17th World Conference on Earthquake Engineering, 17WCEE

Sendai, Japan - September 13th to 18th 2020

THE EARTHQUAKE SIMULATION (EQSIM) FRAMEWORK FOR PHYSICS-BASED FAULT-TO-STRUCTURE SIMULATIONS

D. McCallen(1), N.A. Petersson(2), A. Rodgers(3), M. Miah(4), A. Pitarka(5), F. Petrone(6), H. Tang(7)

(1) *Professor and Director of the Center for Civil Engineering Earthquake Research, University of Nevada, Reno, and Senior Scientist, Lawrence Berkeley National Laboratory, dmccallen@unr.edu*

(2) *Applied Mathematician, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Andersp@llnl.gov*

(3) *Seismologist, Atmospheric, Earth and Energy Division, Lawrence Livermore National Laboratory, rodgers7@llnl.gov*

(4) *Postdoctoral Research Fellow, Energy Geosciences Division, Lawrence Berkeley National Laboratory, mmiah@lbl.gov*

(5) *Seismologist and Deputy Group Leader, Atmospheric, Earth and Energy Division, Lawrence Livermore National Laboratory, pitarka1@llnl.gov*

(6) *Assistant Professor, University of Nevada, Reno and Visiting Faculty Scientist, Lawrence Berkeley National Laboratory, florianapetrone@unr.edu*

(7) *Research Scientist, Computational Research Division, Lawrence Berkeley National Laboratory, htang4@lbl.gov*

Abstract

Earthquake ground motions at a specific site are strongly influenced by the physics of the underlying earthquake processes, which include source rupture mechanics, seismic wave propagation through a heterogeneous medium, local site response and soil-structure-interaction at a particular facility. The complexity of earthquake ground motions results in strong spatial variation in both frequency content and amplitude, which ideally must be understood and quantified for optimal engineering evaluations. The aggregation of records from many different earthquakes and locations in traditional empirically based ergodic ground motion models does not capture such site-specific complexity of ground motions. Over the past decade, interest in utilizing advanced physics-based computer simulations to characterize earthquake ground motions and resulting infrastructure response has accelerated significantly. However, the extreme computational demands required to execute hazard and risk simulations at regional scale have been prohibitive. With recent and pending achievements in High Performance Computing (HPC), advancing regional-scale ground motion simulations from the historical computationally limited frequency range of ~2-3 Hz to the frequency range of interest for a breadth of engineered systems of ~8-10 Hz is becoming feasible. A second opportunity is to implement an HPC framework and workflow that directly couples earthquake hazard and risk assessments through an end-to-end simulation framework that spans from earthquake rupture to structural response, thereby capturing the complexities of interaction between incident three-dimensional seismic wave fields and infrastructure systems.

The EQSIM computational framework is an application development under the U.S. Department of Energy Exascale Computing Project (ECP), and is focused on creating a simulation workflow and computational toolset for earthquake hazard and risk assessment. Starting with a set of the existing application codes, EQSIM is building an end-to-end capability to simulate from fault rupture to surface ground motions (hazard) and ultimately to infrastructure response (risk). The goal of the EQSIM development is to remove computational limitations as a barrier to scientific exploration and understanding of earthquake phenomenology, and to enable comprehensive simulation-based earthquake hazard and risk assessments. To achieve the overall goals, regional-scale ground motion simulations must be executed at unprecedented frequency resolution using large models with faster run times to allow the parametric variations necessary to span critical problem parameters (e.g. multiple fault ruptures). In addition, as the ability to compute to higher frequencies progresses, enhanced characterization of subsurface geologic structure at finer scales is needed, thus a companion schema for representing fine-scale geologic heterogeneities in massive computational models must be developed.

In this article recent computational advancements in regional-scale simulations are described and representative simulations of ground motion hazard and earthquake risk to buildings are presented. The ability to successfully perform regional simulations at frequencies of engineering relevance, including up to 10 Hz, is demonstrated. Simulation results provide insight into the complex distribution of earthquake risk and the degree to which risk for a particular structure can vary for a given earthquake event.

Keywords: Regional simulations, Exascale computing, Physics-based models, Infrastructure risk

1. Introduction

Cutting-edge capabilities in HPC platforms and computational ecosystems have demonstrated continuous advancement at an impressive pace. The inexorable evolution of hardware and computational schemas have overcome performance barriers and sustained performance advancements. As indicated in Fig. 1, which documents the leading scientific compute platform over the last twenty-five years, hardware performance has demonstrated six orders of magnitude increase. To push to the next level of Exaflop platforms currently under development, the U.S. Department of Energy Exascale Computing Project (ECP) is developing a breadth of scientific software applications and enabling software stack (efficient I/O, parallel work distribution on accelerator platforms, exploitation of new accelerator hardware etc.) to be prepared to exploit accelerator exascale platforms in the 2021-2023 time-frame as indicated in Fig. 1.

One of the applications under development for the ECP is an advanced framework for regional-scale assessments of earthquake hazard and risk. The objectives of the Earthquake Simulation (EQSIM) framework is to develop a multidisciplinary workflow for end-to-end, fault-to-structure simulation of earthquake processes as shown in Fig. 2. The EQSIM application development is providing a computational capability to compute from earthquake source to infrastructure response, which requires multidisciplinary coupling of geophysics and engineering simulation codes. The ultimate objective is to remove the current significant computational barriers to allow regional-scale, fault-to-structure simulations at frequencies of relevance to engineered systems. Successful development of this framework will provide a computational capability that can address existing uncertainties and two important questions:

- How do earthquake ground motions spatially vary over a region of interest and what are the implications for infrastructure systems?
- How do complex, three-dimensional earthquake wavefields actually interact with the local superstructure/soil system for specific engineered infrastructure?

