**Title**

Transformational Regional-Scale Earthquake Simulations with the DOE EarthQuake SIMulation Exascale Framework

**Journal**

Computing in Science & Engineering, 26(2)

**ISSN**

1521-9615

**Authors**

McCallen, David
Pitarka, Arben
Tang, Houjun
et al.

**Publication Date**

2024-04-01

**DOI**

10.1109/mcse.2024.3397768

# Transformational Regional-Scale Earthquake Simulations with the DOE EarthQuake SIMulation Exascale Framework

David McCallen, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*
Arben Pitarka, *Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA*
Houjun Tang, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*
Ramesh Pankajakshan, *Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA*
Anders Petersson, *Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA*
Mamun Miah, *Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA*

*Earthquakes present worldwide risk to economic and human safety. The 2023 earthquakes in Türkiye provided a reminder of the potential for catastrophic consequences with 50,700 deaths and 15.7 million people affected. The ability to predict ground motions and infrastructure damage for earthquakes continues to be a challenging problem for scientists and engineers. Until now, estimates of ground motions have been performed empirically by looking at sparse data from past earthquakes. This approach can provide statistical information on intensity amplitudes but cannot inform site-specific ground motions essential to developing the most effective resilience. Interest has grown in large-scale computational models to simulate earthquakes at regional scale. The U.S. Department of Energy EarthQuake SIMulation (EQSIM) framework was developed for regional-scale earthquake simulations at unprecedented fidelity, taking advantage of emerging GPU-accelerated systems. This article describes the EQSIM workflow and demonstrates regional-scale simulations with the new computational capability available to scientists in their quest to mitigate future disasters.*

The existing state of practice for estimating the ground motions during future damaging earthquakes has been one of looking to the past to develop estimates for the future. This practice consists of a statistical analysis of previous earthquake ground motion observations to create an empirically based statistical model of ground motions as a function of earthquake magnitude and geological parameters. One of the principal challenges of such an approach, which introduces significant uncertainty in ground motion estimates, is the major limitations in the historical observational database of earthquake motions. The information in Figure 1 is representative of the data limitation problem. This figure, generated from worldwide measured earthquake ground motion data available from the Pacific Earthquake Engineering Research Center (PEER) at the University of California Berkeley, illustrates the ground motion data currently available from throughout the world where the left figure shows the various measurement site locations. The figure on the right shows each ground motion measurement in the existing database, shown as red and blue dots, plotted as a function of the magnitude of the earthquake generating the ground motion (vertical axis) versus the distance of the site at which the ground motion was

---

Digital Object Identifier 10.1109/MCSE.2024.3397768
Date of publication 8 May 2024; date of current version 30 September 2024.

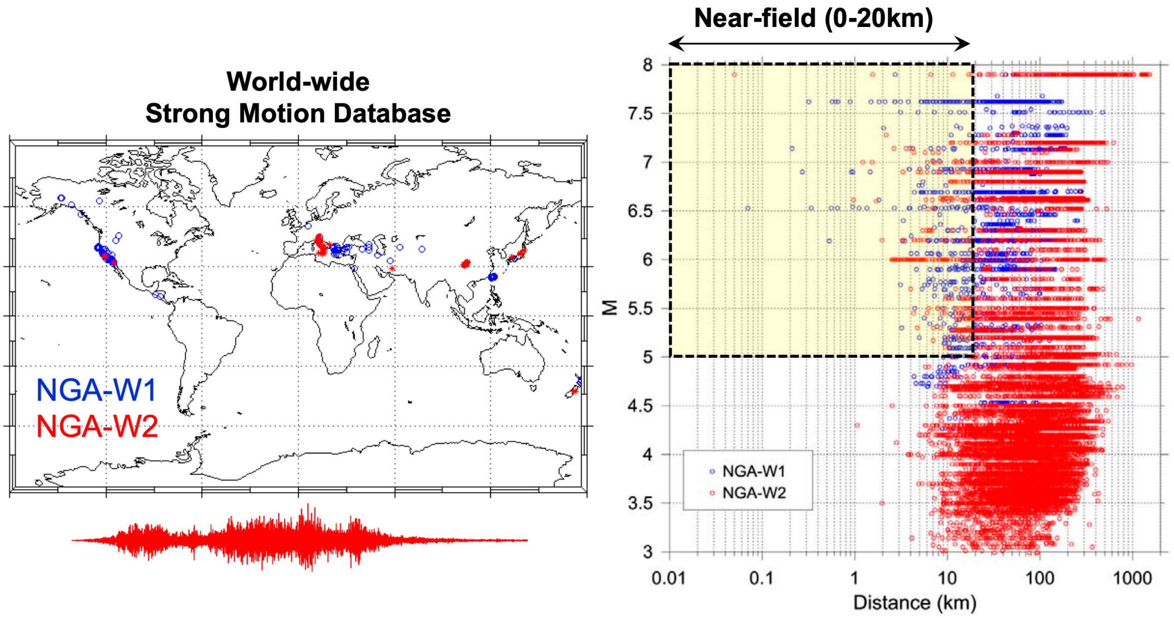

**FIGURE 1.** PEER worldwide database of historical earthquake records showing sparse ground motion data in the near field of large magnitude earthquakes (yellow box), courtesy of PEER.

measured from the causative earthquake fault (horizontal axis). Inspection of this figure illustrates a major sparse-data problem in the region of biggest concern, specifically for sites located near to a fault (i.e., within 20 km of the fault) on which a large earthquake occurs, which is indicated by the yellow box in the figure.

