Urgent Supercomputing of Earthquakes: Use Case for Civil Protection

Josep de la Puente
 Juan Esteban Rodriguez
 Marisol Monterrubio-Velasco
 Otilio Rojas
 Arnau Folch
 josep.delapuate@bsc.es
 juan.rodriguez@bsc.es
 marisol.monterrubio@bsc.es
 otilio.rojas@bsc.es
 afolch@bsc.es
 Barcelona Supercomputing Center
 Barcelona, Spain

ABSTRACT

Deadly earthquakes are events that are unpredictable, relatively rare and have a huge impact upon the lives of those who suffer their consequences. Furthermore, each earthquake has specific characteristics (location, magnitude, directivity) which, combined to local amplification and de-amplification effects, makes their outcome very singular. Empirical relations are the main methodology used to make early assessment of an earthquake's impact. Nevertheless, the lack of sufficient data registers for large events makes such approaches very imprecise. Physics-based simulators, on the other hand, are powerful tools that provide highly accurate shaking information. However, physical simulations require considerable computational resources, a detailed geological model, and accurate earthquake source information.

A better early assessment of the impact of earthquakes implies both technical and scientific challenges. We propose a novel HPC-based urgent seismic simulation workflow, hereafter referred to as Urgent Computing Integrated Services for EarthQuakes (UCIS4EQ), which can deliver, potentially, much more accurate short-time reports of the consequences of moderate to large earthquakes. UCIS4EQ is composed of four subsystems that are deployed as services and connected by means of a workflow manager. This paper describes those components and their functionality. The main objective of UCIS4EQ is to produce ground-shaking maps and other potentially useful information to civil protection agencies. The first demonstrator will be deployed in the framework of the Center of Excellence for Exascale in Solid Earth (ChEese, <https://cheese.coe.eu/>, last access: 12 Feb. 2020).

CCS CONCEPTS

• **Applied Computing**; • **Physical Sciences and Engineering**; • **Physics**; • **Networks** → Network reliability;

KEYWORDS

earthquakes, urgent supercomputing, civil protection

ACM Reference Format:

Josep de la Puente, Juan Esteban Rodriguez, Marisol Monterrubio-Velasco, Otilio Rojas, and Arnau Folch. 2020. Urgent Supercomputing of Earthquakes: Use Case for Civil Protection. In *Proceedings of the Platform for Advanced Scientific Computing Conference (PASC '20)*, June 29-July 1, 2020, Geneva, Switzerland. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3394277.3401853>

1 INTRODUCTION

Since 2010 and until the writing of this document, 363,250 people have died because of earthquakes worldwide. In terms of economic losses, three earthquakes are considered the costliest natural disasters ever with an aggregate cost of almost a thousand billion US dollars [cite10.3389/fbuil.2017.00030]. Even though typically ranked in terms of their moment magnitude (M_w), the most severe events, as for example the Haiti 2010, M_w 7.0 earthquake that caused 300,000+ deaths, may have only moderate magnitude. For example, Japan's event of 2011, M_w 9.1, released about 100 times more seismic energy, yet it caused less than 10% of fatalities when compared to the Haiti event. This is largely due to a number of factors such as the earthquake's depth, epicentral location with respect to populated areas, population density or vulnerability of buildings and infrastructures.

Unfortunately, earthquakes are unpredictable. Albeit better building, popular awareness and early warning can mitigate their impact, we are likely to see similar casualty figures in years to come. On average more than 150 potentially harmful earthquakes ($M_w > 6.0$) are recorded yearly.

The combination of unpredictable, rare, yet extremely destructive nature of earthquakes calls for response measures that are reactive in nature. In order to be relevant they must have an impact in the associated relief efforts. That is precisely what shake maps

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

PASC '20, June 29-July 1, 2020, Geneva, Switzerland

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7993-9/20/06.

