3D Acoustic-Elastic Coupling with Gravity: The Dynamics of the 2018 Palu, Sulawesi Earthquake and Tsunami

Lukas Krenz
lukas.krenz@in.tum.de
Technical University of Munich
Garching, Germany

Carsten Uphoff
Thomas Ulrich
Alice-Agnes Gabriel
uphoff@geophysik.uni-muenchen.de
ulrich@geophysik.uni-muenchen.de
gabriel@geophysik.uni-muenchen.de
Ludwig-Maximilians-Universität München
Munich, Germany

Lauren S. Abrahams
Eric M. Dunham
labraha2@stanford.edu
edunham@stanford.edu
Stanford University
Stanford, CA, USA

Michael Bader
bader@in.tum.de
Technical University of Munich
Garching, Germany

ABSTRACT

We present a highly scalable 3D fully-coupled Earth & ocean model of earthquake rupture and tsunami generation and perform the first fully coupled simulation of an actual earthquake-tsunami event and a 3D benchmark problem of tsunami generation by a megathrust dynamic earthquake rupture. Multi-petascale simulations, with excellent performance demonstrated on three different platforms, allow high-resolution forward modeling. Our largest mesh has ≈ 261 billion degrees of freedom, resolving at least 15 Hz of the acoustic wave field. We self-consistently model seismic, acoustic and surface gravity wave propagation in elastic (Earth) and acoustic (ocean) materials sourced by physics-based non-linear earthquake dynamic rupture, thereby gaining insight into the tsunami generation process without relying on approximations that have previously been applied to permit solution of this challenging problem. Complicated geometries, including high-resolution bathymetry, coastlines and segmented earthquake faults are discretized by adaptive unstructured tetrahedral meshes. This inevitably leads to large differences in element sizes and wave speeds which can be mitigated by ADER local time-stepping and a Discontinuous Galerkin discretization yielding high-order accuracy in time and space.

CCS CONCEPTS

• Applied computing → Earth and atmospheric sciences.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SC'21, Nov 14–19, 2021, St. Louis, MO

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8442-1/21/11...\$15.00

<https://doi.org/10.1145/3458817.3476173>

KEYWORDS

ADER-DG, Elastic-Acoustic-Coupling, Earthquake Simulation, Computational Seismology, SeisSol, Tsunami Generation

ACM Reference Format:

Lukas Krenz, Carsten Uphoff, Thomas Ulrich, Alice-Agnes Gabriel, Lauren S. Abrahams, Eric M. Dunham, and Michael Bader. 2021. 3D Acoustic-Elastic Coupling with Gravity: The Dynamics of the 2018 Palu, Sulawesi Earthquake and Tsunami. In *SC'21: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, Nov 14–19, 2021, St. Louis, MO. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3458817.3476173>

1 INTRODUCTION

Figure 1 illustrates a 3D fully coupled model of the tsunami caused by the magnitude 7.5 strike-slip supershear earthquake that struck Palu Bay of the island of Sulawesi, Indonesia, on September 28, 2018. Tsunamis occur due to abrupt vertical perturbations to the water column. Devastating tsunamis are therefore rarely caused by strike-slip earthquakes (such as the Sulawesi earthquake), as these produce predominantly horizontal displacements. What was additionally unusual about the Sulawesi earthquake was its ability to produce very fast (“supershear”) movements during the rupture itself. This surprising earthquake-tsunami behavior has led to an extensive debate with important implications for submarine strike-slip fault systems worldwide (Mai [37] and references therein). However, efforts to understand this behavior have been hampered by the necessity of making approximations in the fluid-solid coupling, and hence in the tsunami generation process.

Recent simulations that link physics-based earthquake dynamics and elastic (seismic) wave propagation to shallow water models of tsunami genesis and propagation [60] suggest that the Palu tsunami can be primarily understood from the complex interaction of earthquake rupture and the 3D bathymetry of the bay, and that, while the fast rupture speed strongly influences seismic radiation and ground shaking from the earthquake, it seemingly does not affect the tsunami response. However, the standard practice of

Figure 1: 3D fully coupled acoustic-elastic model with gravity of the 2018 Sulawesi supershear earthquake and local tsunami in Palu Bay. (a) Map view of the vertical sea surface velocity at 15 s simulation time. The star indicates the earthquake epicenter. Black lines mark the traces of a system of complex faults. (b) Sea surface height (m, same color map as in (d)) at 60 s (warped vertical displacements in the bay, amplified $\times 2000$). (c) 3D view of earthquake and tsunami model in Palu Bay. Particle velocity (slip rate) across the faults, vertical sea-surface and Earth ground velocity at 15 s. Dynamic earthquake rupture, propagating faster than the shear-wave velocity of the surrounding rocks, generates a clearly visible Mach cone altering the acoustic response of the ocean. (d) Sea surface height (vertical displacements in Palu Bay) at 15 s.

linking a 3D earthquake simulation, conducted without a water layer, to a 2D tsunami shallow water model simplifies the interaction between the ocean floor and water column [35].

Typical limitations of the one-way linking procedure and use of a shallow water model include neglecting dispersion (which can be somewhat addressed using Boussinesq type equations [52]), neglecting horizontal momentum transfer [34, 55] and assuming that the ocean is incompressible [5]. The latter assumption eliminates ocean acoustic waves, which, as our modeling demonstrates for the first time for an actual earthquake-tsunami event, can create perturbations of the ocean surface of comparable amplitude to tsunami waves but at much higher frequencies. Furthermore, at early times in tsunami generation, acoustic waves and tsunami gravity waves are superimposed, creating a complex wavefield that can be measured with ocean bottom pressure sensors and off-shore cabled sensor networks deployed in the earthquake source region with potential for improving tsunami early warning [26, 53, 67].

In this paper, we present a 3D fully-coupled elastic-acoustic model of the solid Earth and ocean response, with gravity, to earthquake dynamic rupture. Implemented in the earthquake simulation package SeisSol (www.seissol.org), the model rigorously, and without the approximations employed in one-way linking procedures, captures the entire process from earthquake rupture to the generation and propagation of seismic waves, ocean acoustic waves, and tsunamis. We present a large-scale fully-coupled model of the 2018 Sulawesi event that links the dynamics of supershear earthquake faulting to elastic and acoustic waves in Earth and ocean to tsunami gravity wave propagation in the narrow Palu Bay. We demonstrate scalability and performance of the MPI+OpenMP parallelization on three petascale supercomputers. The achieved efficiency of our approach allows high-resolution, data-integrated forward modeling: the dynamic rupture earthquake model fuses seismological,

geodetic, tectonic and field observations and the tsunami model matches wave gauge data and inundation mapping [60].