Given the uncertainties in ground motion estimates created in part by this lack of comprehensive data, there has been significant interest in developing alternative approaches that can provide enhanced insight into the regional variability of earthquake ground motions. With increased understanding of the science of underlying earthquake processes, and the inexorable advance of computer performance, there has been major growing interest in developing a physics-based, regional-scale simulation capability to predict both the amplitude and spatial variability of earthquake ground motions. For a particular region of interest, a regional model will include a large 3-D segment of the earth that encompasses the earthquake faults of interest as well as the urban infrastructure domain of interest, as illustrated in Figure 2. Virtual earthquake scenarios can be created through high-performance computer simulations to assess both the distribution of ground motions and the resulting distributions of infrastructure damage.

Noteworthy international developments in physics-based, broad-frequency band simulations include the Spectral Elements in Elastodynamics with Discontinuous Galerkin (SPEED) spectral element code and the associated extensive validation studies through application to many international regions.[1] The Hercules finite-element code represents another major physics-based development that is focused on resolution of broadband ground motions. Hercules has been utilized in Southern California validation studies[2] and recently applied in regional earthquake risk evaluations.[3] An alternative methodology development is represented by the Southern California Earthquake Center (SCEC) Broadband Platform[4] which is based on a hybrid approach of merging low frequency numerical simulations (simulations to 1–2 Hz) combined with stochastic characterizations of high-frequency motions.

The ultimate achievement of the most computationally advanced regional models requires multidisciplinary expertise in the earth sciences, combined with both geotechnical and structural engineering. In addition, due to the fact that such a regional model must capture a large domain on the order of hundreds of kilometers in extent and must resolve higher frequencies that are associated with the vibration characteristics of typical infrastructure, the computational demands can be extreme and advanced computational schemes for parallel platforms must be employed. The EarthQuake SIMulation (EQSIM) framework,[5,6,7] an application development under the U.S. Department of Energy (DOE) Exascale Computing Project (ECP), has focused on creating a rigorous multidisciplinary workflow for fault-to-structure simulations starting from the release

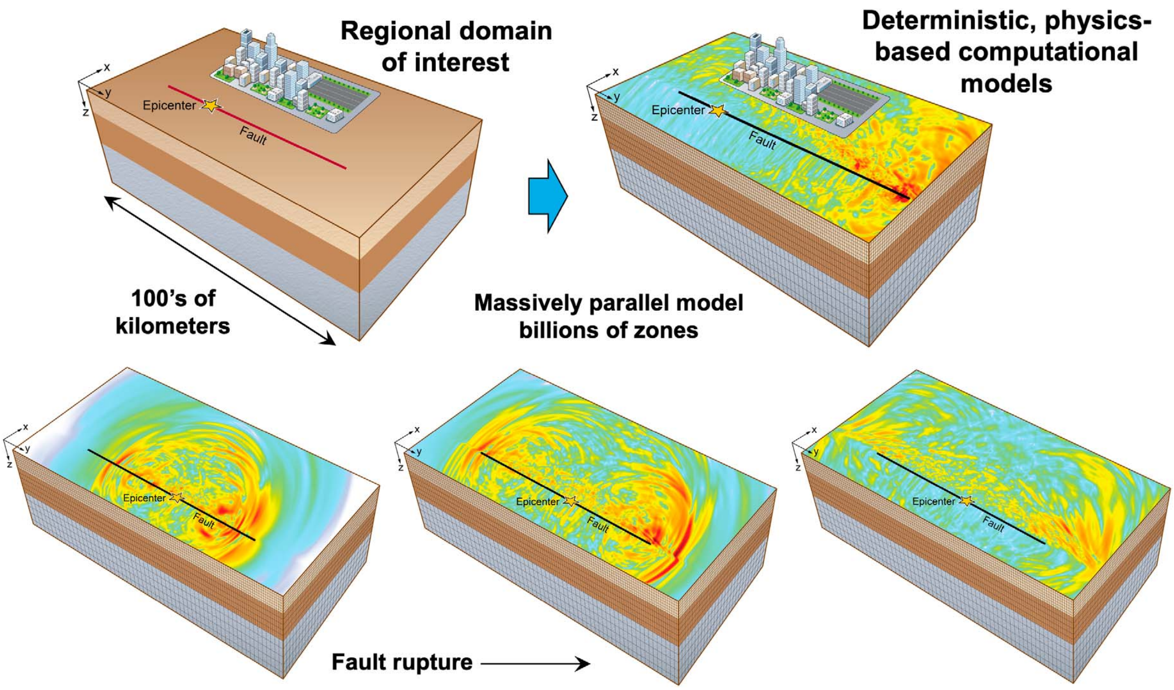

**FIGURE 2.** The concept of large regional-scale earthquake simulations; a model domain that encompasses all relevant earthquake faults, urban areas, and infrastructure of interest.