<https://doi.org/10.1145/3394277.3401853>

Figure 1: peak ground velocity (PGV) values computed at Grenoble valley, France for a hypothetical $M_w=6$ earthquake originating at the fault depicted as a black line. The yellow area displays a sedimentary zone, with different rock properties with respect to the surrounding area. [15]

produced right after the events happens [27], can obtain: accurate distributions of seismic motion caused by the earthquake which, when correlated with vulnerability (e.g. population density, building quality) can help direct the relief efforts. To illustrate what a shake map is, a simulated map of the peak ground velocity (PGV) values computed at Grenoble valley is shown in Fig. 1 [15]. PGV just tells us the maximum velocity of shaking induced by a given earthquake at each point of the map. Hence red zones in Fig. 1 have suffered strong shaking when compared to blue zones. Damage assessment studies typically compare shaking to vulnerability, i.e. the capability of different structures to sustain shaking.

Data from seismometers can be used to record the amount of shaking due to an earthquake, but seismometer networks tend to be very sparse. In order to assess damage at a useful scale, i.e. less than a kilometer resolution, shaking must be computed by other means. To date, existing systems to estimate the impact of an earthquake are based on the so-called ground motion prediction equations (GMPEs). GMPEs are essentially empirical formulas calibrated to each region that give shaking as a function of distance to the epicentre and magnitude [4, 5, 7, 8]. As such, they average out the contributions from several prior earthquakes and provide with a radial shaking distribution, very different to what we observe as the outcome of a single earthquake (see e.g. Fig. 1). The limitations associated to GMPEs are overcome by full physical 3D simulations of earthquakes, yet such simulations typically have long turnaround times. Nevertheless we expect the level of detail attained by seismic simulations, in the order of tens of meters, to be valuable when determining locations that have been most affected by an earthquake.

In this work, we propose an urgent supercomputing workflow for earthquake simulations that can: 1) gather information about potentially lethal earthquakes automatically, 2) decide whether a 3D

simulation should provide valuable information, 3) preprocess all the information needed for the simulation, 4) launch the simulation in a tier-0 PRACE HPC system and, finally, 5) postprocess the results and send them to relevant stakeholders. By taking advantage of the huge parallel efficiency of seismic simulators, we expect to obtain results (i.e. run through the complete workflow) in few hours. This is the first attempt to provide an HPC-based urgent seismic workflow addressed to mitigate earthquake impact. The work below is funded by the Center of Excellence for Exascale in Solid Earth (ChEESE, <https://cheese.coe.eu/>, last access: 12 Feb. 2020).

2 SIMULATION INPUTS AND DATA

An earthquake is consequence of fast rupture, or slip, of a geological fault [21]. Such rupture produces a vibration that is then radiated by means of seismic waves. Such waves are affected by heterogeneity in the subsurface as well as by the topography [3]. Similar to sound or water waves, seismic waves can be trapped in certain structures and suffer from amplification under certain situations. Furthermore their energy fades as they travel away from the source (i.e. from the fault). Earthquake simulations require inputs both for the earthquake's characteristics (e.g. magnitude, location) and the soil properties. Nowadays seismic simulations have reached a high degree of fidelity with respect to actual data [6]. However, instead of being discrete and spatially sparse as are data records, they provide with high resolution both in space and time. In contrast to empirical approximants such as GMPEs [4, 5, 7, 8], seismic simulations are able to reproduce precisely quantities such as spectral acceleration or shake time that inherently require physical modelling of seismic waves. These quantities are fundamental in order to assess whether buildings can withstand the earthquake's energy [14].

Seismic simulations, however, are not without drawbacks. They are, on one hand, very costly computationally, requiring the solution of large-scale three-dimensional physical models, at the very least covering the region encompassing earthquake and the zone of interest. On the other hand, they are fully deterministic and non-linear, hence uncertainties in either the physical model or the source characteristics can have a strong impact on the results. Efficient parallel solvers and distributed-memory compute clusters help in dealing with the first issue. In particular, seismic simulators as those used here allow for almost halving the compute time when doubling the amount of processors available, as we will see in Section 3.3. The second problem requires a proper fine-tuning of the model parameters and, for time-constrained applications, a proper assessment of uncertainties. Sources of uncertainty include:

- Location. Typically measured accurately in well-monitored areas, except for depth, due to seismometers being all located at the Earth's surface.
- Centroid moment tensor, CMT, [11, 19]. These are the source forces equivalent, at long distances, to the earthquake originating at a single point. The CMT accounts for the event's size or magnitude as well as the directivity, which determines in which direction most of the seismic energy is released. CMT computations require a non-negligible calculation time and early access to seismograms from nearby seismograms [24].

