The Third Branch of Science Debuts

Computer simulation has opened a new eye on the world, giving scientists in fields from biology to high-energy physics a way to perform experiments that would be otherwise impossible

Decked out in silly-looking 3-D glasses, the audience looked ready for a showing of that 1950s Hollywood classic, *House of Wax*, with its hideous wax figures leaping out of the screen. But instead of Vincent Price, the leading man was a sinister oxygen atom that had invaded a quiet silicon neighborhood and thrown the local electrons into a tizzy. Granted, it's not the sort of plot likely to win an Oscar at next year's Academy Awards, but the spectators found it gripping nonetheless, and for a good reason: The screenplay was based on a hitherto hard-to-uncover real-life event.

The theater was a room at last month's meeting of the American Physical Society in Indianapolis; the "flick"—actually a series of slides—was "The Enchanting Properties of Oxygen Atoms in Silicon," produced and directed by MIT physicist John Joannopoulos. And although the precise details of what happens to the oxygen atom and his silicon co-stars are of interest mainly to specialists in silicon processing, the very fact that such a documentary could be made is extraordinary. After all, none of science's powerful microscopes or other tools has the ability to peer into a solid piece of silicon, watch as an oxygen atom settles in, and record the reactions of the neighboring electrons. It's simply not possible.

So how has Joannopoulos done it? He's not a magician; he's simply one more researcher who has discovered the power of what is being called the third branch of science: computer simulation of reality. With simulations so accurate that they can take their place alongside experimental data as an object of scientific study, this "computer experimentation" has given researchers a new eye on the world. More and more, scientists are going beyond the data they get from direct observation and watching, for instance, as a drug molecule latches onto a protein or as electrons zip through the guts of a transistor. Spurring them to invest the time to learn to work with new tools and the money to pay for them—the newest supercomputers cost tens of millions of dollars—is the unique intellectual experience a researcher can get from seeing the unseeable. "The visualization can become an extension of the scientist's thinking power," says IBM simulation specialist C.N. Liu. "It allows him to see familiar objects in ways never before possible."

Which is why, unlike Hollywood's 3-D craze, this new way of doing science is just going to grow in popularity: Computers keep getting faster and cheaper, so the experiments will only get better. And plenty of good software is now coming onto the market, so that a scientist no longer has to write his own programs in order to mimic reality in silicon chips. The only difficulty seems to be how to make sense of the huge amount of data generated by these computer experiments, but that's the type of problem that most scientists don't mind attacking.

Joannopoulos, for one, is happy to wrestle with the huge chunks of data his simulations produce about how impurities form in silicon, since nothing in the repertoire of experimental physics could tell him nearly as much. In a typical experiment, he inserts a single oxygen atom into a block of 64 silicon atoms and feeds in the quantum mechanical equations that the system must obey. The computer then traces out the movement of the oxygen atom inside the silicon, offering insight into how the intruder begins the oxidation process that can degrade the silicon used in integrated circuit chips. Already, Joannopoulos has discovered that oxidation does not start in the way that an earlier theory proposed—a finding that could lead to improved processing of silicon.

Such computer discovery is by no means limited to the world of materials science, however. In the warm, living world of carbon molecules, computer experimentation has had even greater cinematic success. Armed with sophisticated molecular modeling programs and visualization software, biologists can now easily create proteins on a computer screen and experiment with them—bringing other molecules up close to them and discovering how they interact, or changing something in a protein's composition and watching how that affects its three-dimensional structure. This capability is particularly valuable for pharmaceutical companies, whose drug designers can test thousands of potential drugs on a computer to find the most promising candidates, and these companies pay top dollar to provide their researchers with the best computers and the most up-to-date molecular-modeling software.

