Lawrence Berkeley National Laboratory

Recent Work

Title

EQSIM—A multidisciplinary framework for fault-to-structure earthquake simulations on exascale computers part I: Computational models and workflow

Permalink

<https://escholarship.org/uc/item/4dh174t0>

Journal

Earthquake Spectra, 37(2)

ISSN

8755-2930

Authors

McCallen, D
Petersson, A
Rodgers, A
et al.

Publication Date

2021-05-01

DOI

10.1177/8755293020970982

Peer reviewed

EQSIM—A multidisciplinary framework for fault-to-structure earthquake simulations on exascale computers part I: Computational models and workflow

Earthquake Spectra

1–29

© The Author(s) 2020

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/8755293020970982

journals.sagepub.com/home/eqs

**David McCallen, M. EERI^{1,2}, Anders Petersson³,
Arthur Rodgers³, Arben Pitarka³ , Mamun Miah²,
Florian Petrone, M. EERI^{1,2}, Bjorn Sjogreen³,
Norman Abrahamson⁴, and Houjun Tang²**

Abstract

Computational simulations have become central to the seismic analysis and design of major infrastructure over the past several decades. Most major structures are now “proof tested” virtually through representative simulations of earthquake-induced response. More recently, with the advancement of high-performance computing (HPC) platforms and the associated massively parallel computational ecosystems, simulation is beginning to play a role in increased understanding and prediction of ground motions for earthquake hazard assessments. However, the computational requirements for regional-scale geophysics-based ground motion simulations are extreme, which has restricted the frequency resolution of direct simulations and limited the ability to perform the large number of simulations required to numerically explore the problem parametric space. In this article, recent developments toward an integrated, multidisciplinary earth science-engineering computational framework for the regional-scale simulation of both ground motions and resulting structural response are described with a particular emphasis on advancing simulations to frequencies relevant to engineered systems. This multidisciplinary computational development is being carried out as part of the US Department of Energy (DOE) Exascale Computing Project with the goal of achieving a computational framework poised to

¹University of Nevada, Reno, NV, USA

²Lawrence Berkeley National Laboratory, Berkeley, CA, USA

³Lawrence Livermore National Laboratory, Livermore, CA, USA

⁴University of California, Berkeley, CA, USA

Corresponding author:

David McCallen, University of Nevada, 1664 North Virginia Street, Reno, NV 89557, USA.

Email: dmccallen@unr.edu

exploit emerging DOE exaflop computer platforms scheduled for the 2022–2023 timeframe.

Keywords

Regional earthquake simulations, high-performance computing, coupled geophysics—engineering simulations, fault-to-structure simulations, infrastructure risk

Date received: 6 May 2020; accepted: 28 September 2020

Introduction

Computational simulations are well established as an essential tool for the seismic design of major infrastructure systems. Both linear and nonlinear simulations of structure and soil systems can now be performed in support of risk-informed, performance-based infrastructure evaluations (Deierlein et al., 2010; Lu and Guan, 2017). These simulations allow exploration of structural response characteristics, quantification of the probability of exceeding established infrastructure performance limits states, and evaluation of structural design and retrofit options.

In the overall evaluation of earthquake risk to major infrastructure, there continue to be major uncertainties in the site-specific ground motions that will be experienced in future earthquakes. Empirically based, probabilistic assessments of ground motions typically utilize measured earthquake motions obtained from many locations, which are homogenized based on an ergodic assumption of earthquake processes, which assumes the median and variability of the ground motion for a given scenario is the same for all data used in developing a ground motion model (Anderson and Brune, 1999). There has been a recognition that there are strong systematic differences in the ground motions for different regions and using the ergodic assumption can significantly underestimate or overestimate the seismic hazard (Abrahamson et al., 2019); however, the data necessary for constructing non-ergodic ground motion models are limited mainly to small and moderate magnitude earthquakes.

As the understanding of earthquake phenomena and the body of observational earthquake ground motion data have increased, it has become clear that site-specific ground motions are highly dependent on the detailed physical characteristics of any particular earthquake (Figure 1). These include the specific manner in which the fault ruptures (source term), how the propagating seismic waves radiate through the heterogeneous earth (path effect), and how the seismic waves are modified and interact with the near-surface, low-wave speed soils and the infrastructure system at a particular site (site response and soil–structure interaction). In light of these site-specific complexities, it would be desirable to have physics-based computational models that are capable of realistically simulating the underlying processes, including capturing the spatial variability and site specificity of ground motions, and representing the complex interactions between ground motion waveforms and structure/soil systems.

Because of the inherently complex nature of the problem, there will always be knowledge gaps related to uncertainties in specific features of the earthquake processes. For example, the mechanics and evolution of fault rupture for a given future earthquake or the influence of fine-scale geologic heterogeneities on ground motion waveforms. However, having a realistic and computationally tractable simulation capability will allow simulation-based exploration of the full parameter space influencing site-specific motions.

Figure 1. Fault-to-structure earthquake processes.

The ability to create representative synthetic ground motions and subsequent infrastructure response through computer simulations can vastly increase the number of ground motion records and augment the limited number of observations of near-fault ground motions from real large earthquake records. In the work described herein, recent developments in creating an end-to-end, fault-to-structure simulation framework are presented. This includes the development of a computationally optimized, high-order of accuracy, explicit finite-difference code for simulating ground motions using kinematic earthquake rupture models, and the subsequent coupling of the geophysics simulations to implicit, nonlinear finite element representations of infrastructure system response as schematically shown in Figure 2.

There has been significant interest from both the engineering and earth science communities in advancing large-scale simulations of strong ground motions and associated infrastructure response. McCallen and Larsen (2003) explored the development of regional-scale computational models for integrated ground motion and building response simulations in studying the effects of large underground explosions. Major advancements in regional-scale simulations for earthquake ground motions include the development of the CyberShake community modeling environment (Deelman et al., 2006; Graves et al. 2011), the development of the Hercules finite element framework and workflow (Taborda and Bielak, 2011), and advancements in multiscale modeling approaches (Hori and Ichimura, 2008; Ichimura et al., 2007). The extreme computational demands of regional earthquake simulations have necessitated special attention to the creation of efficient computational schemas and workflows (Cui et al., 2013; Deelman et al., 2006; Tu et al., 2006). In addition, community-based, large regional-scale simulations have been successfully used as a vehicle for confidence building and verification for selected frameworks (Bielak et al., 2010; Olsen et al., 2009).

More recently, the extension of regional simulations to include coupling of earth and infrastructure systems has been considered (Bijelic et al., 2019; Isbiloglu et al., 2015; Taborda and Bielak, 2011; Zhang et al., 2019), and integrated analyses are now being

Figure 2. Integrated fault-to-structure, regional-scale, multidisciplinary simulations of earthquake processes: (a) overall computational workflow; (b) weak and strong geophysics–infrastructure simulation coupling.

applied to regional simulations (Sahin et al., 2016). However, the frequency resolution required for evaluations of infrastructure response has made direct deterministic simulations at frequencies of engineering relevance elusive. To realize the full potential of regional-scale simulations for infrastructure risk assessments, the importance of resolving higher frequencies relevant to engineered systems has been recognized as an essential objective (Baker et al., 2014).

The principal objectives of the developments described herein are to extend physics-based, deterministic, regional-scale ground motion simulations from the historically realized ~ 1 to 2 Hz resolution (Aagaard et al., 2008, 2010; Harmsen et al., 2008; Stidham et al., 1999), to frequencies relevant to a breadth of engineered structures, and to appropriately couple geophysics and engineering simulations. This drives an ultimate goal of deterministically simulating regional-scale earthquake processes with a frequency resolution of the order of 10 Hz. The focus of the current effort is aimed at creating a fault-to-structure framework and workflow in preparation for the US Department of Energy's (DOE) next generation of graphics processing unit (GPU)-based exascale supercomputers.