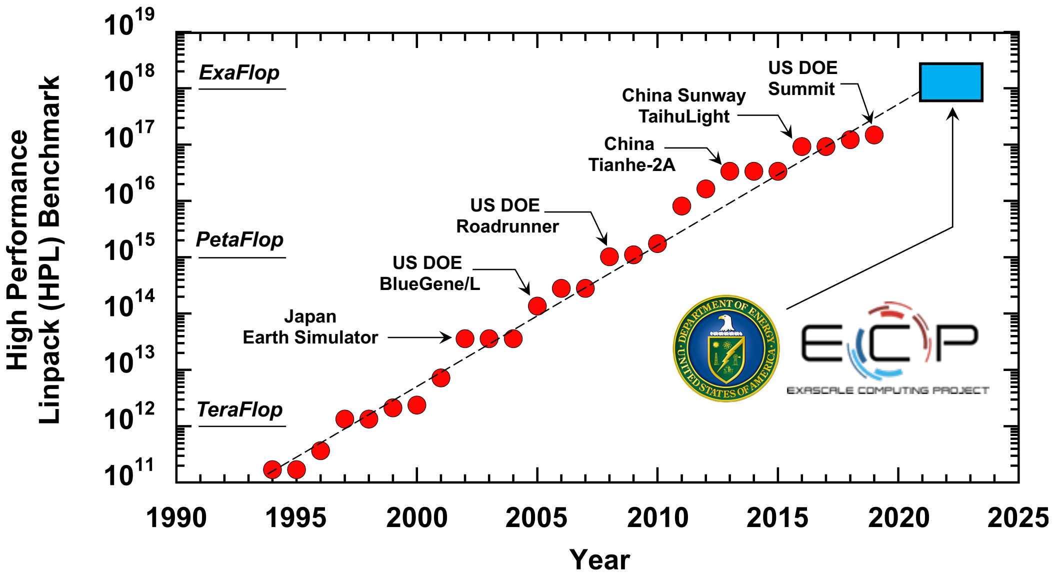

Fig. 1 - Advancements in HPC platforms and the 2023 target objective of the U.S. DOE Exascale Computing Project (source: top500.org).

In this article the details of the EQSIM application and workflow are described, and the code developments and performance metrics achieved to-date are summarized. Early simulation results showing the complexity of the regional-scale distribution of infrastructure risk are presented and the next steps towards exascale readiness are discussed.

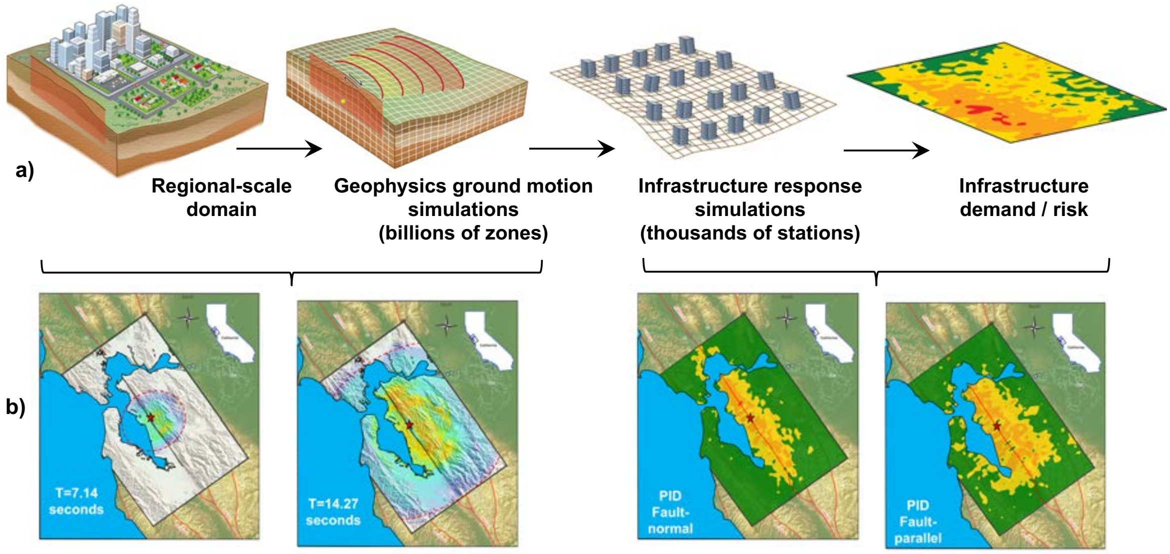

Fig. 2 - End-to-end, fault-to-structure earthquake simulations. a) Linking geophysics and engineering simulations; b) resulting regional scale ground motions (hazard) and infrastructure response (risk).

2. EQSIM workflow and computational strategies at 100's of billions of grid zones

In the EQSIM workflow geophysics and engineering simulation codes are linked to perform end-to-end simulations as shown in Fig. 3. The EQSIM framework provides two options for code coupling, the first option utilizes surface ground motion time histories computed at discrete points on the earth's surface from the geophysics code as input directly into fixed-base structural models assuming the motions are identical across the structure footprint (termed *weakly coupled*) and the second option couples the geophysics and engineering codes through the Domain Reduction Method (DRM) developed by Bielak et. al. (termed *strongly coupled*) [1, 2]. The first approach is representative of traditional fixed-base structural system simulations and is only consistent with an assumption that the ground motions at a discrete point on the earth surface are generated by pure vertically propagating shear (for transverse motion) and compressional (for vertical motion) waves. The strong coupling case allows for the consideration of soil-structure-action, as well as the complex three-dimensional wavefields consisting of arbitrarily oriented body waves and surface waves.

Historically, the assumption of vertically propagating waves for fixed-base structural evaluations was necessary because of the lack of understanding of the true nature of complex incident wavefields, however, with full wavefield simulations to high frequency, the ability to understand infrastructure response to more complex wavefields is now possible. The DRM boundary allows for a detailed assessment of a complex incident wavefield and can be particularly important in understanding seismic wave-structure interactions in the near-fault region.

Each ECP application development must have well defined goals for the exascale challenge problem. For EQSIM the basic goal includes being able to execute regional-scale simulations to high frequency resolution, with sufficient speed to allow a significant number of realizations. In order to develop a realistic model domain, the San Francisco Bay Area (SFBA) of Northern California in the United States was adopted as a numerical testbed for regional-scale simulations as shown in Fig. 4. For the selected domain, the challenge goal for the exascale platforms under development is to execute regional-scale simulations to 10Hz with a wall clock time of approximately 5 hours per earthquake simulation. The desire to reach an execute time ~5 hours will make parametric studies possible to develop risk distributions from different rupture scenarios, different geologic realizations etc. It is noted that this goal is a major jump from previous regional-scale simulations for this extent of a regional domain. Historically, earthquake ground motion simulations of this extent have been limited to the range of on the order of 2Hz, with compute times on the order of 20-30 hours, given the fact that the computational effort varies as the resolved frequency to the fourth power, achieving 10Hz resolution is a very challenging goal. Achieving this goal requires approximately a 5x increase in frequency resolution and a 5x speed-up relative to previously executed simulations (Fig. 4). This ambitious goal can only be accomplished through the combined exploitation of advanced algorithms, code optimization for very large-scale parallel implementations and the utilization of emerging exascale platforms.