Within the heterogeneous solid Earth, we solve the elastic wave equation coupled to non-linear frictional sliding across a complex fault network. Simultaneously, we solve the acoustic wave equation, describing perturbations about an equilibrium hydrostatic state, in a compressible, inviscid ocean of variable depth. The effects of gravitational restoring forces, which are responsible for tsunami propagation, are efficiently incorporated through a modification of the standard free surface boundary condition at the evolving sea surface [32]. The enhanced modeling capabilities and gained geophysical insight come at the cost of unforeseen implementation complexity and substantially higher computational demands. In particular, we make the following contributions:

- In Secs. 3 and 4, we present an Arbitrary high-order DERivative Discontinuous Galerkin (ADER-DG) discretization of the first 3D fully-coupled elastic-acoustic earthquake-tsunami model. We explain how the moving sea surface can be tracked by a linearized free surface boundary condition and discretize it efficiently.
- The fully-coupled model is implemented within the open-source software SeisSol. In Sec. 5, we present node-level performance (esp. NUMA effects) and enhancements to SeisSol's hybrid parallelization scheme, and we discuss load-balancing challenges of the fully-coupled model.
- In Sec. 6.1, we evaluate the fully-coupled model in a 3D benchmark setup of tsunami generation by a dynamic earthquake rupture, and compare against a shallow water one-way linking simulation approach.
- We simulate a large-scale 3D fully-coupled elastic-acoustic scenario of the 2018 Palu earthquake and tsunami (Sec. 6.2)

and show that we capture salient features of tsunami genesis. Our largest setup has 518 million elements (261 billion degrees of freedom). We achieve sustained petascale performance for this setup on two supercomputers, *Shaheen-II* and *SuperMUC-NG*.

- In Sec. 6.3, we demonstrate that our method achieves a parallel efficiency of $>70\%$ on the *Mahti* and *SuperMUC-NG* supercomputers (based on AMD and Intel architectures, from 50 to 700 and from 50 to 1600 nodes, respectively), examining the strong-scaling limit and focusing on NUMA-aware allocation of processes to nodes.

2 RELATED WORK

Tsunami modeling typically uses one-way linking of the earthquake-induced seafloor uplift to the ocean response. This is mostly justified when the earthquake happens so quickly that the tsunami propagation distance during the rupture duration is much smaller than the source dimension. In this approximation, the final, static seafloor uplift is utilized as an initial condition for the tsunami. The seafloor uplift is furthermore commonly simplified by using analytical solutions for the displacement caused by fault slip within a homogeneous elastic half-space (Okada [40]). The long-wavelength components of the seafloor uplift are then assumed to instantaneously uplift the water column. In reality, the seafloor moves dynamically in response to sources such as earthquakes and landslides, generating seismic waves, ocean acoustic waves and tsunami gravity waves. While these waves eventually separate at a sufficiently large distance from the source, they are superimposed in the source region. Offshore instrument deployments in source regions of potential events have motivated the development of methods that capture the full wavefield. Maeda et al. [36] developed the first such coupled model in their study of the 2011 Tohoku, Japan, event, utilizing 3D kinematic earthquake sources. Saito et al. [50] present an alternative method in which the full wavefield is constructed as the superposition of a solution conducted without gravity, but including a compressible ocean, and a second solution to the tsunami problem using forcing from the zero-gravity solution.

Another approach, which we take in this work, is to directly solve the linearized equations of motions governing smaller perturbations of the ocean about a rest state in hydrostatic equilibrium, with rigorous acoustic-elastic coupling to the solid Earth across the deforming seafloor interface. Thus far, this approach, which involves a modification of the free surface condition, has only been realized in 2D [32–34]. These generic 2D studies reveal the fundamentally complex, dynamic nature of the tsunami generation process, which superimposes the initial tsunami waveforms with ocean acoustic waves, oceanic Rayleigh waves, and other seismic waves in simple model setups. Of course, this method must be extended to 3D in order to assimilate and interpret observational data, to better understand the earthquake rupture and tsunami generation process and to improve tsunami early warning capabilities. The fully-coupled method allows for realistic geometries and rupture processes, providing a deeper understanding of the relation between seismic and acoustic amplitudes and periods and tsunami heights, and also serves as a reference solution to which simpler, approximate methods can be compared.

3D high-resolution simulation of earthquakes (including tsunami-genic events) and seismic shaking allow geophysical insight and mitigation of natural hazards, even when complicated by sparse data coverage [39]. Integration with geophysical and geological constraints, such as complex structure and topography, requires massive computational performance, demonstrated in several petascale simulation efforts during the last decade. For example, Roten et al. [49] present non-linear ground motion models with off-fault plasticity on the Blue Waters and Titan supercomputers; Fu et al. [15] and Chen et al. [7] model the Wenchuan Earthquake, incorporating high-resolution surface topography on Sunway TaihuLight. Recent work using the SeisSol software integrates dynamic rupture models of tsunamigenic earthquakes with seismic, geodetic and tsunami data [59, 60, 62]. Ichimura et al. [20, 21] present coupled finite-element simulations of earthquake-induced ground shaking and their impact on urban structures. Rodgers et al. [45, 46] perform broadband anelastic ground motion scenario simulations in the San Francisco Bay Area.

3 MODEL DESCRIPTION

In the following, we lay out the 3D elastic-acoustic model, including interface and boundary conditions, used for our simulations. We use Einstein summation convention, i.e., summation over indices appearing twice is implied.

Earth and rupture model: In the solid Earth, we solve the seismic wave equation in velocity-stress formulation. With the unknowns $\mathbf{q} = (\sigma_{xx}, \sigma_{yy}, \sigma_{zz}, \sigma_{xy}, \sigma_{yz}, \sigma_{xz}, v_x, v_y, v_z)$, i.e., six stress and three velocity components, and using indices $i, j \in \{x, y, z\}$, we can state our set of equations as

$$\begin{aligned} \frac{\partial \sigma_{ij}}{\partial t} - \lambda \delta_{ij} \frac{\partial v_k}{\partial x_k} - \mu \left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} \right) &= 0, \\ \rho \frac{\partial v_i}{\partial t} - \frac{\partial \sigma_{ij}}{\partial x_j} &= 0. \end{aligned} \quad (1)$$

Here, μ, λ and ρ are the spatially-varying material properties and δ_{ij} is the Kronecker delta.

