CyberShake Physics-Based Seismic Hazard Models for Northern California

Scott Callaghan and the CyberShake Collaboration

SC19
Denver, CO
November 17-22

Probabilistic Seismic Hazard Analysis (PSHA)

- What will peak earthquake shaking be over the next 50 years?
- Useful information for:
  - Building engineers
  - Disaster planners
  - Insurance agencies

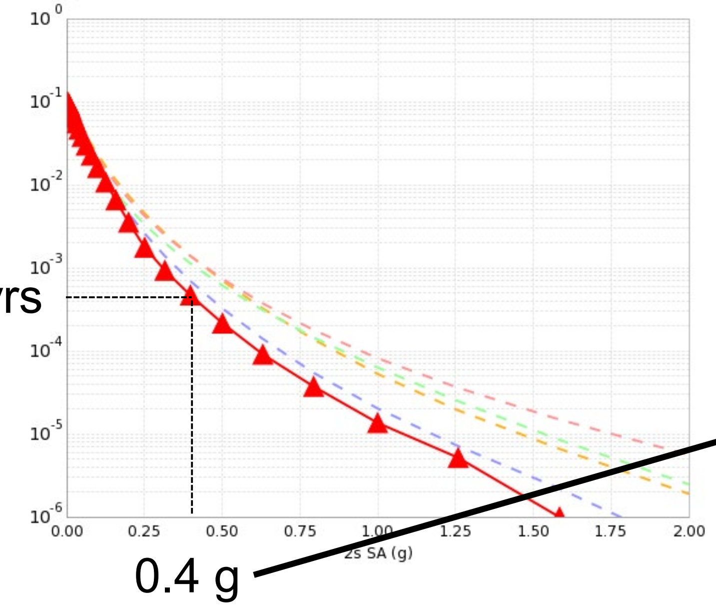

- PSHA performed by
  1. Assembling a list of earthquakes
  2. Determining how much shaking each event causes
  3. Combining the shaking levels with probabilities

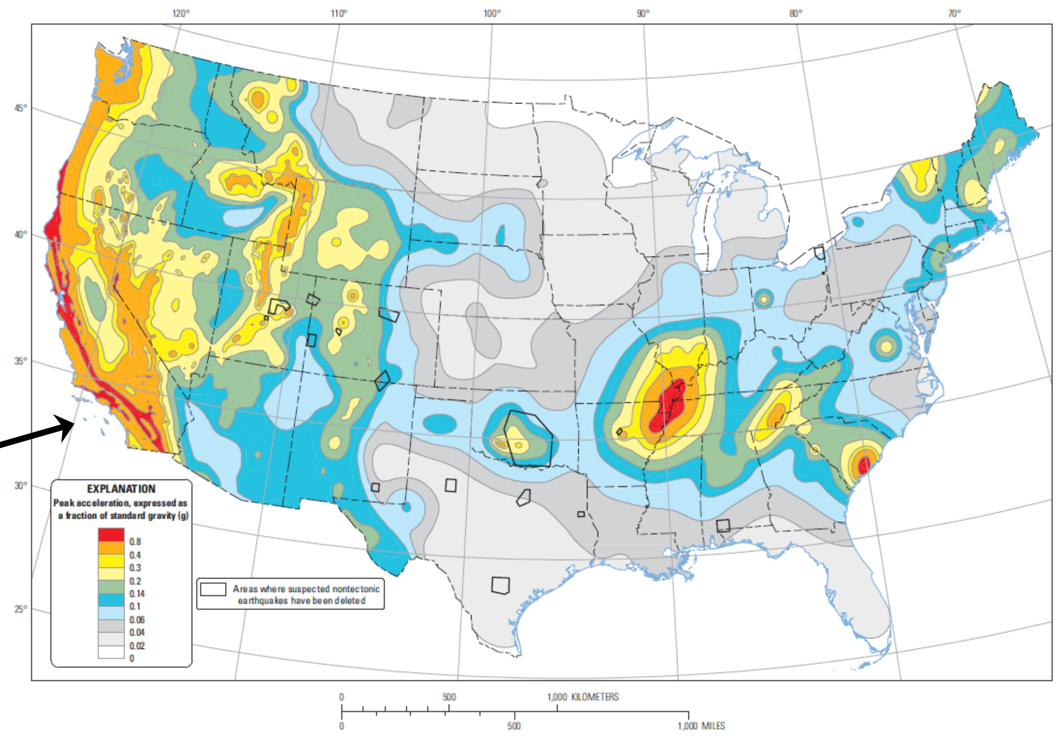

Two-percent probability of exceedance in 50 years map of peak ground acceleration

PSHA Approaches

- Ground Motion Prediction Equations (GMPEs)
  - Equations derived from observed data
  - Mean and standard deviation of ground motion produced from each earthquake
  - Computationally cheap
  - Statistical approach can yield unphysical results

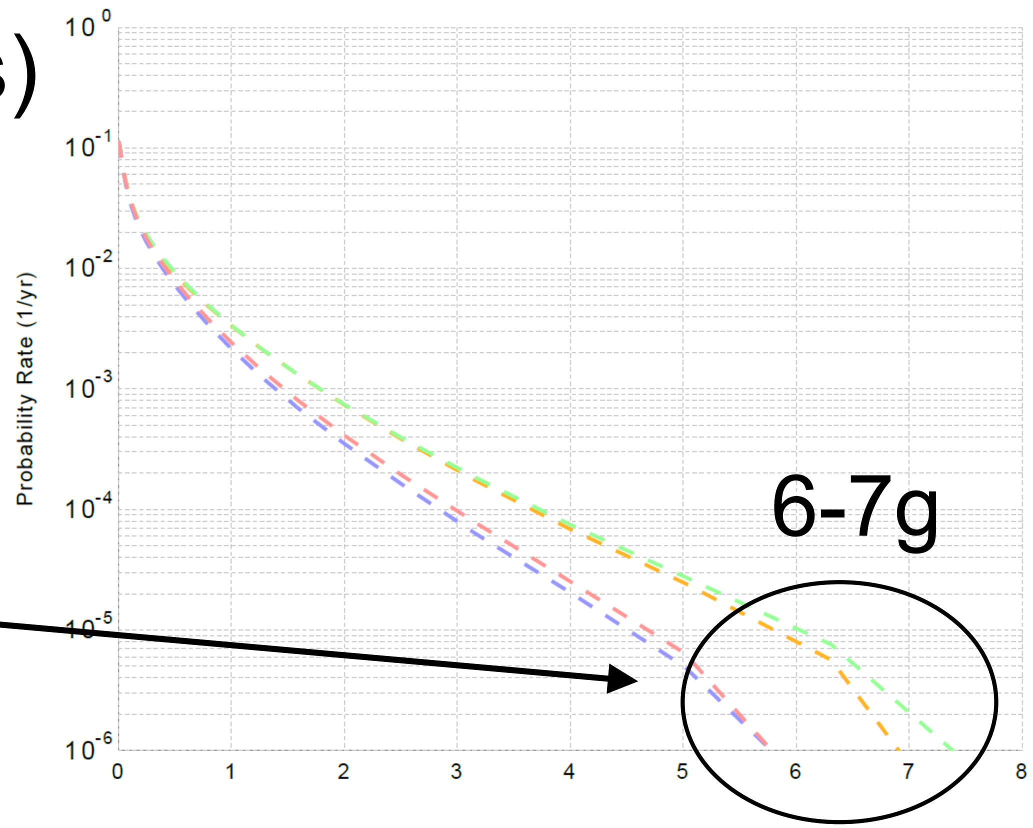

- Simulation-based approach
  - Each earthquake is simulated using wave propagation
  - Can reduce uncertainty by capturing complex physics
  - Computationally expensive