To achieve high-resolution simulations, two parallel advancements are required. First, the ability to execute regional-scale earthquake scenario ground motion simulations resolving much higher frequencies with much faster speed in terms of computer wall clock time. Given the fact that the computational effort in simulating three-dimensional (3D) ground motions increases with the fourth power of the frequency resolved (i.e. a 4-Hz simulation nominally requires 16 times the computational effort of a 2-Hz simulation), this presents a

steep computational hill to climb. Second, the knowledge of subsurface geologic heterogeneities will become increasingly insufficient as the ability to compute to higher frequencies advances. Several studies have investigated the application of 3D seismic waveform tomography for improving existing regional-scale velocity models (Lee and Chen, 2016; Lee et al., 2014; Shaw et al., 2015; Tape et al., 2009). Strong ground motion simulations can benefit from such improvements (Graves et al., 2011); however, these models are poorly constrained for small-scale lengths which can impact the quality of high-frequency simulations. Validation analyses of the Southern California Earthquake Center (SCEC) community velocity models (CVMs) through modeling of recorded earthquakes; for example, Taborda and Bielak (2014) demonstrate that due to the absence of small-scale structural complexities, the performance of the SCEC CVMs tends to decrease at frequencies higher than 1 Hz.

To address this fundamental issue in the Earthquake Simulation (EQSIM) framework, two approaches are being integrated: first, full waveform inversion (FWI) algorithms are being developed in the computational framework to utilize measured ground motions from frequently occurring small earthquakes, or potentially anthropogenic seismic excitations, to provide geologic constraints and geologic model improvements; second, a computational schema for defining and distributing stochastic representations of fine-scale geologic heterogeneities at wavelengths where inversions are unlikely to provide constraints has been developed and implemented (Graves and Pitarka, 2016; Hartzell et al., 2010; Impertori and Mai, 2013).

Ultimately, the objective of the work described herein is to develop an advanced simulation framework for high-performance platforms that, when coupled with emerging exascale computers, eliminates the current significant computational barriers to simulation-based earthquake hazard and risk assessments. In this article the EQSIM solution algorithms and the underlying high-performance massively parallel numerical implementation strategies that enable efficient, high-frequency simulations are described. In a companion paper (McCallen et al., in press), the EQSIM framework is demonstrated for regional-scale simulations of ground motion and building response.

The promise of exascale computer platforms for applications in earthquake science and engineering

The performance evolution of the world's leading scientific computer platforms is illustrated in Figure 3, where a rich history of progressive innovations in machine architectures and associated computational ecosystems have yielded continuous improvement in platform performance over the past 25 years. Currently, the United States is focusing on the development of exascale platforms (i.e. 1×10^{18} flops) through the US DOE Exascale Computing Initiative (ECI). The ECI is aimed at accelerating the delivery of an exascale computing ecosystem that delivers on the order of 50 times more computational science and data analytic application power than available on today's most advanced high-performance computing (HPC) platforms (Alexander et al., 2020). The Exascale Computing Project (ECP) (<https://www.exascaleproject.org>) is a major element of the DOE's exascale initiative with three significant, closely integrated components as follows:

- Hardware and integration—supporting vendor and US DOE National Laboratory hardware R&D activities required to develop node and system designs for at least two exascale systems with diverse architectural features;

Figure 3. The advancement of HPC platforms: world's annual top supercomputer (by LINPACK performance from top500.org) through June 2019 and projection to the first Exaflop platforms. US Department of Energy HPC platforms utilized in the current study are indicated (Summit and Cori).

- Application development—creating physics-based applications for delivering science-based software ready to exploit exascale platforms in problem domains that demand and can fully exploit exascale computing;
- Software technology—developing a comprehensive and coherent software stack that will enable application developers to productively write highly parallel applications with portability across diverse exascale machine architectures.

The work reported on herein is an ECP application development, termed *EQSIM*, focused on creating the computational framework necessary to take full advantage of exascale platforms for earthquake simulations. As described in this article, the intent of the *EQSIM* application development is to exploit exascale platforms to execute physics-based deterministic numerical simulations at frequencies of engineering relevance. The prospect of exascale machines can be transformational for earthquake simulations. Whereas historical regional-scale simulations of earthquake motions have been limited to frequency resolutions of 1–2 Hz, recent frequency resolutions are increasing and exascale platforms offer the promise of ground motion simulations of up to 10 Hz or more, covering the range of interest of many engineered systems. In addition, the throughput of exascale platforms offers the promise of performing exceptionally fast earthquake scenario simulations. Rather than a heroic simulation of an earthquake rupture that may take 30–40 h or more on today's fastest scientific computers, the possibility of routinely simulating high-frequency regional-scale earthquake scenarios of the order of 3–5 h is within grasp on exascale platforms. This will enable conducting the critical sensitivity studies that are necessary to understand the variability of ground motion as a function of parameters such as the fault rupture mechanism, geologic variability, site–structure interaction, and so on.

Each DOE exascale application development project is required to define and track progress toward 5-year exascale goals in terms of both application code performance and science goals. For *EQSIM*, the performance goals are stated in terms of frequency

Figure 4. Performance goals for EQSIM regional-scale simulations: (a) San Francisco Bay Area (SFBA) regional model ($120 \times 80 \times 30 \text{ km}^3$); (b) historical regional scale simulation performance and exascale platform performance goals.

resolution and compute wall clock time for running a single earthquake rupture scenario at regional scale. The EQSIM project has adopted the San Francisco Bay Area (SFBA) as a numerical laboratory with a regional computational grid of an extent of $120 \times 80 \times 30 \text{ km}$ (Figure 4) and the exascale year 2022 objectives that included the following:

- Achieving the ability to execute 10-Hz ground motion and infrastructure response simulations on exascale platforms for the SFBA regional model;
- Achieving a single earthquake rupture scenario simulation within 3–5 h of wall clock time for the SFBA regional model.

Progress toward the exascale goals is monitored and evaluated through an application figure of merit (FOM), which for the case of EQSIM is expressed as follows:

$$FOM = \frac{(Freq)^4}{(Wall\ Clock\ Time \times 7.6)} \left(\frac{500}{V_{smin}} \right)^4 \quad (1)$$

where $Freq$ is the highest frequency resolved on the representative SFBA model; $Wall\ Clock\ Time$ is the computational time required for an $M = 7.0$ Hayward fault simulation on the specific computational platform being utilized; V_{smin} is the lowest shear-wave velocity included in the SFBA model; 7.6 is a factor that normalizes the FOM with respect to the FOM of the first regional-scale simulation performed with the initial EQSIM framework; and 500 is a constant that correlates the model minimum shear wave velocity with the 500 m/s shear wave velocity used in the initial simulations of the SFBA model. The fourth power dependency on the frequency and minimum shear wave velocity reflects the fourth-order increase in computational effort with these parameters and indicates that increasing the model resolution and lowering the model near-surface shear wave velocity to represent softer sediments increases the computational effort substantially.

With the anticipated major increase in computational performance associated with exascale computational platforms, regional-scale simulations with a $5\times$ increase in frequency resolution and $5\times$ speedup (Figure 4) represent realistic goals. Substantial progress has recently been made toward the EQSIM exascale challenge goals and the major simulation performance advancements that have been achieved to date are described in this article.

