From *SIAM News*, Volume 40, Number 5, June 2007

Geosciences Conference Tackles Global Issues

By Barry A. Cipra

The SIAM Activity Group on Geosciences held its eighth biennial conference in Santa Fe, New Mexico, March 19–22. Invited presentations on inverse earthquake modeling, carbon dioxide sequestration, and ocean dynamics were complemented by minisymposia on porous media, multi-phase and multi-physics modeling, and wavefield imaging, to list only a few of the topics covered at the meeting, which brought together nearly 300 mathematicians, engineers, and computational geoscientists from industry, academia, and government labs around the world.

The geosciences activity group "is small but very focused," says Mary Wheeler, one of the founders of the group, which held its first meeting in 1993. (An early SIAM geosciences conference, held in 1989, preceded the formalization of the activity group.) Geoscience was once considered too limited a subject for an activity group, Wheeler recalls: "Many people thought it was too narrow!"—an odd attitude toward a field of literally global scope.

Today, with concern heating up over global warming and impending natural disasters from hurricanes to earthquakes, and with computational resources reaching the point that half a billion square kilometers of surface and a trillion cubic kilometers of volume no longer look unmeshable, more and more researchers from diverse disciplines are gravitating toward simulations of geoprocesses.

Simple, Flexible Models of Subsurface Flow

The combination of rich observational data and fast algorithms on highly refined meshes presents an exciting prospect for scientific computing, says Margot Gerritsen, a numerical analyst in the cleverly renamed Department of Energy Resources Engineering (formerly the Department of Petroleum Engineering) at Stanford University and current chair of SIAG/GS: The improved accuracy exposes weaknesses in the mathematical models. In the good old days, modelers could be content with overall qualitative agreement between computation and observation, excusing discrepancies as the innocent progeny of sparse data and coarse grids. (Indeed, it is clear that many algorithms go astray as mesh sizes tend to zero.) The situation nowadays can be likened to the quandary Kepler faced as he tried fitting circles to Tycho Brahe's meticulous records of planetary wanderings.

Not that the database is anywhere near complete. "In fact, experimental data is often lacking," Gerrritsen points out. "We should never lose track of the need for laboratory experiments and field observations. Too many people nowadays rely on computer models—but they are only as reliable as the physics put in."

In an invited talk at the conference, Gerritsen described some of the issues involved in modeling complex processes in subsurface flow. She is particularly interested in problems of enhanced oil recovery and CO_2 sequestration. (Gerritsen jokes that subsurface studies are natural for her, given her sub-sea-level upbringing in the Netherlands.) Rising oil prices have encouraged petroleum companies to consider ways of squeezing additional crude from "depleted" fields or extracting highly viscous "heavy" oil, which is more plentiful than conventional oil but, given its tar-like consistency, is tricky to coax to the surface. (See "Heavy Oils, Parts I and II," by Zhang-xin (John) Chen, *SIAM News,* April and May 2006; http://www.siam.org/news/archives.php.)

Gas injection, one of the techniques used for enhanced oil recovery, maintains pressure in the reservoir, swells the oil phase, and reduces viscosity. The idea, in essence, is that forcing cheap steam, carbon dioxide, or natural gas down one well will drive an equal volume of valuable oil up another. This in itself involves lots of interesting physics and numerics, Gerritsen points out, but she sees no need to stop there. "Let's make our lives a little bit more exciting by burning things," she says.

In situ combustion, or ISC as it's known, creates a slow, smoldering front; the front sets up a steam flood that drives the injected gas, while burning just a small fraction of the oil—and often upgrading the oil along the way. ISC is a major computational challenge. "The kinetics cause all sorts of problems when you start modeling," Gerritsen says.

For one, the time scale of the chemical reactions is extremely small compared to time scales of other physical processes. This makes the differential equations severely stiff, calling for specialized timestepping. The kinetics are also strongly sensitive to the phase behavior of the fluids, as Gerritsen and colleagues showed in a recent study using highly accurate numerical tools. In all, she says, ISC offers "a wealth of fantastic problems" for numerical analysts.

Gerritsen's group emphasizes adaptivity in the algorithms for grid generation, as well as in the numerical solvers. Models should be simple and flexible, Gerritsen believes; the focus should be on getting the correct physics into the model, rather than on fancy mathematics.

Her group's simulator for gas injection processes takes the mass-balance transport equations and splits them along one-dimensional streamlines, using adaptive Cartesian grids. "I like building simulators," Gerritsen says, noting that computations can lead to improved understanding of the physical processes and can guide the design of experiments for richer, more detailed data.

