From *SIAM News*, Volume 42, Number 3, April 2009

Recreating the Great San Francisco Earthquake

By James Case

Few natural disasters are in greater need of prediction, or seem more difficult to predict, than earthquakes. Even along California's San Andreas fault—surely the world's most intensively studied earthquake zone—confusion reigns. Whereas a study completed in 2006* concludes that the fault is currently under sufficient tectonic stress to generate an earthquake of magnitude 7.0, nobody pretends to know when such an event will occur, where along the 800-mile fault the epicenter will be, or how the rupture will propagate along the fault, bilaterally (as during the great San Francisco quake of 1906) or from south to north (as, apparently, on most other occasions in the recent geological past). "It could be tomorrow, or it could be 10 years from now," said Yuri Fialko, author of the 2006 study.

The great San Francisco earthquake began shortly after five o'clock on the morning of April 18, 1906. It lasted for about a minute and a half, and would have measured roughly 7.9 on either the once popular Richter scale or the more modern moment magnitude scale, had either existed at the time. The area of devastation was about 400 miles long, and extended about 30 miles on either side of the fault line. About 3000 people died as a result of the quake and the four-day fire that ensued. Some 225,000 were left homeless, and about 28,000 buildings were destroyed.

(During the 1970s, improved seismological equipment made it possible to determine after the fact the approximate location at which a particular geological fault had ruptured during an earthquake, as well as the amount of energy released there. The latter quantity is known as the seismic moment. Because the Richter scale was so widely understood, seismologists devised a logarithmic scale comparable to Richter's—called the moment magnitude scale—for recording seismic moments. Although the two scales measure quite different effects, they ordinarily produce virtually the same number, thereby facilitating the comparison of modern earthquakes with those measured in the past using the Richter scale. Many news organizations no longer mention the scale on which a particular earthquake was measured, reporting only the numerical magnitude.)

Rival Models of 1906 Earthquake

Recently, in part to commemorate the traumatic event and in part to assess the state of the art of earthquake modeling, a concerted effort was made to simulate the 1906 disaster. Teams from Stanford University, UC Berkeley, URS Corp., and Lawrence Livermore National Laboratory constructed rival models of the quake, each using their own software (including the open-source software WPP at LLNL) in conjunction with century-old assessments of actual historical damage and geological data compiled over decades by the U.S. Geological Survey. Each team presented its findings in San Francisco in 2006, at the annual meeting of the Seismological Society of America, commemorating both the quake and the 100th anniversary of the founding of the society. The fact that all four teams reached quite similar conclusions suggests that real progress is being made toward an accurate understanding of earthquake phenomena.

An obvious way to compare model findings uses the "tremor histories" calculated for specific locations. Such histories plot the speed at which a specific point on flat ground moves, as if that motion were confined to an imaginary fixed horizontal line. Small preliminary speed oscillations are typically followed by a series of larger but gradually subsiding ones, as indicated in Figure 1. The histories calculated by the four competing models for various locations of interest agreed fairly closely in frequency, in amplitude, and in onset time, both with one another and with the historical record.