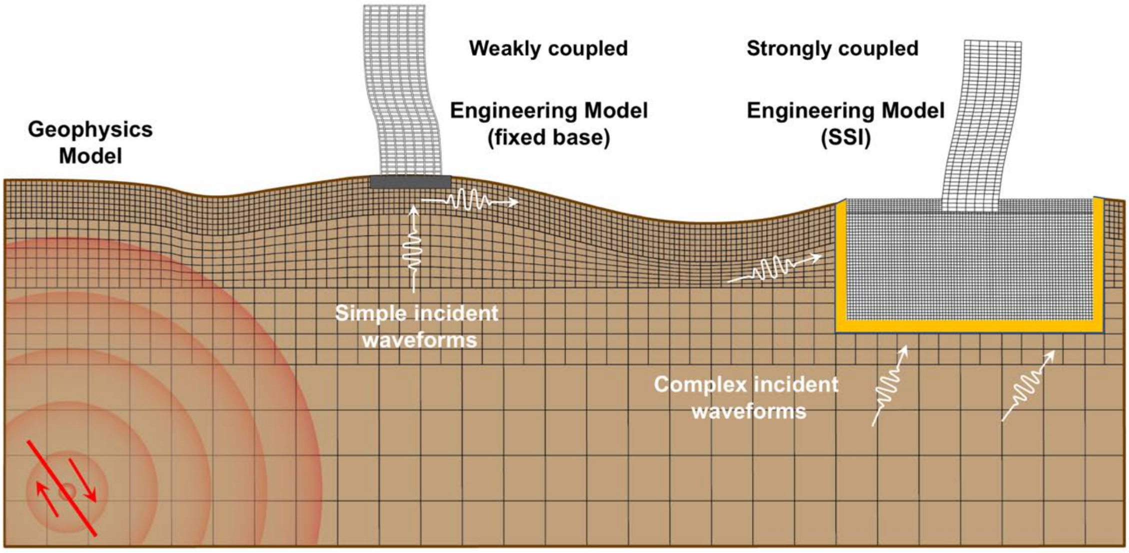

Fig. 3 - Coupling geophysics and engineering simulation codes; weak coupling by applying surface motions to a fixed-base structural model, strong coupling by coupling geophysics and engineering codes through the Domain Reduction Method.

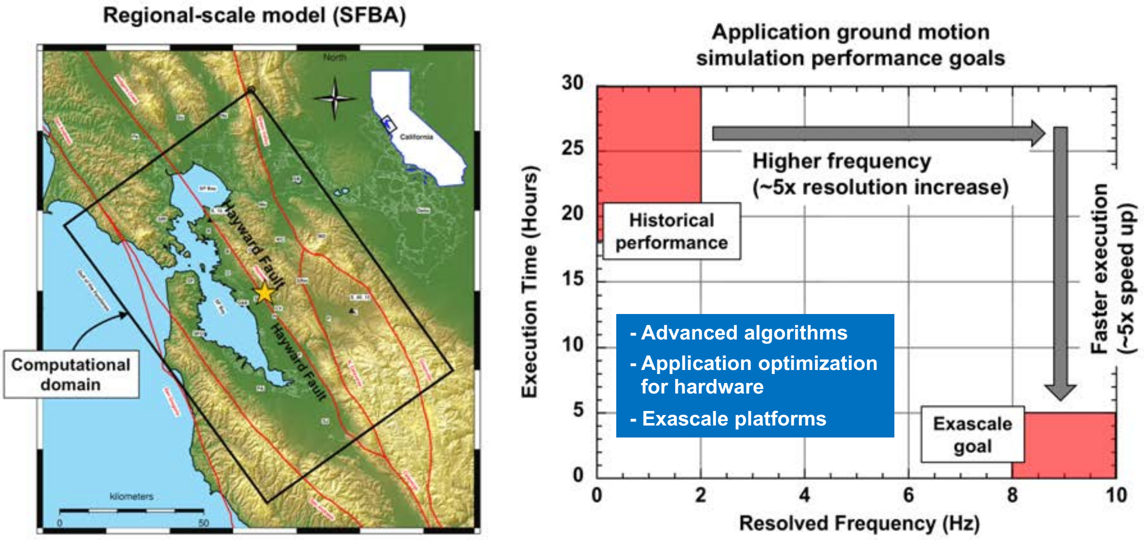

Fig. 4 – Region scale model of the San Francisco Bay Area of California, U.S.A., historical performance of regional scale simulations and EQSIM exascale goal.

The EQSIM development project started from the adoption of a core set of existing simulation codes. In the case of geophysics wave propagation, the SW4 summation by parts finite difference code from the Lawrence Livermore National Laboratory, was adopted [3, 4, 5]. The SW4 code is 4th order accurate in space and time and has a number of existing key features including representation of ground surface topography and effective non-reflecting boundaries. SW4 makes use of a combination of a Cartesian grid at depth and a curvilinear grid near the earth surface to account for elevation changes due to surface topography. To achieve the most efficient capability for regional scale simulations in EQSIM, a number of important new algorithms and capabilities have been implemented in the SW4 code over the past two years including (Fig. 5):

- Mesh refinement in both the curvilinear and Cartesian grids to optimize the computational domain in accordance with the natural depth-dependent variation of geologic material properties [6];
- The implementation of a new semi-stochastic fault rupture model based on the Graves and Pitarka rupture formulation [7];
- Optimization of SW4 for the existing high-performance platform architecture of the CORI computer at Lawrence Berkeley National Laboratory (https://www.nersc.gov);
- Utilization of HDF5 hierarchical data formats to significantly improve the SW4 massively parallel I/O performance both in terms of data size reduction and data read/write time reductions for massively parallel problems [8];
- Utilization of the Lawrence Livermore National Laboratory RAJA encapsulation model for computer architecture portability and transition to GPU based platforms including the SUMMIT computer at Oak Ridge National Laboratory (https://www.olcf.ornl.gov/summit/);
- Parallel implementation of the overlay of fine-scale stochastic geology to model the scattering of high frequency seismic waves [9].