High-order earthquake ground motion simulations at regional scale

The computational demands associated with simulating earthquake ground motions at regional scale are extreme, particularly when the objective is to resolve ground motions at frequencies relevant to engineered infrastructure. The EQSIM framework has utilized and advanced SW4 (*Seismic Waves, 4th order*) a Summation-By-Parts (SBP) finite-difference program for earthquake ground motion simulation, which solves the viscoelastic wave equation with fourth-order accuracy in both space and time. To meet the EQSIM framework exascale goals, SW4 has recently undergone extensive development in terms of enhanced computational features and advanced algorithms, as well as in implementation and optimization on large, distributed memory parallel computers of various architectures.

SW4 is built upon the solid theoretical foundation provided by the SBP principle, which guides the derivation of high-order-of-accuracy finite-difference methods such that the resulting time integration scheme is provably stable, without requiring artificial dissipation or ad hoc filtering (Petersson and Sjogreen, 2015). The SBP methodology puts particular emphasis on how the difference stencils must be modified near boundaries and also prescribes how the boundary conditions must be discretized. The basic discretization of the isotropic elastic wave equation on a single Cartesian grid was initially developed for second-order accuracy (Nilsson et al., 2007), and was later generalized to fourth-order accuracy (Sjogreen and Petersson, 2012).

SW4 has a number of advanced algorithms tailored to the efficient simulation of earthquake ground motions at regional scale (Figure 5). For the EQSIM framework development, the basic SBP approach has been extended to include mesh refinement interfaces (Petersson and Sjogreen, 2010). The code utilizes a hybrid curvilinear/Cartesian grid in which the near-surface curvilinear portion of the grid undergoes a coordinate transformation to conform to the ground surface topography and represent associated wave scattering at the earth surface and to satisfy the free surface boundary conditions along realistic (nonplanar) topographies. The computational grid at depth utilizes an efficient structured Cartesian grid and adaptive grid refinement, which allows grid sizing to be optimized for problem-specific geologic properties. Adaptive grid refinement has been implemented in both the curvilinear and Cartesian grids, as indicated in Figure 6. The ability to adapt the computational grid size to optimally represent low-velocity, short-wavelength near-surface sediments and high-velocity, long-wavelength geologic structure at depth leads to a substantial reduction of the total number of grid points and also enables a longer time step to be used for regional earthquake simulations.

For regional earthquake simulations, it is necessary to truncate the computational domain by a far-field closure. For this purpose, highly accurate supergrid layers have been developed that are provably stable for both isotropic and anisotropic elastic materials with free surface boundaries. SW4 employs the supergrid far-field truncation technique to create a dissipative boundary at the edges of the computational model (Figure 5). The associated

Figure 5. Key features and recent developments in the SW4 ground motion simulation code.

Figure 6. Example SW4 computational grid including surface topography and user-specified mesh refinement in the curvilinear and Cartesian grids.

algorithm adds a “sponge layer” around the edges of the finite-domain regional model. Inside the sponge layer, the wave equation is modified by a combination of grid stretching and high-order artificial damping (Petersson and Sjogreen, 2014). The underlying concept is to utilize a fixed number of grid points to mimic a very large physical domain where reflections from the boundary would require a long transit time to reflect from the far-field boundary and return to the computational domain of interest. The greatest strength of the supergrid approach is that it gives very small reflections, and is provably numerically stable within an explicit time integration framework.

Geologic material energy dissipation is represented by incorporation of a viscoelastic rheological model that results in quality factors, Q_p and Q_s that attenuate compressional

and shear waves, respectively (Petersson and Sjogreen, 2012). The SW4 rheological model utilizes standard linear solid (SLS) elements coupled in parallel. The coefficients in each SLS are determined such that the resulting quality factors Q_p and Q_s become approximately constant as a function of frequency. These quality factors can vary from grid point to grid point over the computational domain and are read as input in the same manner as the elastic properties of the material model.

The earthquake rupture model employed in the EQSIM framework, based on the work of Graves and Pitarka (2016) and Pitarka et al. (2020), allows complex rupture models to be automatically converted into a large number of moment tensor point sources. This rupture model (GP2016) is designed to represent rupture kinematics using a semi-stochastic approach. The fault slip distribution is essentially a two-dimensional (2D) spatially correlated random field with a von Karman wavenumber spectrum. The slip field is scaled such that the earthquake target moment is achieved and the coefficient of variation of slip is 0.85. For a given prescribed hypocenter, the rupture initiation times are determined using a two-step physics-based procedure constrained by rupture dynamics modeling.

First, a background rupture speed is specified to be 80% of the local shear-wave velocity (V_s) for depths below 8 km and the speed is tapered to 56% of the local V_s at 5 km depth and remains at that value up to the ground surface. The reduction above 5 km is designed to be representative of the weaker shallow zone in surface rupturing events and is consistent with statistical analysis of a large number of events in California. The second step of the rupture evolution applies a timing perturbation to each subfault segment that scales with local slip and slows in regions of low slip. The slip-rate function is a Kostrov-like pulse with a total duration (rise time) that scales with the square root of the local slip. The average rise time is constrained to scale in a self-similar manner with seismic moment. Random perturbations are applied to the rupture time and rise time of each subfault following log-normal distributions with standard deviations of 0.2 and 0.5, respectively.

The final step in the rupture generation process is the specification of the slip direction (rake), which is allowed to vary across the fault with a standard deviation of 15° about a prescribed mean value. The spatial distribution of the random rake variations follows a similar von Karman correlation function as that used for the slip variations (Figure 7). GP2016 has been validated through comparisons of simulated broad band ground motions against ground motion prediction equations (GMPEs), as well as direct comparisons with a large number of crustal earthquakes in California and Japan (Graves and Pitarka, 2010, 2016; Pitarka et al., 2020).

Figure 7. Kinematic rupture model for an M6.7 crustal earthquake generated with GP2016 using random perturbations: (a) fully stochastic with correlated spatial variability; (b) fully deterministic asperity-based model; and (c) GP2016 hybrid model including both large-scale asperities and stochastic spatial variability.

Existing regional-scale, 3D geologic structure models tend to be very coarse spatially and where regional geologic models exist, they can typically support deterministic ground motion simulations on the order of 1–2 Hz. For the SFBA, for example, the existing USGS 3D model provides geologic properties at 25 m vertical and 100 m horizontal voxels; however, the geologic data may not even support the resolution implied in 25 m data as it has been generated from interpolations of coarser data. For the EQSIM framework development, the two complementary approaches mentioned above for enhancing existing geologic models have been implemented. The first method relies on FWI that can be used in conjunction with measured earthquake data to perturb existing geologic models to achieve better fit between simulated and measured waveforms. The second, complementary, method includes the representation of a statistically relevant stochastic overlay of geologic material variability at fine scale.

As for the first method, the geologic material inversion program, SW4mopt (*SW4 material optimization*) has been created from the same code base as SW4, with the objective of improving geologic material models to become more representative and accurate. Existing geologic material models have mainly been validated with simulations performed at lower frequencies than what are now becoming resolvable in high-performance simulations. Furthermore, it is not always immediately clear to what frequency level a given material model is valid. Computationally, an objective of SW4mopt is to make the material inversion process as automated as possible, while existing techniques often require manual intervention at various stages in the inversion process.

The inversion problem is considerably more challenging than the forward earthquake simulation problem, both in terms of complexity of the solution algorithms and in terms of computational requirements. The solution algorithm must deal with the fact that the problem of finding the “best” material characterization from a given set of ground motion time-series data at various stations on the surface tends to be underdetermined.