- **Fault slip distribution.** The detailed rupture of the fault, of which the CMT (and an associated time signal) is a simplification. This is largely unknown even after detailed post-event studies [2]. Heterogeneity in the slip distribution, however, is responsible for the frequency content of the earthquake. Such frequency content is key in seismic engineering studies, because different structures or buildings are sensitive to different frequencies. As an alternative to the process of computing the fault slip distribution of a particular earthquake, several distributions can be sampled among random models which are compatible with fault physics, location and CMT.
- **Geological model.** A 3D representation of physical properties of the subsurface which are more relevant for seismic wave propagation, including: compressional (P) and shear (S) wave velocities, density and attenuation, with special attention to shallow properties which deeply affect the shaking locally. This last aspect can only be addressed by means of calibration between data and simulations.

It should be noted that uncertainties in the items above are active research topics in the seismological community (see, e.g. [17, 23]).

3 URGENT EARTHQUAKE SIMULATION WORKFLOW

The proposed workflow is composed of four components that provide specific services, jointly enabling the urgent computing procedure: Automatic Alert Service, Smart Centre Control, HPC access, and Post-process service.

UCIS4EQ requires an integrated workflow manager with the following requirements: manage task dependencies, run parallel simulations on different HPC infrastructures, manage batch jobs (submission, monitoring, cancellation), manage conditional paths and access external data repositories (R/W) such EUDAT. PyCOMPSS [25] is a Barcelona Supercomputing Center in-house developed workflow manager that covers all these requirements and hence is the main candidate to support UCIS4EQ's development.

A conceptual structure of the UCIS4EQ is shown in Fig. 2. In this figure the color blue indicates components requiring HPC infrastructure, as opposed to smaller components that may run in servers, which are depicted in gray. The following subsections describe each component of UCIS4EQ in detail.

3.1 Automatic Alert Service

The activity diagram of the first component is shown in Fig. 3, (AAS). This subsystem is in charge of detecting new earthquake events as soon as they are detected and consists on three subcomponents sequentially connected: *notification*, *registration module* and *triggering*.

The first of the three subcomponents is the general entry point for the UCIS4EQ system. It has been designed as a service that enquires earthquake data centers via the International Federation of Digital Seismograph Networks (FDSN) web service interface [22]. Such a federation is composed of the global and regional online seismic data centres, such as the Incorporated Research Institutions for Seismology (IRIS), United States Geological Survey (USGS), Icelandic Meteorological Office (IMO), National Institute of Geophysics

and Volcanology (INGV), or GEOFON. All of them share the single mission of extracting the earthquake information related to each new event in the minimum possible time. Additionally, the FDSN has developed a data exchange format called Standard for the Exchange of Earthquake Data (SEED) for the distribution of seismic information that has become the main method for seismological data exchange between independent networks and data centers worldwide via web servers [22]. Hence, the notification subcomponent acts as a contact point between the agencies and the UCIS4EQ system, by continuously enquiring data centers for either new events or updated information on recent ones. Once the AAS receives a new event, it must be registered internally for monitoring purposes. However, prior to this, incoming events must be disambiguated. Quite often, an event is registered by more than one FDSN agency and each agency identifies the event's properties using their own methodologies. This results in discrepancies in the information received regarding a single event. The registration subcomponent is in charge of disambiguation, by assigning to each earthquake a unique universal identification code (UUID) and registering all related information, consistent or not, on an internal database.

After an event is correctly registered, the trigger subcomponent is activated. This subcomponent determines whether a new earthquake has the potential (size, proximity to populated areas) to trigger an urgent simulation. In affirmative scenarios, this component activates the next UCIS4EQ system, the Smart Centre Control (SCC). The heuristics for triggering include regional constraints and magnitude thresholds but might include more precise means of evaluation in the future.

