation dose planning, Drebin explored how particles would scatter in a child's brain. Planning is now done on CT scans

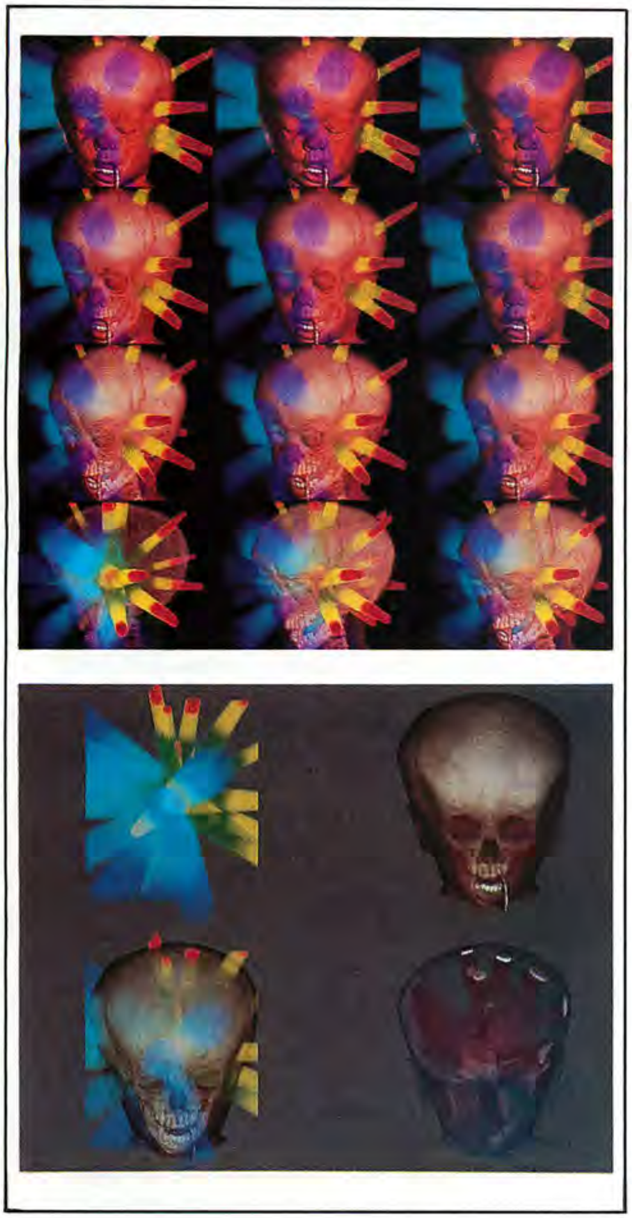

Volume rendered studies of radiation dosage can be used to pinpoint the location of tumors more accurately, allowing physicians to protect healthy tissue. Shown here are studies of radiation dosage in the head of a healthy child.

displayed on a two-dimensional plane. Matte algebra and choosing the right translucence and surface characteristics help elucidate the fuzzy radiation. That could improve targeting a tumor, for example, and protect healthy tissue.

NEURONS AND A NEW DATA SOURCE

Argiro started building graphics tools to help satisfy his curiosity about how the brain functions at the cellular level. He wanted to see how live neurons take on the special structural and functional properties that allow them to participate in the "extraordinary network that gives rise to our awareness, our capacity to learn and create and so forth." And then he wanted to "build beautiful, three-dimensional pictures of neurons blindingly fast." Recognizing how ambitious such an undertaking is, he says, "It's one of the most difficult problems in science. The brain is the most structurally complicated object in the universe and it's the most complicated volume that one could ever volume render."

A participant in Silicon Graphics Inc.'s Geometry Developers Group, Argiro's company, Vital Images, Inc., is developing software for SGI's machines, including the most advanced model, the Iris 4D-120 GTX®. His work is funded for three years for half a million dollars by the National Science Foundation and Iowa State Department of Economics, which wants to see its "Silicon Cornfields" grow.

