


Exploring a virtual universe

Cosmologists have already created entire universes within computers. Now astrophysicists are focusing on the fine details of asteroid collisions and supernovae. Govert Schilling investigates.



Heckling is not something scientists expect when they present results to their peers. But when Simon White revealed the image above — which shows a cluster of galaxies — during a meeting on galaxy formation held in Leiden last May, a member of the audience shouted out: “Is that an observation or a simulation?”

Such a question would have been unlikely 20 years ago. Then, many researchers viewed computer simulations with suspicion. But in the 1980s, computer models of how large-scale structures such as galaxy clusters evolved began to produce impressive results. More recently, new computing techniques have allowed smaller-scale processes, from

asteroid collisions to supernova explosions, to be modelled with unprecedented accuracy. White’s image was a simulation, but the fact that the question needed to be asked shows just how successful computational studies have become.

The early simulations helped to shed light on some of the mysteries of the Universe. Galaxies and galaxy clusters evolved from the soup of particles and radiation present after the Big Bang. The laws that control this evolution, such as Newton’s laws of motion, were well understood. But other factors, such as the role of particles of cold dark matter (CDM), were less well defined. Researchers suspected that CDM made up a significant proportion of the mass of the Universe, and so affected the way it had evolved. But because the particles do not emit light or electromagnetic radiation, they had never been observed directly.

To test different ideas about CDM’s role, cosmologists needed an experimental universe. Tweaking the initial conditions of virtual universes, or the laws that govern them, changes the way in which they evolve. By searching for the simulations that resulted in universes similar to our own, researchers were able to test different theories.

From the early 1980s onwards, computers were powerful enough to simulate useful vir-


Growth factor: a cosmology simulation showing galaxy clusters forming from cosmic dust.

tual universes. Researchers found that large-scale structures, such as clusters of galaxies, would only evolve if their simulations included large amounts of CDM. The studies gave cosmologists confidence in their assumptions, and most now believe that CDM makes up about 95% of the mass of the Universe.

In the dark

Despite these successes, many issues in cosmology and astrophysics, such as the formation of individual galaxies, remain difficult to simulate. “It’s relatively simple to trace the distribution of dark matter in the Universe,” says White, who directs the Max Planck Institute for Astrophysics in Garching, Germany. “But with luminous matter, there’s a lot more physics involved, such as star formation processes and supernova explosions.”

Fortunately, researchers such as White will soon be able to call on a new generation of supercomputers. This July, the Virgo Consortium of European researchers inaugurated its new Cosmology Machine, based at the University of Durham in England, which can perform 228 billion calculations per second (228 gigaflops). In the United States, a con-



Simulations are no longer a 'black art' for cosmology and astrophysics, says Mike Norman.

sortium headed by the San Diego Supercomputer Center at the University of California, San Diego, and the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, has also recently announced plans to build a system that can do a staggering 13.6 trillion calculations per second (13.6 teraflops). The system will be used by researchers in several fields, but US cosmologists and astrophysicists expect to be allocated some time on it.

Using these new machines, it may be possible to fix some of the problems with simulating galaxy formation. All large galaxies, for example, appear to be embedded in a halo of CDM. The halo's structure can be inferred by measuring the effect its gravitational field has on the motion of the galaxy it surrounds. But the results of observations of this effect do not tally with those produced by simulations.

All for one

Both of the new computing systems are examples of a recent trend in computer design. By linking networks of less powerful machines together to form clusters of processors, researchers are creating systems that can match the speed and power of bigger machines, but at a fraction of the cost. The NASA scientists who created the first such network in 1994 christened their cluster Beowulf — the title of an eighth-century English poem — and the name has stuck.

The San Diego and Durham systems are examples of relatively expensive Beowulf systems. More usually, such clusters are deployed by those with smaller budgets, and some of these less expensive set-ups are now being used to tackle the smaller-scale processes that have yet to be simulated successfully.

"Beowulfs are the platform of choice for astronomy departments that don't have unlimited money," says Derek Richardson, an astrophysicist at the University of Maryland in College Park. Together with colleagues at the University of Washington in Seattle, where he worked until last year, Richardson has used Beowulfs to model the way asteroids behaved in the early Solar System¹.