of energy along the rupture of the causative fault, the subsequent propagation of seismic waves radiating away from the fault, and finally the interaction of the seismic waves with soil/structure systems, as shown schematically in Figure 3.

The EQSIM workflow was purpose-built to overcome existing computational barriers to regional simulations by taking full advantage of DOE's new GPU-accelerated exaflop computer platforms coming online in 2023 and 2024. This includes the Frontier computer at Oak Ridge National Laboratory and the Aurora computer at Argonne National Laboratory. Frontier was completed in early 2023 and has been certified as the first exaflop system at 1.194 exaflops, and Aurora is scheduled for completion in 2024 with a projected performance of approximately 2 exaflops.

## OBJECTIVES OF THE EQSIM DEVELOPMENT

The principal simulation advancement objectives of the EQSIM framework development defined at the project outset included the following:

1) the ability to perform regional-scale ground motion simulations at frequency resolutions relevant to a breadth of engineered systems, which includes regional ground motion simulations on the order of 10-Hz frequency resolution
2) the ability to represent soft, near-surface sedimentary soils that can locally amplify earthquake ground motions
3) rigorous coupling of global geophysics models with local engineering soil/structure models to appropriately account for infrastructure interactions with complex incident seismic wavefields
4) achievement of fast earthquake simulation wall clock times so that many realizations of a specific earthquake scenario can be executed to span the space defining the range of potential fault rupture possibilities
5) ensure that EQSIM can efficiently and fully exploit the performance potential of the new exaflop systems
6) develop an interactive visual-based environment that permits efficient data mining for the very large datasets resulting from regional earthquake simulations.

The ECP project advancements realized toward all these goals have been summarized in detail in a recent journal article.[5]

The final EQSIM workflow and typical regional-scale simulation metrics are illustrated in Figure 4, where

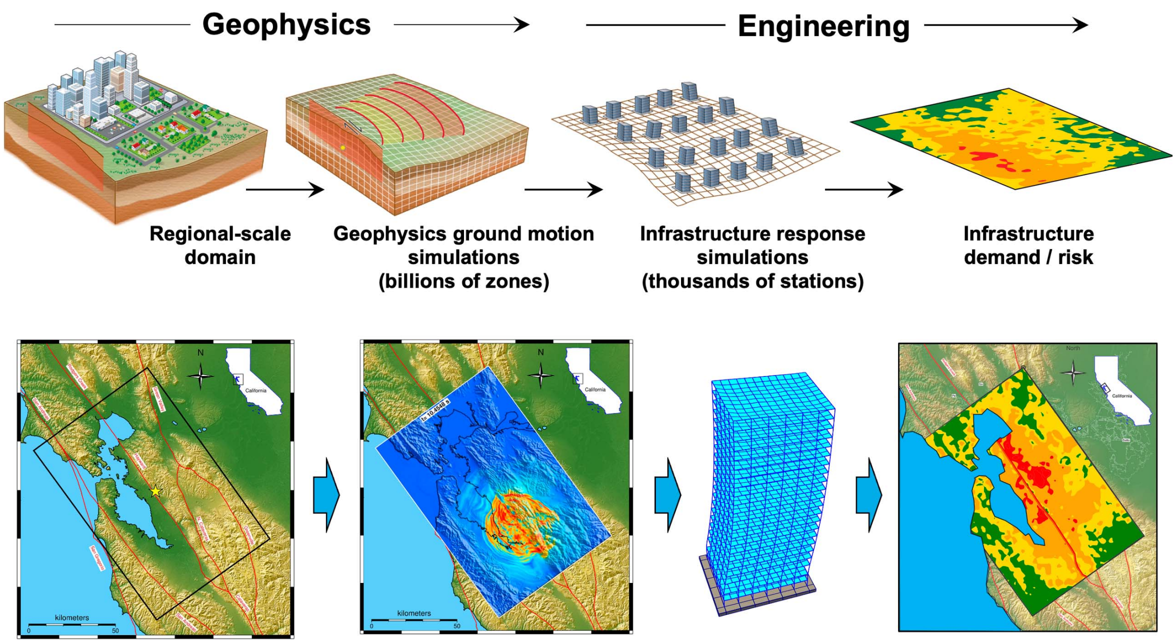

**FIGURE 3.** Fault-to-structure simulations; end-to-end elements (top) and a representative application example in the San Francisco Bay Area, with the final products of site-specific ground motions and infrastructure risk.