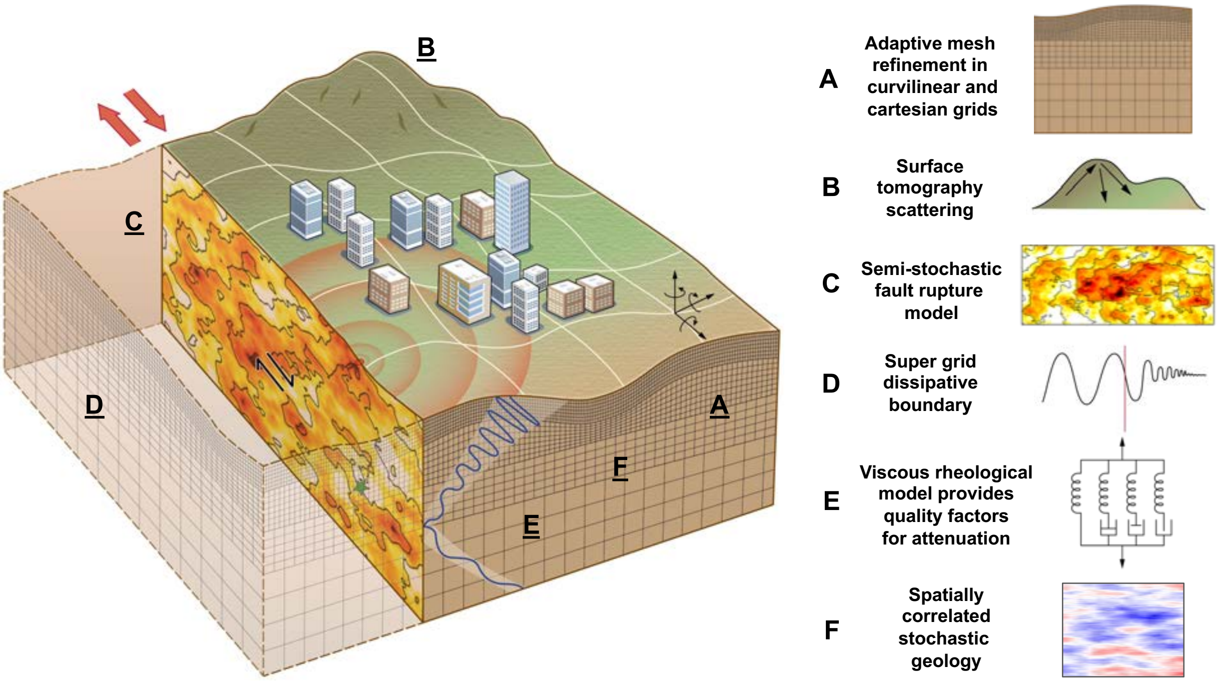

Fig 5. – Advanced algorithms and features of the SW4 geophysics code.

For the simulation of engineering systems, the capability of linking to multiple finite element-based engineering codes has been incorporated in the EQSIM workflow. These include the NEVADA program for the nonlinear analysis of steel frame buildings [10, 11], the OpenSees program for nonlinear analysis of soil-structure systems (opensees.berkeley.edu) and the ESSI program for nonlinear soil-structure-interaction (http://sokocalo.engr.ucdavis.edu/~jeremic/Real_ESSI_Simulator/). The NEVADA program is used to efficiently perform regional scale nonlinear building simulations based on fixed-base input motion and the OpenSees and ESSI programs are both used for soil-structure systems with coupling to SW4 provided through the DRM.

Substantial effort has focused on the development of the computational workflow necessary to efficiently execute fault-to-structure simulations in a massively parallel environment. When executing regional models with the 200-300 billion grid points necessary to resolve 5-10Hz frequencies, every aspect and step of the workflow must be tuned for efficiency and high performance. Features of the EQSIM workflow include (Fig. 6):

- The development of HDF5 binary databases for storing geologic material properties, surface topography and the fault rupture parameters;
- At run time, reading of HDF5 data files and the automatic generation of the SW4 input file for the problem being executed, it is prohibitive to store massive input files associated with 200 billion zone problems, the SW4 input file is autogenerated at run time based on problem command data as a core element of model execution;
- Execution of the ground motion simulation time stepping loop with the creation of very large checkpoint (restart) files at user specified time steps to provide a restart state in the event of a core or node failure during a massively parallel run;
- The creation of an earthquake ground motion database in HDF5 format either in terms of ground surface motions throughout the region at a specified station density (for a weakly coupled problem) or in a specified three-dimensional volume surrounding a particular facility (for a strongly coupled problem);
- Execution of infrastructure response simulations for a weakly or strongly coupled problem.

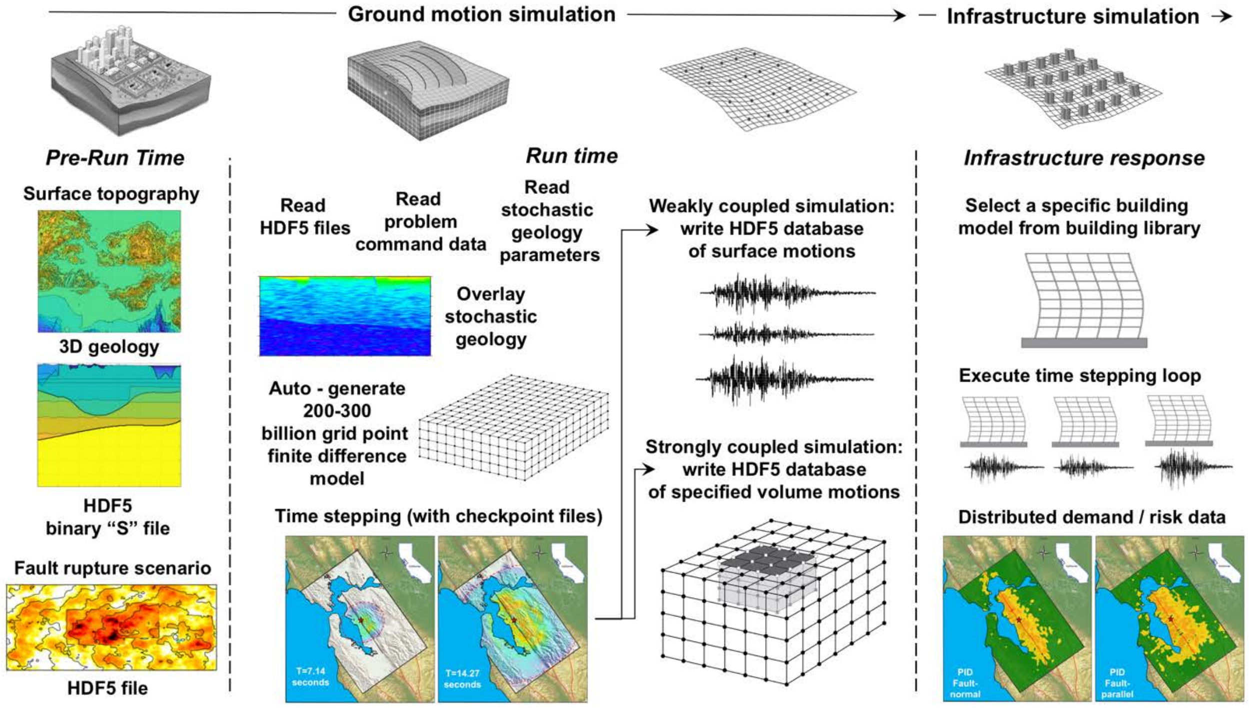

Fig 6. – EQSIM computational workflow including on-the-fly generation of large input files at run time.