During the triggering process, the information obtained from the FDSN agencies which is relevant for the simulations is collected into JSON documents. Each single document contains the input parameters that the SCC needs to start running.

3.2 Smart Centre Control

The Smart Centre Control (SCC) triggers an internal workflow manager. The set of tasks that are executed is split into three different branches, which run in parallel with mild dependencies, as shown in Fig. 4.

- (1) **Computational resources.** Requests (urgent) access to HPC resources. This execution path is crucial since the computational resources for simulations must be ready as fast as possible. Clear and fast urgent computing policies must be enabled at HPC centers in order to minimize time to solution.
- (2) **Simulation input parameters.** Builds the input files required by seismic simulators at upcoming stages of the workflow. Must deal with variations resulting from uncertainties in the input parameters, potentially generating a set of input parameters.
- (3) **Streaming service.** Inquires external databases for updates or corrections in earthquake data, and determines whether new information renders the old information useless (hence restarting the workflow at the SCC) or uncertainty estimates must be modified.

Figure 2: Conceptual diagram of the Urgent Computing Integrated Services for EarthQuakes (UCIS4EQ) workflow. UCIS4EQ is developed in four components, Automatic alert service, Smart Centre control, HPC facilities, and Post-process service. The legend describes the meaning of the shapes and colors of the activity diagram components in Figs. 3,4, and 5.

Figure 3: Activity diagram of the Automatic Alert service (AAS) component.

3.2.1 Computational resources. This SCC branch is in charge of requesting computational resources, and also determining the priority level of that job as well as its expected size in order to reach results in a useful time. A pivotal step for a fully operational urgent computing system based upon public tier-0 supercomputers is the establishment of protocols that must be matched to the HPC system's governance in order to automatize the process of launching simulations [18]. The Urgent Seismic Computing protocol (USCp) is a technical guide to execute seismic simulations in a High Performance Computing (HPC) environment considering strong time constraints. The USCp is being developed parallel to the technical developments of this workflow under the ChEESE framework. An

urgent computing (UC) application must manage complex interactions between users, computational resources, elevated executed priority policies, and working sessions. UC applications must provide relevant information before a deadline since later information may be useless. The urgency usually implies that there is no second chance, since the job execution after a deadline makes the result useless or of little relevance. Nowadays, HPC infrastructure provide virtually no support for UC, where users require efficient and prompt access to the necessary computational resources. Although some effort has been done in order to specify UC service inside the PRACE Research infrastructure [18], a specific protocol to execute Urgent Seismic Computing is in progress. The USCp policies within the PRACE Research Infrastructure require that those executions

Figure 4: Activity diagram of the Smart Centre Control subsystem. The blue part depicts tasks that require HPC hardware. Three branches concur in parallel horizontally, related to computer resources, simulation parameters and streaming service. The latter could force a stop in the workflow if data updates result in too large differences with respect to early data.

comply with all preset statements. Moreover, the USCp requires technological developments ensuring the QoS before, during, and after an urgent job execution. USCp must consider resource selection, submission policies and compensation schemes for affected users.

3.2.2 Simulation input parameters. In order to produce proper data for simulations, detailed inputs are required. The tasks that are developed for this branch must determine such inputs and encapsulate them for the following stage. In our application simulation inputs are:

- (1) CMT, as described in previous Sections. A CMT provides information on the earthquake's rupture directivity [26]. The frequent lack of a fast official CMT information, provided by seismological agencies is thus an important issue. Therefore, within UCIS4EQ we enable several alternatives towards obtaining a fast estimate of the CMT. In particular:
 - A comparison to nearby historical events which may have been generated at the same fault and thus share similar rupture mechanisms, i.e. similar CMT characteristics. The historical CMT database is available for different regions by means of online repositories such as IRIS or GLOBAL-CMT project [10, 11]. These repositories are freely available and contain more than 25,000 moment tensors for earthquakes since 1960.

