MACHINE DREAMS

Seeing is believing. These days the images on a computer screen are often the only path to truth.

BY BRIAN HAYES

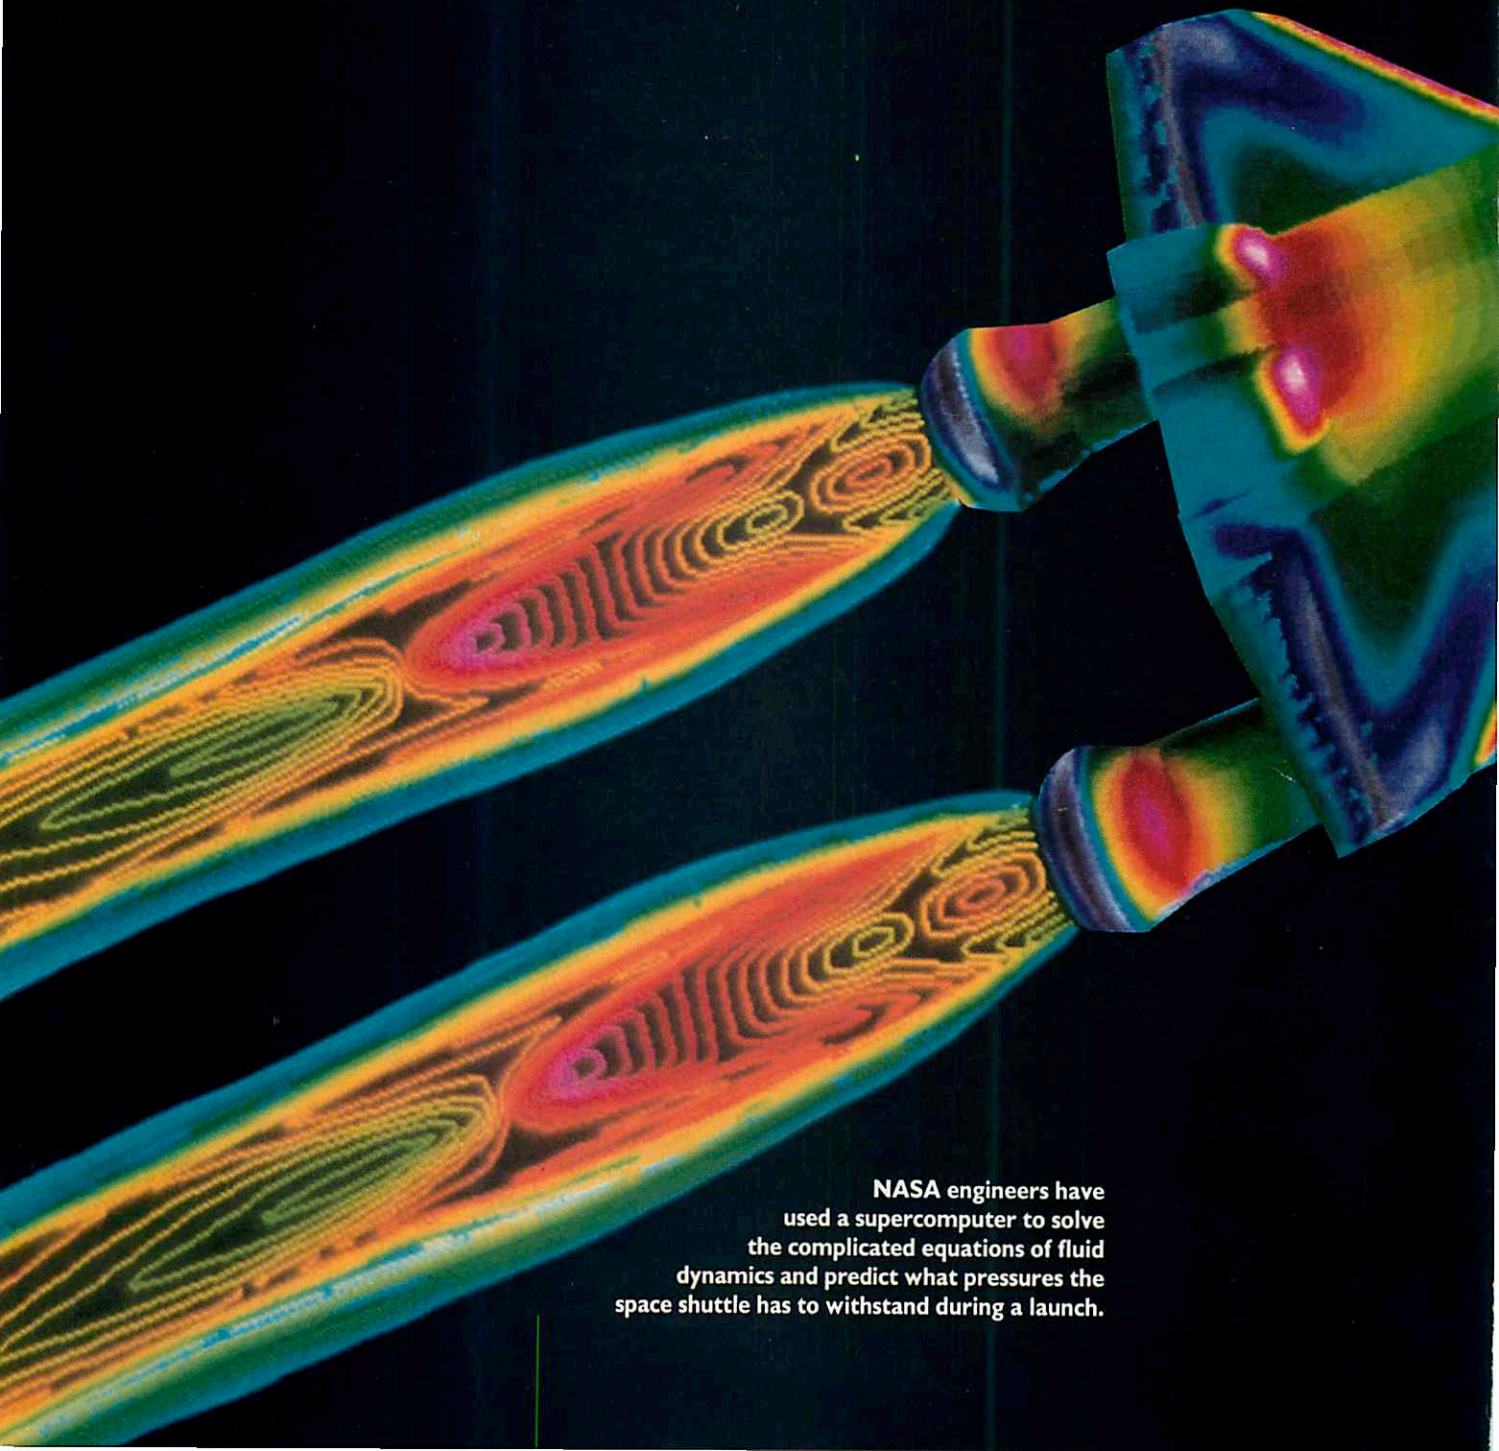

NASA engineers have used a supercomputer to solve the complicated equations of fluid dynamics and predict what pressures the space shuttle has to withstand during a launch.

PHOTOGRAPH COURTESY NASA

Science is generally conceived of as a dialogue between theory and experiment, but in recent years a third party has joined the conversation. The intruder is computer simulation, and it proposes nothing less than a new method of discovering truth—a way of understanding the world by reinventing it. In the 1980s this third way of doing science has penetrated most areas of research.

Often enough it is the only way that works. As researchers advance into ever murkier corners of the natural world, they are finding details and complexities that were not imagined a few decades ago, and they are constantly confronting the limitations of their old methods and of their own unassisted brains. Let me give an example of what I mean.

Astronomers have known for some time now what a supernova is. In the 1930s Subrahmanyan Chandrasekhar showed that when a large star exhausts its fuel, it collapses under its own weight. In some cases, it was later realized, the collapse triggers an explosion—a supernova—that blows off much of the star's mass, leaving behind a dense neutron star or perhaps even a black hole. All this was revealed by the methods of paper-and-pencil physics. But those methods bog down in tracing the events of the crucial milliseconds in which the core of a star contracts, "bounces," and then explodes. There is so much going on in that brief span that theorists cannot cope; and laboratory experiments, of course, are unthinkable.