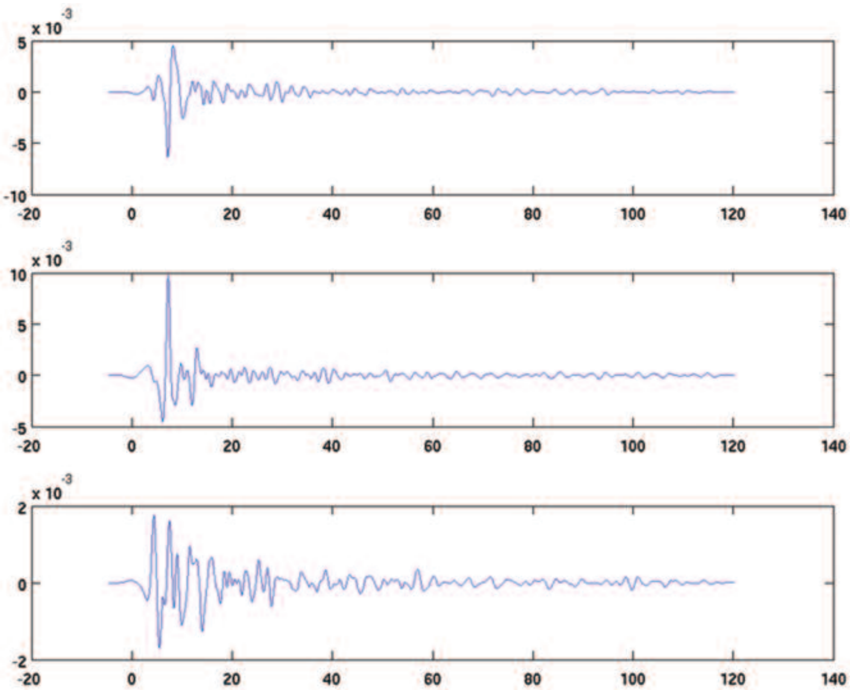

Figure 1. *Recorded ground velocities at seismic station SCCB for the magnitude 5.4 Alum Rock, California, earthquake, which occurred in October 2007. From top to bottom: velocities in the east, north, and vertical directions as a function of time, after low-pass filtering to frequencies up to 1 Hz. From Anders Petersson.*

The San Andreas fault forms the tectonic boundary between the Pacific and North American plates. It extends from the vicinity of Coos Bay, Oregon, to that of the Salton Sea in southern California. The quake of 1906 occurred along the northern section, which follows the shoreline south to San Francisco Bay, then runs inland through the Santa Cruz mountains to the neighborhood of Hollister, California. The

*Yuri Fialko, "Interseismic Strain Accumulation and the Earthquake Potential on the Southern San Andreas Fault System," *Nature*, 441 (2006), 968–971.

two plates slid several meters relative to one another during the 1906 quake; the long-term relative slippage rate appears to average only 33–37 mm per year. Because the North American plate is moving roughly south–southeast relative to the Pacific plate, the coastal mountain ranges are gradually growing taller. Detailed measurements of the plates' relative motion reveal that movements along the San Andreas and its various branch faults account for only about 75% of the plate displacement, the rest taking place in the less studied area east of the Sierra Nevada.

Computational Challenges

The team from LLNL was led by Anders Petersson, an applied mathematician, and Arthur Rodgers, a seismologist. In addition to the centennial study, which was conducted over a period of roughly two years, most of the mathematicians involved had participated in the serpentine wave propagation (SWP) project, headed by Petersson. That project was a broad study of the propagation of waves in nature, be they sound waves, water waves, electromagnetic waves, or seismic shear and compression waves. To accommodate their model of the San Francisco Bay area, the SWP team used two LLNL supercomputers, the smaller of which has 2300 processors and runs at a peak speed of about seven teraflop/s. The larger, nicknamed "Thunder," has 4096 processors and runs at 21 teraflop/s. The Department of Energy's Office of Advanced Scientific Computing Research contributed financial support. The 1906 earthquake simulations were also supported by an internally funded laboratory research and development project at Lawrence Livermore.

In a plenary lecture at the 2008 SIAM Annual Meeting in San Diego, Petersson described computational challenges encountered by the LLNL team in simulating the 1906 quake. Most of the difficulty, he said, resulted from the irregular topography and the heterogeneous media, from solid bedrock to loose sedimentary soil, characteristic of the earth's crust in the San Francisco Bay area. As a consequence, the shear and compression waves propagated at speeds that varied by a factor of eight within the three-dimensional box-shaped region—containing the whole of the San Francisco Bay area—to which the model applied. Conventional finite difference methods lose numerical stability when confronted with speed variations of such magnitude. The team invented a new finite difference scheme, satisfying the summation-by-parts principle, which enabled them to retain numerical stability at all times and places.

The basic task was to solve the elastic wave equation inside the box (550 km long, 200 km wide, and 40 km deep) within which the 1906 tremors were destructive. After discretization of the space variables, the elastic wave equation can be reduced to a system $\mathbf{v}_{tt} = A\mathbf{v}$ of ordinary linear differential equations for the unknown displacement vector $\mathbf{v} = (u_1, u_2, \ldots u_N)^T$. It is easily shown that the solutions of this system are uniformly bounded if and only if the eigenvalues of A are all real and negative, while the eigenvectors form a complete set. Such is the case if A is negative definite and symmetric.

With a grid spacing of 125 meters, the model contains 2.3 billion grid points; with three degrees of freedom at each one, there are 6.9 billion unknown functions of time to be evaluated. In fact, the LLNL group discretized the time dimension as well as the space dimensions, by breaking the first 300 seconds of earthquake motion into 30,000 subintervals of about 0.01 second each. That left them with roughly 200 trillion highly structured linear algebraic equations to be solved for an equal number of unknowns. Small wonder supercomputers were needed. The results obtained are perhaps best exemplified by three pictures depicting the spread of the tremors away from the epicenter and along the San Andreas fault (Figure 2).