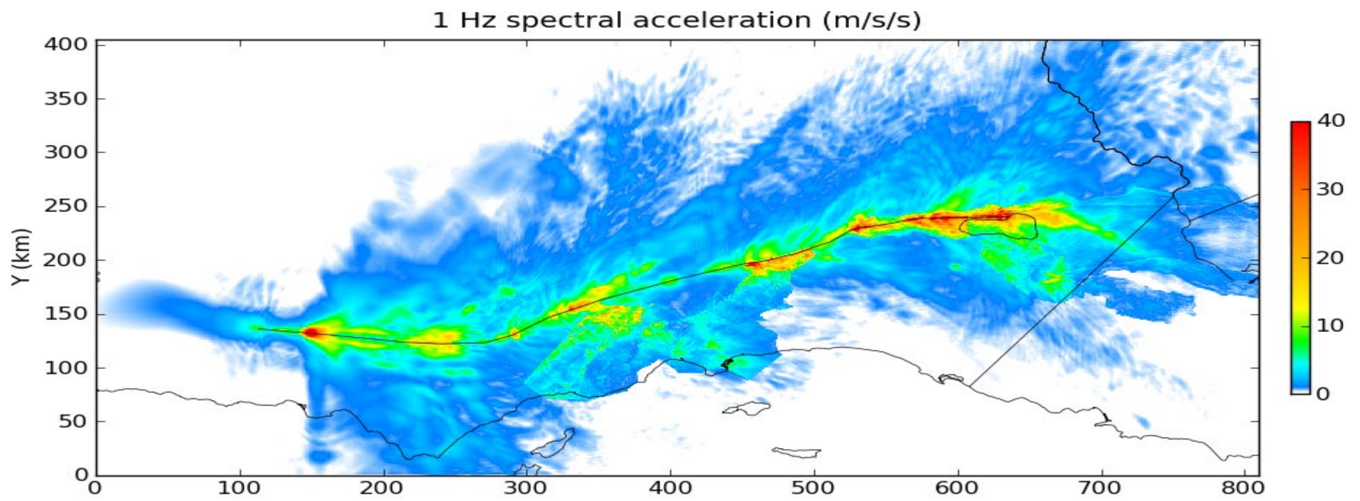

SCEC CyberShake Project

- 3D physics-based platform for PSHA
- For each site of interest:
  - Determine nearby (<200 km) earthquakes
  - Add variability to earthquakes
  - Simulate each of 500,000 earthquakes
  - Determine maximum shaking from each
  - Combine with probabilities to produce curve
- Repeat process for multiple locations
- Continual improvement since 2007

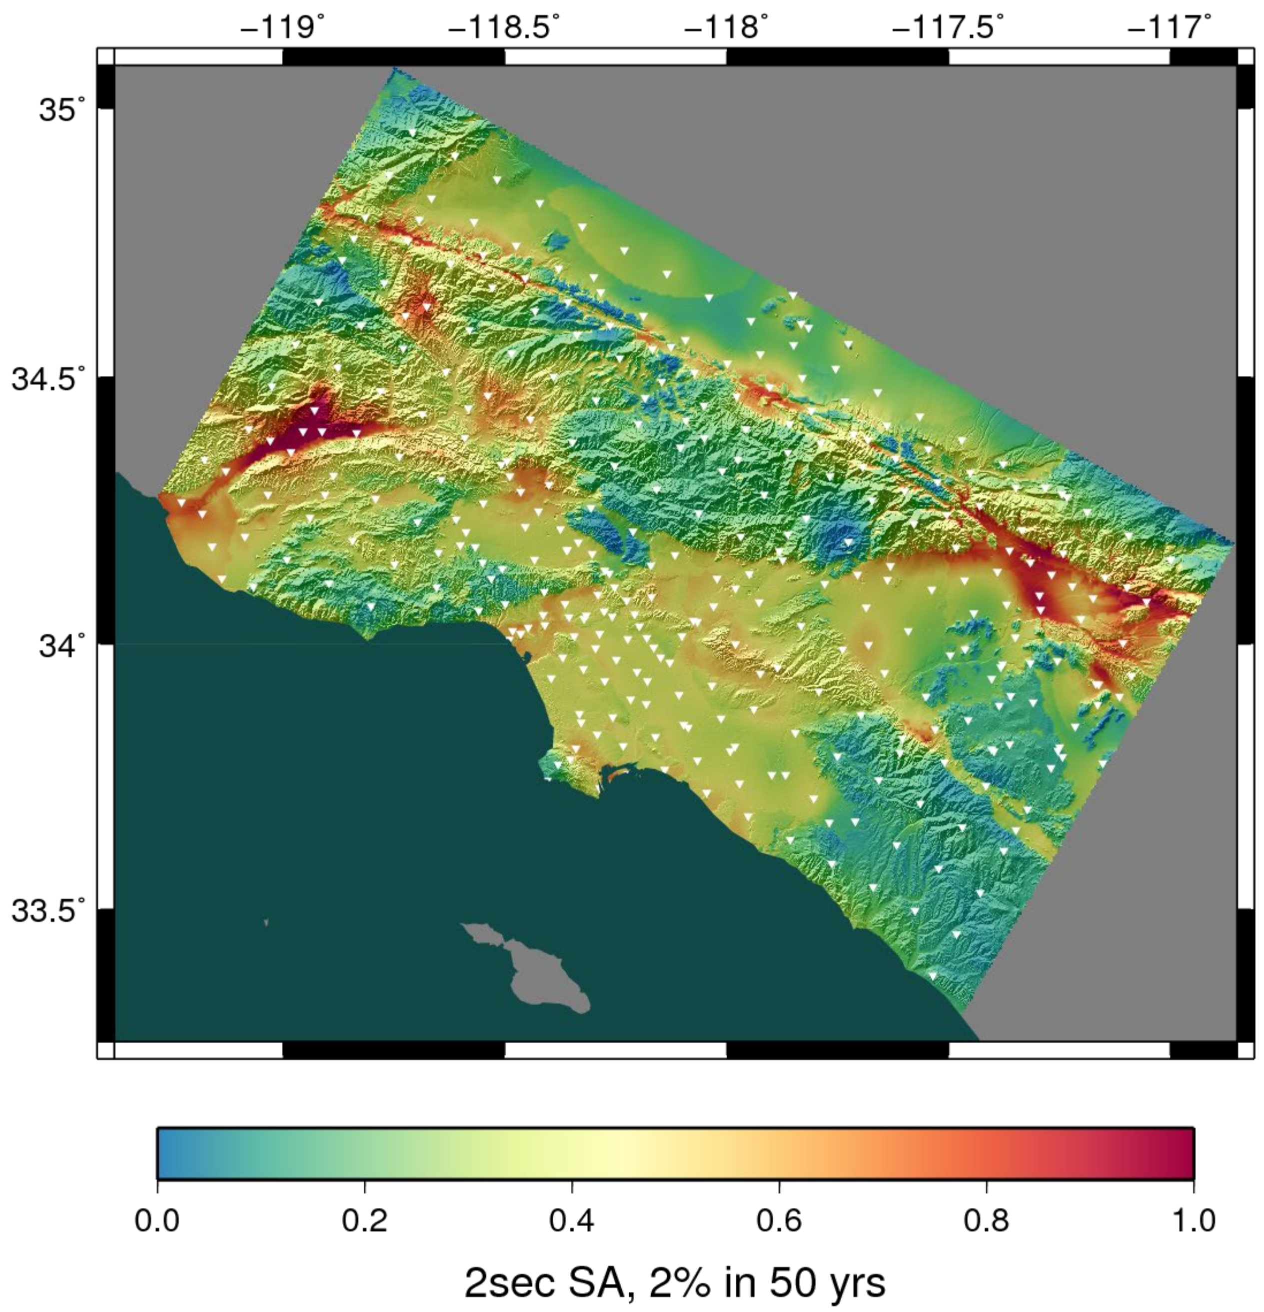

Northern California CyberShake: Study 18.8

- Build upon previous CyberShake results
- 800+ new locations
- New velocity model of earth's crust
- Statewide simulation volumes
- Largest suite of CyberShake simulations to date