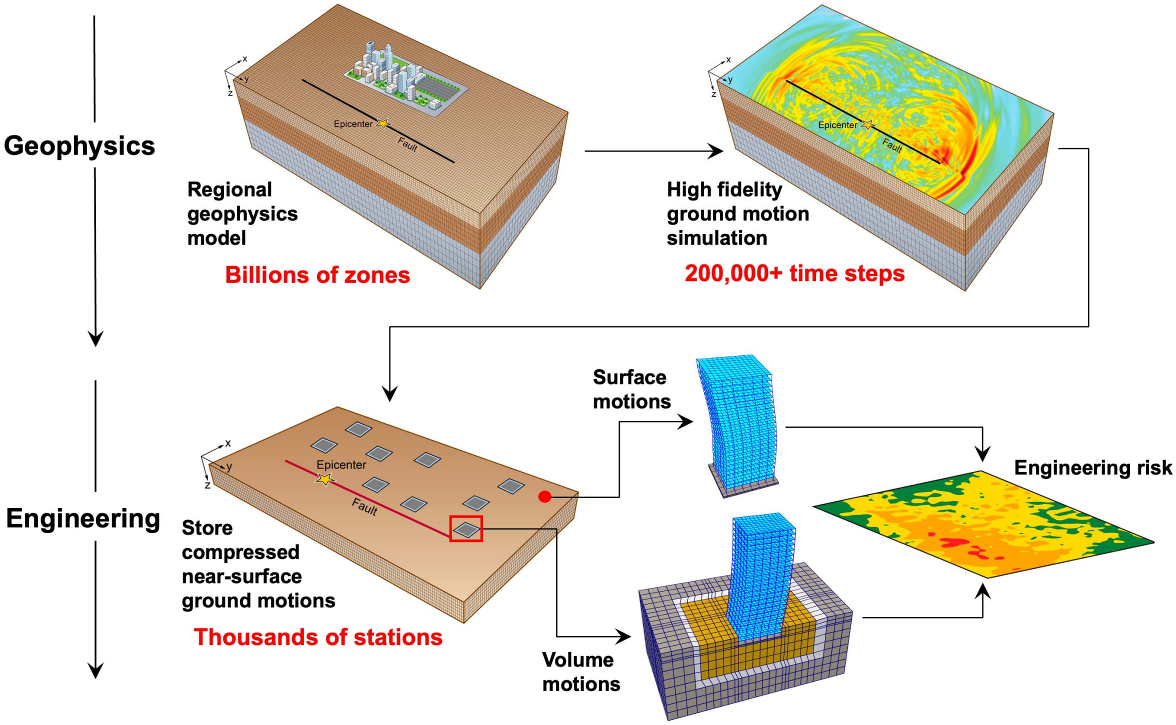

**FIGURE 4.** Sequential steps and workflow in the EQSIM regional fault-to-structure simulations: Step one (geophysics) consists of regional scale wave propagation simulation, step two (engineering) consists of building infrastructure system nonlinear simulations throughout the domain.

step one includes the geophysics simulation of the fault rupture and seismic wave propagation throughout the domain. This step must be executed on the most powerful compute engines available and for the case of EQSIM the Frontier exaflop system. As illustrated next for a San Francisco Bay area regional simulation domain, the geophysics finite difference models[5,6,7] can contain hundreds of billions of model grid points, and a specific large earthquake simulation can consist of a total number of time steps of order 200,000 to 500,000. The ground motion data for a near-surface volume of the regional model is stored for subsequent infrastructure response simulations in step 2. For the infrastructure response simulations in step 2, nonlinear finite-element models are created that can consist of, for example, a building superstructure or a coupled building and soil system in cases where soil–structure interaction may be an important consideration. The response of infrastructure elements can then be executed at thousands of surface sites throughout the regional domain to develop a regional map of earthquake risk for structures of interest.

The computational demands for such regional-scale earthquake simulations are extreme. For the regional-scale ground motion simulations, it can be shown that the computational effort is proportional to

$$\text{Computational effort } \propto (\text{Model volume}) \times (\text{Earthquake duration}) \times \left(\frac{\text{Fmax}}{\text{Vsmin}}\right)^{4}.$$

The fourth-order dependency of the last term illustrates that it is very computationally challenging to increase the maximum frequency resolution of the simulation (Fmax) and to decrease the minimum shear wave speed (Vsmin), which is an important seismic parameter that characterizes the strength of the soft near-surface soil layers. In the EQSIM development, the integrated computational advancements were achieved through the implementation of advanced efficient computational algorithms, including computational grid refinement optimally matched to the geologic properties,[8] development of efficient and scalable parallel input–output and data reduction methods for extreme-scale datasets,[9] and efficient parallelization and optimization on thousands of GPU-accelerated nodes.[10]