The material inversion capability strives to minimize the misfit between observed seismograms and computed synthetic seismograms with respect to the material properties. SW4mopt allows inversion with simultaneous utilization of more than one event. The minimization is performed by a gradient-based search algorithm. Two different methods are implemented in SW4mopt: a nonlinear conjugate gradient method and a quasi-Newton method. The gradient of the misfit with respect to the material is computed by solving the adjoint of the elastic wave equation. The computational cost to compute the gradient is approximately equal to the cost of three forward solves per event. Among the capabilities currently available in SW4mopt (Sjogreen and Peterson, 2014) are the following:

- Selection between two different misfit functions;
- Selection between two different regularization functions added to the misfit function;
- Selection between inversion for one (p -wave velocity), two (s -wave velocity, p -wave velocity), or three (density, s -wave velocity, p -wave velocity) unknown fields;
- Utilization of an arbitrary number of earthquake events, providing greater spatial coverage of the computational domain;
- Specification of windows on the seismograms, to focus the optimizer on the most interesting portions of the data.

For a selected time-window, a two-step process is followed whereby the first step includes maximizing the cross correlation between synthetic and real data waveforms to achieve agreement in arrival times and the second step includes minimizing the L2 norm of the difference. This approach addresses both phase and amplitude differences between synthetic and measured waveforms. The usage of SW4mopt is very similar to the usage of SW4. The user provides an input file in the same format as an SW4 input file, specifying the computational domain and finite-difference grid. The stations where observed data are available are specified by a new command (“observation”). In the case of multiple events, each source and each observation station have to be provided with a label indicating to which event it belongs. Windows on the seismograms can also be set directly in the input file. Before running SW4mopt, it is necessary to prefilter the observed data to a frequency range that can be resolved on the computational grid. SW4mopt iterates until the norm of the gradient of the misfit is less than a prescribed tolerance, or until a maximum number of iterations is reached. Given the large size of data sets (events, windows, and waveforms) an automated approach has been developed to minimize user intervention.

SW4mopt is a separate executable which, to a large extent, uses the same source code as SW4, with a few SW4mopt-specific functions in the SW4 source directory (e.g. the adjoint solver). The forward solver in SW4mopt uses identical source code as the solver in SW4. This unified code base simplifies the development of SW4mopt, and ensures that all improvements made to SW4 will also be directly available in SW4mopt. To date, work has been focused on initial investigations of improving the USGS geologic model for the SFBA model. The code capabilities have been extensively verified on synthetic test problems, and preliminary inversions with real earthquake data have been made, with encouraging results; however, more comprehensive testing and algorithm tuning is underway.

Small-scale variations in shallow crustal geology create wave scattering over a broad range of frequencies. The second method for addressing fine-scale geologic structure which has been implemented directly in the EQSIM framework consists of adding stochastic variability to the underlying 3D geologic model, for example, the SFBA USGS geologic model (<https://earthquake.usgs.gov/data/3dgeologic>). A statistical approach has been implemented for creating fine-scale model variability to produce wave scattering at frequencies greater than approximately 1 Hz. In this approach, 3D correlated statistical perturbations are applied to the entire background seismic velocity structure. The perturbation field is constructed in the wavenumber domain by first generating a 3D complex array of Gaussian random numbers having zero mean and unit variance. This array is then filtered such that its power spectrum follows a von Karman correlation function given as follows:

$$P_V(k) \propto [1 + k_r^2]^{-(H+D/2)} \quad (2)$$

where $k_r = \left(a_x^2 k_x^2 + a_y^2 k_y^2 + a_z^2 k_z^2\right)^{1/2}$ with k_x , k_y , and k_z representing the wavenumbers in the three spatial directions, a_x , a_y , and a_z are the spatial correlation lengths in these three directions, D is the Euclidean dimension (set to 3 for a 3D medium) and H is the Hurst exponent, which is set to values between 0 and 0.5. After transforming the perturbation field to the spatial domain and retaining only the real portion, the amplitudes of the perturbations are scaled to match a target standard deviation with a mean value of zero. The full 3D seismic velocity model is then generated by adding the perturbed values of V_p and V_s to the V_p and V_s values taken from the background 3D velocity structure at each grid point.

Figure 8. Resulting generated stochastic variability in subsurface geology: (a) regional SFBA model and surface geology; (b) geologic properties provided by the existing USGS 3D model; and (c) USGS model with stochastic geology perturbations.

With this procedure, the large-scale characteristics of the baseline velocity model are honored and preserved, while adding correlated stochastic velocity perturbations. The typical correlation lengths, in the horizontal directions, a_x and a_y , are of the order of 1 km, and correlation in the vertical direction is taken as one-fourth of that in the horizontal direction. The anisotropy of the correlation lengths with depth for shallow crustal layers is the subject of current research. A graphical representation of the resulting velocity model for the SFBA is illustrated in Figure 8.

Coupled geophysics and engineering simulations

In defining the interaction between ground motion and structural response, two forms of coupling between the ground motion and structural system are considered, each with an associated computational workflow, as indicated in Figure 9. In the weak-coupling option, any soil–structure interaction is assumed to have a negligible effect and the simulated ground motions at a point on the earth surface are applied directly to the structural system in a fixed-base structural analysis. In the strong-coupling option, the geophysics and soil/structural system models are directly coupled through the domain reduction method (DRM) developed by Bielak et al. (2003) and Yoshimura et al. (2003) to allow representation of soil–structure interaction and the consideration of a complex, fully 3D incident wavefield impinging on the soil–structure system.

In the case of weak coupling, simulated earthquake motions at the ground surface are written to a database and accessed for structural response simulations. In the classical case in which ground rotations are neglected (Figure 9a), the horizontal and vertical ground motions from a point on the earth surface are assumed to be uniform across the footprint of the infrastructure system, which is consistent with an assumption of the incident

Figure 9. Coupling between geophysics and infrastructure models. Weak coupling: (a) ground motion horizontal and vertical motions only; (b) ground motion horizontal and vertical motions plus a ground surface rotation associated with a gauge length equal to the building foundation width. Strong coupling: (c) input of the complex 3D incident wavefield from the geophysics model to the soil-structure system through the DRM.

wavefield consisting of pure vertically propagating compressional and shear waves. Rotations of the ground surface are very challenging to measure in the field and few reliable observations exist; however, various studies with simplified models have indicated that ground rotations can potentially contribute significantly to structural response and structural demands (Trifunac, 2009a, 2009b). With a high-fidelity geophysics model, numerical estimates of ground rotations can be obtained and applied to a fixed base structure in addition to the ground surface translations (Figure 9b), and that option is included in the EQSIM framework.

In the weakly coupled case (first workflow), a building model is subjected to the selected components of ground motion at each ground surface “station” in the geophysics model in a fixed-base building analysis. For each building simulation, selected structural response measures over the duration of the earthquake are saved and written to a file. This process is repeated for every ground station throughout the domain to develop a regional map of demand/risk for the particular building class being evaluated.

To support realistic earthquake response simulations for a breadth of geologic and geotechnical conditions, the ability to represent potential nonlinear behavior in near-surface sedimentary layers is an essential capability that is currently undergoing development in the EQSIM framework through extended application of the DRM to couple the geophysics and engineering computational domains. An extended horizontal DRM boundary is being implemented between geophysics and engineering models to allow the implementation of nonlinear models for near-surface soils.

For the fault-to-structure simulations performed in the current study, representative steel moment frame buildings of 3, 9, 20, and 40 story heights were utilized as shown in

Figure 10. Representative steel moment frame buildings in the infrastructure model library; 40-, 20-, 9-, and 3-story building models, Nevada finite-deformation, and co-rotational fiber-based nonlinear beam element for elasto-plastic materials.