Earthquake Inversion—Algorithmic Challenges

The impending petaflops era both excites and concerns Omar Ghattas of the University of Texas at Austin. Will modelers be able to properly leverage all that computational power? he wonders. In 2002, when the Japanese Earth Simulator seized the computational speed record, at a then-blazing 35.86 teraflops on 5120 processors, U.S. scientific policy circles saw it as a sort of silicon Sputnik. The response has been a suite of initiatives at the Department of Energy, the National Science Foundation, and other agencies, aiming to restore U.S. computational bragging rights. Hardware-wise, the effort has succeeded: In the November 2006 Top 500 supercomputer list, U.S. machines occupied the top four places, and seven of the top ten, with IBM's BlueGene/L at Lawrence Livermore National Laboratory in first place, running at 280.6 teraflops on 131,072 processors. The challenge now is algorithmic—making sure that the power of these elegant brutes can be harnessed to solve scientific grand challenges.

For Ghattas, one of those grand challenges is earthquake inversion. Given a detailed model of the earth's properties in a region of interest, it's straightforward, if laborious, to simulate the effects of a large rupture, say along the San Andreas Fault, by solving well-known seismic wave-propagation equations.

But "that sort of begs the question, where do you get the earth model from?" Ghattas says. Short of aerating Southern California by taking zillions of core samples, the obvious answer is to infer what's underground from historical earthquake records. If you know how cause leads to effect via the earth model, you can hope to work backward to recover the model.

That's an inverse problem—and a seemingly intractable one at that. Algorithms for solving inverse problems often require running thousands of "forward" computations. For earthquakes, a single high-resolution forward computation can take the better part of a day on a parallel machine with a few thousand processors. Three years is a long time to wait for a result that's likely to be questionable. But on a few hundred thousand processors—and with the right kind of algorithm—the wait time might reduce to a matter of days and the result might be more accurate.

Ghattas and his collaborators in the Quake Project will soon have access to their own "earth simulator," a Sun/AMD supercomputer named Ranger being installed with NSF support at the Texas Advanced Computing Center. Expected to begin operating later this year, Ranger will run at a peak of 530 teraflops on more than 62,000 processing cores, with 125 terabytes of memory and a 1.9-petabyte disk. "There's a tremendous sense of excitement" at the prospect of doing science on such a behemoth, Ghattas says. He and his colleagues are gearing up to perform inversions of large 3D models of southern California with available earthquake observations. The calculations are practical only because of the new machines and advances in algorithms; the technology of just a few years ago would be hard pressed to solve problems of these types.

A typical model covers a region of dimension 600 km × 300 km × 80 km, and is subdivided into a mesh of 140 million cells. Initial computations on machines with a few thousand processors, such as Lonestar at TACC and Big Ben at the Pittsburgh Supercomputing Center, indicate parallel efficiencies in the range of 85%. Ghattas anticipates similar parallel efficiencies on machines with tens of thousands of cores. (Parallel processing inevitably entails efficiency losses, due to the processors' need to communicate with each other and coordinate their computations. One of the chief challenges facing computational science is to make the transition from serial to parallel computation. Computational scientists cannot afford to ignore this challenge, Ghattas points out, as even low-end laptops now come with multiple cores.)

For the inverse problem, the team is taking a nonlinear least-squares optimization approach, minimizing the data–model mismatch subject to the requirement that the model adhere to the usual physical equations. The numerics present Scylla-and-Charybdis challenges: High-frequency waves passing through relatively smooth media produce lots of local minima for the optimization problem, while low-frequency waves in rough, heterogeneous media lead to ill-conditioned calculations. As Ghattas puts it, "you can have the world's greatest model, but if you have an algorithm that doesn't scale, the game is over." He hopes they'll have some answers before the Big One makes the calculation moot.

Computational Challenges of Wind and Water

Major earthquakes aren't the only catastrophes calling for computational study. Wind and water wreak havoc as well. In a session on multiphysics modeling, Joannes Westerink of the University of Notre Dame presented an analysis of storm surges from hurricanes Katrina and Rita. Westerink's group combined a detailed topographic model of the Gulf Coast, from Sabine Lake on the Texas–Louisiana border to Mobile Bay, Alabama, with extensive measurements from the hurricanes themselves. (Katrina in particular was a "very data-rich event," Westerink says.) The model uses a mesh with more than two million nodes (2,137,978, to be precise) and includes information on the frictional resistance to storm surge arising from vegetation (e.g., marsh, cypress forest) and other land usage. "We're trying to put as much physics into the model as we can," Westerink says.