To provide visual demonstration of the significance of the push for higher resolution, Figure 5 illustrates the ground acceleration associated with the earthquake-induced motions at a point on the regional domain surface. The plots show the increase in simulated ground motion amplitude as the frequency resolution of the regional simulation progresses from 1 to 10 Hz. This illustrates the importance of the regional simulations sufficiently resolving higher frequencies to fully capture ground motions that generate the risk to infrastructure. Similarly, Figure 6 visually illustrates the change in the EQSIM model surface geologic structure for the San Francisco Bay Area model as the minimum shear wave speed in the model is reduced from 500 m/s to 140 m/s. As indicated in this figure, the achievement of resolving a 140-m/s seismic wave speed is necessary to honor the geophysical material properties from the United States Geologic Survey geologic database. Numerical achievements of both the frequency resolution and soft sediment representation features are key to developing simulated ground motions that are appropriate for earthquake risk assessment of infrastructure systems.

As the ability to computationally resolve high-frequency ground motions has been realized, there is motivation to continue to improve regional geologic data so that these simulation models are well constrained with the best geologic data for high-frequency simulations.

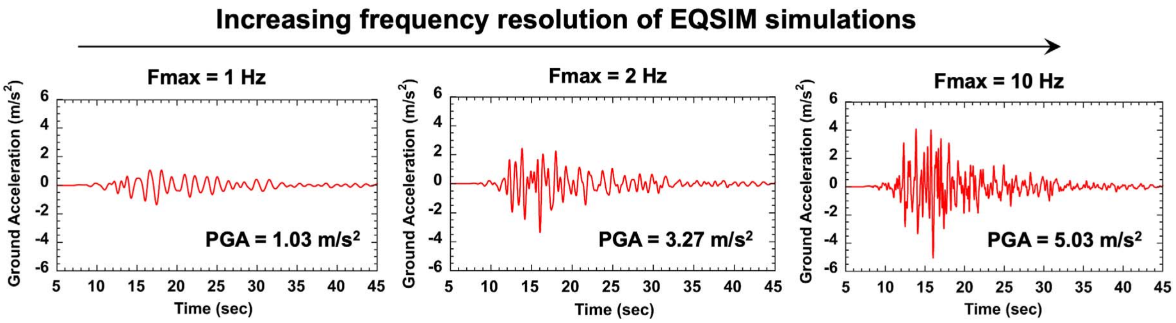

**FIGURE 5.** Increase in ground motion accelerations as the frequency resolution of the ground motion simulation increases from 1 to 10 Hz with a required computational effort increase of 10,000X.

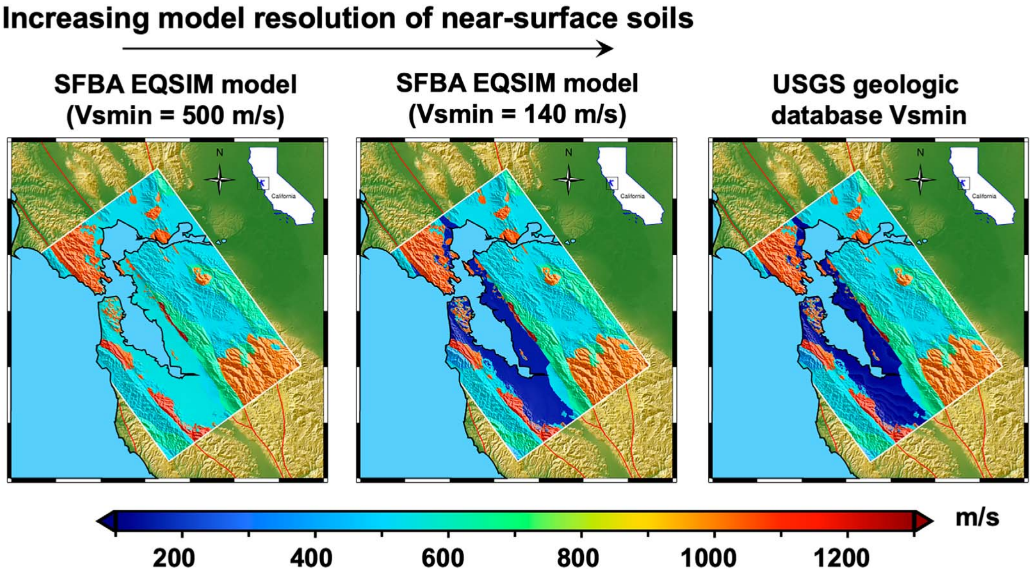

**FIGURE 6.** Increased resolution of near-surface soft sediments as the minimum resolved shear wave speed in the ground motion simulation model decreases from 500 to 140 m/s, which is required to honor the USGS Bay Area velocity data, with a required computational effort increase of 162X.