This brave new world of scientific experimentation isn't problem free, of course. Just a few years ago, if a scientist wanted to get started in so-called rational drug design or in the related area of protein engineering, he faced an unpleasant choice: either write his own software, a tedious and difficult job, or else use programs written by another researcher, which would most likely be idiosyncratic and difficult to learn. Now there are two good, widely used commercial software packages for molecular modeling—Discover by Biosym Technologies in San Diego, and Charmm by Polygen in Cambridge, Massachusetts—that would appear to have an-

Computer visualization lets scientists see the unseeable: winds and temperatures in a global atmospheric model.

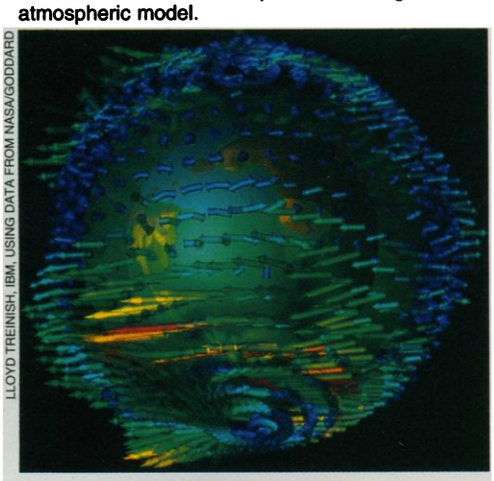

LLOYD TREINISH, IBM, USING DATA FROM NASA/GODDARD

swered the most difficult part of the challenge. Suddenly, anybody with $100,000 can set up a mini drug company. The software will cost you around $25,000, you'll need about $50,000 for a good computer workstation to perform the calculations, and the leftover $25,000 will buy a top-of-the-line graphics workstation to display the results in breathtaking, three-dimensional detail. But, oddly enough, the ready availability of commercial software may be a mixed blessing.

Klaus Schulten, a University of Illinois biologist who uses computers for rational drug design, acknowledges that it certainly brings more researchers into the field and helps the drug companies, which don't have to develop their own software. On the other hand, he warns, "we've reached the situation where some of the progress is dependent on the release of the next generation of software." Schulten, whose group still designs its own software (although it is careful to keep it compatible with Polygen's Charmm), isn't comfortable with scientific breakthroughs having to wait for one of two companies to add another few thousand lines of programming code to their software.

Aside from these immediate, nagging concerns, most digital experimenters see an even grander challenge on the horizon. Yes, the incredibly vivid images on a computer screen can make a computation come to life, but they are most effective for only a limited set of problems that can be easily visualized—such as the three-dimensional structure of a protein. Over the next few years, the experts say, one limiting factor in the expansion of computer experimentation is likely to be the ability to display the growing variety of computer-generated data in forms the researcher can use.

How to see a cloud. The history of Joannopoulos' work summarizes the problem nicely. When he first started his work on silicon, he didn't have a good way to visualize the results, although he knew what he wanted: "something that really gave me the sense that I was in the material—to imagine that I was an electron inside the material." The important features to understand were the clouds of electrons around the silicon and oxygen atoms, which determine how the atoms bond with each other. But how, he wondered, does one picture an electron cloud, which is nothing more than a set of probabilities that an electron is in a particular spot?

DAVID VOSS

To answer that question, Joannopoulos took his data—a set of numbers from 6 to 1800 specifying the electron density at different points inside a three-dimensional bit of silicon—to Robert Wolfe, a visualization specialist at IBM. Wolfe's assignment: to create from those numbers an image that would clearly display what was going on inside the silicon. It wasn't easy, but after peering at the data from every angle, Wolfe settled on a scheme that would emphasize key features and ignore others as unimportant.

The bonds between the silicon atoms had certain characteristic electron densities (between 110 and 140), Wolfe discovered, so he colored those orange; the most intense spots in these bonds reached densities as high as

A Visit to a Virtual World

Put on the goggles, grasp the control arm, and suddenly you're a captain trying to dock your ship into an irregular, three-dimensional pier. It all seems like some futuristic video game, but it's much more serious: Your "vessel" is actually a molecule of methotrexate, the pier is an active site on the protein dihydrofolate reductase, and if methotrexate docks well, it may prove useful in fighting cancer.