Figure 10 (Astaneh-Asl, 2018). Detailed nonlinear building models have been developed and carefully tested through extensive code-to-code comparisons (Wu et al., 2018). The models adopted detailed, fiber cross-section section models of the wide flange sections and were developed for the special-purpose Nevada finite element program as shown in Figure 10 (McCallen and Larsen, 2003; Miah et al., 2018; Petrone et al., 2016). The Nevada program was written specifically for efficient nonlinear seismic analysis of building systems and its simplicity and self-contained architecture make it ideally suited to expedient implementation on massively parallel computer platforms with various node and core architectures. The steel frame models utilize a classical elastoplastic material representation with kinematic hardening and the beam elements employ continuously updated co-rotational coordinates to explicitly track the displacements and deformed shape at each instant in time. The end nodes of each beam element have a local coordinate system in addition to an element co-rotational coordinate system associated with the beam element overall motion. This approach allows incremental updates of the element end rotations for large-displacement problems where rotations cannot be combined vectorially. The beam element employs three-point Lobatto integration with integration points at the end of each beam element so that initial yielding near the joint is captured. These features of the fiber model ensure that interaction between stress resultants are explicitly accounted for in the determination of plasticity evolution, and the building geometry is explicitly represented in space and time so that p -delta effects are rigorously captured for tall buildings.

The Nevada code utilizes an implicit incremental/iteration time stepping scheme based on Newmark-Beta time integration whereby, within each time step, a Newton-type

Figure 11. Computational workflow for fault-to-structure simulations: (a) for weakly coupled simulations, surface motions from the regional simulation are saved in an HDF5 data container and applied directly to fixed-base structural models; (b) for the strongly coupled simulations, grid point motions from user specified 3D “blocks” of the SW4 wave propagation model are automatically interpolated to the boundary of the local structure/soil island model.

iteration scheme is employed with equilibrium iterations to drive the L2 norm of the vector of nodal residual forces to within a convergence tolerance of zero (McCallen and Larsen, 2003). For the formation of the tangent stiffness in the Newton iterations, both full and modified Newton options are available with associated tradeoffs between convergence speed and convergence reliability. In this study, tens of thousands of nonlinear building earthquake response histories were computed, which provide insight on nonlinear building model convergence.

In the strong coupling case (second workflow; Figure 9c), the earthquake ground motions from the SW4 run are saved and stored for user-specified 3D SW4 subdomain volumes that encompass a soil island and structure of interest. The motions at each grid point of the SW4 subdomain are then interpolated to the soil island boundary to provide input motions for a soil–structure interaction simulation. The interface between the geophysics and infrastructure models is represented through the DRM that essentially transfers time-dependent interface tractions between the global and local models. This approach allows for the explicit inclusion of soil–structure interaction, and permits full spatially and temporally varying input motions that are consistent with the incident seismic waves from SW4. This option has been implemented by developing the workflow between SW4 and the ESSI (http://sokocalo.engr.ucdavis.edu/~jeremic/Real_ESSI_Simulator/) and OpenSees (<http://opensees.berkeley.edu>) codes for simulation of a single specific structure.

The weak- and strong-coupling options spawn two alternative computational workflows for the massively parallel simulations in the EQSIM framework as indicated in Figure 11. In the weak-coupling case surface motions, including horizontal and vertical

ground translations plus a potential ground rotation, are saved for all locations on the earth surface of the model domain for subsequent use in structural response simulations. In the strong-coupling case, motions at the grid points of a 3D subdomain are saved for interpolation to the DRM boundary of a local structure/soil island model. In each case a standard Hierarchical Data Format—HDF5 (<https://www.hdfgroup.org>) data container is written from the SW4 code to store either the ground surface motions (for weak coupling) or the motions from a 3D near-surface volume (for strong coupling) for structural evaluations. This simple integrated approach to code coupling is developed in a general manner so that ground motions from SW4 can be readily coupled to any existing finite element code with the necessary capabilities for traditional fixed-base analysis or soil–structure interaction through the DRM.

Computational strategies and workflow at hundreds of billions of grid zones

The successful execution of regional-scale simulations that can resolve high frequency ground motions demands significant computational effort and an extreme number of zones in the discretized model. For the SFBA model, using the computational domain indicated in Figure 4, an SW4 finite-difference grid that can resolve 10-Hz motions requires over 200 billion grid points in the computational domain. The execution of such a model can only be accomplished through the efficient and effective implementation of massive parallelism in the calculations, which required development of strategies for efficiently handling massive data input and output (I/O), efficient creation of the computational model grid definition and run control information for 200 billion grid points, and practical runtime management of the model execution. Runtime considerations include appropriate use of burst-buffer memory for on-the-fly creation of restart (checkpoint) files to guard against run upsets such as potential core failures in simulations using upwards of 500,000 cores that require a restart of the analysis.

In the EQSIM framework, to distribute effort associated with regional-scale computational tasks on massively parallel computers, computations associated with the SW4 geophysics model grid are distributed across thousands of nodes on a parallel computer platform by subdividing the overall model domain into vertical pencil-shaped subdomains which extend from the earth surface to the bottom of the computational domain (Figure 12). Each pencil is distributed to separate machine nodes using a parallel load-balancing algorithm that decomposes the computational tasks that are distributed across the parallel computer platform. The associated message passing interface (MPI) implementation must be optimized for the particular node and core architecture associated with any specific computer platform. For building simulations associated with weakly coupled simulations, once a specific building model is selected from the existing building library, the building simulation for each ground motion station of the regional scale model is distributed to a single core on a massively parallel platform as shown in Figure 12. This allows very simple parallelization for thousands to hundreds of thousands of building simulations, and provides for very effective scaling, that is, 100,000 building simulations can be executed within nearly the same wall clock time as 1000 building simulations, simply by scaling the number of cores utilized by the number of building simulations desired.

Contrary to traditional simulation approaches, the SW4 grid and 3D material input file for a specific simulation is not created and stored prior to problem run time due to the prohibitive size of a full file, rather the input file is created on the fly at runtime as the

Figure 12. Partition of computations on distributed computer platforms: (a) distribution of the regional geophysics model using node-dependent MPI on massively parallel platforms; (b) distribution of building simulations in simple core-dependent parallelization.

front-end of the run execution. This is achieved through an in-situ parallel mesh generator that constructs the computational grid at runtime based on user supplied problem topographic elevation of the earth surface, a 3D geologic model, and prescribed depths of the mesh coarsening interfaces (see Figure 6). This precludes the challenges of storing massive, prohibitively large data files. For creating the SW4 input file at runtime, geologic data from a 3D database is stored in an efficient, newly developed binary *S file* in the HDF5 format. At runtime, the SW4 code reads the user supplied problem control information, the definition of the specific fault rupture scenario and the geologic material data file prior

Figure 13. Three major blocks of computational workflow and strategies for massively parallel fault-to-structure regional-scale simulations.

to initiating the time stepping starting with the earthquake rupture and subsequent period of seismic wave propagation as indicated in Figure 13.

As advanced computer system node architectures evolve with vendors striving to deliver increasing levels of performance gain, the porting and transformation of high-performance codes to run efficiently on a multiplicity of platforms can become time-consuming and difficult. One of the objectives of the ECP project is to develop alternative architectures for ECP platforms to provide multiple technology paths for future cutting-edge machines

Table 1. Progression of EQSIM ground motion simulations with SW4

Benchmark simulation (platform)	Code attributes	Frequency resolution (Hz)	Number of compute nodes	Wall clock time (h)	Figure of merit
A (Cori)	Initial run of SW4 ported to Cori	3.67	2048	23.9	1.0
B (Cori)	SW4 with optimized hybrid MPI/OpenMP loops	4.17	6528	12.0	3.32
C (Cori)	SW4 with Cartesian mesh refinement	4.17	4000	6.0	6.63
D (Cori)	SW4 using all of the Cori computer	5.0	8192 (all of Cori)	9.2	8.95
E (Summit)	Initial run of SW4 ported to the Summit computer	10.0	1200 (1/4 of Summit)	19.9	66.2
F (Summit)	Most recent run of SW4 including enhance I/O, curvilinear, and Cartesian mesh refinement	10.0	1024 (<1/4 of Summit)	6.9	189

SW4: Seismic Waves, 4th order.