Argiro had been developing a technique for surface rendering living cells' "foot prints" in tissue culture under an interference reflection microscope. Colleagues attempting to recover three-dimensional geometry were, and still are, manually tracing slices and then building pictures. For the most part, Argiro has eschewed the two primary data sources, CT and MRI, in favor of confocal data. In CT, there is a straightforward relationship between tissue density and the brightness of the image. Although most agree that MRI will replace CT, for volume rendering it presents difficulties. MRI data yields much higher spatial resolution and very subtle contrasts particularly in soft tissue, but the relationship between the signals and the underlying nature of the material is not simple, and that is an active research topic now.

Argiro saw the first commercial laser scan confocal microscope, the Phoibus 1000 made by the Swedish company Sarastro A.B. (U.S. patent number 4,631,581), at a 1987 neuroscience meeting. It was hooked up to a Compaq 386 and the picture quality was "very elementary." It took half an hour to build one picture on a five-MIP machine, but it increased resolution in the horizontal plane, and more important, greatly increased resolution in the vertical plane. "The depth of field of the microscope is very, very small," he explains, "so it's possible to take a half-a-millimeter-deep slice through an intact volume of living tissue without actually having to disturb the structure of the tissue." The laser is focused onto the specimen plane and scans across that plane. Instead of acquiring the image as a plane image in parallel, as in a normal optical microscope, the image is built up point by point, like a television or a raster. Light coming back from the sample is acquired by a photo multiplier, digitized, and then a number is recorded for each point in the image. This information builds up in the frame buffer. Then, to section finely in the vertical dimension, a stepper motor on the focus control takes a series of 2-D sections, optically slicing up a volume of tissue. The resolution is 0.25 microns in the x and y plane and 0.75 microns in the vertical plane. The Phoibus 1000 can scan 1024 × 1024 × 512 voxels. The maximum capability could have half a gigabyte of data generated from one experiment. "Unless you can render several hundreds of thousands

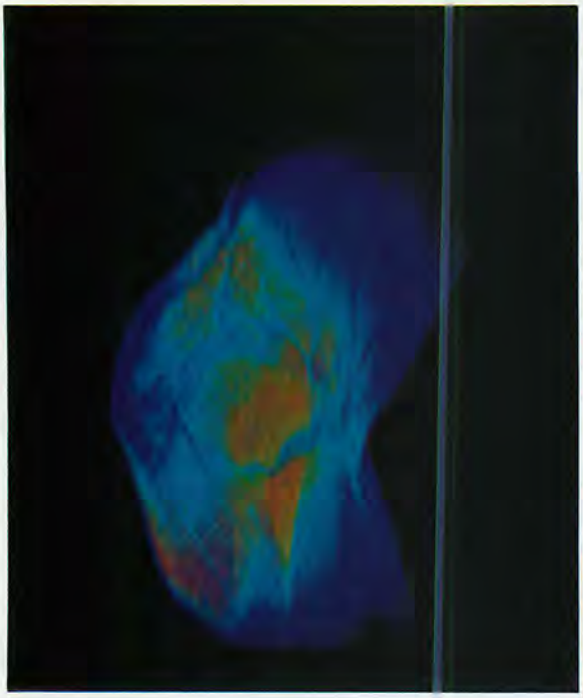

This pseudo-color rendering of a knee from MRI data shows the three-dimensional relationships between fine veins running across the knee (blue), musculature, including fibers, and bone (green). There are 256 × 256 pixels × 64 slices of 2 mm resolution MRI data, or 8,388,608 voxels. It took four to five seconds to render on a Silicon Graphics 4D70 GT®. Data courtesy of Picker International, Inc.

of voxels per second, you're not going to get near one or a few seconds of rendering time, and physicians can't wait more than five minutes to build a picture," says Argiro. On the SGI GT Series some pictures can be rendered in 0.5 seconds. So far the highest rate has been 240,000 voxels per second on MRI data of a human head. "No one else has gotten anywhere near that," says Argiro, "It has raised a lot of eyebrows."

FINE-TUNING MICROCODE

Besides making volume rendering an interactive rather than a batch process, the challenge has been to determine what tasks could be done using SGI's existing architecture, and to fine-tune existing algorithms. To make algorithms faster, Argiro focused on two areas. First, he looked at how to organize a volume database so that it could be traversed very rapidly, and useful information retrieved very quickly. Memory access time in a raw volume database of 10 megabytes, which is the size of the neuron set, means the loop in a program takes a substantial amount of time all by itself. Second, he developed a shading or contrasting method that be believes is unique.

