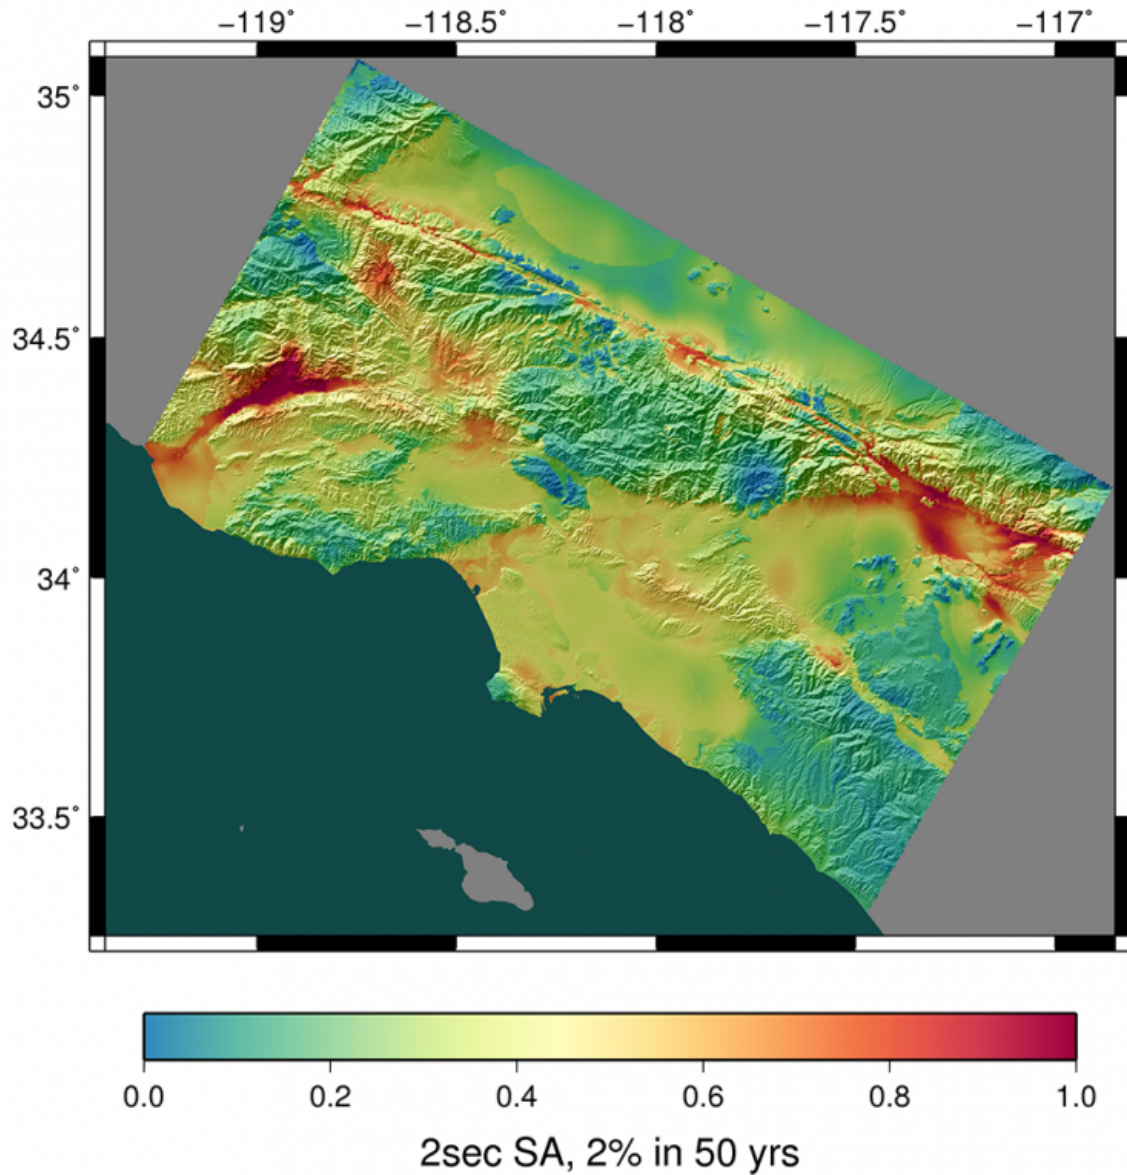


Titan takes on the big one

November 11 2015



The CyberShake seismic hazard map shows the magnitude, or level of shaking,

for the Los Angeles region, defined by the amount of change of a surface or structure in a 2-second period, with a 2% probability of increasing within the next 50 years. The map provides engineers with vital information needed to design more seismically safe structures Credit: Scott Callaghan, Kevin Milner, and Thomas H. Jordan (Southern California Earthquake Center)

The San Andreas Fault system, which runs almost the entire length of California, is prone to shaking, causing about 10,000 minor earthquakes each year just in the southern California area.