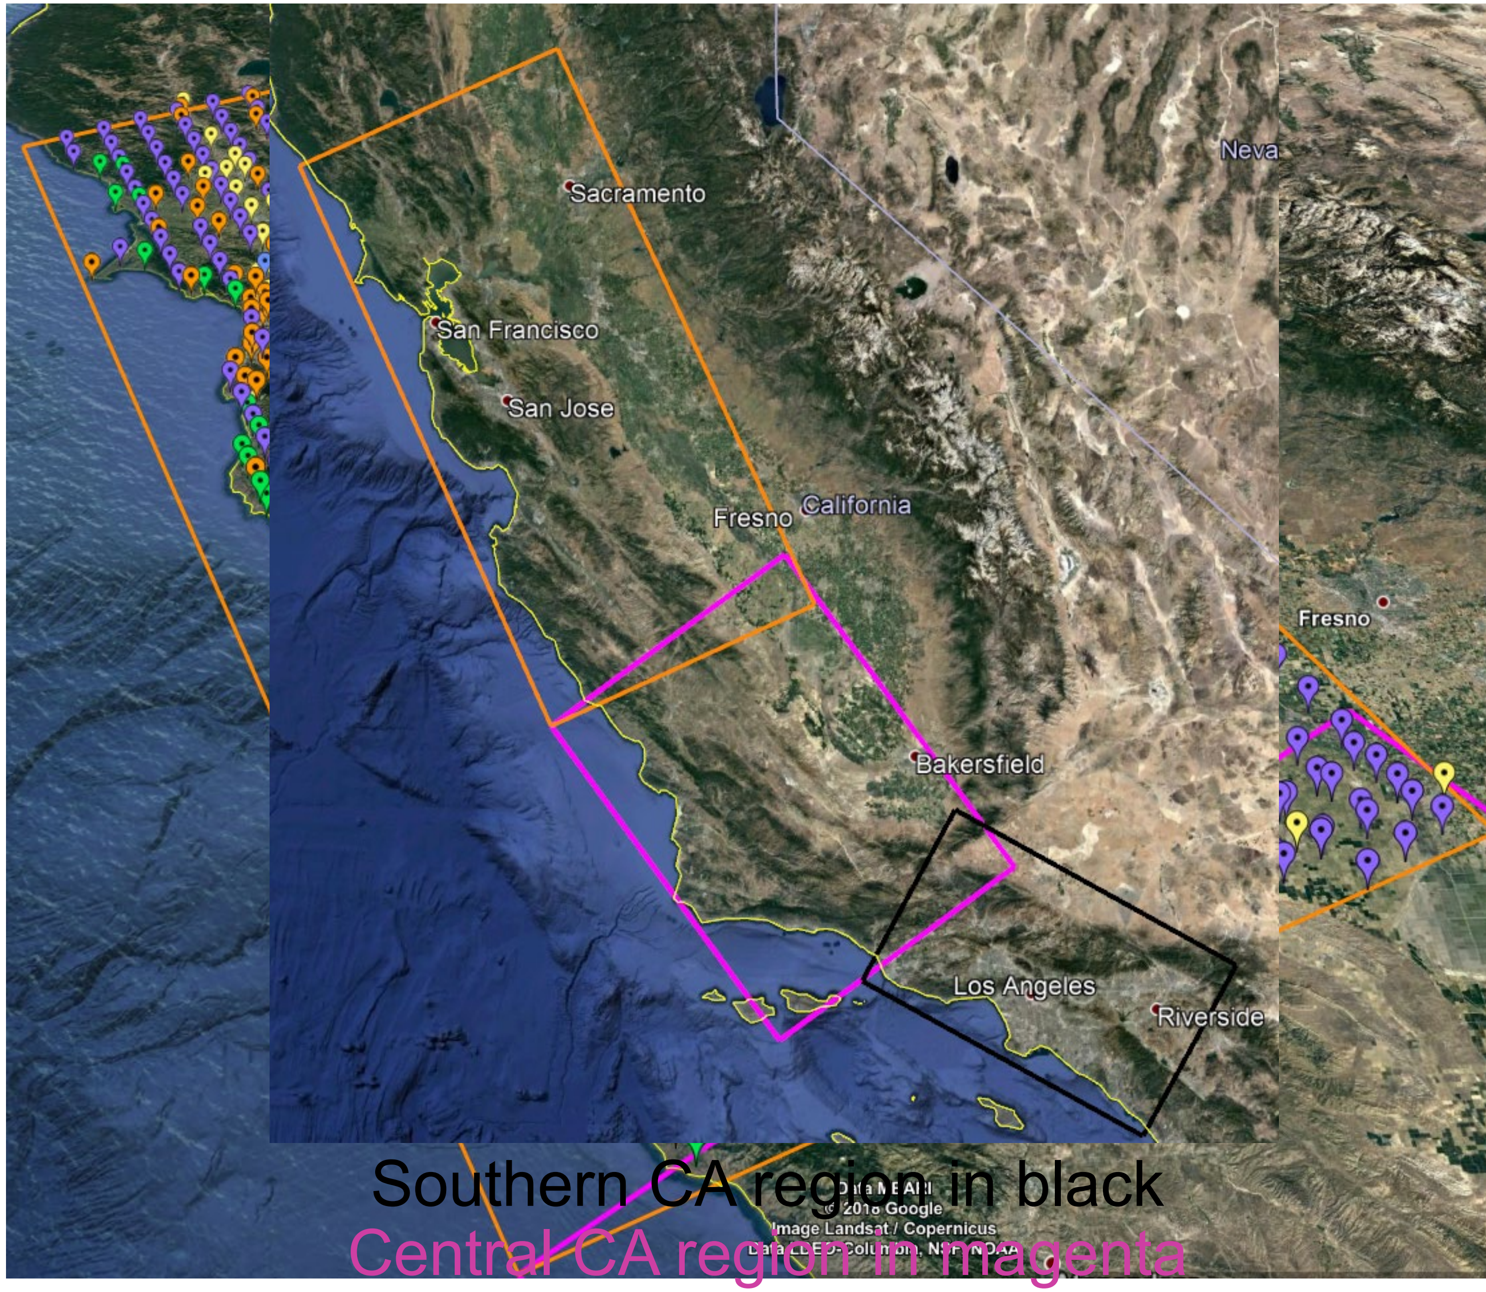

Southern CA region in black
Central CA region in magenta
Bay Area region in orange
869 sites, densest near San Francisco Bay

Combined Velocity Model

- No single 3D model large enough for large volumes
- Stitch together multiple models
  - Central California (blue)
  - USGS Bay Area (green)
  - Southern California (red)
  - 1D background model (white)
- Apply smoothing along model interfaces
  - Average of neighbor values