The ability to effectively distribute the computations in a massively parallel environment with efficient scaling is a core element of running high resolution regional scale models. The SW4 code utilizes a "pencil" strategy whereby vertical pencil-shaped segments of the computational domain are distributed to the various compute nodes of a massively parallel computer as shown schematically in Fig. 7. For the case of regional-scale simulations of building response, the nonlinear analyses of a selected representative building at each specified station on the ground surface are distributed with one building to one machine core, as shown in Fig. 8. In the current EQSIM framework, a planar building model can be selected from an existing building library and the building response is analyzed at each surface station in the domain for fault parallel and fault normal ground motion components. This trivial parallelization is very effective from the standpoint of achieving very efficient scaling with the number of building analyses. For example, as indicted in Table 1, the one building – one core model allows 10,000 building simulations to be executed with only 1.5 times the wall clock time of single building simulation. On platforms with on the order of a half million cores, this parallelization strategy is very effective and does not demand excessive portions of the available compute nodes.

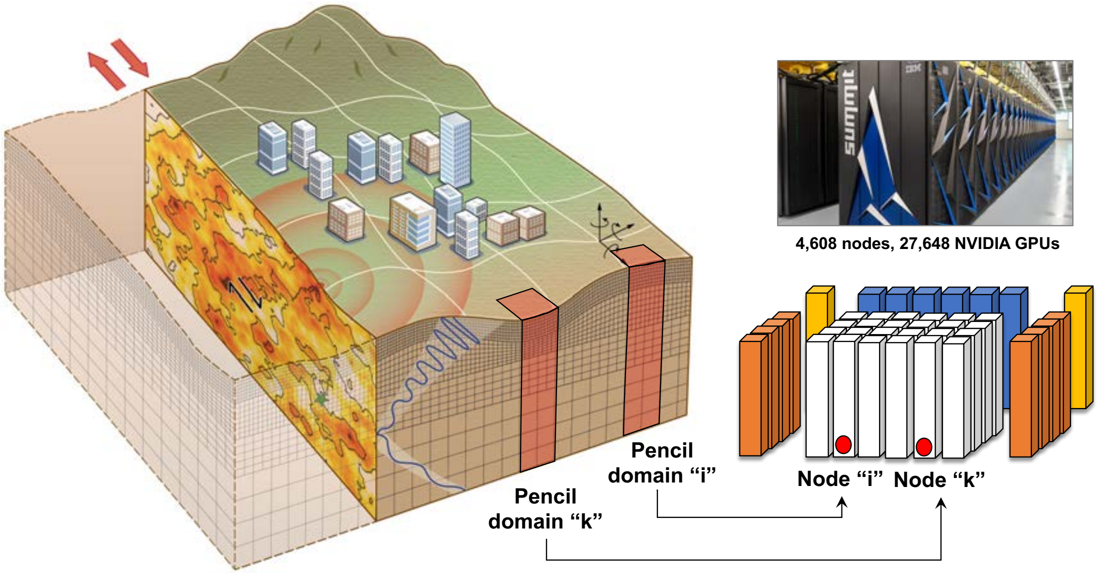

Fig 7. – SW4 ground motion model parallel distribution on SUMMIT, "pencil" domains distributed across different computer nodes.

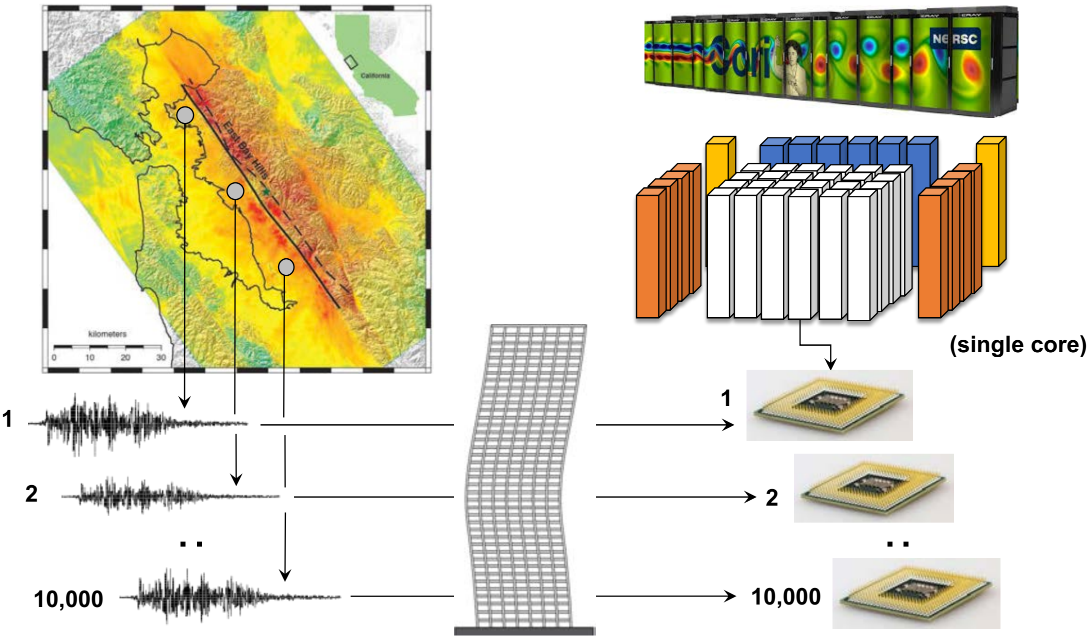

Fig 8. – NEVADA building simulations parallel distribution on CORI, one nonlinear building simulation distributed to one core.

3. Regional-scale simulations and application performance

As part of the EQSIM development, the simulation performance is tracked against the ultimate exascale challenge goal. Each application establishes a Figure of Merit (FOM) that quantitatively measures progress. For EQSIM a FOM for massively parallel regional ground motion simulations was established as,

$$\text{FOM} = \frac{(\text{Freq})^4}{(\text{Wall Clock Time x } 7.6)} \left(\frac{500}{\text{Vsmin}}\right)^4 \tag{1}$$

Where *Freq* is the highest frequency resolved in the ground motion simulation, *Wall Clock Time* is the run time of the regional scale simulation and *Vsmin* is the minimum shear wave velocity in the regional scale geologic model. As the various algorithm advancements and platform optimizations have progressed since the EQSIM development started in 2017, significant improvements in SW4 performance have been achieved as shown in Fig. 9. The left curve in Fig. 9 shows the step to a new performance curve with each successive enhanced performance development in SW4. The first four performance curves, which are anchored by the specific performance data points achieved (colored boxes), and vary with frequency to the forth power, illustrate the performance improvements on the LBNL CORI computer (~30 PF machine). Finally, the last performance curve, corresponding to the run shown as "E" on the plot, illustrates the performance increase from transitioning to the SUMMIT computer, currently ranked number one in the world (Fig. 1) with performance in the 200 PF range. The graph on the right of Fig. 9 illustrates the corresponding increase in the FOM for these performance steps.