In the past decade or so computer simulations have begun to fill in the missing details of star death. Such simulations begin with a mathematical model of a star

The results are worth waiting for. Computer simulations have revealed that it is not always easy to blow up a star.

that describes its mass and chemical composition and internal structure. Starting up the simulation has the effect of turning on the laws of physics, so that the model star begins to cook. Nuclear reactions are ignited; convection currents begin to flow; gravitation compresses the stellar core; neutrinos and other particles stream through the layers of the atmosphere. Meanwhile the physicist waits to see what will happen. The wait is not negligible: even with a fast computer, the simulation may take several minutes for each millisecond of real time.

But the results are worth waiting for. Computer simulations have revealed, for example, that it is not always easy to blow up a star: the shock wave that carries momentum from the collapsed core to the outer layers tends to stall, dissipating its energy in breaking up atomic nuclei, until it is revived by neutrinos from the core. Conversely, observations of real supernovas—particularly the one that erupted in the Large Magellanic Cloud in 1987—have inspired computer modelers. The dialogue of science has become a noisier, three-way klatch, from which all parties emerge the wiser.

Actually the use of computers to work out the detailed consequences of physical laws is not entirely new. In the eighteenth century, orreries with elaborate clockwork gears—mechanical computers, in effect—simulated the motions of the planets and their satellites; if you wanted to know the configuration of the solar system some years hence, you could just turn the crank. Some medieval astrolabes served a similar purpose. Indeed, the technology of such simulations is probably ancient: a hunk of corroded metal found in a shipwreck near the Greek island of Antikythera is thought to be the remains of a planetary computer built in the first century B.C.

But those devices were mere labor savers; they didn't solve any problems that couldn't be solved by other means. Only when mechanical computers were replaced by more powerful electronic ones, in the middle of this century, did simulation emerge as an alternative path to genuinely new knowledge.

One of the earliest and most famous instances of science by simulation happened just after the Second World War at the Los Alamos laboratories. Researchers there were still hard at work on better atomic bombs, and they needed to predict the fate of neutrons traveling through various materials. When a neutron struck an atom, it could be scattered or absorbed, or in some cases the atomic nucleus could undergo fission, thereby liberating more neutrons. For any given collision the probabilities of the various outcomes were known, but the overall problem remained a challenging one because each neutron took part in a vast number of collisions.

Stanislaw Ulam, a mathematician by training, devised a simple solution—or at least it seems simple in retrospect. The outcome of each collision was determined by choosing a number at random, according to the known probabilities. The simulation could in principle be performed with the aid of a roulette wheel. For example, it might be decided that whenever a spin of the wheel produces an even number, the neutron is scattered, whereas an odd number corresponds to absorption; the rare fission events might take place when a zero turns up. (The actual probabilities are rather different from these.) After thousands of spins, the average properties of the neutrons' trajectories would begin to emerge.

Because of the element of chance in this procedure, it was dubbed the Monte Carlo method. Over the years it has been applied to many problems other than tracing the paths of neutrons. In practice, a roulette wheel is not the instrument of choice for Monte Carlo studies. For these problems and others Ulam and his colleagues employed some of the earliest electronic digital computers, namely ENIAC and MANIAC. Today, the simulations run on the fastest machines available; and it is the wider availability of such machines that has enabled computer simulation to blossom in the 1980s.

It is hardly surprising that simulated science has been adopted most eagerly in those fields where other methods run into the most forbidding obstacles. Supernovas are simulated because the real thing is rare, remote, and inaccessible. Similarly, there are few alternatives to simulations in trying to understand the large-scale structure of the universe. Sim-