Of the microcode, Argiro says, "We found that, inadvertently and fortunately, SGI had put most of the operations for voxel rendering into their microcode." For example, because volume rendering is compositing, you have to sum contributions from different planes that would be in effect in front of or behind the viewing screen. That means compressing them into a flat picture, an arithmetic operation done at the last minute before going into the frame buffer. Consequently, there must be a very fast adder/multiplier (in SGI's machine it's called a "blend function") right at the frame buffer. Another advantageous function is the complex lighting model, which allows you to set up a complex lighting equation that takes into account up to eight infinite or point light sources. In addition, material properties can be assigned to characterize the object in question and then executed on the fly at a very high rate. In principle that could be done up to 400,000 voxels per second.

The voxel rate is not limited by the graphics hardware, however. Argiro explains,

> It's actually limited by the CPU and how fast one can traverse a huge volume database and then stuff the information into the graphics hardware. If you're talking about rendering half a million voxels per second, you have two microseconds to do everything for that voxel. Now, two microseconds, even on a very fast, current generation RISC CPU is still only about 20 machine cycles, which means even on a very efficient processor like the MIPS that SGI uses, you've got 15 to 18 machine instructions to turn that voxel around. If you use C, you've got about four or five lines of code to handle that voxel. And then you've got to be doing the next one. That's how tight your program's inner loop has to be. And so that's where the performance tuning has gotten very involved. There just isn't a whole lot of time if you want to keep that rate that high. So we've sat many long hours staring at very small pieces of code, figuring out how to make them absolutely as fast as possible, coming up with data structures, coming up with approximation methods, and so forth, that give us a very heuristically pleasing image on screen, but still do it as blindingly fast as possible.

Argiro hopes to break the one million voxel barrier by adding microcode to the geometry pipeline. Theoretical calculations on dataflows on a number of operations that can be done per voxel, may hit rates as much as 5 million voxels per second on the existing architecture. But he admits that his algorithm does not work equally well on all databases. They work best on inherently high resolution, dense, data (256 × 256 × X slices and up), which would include most volume data. People have used it on very sparse or very low resolution data, (50 × 50 × 50), or just a few thousand voxels, but that requires interpolation to generate intermediate values, which the algorithm is not intended to do.

DISCOVERIES COMPARABLE TO THE DOUBLE HELIX?

During the past 15 to 20 years, neurobiologists have spent hundreds of thousands of dollars on histological and primitive computer graphics methods that would take weeks to recover the geometry of one neuron. The work has been as painstaking as it must have been for Watson and Crick to analyze X-ray crystallography data without the the aid of computer graphics. For Argiro, the confocal microscope has opened up hundreds of experiments and will help answer questions that are "just dying to be answered." He expects a huge flurry of cell biology papers because research that was previously impractical now is. Does he expect a discovery in neurobiology that is analogous to the double helix? "It hasn't happened yet, but it could. There are many, many cases in science where opening up a new experimental window opens up possibilities for people's imaginations and then—boom—something comes out the mist. The rate of understanding three-dimensional relationships in biological tissue can accelerate dramatically. That's been very exciting to me."

MACROMOLECULES: FROM DNA TO THE POLIO VIRUS

Working at the macromolecular level is Arthur Olson of the Scripps Clinic Molecular Biology Department in La Jolla, California. Trained as an X-ray crystallographer, he is now a molecular modeler, exploring how scientists interact with molecular models that they build based on data derived either from experimentation or computation. Interactivity goes hand in hand with interpretability, says Olson, so his team explores the best ways to represent a molecule's properties or characteristics. The goal is to make the process of examination and understanding as clear as possible.

Volume rendering macromolecules like proteins involves several hundred to several thousand atoms. Viruses, the next level of organization, have hundreds of proteins. The first step is to understand how molecules fold and take on their particular shapes. The next step is to determine how they interact with one another to produce larger assemblies that then go on and perform particular functions. The positions, charges, and other chemical characteristics of those atoms play a role in how these molecules function.