The "hindcast" results are encouraging. Using one-second timesteps and algorithms tailored to high velocities and large spatial gradients, the model, which computes a day's worth of surge in just under an hour, postdicts water levels to within about 10% of measured values.

Moreover, the largest errors occur in areas where physical effects currently unaccounted for in the model are clearly important. The researchers plan to make the mesh at least partially three-dimensional, couple the model with models of precipitation and run-off, and incorporate air–sea momentum trans-fer. All this and more should be possible as com-puting itself surges from tera- to peta-scale.

Analysis of the air–sea interface in a hurricane could be a bit tricky, notes William Dewar of Florida State University, who gave an invited presentation on open problems in ocean dynamics. Surface water beneath a hurricane is so whipped up that a sharp boundary may be replaced by a density continuum, from liquid to vapor. Measurements are obviously hard to come by, although Dewar mentions plans for (unmanned) submersibles that can cruise at a comfortable depth and poke up instruments to test the waters.

Winds and tides are the usual suspects in ocean energetics, Dewar says, but it's possible that ocean life—especially zooplankton, which exist in huge quantities—also plays an important role in turbulent mixing.

"Surprisingly little power is involved in mixing the modern ocean," Dewar says. A single, hundred-watt kitchen blender would suffice to mix a cubic kilometer at observed rates. The incessant motion of myriad critters, such as salps (filter feeders similar to jellyfish, but with the beginnings of a backbone), which propel themselves vertically by enveloping and ejecting tiny quantities of water—not to mention the titanic battles between sperm whales and giant squid (of which there may be a billion worldwide)—contribute terawatts of power. "Clouds" of zooplankton, for example, migrate daily from the surface to as far as a kilometer below, the equivalent of a 400-gigawatt blender. All this bioenergetics raises an interesting question: Could overfishing be yet another human mechanism of climate change?

"The Modeling Problem of a Lifetime"

The 25-gigaton gorilla of climate change is, of course, carbon dioxide.

Humans currently add roughly 25 gigatons of CO_2 to the atmosphere each year, primarily from coal-burning power plants and gas-guzzling vehicles. (Today's atmosphere contains approximately three teratons of CO_2, about 800 gigatons more than it did before the Industrial Revolution. The numbers are usually given in parts per million by volume: 383 ppmv today versus an estimated 280 for the late 18th century.) The annual increase is anticipated to rise to 50 gigatons by 2050, if "business as usual" persists.

Climatologists have all but given up on returning to pre-industrial levels of CO_2, or even holding the level to its present parts per million; they would be content to keep the *derivative* constant, at today's 25 gigatons/year. Among the proposed schemes for accomplishing this is the notion of CO_2 sequestration: pumping carbon dioxide deep underground in the hope that it will stay put for thousands of years.

Sequestration is feasible, at least in principle. There's a lot of room underground, in depleted oil and gas reservoirs, unmineable coal seams, and, especially, deep saline (nonpotable) aquifers. It's been estimated that U.S. saline aquifers have a CO_2 storage capacity of 500 gigatons. In theory, CO_2 injected as a supercritical gas would be held in place by overlying cap rock, slowly dissolve into the brine, and eventually react chemically with the mineral matrix to form, for example, calcium carbonate (aka limestone). But will it work in practice? That's where analysis and simulation come in.

"It's the modeling problem of a lifetime," says Wheeler. CO_2 sequestration is a multiphase fluid flow problem in porous media, involving both physics and chemistry at scales ranging from micrometers to kilometers spatially and seconds to centuries temporally. One of the key concerns is leakage: Being a gas, carbon dioxide will try to work its way upward. How long will it take to do so, and how much will remain behind? (CO_2 isn't toxic per se, so a little leakage is okay, although determining what constitutes a "little" is a policy decision. The possibility of massive, gas-spewing ruptures is another matter.)

Several speakers in a string of minisymposia described computational methods and analytic models for estimating CO_2 leakage. Susan Minkoff of the University of Maryland, Baltimore County, presented a one-dimensional, single-phase model for lateral leakage of CO_2 escaping from a reservoir straight up a vertical fault. She stressed the importance of developing a certification framework for CO_2 storage sites that is transparent and simple. The policy procedures and the models that support them must be understandable and convincing to the regulators, the stakeholders, and the public at large, who tend to be suspicious of complex models, Minkoff says. (In her own group's model, she notes, "We used a very simple Taylor series: constant!")