- Determining the major fault nearest to the hypocenter, and assuming a good knowledge of the regional tectonics by means of previous geological studies, we may have a good guess of likely mechanisms for our earthquake. The reason for this is that it is likely that large events (i.e. those relevant for civil protection purposes) happen in well known faults, whose geometry strongly constrains the CMT characteristics of an event.
- Running a fast CMT estimation algorithm, which requires fast access to seismological records of the event. This is still under development within UCIS4EQ.

Our algorithm uses all of the strategies above and estimates the most likely CMT based upon the reliability of each strategy for each region, as well as quantifying uncertainties in the computation.

- (2) The temporal evolution of earthquake slip along the fault surface can be estimated by using a CMT approximation and the fault geometry information by means of stochastic methods. State-of-the-art kinematic rupture generators employ physical-based statistical distributions of fault slip and additional rupture variables, inspired and constrained by the dynamics of real earthquakes. In our case, we incorporate the Graves-Pitarka rupture generator method [12, 13] with more than ten years of evolution and testing. Because of the random component of the generated ruptures, we simulate a

set of ruptures and combine their results. Nevertheless, each rupture must be simulated as an independent simulation.

- (3) The 3D velocity and density models are assumed known for a particular region. Files can be large, thus we plan to employ EUDAT data management services to handle their accessibility across servers and HPC centers. This is considered external "static" data, trimmed and loaded by means of a Query process. Because of the size of some datasets (especially 3D models) and the need for preprocessing (e.g. meshing), several tasks related to preparing simulation inputs are performed asynchronously prior to USCp approval or resource allocation. As shows Fig. 4 the model building is the only component that must be executed in an HPC environment at this point to reduce overheads related to large data movement.

Two of the the three parallel branches in the SCC component (Computational resources and simulation input) reach a synchronization point prior to performing the simulation on the HPC environment. The final output files generated by the SCC component are written in YAML format, which can be easily interfaced with the different seismic simulation codes (see Section 3.3).

3.2.3 Streaming service. The Streaming service includes an additional set of tasks that checks for new information on the global and regional seismic databases (through the FSDN). For a specific event, whenever an external provider modifies any previously registered data, the SCC loads the new inputs, updates the internal state of such event on the database and evaluates the impact of the incoming information on the quality of the results comparing to the already running tasks. As a result, the streaming service may decide to reconfigure, restart or omit the discrepancies. An effective change evaluation methodology is critical at this point for providing reliable results. The fundamental component is the analysis of the impact of modifications both in the accuracy/relevance of the results and the time to solution. Present configurations only consider the omission scenario, but more complex heuristics will be developed in the future.

3.3 HPC facilities and tools

Key to the success of urgent earthquake supercomputing are efficient and accurate simulation codes. Due to the time criticality of the problem, being able to rely on massive parallelism to cut turnaround times is a very desirable feature.

There exist several software packages that offer the physical capabilities (viscoelastic, accurate topography representation and fault rupture description) and computational efficiency (good scalability up to tens or hundreds of thousand cores). In particular, Seissol [9], SPECFEM 3D [16], ExaHyPE [20] and Salvus [1] are natural choices and their main developers are part of the ChEESE consortium. All the codes considered rely on explicit high-order finite elements, of the discontinuous Galerkin or spectral types. They all support unstructured meshes in order to incorporate realistic geological and topographic features. Within the ChEESE project all codes will be interoperable by means of simple translation of the YAML descriptors of the SCC into particular input files for each package, albeit an extended effort will aim at determining optimal software/hardware combinations in terms of speed and reliability.

Further optimization efforts within the project aim at obtaining the most efficient simulations possible within PRACE HPC systems.

A major aspect that needs to be taken into account is fast and reliable interoperability between the codes, in particular aspects related to building meshes, importing fault models and outputting the desired information, hence a complete API is being designed to ensure a seamless interaction.