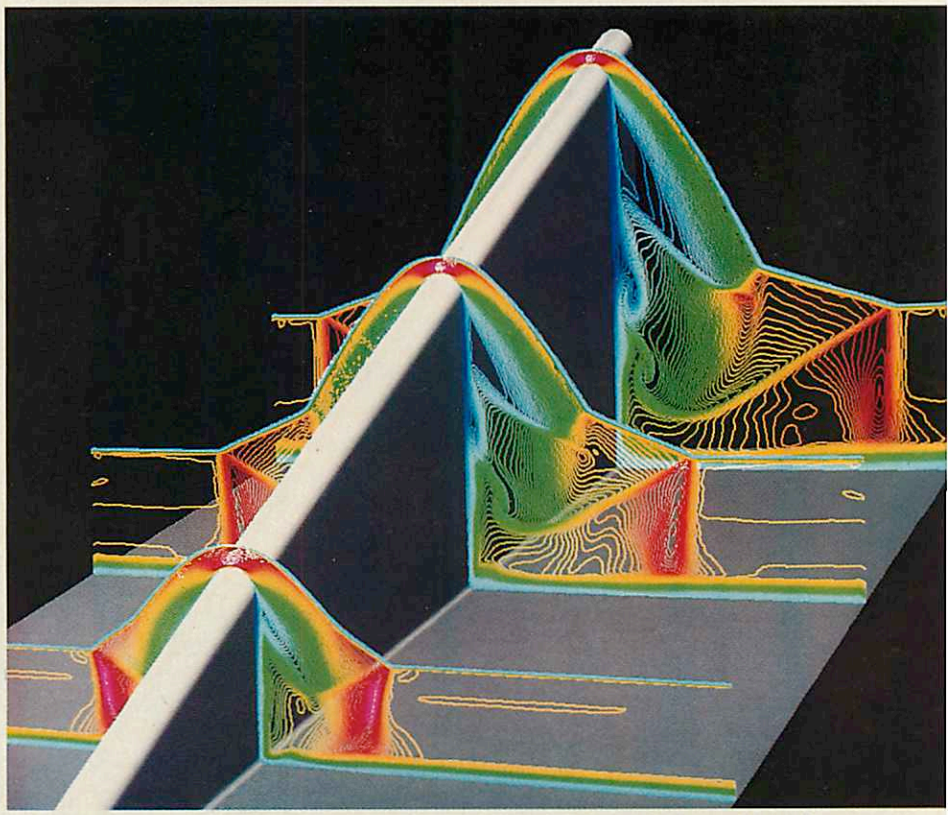

PHOTOGRAPH COURTESY NASA

A simulation reveals the pressures on the wing or tail of a jet flying at Mach 8.