(<https://www.exascaleproject.org/research-group/pathforward/>). To prepare applications for changes and alternatives in hardware architecture, ECP software stack developments are creating high-level software that encapsulates platform-specific code to enable applications to be portable across multiple platforms with a minimum of source code disruption. For the EQSIM framework, the RAJA C++ libraries (<https://github.com/LLNL/RAJA>) have been implemented in the SW4 source code to simplify the process and level of effort associated with porting to advanced architectures such as GPU-based systems. This has yielded rapid porting to advanced platforms as the EQSIM development progresses, and is an essential foundation for transition to emerging exascale platforms (Pankajakshan et al., 2019).

Regional-scale simulations and framework performance

In support of the preparation for exascale platforms, the EQSIM framework has undergone progressive advancements that have included the development and implementation of advanced algorithms, optimization of large volume I/O, and preparation of efficient workflow for code execution and high performance on massively parallel computers. These advancements have resulted in significant performance increases and the ability to achieve high-frequency simulations on today's leading platforms. The major performance enhancements to date are summarized in Figure 14 and Table 1. Each curve in Figure 14a illustrates the performance curve of the specific version of SW4, with the features available at the time of the performance evaluation and with a computational effort varying with the fourth power of the frequency being resolved. Starting with the initial ground motion simulation runs on Cori in January 2017, with a realized frequency resolution with the SFBA model of 3.67 Hz and an application FOM of 1.0 (point "A" in Figure 14), optimization of SW4 for the Cori architecture and the implementation of adaptive mesh refinement in the Cartesian portion of the grid resulted in a significant performance increase to a frequency resolution of 5.0 Hz and an application FOM of 8.95 (point "D" in Figure 14).

Next, a very large performance boost was achieved by appropriate porting of the SW4 program to the Summit computer at Oak Ridge National Laboratory, the currently

Figure 14. Successive EQSIM performance increases during the EQSIM project: (a) progress toward exascale goals (benchmarks described in Table 1); (b) increase in the FOM over time.

number one ranked scientific computer in the world (Figure 3). Rapid porting and successful execution on Summit were accomplished through utilization of Raja libraries that support seamless transition to GPU-based architectures like those on Summit. The first science demonstration runs on Summit, which utilized one-quarter of Summit, resulted in a major performance leap: SFBA regional model simulations to 10-Hz resolution and a major increase in FOM to 66.2 (point “E” in Figure 14). Finally, for the most recent performance testing, I/O enhancements were made and efficient mesh refinement was implemented for both the curvilinear and Cartesian portions of the grid, providing another large performance boost (point “F” in Figure 14).

The advancements over the past 3 years achieved major increases in computational performance with EQSIM. This represents a major step toward overall exascale goals and the results have indicated the potential step function performance increases that can be achieved with transition to state-of-the-art computer platforms. In terms of future performance increase, all of the key algorithmic improvements in SW4 are nearing completion and future attention will be focused on the preparation for exploitation of exascale platforms and the anticipated substantial computational performance increases.

An example of the successful execution of a 10 Hz simulation for an $M = 7.0$ Hayward Fault event in the SFBA is illustrated in Figures 15–19 with problem simulation parameters shown in Table 2. In this simulation, a fault rupture containing defined asperities was created to represent the Hayward fault rupture scenario. Once the ground surface motions were obtained throughout the domain, representative planar building models were analyzed at each ground surface site for fault-normal and fault-parallel directions, respectively. The peak interstory drift contour plots, color coded in accordance with the American Society of Civil Engineers building limit states defined in standard ASCE 43-05 (2005), are plotted for 3-, 9-, 20-, and 40-story buildings, respectively.

The building drift plots shown in Figures 16–19 were developed based on 76,800 non-linear time history building simulations utilizing the weak coupling option (i.e. a traditional fixed-based building simulation), with each simulation executed for 90 s of earthquake motion. The simulation results illustrate a complex distribution of peak

Table 2. Parameters of 10-Hz SFBA ground motion simulations on the Summit computer

	Benchmark E	Benchmark F
Frequency resolved	10 Hz	10 Hz
V_{smin}	500 m/s	500 m/s
Number of grid points	203 billion	63 billion
Smallest cell size	6.25 m	6.25 m
Time step size	7.119e-4	8.491e-4
Total time steps	126,430	106,000
Platform	Summit	Summit
Number of compute nodes	1200 (1/4 of Summit)	1024
Wall clock time	19 h 52 min	6 h 58 min

building drift as a function of building size and geographic location. For all buildings, there is a near-fault region with significant building demand and damage, and drifts associated with significant inelastic action are widespread. For the 9- and 20-story buildings, localized areas of extensive inelastic action associated with ASCE 43-05 Limit State A (large permanent distortion) are evident as indicated by the red zones shown in Figures 17 and 18. For all four buildings, the simulation results indicate that the fault-parallel motion tends to create a somewhat wider zone of building damage compared with the fault-normal component of motion. This SFBA simulation utilized a minimum shear-wave velocity cutoff of 500 m/s. Parametric studies have indicated the desirability of resolving even lower shear-wave velocities of the order of 250 m/s (or lower) in order to appropriately characterize the risk to buildings located on class D soil sites (Rodgers et al., 2019a). However, the simulation effort associated with V_{smin} of 250 m/s is not yet achievable even on the Summit platform for long-duration 10 Hz regional-scale simulations and will have to await the build-out of exaflop machines.

Discussion and summary

Significant progress has been made in developing and implementing the computational models and integrated workflow necessary for effective massively parallel fault-to-structure earthquake simulations. In alignment with the goals of the DOE Exascale Computing Initiative, the work on developing the EQSIM framework is laying the foundation for utilizing emerging DOE exascale platforms that will be arriving in the next 2–3 years (<https://www.energy.gov/articles/us-department-energy-and-intel-build-first-exascale-supercomputer>). The computational advancements that have been achieved to date, in terms of frequency resolution and computational effort in a regional-scale model of the SFBA, are summarized in Figure 20. The historical simulations shown were computed with somewhat different physical domain sizes and problem parameters; nevertheless, the rapid increase in frequency resolution with the EQSIM framework is indicative of the improved performance achieved with leading massively parallel computer platforms and computational ecosystems. The computational effort for a regional-scale simulation is proportional to the volume of the model, the duration of the earthquake simulation, and the ratio of the maximum resolved frequency to the model minimum shear wave speed raised to the fourth power. In Figure 20b, the computational effort associated with each of the previous SFBA simulations is shown, computed using the frequency resolution and minimum shear wave speed of each simulation. The computational effort is normalized to a uniform regional volume and earthquake simulation duration, and anchored at 1.0 for the earliest simulation (Stidham et al., 1999). Figure 20 illustrates the

Figure 15. Time snapshots of $M = 7$ Hayward fault rupture simulation: (a) wave front at 7.14 s into the earthquake; (b) wave front at 14.27 s into the earthquake.

Figure 16. A 3-story steel frame building peak interstory drift: (a) building subjected to fault-normal ground motions; (b) building subjected to fault-parallel ground motions.

Figure 17. A 9-story steel frame building peak interstory drift: (a) building subjected to fault-normal ground motions; (b) building subjected to fault-parallel ground motions.

Figure 18. A 20-story steel frame building peak interstory drift: (a) building subjected to fault-normal ground motions; (b) building subjected to fault-parallel ground motions.

dramatic increase in computational effort that has been achieved with recent simulation advancements.

Figure 19. A 40-story steel frame building peak interstory drift: (a) building subjected to fault-normal ground motions; (b) building subjected to fault-parallel ground motions.