An earthquake in the domain Ω_E is modeled as frictional shear failure on a pre-defined infinitesimally thin but potentially geometrically complex fault $\Gamma_F \subset \partial\Omega_E$. The following boundary conditions have to be satisfied on Γ_F :

$$\begin{aligned} |\tau| &\leq \tau_S, \\ \tau_S V_i - \tau_i |V| &= 0, \\ \tau_S &= s f(|V|, \psi), \\ \frac{d\psi}{dt} &= g(|V|, \psi), \end{aligned} \quad (2)$$

where $s = n_i \sigma_{ij} n_j$, $\tau_i = (\delta_{ij} - n_i n_j) \sigma_{jk} n_k$, $V_i = (\delta_{ij} - n_i n_j) (v_j^+ - v_j^-)$, n is the unit normal pointing from the “−”-side to the “+”-side, and $|\cdot|$ is the vector length. The function f returns the coefficient of friction that depends on the slip-rate $|V|$ and the state variable ψ . The evolution of ψ in time is controlled by the function g [41].

Ocean model: Following Lotto and Dunham [32], the ocean dynamics are modeled as a small perturbation about the hydrostatic state, that is, pressure is given by $p(x, y, z, t) = p_0(z) + p'(x, y, z, t)$. We adopt an Eulerian description with the sea level at $z = 0$. The positive z -axis points up. The equations of momentum and mass

balance are linearized about the hydrostatic state such that

$$\begin{aligned} \frac{\partial p'}{\partial t} + K \frac{\partial v_k}{\partial x_k} &= \rho g v_z, \\ \rho \frac{\partial v_i}{\partial t} + \frac{\partial p'}{\partial x_i} &= -\delta_{iz} \frac{\rho g}{K} p, \end{aligned} \quad (3)$$

where K is the bulk modulus. Note that in Eq. (3) we use the sign convention that pressure is positive in compression; in contrast, stress is positive in tension in Eqs. (1) and (2). We identify Eq. (3) as the equations of linear acoustics, if the right-hand side is set to zero. The terms on the right-hand are neglected because they are several orders of magnitude smaller than the terms that are retained [32].

The pressure on the moving sea surface must be equal to the atmospheric pressure p_a , i.e.,

$$p(x, y, \eta(x, y, t), t) = p_a, \quad (4)$$

where η is the displacement of the sea surface. By linearizing and inserting the pressure gradient in the hydrostatic state, we can simplify Eq. (4) as [32]:

$$p(x, y, 0, t) \approx p_a + \rho g \eta(x, y, t), \quad (5)$$

where $g = 9.81 \text{ m/s}^2$. Thus, we impose the following boundary conditions at $z = 0$:

$$p'(x, y, z = 0, t) = \rho g \eta, \quad (6)$$

$$\frac{\partial \eta(x, y, z = 0, t)}{\partial t} = u_z(x, y, z = 0, t). \quad (7)$$

In this way, our model incorporates gravitational effects via a modification of the surface boundary condition, which is applied on the equilibrium ocean surface, thereby avoiding the need for a time-dependent mesh that tracks the moving surface.

In summary, our 3D fully-coupled model includes earthquake rupture dynamics, wave propagation in both elastic and acoustic media, as well as tsunami propagation. The model naturally accounts for effects of compressibility, as well as non-hydrostatic ocean response that leads to dispersion and related effects during tsunami propagation and generation.

4 DISCRETIZATION & IMPLEMENTATION

We implement the model from Sec. 3 in the high-order earthquake simulation software SeisSol (e.g., [17, 62]). We first summarize the ADER-DG method used in SeisSol and explain how we extend it to include the acoustic-elastic coupling and gravitational free surface boundary condition. Then, we explain details of SeisSol's local-time-stepping approach and optimized parallel implementation, and how we utilize these for our coupled model. The resulting implementation is the first to allow for 3D fully-coupled elastic-acoustic simulations that capture the full dynamics of the earthquake process, seismic wave propagation, the displacement of seafloor and surface and the onset of tsunami propagation.

4.1 The ADER-DG method

The acoustic wave equation (3) becomes a special case of Eq. (1) by setting $K = \lambda$, $\mu = 0$, and $\sigma_{ij} = -p\delta_{ij}$. Hereafter we use p to denote the pressure perturbation p' . We may thus write both

Eqs. (1) and (3) in the general non-conservative hyperbolic form

$$\frac{\partial \mathbf{q}}{\partial t} + \mathbf{A} \frac{\partial \mathbf{q}}{\partial x} + \mathbf{B} \frac{\partial \mathbf{q}}{\partial y} + \mathbf{C} \frac{\partial \mathbf{q}}{\partial z} = 0, \quad (8)$$

where $(\mathbf{A}, \mathbf{B}, \mathbf{C}) \in \mathbb{R}^{9 \times 9}$ are the space-dependent Jacobian matrices. We note that embedding the acoustic equations into the equations of elasticity entails a computation and memory overhead, however, existing data-structures in our inherited code-base are unaffected and as such the incorporation of the ocean model is less invasive.

We describe the Discontinuous Galerkin (DG) method briefly in the following: The domain Ω is approximated using conforming tetrahedral meshes $\mathcal{T}_h = \{K\}$. We define the usual broken finite element space

$$W_h := \{\mathbf{p} \in [L^2(\Omega)]^9 : \mathbf{p}|_K \in [\mathcal{P}_N(K)]^9 \quad \forall K \in \mathcal{T}_h\}, \quad (9)$$

where $\mathcal{P}_N(K)$ is the space of polynomials on K of maximum degree N . The space-dependent Jacobians $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are element-wise constant. We derive the semi-discrete form of Eq. (8) by multiplying with a test function $\mathbf{p} \in W_h$, integrating over an element K and by integration by parts:

$$\begin{aligned} \int_K \mathbf{p} \cdot \frac{\partial \mathbf{q}}{\partial t} dV - \int_K \left(\frac{\partial \mathbf{p}}{\partial x} \mathbf{A} + \frac{\partial \mathbf{p}}{\partial y} \mathbf{B} + \frac{\partial \mathbf{p}}{\partial z} \mathbf{C} \right) \mathbf{q} dV \\ + \int_{\partial K} \mathbf{p} \cdot (n_x \mathbf{A} + n_y \mathbf{B} + n_z \mathbf{C}) \mathbf{q}^* dS = 0, \end{aligned} \quad (10)$$

where $(n_x \mathbf{A} + n_y \mathbf{B} + n_z \mathbf{C}) \mathbf{q}^*$ is the numerical flux, which we discuss later in more detail, (n_x, n_y, n_z) is the outward-pointing normal vector. The trial space is the same as the test space. Thus, using Dubiner's basis ϕ_l for $\mathcal{P}_N(K)$ we can expand $\mathbf{q}$ as

$$\mathbf{q}(t, \mathbf{x}) = \sum_{l=1}^{B_N} \mathbf{Q}_l(t) \phi_l(\mathbf{x}), \quad \text{where } B_N = \binom{N+3}{3}. \quad (11)$$

For the efficient implementation of Eq. (10), we map every element to a reference element. The resulting integrals only depend on the reference element and are pre-computed, leading to a quadrature-free computation of stiffness and flux terms. We refer the interested reader to [11, 62] for further details.