## EXAMPLE REGIONAL EARTHQUAKE SIMULATIONS FOR THE SAN FRANCISCO BAY AREA

The value proposition of regional simulations lies in the insight that can be gained about the region-wide, site-specific distribution of ground motions and resulting infrastructure risk for a specific scenario earthquake. The first simulation example in Figure 7 illustrates the result of fault-to-structure simulations for a major M7 earthquake on the Hayward Fault in the San Francisco Bay Area region of northern California.[5,11] The seismic waves emanating from the fault rupture along with time-synchronized damage to archetype modern 12-story reinforced concrete frame buildings at each point on the ground surface is shown. This simulation illustrates the importance of the fault directivity effect in the sustainment of major infrastructure damage. Rupture directivity occurs when the speed of the fault rupture propagation is close to the speed of seismic wave propagation resulting in the effective stacking-up of seismic waves in the direction of propagation of the fault rupture. This specific rupture scenario has a hypocenter (point on the fault plane where the fault rupture initiates) located on the southern end of the Hayward Fault. As the fault rupture progresses, the ground motion increases toward the north, leading to a building damage pattern that is significantly more severe in the north. In addition to the directivity of the building damage, the complex site-dependent character of the damage is evident from the geometrically complex contour plots of building damage. This type of insight into site-specific ground motions is essential information for carrying out infrastructure risk assessments. For validation of simulation results, extensive validation comparisons have been made between EQSIM simulated ground motions and measured ground motions from small Bay Area earthquakes as well as comparisons with empirical ground motion models.[12]

A second simulation example, which is illustrative of the current cutting edge of regional simulations, includes an M7.5 San Andreas Fault rupture of much larger extent (Figure 8). This event, executed on the Frontier platform, illustrates the seismic energy propagation as the fault rupture progresses. The ability to simulate these very rare, large events and generate region- and site-specific ground motions can significantly inform the understanding of earthquake motions and fill in the "yellow box" void in Figure 1. This simulation utilized a regional geophysics model with over 500 billion grid points and a 10-Hz resolution simulation. The wall clock time for a M7.5 earthquake of 200-s duration was approximately 40 h on the Frontier system.

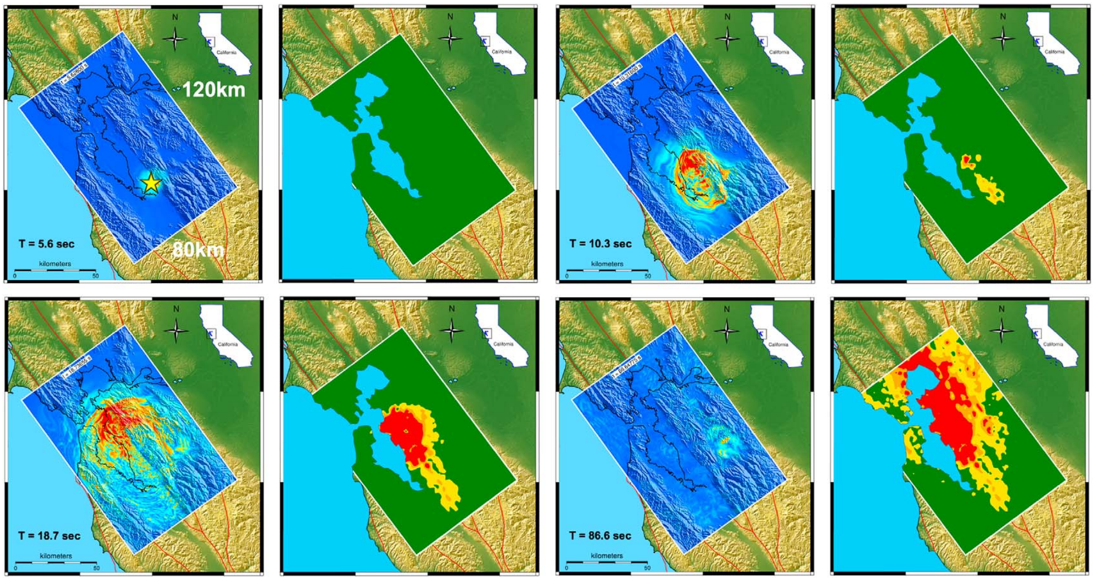

**FIGURE 7.** Magnitude 7.0 Hayward Fault regional-scale earthquake simulation; snapshots of seismic wave propagation and corresponding damage for a typical 12-story reinforced concrete building at each ground surface site (building damage limit states; green = no damage, yellow = minor damage, orange = moderate damage, red = significant damage).