Marvelling at the complexity in nature, Olson says,

> The striking thing to think about is that, on the one hand, these are chemicals; they're just assemblies of atoms. But on the other hand, these aren't things that were just thrown together in a test tube; this is an assembly of large numbers of atoms that have evolved over billions of years, to perform particular functions. So that their design is incredibly sophisticated, incredibly intricate. In order to have a hope of understanding the relationship between the placement of all these atoms and the biological function at the other end, we have to be able to visualize the whole, and understand it in terms of all the parts, because certainly the whole is much, much greater than the sum of all of the parts.

Rather than choose between surface and volume rendering, Olson combines the two so that there are "landmarks," which are necessary for understanding and interpreting the many twists and folds of these very complex structures. Volumetric information is mapped

This neuron from a lamprey eel appeared round in previous renderings. Here, data rendered in a few seconds at 512 × 512 pixels × 160 slices using a Phoibus 1000® laserscan confocal microscope shows that the dendrite is a broad ribbon. The fine resolution shows a nissl substance near the nucleus. (If you imagine the dendrite is the shape of a fish, the nucleus would be the fish's eye.) Data courtesy of Peter Wallen, Karolinska Institute, Stockholm, Sweden.

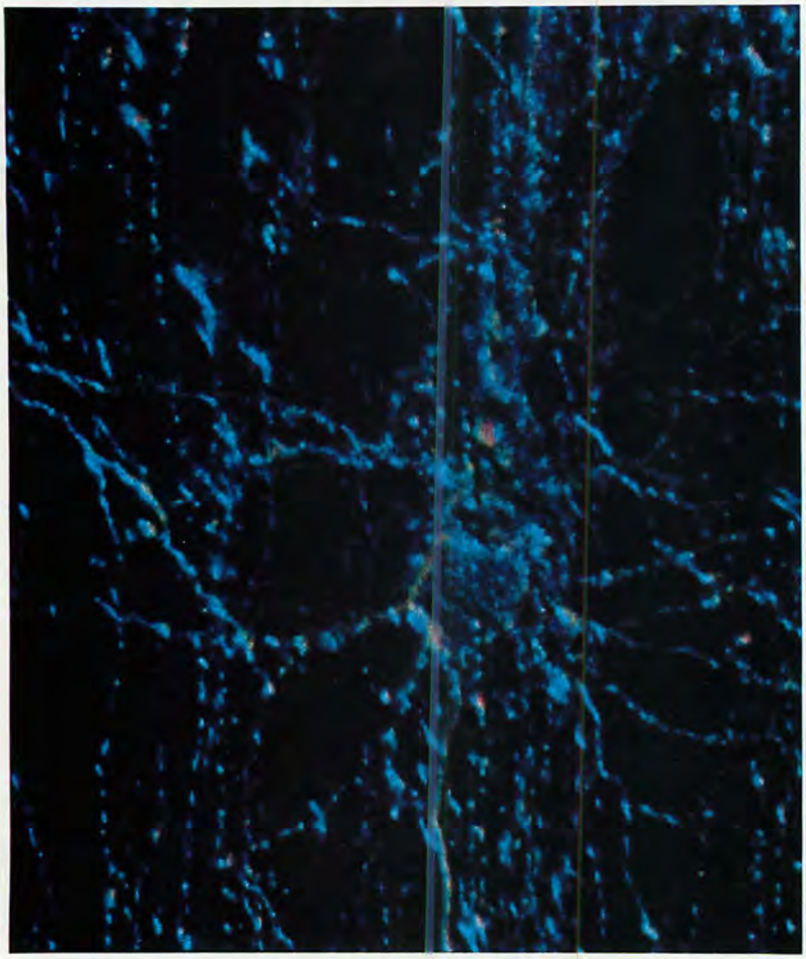

Neurons seen intertwining under a confocal microscope create this mesh. On the lower right is a paramital cell surrounded by fiber tracks of axons making contact with it. Above it are orange dots that show vericosity—thickening of dendrites or axons. The shininess is due to the lighting system, which helps pull out details (256 × 256 × 150 slices). Data courtesy of Stephen L. Senft, Washington University School of Medicine, St. Louis, Mo.