However, cities that line the fault, like Los Angeles and San Francisco, have not experienced a major destructive earthquake—of magnitude 7.5 or more—since their intensive urbanizations in the early twentieth century. With knowledge that large earthquakes occur at about 150-year intervals on the San Andreas, seismologists are certain that the next 'big one' is near.

The last massive earthquake to hit San Francisco, having a 7.8 magnitude, occurred in 1906, taking 700 lives and causing \$400 million worth of damage. Since then, researchers have collected data from smaller quakes throughout California, but such data doesn't give emergency officials and structural engineers the information they need to prepare for a quake of magnitude 7.5 or bigger.

With this in mind, a team led by Thomas Jordan of the Southern California Earthquake Center (SCEC), headquartered at the University of Southern California (USC) in Los Angeles, is using the Titan supercomputer at the US Department of Energy's (DOE's) Oak Ridge National Laboratory (ORNL) to develop physics-based earthquake simulations to better understand earthquake systems, including the potential seismic hazards from known faults and the impact of strong

ground motions on urban areas.

"We're trying to solve a problem, and the problem is predicting ground shaking in large earthquakes at specific sites during a particular period of time," Jordan said.

Ground shaking depends upon the type of earthquake, the way a fault ruptures, and how the waves propagate, or spread, through all 3-D structures on Earth.

Clearly, understanding what might happen in a particular area is no simple task. In fact, the prediction involves a laundry list of complex inputs that could not be calculated all together without the help of Titan, a 27-petaflop Cray XK7 machine with a hybrid CPU-GPU architecture. Titan is managed by the Oak Ridge Leadership Computing Facility (OLCF), a DOE Office of Science User Facility located at ORNL.

Running on Titan, the team uses SCEC's CyberShake—a physics-based computational approach that integrates many features of an earthquake event—to calculate a probabilistic seismic hazard map for California. In May, Jordan's team completed its highest resolution CyberShake map for Southern California using the OLCF's Titan.

Shaking It Up

One of the most important variables that affects earthquake damage to buildings is seismic wave frequency, or the rate at which an earthquake wave repeats each second. With greater detail and increases in the simulated frequency—from 0.5 hertz to 1 hertz—the latest CyberShake map is the most useful one to date and serves as an important tool for engineers who use its results to design and build critical infrastructure and buildings.

Building structures respond differently to certain frequencies. Large structures like skyscrapers, bridges, and highway overpasses are sensitive to low-frequency shaking, whereas smaller structures like homes are more likely to be damaged by high-frequency shaking, which ranges from 2 to 10 hertz and above.

High-frequency simulations are more computationally complex, however, limiting the information that engineers have for building safer structures that are sensitive to these waves. Jordan's team is attempting to bridge this gap.

"We're in the process of trying to bootstrap our way to higher frequencies," Jordan said.

Let's Get Physical

The process that Jordan's team follows begins with historical earthquakes.

"Seismology has this significant advantage of having well-recorded earthquake events that we can compare our simulations against," said Philip Maechling, team member and computer scientist at USC. "We develop the physics codes and the 3-D models, then we test them by running a simulation of a well-observed historic earthquake. We compare the simulated ground motions that we calculate against what was actually recorded. If they match, we can conclude that the simulations are behaving correctly."

The team then simulates scenario earthquakes, individual quakes that have not occurred but that are cause for concern. Because seismologists cannot get enough information from scenario earthquakes for long-term statements, they then simulate all possible earthquakes by running ensembles, a suite of simulations that differ slightly from one another.

"They're the same earthquake with the same magnitude, but the rupture characteristics—where it started and how it propagated, for example—will change the areas at the Earth's surface that are affected by this strong ground motion," Maechling said.

As the team increased the maximum frequency in historic [earthquake simulations](#), however, they identified a threshold right around 1 hertz, at which their simulations diverged from observations. The team determined it needed to integrate more advanced physics into its code for more realistic results.

