News Editor: Scott L. Andresen, sandresen@computer.org

Where Have All The Comets Gone?

By Pam Frost Gorder

A chunk of ice and rock 10 kilometers wide drifted through space. As part of leftover cosmic debris on the outskirts of a solar system billions of years old, the rock circled until a passing planet's gravity nudged it off course just enough to trace a new orbit, one that periodically carried it closer to its central star. Solar radiation washed over the rock, setting it aglow.

A close encounter with Jupiter tore the rock to pieces, and its remains spiraled inward at 60 kilometers per second—a speed that would cover the distance from New York City to Los Angeles in less than a minute and a half. When the fragments struck Jupiter, they detonated with more energy than all of the Earth's nuclear weapons combined.

It was a cataclysmic event like countless others that helped shape our solar system so long ago. What was different about this event was that it happened in 1994, the rock was Comet Shoemaker-Levy 9, and, for the first time, the world was watching.

This is a composite photo from separate images of jupiter and Comet Shoemaker-Levy 9, as imaged by NASA's Hubble Space Telescope. Photo courtesy of NASA.

Comet Modeling

Comets are the fossils of the solar system. From their chemical composition to the path they cut across the sky, they convey a wealth of information about how the planets first formed. Events such as the demise of Shoemaker-Levy offer clues about whether a similar collision is in Earth's future. But because sightings of these objects are so rare, astronomers must use computer simulations to fill the knowledge gap.

Planetary scientist Hal Levison and his colleagues at the Southwest Research Institute (SWRI) in San Antonio, Texas, recently drew on two computer models to solve a mystery: Why do we see so few comets?

Theoretically, the solar system is surrounded by billions of comets in a swarm called the Oort comet cloud. Of the thousands of Oort comets expected to fall into near solar orbit every year, we only spot about a dozen.

Levison and his team modeled the fate of thousands of fictitious comets. After calculating how many would hit a planet, hit the Sun, or be ejected from the solar system, some 99 percent of comets remained unaccounted for. The astronomers concluded in the June 2002 issue of the journal *Science* that these comets must simply break up—or, as Levison put it, "go poof."

What has Levison been working on since publishing the paper?

"Well, this may sound funny, but I am trying to prove it wrong," he says. "I think we are missing something very fundamental about the dynamics of comets."

As astronomers confront the question of the missing comets, they are forced to wonder whether their ideas about the solar system's origin and structure are lacking. The necessary calculations can involve large numbers of objects, complex gravitational interactions, or both. But most of the work can be done on a small computer cluster or even a single PC.

Levison's team modeled its fictitious comets on a small cluster of Unix workstations. "We really don't need any fancy hardware to do this work. Our problem is one of CPU speed and software efficiency, since the simulations usually take weeks to months of CPU time," he says.

To Levison, most astronomers are modeling problems that are "wide and shallow," meaning they want to track the behavior of millions of objects for only a few iterations. His problems are "narrow and deep," because he tracks a small number of objects over many iterations—for instance, the orbits of 1,000 comets over 40 billion iterations, or 4 billion years.

Comets trace such long orbits because they have been around since the solar system began. According to astronomers' current view, bits of rock and frozen chemicals circled the early sun, condensing from an amorphous sphere into a disk. Rocks collided, fusing into larger chunks with the chemicals buried inside. Some of the chunks grew bigger, eventually becoming Earth and the other terrestrial planets. The new planets' gravity kicked around the leftover bits until they were knocked into the outskirts of the solar system. Barely held in orbit by the distant Sun's gravity, comets in the Oort cloud feel the tugs of other passing stars and even the slow rotation of our Milky Way galaxy.

We must take all these gravitational effects into account when modeling comets, explains Paul Thomas, professor of physics and astronomy at the University of Wisconsin-Eau Claire (UWEC).

"The modeling of cometary orbits faces the worst challenges of solar system dynamics," he says. "Given that the long-term study of orbits for even sedate bodies like planets is limited by chaos considerations, evaluating the orbit of any single comet for even one period into the future is a difficult task."

Comet structure further complicates things, because the frozen chemicals trapped in the rock can vaporize when the comet ventures close to the Sun. Streams of gas and dust jet from openings in the rock, pushing the comet like rocket engines. What's more, the jets could "fire" differently each time a comet passes the Sun, because pockets of volatile material exist at different places inside the nucleus.