The docking exercise is one of many "virtual environments," or computer-generated worlds, created at the Department of Computer Science at the University of North Carolina (UNC) at Chapel Hill. By combining sight, sound, and even touch, the UNC researchers are finding new ways for humans to absorb the tremendous amounts of information being generated by computer modeling and simulations (see main story). In the future, these seers say, the two-dimensional screen could go the way of punch tape.

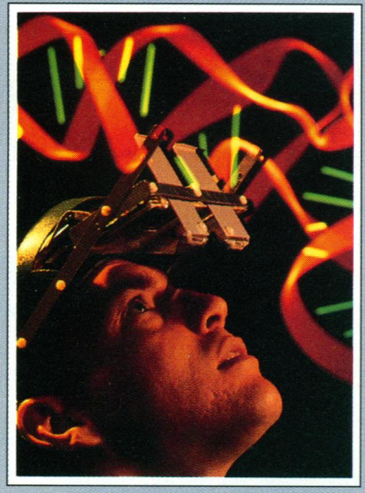

© 1990 PETER MENZEL

Head-mounted displays open a window into a virtual world.

Take drug design. Today, scientists use graphics workstations to visualize a protein in three dimensions and determine which drugs will fit best into its active sites. Such visualization has proved to be so powerful that no pharmaceutical researcher would try to do the job without it, but UNC computer scientists think they can do even better by literally giving a researcher the "feel" of a drug molecule sliding into a protein.

Led by virtual environments guru Frederick Brooks, the UNC group took a control arm, once used at Argonne National Laboratory to operate a remote arm for handling radioactive material, and hooked it up to a computer that simulates drug-protein interactions. Pushing, pulling, and turning the control arm moves the drug around the protein, while simultaneously the arm is pushing back on the user with a force that mirrors the electromagnetic force between the drug and the protein. At the same time, the researcher gets a three-dimensional view of the drug and protein in the goggles. Ask the biochemists who have tested the system and you hear that the combination of vision and direct physical feedback gives a better understanding of the forces between the drug and protein than sight alone.

The UNC researchers say that their crude prototype offers only a taste of the power that fully developed virtual environments will provide. "The ultimate goal," says William Wright of the UNC group, "is to give the user the illusion he is handling a real physical model that moves as it's supposed to." To reach that goal, the team wants to make the molecules more realistic—for instance, by having the now-rigid proteins deform in response to the presence of a drug. And the team also plans to give operators more sensitive physical control. A new whole-arm glove will respond to elbow and wrist movements and will permit thumb-second finger grasping—a big improvement over the simple hand grip of the current control arm.

And even that is considered just a beginning. The group is working on "walk-around" virtual environments, in which the user wears a three-dimensional head-mounted display that is sensitive to both the user's head position and place in a room. As the user moves around, turns his head, or kneels down to get a closer look at something, the view through the goggles changes. By combining that walk-around ability with molecular modeling, says Henry Fuchs, another member of the UNC team, the group hopes to create room-sized proteins. Then a researcher could inspect molecules from all angles—and get up close for greater detail. "Walking around gives a much stronger impression than moving the object around or flying [moving about at the push of a button instead of physically walking]," he says. The challenge will be to produce the force feedback on the arm while also giving the user complete freedom of movement. "We haven't a clue of how to do that," Fuchs admits. But that's the fun of it for these video game visionaries.

The Hunting of the Quark—Computer-Style

No one has ever seen a quark, and, if current wisdom is correct, no one ever will. The closest anyone will get is simulated quarks stalking around inside the supercomputers of researchers like IBM's Don Weingarten, who studies how these semi-mythical beasts combine to form such tangible particles as protons and neutrons. But if the computer safaris of Weingarten and other physicists prove successful, the trophy they bring home will be every bit as real as if they had snared it with a particle accelerator or some other traditional tool of experimental physics.