Table 1. NEVADA building simulation scaling on CORI (three story building).

Number of Nonlinear Building Simulations (2000 time steps)	**Run Time (Min:Sec)**
1	5:56
10	7:10
100	8:08
1000	9:16
10,000	9:55

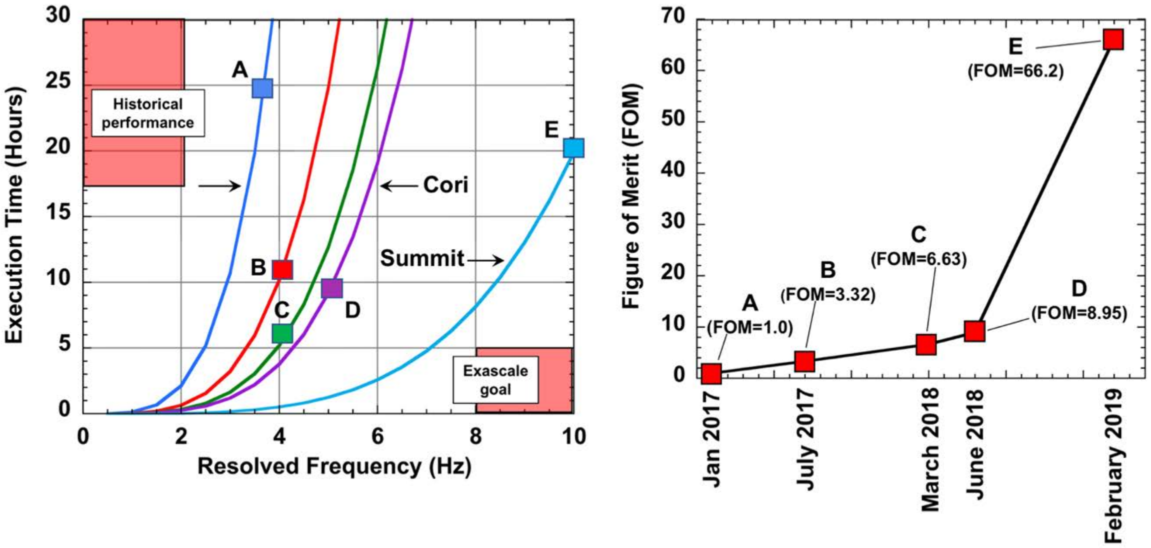

Fig 9. – Progress in the SFBA regional scale ground motion simulations since the inception of EQSIM development.

With the major performance increase that was realized in transition to the GPU-based SUMMIT platform, the EQSIM team was able to execute the first 10Hz ground motion simulation with the SFBA model. The metrics of this ground-breaking simulation are summarized in Table 2. The 10Hz run utilized 1,200 compute nodes, approximately one quarter of the entire computer, and took approximately 21 hours of wall clock time for a regional model of 203 billion grid points.

Table 2. SFBA M7 Hayward Fault event simulation to 10Hz on SUMMIT.

Ground Motion Simulations for the SFBA Model on Summit	
Frequency resolved	10 Hz
Vsmin	500 m/s
Number of grid points	203 Billion
Smallest cell size	6.25 m
Time step size	7.119e-4
Total time steps	126,430
Number of nodes	1,200
Wall clock time	~21 hours (one check point file creation)

To provide some historical perspective to this simulation achievement, the history of regional-scale SFBA ground motion simulations performed by a number of researchers are summarized in Fig. 10. The dramatic increase in frequency resolution since the start of the EQSIM project in comparison to previous simulations is evident.

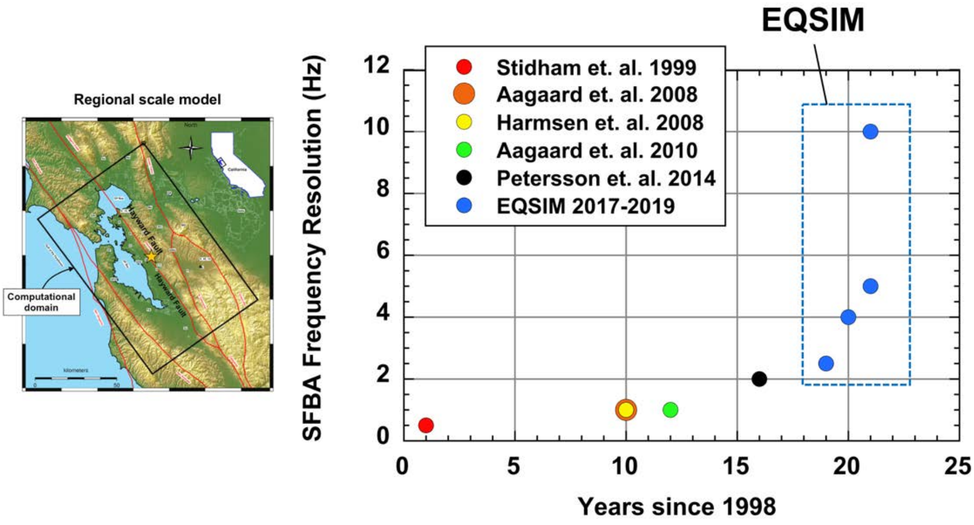

Fig 10. – History of maximum frequency resolution in San Francisco Bay Area ground motion simulations.

With the computational advancements achieved to-date, the SFBA model is being evaluated and utilized to perform unprecedented simulations and provide new insight into SFBA earthquakes [12, 13]. Once the regional scale ground motion simulations are completed, the ground surface motions can be utilized to evaluate structural risk for representative building structures. Fig. 11 shows the distribution of demand/risk, as quantified by the peak interstory drift that occurs anywhere in the building for the duration of the earthquake with weak coupling type of analysis, for 9 and 20 story steel moment frame buildings for a Mw 7 Hayward fault rupture [14]. The drift limit plots, separated into building response to fault-parallel and fault-normal ground motion components, illustrate the complex distribution of building risk as compared to traditional hazard assessments where the ground motions are assumed equivalent for all sites at equal distance from the fault. The variability of demand / risk is further illustrated in Fig. 12 where the variation of peak drift demand for a twenty-story steel moment frame building is shown for sites located along lines parallel to, and at equal distance from the Hayward fault trace at 4Km and 20Km distance respectively from the fault. For the synthetic ground motion records at stations along these lines it is observed that there is significant variability in building demand for sites located close to the fault – up to a factor of almost six for buildings located 4Km from the fault.