We employ ADER time-stepping, which is a predictor-corrector approach that results in an efficient one-step update [11, 23, 57]. In the predictor step, we expand the cell-local solution in time

$$\mathbf{Q}_l(t) \approx \sum_{k=0}^N \frac{(t - t_0)^k}{k!} \frac{\partial^k \mathbf{Q}_l}{\partial t^k}(t_0) \quad (12)$$

around the beginning of our timestep t_0 . We compute the time derivatives of $\mathbf{Q}_l$ with the discrete Cauchy-Kowalesky procedure [28], i.e., time-derivatives are replaced by spatial derivatives using the PDE recursively. We obtain the corrector step by inserting Eq. (12) into the time-integrated Eq. (10).

4.2 Elastic-acoustic coupling

In Eq. (10), the numerical flux, i.e., the term $(n_x \mathbf{A} + n_y \mathbf{B} + n_z \mathbf{C}) \mathbf{q}^*$, must be chosen carefully to ensure consistency and stability of our ADER-DG scheme. We choose the Godunov numerical flux, i.e., we solve the one-dimensional Riemann problem exactly for all basis functions on the faces [30, 58].

Figure 2: Solution structure of the Riemann problem in coupled elastic-acoustic media. The eigenstructure implies that we have one left-going P-wave and two left-going S-waves in the elastic medium. In contrast, only one right-going acoustic wave is present in the acoustic medium.

Due to the eigenstructure of Eq. (1), we have one (compressional) P-wave and two (shear) S-waves in the elastic domain, traveling with speeds c_p^- and c_s^- , respectively. In the acoustic domain, we have one P-wave (acoustic wave), traveling with speed c_p^+ . The resulting solution structure at an elastic-acoustic interface is shown in Fig. 2. The jumps $q^a - q^-$ and $q^b - q^+$ have to satisfy the Rankine-Hugoniot conditions

$$\begin{aligned}\hat{A}^-(q^- - q^a) &= -c_p^-(q^- - q^a), \\ \hat{A}^-(q^a - q^b) &= -c_s^-(q^a - q^b), \\ \hat{A}^+(q^c - q^+) &= c_p^+(q^c - q^+),\end{aligned}\quad (13)$$

where $\hat{A}^\pm = n_x A^\pm + n_y B^\pm + n_z C^\pm$. Note that $\hat{A}^-$ denotes the rotated Jacobian of the elastic domain with Lamé parameters λ^-, μ^- , whereas $\hat{A}^+$ denotes the rotated Jacobian of the acoustic domain with bulk modulus $\lambda^+ = K^+$.

Equation (13) is an eigenproblem. Therefore, we can state the jumps as a linear combination of eigenvectors

$$\begin{aligned}q^- - q^b &= \alpha_1 \hat{r}_1^- + \alpha_2 \hat{r}_2^- + \alpha_3 \hat{r}_3^-, \\ q^c - q^+ &= \alpha_4 \hat{r}_4^+, \end{aligned}\quad (14)$$

where $(-c_p^-, \hat{r}_1^-)$, $(-c_s^-, \hat{r}_2^-)$, $(-c_s^-, \hat{r}_3^-)$ are eigenpairs of $\hat{A}^-$ and $(c_p^+, \hat{r}_4^+)$ is an eigenpair of $\hat{A}^+$. In practice, it is not necessary to find these eigenpairs for every instance of the normal vector $\mathbf{n}$, as both acoustic and elastic wave equations are rotational invariant. The resulting similarity transform, represented by the matrix T [11], leads to the identity

$$\hat{A} = n_x A + n_y B + n_z C = T(\mathbf{n}) A T^{-1}(\mathbf{n}). \quad (15)$$

Thus, $\hat{r}_i = T \mathbf{r}_i$, where $\mathbf{r}_i$ is an eigenvector of A . Multiplying Eq. (14) with T^{-1} and letting $\mathbf{w} = T^{-1} \mathbf{q}$ gives

$$\begin{aligned}\mathbf{w}^- - \mathbf{w}^b &= \alpha_1 \mathbf{r}_1^- + \alpha_2 \mathbf{r}_2^- + \alpha_3 \mathbf{r}_3^-, \\ \mathbf{w}^c - \mathbf{w}^+ &= \alpha_4 \mathbf{r}_4^+.\end{aligned}\quad (16)$$

Equation (14) constitutes an under-determined linear system of equations. In order to close the system of equations, we require continuity of traction and continuity of the normal component of the velocity, as appropriate for coupling an elastic solid to an inviscid fluid [54]. Here, these interface conditions are $w_i^b = w_i^c$, $i = 1, 7$ and $w_i^b = w_i^c = 0$, $i = 4, 6$.

Solving the Riemann problem gives the boundary state

$$\begin{aligned}w_1^b &= w_1^- - \alpha_1, & w_7^b &= w_7^- - \frac{1}{Z_p^-} \alpha_1, \\ w_4^b &= 0, & w_8^b &= w_8^- - \frac{1}{Z_s^-} \alpha_4, \\ w_6^b &= 0, & w_9^b &= w_9^- - \frac{1}{Z_s^-} \alpha_4,\end{aligned}\quad (17)$$

where $Z_p = c_p \rho$, $Z_s = c_s \rho$, and

$$\alpha_1 = \frac{Z_p^- Z_p^+}{Z_p^- + Z_p^+} \left(\frac{w_1^- - w_1^+}{Z_p^+} + w_7^- - w_7^+ \right). \quad (18)$$

The components w_8^b and w_9^b are the tangential particle velocities. We observe that the tangential traction components w_4^b, w_6^b penalize the tangential particle velocities, thus, weakly enforcing that shear tractions must be zero on the elastic-acoustic interface. Note that the components w_2^b, w_3^b, w_5^b are irrelevant as they do not contribute to the numerical flux due to the non-zero structure of A .

It is straightforward to see that the boundary state is consistent with the physical interface conditions by inserting $w_i^- = w_i^+$, $i = 1, 7$, and $w_i^- = w_i^+ = 0$, $i = 4, 6$ in Eq. (17). Failing to ensure consistency, e.g., if a flux is used that only takes material parameters from one side into account, as e.g., in [11, 29], may lead to a non-converging scheme when coupling elastics and acoustics [64].

In total, the numerical flux is given by

$$\hat{A}^- \mathbf{q}^* = T A^- \mathbf{w}^b (T^{-1} \mathbf{q}^-, T^{-1} \mathbf{q}^+). \quad (19)$$

In our implementation, we exploit that $\mathbf{w}^b$ is bilinear and rewrite Eq. (19) as

$$\hat{A}^- \mathbf{q}^* = F^- \mathbf{q}^- + F^+ \mathbf{q}^+. \quad (20)$$

The 9×9 matrices F^- and F^+ can be pre-computed for each face. Thus, we can use the exact Riemann solver for little additional cost.