onto, or rendered in the context of, a geometrical model, like a molecular skeleton or surface. The model is based on some data, and then the model itself is the basis for the synthetic image of the molecule. Unlike those working strictly on data interpretation, Olson is looking at it from the model analysis end. The model and volumetric data both are rendered in the same frame—the program takes two different types of information and integrates it into a single view. Their algorithm is in the frame buffer, is written in FORTRAN, and runs on the Convex C1 computer. Generally, it calculates and renders images that are 500 pixels squared at 24 bytes per pixel for four color separations in 20 to 40 minutes. Olson would like to get simulations of optical interaction due to chemical properties in real time, but "that's asking a lot," he says. "We work at both ends of the problem." In this area, intelligibility of images and picture quality have been the priorities and they have not been sacrificed for faster computation.

One unanticipated discovery that was extremely useful for understanding the self-assembly of a virus particle resulted from volume rendering the polio virus. Inside the molecule one would expect to find an even, relatively neutral inner surface because the phosphate groups on the nucleic acids have a lot of negative charges. In turn, one would expect the protein code to have a lot of positive charges around it. But when Olson and post-doctoral fellow David Goodsell did the calculation, it showed a large lobe of negative density. This helped them see the complementarity between one side of an assembly unit and its mate. "So not only are the shapes complementary but you can easily see from these renderings that the electrostatic fields are complementary," Olson explains "Often you see things you weren't looking for. It's when you don't know what questions to ask that graphics takes on importance."

Are artifacts a problem at this scale? Since no life threatening situations can arise, the stakes are of course not as high as in medicine. Further, these macromolecular researchers are not looking at individual cases, but at the generic case. And the properties they are looking at are calculated based upon a grid for descretizing and sampling points in space. For example, to calculate the electrostatic potential at a point, the charges in space are summed *vis-à-vis* a test charge at a particular point. Then they move the point a discrete amount in space, and do the calculation again. Upon volume rendering, they cast a ray of light through that grid, and perform interpolations based upon those values computed at that grid. The researchers define that grid, but in medi-

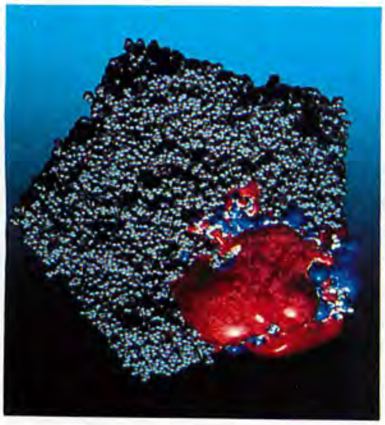

The red and blue bubbles show an isocontour of negative and positive electrostatic potentials of one of five protein building blocks that form the polio virus pentagon. The predominance of red areas indicates that the inner surface of the virus is mainly negative, which is surprising because the RNA found there is also negative. Data provided by J. Hogle, Scripps Clinic. Computer graphics courtesy of D. Goodsell and A. Olson, © 1989, Scripps Clinic.

This image data of ribonuclear protein from a mosquito was collected by an electron microscope. The yellow strand of DNA has been added to give a sense of scale. The ribonuclear protein particle transports ribonucleic acid out of a cell nucleus into the cell body. The horizontal lines are 13.1 nanometers apart. Data provided by D. Skoglund, Karolinska Institute, Stockholm, Sweden. Computer graphics courtesy of D. Goodsell and A. Olson, © 1989, Scripps Clinic.

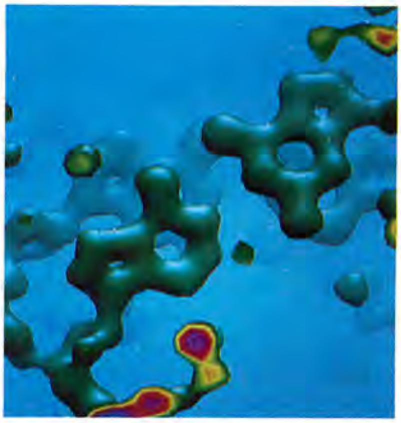

Experimental data derived from X-ray diffraction analysis of DNA shows a close-up of one part of the electron density map. Here, the electron density is so clearly defined that you can see the chemical rings of this base pair mismatch. Data provided by G. Prive and R. Dickerson, UCLA. Computer graphics courtesy of D. Goodsell and A. Olson, © 1989, Scripps Clinic.

cal imaging, the scanning equipment dictates the grid and therefore the sampling of measurements. "They don't have much of a choice," explains Olson, "If we feel that we're getting artifactual results, we can adjust the grid and make it smaller."