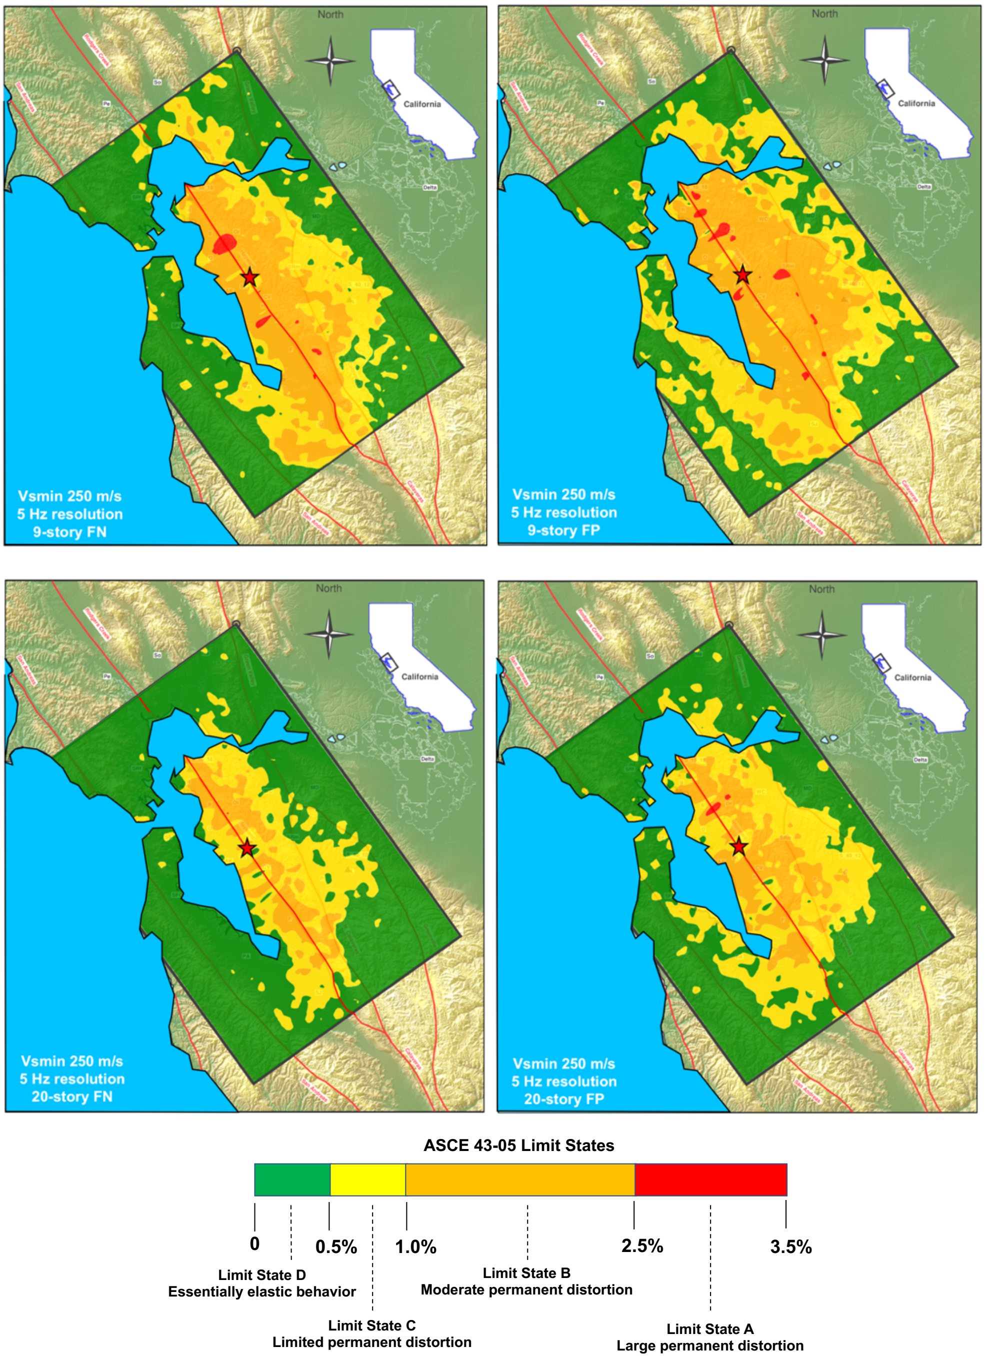

Fig 11. – Distribution of demand/risk for 9 and 20 story steel moment frame buildings for a Mw 7 Hayward fault earthquake, contours of peak interstory drift color coded by ASCE 43-05 drift limit states.

The simulation results indicate that for the same earthquake and the same building the realized risk variability is substantial in the near-fault region, which is in significant contrast to historical equal hazard at equal distance assumptions and raises some challenging questions for engineering design.

