Urgent Needs in Creating A Cyberinfrastructure for Extreme-scale Earthquake Software and Data

Yifeng Cui, San Diego Supercomputer Center, January 2020 (*yfcui@sdsc.edu*)

Introduction. Earthquakes and their effects are notoriously difficult to predict primarily because we do not fully understand the complex interaction between matter and energy within the lithosphere. Thanks to dramatically increased computational and data resources in recent years, physics-based simulation has become an integral part of seismic hazard research, enabling transformative research including integrated simulations of 3D dynamic fault ruptures and seismic wave propagation. The community's computational needs are, however, growing more rapidly than our nation's open-science supercomputer resources, as we cannot yet achieve the physical scale range needed to simulate more realistic earthquake rupture dynamics and resulting ground motion at high frequencies relevant for seismic hazard and risk. This scale of computing will be necessary to prepare simulation-based seismic hazard models for use in national seismic hazard maps, and in the American Society of Civil Engineers (ASCE) building code recommendations, to guide the long-term construction of a seismically safe built environment.

Notably, the status quo in the current funding model supporting earthquake HPC work - relative short-term project grants and limited strategic planning for R&D - has led to significant fragmentation, inhibiting efforts to solve community exascale computing challenges related to concurrency, data movement, energy efficiency and resilience [1]. A unified and synergistic umbrella for community practices is critical to serve these requirements through software tools, utilizing existing and preparing for future HPC systems [2], to address the community's future needs for a balanced ecosystem that supports both computationally-intensive and data-centric computing. Achieving a national priority to advance high-performance research computing for earthquake research will require rethinking the algorithms and software used in HPC applications, and a co-design approach, in which next generation hardware and software are developed together, creating effective new techniques for extracting useful information from the explosion of data in sensor networks, earthquake instruments, and numerical simulations themselves.

To exploit dramatic shifts in available petascale computing power, NSF successfully advanced a petascale collaboratory for the geosciences [3] as early as 2005. However, NSF is now behind DOE in providing funding to address major challenges involved in the transition to HPC at the exascale. *Earthquake system science as a whole would benefit from a tailored computational and data cyberinfrastructure center that collaborates with the supercomputer facilities in extreme-scale co-design activities, thereby enabling rapid major advances in this vibrant area of research.*

Where We Stand Now. The usage of earthquake computations has increased dramatically over the last decade. SCEC's awarded computing core-hours, for example, have increased by 900x since 2004 (**Fig. 1**), which have not been normalized by the respective growth in architectural computing power. The current mean rate of computer usage by computational seismology collaborators in the U.S. exceeds 1.5 million core-hours per day on average.

The 3D wave propagation code AWP-ODC is a good example. This ACM Gordon Bell winning application evolved from a relatively inefficient research code to a top HPC application used by the SCEC community, achieving a 500x performance enhancement between 2004-2017 (**Fig. 2**).

SCEC Computational Allocation Increases

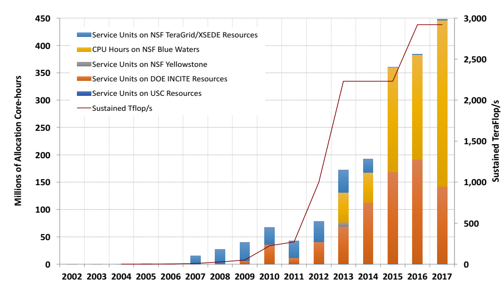

Figure 1. Computing service units per year allocated to SCEC on HPC facilities operated by NSF, DOE and USC

While access to computational resources have been critical, there has been no fixed accompanying support for software development efforts. The current software R&D in the earthquake science community doesn't have fixed person-hour budget support matching the allocation awards, which are *valued at annual $26.5M federal investment*, *based on Amazon EC2 m4. large pricing of $0.05 per core/hour* [4]. SCEC's collaborative computational research is a well justified use of these shared research facilities. Historically, funding for SCEC HPC software development was from NSF awards received from the OAC/GEO directorates, including ITR, SI2 and Geoinformatics grants. Earthquake

---

[1] DOE ASCAC Data Subcommittee Report: Data-intensive Science and Exascale Computing, Office of Science, DOE, March 2013.
[2] Alan Levander, Jeroen Tromp, Ray Willemann, A Center for Computational Solid Earth Science, White Paper submitted to NSF EarthCube Program, pp. 1-4, 2011.
[3] Bryan F. (chair) et al., Ad Hoc Committee and Technical Working Group for a Petascale Collaboratory for the Geosciences, 2005: Establishing a Petascale Collaboratory for the Geosciences, 2005: Scientific Frontier. A Report to the Geosciences Community. UCA/JOSS, Boulder, Colo., 80pp.
[4] Amazon EC2 Pricing, https://aws.amazon.com/ec2/pricing/on-demand/

simulations hazard research results were widely used by NSF and DOE resource providers as examples of computational results with broad beneficial societal impact.