Whereas for Argiro, volume rendering unifies inquiry across scales, Olson sees it as a unified paradigm for different types of data within the macromolecular level. "We can use the same algorithm to show contact surface, electron density, electrostatic potential, and any calculated or observed properties in three dimensions," says Olson. "So it has the potential of unifying a lot of different types of analysis for us."

OBJECT-ORIENTED GRAPHICS

Recognizing the need for scientists to do both geometric and nongeometric modeling, and to help them reuse their code, several companies are creating object-oriented programming (OOP) graphics tools and libraries. Among these are Apollo Computer, Inc. in Chelmsford, Massachusetts, Ardent Computer Corporation, in Sunnyvale, California, and Stellar Computer, Inc. in Newton, Massachusetts. These companies also see OOPs as a way of distributing graphics and providing data abstraction to help process, and determine where best to process, the torrents of data volume rendering involves.

At Apollo, OOPs has grown out of a desire to improve the efficiency of distributed processing done at workstations linked by networks. In such "network computing," as Al Lopez, Apollo's Director of Graphics and User Environment R & D, calls it, data exists at many locations on the network and the problem of getting the processing out to the data must be addressed. "How do you process it?" he asks, "Do you move the data around the network or do you position it strategically next to the computing resource that's going to be processing the data? Then, how do you get the computing, or that part of the application that is going to operate on the data, out to that resource?" Because scientists want to control or steer components of the application, they no longer want a batch process. There is a demand for interactive, low-end systems on high bandwidths, but it is not enough to simply move and process data faster, says Lopez. Interactive, distributed applications must be supported in such a way that the network is transparent to users and that they are immune to changes in the network. The next generation of graphics environments will be object-oriented and coupled with object-oriented network computing systems, like Apollo's Network Computer System, he says. To such a graphics interface you would be able to say, "Rotate yourself 30 degrees" and not care about whether it's a bit-mapped image or some geometry. In addition, OOPs would provide extensibility so that users would not have to wait for a new release from a vendor. And it would also provide distributed access to shared objects anywhere on the network. Says Lopez, "If you say to an object, 'Ray trace yourself,' [you should] have the method for that object to use the distributed computing environment to spawn out copies of the ray tracer on multiple nodes in parallel."

One object-oriented graphics application interface is Ardent's Doré® (Dynamic Object-Oriented Environment), which was written to accommodate the Ardent mini-supercomputer's closely coupled computational and graphics data architecture. To satisfy that kind of abstraction, it is built at the macro level, so that primitives, attributes, views, and devices are objects. The user does not have to understand rendering, know where the processing is occurring, what decomposition of geometry is being performed at a lower level, or what attribute assignment mechanisms are in use. Instead of saying how to render an image, Doré directs different renderers, like Pixar's RenderMan®, to obtain the best output for the devices being used. Doré is also extensible, so you can, as Ardent did, take a photograph of Yosemite Valley, map it onto a three-dimensional height field, make a mesh of triangles, color them and then rotate them in real time. Then you could add those primitives to Doré and reuse them. Programmed in Standard C, Doré is portable to other systems.

LINKING CAD AND VOLUMETRIC DATABASES

In the nondestructive testing field, engineers are saving money by using volumetric data to analyze and inspect parts and to compare them to the idealized versions that they designed. Donald Jones, Director of Engineering Animation and Visualization at Failure Analysis Associates, Inc.® in Palo Alto, California, has used CT, MRI, and ultrasonic data to check dimensions, and eddy current data to measure conductivity. In the past, Jones was able to look at only a subset of data after preprocessing, like polygons that then had to be interpolated; now he can look at all the raw data. Occasionally, preprocessing causes inspectors to miss minor defects, but with volume rendering, "they don't have to throw [any data] away." Since industrial CT can take far more slices than can be taken from humans (who cannot be exposed to the radiation as long), the resolution is considerably finer: 5 mil (1 mil = 0.001 inch; or 1 pixel is about equal to 5 mil) compared to between 40 mil and 50 mil for humans. Jones, a Pixar user, performs both surface extraction and assigns transparency to voxels because sometimes they are interested in the form and shape inside a container. Putty inside a part will have its own form, for example. Since different types of materials adhere differently to surfaces, that has to be taken into account if the material and the part are processed together.