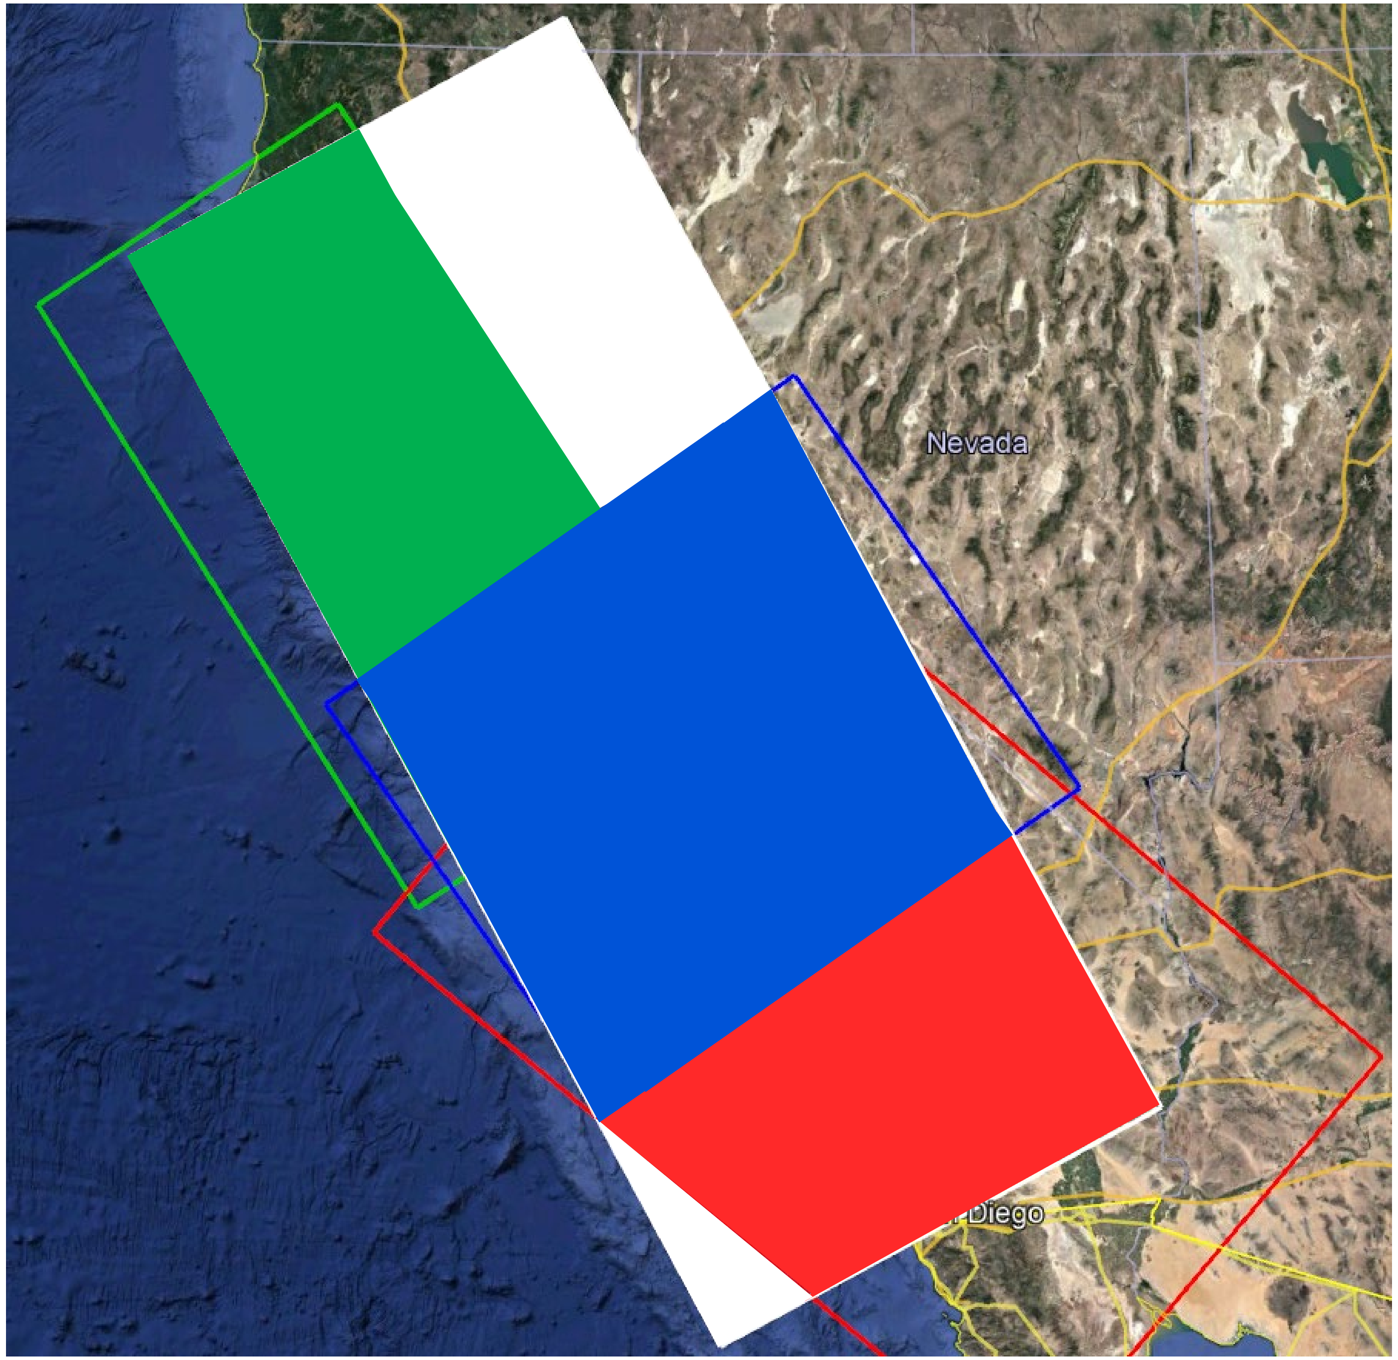

Study 18.8 Computational Requirements

CyberShake stage	Compute-hours	Output data
Velocity mesh creation	3,700 CPU	177 GB
Wave propagation (“SGT”) simulations	2,500 GPU	345 GB
Seismogram synthesis	61,000 CPU	17 GB
Total, 1 site	64,700 CPU; 2,500 GPU	539 GB
Total, Northern California Study	**56 million CPU, 2.1 million GPU**	**457 TB**

- Very large computational requirements
- Targeted NCSA *Blue Waters* and OLCF *Titan* supercomputers
- High degree of automation required for around-the-clock execution
  - Rely heavily on scientific workflow tools
  - Workflows orchestrated from USC

Scientific Workflow Tools

- Pegasus-WMS
  - Use API to create description of workflow
    - Tasks with dependencies
    - Input/output files
  - Plans workflow for execution on specified systems
  - Adds jobs to manage data
  - Wraps executables to track metadata
- HTCondor
  - Manages real-time execution of jobs
  - Submits jobs to remote systems, checks on success
  - Monitors dependencies
  - Checkpoints workflow
- GridFTP used to transfer data