While earthquake HPC has essentially no fixed NSF funding for computational research, NCAR receives an approximately annual \$33M fixed HPC operation budget through its Computational Information Systems Laboratory (CISL) program, including a \$7-8M annual hardware upgrade. NCAR's Cheyenne system allocates 1 billion core-hours per year, roughly doubled size of the SCEC allocated core-hours [5]. DOE's ECP provides a baseline cost range of \$1.7B-\$2.1B, including \$620M funding in application development [6], approximately 30% of exascale investment for application software development. Within DOE, earthquake research has not been identified as a strategic application area, so while we use their supercomputers, we receive no funding support from DOE for our application software development.

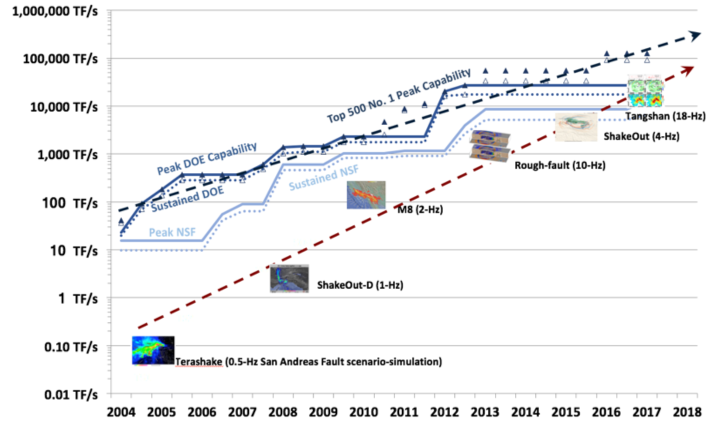

Figure 2. Timeline of the AWP-ODC evolution in comparison to sustained performance with to DOE, NSF and Top 500 Nr. 1 systems. The AWP-ODC trend is steep, indicating the need for extreme-scale computing. Points are the milestone simulations including TeraShake, ShakeOut, M8, simulation using rough-fault dynamic rupture, non-linear version of ShakeOut, and the latest Gordon-Bell-prize-winning Tangshan simulations, corresponding to a 10^4 increase in the 4D outer/inner scale ratio.

The imbalance within NSF in terms of supporting earthquake community application software activities provides a context for SCEC's funding requests as special projects. We believe that extreme-scale seismic hazard software and research should be funded via sustained investment, rather special projects, to reflect better the increased use of seismic hazard modeling by public agencies and civil engineering communities. An HPC program like CISL dedicated to seismic hazard model development, including long-term staff and software development support, is needed to establish and maintain an effective U.S. computational earthquake research program.

We propose to establish a dedicated computational earthquake science center for research, development, education and training. This CI center would provide the following services: **1**. Acceleration of existing community software development, that supports exascale computing for comprehensive earthquake simulation systems designed to model various aspects of an earthquake disaster. **2.** Architecture-aware numerical approaches and implementation for modeling/simulation and data centric systems, high-dimensional seismic optimization and uncertainty quantifications for efficiency and accuracy. **3.** Coherent data analytics and data-driven visualization technologies, for visual interactive supercomputing, data-mining, and machine-learning with emphasis in integration of simulation, data and learning stack that support next generation neural network architectures. **4.** Operational earthquake forecasting and risk assessment, including near real-time seismological simulations and urban response to great and significant events, and on-demand simulations. **5.** Small to middle-size hardware to support for on-demand computing needs, software development, co-design of next-generation applications, continuous verification and deployment, and training (10s-100s of nodes, 100s of TB fast memory, high bandwidth access). **6**. Education, Outreach & Training for the user community via workshops, summer schools, and on-line forums, connected to a professional Ph.D program in computational earthquake simulation and software engineering that trains the next generation of computational and data scientists.

The center would become a focus of scientific research, education and scientific collaboration for the worldwide community of computational earthquake scientists, to catalyze significant advances in software infrastructure, education, standards, and best-practices (**Fig. 3**). Innovative ideas in software and data analytics will not languish in the "get the science right" mode, but address the next set of grand challenges and enable bigger and better earthquake science. To meet the earthquake community's simulation needs, the center would design

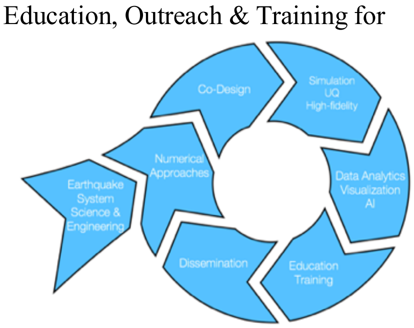

Figure 3. Data and Software activities

[5] Computational & Information Systems Lab (CISL) at NCAR, https://www2.cisl.ucar.edu
[6] Office of Advanced Scientific Computing Research and Office of Science, DOE Exascale Computing Project, Project Execution Plan, Nov 2017

advanced technologies with next-generation hardware configurations, and applications optimized for 3P (*Performance, Portability and Productivity*), which allows for user interaction between iterations in compute-intensive calculations. The center would provide data-intensive computing guidance, enabling "Big Data" approaches to deal with datasets, e.g. LiDAR and InSAR observations, and create micro-earthquake catalogs of unprecedented resolution and large-scale ambient-wavefield analysis for broader implementation.