Weingarten studies quantum chromodynamics (QCD)—the physical theory that explains the strong force and all the particles affected by it: protons, neutrons, pions, and so forth. "Most particle physicists now consider QCD correct," he notes, "but a clean, quantitative test of QCD has still not been made." The reason is simple: The computations demanded by QCD are so difficult to perform that no one has ever managed to calculate a prediction that could be tested decisively. Now, with the help of massively parallel computers (see story, p. 50), Weingarten and others are closing in on a test of QCD and the existence of quarks.

According to QCD, all hadrons (particles that feel the strong force) consist of quarks and their antiparticles, called antiquarks, bound up by the "chromoelectric field." Protons, neutrons, and other baryons contain three quarks, while pions, kaons, and other mesons have a quark and an antiquark; the quarks are permanently confined inside the particles and can never be seen as separate entities.

Weingarten wants to apply QCD theory to make a very simple prediction: the relative masses of the proton, neutron, and other less well-known baryons, such as the delta. Such a prediction would be analogous, he notes, to the first real success of quantum mechanics: the explanation of the spectrum of the hydrogen atom. Explaining the hydrogen atom's spectrum was a matter of calculating the energy levels of the proton and electron making up the atom; similarly, predicting the relative masses of baryons will depend on computing energy levels of different three-quark combinations. If he can show that QCD theory does indeed explain the masses of the proton, neutron, etc., he will have bagged his quark.

To perform the QCD calculations, Weingarten uses a standard technique in which he breaks space-time into a finite set of points, or "lattice," and then does the computations on the lattice instead of in continuous space-time. Even with this simplification, a typical computer run involves calculations on a 24×24×24×24 lattice, and since Weingarten must keep track of 32 different components of the chromoelectric field at each lattice point, the computer is actually juggling some 10 million field components. Although the computations involving these field components are relatively simple, there are so many of them that a typical QCD calculation takes many months, even on the fastest computers.

IBM

IBM's Weingarten displays part of his quark-hunting equipment: one of 480 parallel arithmetic processors.

Weingarten says that physicists are getting close to being able to predict proton and neutron masses with QCD. He and others have been working with simplified QCD approximations to figure out first how large the lattice should be to give a reasonable approximation to continuous space-time; they have determined that lattices about 2×10^{-13} centimeters across—or about twice as wide as a proton—are good enough for their purposes. This leaves the question of how fine the lattice spacing has to be. "I'm hoping that we will have something sensible to say about that this year," Weingarten says. He and others are now doing calculations on lattices whose points are spaced about 10^{-14} centimeters apart, and if it turns out that this spacing is sufficient, "then we go to the full theory." In that case, they could have some predictions with another year's worth of calculations.

If the spacing needs to be cut in half, however, they would have to move to a 48×48×48×48 lattice, which Weingarten says would increase the computational demands by a factor of 100. Fortunately, it's likely that there will be massively parallel machines 100 times faster than today's best within a few years (see story, p. 50), so Weingarten or some other quark hunter shouldn't have to wait too long to hang that trophy on the wall.

180, so anything above 174 he made green to highlight those areas. The single oxygen atom, with its very high electron density, got a color to itself: blue. Everything else he kept transparent so as not to clutter up the image. The finishing touch was making videotapes of the oxygen atom moving through the silicon structure and adding a three-dimensional feel with special shuttered glasses synchronized with a video display so that the right eye saw a slightly different view than the left. Joannopoulos was pleased: "Your eye immediately detects what's going on." (The 3-D images shown at the American Physical Society were slides instead of a videotape, because Wolfe and Joannopoulos haven't yet made movies that work with the cheap polarized glasses instead of the expensive shuttered ones, but with the aid of synchronized movie projectors—similar to those used for Hollywood's 3-D movies in the 1950s—it will be easy to take that last step.)