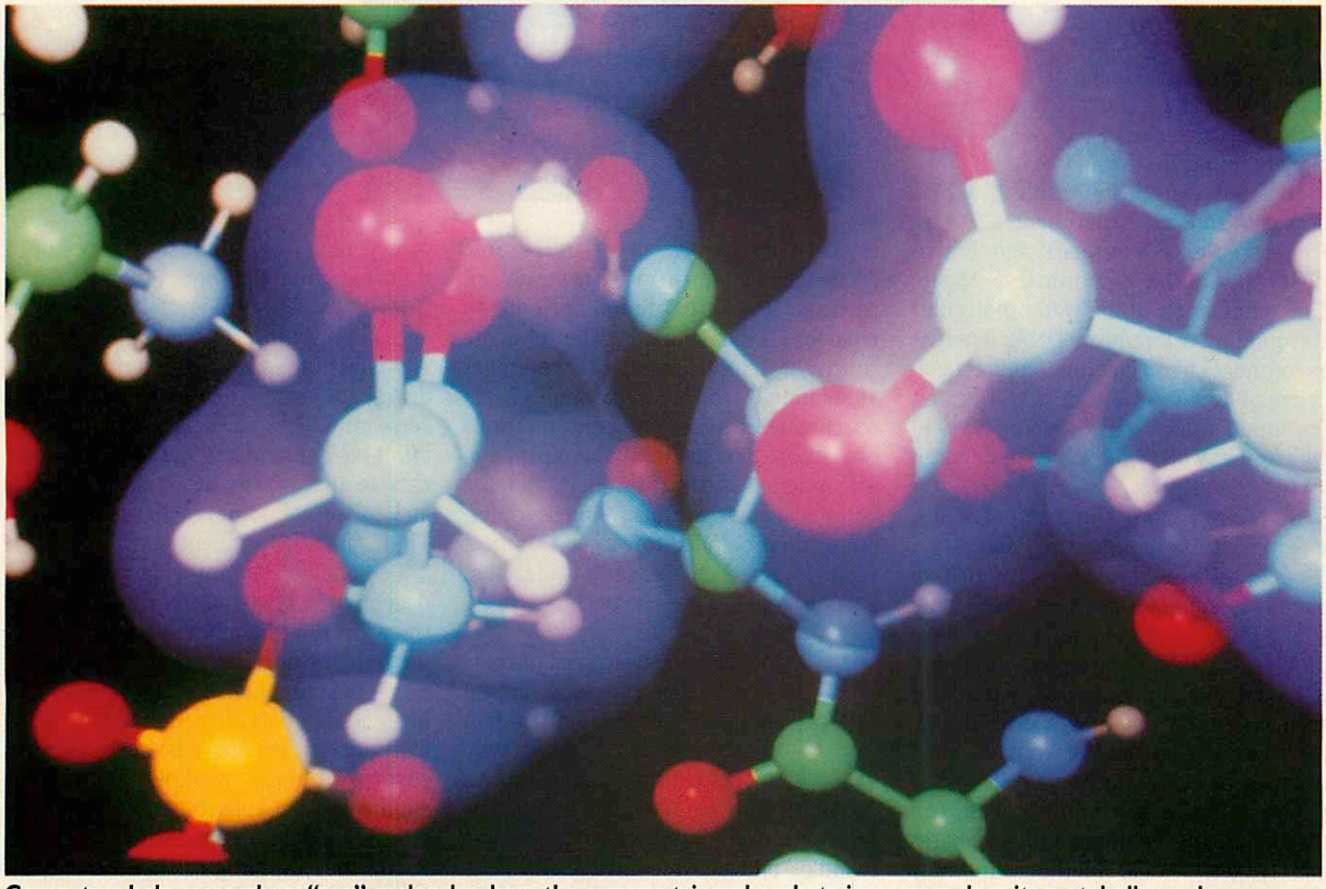

Computers help researchers "see" molecules; here the enzyme triosephosphate isomerase does its metabolic work.

ulations are the only way to get an outsider's view of the universe. They show vast wispy filaments and cobwebs that resemble the tentative maps produced by observational astronomers, and they give at least preliminary hints about how these structures might have evolved.

Simulation can fill in not only for impractical experiments but also for inadequate theories. In the 1970s one of the hardest problems in theoretical physics was explaining the "confinement" of quarks (the component particles of protons, neutrons, and the like). Inside a proton, quarks seemed to move around without constraint, but any attempt to pull them apart in particle accelerators revealed unbreakable bonds between them. The bonds were those of the strong nuclear force, which is described by a theory called quantum chromodynamics; but theorists were at a loss to find in quantum chromodynamics an explanation of why quarks should be so free and yet so inseparable. Essentially the problem was that the force field between two quarks was too complicated to calculate.

A definitive solution to this problem has still not been found, but the best approximation so far is a type of computer simulation devised by Kenneth Wilson of Cornell. In this technique the force-field calculation is simplified by replacing continuous space and time with a fictitious four-dimensional lattice, a kind of space-time jungle gym. The quarks can only occupy intersection points in the lattice, and the links connecting these points represent the forces acting on the quarks. Each link is variable, and so the number of possible configurations of the overall force field is enormous; it's as if each bar in the jungle gym could be painted one of many different colors. But in the approach followed by workers who have elaborated on Wilson's original idea, the computer considers only a representative sample, which it selects at random by means of the Monte Carlo method. It then computes an average configuration.

What such simulations have done is to vindicate quantum chromodynamics, by showing that it does indeed predict that quarks should remain confined by the strong force in protons and neutrons. That in itself is a scientific triumph. Beyond that, however, the simulations have also made a fascinating prediction. At temperatures of several trillion degrees, they suggest, quarks might be liberated from the strong ties that bind them and swim freely in a kind of "quark soup." Many investigators now think that quark soup is what the universe was made of in the immediate aftermath of the Big Bang. It may still be around today inside the neutron stars left behind by supernovas.