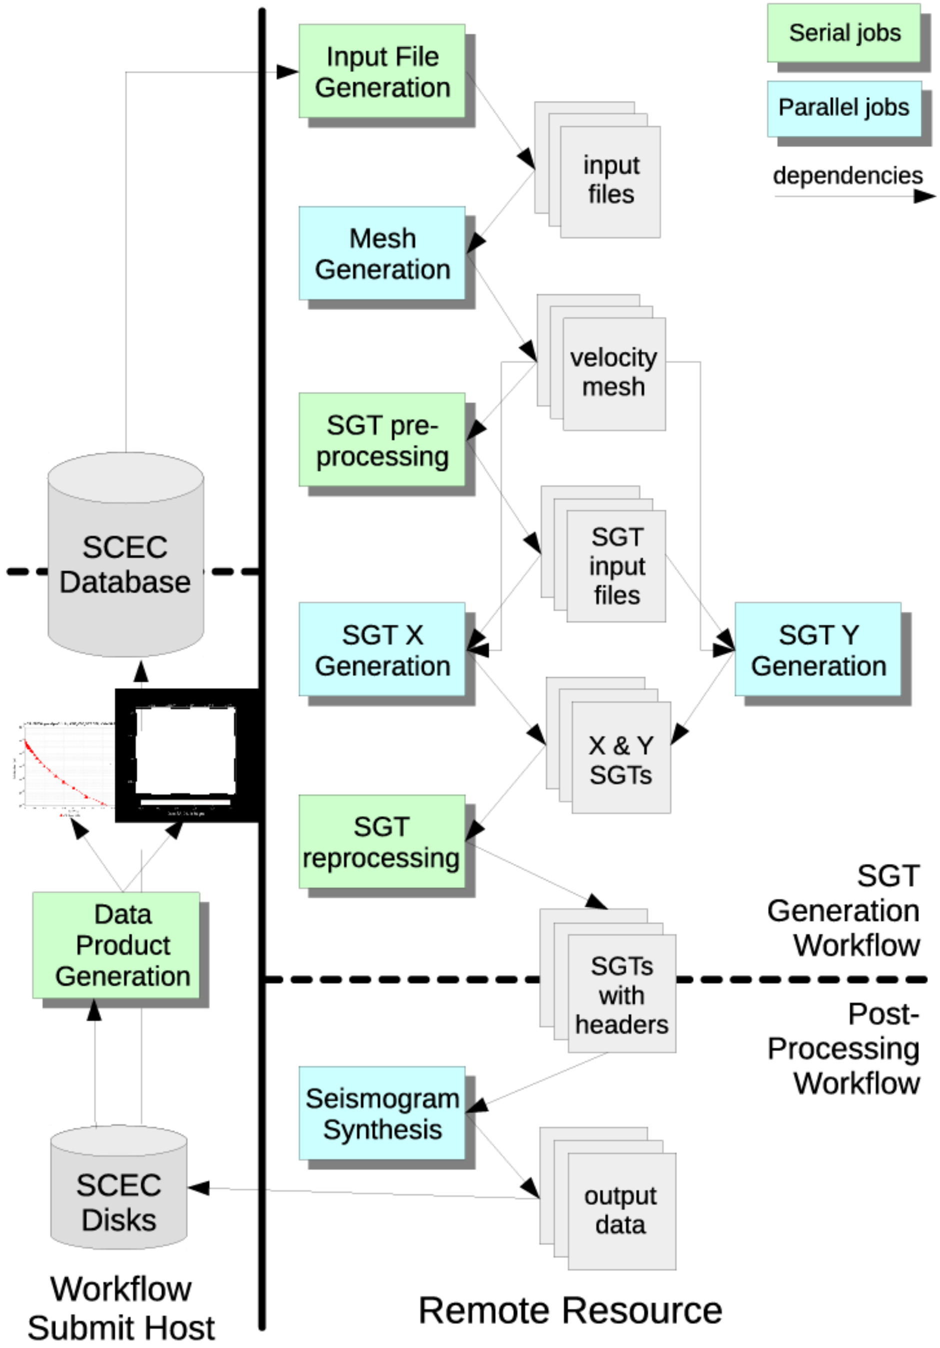

Schematic of CyberShake workflow

Automated Remote Job Submission

- Push-based
  - When jobs are ready to run, send them over the network to wait in queue
  - SSH: keys must be accepted on remote system
  - Globus GRAM: protocol for job submission, requires support on remote system
  - rvGAHP: daemon on remote system connects to workflow submit host
    - Can be used on systems with two-factor authentication
- Pull-based ("pilot jobs")
  - Submit job on remote system first
  - After job starts up, advertises to workflow submit host
  - Results in additional overhead
  - Can take advantage of scheduling policies

Bin	Node Range	Aging Boost
1	11250+	15 days
2	3750 – 11249	5 days
3	313 – 3749	0
4	126 – 312	0
5	1 – 125	0

OLCF *Titan* Scheduling Policy

Dynamic Workflow Assignment

- To accomplish CyberShake study efficiently, must be able to use resources when available
- Job throughput on large clusters varies widely
- Designed workflow metascheduler to submit workflows
  - Split workflows into SGT and post-processing
  - Ability to run each part on separate systems