As difficult as Joannopoulos' data visualization problems seemed at first to IBM's Wolfe, they are pretty tame compared to those raised by climate data. Imagine the challenge of picturing temperature, humidity, air pressure, and wind, all varying over time in the same movie—and doing it in three dimensions. The human mind doesn't easily take in this much information, so researchers like IBM's Lloyd Treinish are trying to find ways to make it more digestible. He experiments with different imaging techniques to make sense out of climate data, particularly ozone measurements. Much of his work is aimed at letting people visually compare multiple sets of data—observed versus calculated ozone levels, for instance, or the variation of ozone levels over time. If the data set is a simple two-dimensional one, such as ozone measurements over various points on the earth's surface, he can "stack" the data sets to make a three-dimensional image, in which the third dimension shows the changes between the two-dimensional sets. But it gets harder when more information has to be crammed in.

Suppose, Treinish says, you want to understand how the polar vortex affects ozone levels—a challenge that calls for watching atmospheric temperatures, wind patterns, and ozone all at once. His solution would be to use shape, color, size, translucency, and every other trick he can think of. To show wind patterns, for instance, flow lines are an obvious choice; for temperature, it might be different colorings or a changing opacity. You have to "figure out ways to mix and match different visualization techniques in one application," he says. "We are pushing the limits of available graphics techniques."

Even more information can be added to a simulation by turning to the other senses—hearing and touch, mostly—to avoid overloading the eyes. Researchers at places like IBM and the University of North Carolina

Materials Scientists Go to Work in the Family Business

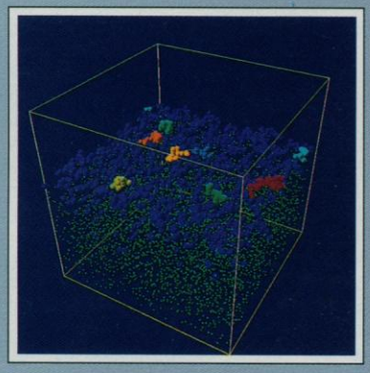

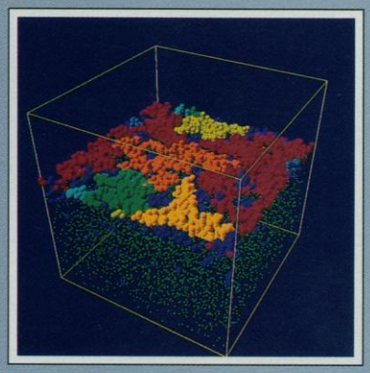

DAVID SILVERMAN/DAN PLATT, IBM

Copper atoms (colored areas) being deposited on a polyimide surface. Each color depicts an electrically conducting region.

Probably no area of science has been hit any harder by the computer experimentation craze (see main story) than materials science. Researchers can assemble tens of thousands of atoms in any configuration and then see how they behave, and out of these computer simulations have come tougher plastics, better lubricants, improved insights into how cracks form in metals and ceramics, and much more. As if this weren't enough, in an ironic twist this child of the computer revolution known as computational materials science has secured an important job in the family business: armed with computers, materials scientists are coming up with many of the keys to tomorrow's faster, more powerful computers.

At Georgia Tech, for instance, Uzi Landman's studies of thin films of lubricants should help computer designers better protect the magnetic disks that hold a computer's tens of millions of bytes of memory. In operation, such disks spin so fast that the head used to read and write data on each "flies" on a cushion of air some 200 nanometers above the disk. But when the disk is still, the head comes to rest on it, and if the disk isn't guarded by an effective lubricant, it can be damaged when the head is starting up or stopping.