4.3 Boundary conditions

We impose the boundary conditions weakly, i.e., through the numerical flux. We construct ghost values $\mathbf{q}^+$ such that the solution of the Riemann problem results in the correct state at the boundary. The solution of the “inverse Riemann problem” for boundary condition Eq. (6) is given by

$$u_n^+ = v_n^-, \quad p^+ = 2\rho g \eta - p^-, \quad (21)$$

where v_n is the velocity in normal direction. Solving the Riemann problem with the ghost state and interior state as input, we obtain

$$v_n^b = v_n^- - \frac{1}{Z} (\rho g \eta - p^-), \quad p^b = \rho g \eta, \quad (22)$$

where the superscript b denotes the middle state (cf. Fig. 2) and $Z = c_p \rho$ is the impedance. The value of v_n^b can be interpreted as an extrapolation of the interior value v_n^- and a penalty term that measures the difference between the pressure at the boundary p^- and the actual pressure $\rho g \eta$.

The displacement η , cf. Eq. (7), is evolved with

$$\frac{\partial \eta}{\partial t} = v_n^b = v_n^- - \frac{1}{Z} (\rho g \eta - p^-). \quad (23)$$

It is critical to use the velocity v_n^b here (instead of, e.g., v_n^-) as only then we have a stable scheme.

For the ADER-DG scheme, we need to integrate the boundary state in time. We do so by integrating the coupled ODE system

$$\begin{aligned}\frac{\partial \eta}{\partial t} &= v_n^- - \frac{1}{Z}(\rho g \eta - p^-), \\ \frac{\partial H}{\partial t} &= \eta,\end{aligned}\quad (24)$$

with initial conditions for each timestep

$$\eta(t_n) = \eta^n, \quad H(t_n) = 0, \quad (25)$$

where η^n is the displacement at the beginning of the timestep. In this way, we can compute the time-integrated displacement at the same time as the displacement. The time-integrated boundary states are obtained with

$$\int_{t_n}^{t_{n+1}} u_n^b dt = \eta^{n+1} - \eta^n, \quad \int_{t_n}^{t_{n+1}} p^b dt = \rho g H(t^{n+1}). \quad (26)$$

The time-integration of Eq. (24) can be performed using standard ODE solvers. We use Verner's "most efficient" Runge-Kutta scheme of order 7 [63]. In the Runge-Kutta stages, we need to evaluate the right-hand side of Eq. (24) at several points in time. In order to evaluate v_n^- and p^- at arbitrary times in the interval $[t_n, t_{n+1}]$, we use Eq. (12) to predict the evolution of v and p in the volume and then extrapolate the volume data to the boundary.

Note that this implementation is more efficient than computing the analytical solution of the boundary ODE in Eq. (23) and subsequent integration. This would require computing integrals (e.g., of the pressure) for each time quadrature point, i.e., a nested quadrature scheme. As the evaluation of v^- , p^- is the most expensive part of computing the boundary condition, this would incur significant computational costs without increasing the accuracy.

4.4 Local time-stepping

The stability of the explicit ADER-DG scheme is controlled by the (Courant–Friedrichs–Lewy) CFL condition [8, 10]

$$\Delta t \leq C(N)h(|\lambda^{\max}|)^{-1}, \quad (27)$$

where $C(N) \leq (2N+1)^{-1}$ is a constant that depends on the polynomial order and h is the maximum diameter of the insphere of the tetrahedron. The maximum eigenvalue of the Jacobian $\lambda^{\max} = \sqrt{(\lambda + 2\mu)/\rho}$ is the P-wave speed (speed of sound in the acoustic layer) and thus only depends on the material parameters. Considering geophysically motivated realistic scenarios with material heterogeneity, geometric complexity and statically adaptive meshes leads inevitably to large differences in element sizes and wave speeds which can be mitigated by local time-stepping [6, 43, 44, 47, 62]. Our proposed acoustic-elastic coupling adds to this challenge, as the wave speeds in the Earth and the ocean differ significantly. Furthermore, incorporating shallow water depths, as in the case of the Palu tsunami, demand locally fine spatial discretizations. It is easy and efficient to implement local time-stepping (LTS) with ADER due to the inherent polynomial-in-time representation of our discrete solution [12]. LTS was also shown to reduce dispersion and dissipation errors associated with unnecessarily small timesteps [12, 13]. SeisSol provides a multi-rate local time-stepping implementation, which clusters elements with similar timesteps [6, 62], such that the CFL condition of no element is violated and that each element

is within the cluster with the largest admissible timestep size. Parallelization thus works over batches of elements and avoids the otherwise erratic element-wise time-stepping. We chose rate-2 LTS, such that the first cluster contains all elements with the smallest overall timestep size, the second cluster all elements with a timestep size larger than $2(\Delta t)_{\min}$, and so on. In our experiments (see Sec. 6), we find that LTS has an even stronger influence on time-to-solution than reported previously for earthquake simulations, which we attribute to the discrepancy of wave speeds and mesh resolution in the acoustic and elastic layer. See also the histogram of timestep sizes in Fig. 4.

5 HPC ASPECTS

Our coupled elastic-acoustic model is integrated in the software package SeisSol that is optimized for earthquake simulations at petascale [17, 62]. Here, we briefly summarize SeisSol's optimization and parallelization strategy and outline which additional contributions were necessary to realize the elastic-acoustic landmark simulations of Sec. 6: (i) we demonstrate strong NUMA effects of node-level performance, especially on AMD Rome, (ii) we improve support for using multiple MPI ranks per node, via careful pinning of OpenMP threads and POSIX threads used for communication, and (iii) we implement performance-aware static load-balancing on heterogeneous supercomputers.

5.1 Node-level performance

High-order ADER-DG schemes spend the majority of the computational time in element-local kernels, most of which consist of linear algebra operations of small size. These kernels can be expressed as a sequence of tensor operations, potentially involving sparse matrices. SeisSol implements these tensor expressions using the code generator YATeTo [61]. YATeTo transforms tensor expressions into sequences of small GEMMs whenever possible.