These examples may give the impression that computer simulation is used only in extreme situations, when traditional theory and experiment fail. But in fact simulation has also found a place in sciences that have no such methodological barriers. Chemistry, for example, is the prototypical experimental science—the chemist's test tube serves as an emblem for all of science—but during the past decade computational chemistry has become an important subdiscipline. Some computational chemists labor to calculate the structure of a molecule from the fundamental equations of quantum mechanics. Others employ less rigorous computational methods to describe large molecules, such as polymers, or to investigate complex systems of linked

PHOTOGRAPH BY DONNA RAIMONDI, ALLIANT COMPUTER SYSTEMS

reactions, as in the combustion of fuels. A big advantage of the computational methods is that they allow the chemist to see directly things that would have to be inferred from the results of a conventional experiment. For example, a few years ago Enrico Clementi of IBM performed a Monte Carlo simulation of the water molecules surrounding a strand of DNA; whereas experiments had given only indirect evidence of how DNA interacts with water, the simulation revealed specific ways in which the water molecules tended to arrange themselves.

In biology, too, the computer has elbowed its way into the laboratory. Among the first biologists to turn to computer simulations were those who study population dynamics. A typical problem in that discipline is to figure out how the population of a predator species and that of a prey species fluctuate when the two are in contact. Field observations can take years, but in just a few minutes a computer can run through many generations of, say, hares and lynxes. Similar methods are now being applied to the study of epidemic disease, including the spread of AIDS through the human population.

Climate forecasting is another area in which computer simulations are touching on issues of immediate social concern. In the 1980s all of us became aware that the human race has the power, through the greenhouse effect, to change Earth's climate. The ultimate source of that realization was computer simulations, which forecast the effects on climate of increasing concentrations of carbon dioxide in the atmosphere. The science is still primitive, and the various models are distressingly different in their detailed forecasts. But most seem to agree on the central idea that Earth is going to get warmer.

I do not believe that deceptive models provide much of a threat to the pursuit of knowledge.

Not everyone welcomes the new role of the computer in science. One cause of discontent is cultural: a biologist who works with furry animals or a geologist who hammers rocks may be reluctant to acknowledge someone who twiddles bits as a member of the same fraternity. But there are also more substantial questions about the prudence of trusting answers that come from a machine.

The hazard that gets the most attention is what might be called the Ptolemaic fallacy. A computer model could work smoothly, reproduce experimental results in great detail, even give accurate predictions of future observations, and yet still be totally wrong. A case in point is an orrery built to represent a geocentric universe, like the one Ptolemy envisioned. By fine-tuning the gear ratios and adding epicycles, we could continually refine the machine, covering up any discrepancies between the simulated universe and the real one. With enough effort we could match the model's accuracy to that of any available telescope, and so we would never detect a failure. That is just the problem.

Personally, I do not believe that such deceptive models present much of a threat to the pursuit of knowledge and the integrity of science. After all, conventional experiments can also be misleading, and theorists are certainly fallible; if institutional science has been able to cope with these weaknesses, then it should also be able to handle the occasional simulation gone astray. Indeed, a third mode of doing science—even if it is very imperfect—ought to improve overall reliability by providing an additional check on the two existing modes.

Another hazard of simulation may be more insidious. Suppose we set out to build a computer model of the biochemistry of a living cell. We might start with a few fundamental reactions, say the Krebs cycle for extracting energy from nutrients. Then we could add some finer details, such as the "pumping" of ions across cell membranes. Next would come the enzymes that regulate the metabolic process, the genes that regulate the enzymes, and finally the feedback loops by which the metabolic products regulate the genes that regulate the enzymes. Eventually our model might become so accurate that it could mimic the behavior of the real cell in full detail. Such an achievement would have to be counted a success, and yet there is something troubling about it. At some point the model inevitably becomes so complex that we cannot understand it any better than we can understand the real cell. Then we need a new science—a metabiology—to help us interpret our own models.