	BW SGTs	Titan SGTs	Total
BW PP	444	290	734
Titan PP	0	135	135
Total	444	425	869

Systems used for SGT and post-processing workflows

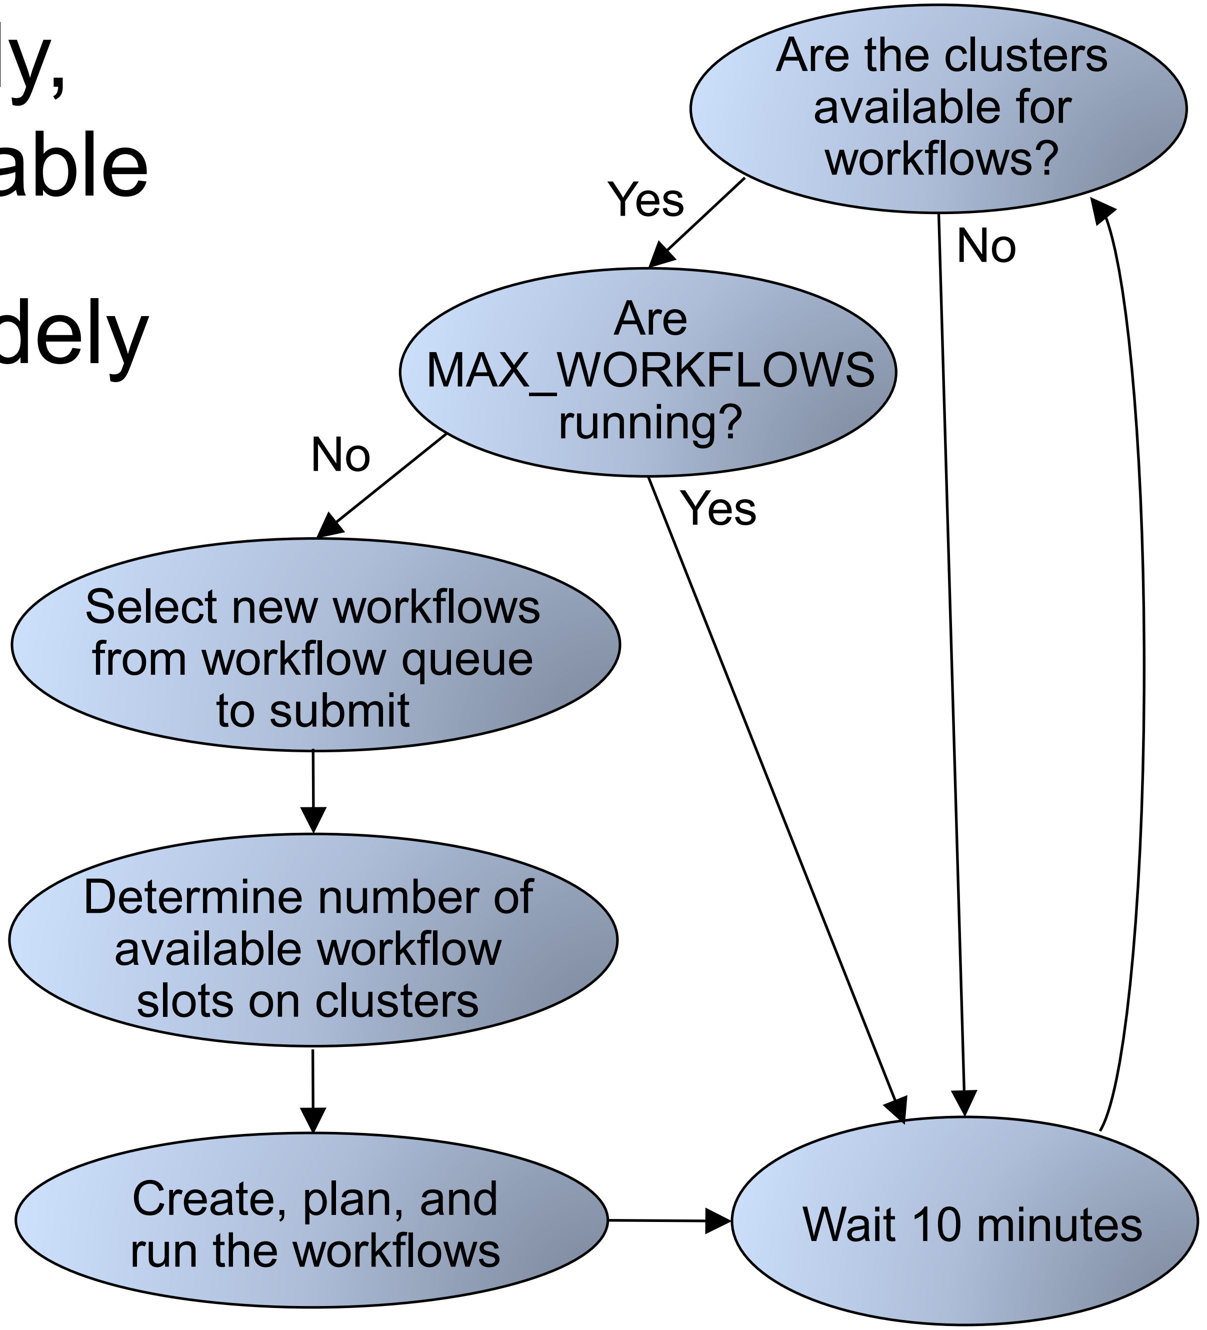

CyberShake Study 18.8 Metrics

- Study conducted over 128 days
- Consumed 6.2 million node-hours (120M core-hours/13,650 core-years)
  - Averaged 2,018 nodes / 38,850 cores
  - Max of 16,219 nodes / 279,984 cores
- Ran 21,220 jobs at USC, 10,308 at Blue Waters, and 7,757 jobs at Titan
- 1.2 PB of data generated
  - 157 TB of data automatically transferred
  - 14.4 TB of final data products staged to USC HPC
- Simulated 203 million seismograms
  - 30.4 billion shaking values