Small-GEMM subroutines are generated by hardware-specific backends, such as LIBXSMM [18] or SeisSol's PSpaMM¹. On the Sky-lake architecture, SeisSol uses both generators: PSpaMM for dense-sparse matrix-matrix multiplications and LIBXSMM for dense-dense matrix-matrix multiplications. On the AMD Rome architectures we use LIBXSMM and generate code for the Intel Haswell architecture which results in efficient assembler code.²

Performance is evaluated by a performance reproducer for the wave propagation part. We expect to see NUMA effects in the corrector step due to unstructured memory access over NUMA boundaries, but not in the predictor step, which only requires element-local data. We tested this hypothesis on a dual-socket AMD Rome 7H12 system with base frequency of 2.6 GHz. Each CPU has 64 cores and 4 NUMA nodes. For all benchmarks the entire node is used to ensure that the boost frequency (3.3 GHz) is not triggered. We thus assume a peak performance of 5325 GFLOPS per node. Running only the predictor step of our scheme, we get a performance of 3360 GFLOPS (63% of peak) on the entire node and a performance of 428 GFLOPS on a single NUMA node. Extrapolating the result on a single NUMA node predicts a performance limit of 3424 GFLOPS (64% of peak), i.e., NUMA has little effect on the predictor step. Running both

¹<https://github.com/peterwauligmann/PSpaMM>

²PSpaMM only supports code generation for AVX512.

predictor and corrector step of the scheme, we get a performance of 2053 GFLOPS (38% of peak) on the entire node and 376 GFLOPS on a single NUMA node. Extrapolation of the latter result predicts a performance limit of 3008 GFLOPS (56% of peak). Note that we also ran this on one entire socket and get a performance of 1390 GFLOPS (52% of peak). Thus, we confirm strong NUMA effects in the corrector step on this architecture.

5.2 Hybrid parallelization

SeisSol is parallelized using MPI, OpenMP, and POSIX threads (pthreads). On each MPI partition, OpenMP is used to parallelize loops on the level of time clusters, i.e., bulk-synchronous loops over elements and dynamic rupture faces are distributed with OpenMP for all elements or dynamic rupture faces in one time cluster. In addition, we use a communication thread (implemented with pthreads) to ensure MPI progression and truly asynchronous communication [6, 19]. Pthreads are also created for asynchronous I/O [42].

The default operation mode of SeisSol has been to use one rank per node with one communication thread and NUMA-aware memory initialization. This mode of operation resulted in best performance on previous generation supercomputers, but the results in Sec. 5.1 suggest that multiple ranks per node increase performance on systems with many NUMA domains.

Communication (and I/O) threads must not be pinned to worker threads, as the heavy use of MPI_Test within the communication thread would interfere with the worker. We develop a simple and portable algorithm to correctly pin communication threads when using multiple ranks per node. Affinity of worker threads is controlled via the OMP_PLACES and OMP_PROC_BIND environment variables. We set the number of OpenMP threads to leave one physical core per MPI rank unused. Next, we pin communication threads to free cores using sched_setaffinity. To determine free cores, each rank first determines which CPU numbers are occupied by worker threads and saves them in a CPU mask. The CPU mask is reduced on each compute node by splitting the MPI communicator with MPI_COMM_TYPE_SHARED. We then use libnuma to compute the list of NUMA domains covered by the worker threads. Finally, we pin the communication and I/O threads to all free logical cores that are also in the list of used NUMA domains. We thus ensure that pinning of pthreads is NUMA aware and does not interfere with the worker threads. Note that throughout this study, we use simultaneous multithreading, i.e. two threads per physical core, as this results in better performance for our code.

5.3 Load balancing

SeisSol's static load balancing relies on graph-based partitioning. The dual graph of a tetrahedral mesh is built by defining a vertex for each element and an edge for any two elements that share a face. Vertex and edge weights model computation and communication cost, respectively. SeisSol uses ParMETIS to solve the resulting graph partitioning problem [51]. Vertex and edge weights are crucial for SeisSol's local time stepping scheme, as the element's update-rate is directly proportional to the element's computation cost [6].

New challenges raised by the fully-coupled model include:

- Our friction law requires the solution of a non-linear system of equations, cf. Eq. (2). The number of Newton iterations varies over time such that the computational load is dynamic.
- The gravitational boundary requires additional operations due to the integration of the displacement in time. As discussed in Sec. 4.3, we have to solve an ODE on the faces. Each Runge-Kutta stage requires an evaluation and extrapolation of the numerical solution.

The resulting variable performance, however, can only be addressed by assigning static (vertex) weights for partitioning. To determine these weights, we count for each vertex the dynamic rupture faces (n_{DR}) and the number of gravitational boundary faces (n_G) of the corresponding element. The total weight is then given by

$$2^{c_{\max}-c_v} (w_{\text{base}} + w_{DR} n_{DR} + w_G n_G). \quad (28)$$

The update rate of the element is reflected by $2^{c_{\max}-c_v}$, where c_v is the cluster number of vertex v and $c_{\max}$ is the number of the largest time cluster [6]. As vertex weights are integers in ParMETIS, we assign a base weight of $w_{\text{base}} = 100$ and determine w_{DR} and w_G relative to w_{base} .

We varied the weights w_{DR} and w_G in the range between 50 and 500 and ran a production simulation for 0.2 s. For w_G , we found that the performance generally increases with weight, indicating that a weight in the range of 300–500 is appropriate. For w_{DR} , a clear trend is not apparent. Moreover, we face the following conceptual problems for w_{DR} : First, the short simulated time for such test runs might not be indicative for a whole production run as the load is dynamic and expected to increase after earthquake nucleation. Second, when a dynamic rupture face is adjacent to two partitions, the slip-rate on a fault face is computed twice. That is, the load depends not only on the weights but on the partition itself. Improving the dynamic rupture load balancing hence likely requires dynamic load balancing techniques. We therefore set weights heuristically to $w_{DR} = 200$ and $w_G = 300$ in this work, which leads to an adequate parallel efficiency as shown in Sec. 6.

Petascale supercomputers can suffer from strong performance fluctuations and few nodes may be substantially slower than others. Thus, treating supercomputers as homogeneous can lead to a slowdown of the entire simulation [38, 66]. Hence, we assume that their performance is inherently heterogeneous. For each simulation run, we perform a small benchmark of our computational kernels prior to partitioning. The inverse of the measured computational time for the small benchmark is taken as node weight. Normalized weights are then fed into ParMETIS (tpwgts). The performance impact of node weights is discussed in Secs. 6.2 and 6.3.

6 RESULTS

Our fully-coupled elastic-acoustic model establishes a novel method for solving an entirely new class of earthquake-tsunami problems well beyond established earthquake and shallow water tsunami solvers. In particular, our model is not just an alternative way to compute the initial ocean displacement or even the tsunami wavefield. While the tsunami solution from 3D fully coupled modeling is very similar to that of the more traditional one-way linked modeling under certain simplifying conditions (cf. Sec. 6.1), the fully coupled model always computes the full wavefield that includes seismic and

acoustic waves in addition to tsunami waves and allows non-linear coupling to earthquake physics. In important parts of the parameter space, a fully coupled solution is essential since (i) the approximations that form the basis of the shallow water equations are invalid, (ii) the shallow water equations simply cannot capture seismic and acoustic waves that can be the dominant features in seismograms and pressure data recorded by offshore instruments over the source region and (iii) the complexity of the tsunami, acoustic and seismic ocean response phases and the influence of geometry, geologic structure and earthquake dynamics requires coupled solutions to fully understand tsunami generation processes [4, 14].