The mathematics is also complex, explains John Anderson, an astronomer at NASA's Jet Propulsion Laboratory (JPL). The equations are nonlinear and require numerical approximations. Also, the equations must account for movement with six degrees of freedom—three degrees each for the position and velocity of a comet. "Levison's comet project would have involved thousands of degrees of freedom," Anderson says.

Changing Views of the Solar System

Astronomers once explained away missing comets by suggesting that the comets had gone dormant—meaning that the Sun's radiation had stripped away the ice that creates the stereotypical comet glow in the sky as seen from Earth. Without their ice, comets would be small, dark, and much harder to see.

But Levison's latest work refutes the idea that many comets become dormant. He modeled not only the paths of thousands of hypothetical comets, but also the number of sightings that should hypothetically result. His conclusion: if many dormant comets were floating around the solar system, more of them would have been seen by now.

That's what led him and his colleagues to wonder if comets do indeed "go poof." They suggested in their paper that perhaps the frozen comets heat too quickly when approaching the Sun and burst into pieces.

For scientists who are trying to calculate the chances of Earth suffering a comet impact, Levison's work would seem to suggest that the danger is quite low.

At UWEC, Thomas has modeled how comets might have impacted the Earth in the early days of the solar system. His work grew into the book *Comets and the Origin and Evolution of Life* (Springer-Verlag, 1996). The evidence of collisions can be seen in the craters of the moon, and even craters here on Earth, he explains. But because scientists know so little about the dynamics of the system, Thomas and his collaborators try to estimate plausible fractions of impacts that may have been of cometary origin.

"If Levison's analysis stands up—and it seems very thorough to me—it might help reduce our estimates for impact risk from near-Earth asteroids that originated as dormant comets," Thomas says.

JPL's Anderson says Levison's work definitely influences his own study of cometary dynamics, particularly because of the problems it raises with the Oort cloud. Perhaps the comet deficit occurs because the cloud is disk shaped not spherical; Levison's analysis addresses this possibility. Or perhaps some important component of orbit dynamics is missing from the analysis.

At the January 2003 meeting of the American Astronomical Society (AAS) in Seattle, Anderson presented work on a phenomenon known as the Pioneer anomaly—an unexplained acceleration of the two Pioneer spacecraft toward the Sun—and the effect such a force could have on comet orbits.

While most of his work uses NASA's long-established software programs for spacecraft navigation, Anderson is able to model the Pioneer anomaly on a single PC using commercially available software such as Mathematica from Wolfram Research.

"We don't even know whether we should model the Pioneer anomaly as a constant force directed toward the Sun, or as a drag force in some sort of resisting medium beyond the orbital radius of Uranus," Anderson said at the conference.

If a previously unknown drag force exists, it could change the appearance of a comet's orbit, causing astronomers to interpret its location of origin incorrectly—again, altering scientists' picture of the solar system.

"Because we are far from convinced that comets experience the same anomalous acceleration as the Pioneer spacecraft, it is premature to claim that the Pioneer anomaly can solve some of the problems with the current observed comet distribution," Anderson said. "We simply suggest that this is an important idea that should be explored in future cometary dynamical simulations, perhaps using a brute-force approach with Beowulf clusters."

Beowulf Cuts Its Teeth—On a Comet

At least astronomers don't have to model one comet's tiny interactions with another. Shoemaker-Levy broke into more than 20 separate pieces during its fatal close encounter with Jupiter. Thomas Sterling, now a Faculty Associate at the Center for Advanced Computing Research at the California Institute of Technology, was working at NASA Goddard Spaceflight Center in 1994 when he and Donald J. Becker designed the first ever Beowulf computer cluster. NASA scientists used that first Beowulf to predict where and when the comet string would strike the planet.

They had to calculate Jupiter's pull on each individual fragment, as well as the potential gravitational effects of the Galilean moons and the Sun. But because the fragments didn't exert much gravity on each other, the scientists were able to divide up the problem with minimal communication between parallel nodes on the cluster.

"The medium-scale Beowulf systems could exploit the coarse-grain parallelism available in this problem. But beyond a certain scale, larger Beowulfs would not have been able to accelerate the problem much more," Sterling says.

If all the nodes on a cluster had to communicate with each other, as they normally do on a Beowulf, that would effectively kill the comet simulation, Levison says.

"As long as figuring out where the planets are is a trivial part of the simulation, we can let each CPU do that. Besides, it is much faster to let each node do the planet calculation redundantly than for them to communicate with each other," he adds.