Figure 20. Advancements in SFBA regional scale model frequency resolution and computational effort: (a) increase in frequency resolution of regional-scale models; (b) increase in computational effort, EQSIM advancements are represented in blue.

As leadership computer platforms continue to advance and the ability to compute earthquake phenomenon continues to rapidly increase, it is essential to critically assess the realism of the simulations in multiple dimensions. Previous work has evaluated the realism of ground motion simulations to 5 Hz by comparison with existing GMPEs (Rodgers et al., 2019b; Rodgers et al., 2018a), and recent work illustrates SFBA regional simulations to 10 Hz (Rodgers et al., 2020). A companion paper provides the first assessments of the

infrastructure response with fault-to-structure simulations utilizing the EQSIM framework (McCallen et al., in press). Rigorous computational validation of the 3-D velocity structure and the resulting path effects on the ground motion predictions are also underway through the simulation of actual small earthquakes and comparison with measured waveforms (Rodgers et al., 2018b). Small events are particularly well suited to understanding the ability to capture path and site effects as a result of the simplicity of the source functions.

The overarching focus of the EQSIM development is one of computational enablement and the removal of the computational limitations and barriers to executing regional-scale simulations. The full potential of regional-scale simulations for practical hazard and risk assessments will only be realized when such simulations can be computed routinely for many realizations as opposed to historical heroic, one-off simulations. The work to date demonstrates the promise of achieving breakthrough simulations on emerging exascale platforms, and providing the computational toolkit necessary to explore complex multidisciplinary interactions between geophysical phenomenon and engineered systems.

Acknowledgements

The careful review and comments provided by Dr Ricardo Taborda and an anonymous reviewer improved the content of this manuscript, and they are gratefully acknowledged by the authors.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the Exascale Computing Project (ECP; Project No. 17-SC-20-SC), a collaborative effort of two US Department of Energy (DOE) organizations—the Office of Science and the National Nuclear Security Administration. The work was performed at the Lawrence Berkeley National Laboratory and the Lawrence Livermore National Laboratory under DOE Contract No. DE-AC52-07NA27344. Computer access and exceptional support from the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Lab and the Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Lab are gratefully acknowledged.

ORCID iD

Arben Pitarka <https://orcid.org/0000-0001-6026-8859>

References

- Aagaard B, Brocher T, Dolenc D, Dreger D, Graves R, Harmsen S, Hartzell S, Larsen S, McCandless K, Nilsson S, Petersson NA, Rodgers A, Sjogreen B and Zoback ML (2008) Ground motion estimates for the 1906 earthquake and scenario events. *Bulletin of the Seismological Society of America* 98: 1012–1046.
- Aagaard B, Graves R, Rodgers A, Brocher T, Simpson R, Dreger D, Petersson NA, Larsen S, Ma S and Jachens R (2010) Ground motion modeling of the Hayward fault scenario earthquakes, Part II: Simulation of long-period and broadband ground motions. *Bulletin of the Seismological Society of America* 100: 2945–2977.

- Abrahamson N, Kuehn N, Walling M and Landwehr N (2019) Probabilistic seismic hazard analysis in California using nonergodic ground motion models. *Bulletin of the Seismological Society of America* 109: 1235–1249.
- Alexander F, Almgren A, Bell J, Bhattacharjee A, Chen J, Colella P, Daniel D, DeSlippe J, Diachin L, Draeger E, Dubey A, Dunning T, Evans T, Foster I, Francois M, Germann T, Gordon M, Habib S, Halappanavar M, Hamilton S, Hart W, Huang Z, Hungerford A, Kasen D, Kent PRC, Kolev T, Kothe DB, Kronfeld A, Luo Y, Mackenzie P, McCallen D, Messer B, Mniszewski S, Oehmen C, Perazzo A, Perez D, Richards D, Rider WJ, Rieben R, Roche K, Siegel A, Sprague M, Steefel C, Stevens R, Syamlal M, Taylor M, Turner J, Vay JL, Voter AF, Windus TL and Yelick K (2020) Exascale applications: Skin in the game. *Philosophical Transactions of Royal Society A*. Epub ahead of print 20 January. DOI: 10.1098/rsta.2019.0056.
- American Society of Civil Engineers (ASCE) (2005) *Seismic Design Criteria for Structures, Systems and Components in Nuclear Facilities*. Reston, VA: ASCE/SEI 43-05.
- Anderson JG and Brune JN (1999) Probabilistic seismic hazard analysis without the ergodic assumption. *Seismological Research Letters* 70: 19–28.
- Astaneh-Asl A (2018) *Designs of 3-, 9-, 20- and 40-Story Frames*. Technical report to the Lawrence Berkeley National Laboratory, Berkeley, CA.
- Baker JW, Luco N, Abrahamson NA, Graves RW, Maechling PJ and Olsen KB (2014) Engineering uses of physics-based ground motion simulations. In: *Proceedings of the Tenth U.S. national conference on earthquake engineering*, Anchorage, Alaska, 21–25 July.
- Bielak J, Graves R, Olsen K, Taborda R, Ramirez-Guzman L, Day S, Ely G, Roten D, Jordan T, Maechling P, Urbanic J, Cui Y and Juve G (2010) The shakeout earthquake scenario: Verification of three simulation sets. *Geophysical Journal International* 180: 375–404.
- Bielak J, Loukakis K, Hisada Y and 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: 817–824.
- Bijelic N, Lin T and Deierlein G (2019) Evaluation of building collapse risk and drift demands by nonlinear structural analyses using conventional hazard analysis versus direct simulation with CyberShake seismograms. *Bulletin of the Seismological Society of America* 109: 1812–1828.
- Cui Y, Poyraz E, Olsen K, Zhou J, Withers K, Callaghan S, Larkin J, Guest C, Choi D, Chourasia A, Shi Z, Day S, Maechling J and Jordan T (2013) Physics-based seismic hazard analysis on petascale heterogeneous supercomputers. In: *Proceedings of the international conference on high performance computing, networking, storage and analysis*, Denver, CO, 17–22 November.
- Deelman E, Callaghan S, Field E, Francoeur H, Graves R, Gupta N, Gupta V, Jordan T, Kesselman C, Maechling P, Mehlinger J, Mehta G, Okaya D, Vahi K and Zhao L (2006) Managing large-scale workflow execution from resource provisioning to provenance tracking: The CyberShake example. In: *Proceedings of the second IEEE international conference on E-science and grid computing*, Amsterdam, The Netherlands, 4–6 December.
- Deierlein GG, Reinhorn AM and Willford MR (2010) Nonlinear structural analysis for seismic design. NEHRP Seismic Design Technical Brief No. 4, NIST GRC 10-917-5, 15 October.
- Graves R, Jordan T, Callaghan S, Deelman E, Field E, Juve G, Kesselman C, Maechling P, Mehta G, Milner K, Okaya D, Small P and Vahi K (2011) Cybershake: A physics-based seismic hazard model for southern California. *Pure and Applied Geophysics* 168: 367–381.
- Graves R and Pitarka A (2010) Broadband ground motion simulation using a hybrid approach. *Bulletin of the Seismological Society of America* 100: 2095–2123.
- Graves R and Pitarka A (2016) Kinematic ground motion simulations on rough faults including effects of 3D stochastic velocity perturbations. *Bulletin of the Seismological Society of America* 106: 2136–2153.
- Harmsen S, Hartzell S and Liu P (2008) Simulated ground motions in Santa Clara Valley, California and vicinity for $M > 6.7$ scenario earthquakes. *Bulletin of the Seismological Society of America* 98: 1243–1271.
- Hartzell S, Harmsen S and Frankel A (2010) Effects of 3D random correlated velocity perturbation on predicted ground motion. *Bulletin of the Seismological Society of America* 100: 1415–1426.