As far as I know, simulations in the physical and biological sciences have

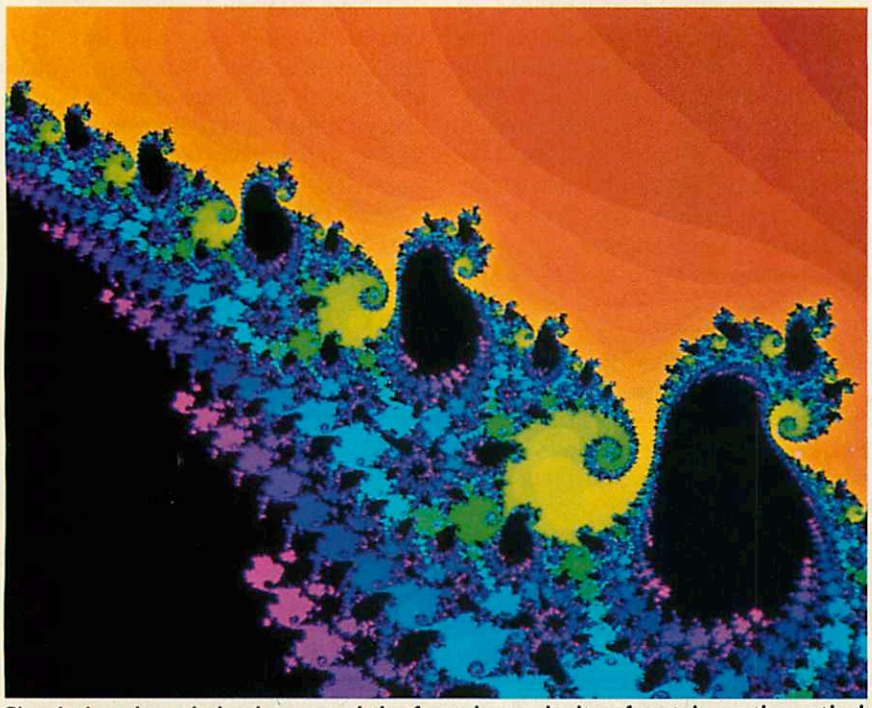

Simulations have helped to reveal the fractal complexity of certain mathematical (and natural) structures. The structure shown here is called the Mandelbrot set.

PHOTOGRAPH BY JANE STALLER/ART MATRIX

not yet reached this impasse, but mathematicians have had to confront a similar problem. Certain mathematical theorems have been proved only with the aid of a computer, the most notable example being the four-color-map conjecture, which was proved in 1976 by Wolfgang Haken and Kenneth Appel of the University of Illinois (using three computers for 1,200 hours). Like other computer-aided mathematical proofs, this one is so long and convoluted that it can be checked only with the assistance of another computer program. The notion that something can be considered proved and yet remain beyond the reach of the human intellect leaves some mathematicians deeply unsatisfied.

There is still another objection sometimes raised against computational science: that we cannot learn anything fundamentally new from a simulation. The argument runs as follows: all a computer can do is reshuffle its inputs in various ways and eventually return some permutation of them; thus whatever answer comes back from the computer must have been immanent in the data to begin with. Strictly speaking, this proposition must be true. But, then again, strictly speaking any answer that comes out of an ordinary experiment must have been immanent in nature to begin with. So why can't we see the answer without bothering to run the experiment?