Levison's research has inspired some innovation in Beowulf clusters, however. He, Peter Tamblyn, and other colleagues at SWRI are working on planet formation simulations. In such cases, the gravitational forces between the rocks, called planetesimals, are paramount. They had originally hoped that a Beowulf cluster would speed the simulations, but because the problem is still "narrow and deep," it created too much crosstalk between the nodes.

He described the project at the January 2003 AAS meeting. The new strategy is to multicast messages to all the nodes and optimize the acknowledgment portion of communication between them. Normally, a node sends an acknowledgement message to prove that it received data. But if the user assumes communication between nodes is nearly reliable, infrequent acknowledgements are sufficient for catching errors and recovering data, Tamblyn explains.

Sterling approves of the idea. "They are right to observe the inefficiencies that come from attempting to scale up collective operations and multicast communication," he says. "This is an area that can impose severe bottlenecks for some scientific simulations, and improving the overall communication behavior in this area could be a significant contribution to the community as a whole."

Key Codes

"The key is the development of fast codes—that is where we spend a lot of our time," Levison says. Following the orbit of a comet for 4 billion years is complicated enough in itself, but once the comet enters a solar system, its orbit is dominated by gravitational encounters with that system's planets. Even though the encounter might only last one day, the code must suspend its normal activity and deal with each planet separately.

"It has been very difficult to develop a code that can do this and preserve the simulations for long time scales," he says.

These kinds of problems are common to many large-scale simulations, including those that model climate patterns or interactions between atoms in a material. When coding systems emerge that make simulations easier, astronomers will be able to better understand how the solar system formed, where the comets come from, and where the comets have gone.

Pam Frost Gorder is a freelance science writer based in Columbus, Ohio.

BRIEFS

Mathematical Models Correctly Predict Cell Cycle Behavior

By Anne Jacobson

In finally verifying a 10-year-old mathematical model of cell behavior, both experimental and computational biologists are recognizing how valuable their collaborations can be. With the discovery, researchers from the disparate fields of cell biology and computer science are learning from each other and working to strengthen their partnerships.

"Biology is a very messy system, and working with mathematical models requires us to be much more quantitative and precise than we are used to," explains Virginia Tech cell biologist Jill Sible. For modelers, who can be very literal when working with numbers, the challenge is learning to accept the abundant noise inherent in experimental data, she says.

In this collaborative effort, Sible teamed with fellow Virginia Tech professor John Tyson, whose comprehensive model describing the molecular switch that governs cell division was published in 1993. To test Tyson's model, Sible and colleagues manipulated levels of cyclin, a regulating protein, and observed how the fluctuating worked to start, maintain, and stop mitosis, or cell division.

As Tyson's model predicted, the molecular switch required a significant push to go from on to off but was fairly stable in either position. Triggering mitosis (turning the switch on) required a push of at least 40 units of cyclin, but once the switch was on, cyclin levels could drop as low as 30 or 20 units without interrupting cell division. Likewise, stopping mitosis (turning the switch off) required a drastic fall to 16 units or less, but once off, cyclin levels could inch upward without effect.

The cell-cycle regulation behavior is more like that of a light switch, which requires equal force to move from off to on and back again than that of a buzzer, whose default position is off. In engineering circles, this lever-like behavior is called hysteresis.

"For cell biologists, there is nothing intuitive about hysteresis," Sible says. "Other groups have published models of the cell cycle that function without hysteresis, but it was not until we tested this model that we could determine that hysteresis is indeed the basis of the switches that start and stop mitosis." Sible, Tyson, and their team published the results in the 30 December 2002 issue of *Proceedings of the National Academy of Sciences*.

"Our recent paper verified a model that had already stood the test of time, but now we want to put models into a more workable context," Sible says. "We have lots of unexplained experimental data, and we'd like to use mathematical modeling as a tool for building hypotheses."

With collaboration from Tyson and other computer scientists, Sible also hopes to build a comprehensive mathematical model of the human cell cycle that is several orders of magnitude more complex than existing databases. Such a comprehensive and detailed model could predict behaviors that would be impossible to discover otherwise.

"Mathematical models reveal subtle interactions that we cannot capture experimentally," Sible explains. For example, some models show that modest two-to-threefold differences in concentrations—differences well below the radar of experimental biologists—can have an enormous effect on cell behavior. In addition to describing basic cell behavior, these types of quantitative predictions can help pinpoint novel drug targets.

"Our perspective has certainly been broadened by this collaboration," Sible says. "With the type of data that we cull from day-to-day experiments, I wouldn't have known to ask these questions."

Anne Jacobsen is a freelance science writer based in Washington, DC.