"One of the simplifications we use in low-frequency simulations is a flat simulation region," Maechling said. "We assume that the Earth is like a rectangular box. I don't know if you've been to California, but it's not flat. There are a lot of hills. This kind of simplifying assumption worked well at low frequencies, but to improve these simulations and their results, we had to add new complexities, like topography. We had to add mountains into our simulation."

Including topography—the roughness of the earth's surface—the team's simulations now include additional geometrical and attenuation (gradual dampening of the shaking due to loss of energy) effects—near-fault plasticity, frequency-dependent attenuation, small-scale near-surface heterogeneities, near-surface nonlinearity, and fault roughness.

On Titan, the team introduced and tested the new physics calculations individually to isolate their effects. By the end of 2014, the team updated the physics in its code to get a complete, realistic simulation capability that is now able to perform simulations using Earth models near 4 hertz.

"The kind of analysis we're doing has been done in the past, but it was using completely empirical techniques—looking at data and trying to

map observations onto new situations," Jordan said. "What we're doing is developing a physics-based seismic hazard analysis, where we get tremendous gains by incorporating the laws of physics, to predict what will be in the future. This was impossible without high-performance computing. We are at a point now where computers can do these calculations using physics and improve our ability to do the type of analysis necessary to create a safe environment for society."

Movers and Shakers

With the new physics included in SCEC's earthquake code—the Anelastic Wave Propagation by Olsen, Day, and Cui (AWP-ODC)—Jordan's team was able to run its first CyberShake hazard curve on Titan for one site at 1 hertz, establishing the computational technique in preparation for a full-fledged CyberShake map.

A seismic hazard curve provides all the probabilities that an earthquake will occur at a specific site, within a given time frame, and with ground shaking exceeding a given threshold.

The team used the US Geologic Survey's (USGS's) Uniform California Earthquake Forecast—which identifies all possible earthquake ruptures for a particular site—for generating CyberShake hazard curves for 336 sites across southern California.

This May, the team calculated hazard curves for all 336 sites needed to complete the first 1 hertz urban seismic hazard map for Los Angeles. With double the maximum simulated frequency from last year's 0.5 hertz map, this map proves to be twice as accurate.

The map will be registered into the USGS Urban Seismic Hazard Map project, and when it passes the appropriate scientific and technical review, its results will be submitted for use in the 2020 update of the

Recommended Seismic Provisions of the National Earthquake Hazards Reduction Program.

This major milestone in [seismic hazard](#) analysis was possible only with the help of Titan and its GPUs.

"Titan gives us the ability to submit jobs onto many GPU-accelerated nodes at once," Jordan said. "There's nothing comparable. Even with other GPU systems, we can't get our jobs through the GPU queue fast enough to keep our research group busy. Titan is absolutely the best choice for running our GPU jobs."

Yifeng Cui, team member and computational scientist at the San Diego Supercomputer Center, modified AWP-ODC to take advantage of Titan's hybrid architecture, thereby improving performance and speed-up. He was awarded NVIDIA's 2015 Global Impact Award for his work.

"It's fantastic computer science," Jordan said. "What Yifeng has done is get in and really use the structure of Titan in an appropriate way to speed up what are very complicated codes. We have to manipulate a lot of variables at each point within these very large grids and there's a lot of internal communication that's required to do the calculations."

Using Cui's GPU-accelerated code on Titan, the team ran simulations 6.3 times more efficiently than the CPU-only implementation, saving them 2 million core hours for the project. Completion of the project required about 9.6 million core hours on Titan.

"The computational time required to do high-frequency simulations takes many node hours," Maechling said. "It could easily take hundreds of thousands of node hours. That's a huge computational amount that well exceeds what SCEC has available at our university. These pushing-to-higher-frequency [earthquake](#) simulations require very large computers

because the simulations are computationally expensive. We really wouldn't be able to do these high-frequency simulations without a computer like Titan."

With Titan, Jordan's team plans to push the maximum simulated frequency above 10 hertz to better inform engineers and emergency officials about potential seismic events, including the inevitable "big one."

"We have the potential to have a positive impact and to help reduce the risks from earthquakes," Maechling said. "We can help society better understand earthquakes and what hazards they present. We have the potential to make a broad social impact through safer environment."

Provided by Oak Ridge National Laboratory

Citation: Titan takes on the big one (2015, November 11) retrieved 20 April 2024 from <https://phys.org/news/2015-11-titan-big.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.