Jones emphasizes the interface between his inspection and CAD databases. "Now that we can reconstruct, three-dimensionally the *actual* part, not an idealized part, we can take the CAD database that the designer has used, and look at that data to cross check whether our part is dimensionally tolerant." Explains Jones, "We take an existing finite element model and look at our CT and then warp or scale the finite element model to match the CT model. Now we have an exact model and can run an analysis on that and see if it will fail, based on the mathematical geometric properties of that particular part." If it is far off, they reject it. But if it is just out of tolerance, they can store it and wait until an accompanying part appears with compatible tolerances. This can lead to tremendous savings. Says Jones, "These parts are not cheap. They're the cost of a couple of California houses—$300,000 each. You don't want to just throw them away if you don't have to." Eventually, Jones hopes to automate his inspection system with neural networks that will look at patterns of defects and decide what to do with defective parts.

MOTION AND PHOTOREALISM VERSUS FAST IMAGE PROCESSING

Volumetric images offer the most information when they are seen moving on the screen. Argiro says rotating an image on the screen is essential to the perception value of the computer graphics. "A lot of the brain deals with perception of motion," says Argiro, "and if you have a static picture on the screen there's a whole part of the brain that just goes to sleep. As soon as something moves on the screen a whole part of the brain wakes up." According to Argiro, interacting with a moving image is "SGI's angle." "Let's make [the visualization tool] interactive, and then we'll tune the quality and photorealism up." As the technology advances, he says, SGI will "make the pictures prettier and prettier and prettier, but without sacrificing interactivity. [The technology will get to] the point where one could make photorealistic pictures in real time, but one still wants the real time response."

It's computer graphics on the one hand, but it's image processing on the other. You're dealing with real data about real structures. These are not synthesized pictures that just come out of a mathematical equation.—Vincent Argiro

At Pixar, the philosophy is different. On the photorealism/fast image processing spectrum, Drebin himself leans toward photorealism, but says Pixar falls somewhere in the middle. "If you can approach realism, then you can back off again," he says, preferring to provide a superset of information first. "Then you have much more control over the abstract way of representing information. That's one way we differ significantly from other approaches."

Although Drebin agrees that the next quantum leap in volume rendering will come when very high quality images are produced in real time, much is still not understood: "We're getting a lot closer to understanding why images come out the way they do, but because you don't know the data you're working with, it sometimes is harder to control than when you're working with a geometric model where you chose where the type of light, your depth of field, the surface color, and the surface texture. With volume rendering you don't quite have that control yet. One of the big challenges is to get volume rendering to the point where you have all the controls you have in traditional rendering, so that you can predictably say, 'I want the bone to have this surface characteristic, with the proper reflections, the proper roughness.' One [end of the spectrum] is getting closer to traditional rendering, or getting closer to realism. And the other side is just making the environment in which it is possible to process that data much faster."

Olson believes that photorealism and merely interpretable images will converge. "You need data from both ends," he says. "Different types of people are oriented towards different things. I'm visually oriented. A lot of people say 'Why bother taking pictures when really the constructs are in your mind?' But, to me and to a lot of other people, things become much more real and concrete when you can have a view, a vision, of what they are. I'm not saying it's necessary for every scientist to look at things in this way, but I know it's necessary for a lot of scientists." And a "spin-off" benefit is that it helps convey scientific principles to others: "It's important in terms of communicating some of the science to other people."

CR Categories and Subject Descriptors: I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling—*geometric algorithms, languages, and systems*; J.3 [**Computer Applications**]: Life and Medical Sciences—*Medical Information Systems*

General Terms: Algorithms, Design, Performance

Additional Key Words and Phrases: Computed axial tomography, failure analysis, magnetic resonance imaging, neurobiology, nondestructive testing, radiology, tomography, volume rendering, voxels