We therefore first validate our 3D fully-coupled Earth and ocean model against one-way linked earthquake-shallow-water modeling in a setup where the latter approximation is valid. We then present results of an extreme-scale application of the 2018 Sulawesi earthquake and tsunami, which showcases the added complexity of results obtained via the fully-coupled model. We present strong scaling and performance analysis of the end-to-end application scenario. We use maximum polynomial degree $N = 5$ and $C(N) = 0.35(2N + 1)^{-1}$ for all simulations. The performance and scaling tests and the production runs were executed on three different petascale systems:

Shaheen-II 6174 nodes with dual socket Intel Xeon E5-2698 v3. Nodes are connected with an Aries interconnect and Dragonfly topology.

SuperMUC-NG 6336 nodes with dual socket Intel Skylake Xeon Platinum 8174 (24 cores each). The nodes are organized in 8 islands. A fat tree OmniPath network is used, between each island the connection is pruned (pruning factor 1:4)

Mahti 1404 nodes with dual socket AMD Rome 7H12 (64 cores each). Nodes are connected by a Mellanox HDR InfiniBand interconnect using a Dragonfly+ topology.

6.1 Earthquake-Tsunami benchmark scenario

Verification of the correctness of the acoustic coupling and interface condition implementation has been demonstrated in preliminary convergence analyses with respect to analytic solutions [2, 27]. In the following, we compare the fully-coupled model with linked modeling for the “Scenario A” dynamic megathrust earthquake benchmark of Madden et al. [35]. This setup includes an idealized 3D megathrust dynamic rupture earthquake source on a planar fault, tsunami generation and propagation in a basin with flat bathymetry and inundation on a linearly sloping beach. The benchmark is run with both our fully coupled model and a one-way linked 3D-2D modeling approach using a shallow water solver for tsunami propagation (as in [35]). While we anticipate some level of agreement between the two models, differences in ocean response are expected as a consequence of the approximations employed in the linked model. Specifically, we expect the excitation of ocean acoustic modes in the fully coupled simulation within the deep water (2 km depth) by the high-frequency seismic waves resolved in the earthquake model [1]. Further differences are expected since the fully coupled model does not include the sloping beach.

In Madden et al. [35]’s Scenario A, the fault is 200 km wide and extends from the surface to 35 km depth at a 16° dip. We discretize the fully-coupled scenario with tetrahedral edge length of 400 m

at the fault and gradually coarsen the mesh away from the source. We add a 2 km water layer atop the Earth model. Since we use conforming meshes, the resolution of the water layer is higher near the earthquake source (locally refined to conform with a 66 m edge length at the megathrust fault). Adding the water layer increases the mesh size from roughly 16 million (for the earthquake model) to 29.5 million elements for the fully coupled model.

On-fault earthquake dynamics are governed by a linear-slip weakening friction law [3], which is computationally less demanding than rate-and-state friction employed in the complex Palu, Sulawesi application (Sec. 6.2). Higher fault strength near the seafloor smoothly stops the rupture as it approaches the Earth surface. The earthquake magnitude is $M_w = 8.5$ and the rupture propagates at 3.5 km/s on average. The solid has homogeneous elastic properties representative of oceanic crust in a subduction zone [56]: p-wave speed $c_p = 7639.9$ m/s, shear wave speed $c_s = 4229.4$ m/s and density $\rho = 3775$ kg/m³. The ocean has an acoustic wave speed of $c_p = 1500$ m/s and density $\rho = 1000$ kg/m³.

We compare the fully-coupled simulation to the one-way linked reference solution in Fig. 3. In the one-way linked approach, the seafloor displacement recorded on the unstructured mesh of the earthquake model is bilinearly interpolated to an intermediate uniform Cartesian mesh with a resolution of 1000 m, which is subsequently used as a time-dependent source in the hydrostatic non-linear shallow water tsunami model. The time-dependent forcing term in the shallow water mass balance is here taken as the unfiltered seafloor displacement field. Then the second-order Runge-Kutta DG non-linear shallow water solver sam(oa)²-flash is used to perform the tsunami simulation using dynamically adaptive mesh refinement (see Madden et al. [35] for details).

The sea surface height from our fully coupled solution matches the one-way linked approach at the low frequencies characterizing the tsunami response (Fig. 3, b). As expected, differences appear at $x > 240$ km, at the right-hand side of the domain, where the reference simulation includes a beach which is not contained in our fully coupled model. The short wavelength and high frequency oscillations (periods < 5.3 s) trailing the leading seismic wavefronts (Fig. 3, a) are reverberating acoustic wave modes in the ocean, which are captured only in our fully coupled model.

6.2 Dynamics of the 2018 Palu, Sulawesi tsunami generation

Next, we present the first fully-coupled 3D simulation of a real event: the 2018 magnitude 7.5 strike-slip supershear Sulawesi earthquake and associated local tsunami in Palu Bay (Fig. 1). The dynamic rupture earthquake source model features sustained supershear rupture propagation across a multi-segment fault model (Fig. 1, c) and matches key observed earthquake characteristics, including the moment magnitude, rupture duration, fault plane solution, teleseismic waveforms and inferred horizontal ground displacements from observations of optical and radar satellites [60]. It incorporates a fast velocity weakening rate-and-state friction law [16, 41] which is motivated by laboratory rock physics experiments. The time-dependent 3D seafloor displacements produce a mean vertical uplift of 1.5 m underneath Palu Bay. This sources a tsunami

Figure 3: Sea surface height in a 3D megathrust earthquake-tsunami benchmark (“Scenario A” of [35]). (a) Snapshot of the sea surface height (vertical displacement η) of the fully coupled model at $t=120$ s. (b) Comparison of the fully coupled model to a one-way linked model along a cross-section at $y=0$ km (indicated in (a) as black line) through the epicenter at $t=120$ s. Blue: fully coupled sea surface height, including ocean response. Red: one-way linked model with a non-linear shallow water solver using the dynamic seafloor displacement as forcing term and including a linearly sloping beach at $x=240$ km. Both models use the same dynamic rupture earthquake source performed with the same polynomial order 5.

with wave amplitudes and periods that match those measured and inundation that reproduces observations from field surveys.