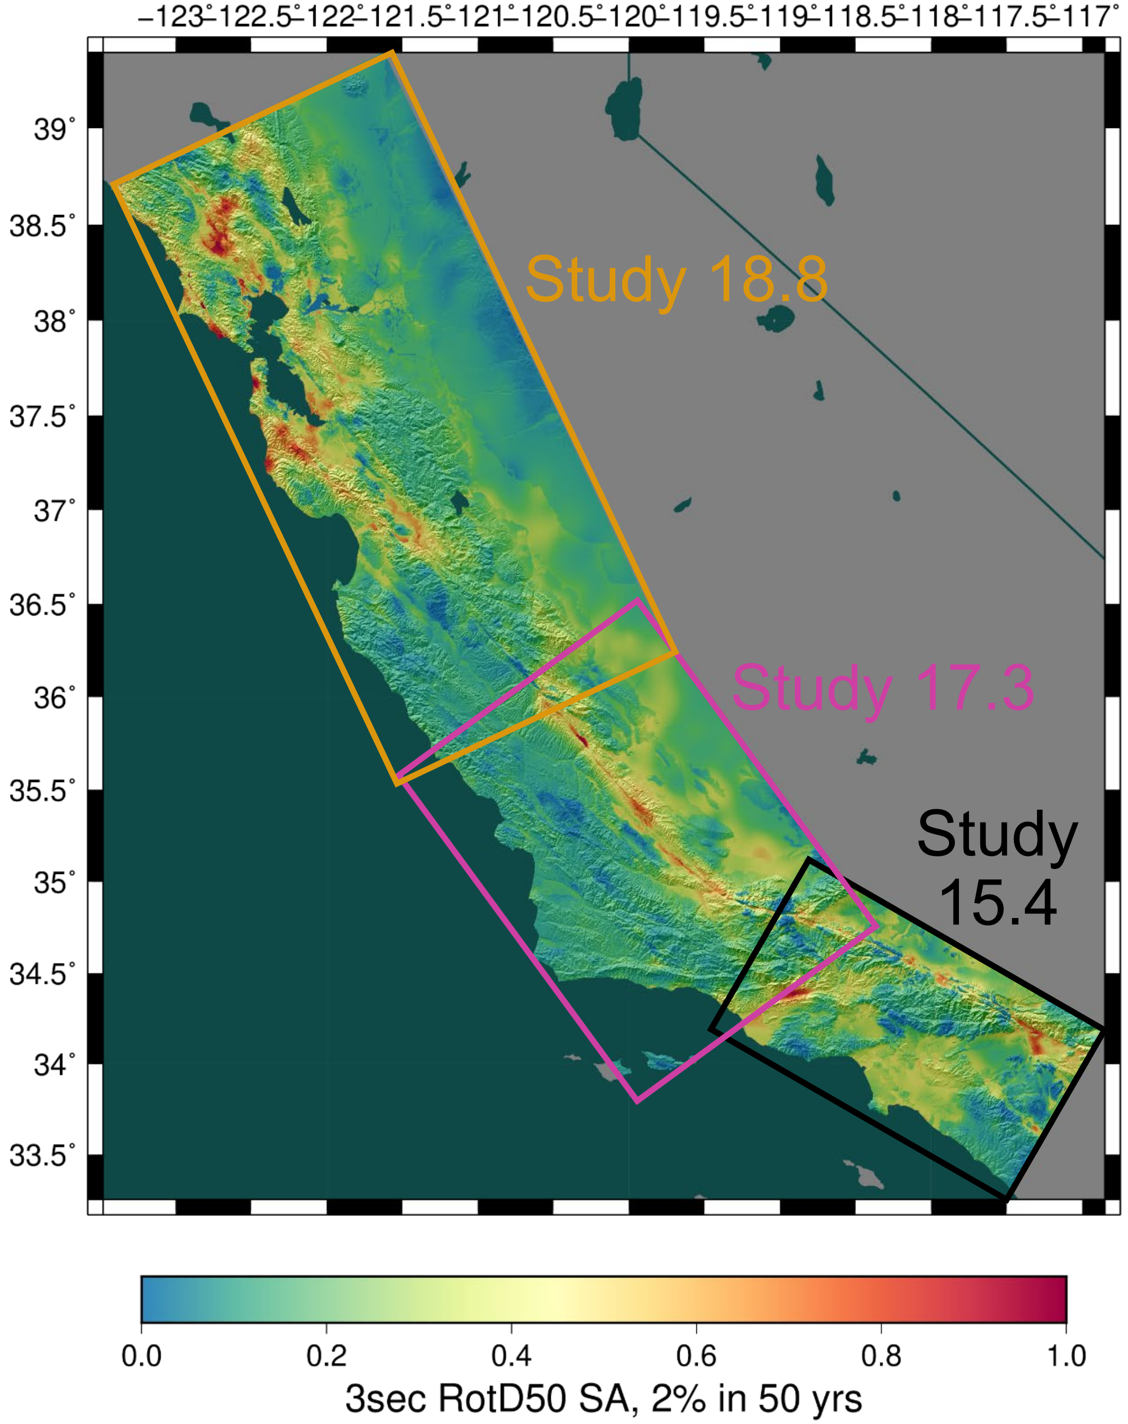

Study 18.8 Results

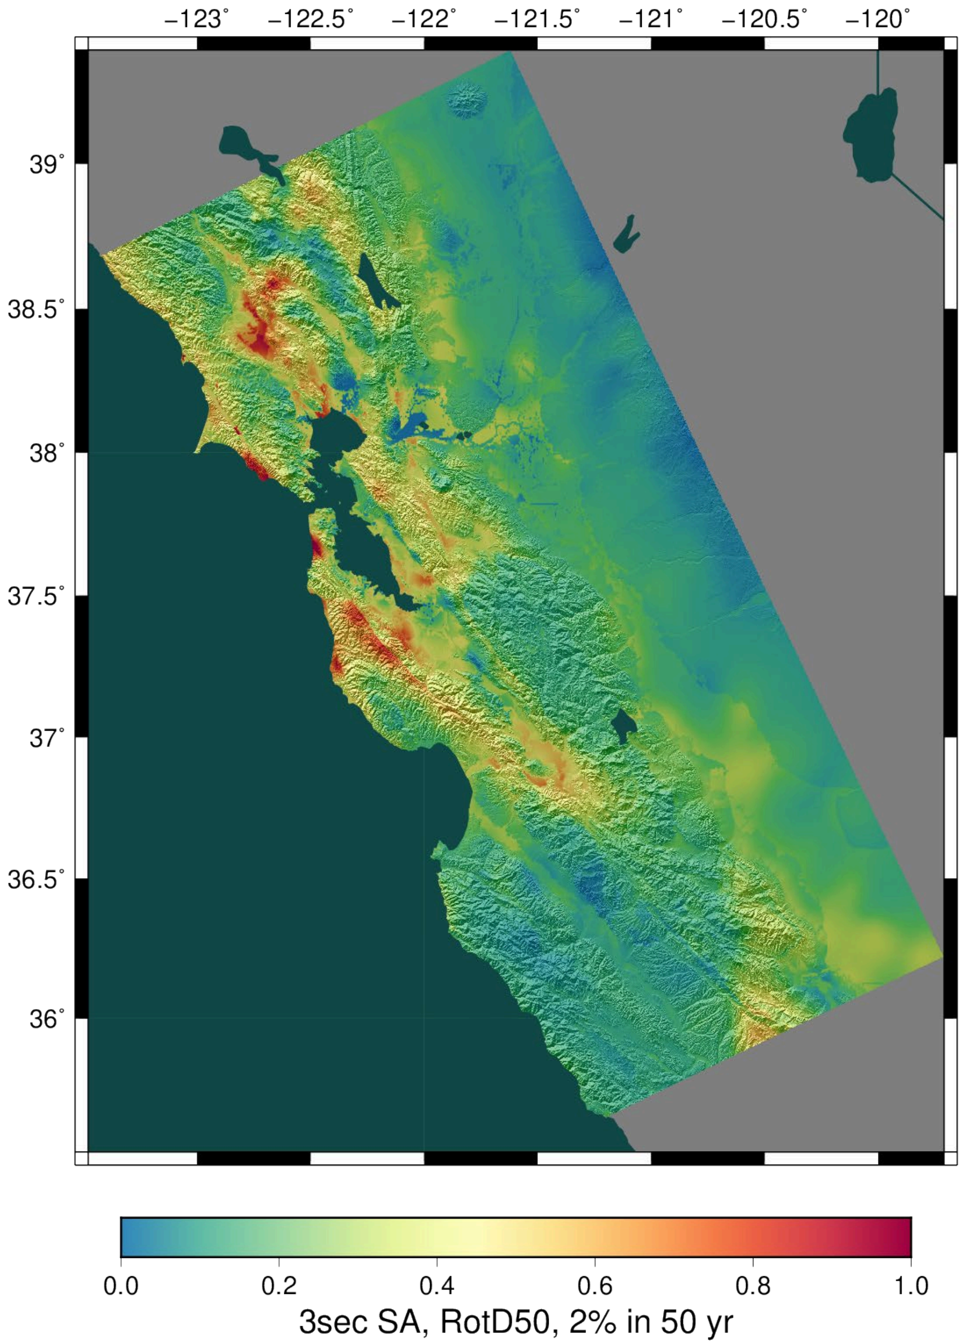

CyberShake

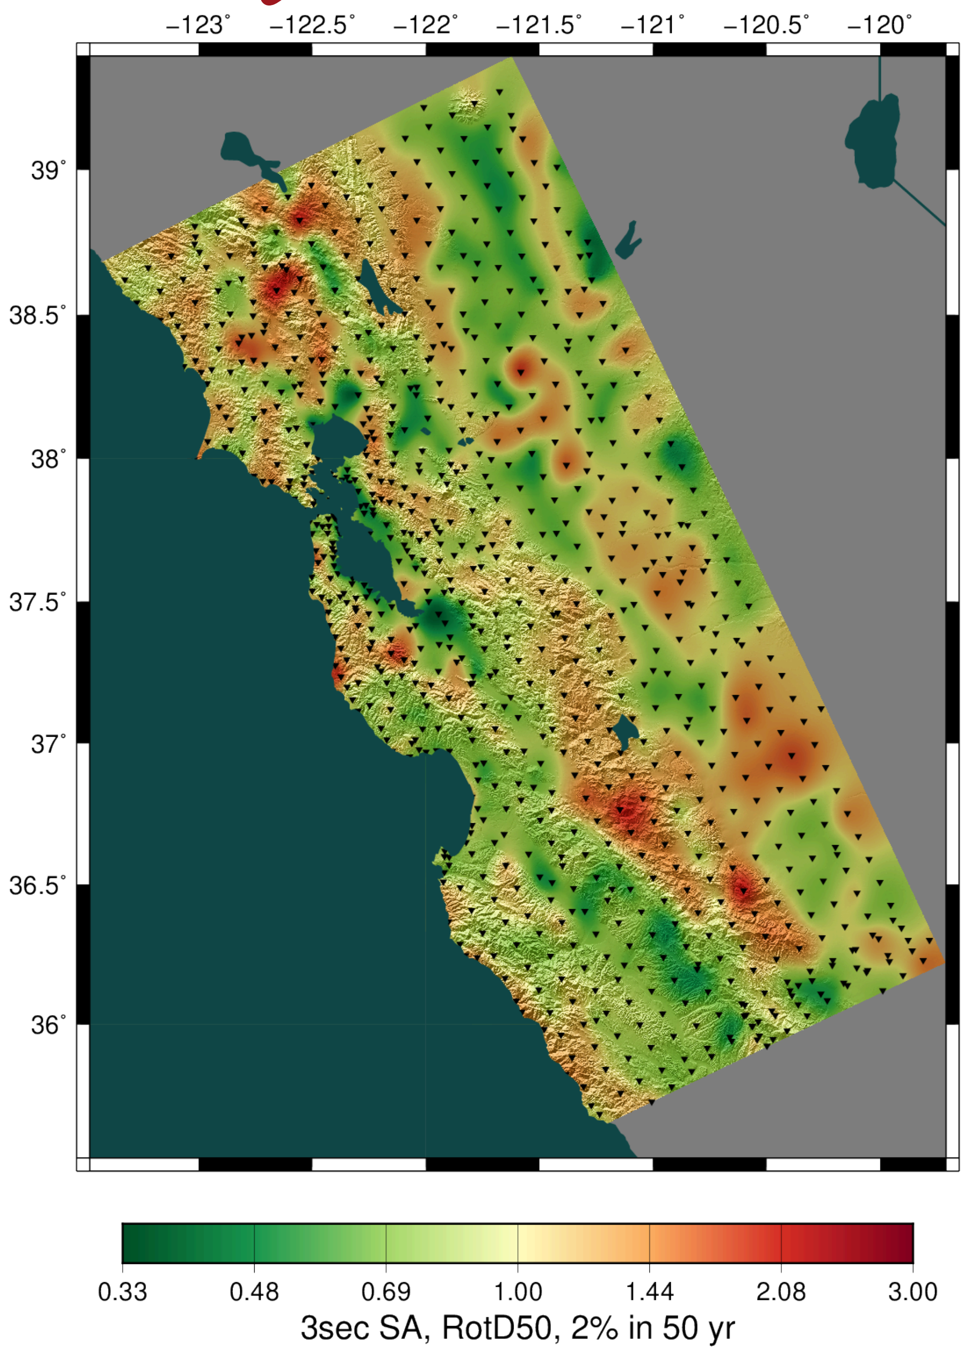

Ratio of CyberShake/GMPEs

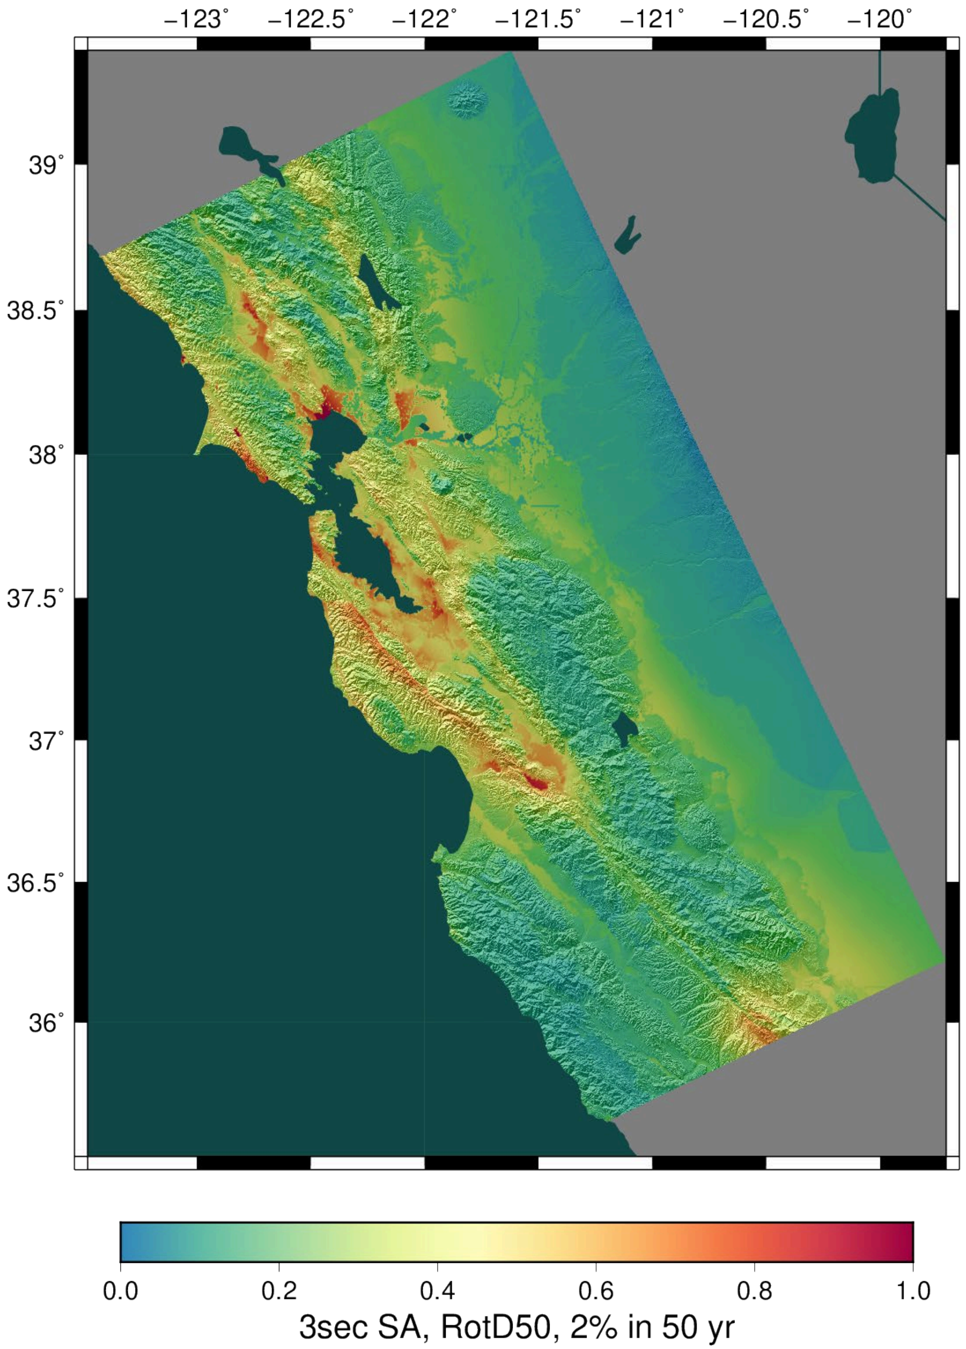

Average of 4 GMPEs

Future Directions

- Moving to next-generation systems
  - OLCF *Summit:* CyberShake fully verified
  - TACC *Frontera*: Verification underway