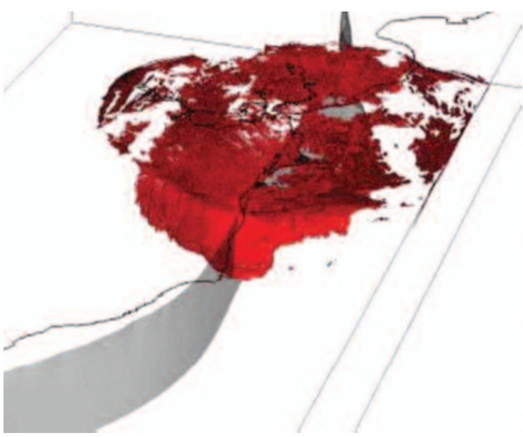

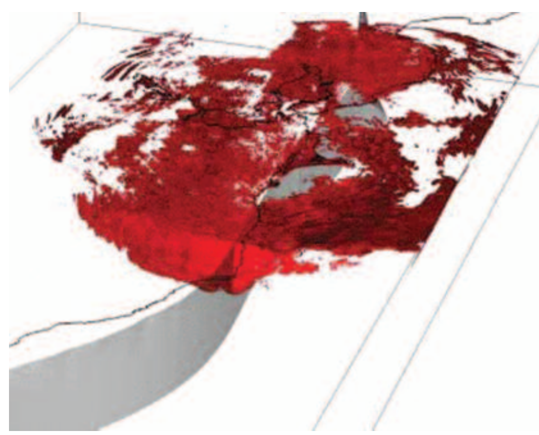

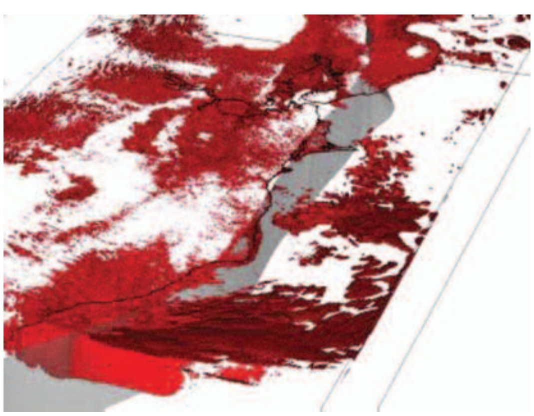

Figure 2. *The propagation of shear waves (in red) through the three-dimensional model of the greater San Francisco Bay area, looking southeast toward the Bay Area from the Pacific Ocean. The San Andreas fault surface is shown in gray, the coastline of northern California in black. Clockwise from top left: 22.5 seconds, 30 seconds, and 60 seconds after the start of the earthquake. The illustrations, from Anders Petersson, also appeared in DEIXIS, 2007–08, U.S. Department of Energy.*

Although topographic features have little effect on the propagation of longer waves, peaks and valleys do tend to interfere with shorter (higher-frequency) waves. The LLNL team was therefore pleased to discover that they could extend their summation-by-parts technique to non-planar free surfaces via the introduction of curvilinear coordinates, which map the original domain with the irregular topography to a domain with a flat top surface. The modified procedure was validated through numerical experimentation.

A particularly interesting experiment involved the break-up and delay of a rectilinear wave front advancing across a flat elastic plane upon which stands an isolated three-dimensional Gaussian hill. In the absence of any hill, such a wave would cause the earth beneath an observer's feet to rise above, then fall below, then rise again above its initial elevation, before permanently settling back to its original level. A short distance from such a hill, the observer would experience a slightly weaker version of the same sequence, followed by one or more similar (but still weaker) aftershocks. The timing would depend on the distance from the hill. On or very close to it, the various shocks would be jumbled together into a more complicated motion. These experimental results were in good agreement with those obtained with other numerical techniques.

Other computational refinements included the use of local mesh refinements in critical portions of the model region, and the substitution of visco-elastic for elastic modeling. The former sped up computation by as much as 100-fold; the latter was found to attenuate oscillations in a realistic manner.

To calibrate the model, the team ran it for a variety of more recent, less violent, and more accurately measured earthquakes in the Bay Area. Pointing out that more sophisticated rupture models are needed to simulate larger quakes, Petersson expressed the hope that earthquake models would in time be linked to structural damage models for buildings, bridges, and the like.

James Case writes from Baltimore, Maryland.