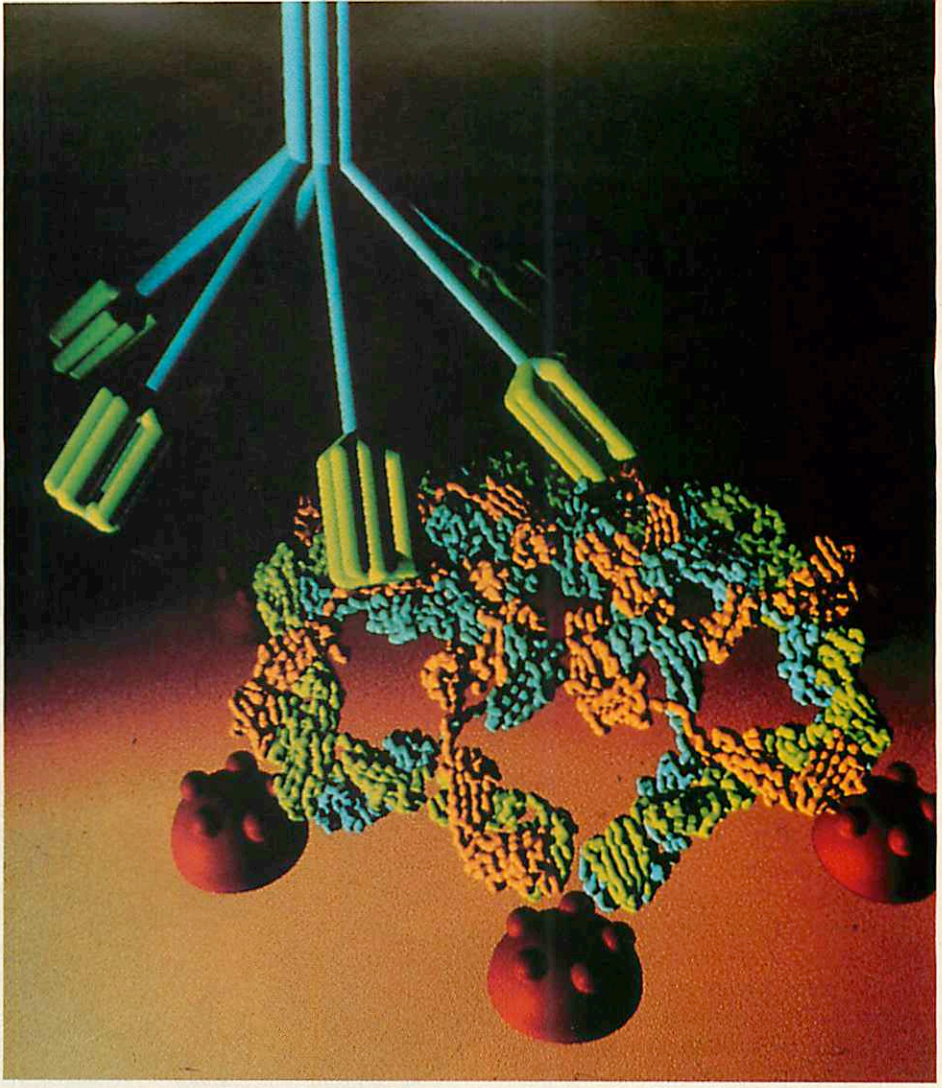

In this model six antibodies have grabbed molecules on a foreign cell (red spheres) and are linking it to the six-stemmed enzyme that will start its destruction.

PHOTOGRAPH BY GEOFFREY FORD, KREBS INSTITUTE FOR BIOMOLECULAR RESEARCH, UNIVERSITY OF SHEFFIELD

Setting aside these philosophical quibbles, computer simulations have produced a number of results that give a strong impression of novelty and surprise. A well-known instance came from the atmospheric studies of Edward Lorenz of MIT in the 1960s. Lorenz was running a computer simulation of the weather when he discovered by accident that his model of the atmosphere was "chaotic." Here the term "chaotic" has a special sense: it signifies that even a minuscule change in initial conditions can make an arbitrarily large difference in the outcome of a simulation. The fluttering of a butterfly's wings on Friday could change what would have been a sunny day next Thursday into a downpour. This exquisite sensitivity to disturbances was not a flaw of Lorenz's model; on the contrary, the real atmosphere is chaotic in exactly the same sense, and thus the model's instability could be considered its most realistic aspect. The presence of chaos in the computer simulation was nevertheless a surprise, because the equations defining the model were deterministic; they included no element of randomness. The study of "deterministic chaos" has since become a small industry in its own right.

Chaos also has a place in another surprising result of computational science, one that improves on the first simulations ever done. A group of workers at MIT, including Gerald Jay Sussman and Jack Wisdom, have built a "digital orrery"—an electronic computer specialized for the single purpose of calculating planetary orbits. With a mechanical orrery, inaccuracies in the gear train limit predictions to a few centuries at most. The digital orrery, on the other hand, was able to track the motions of the planets for 845 million years (which is roughly a fifth of the age of the solar system). Sussman and Wisdom found that the solar system, like Earth's atmosphere, is chaotic. In particular, the orbit of Pluto is unpredictable over intervals of more than a few hundred million years.

Computer simulation is not going to supplant either theory or experiment; there is no immediate prospect of a science in which computers dream up experiments, carry them out in their own imagination, and then announce their conclusions. Computer-aided science, however, surely has a future. Without it, most scientists necessarily focused on the tidiest problems, the simplest examples, the special limiting conditions, the isolated systems. They had to search for oases where the equations "come out even," because those were the only equations they could solve. Computer methods have opened science to a wider and freer and messier world. ❑