- Enhancing and improving physics in CyberShake simulations
  - Updated rupture generation code with better observational agreement
  - Discontinuous mesh version of wave propagation code
  - Topography
  - Higher frequencies
- Performance optimizations
  - Machine learning to eliminate some events

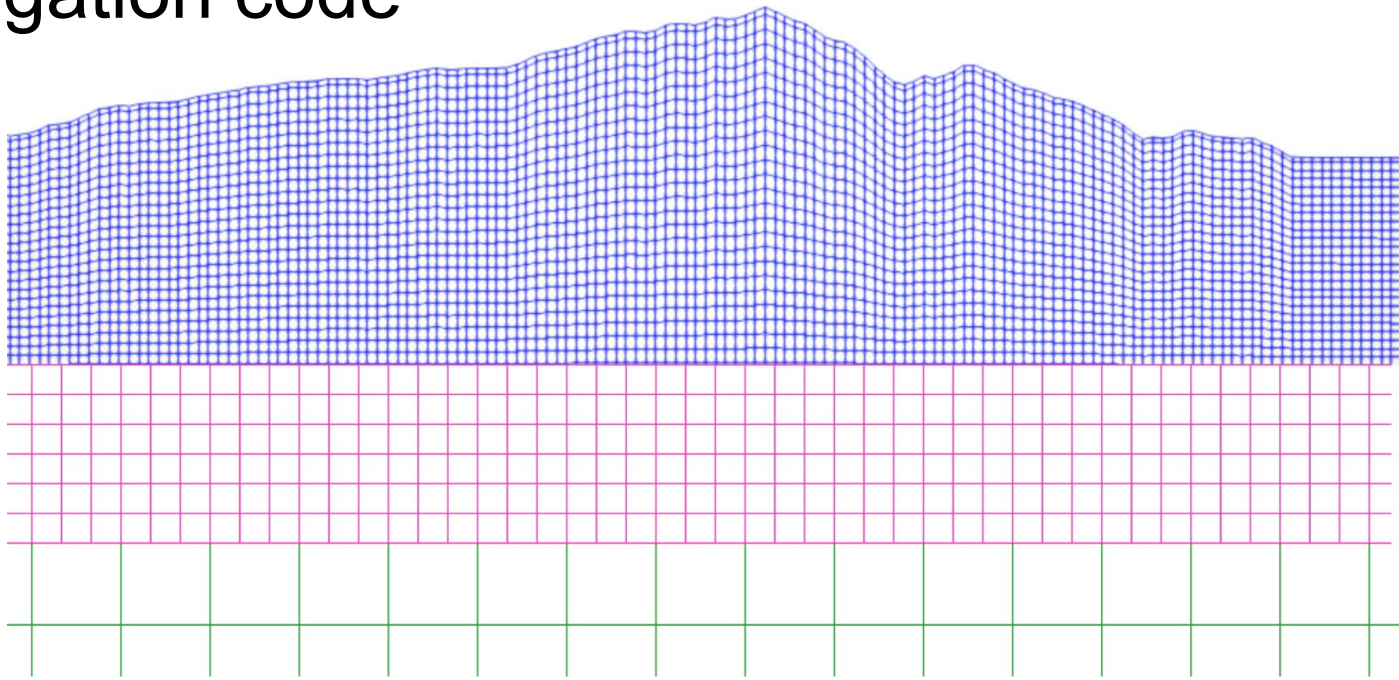

Questions?