Small bodies, such as asteroids and comets, grouped together over millions of years to form planets. Richardson simulated

collisions between one type of asteroids and showed that only rarely — in around 30% of cases — do they stick together when they collide. His work suggests that collisions between equal-sized bodies played only a minor role in planetary formation, and that the gravitational pull of large asteroids on smaller bodies is the more important process.

But the Beowulf concept does have its drawbacks. Beowulf software has to be designed so that it can run on a large number of processors at once. A grid-like code, where space is divided into tiny sub-regions, is ideal. To model the fluid dynamics of the gases inside the Sun, for example, the Sun is divided into a series of cells. Each individual processor simulates the behaviour of gases in a small number of cells by modelling the physics taking place inside them and assessing the influence of their immediate neighbours.

Such problems lend themselves to Beowulf systems because each spatial cell is influenced only by neighbouring cells, and the amount of communication between processors is limited. But in other simulations, things are not so simple.

Erik Asphaug of the University of California, Santa Cruz, has used a low-cost Beowulf system to model the collision between the young Earth and a second proto-planet which led to the formation of the Moon². He divides the system up into a series of 'test particles', groups of which are simulated by individual processors. But the most important force in the system — gravity — acts across all the test particles, so each processor has to communicate with all of the others. This slows the simulation down. Asphaug admits that he is envious of people who have code that can be run in parallel. Richardson's asteroid simulations, in which gravity is also very important, suffer from similar problems.

Globular clusters — spherical groups of


stars that lie within, or orbit, galaxies — are also difficult to simulate using Beowulf clusters. A Beowulf network could, in principle, be used to model them. But, unlike the rocks in Asphaug's simulations, stars within the clusters cannot be lumped together as single test particles. Stars of different masses, for example, evolve in different ways. And because clusters contain hundreds of thousands or even millions of stars, simulations of them require an equivalent number of test particles.

Fruitful arrangement


Beowulf networks of millions of processors are not feasible, so researchers modelling globular clusters at the Institute for Advanced Study in Princeton have taken a different approach. Rather than writing software to run on general-purpose machines or networks, they are using machines that are hard-wired to solve the equations describing the gravitational interactions. The machines, known as GRAPE-6s, were designed by researchers at the University of Tokyo and, with a peak performance of tens of teraflops, were last year awarded the title of fastest computers in the world.

The Princeton team says that it will soon be able to run simulations of a few hundred thousand stars. The group plans to study how conditions in the early Universe influenced the evolution of globular clusters. For example, it hopes to explain why some clusters contain many millisecond pulsars — binary stars in which one member of the pair has sucked in matter from its companion star — whereas others contain only a few such systems.

But GRAPE-6 users can forget about using them to play games during their lunch-break. The machines are designed to tackle specific problems, and cannot be used for anything else. For some researchers, this narrow focus is a disadvantage. "For special



Starring role: a GRAPE-6 machine (above), custom-built to model globular clusters (left).



Crash course: Derek Richardson (left) has modelled colliding asteroids (above).




areas GRAPes are great,” says David Clarke, a computational astrophysicist at Saint Mary’s University in Halifax, Canada. “But many areas in astrophysics are messy, so you really need a general-purpose machine.”

For some of these messy problems, researchers are turning to individual supercomputers that can be adapted to simulate a range of different physical processes. The new 5-teraflops machine at the National Energy Research Scientific Computing Center, run by the Lawrence Berkeley National Laboratory in California, is one such example. Named after the late nuclear physicist Glenn Seaborg, the Nobel laureate and former associate-director at the laboratory, the machine is the most powerful non-classified, general-purpose supercomputer in the world.

Seaborg is currently being used to study supernovae, which are among the most physically complicated events in the Universe. The explosions involve extreme conditions, including very high densities and temperatures, as well as phenomena that are not well understood, such as strong magnetic fields and the production of gravitational waves. They also evolve extremely fast.

“Many years ago, we started with crude one-dimensional simulations,” says Adam Burrows of the University of Arizona in Tucson, one of a team of 12 astrophysicists working on the supernova study. “But one-dimensional supernovae don’t seem to explode.” Improved two-dimensional simulations have shed light on some processes, such as convection currents in the exploding star. Burrows’ collaboration, as well as a rival team at the Oak Ridge National Laboratory in Tennessee, now



Big bangs: Adam Burrows plans to run 3D simulations of supernova explosions.

plans to use three-dimensional versions to explore the problem in depth.