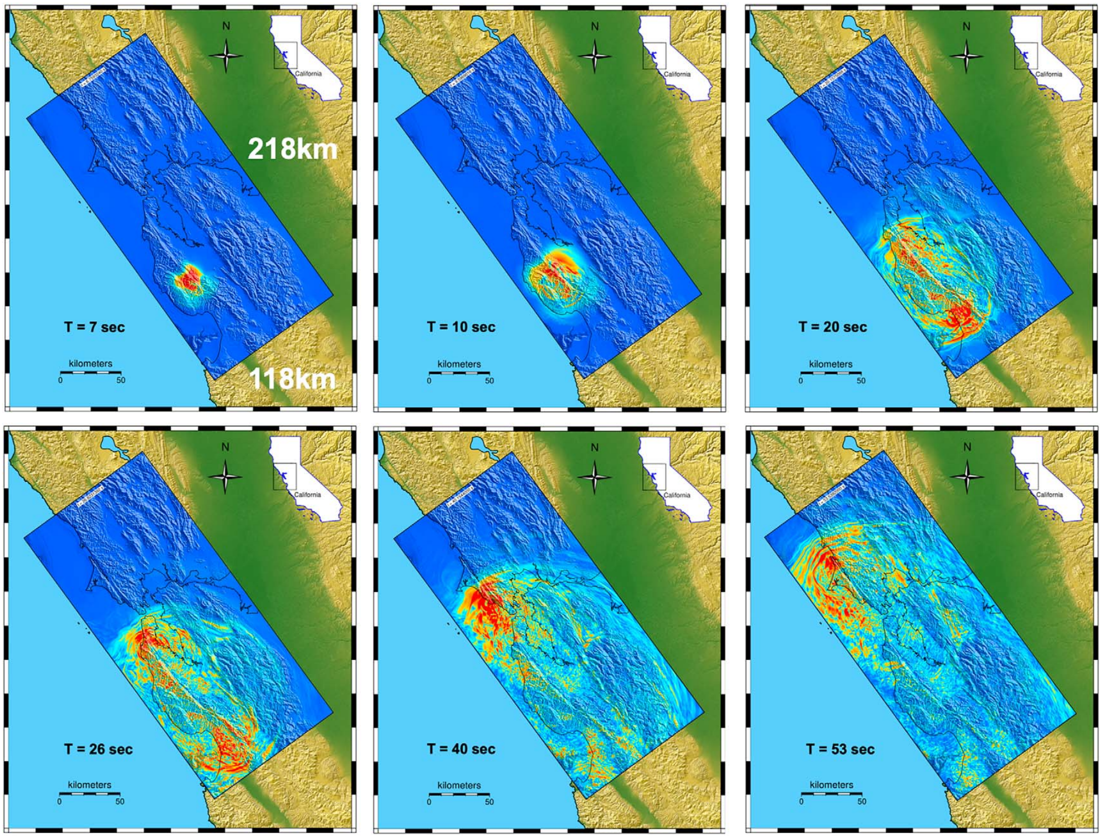

**FIGURE 8.** Magnitude 7.5 San Andreas Fault regional-scale earthquake simulation (red contours indicate a ground velocity $> 0.5$ m/s).

## SUMMARY AND CONCLUSION

The major advancements in high-performance computing platforms and efficient software workflows are enabling transformational multidisciplinary simulations of end-to-end earthquake processes. The numerical simulation initiates from the rupture of the earthquake fault, tracks the seismic wave propagation through the region, and finally the response of infrastructure systems. EQSIM provides scientists and engineers with a powerful new toolset to attack the existing uncertainties in earthquake processes and is particularly relevant to exploring regions and earthquake scenarios where there are no existing historical observations and measurements of large earthquake events. Analysis of simulation results is yielding new insight on the characteristics of ground motions and the variability of infrastructure response (e.g., the pervasive impact of rupture directivity as shown in Figure 7). The application of these simulation capabilities is in its infancy, but with the potential payoff of deeper understanding and enhanced predictions of earthquake effects, there is significant motivation for widespread adoption.

## ACKNOWLEDGEMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. The authors also gratefully acknowledge the critical access to the Frontier exascale system provided by the DOE Office of Science Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. A portion of this work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under 764 Contract DE-AC52-07NA27344.

Dr. Maryam Tabbakhha performed the simulation of reinforced concrete building response that contributed to the illustration of the distribution of building damage in Figure 7 on the Perlmutter computer at Berkeley Lab and that effort is gratefully acknowledged.