Sally Benson of Lawrence Berkeley National Laboratory described models tailored to the Frio Brine experiment, in which 1600 tons of CO_2 were injected at a depth of 1540 meters into the Frio Formation on the Gulf Coast of Texas over ten days in October 2004. Measurements were taken in an observation well 33 meters from the injection well. (The experiment was led by the Bureau of Economic Geology at the University of Texas, in collaboration with GEO-SEQ, a consortium led by Benson and Larry Myer of LBNL.) According to Ben-son, one-dimensional simulations do a decent job of matching the data, but they require the assumption of high residual saturation. Two-dimensional simulations that include gravity and heterogeneity do better. (Gravity turns out to be the more important factor.) The dissolution of CO_2, Benson says, depends strongly on flow geometry.

Ruben Juanes of MIT presented a multiscale method for simulating storage and buoyancy-driven migration in highly heterogeneous formations. A coarse grid handles the pressure equation, while the saturation equation is solved on the fine grid. The process of "imbibition"—the displacement of one fluid by another—features a hysteresis effect at the pore scale: Even if you drive brine back in, clusters of residual CO_2 are left behind—in essence, small fizzy bubbles are trapped in watery rock.

Buoyancy is a big concern, especially if the aquifer slopes: CO_2 in such a setting could migrate hundreds of miles from its injection site. Surprisingly, though, as Marc Hesse of Stanford showed, a gentle slope could be a CO_2 sequestor's best friend.

The key is an interplay of "slumping" and "sliding." A plume of gas trapped under perfectly horizontal cap rock will expand laterally as it slumps toward the ceiling; the evolution can, in fact, be modeled analytically with similarity solutions. The long-term storage takes the form of residual saturation in the region vacated by the plume. If the cap rock slopes, the plume slides as well as slumps. Initially, and for small slopes, slumping dominates and the evolution of the plume volume obeys a power law. But eventually (in the analytic model, at least), there is a transition to superexponential behavior with sliding dominant (see Figure 1).

One of the easiest ways for the CO_2 in a plume to get back into the atmosphere is through an abandoned well. Even a filled-in or capped well might provide a direct route upward. Texas alone has more than a million such manmade fissures. The Alberta Basin in Canada is another densely welled area. Michael Celia of Princeton

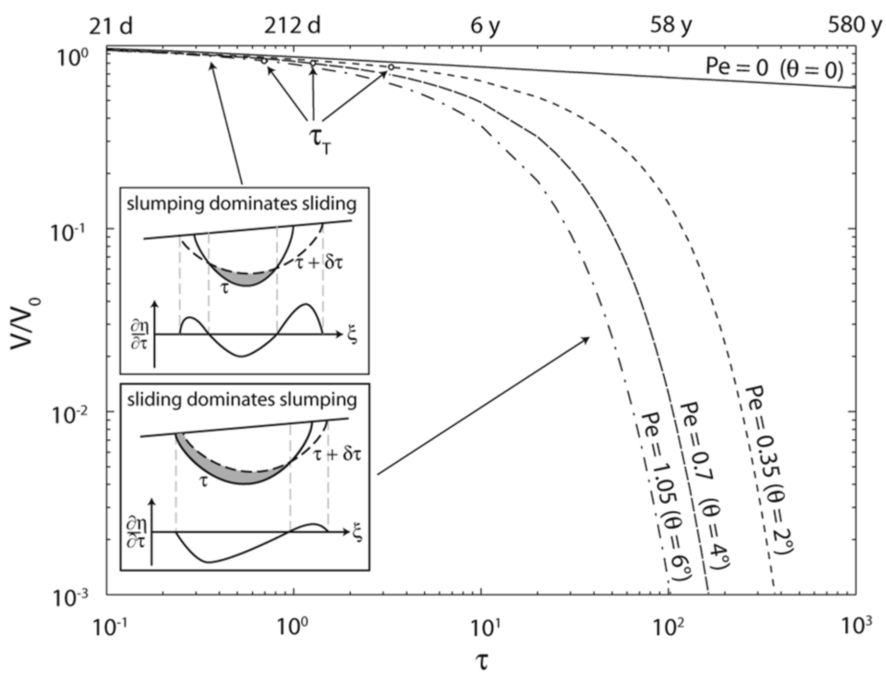

Figure 1. *Evolution of CO_2 plume volume: A transition from power-law to super-exponential behavior is seen as "slumping" gives way to "sliding." (Figure courtesy of Marc Hesse, adapted from "Scaling Analysis of the Migration of CO_2 in Saline Aquifers," by Marc Hesse, Hamdi Tchelepi, and Franklin Orr, Jr., Society of Petroleum Engineers Technical Conference, 2006.)*