Linked to the protocols discussed in the previous section should be access to special priority queues in the HPC systems. Similarly, a QoS of each available system should be the final component in deciding how many resources on which HPC system results in the shortest time-to-solution. Fig. 5 displays an activity diagram at the HPC facilities. One or more executions occur in parallel depending on the number of simulations defined in the SCC subsystem. As a requirement, seismic simulation packages must be pre-installed at the HPC machines and accessible for users for a correct execution once the resources are available.

3.4 Post-process service

A critical aspect of the system is its capacity to provide suitably tailored information to all stakeholders involved. In order to have the best impact in relief efforts, information must be delivered to each stakeholder's preferred communication channel and in its preferred form. The post-process service performs the associated tasks. Among the immediate outputs are PGV (peak ground velocity) and PGA (peak ground acceleration) maps for the region of interest, as well as spectral acceleration at selected sites (e.g. infrastructures). Our current efforts rely on the ChEESE project's partners as early adopters or initial stakeholders, including INGV in Italy and the Iceland Meteorological Office, both in charge of their respective national monitoring networks. A crucial aspect of post-processing will be harnessing the uncertainty in the results coming from exploration. Uncertainty analysis typically requires a profound exploration of the results. Such effort somehow collides with the need for urgent results in decision making, which is the ultimate product that disaster management agencies require. Hence novel approaches to provide with a clear picture of the affectionation while embracing the uncertain nature of results will need to be developed in the long term. Similarly, we will need to study the need for calibration or amplification due to site effects, i.e. contributions from the shallowmost part of the Earth's crust which are difficult to include in physical simulations but can be adopted from data analysis from seismic data analysis.

3.5 Workflow summary and current stage

UCIS4EQ is the first package that aims at using (urgent) HPC simulations to analyze the impact of earthquakes. Composed of four systems (AAS, SSC, HPC computation and post-processing), we envision a system capable of autonomously determining the need for analyzing a particular earthquake, gathering the necessary computational resources and outputting informative data on the event's affectionation. UCIS4EQ uses the PyCOMPSs workflow manager [25], a flexible, fault tolerant and parallel manager for task-based executions.

Currently we work on a demonstrator for the technology that will produce results for a past event in Iceland for which there is

Figure 5: Activity diagram of HPC resources and tools component.

extensive data. The demonstration should give us hindsight on its potential capabilities in terms of accuracy, but also in terms of its cost. We will focus on analyzing execution bottlenecks and making a sensitivity analysis of the uncertain parameters. At a later stage we will expand the demonstration to other areas both in Europe.

We remark that the current implementation still has several manual overrides, for example in the decision for an urgent simulation or in determining the best CMT estimates. It is expected that as the workflow matures automation becomes more prevalent until eventually the workflow can run without the need for human intervention, albeit that would require, for example, clear protocols for urgent HPC access, similar to what we propose in USCp.

4 WORK AHEAD AND CONCLUSIONS

The workflow discussed in the present work is set for an initial testing phase in mid-2020. We expect the demonstration to be an enabler for further discussion with PRACE to try defining clear policies and special queueing systems for urgent supercomputing (see, e.g. [18]). Similarly the demonstration should showcase UCIS4EQ's capabilities to potential end users, including our early adopters, but also other institutions that might have interest in exploring the possibilities that urgent simulations can give them. We can only guess at the modifications and suggestions that such stakeholders might suggest, but the design of the workflow has been developed with sufficient flexibility to accommodate virtually any strategies towards further exploiting the richness of full physical earthquake simulations.

We further acknowledge the synergies between our work and current trends in seismology that are key technology enablers for UCIS4EQ. This includes novel techniques for seismic tomography that produce 3D maps of the subsurface with unprecedented level of accuracy, meshing efforts that reduce significantly the pre-processing time for simulations or theoretical developments in uncertainty quantification that result in a solid foundation for rigorous studies of sensitivity and accuracy in our results. All of these

developments will further strengthen the potential usability of the workflow in the future.