Enter Landman, who studies such frictional forces with computer models of tens of thousands of atoms. In recent work, he simulated a surface coated with a film of hexadecane ($C_{16}H_{34}$, a lubricating hydrocarbon) up to 50 angstroms thick—about the thickness of the protective layer on a magnetic disk—and then brought a tip down toward the surface to see how the lubricant reacted. One finding in particular may be of interest to hard disk manufacturers: The 200-miles-per-hour speeds of a spinning disk push the long hexadecane molecules to align themselves parallel to each other, which, in turn, lowers the film's viscosity and makes it easier for the head to damage the disk. So Landman is trying different molecules to find which work best on a spinning disk. Already he's learned that changing hexadecane's long, straight structure by adding side chains, or small arms of carbon atoms, dramatically alters the lubricating properties. "We are really laying down the principles that have to go into the book for the design of these molecules," he says.

At IBM, David Silverman is one of many materials scientists who are using the computer to help the computer. Silverman "experiments" with polyimides (polymers closely related to the Kevlar of bulletproof-vest fame), which are used as insulators on integrated circuit chips. They are excellent insulators, especially suited for the advanced chips, with transistors crammed as close as they can get, but the polyimides have a weakness: They don't match up well with the copper that is used for carrying the electrical currents around the chip, and when copper is laid down on a polyimide layer, it can create a weak spot on the chip. So Silverman is studying how copper atoms settle onto the polyimide layer in hopes of figuring out some way to fix this shortcoming.

Because Silverman's work is all done on a computer, it's easy to vary the temperature and deposition rate to look for the best method of bringing the two contentious materials together, and it's possible to view the results via animated films created in collaboration with IBM colleague Dan Platt. There are no answers yet, but Silverman seems confident he'll uncover some way to improve this facet of the computer chip, "perhaps by getting the copper to channel into the polyimides and anchor the copper layer."

Meanwhile, Massimo Fischetti, another IBM researcher, is building smaller, faster transistors—on a computer, of course. In shrinking the transistors and other parts of computer chips past the 0.5-micron size of today's smallest components, IBM and other manufacturers are sailing into uncharted waters: In larger transistors, the electrons that carry signals can be thought of simply as charged particles moving according to the laws of classical mechanics, but as the transistors shrink, quantum effects become important and the electron's wavelike nature must be taken into account. "The devices begin to function in ways that are not understandable in the simple way of larger devices," Fischetti says. So, to save time for the chip designers working on the next generation, he builds transistors and other components on his supercomputer and checks their performance.

That effort may already have saved chip designers a lot of frustration over gallium arsenide transistors. Although gallium arsenide is more expensive and harder to work with than silicon, transistors made with gallium arsenide are faster, so many state-of-the-art computer chips are made with that balky material. But the gallium arsenide edge doesn't hold as the devices get smaller, Fischetti discovered. As the dimensions shrink, the electric fields get larger (supposing that the applied voltage remains the same), and at these higher fields, gallium arsenide loses most of its speed advantage. "When you reach the 0.1-micron point," he says, speaking of a size that has been achieved in laboratory silicon devices, "the gallium arsenide advantage disappears." Without Fischetti's discovery, IBM or some other company might have spent millions of dollars trying to shrink gallium arsenide circuits to 0.1 micron, only to find that they didn't live up to their billing. The moral: A little computer simulation can make a lot of difference, which augers well for a "family business" blessed with progeny like computational materials science.

are experimenting with "virtual reality," in which an observer can "explore" the images by turning his head and seeing different things or physically moving around them, and even interact with the images by touching them and manipulating them like real objects (see box, p. 45). But much greater levels of computing power must be invoked than are available today to make yet more complex processes seem real. Needed is more than dramatic increases in image resolution; the simulations will have to be able to respond to an observer's input within a few milliseconds to avoid noticeable time lapses and give researchers the feel of working with a real physical system. Of course, computer aficionados say all this is coming. It's only a matter of time, they say, before observing a computer-generated experiment will seem even more real than watching *The House of Wax* in 3-D sensurround.

—R.P.