## REFERENCES

1. R. Paolucci, C. Smerzini, and M. Vanini, “BB-SPEEDset: A validated dataset of broadband near-source earthquake ground motions from 3D physics-based numerical simulations,” *Bull. Seismological Soc. Amer.*, vol. 111, no. 5, pp. 2527–2545, 2021, doi: 10.1785/0120210089.
2. R. Taborda, and J. Bielak, “Ground-motion simulation and validation of the 2008 Chino Hills, California, earthquake,” *Bull. Seismological Soc. Amer.*, vol. 103, no. 1, pp. 131–156, 2013, doi: 10.1785/0120110325.
3. W. Zhang et al., “A computational workflow for rupture-to-structural-response simulation and its application to Istanbul,” *Earthq. Eng. Structural Dyn.*, vol. 50, no. 1, pp. 177–196, 2021, doi: 10.1002/eqe.3377.
4. R. Graves et al., “CyberShake: A physics-based seismic hazard model for southern California,” *Pure Appl. Geophys.*, vol. 168, nos. 3–4, pp. 367–381, 2011, doi: 10.1007/s00024-010-0161-6.
5. D. McCallen et al., “Regional-scale fault-to-structure earthquake simulations with the EQSIM framework: Workflow maturation and computational performance on GPU-accelerated exascale platforms,” *Earthq. Spectra*, early access, 2024, doi: 10.1177/87552930241246235.
6. D. McCallen et al., “EQSIM - A computational framework for fault-to-structure earthquake simulations on exascale computers part I: Computational models and workflow,” *Earthq. Spectra*, vol. 37, no. 2, pp. 707–735, 2021, doi: 10.1177/8755293020970982.
7. D. McCallen, F. Petrone, M. Miah, A. Pitarka, A. Rodgers, and N. Abrahamson, “EQSIM - A multidisciplinary framework for fault-to-structure simulations on exascale computers part II: Regional simulations of building response,” *Earthq. Spectra*, vol. 37, no. 2, pp. 736–761, 2021, doi: 10.1177/8755293020970980.
8. S. Wang, and N. A. Petersson, “Fourth order finite difference methods for the wave equation with mesh refinement interfaces,” *SIAM J. Sci. Comput.*, vol. 41, no. 5, pp. A3246–A3275, 2019, doi: 10.1137/18M1211465.
9. H. Tang, S. Byna, N. A. Petersson, and D. McCallen, “Tuning parallel data compression and I/O for large-scale earthquake simulation,” in *Proc. IEEE Int. Conf. Big Data*, Orlando FL, USA, 2021, pp. 2992–2997, doi: 10.1109/BigData52589.2021.9671876.
10. R. Pankajakshan, P.-H. Lin, and B. Sjögreen, “Porting a 3D seismic modeling code (SW4) to CORAL machines,” *IBM J. Res. Develop.*, vol. 64, no. 3/4, pp. 17:1–17:11, 2020, doi: 10.1147/JRD.2019.2960218.
11. A. Rodgers, N. Anders Petersson, A. Pitarka, D. B. McCallen, B. Sjogreen, and N. Abrahamson, “Broadband (0–5 Hz) fully deterministic 3D ground-motion simulations of a magnitude 7.0 Hayward fault earthquake: Comparison with empirical ground-motion models and 3D path and site effects from source normalized intensities,” *Seismological Res. Lett.*, vol. 90, no. 3, pp. 1268–1284, 2019, doi: 10.1785/0220180261.
12. F. Petrone, N. Abrahamson, D. McCallen, A. Pitarka, and A. Rodgers, “Engineering evaluation of the EQSIM simulated ground-motion database: The San Francisco Bay Area region,” *Earthq. Eng. Structural Dyn.*, vol. 50, no. 15, pp. 3939–3961, 2021, doi: 10.1002/eqe.3540.

**DAVID MCCALLEN** is the Critical Infrastructure Program leader in the Energy Geosciences Division at Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA. His research interests include advanced simulations at the nexus of earth sciences and engineering. McCallen received his Ph.D. degree in civil engineering from the University of California Davis. Contact him at dbmccallen@lbl.gov.

**ARBEN PITARKA** is the Seismology Group leader at the Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA. His research interests include development of physics-based earthquake rupture models through dynamic rupture modeling and numerical wave propagation modeling applied to strong ground motion simulations. Pitarka received his Ph.D. degree in seismology from Kyoto University. Contact him at pitarka1@llnl.gov.

**HOUJUN TANG** is a computer research scientist in the Scientific Data Division at Lawrence Berkely National Laboratory, Berkeley, CA, 94720, USA. His research interests include data management, storage systems, parallel input–output, and high-performance computing. Tang received his Ph.D. degree in computer science from North Carolina State University. Contact him at htang4@lbl.gov.

**RAMESH PANKAJAKSHAN** is a computer scientist in the Livermore Computing Division at Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA. His research interests include high-performance computing, advanced computer architectures, modeling, simulation, and performance optimization. Pankajakshan received his Ph.D. degree in computational engineering from Mississippi State University. Contact him at pankajakshan1@llnl.gov.

**ANDERS PETERSSON** is a senior scientist at the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory, Livermore, CA, 94550, USA. His research interests include numerical analysis, scientific computing, wave propagation phenomena, and numerical optimization. Petersson received his Ph.D. degree in numerical analysis from KTH Royal Institute of Technology. Contact him at petersson1@llnl.gov.

**MAMUN MIAH** is a project scientist in the Energy Geosciences Division at Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA. His research interests include the nonlinear analysis of infrastructure systems subjected to earthquake motions. Miah received his Ph.D. degree in civil engineering from the University of Mississippi. Contact him at mmiah@lbl.gov.