Last but not least, it should be noted that albeit earthquakes are a clear use-case for urgent simulations, there are other natural and man-made disasters that might benefit from fast access to tier-0 supercomputing centers (e.g. forest fires, tsunamis, volcanic eruptions, accidental release of toxic pollutants such as radionuclides). We believe that an urgent usage of supercomputing resources, for such events, far from being a misuse of resources due to their impact in the systems's regular queue, might improve the system's efficiency. By running actual seismic events, rather than hypothetical or past events, we can have a direct assessment of impact on the wellbeing of the population being affected by such disasters. Furthermore, it might be a vehicle for normalizing the usage of HPC simulations in civil protection activities, a field with a huge potential ahead. This is of particular value in situations where past data is not complete enough to reproduce all potential future scenarios, such as in seismology.

ACKNOWLEDGMENTS

The ChEESE project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement Number 823844. Further funding includes the Spanish Ministry project GAP TIN2016-80957-P, the EU project MATHROCKS Number 777778. The Graves-Pitarka Generator is available through the SCEC Broadband Platform version 19.4. We acknowledge the assistance by Favio Silva and Phillip Maechling on its compilation and execution.

REFERENCES

- [1] M Afanasiev, C Boehm, M van Driel, L Krischer, DA May, M Rietmann, and A Fichtner. 2018. Salvus: An open-source high-performance package for full waveform modelling and inversion from laboratory to global scales. *Geophysical Journal International* (2018).
- [2] RMJ Amey, A Hooper, and RJ Walters. 2018. A Bayesian method for incorporating self-similarity into earthquake slip inversions. *Journal of Geophysical Research*:

- Solid Earth* 123, 7 (2018), 6052–6071.
- [3] Ari Ben-Menahem and Sarva Jit Singh. 2012. *Seismic waves and sources*. Springer Science & Business Media.
 - [4] Julian J Bommer, John Douglas, Frank Scherbaum, Fabrice Cotton, Hilmar Bungum, and Donat Fäh. 2010. On the selection of ground-motion prediction equations for seismic hazard analysis. *Seismological Research Letters* 81, 5 (2010), 783–793.
 - [5] David M Boore and Gail M Atkinson. 2008. Ground-motion prediction equations for the average horizontal component of PGA, PGV, and 5%-damped PSA at spectral periods between 0.01 s and 10.0 s. *Earthquake Spectra* 24, 1 (2008), 99–138.
 - [6] Yifeng Cui, Kim B Olsen, Thomas H Jordan, Kwangyoon Lee, Jun Zhou, Patrick Small, Daniel Roten, Geoffrey Ely, Dhabaleswar K Panda, Amit Chourasia, et al. 2010. Scalable earthquake simulation on petascale supercomputers. In *SC'10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–20.
 - [7] J. Douglas. 2011. *Ground-motion Prediction Equations 1964–2010*. Pacific Earthquake Engineering Research Center. <https://books.google.es/books?id=LUxwnQAACAAJ>
 - [8] John Douglas and Hideo Aochi. 2008. A survey of techniques for predicting earthquake ground motions for engineering purposes. *Surveys in geophysics* 29, 3 (2008), 187.
 - [9] M. Dumbser and M. Käser. 2006. An Arbitrary High Order Discontinuous Galerkin Method for Elastic Waves on Unstructured Meshes II: The Three-Dimensional Case. *Geophys. J. Int.* 167, 1 (2006), 319–336.
 - [10] AM Dziewonski, T-A Chou, and JH Woodhouse. 1981. Determination of earthquake source parameters from waveform data for studies of global and regional seismicity. *Journal of Geophysical Research: Solid Earth* 86, B4 (1981), 2825–2852.
 - [11] Göran Ekström, Meredith Nettles, and AM Dziewoński. 2012. The global CMT project 2004–2010: Centroid-moment tensors for 13,017 earthquakes. *Physics of the Earth and Planetary Interiors* 200 (2012), 1–9.
 - [12] Robert Graves and Arben Pitarka. 2016. Kinematic ground-motion simulations on rough faults including effects of 3D stochastic velocity perturbations. *Bulletin of the Seismological Society of America* 106, 5 (2016), 2136–2153.
 - [13] Robert W Graves and Arben Pitarka. 2010. Broadband ground-motion simulation using a hybrid approach. *Bulletin of the Seismological Society of America* 100, 5A (2010), 2095–2123.
 - [14] Yu-Xian Hu, Shih-Chi Liu, and Weimin Dong. 1996. *Earthquake engineering*. CRC Press.
 - [15] Martin Käser, Michael Dumbser, and Josep de la Puente. 2006. An efficient ADER-DG method for 3-dimensional seismic wave propagation in media with complex geometry. In *ESG 2006, Third Intl. Symposium on the Effects of Surface Geology on Seismic Motion*. Vol. 1. LCPC Editions, 455–464.
 - [16] D. Komatitsch and J. Tromp. 1999. Introduction to the spectral-element method for 3-D seismic wave propagation. *Geophys. J. Int.* 139 (1999), 806–822.
 - [17] Lion Krischer, Andreas Fichtner, Christian Boehm, and Heiner Igel. 2018. Automated Large-Scale Full Seismic Waveform Inversion for North America and the North Atlantic. *Journal of Geophysical Research: Solid Earth* 123, 7 (2018), 5902–5928.
 - [18] Mirosław Kupczyk. 2017. Urgent Computing service in the PRACE Research Infrastructure. <https://doi.org/10.5281/zenodo.832029>
 - [19] S Pondrelli, S Salimbeni, A Morelli, G Ekström, L Postpischl, G Vannucci, and E Boschi. 2011. European–Mediterranean regional centroid moment tensor catalog: solutions for 2005–2008. *Physics of the Earth and Planetary Interiors* 185, 3–4 (2011), 74–81.
 - [20] Anne Reinarz, Dominic E Charrier, Michael Bader, Luke Bovard, Michael Dumbser, Kenneth Duru, Francesco Fambri, Alice-Agnes Gabriel, Jean-Mathieu Gallard, Sven Köppel, et al. 2019. ExaHyPE: an engine for parallel dynamically adaptive simulations of wave problems. *arXiv preprint arXiv:1905.07987* (2019).
 - [21] Christopher H Scholz. 2019. *The mechanics of earthquakes and faulting*. Cambridge university press.
 - [22] Gerardo Suárez, Torild van Eck, Domenico Giardini, Tim Ahern, Rhett Butler, and Seiji Tsuboi. 2008. The international federation of digital seismograph networks (FDSN): An integrated system of seismological observatories. *IEEE Systems Journal* 2, 3 (2008), 431–438.
 - [23] Shunsuke Takemura, Takanori Matsuzawa, Takeshi Kimura, Takashi Tonegawa, and Katsuhiko Shiomi. 2018. Centroid moment tensor inversion of shallow very low frequency earthquakes off the Kii Peninsula, Japan, using a three-dimensional velocity structure model. *Geophysical Research Letters* 45, 13 (2018), 6450–6458.
 - [24] Stefania Tarantino, Simona Colombelli, Antonio Emolo, and Aldo Zollo. 2019. Quick Determination of the Earthquake Focal Mechanism from the Azimuthal Variation of the Initial P-Wave Amplitude. *Seismological Research Letters* (2019).
 - [25] Enric Tejedor, Yolanda Becerra, Guillem Alomar, Anna Queral, Rosa M Badia, Jordi Torres, Toni Cortes, and Jesús Labarta. 2017. PyCOMPSs: Parallel computational workflows in Python. *The International Journal of High Performance Computing Applications* 31, 1 (2017), 66–82.
 - [26] Nikolaos Triantafyllis, Efthimios Sokos, Aristidis Ilias, and Jiri Zahradnik. 2016. Scisola: automatic moment tensor solution for SeisComp3. *Seismological Research Letters* 87, 1 (2016), 157–163.
 - [27] David J Wald, C Bruce Worden, Vincent Quitoriano, and James Goltz. 2002. Shakemap: its role in pre-earthquake planning and post-earthquake response and information. In *Proceedings SMIP02 Seminar on Utilization of Strong-Motion Data*. Citeseer, New York, 1–20.