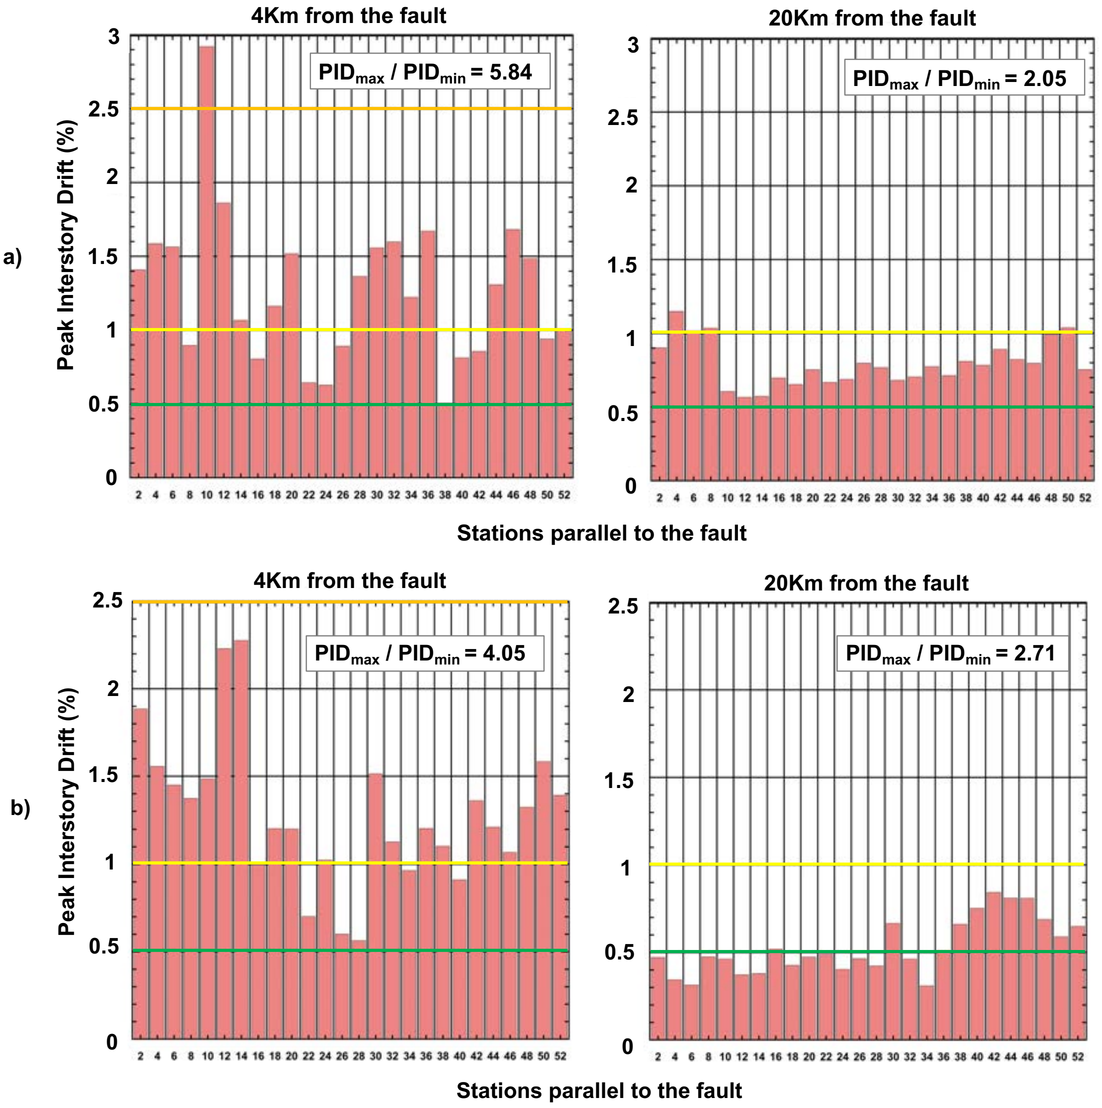

Fig 12. – Distribution of demand/risk for stations equidistant from the fault. a) Twenty story building subjected to fault parallel motions; b) twenty story building subjected to fault normal motions.

4. Conclusions

The ability to efficiently simulate earthquake ground motions and infrastructure risk at regional scale will provide an opportunity to gain new insights and understanding of earthquake phenomenon and to develop transformational simulation-based approaches to earthquake hazard and risk assessments. Much remains to be done in terms of simulation model validation and comprehensive investigation of complex interactions between incident earthquake waves and infrastructure systems, but the emerging simulation advancements of EQSIM will provide the computational tools necessary for deep numerical exploration in a regime of frequency resolution relevant to engineering applications.

5. Acknowledgements

This research was supported by the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of two U.S. Department of Energy organizations - the Office of Science and the National Nuclear Security Administration. This work was performed under the auspices of the U.S. Department of Energy at the Lawrence Berkeley National Laboratory, and at the Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The access and support of the supercomputer facilities at Lawrence Berkeley and Oak Ridge National Labs was essential to this application development work and gratefully acknowledged.

5. References

[1] Bielak J, Loukakis K, Hisada Y, Yoshimura C (2003): Domain reduction method for three-dimensional earthquake modeling in localized regions, part I: theory, *Bulletin of the Seismological Society of America*, **93** (2), 817-824.

[2] Yoshimura C, Bielak J, Yoshiaki H, Fernandez A (2003): Domain reduction method for three-dimensional earthquake modeling in localized regions, part II: verification and applications, *Bulletin of the Seismological Society of America,* **93** (2), 825-840.

[3] Sjogreen B, Petersson NA (2012): A forth order accurate finite difference scheme for the elastic wave equation, *Journal of Scientific Computing*, **52** (1), 17-48.

[4] Petersson NA, Sjogreen B (2014): Super-grid modeling of the elastic wave equation in semi-bounded domains, *Communications in Computational Physics*, **16**, 913-955.

[5] Petersson NA, Sjogreen B (2015): Wave propagation in anistropic elastic materials and curvilinear coordinates using a summation-by-parts finite difference method, *Journal of Computational Physics*, **299**, 820-841.

[6] Wang S, Petersson NA (2019): Fourth order finite difference methods for the wave equation with mesh refinement interfaces, *SIAM Journal of Scientific Computing*, **41** (5), 3246-3275.

[7] Pitarka A, Graves R, Irikura K, Miyakoshi K, Rodgers A (2019): Kinematic Rupture Modeling of Ground Motion from the M7 Kumamoto, Japan Earthquake, *Pure and Applied Geophysics*, https://doi.org/10.1007/.

[8] Byna S, Breitenfeld M, Dong B (2020): ExaHDF5: Delivering efficient parallel I/O on exascale computing systems, *Journal of Computer Science and Technology*, 35 (1), 145-160.

[9] Hirakawa E, Pitarka A, Mellors R (2016): Generation of shear motion from an isotropic explosion source by scattering in heterogenous media, *Bulletin of the Seismological Society of America*, **106** (5), 2313-2319.

[10] McCallen D, Larsen S (2003): NEVADA – A simulation environment for regional estimation of ground motion and structural response, Lawrence Livermore National Laboratory Directed Research and Development *Report UCRL-ID-152115*.

[11] Miah M, Petrone F, Wong J, McCallen D (2018): Regional scale earthquake risk estimation based on broadband ground motion simulations, *Proceedings of the 11th U.S. National Conference on Earthquake Engineering, Los Angeles Ca.*, paper 1278.

[12] Rodgers A, Pitarka A, McCallen (2019): The effect of fault geometry and minimum shear wave speed on 3D ground motion simulations for an Mw 6.5 Hayward fault scenario earthquake, San Francisco Bay Area, Northern California, *Bulletin of the Seismological Society of America*, **109** (4), 1265-1281.

[13] Rodgers A, Petersson NA, Pitarka A, McCallen D, Sjogreen B, Abrahamson N (2018): Broadband (0-5Hz) fully deterministic 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 Research Letters*, **90** (3), 1268-1284.

[14] American Society of Civil Engineers (ASCE), Seismic Design Criteria for Structures Systems and Components in Nuclear Facilities, *ASCE/SEI 43-05*.