- Hori M and Ichimura T (2008) Current state of integrated earthquake simulation for earthquake hazard and disaster. *Journal of Seismology* 12: 307–321.
- Ichimura I, Hori M and Kuwamoto H (2007) Earthquake motion simulation with multiscale finite-element analysis on hybrid grid. *Bulletin of the Seismological Society of America* 97: 1133–1143.
- Imperator W and Mai P (2013) Broad-band, near-field ground motion simulations in 3-dimensional scattering media. *Geophysical Journal International* 192: 725–744.
- Isbiloglu Y, Taborda R and Bielak J (2015) Coupled soil-structure interaction effects of building clusters during earthquakes. *Earthquake Spectra* 31: 463–500.
- Lee E and Chen P (2016) Improved basin structures in southern California obtained through full-3D Seismic Waveform Tomography (F3DT). *Seismological Research Letters* 87: 874–881.
- Lee E, Chen P and Jordan T (2014) Testing waveform predictions of 3D velocity models against two recent Los Angeles earthquakes. *Seismological Research Letters* 85: 1275–1284.
- Lu X and Guan H (2017) *Earthquake Disaster Simulation of Civil Infrastructures From Tall Buildings to Urban Areas*. New York: Springer.
- McCallen D and 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, Livermore, CA.
- McCallen D, Petrone F, Miah M, Pitarka A, Rodgers A and Abrahamson N (in press) EQSIM—A multidisciplinary framework for fault-to-structure simulations on exascale computers Part II: Regional simulations of building response. *Earthquake Spectra*.
- Miah M, Petrone F, Wong J and McCallen D (2018) Regional scale earthquake risk estimation based on broadband ground motion simulations. In: *Proceedings of the 11th U.S. National conference on earthquake engineering*, Los Angeles, CA, June 2018.
- Nilsson S, Petersson NA, Sjogreen B and Kreiss H (2007) Stable difference approximations for the elastic wave equation in second order formulation. *SIAM Journal on Numerical Analysis* 45: 1902–1936.
- Olsen K, Day S, Dalgner L, Mayhew J, Cui Y, Zhu J, Cruz-Atienza V, Roten D, Maechling P, Jordan T, Okaya D and Chourasia A (2009) ShakeOut-D: Ground motion estimates using an ensemble of large earthquakes on the southern San Andreas fault with spontaneous rupture propagation. *Geophysical Research Letters* 36: 1–6.
- Pankajakshan R, Lin PH and Sjogreen B (2019) Porting a 3D seismic modeling code (SW4) to Coral machines. *IBM Journal of Research and Development* 64: 1–11.
- Petersson NA and Sjogreen B (2010) Stable grid refinement and singular source discretization for seismic wave simulations. *Computer Physics Communications* 8: 1074–1110.
- Petersson NA and Sjogreen B (2012) Stable and efficient modeling of an elastic attenuation in seismic wave propagation. *Computer Physics Communications* 12: 193–225.
- Petersson NA and Sjogreen B (2014) Super-grid modeling of the elastic wave equation in semi-bounded domains. *Computer Physics Communications* 16: 913–955.
- Petersson NA and Sjogreen B (2015) Wave propagation in anisotropic elastic materials and curvilinear coordinates using a summation-by-parts finite-difference method. *Journal of Computational Physics* 299: 820–841.
- Petrone F, Wong J, McCallen D and McKenna F (2016) Development of advanced numerical tools for the nonlinear seismic analysis of nuclear systems. *Transactions of the American Nuclear Society* 115: 765–768.
- Pitarka A, Graves R, Irikura K, Miyakoshi K and Rodgers A (2020) Kinematic rupture modeling of ground motion from the M7 Kumamoto, Japan earthquake. *Pure and Applied Geophysics* 177, 2199–2221.
- Rodgers A, Petersson NA, Pitarka A and McCallen D (2018b) Simulation of large and recorded moderate earthquakes in the San Francisco bay area. In: *Proceedings of the 11th national conference on earthquake engineering*, Los Angeles, CA, June 2018.
- Rodgers A, Petersson NA, Pitarka A, McCallen D, Sjogreen B and Abrahamson N (2019b) 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 Research Letters*, 90, 1268–1284.

- Rodgers A, Pitarka A and McCallen D (2019a) 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: 1012–1046.
- Rodgers A, Pitarka A, Pankajakshan R, Sjögreen B and Petersson NA (2020) Regional-scale 3D ground-motion simulations of Mw 7 earthquakes on the Hayward fault, northern California resolving frequencies 0– 10 Hz and including site-response corrections. *Bulletin of the Seismological Society of America*. Epub ahead of print 11 August. DOI: 10.1785/0120200147.
- Rodgers A, Pitarka A, Petersson NA, Sjögreen B and McCallen D (2018a) Broadband (0-4 Hz) ground motions for a magnitude 7.0 Hayward fault earthquake with 3D structure and topography. *Geophysical Research Letters* 45: 739–747.
- Sahin A, Sisman R, Askan A and Hori M (2016) Development of integrated earthquake simulation system for Istanbul. *Earth, Planets and Space* 68(115): 1–21.
- Shaw J, Plesch A, Tape C, Suess P, Jordan T, Ely G, Hauksson E, Tromp J, Tanimoto T, Graves R, Olsen K, Nicholson C, Maechling P, Rivero C, Lovely P, Brankman C and Munster J (2015) Unified structural representation of the southern California crust and upper mantle. *Earth and Planetary Science Letters* 415: 1–15.
- Sjögreen B and Petersson NA (2012) A forth order accurate finite difference scheme for the elastic wave equation. *Journal of Scientific Computing* 52: 17–48.
- Sjögreen B and Petersson NA (2014) Source estimation by full wave form inversion. *Journal of Scientific Computing* 59: 247–276.
- Stidham C, Dreger D, Larsen S and Romanowicz B (1999) Three-dimensional structure influences on the strong-motion wavefield of the 1989 Loma Prieta earthquake. *Bulletin of the Seismological Society of America* 89: 1184–1202.
- Taborda R and Bielak J (2011) Large-scale earthquake simulation: Computational seismology and complex engineering systems. *Computing in Science and Engineering, July/August 2011*: 14–26.
- Taborda R and Bielak J (2014) Ground-motion simulation and validation of the 2008 Chino Hills, California earthquake using different velocity models. *Bulletin of the Seismological Society of America* 104: 1876–1898.
- Tape C, Liu Q, Maggi A and Tromp J (2009) Adjoint tomography of the southern California crust. *Science* 325: 988–992.
- Trifunac MD (2009a) Rotations in structural response. *Bulletin of the Seismological Society of America* 99: 968–979.
- Trifunac MD (2009b) The role of strong motion rotations in the response of structures near earthquake faults. *Soil Dynamics and Earthquake Engineering* 29: 382–393.
- Tu T, Yu H, Ramírez-Guzmán L, Bielak J, Ghattas O, Ma K and O'Hallaron D (2006) From mesh generation to scientific visualization: An end-to-end approach to parallel supercomputing. In: *SC'06 proceedings of the 2006 ACM/IEEE international conference for high performance computing, networking, storage and analysis*, IEEE Computer Society, Tampa, FL, 11–17 November.
- Wu S, Miah M and McCallen D (2018) Four canonical steel moment frame buildings and inter-code comparisons of nonlinear building response. University of Nevada, Center for Civil Engineering Earthquake Research Center Report CCEER 19-05, Reno, NV.
- Yoshimura C, Bielak J, Hisada Y and 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: 825–840.
- Zhang W, Esmaeilzadeh Seylabi E and Taciroglu E (2019) An ABAQUS toolbox for soil-structure interaction analysis. *Computers and Geotechnics* 114: 1–21.