The center would help expand the volume and sustainability of open-source software, and to establish standards designed to insure software portability across platforms. The scientific verification and validation efforts would tightly couple basic earthquake research to the practical needs of probabilistic seismic hazard analysis, operational earthquake forecasting, earthquake early warning, and rapid earthquake response. The center would support and mentor a cohort of Software Fellows actively developing code infrastructure in research groups across the U.S., and in turn, they will engage in outreach and education activities within the larger earthquake science community. The center would sustain an ecosystem with the evolution of broader ambitious system architectures incorporating data analytic attributes and capabilities, for bringing the data analytics, artificial intelligence and HPC development together.

The approximate investment for the center would require *5-10 FTEs* with fixed annual installments over five fiscal years, sustained by a mix of personnel contributions through NSF GEO/EAR (facility), CMMI (engineering), and OAC (HPC and Data) programs, USGS (NEHRP) program, DOE (ECP) program, and FEMA (emergency response) as well as industrial contributions such as PG&E, and hardware vendors. This investment would help to maintain a healthy and vital software ecosystem for the community (**Fig. 4**).

Community Involvement. The center would serve the broad earthquake science community, executed using the partnering mechanism involving multiple federal programs with the biggest stake in a balanced ecosystem, including: **1.** Earthquake Hazard Seismology: SCEC continuation**. 2.** HPC and Data Analytics: XSEDE at an NSF Supercomputer center. **3.** Education & Training: SCEC UseIT, CIG, CIDER and XSEDE educational programs**. 4.** Broader Collaborations with mix of personnel contributions: Engineering Research (e.g. NHERI, UCLA), Observation (e.g. USGS, NASA, IGPP, GAGE), Geodynamics (e.g. CIG), Global Seismology (e.g. SAGE, IRIS), Earth System (e.g. USGS, NCAR, NASA), Supercomputing (e.g. SDSC, TACC, NERSC), Industry (e.g. PG&E, Intel, AMD, NVIDIA, Amazon), Risk Management (e.g. RMS, SwissRe), and Public (e.g. USGS, FEMA, CEA, NHERP).

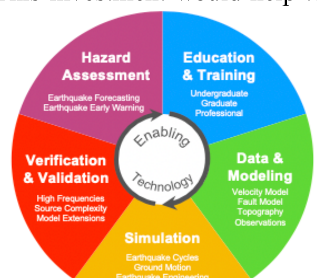

Figure 4. Enabling technology and education of earthquake system science that support simulation-based predictions, validation against observations, verification for system-model, and hazard assessment

Why Now? The proposed center is timely, giving multi-agency strategic investment in National Strategic Computing Initiative (NSCI) and sustained investments in geoscientific software and data, and the decade's growth in computational size and data volumes in earthquake science. Given the fact that U.S. is already behind Japan and China last few years in supporting end-to-end earthquake simulations, the importance of establishing a computational center for accelerated innovation and enhanced competitiveness is hard to overstate. In view of today's fiscal constraints, we must apply the lessons learned in our past successes to strategically target balanced investments in both ongoing earthquake system science core programs and breakthroughs through the enabling technology. An integrated, sustained, and decadal program of technology development and testing would maintain a healthy U.S. earthquake HPC base.

The system needed for HPC and Big Data in the earthquake sciences has specific requirements to optimize productivity running open-source software in use by and under development in the computational earthquake community. We would only add minor hardware and connectivity infrastructure, counting on continuing allocation of very large compute resources through NSF and DOE allocation awards. Software engineering and a professional manager attuned to the needs of HPC for the earthquake sciences are needed to optimize community codes for the specific or general exascale hardware. We believe that such a computational center can be set up at one of the many existing HPC centers in the U.S. This would allow us to leverage an existing concentration of HPC knowledge, avoid the costs of acquiring the specialized physical infrastructure for the hardware, and have easy access to faster network and petabytes data storage capabilities.

In summary, a key to future success will be a rigorous requirement for developing, re-implementing and mapping of advanced algorithms and applications to exascale systems, and integrating simulation, data and learn stack in the community. We believe that the practical societal benefits of the improved seismic hazard information from the center will attract financial support from additional stakeholders, and the center research activities will scale-up as additional resources are added.

Acknowledgements. Yehuda Ben-Zion & Philip Maechling of SCEC, and Michael Norman & Shawn Strande of SDSC provided helpful comments for improvements of this white paper, which was originally initiated in Jan 2018.