and Jan Nordbotten of the University of Bergen, Norway, described semi-analytic models and Monte Carlo simulations of the potential for CO_2 leakage in a square region 50 kilometers on a side near Wabamun Lake, southwest of Edmonton, which is dotted with oil wells, at an average spatial density of about 0.5 wells per square kilometer (see Figure 2). To a depth of nearly 3000 meters there are a dozen sandstone and limestone aquifers, separated by shale and other aquitards. The locations and depths of the wells are known (most are between one and two kilometers deep). What's uncertain is the condition of the wells, in particular properties like their effective permeability.

"Given that we know very little about abandoned wells, we have to assume something about them," Nordbotten says. The researchers produced a probability distribution for the wells' permeability, then ran Monte Carlo simulations for 32 years worth of plume migration, obtaining a histogram for the likelihood of leakage. They conclude that existing wells (and more are being drilled every year) are potentially important pathways, and stress the need for an experimental program to better assess their properties.

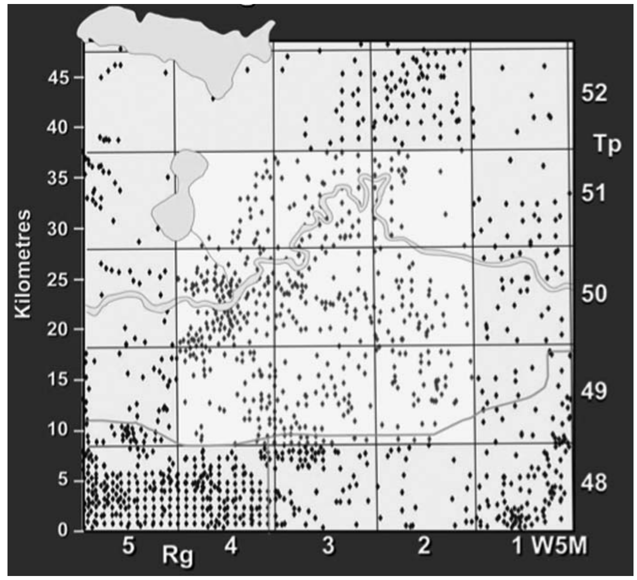

Figure 2. *The oil-rich area near Wabamun Lake in Alberta, Canada, is dotted with wells. Could they pose a problem for CO_2 sequestration? Michael Celia and colleagues hope to find out. (Figure courtesy of Michael Celia, prepared by Stefan Bachu, Alberta Energy and Utilities Board.)*

A Call for Caution in Geoscientific Modeling

As rich as the data is, uncertainties abound when it comes to putting carbon back in the ground, and limitations of the models need to be acknowledged. In an invited talk, Gary Pope, director of the Center for Petroleum and Geosystems Engineering at the University of Texas at Austin, sounded a clarion call for caution in geoscientific modeling.

The danger, he says, is falling in love with limited—and often inaccurate—data, being beguiled by the beautiful color images computers use these days to summarize results, and overlooking or underplaying the models' deficiencies. Subsurface flow models, some of which have accreted code like so much binary sedimentation, are notorious for incorporating flawed physics, Pope points out. As for their technicolor output, they not only don't represent reality, in many cases they aren't even faithful to the models that produce them.

Merely running simulations and publishing pretty pictures, Pope says, "is a form of intellectual dishonesty."

Pope urged researchers to continually revisit the assumptions of their models and to "use your imagination" not just to solve problems, but to change them. He described studies of surfactant-enhanced aquifer remediation (SEAR), in which wetting agents are introduced along with injected water to flush out dense nonaqueous-phase liquid (DNAPL) contaminants. "Natural attenuation is very hard to predict, very uncertain and very slow, whereas injecting water following SEAR is simple, easy to model and polishes off the contaminant in months rather than years," he says. In conjunction with experiments conducted at Hill Air Force Base in Utah, Pope and colleagues ran simulations that made actual predictions—that is, numerical results obtained in advance of the experiments—in good agreement with field observations.

Small pilot studies are one thing, large-scale operations quite another, Pope stresses. It's been hard enough to characterize large aquifers in oil reservoirs, where the economic incentives are great, he says. Modelers should keep their models as simple as possible and invest time in sensitivity analyses to quantify the uncertainties in their results. As Spiderman might put it, with great computational power comes great scientific responsibility.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.