“We’ll take existing stellar evolution models and try to explode them,” says Burrows. He plans to run models of stellar evolution until the simulations reach the end of the stars’ life. At this point, a real star would start to collapse in on itself. By adding extra detail to the models, such as better descriptions of the nuclear processes within the star, he hopes to create simulations that do the same. The team will then analyse details of the virtual explosions and compare them with data from real supernovae.

In the process, the researchers hope to shed light on some of the unknowns in supernova research, such as the processes that control the production of neutrinos—subatomic particles with very low mass which are generated in the explosions. Neutrinos from explosions close to our Galaxy, such as from the supernova 1987A, can be detected on Earth and linked to their source. Researchers expect similar explosions to occur in our Galaxy every few decades or so. Different types of supernovae probably produce neutrinos in different numbers and with different energies, so better knowledge of these factors could help to classify future explosions.

Planet builders

Tom Quinn of the University of Washington in Seattle is using a Cray supercomputer to study another little-understood issue—that of how the Earth got its supplies of carbon and water. Like other planets, the Earth formed when solid matter in the disk of gas and rubble left over from the Sun’s formation grouped together over millions of years. But the part of the disk in which the Earth formed was too close to the Sun for water or carbon to have existed in solid form.

Some researchers have suggested that asteroids containing carbon and water collided with the Earth after being knocked out of orbits further away from the Sun. But initial results from Quinn’s simulations suggest that there was little transfer of matter between different orbits in the middle stages of planet formation, around 4.5 billion years ago³. This lends support to the alternative view that water and carbon arrived when the Earth was bombarded by comets later in its life, around 3.8 billion years ago.

Groups that do not have a Seaborg or a Cray of their own might still have the chance to access one, thanks to emerging techniques for sharing computing power. Clarke and col-

leagues at Saint Mary’s have recently bought a 10-gigaflops computer. Government funding was conditional on a fraction of the machine’s computing time being made available to other Canadian science institutions. This will be done through a network that links some of the country’s most powerful computers. Tasks sent to the network are distributed over whichever machines are free at the time, and users can select the type of machine most suited to their simulation.

Clarke’s new component for this network will be used by researchers at the new Institute for Computational Astrophysics, based at Saint Mary’s. Funding for the institute has been secured, and Clarke says he aims to have five or six staff in place by this time next year. The precise direction of the new team has yet to be decided. “We can’t cover the entire field,” says Clarke. “We might adopt a theme, such as the origins of stars and galaxies or the life cycle of stars.”

Simulation stimulation

The flurry of fresh activity at Saint Mary’s and elsewhere illustrates how important simulations have become. More powerful computers have undoubtedly played a role. But for cosmologists, better observational data—such as the precise measurements of the cosmic microwave background, the radiation that dates from just half a million years after the Big Bang—have also been vital. “These data have really pinned down the cosmological parameters,” says Mike Norman of the University of California, San Diego. “A few years ago, you had to run numerous simulations with different parameters. Now you can just take one model and compute the consequences in great detail.”

But such success can bring problems. As simulations become ever larger, finding the human resources to interpret them is becoming difficult. “Post-processing and analysing the huge amounts of computer data is very person-intensive,” says White. “And it’s almost impossible to automate.” There is also a need for new software. Computers are getting faster, says Richardson, but computer algorithms have to keep pace.

Such issues were probably the least of the worries for researchers performing the early simulations. “It used to be regarded as a black art, performed by people who possessed specific tricks,” says Norman. Many theoreticians refused to believe the results of simulations unless they had been verified by other means.

Now ideas for new observational projects often stem from computer models, showing just how trusted they are. For many in the field, simulations have joined observation and theory as the third pier of astronomy. ■

Govert Schilling is a writer in Utrecht, the Netherlands.

1. Leinhardt, Z. M., Richardson, D. C. & Quinn, T. *Icarus* 146, 133–151 (2000).
2. Canup, R. M. & Asphaug, E. *Nature* 412, 708–712 (2001).
3. Richardson, D. C., Quinn, T., Stadel, J. & Lake, G. *Icarus* 143, 45–59 (2000).