The fully coupled model includes a water layer of variable depth atop BATNAS bathymetry and topography data of horizontal resolution of 6 arc seconds (≈ 190 m, [9]). Due to the bathtub-like, steep geometry of Palu Bay and the shallow water depth (on average 600 m), the locally adaptive unstructured tetrahedral mesh is crucial for rendering this model feasible. Adaptive local refinement as fine as 200 m element edge length ensures accurate resolution of rupture dynamics across the complex fault system [65] – we observe high frequency seismic wave radiation of up to 10 Hz close to the faults within Palu Bay. We set a maximum global element size of 5000 m and refine the resolution in the water layer and in our region of interest: a cuboid centered at $x = 5$ km, $y = 0$ km, $z = -8$ km with extent $x = 70$ km, $y = 180$ km, $z = 34$ km.

We present “production runs” using the following two meshes. The respective numbers of degrees of freedom using a polynomial order of 5 are given in parentheses.

- M** Our medium-sized mesh with ≈ 89 million elements (≈ 46 billion degrees of freedom) has a water layer resolution of 100 m and resolves seismic wave propagation with elements of a maximum size of 1000 m.
- L** Our large mesh has ≈ 518 million elements (≈ 261 billion degrees of freedom). Its resolution in both the water layer with 50 m and in the seismic wave refinement zone with 500 m is twice as fine as in the **M** mesh.

In the **L** mesh, 453.7 million grid cells are used to discretize the ocean (acoustic layer), which increases the total mesh size by a factor of 8. This makes the simulation of the **L** mesh almost twice as large in terms of degrees of freedom as the so-far largest reported SeisSol simulation [62].

We run the **M** mesh for 100 s which is long enough for the localized tsunami to interact with the coastline of the bay. The higher resolution **L** model is run for 30 s during which we capture details

of the entire earthquake rupture process and the generation and interaction of both acoustic waves and tsunami. Within the water layer, we prescribe a vertical and horizontal spatial discretization of maximum 50 m element edge length in the **L** mesh, resolving a maximum frequency of at least 15 Hz of the acoustic wave field (ensuring that 2 elements of polynomial order 5 space-time accuracy sample one wavelength propagating at an approximate speed of 1483 m/s [29]). With this setup, we measure wave excitation of up to 30 Hz in the Fourier spectra of the recorded acoustic velocity time series which we attribute to the highly variable water depth.

The high-resolution simulation using the **L** mesh is illustrated in Fig. 1. The velocity field snapshot in Fig. 1 (a) reveals a very high degree of complexity at 15 s simulation time. A sharp shear Mach front is stacked up due to the fast moving dynamic rupture source. Seismic waves within Earth, acoustic waves within the ocean (expected at periods shorter than ≈ 1.6 s) and wave conversions superimpose. At this point in time during the simulation, the unilaterally propagating southward rupture arrives at the end of Palu Bay (Fig. 1, b). The vertical displacement field of our high-resolution fully-coupled model in Fig. 1 (c) shows that seismic waves (including the sharply imprinting supershear Mach front) result in transient motions of the sea surface and affect the ocean response but do not appear to contribute to tsunami generation.

Time to solution, specifically of the large production run, benefits greatly from local time-stepping: The chosen clustering (Fig. 4) reduces the total number of required element updates by a factor of ≈ 30 . More than 86% of all elements fall within the cluster with timestep $32(\Delta t)_{\min}$. Running with global time-stepping is therefore not feasible and we thus do not report respective performance or simulation results. We ran the **L** mesh on 3072 nodes of *SuperMUC-NG* and achieved 3.14 PFLOPS on average, i.e., a bit more than 1 TFLOPS per node. The run took 5 hours and 30 minutes to completion (time reported by Slurm), hence running it for

Figure 4: Distribution of elements to the local time-stepping clusters for mesh L. The x-axis gives the timestep size of the clusters in terms of the minimum timestep size $(\Delta t)_{\min}$. Note that the y-axis is scaled logarithmically.

the same time as the **M** mesh would take 18 hours and 15 minutes, which is computationally feasible. We ran the same setup on 6144 nodes of *Shaheen-II*, reaching an estimated performance³ of 2.3 PFLOPS. These simulations (and performance numbers) include receiver output (every 0.01 s) and free surface output (including the elastic-acoustic interface, every 0.1 s). The merits of assigning node weights become clear for these large-scale runs. For the **L** run on *SuperMUC-NG*, we see node weights of $\approx 4.54 \pm 0.087$ (mean $\pm$ standard deviation) with a minimum weight of just 2.74. That means the slowest node had a performance of only 60.4% of the average node. On *Shaheen-II*, performance variability was less severe, with node weights of 3.34 ± 0.023 with a minimum of 3.19.

Figure 5 shows a series of snapshots in time of the vertical displacement fields evolving during the **M** mesh earthquake-tsunami simulation in the upper row. The lower row shows the same snapshots generated with a one-way linked method that uses a non-linear shallow water solver. The ocean experiences subsidence (southeast of the fault) and uplift (northwest of the fault) from the static earthquake displacements (Fig. 1, d) as modulated by topography and bathymetry. As the rupture propagates much faster than the tsunami wave speed, the tsunami is sourced with nearly planar wavefronts parallel to the fault trace (Fig. 5, 40 s). Deviations are most pronounced where the fault enters and exits the bay and are primarily caused by the slower wave speed in the shallow water along the edge of the bay. Such clear patterns are not visible anymore at 100 s. The local extrema in vertical displacement along the coast illustrates the complex wave reflections and refractions within the bay caused by shallow bathymetry as well as funnel and focusing effects. We observe continuous reflections of waves from the coast, for example, on the lower-left corner.

While most wavefield features are quite similar, as are predicted wave heights, there are also clear differences. Both simulations are performed with the same bathymetry and coastlines and similar element sizes, suggesting that differences may be physical and may arise from the inclusion of additional processes, such as compressibility, in the fully coupled model. The one-way linking approach produces a tsunami with much sharper wavefronts, which we speculate may be caused by the hydrostatic approximation used in the shallow water solver and aided by the very complex bathymetry

³The simulation timed out due to the allocated job duration after reaching ≈ 16 s. We estimate performance by interpolation from the FLOP count of the *SuperMUC-NG* run and from the measured execution time.

Figure 5: Snapshots of vertical displacement of the ocean surface within Palu Bay in the fully-coupled Sulawesi earthquake-tsunami simulation using mesh **M (upper row) compared to a one-way linked 2D shallow water tsunami simulation (lower row). While both models capture the overall dynamics of the dynamic earthquake-ocean response we observe differences in the sharpness of wave fronts (see text) in addition to the excitation of acoustic waves as was observed in the deep ocean benchmark (fast seismic and acoustic waves are not shown here but visible in Figs. 1 and 3). Fully-coupled 3D simulations open the possibility to explore the effects of variable depth ocean non-hydrostatic response, compressibility and acoustic wave generation across complex bathymetry.**

in Palu Bay. The wavefield is notably smoother in the fully coupled model, perhaps due to non-hydrostatic effects that filter short-wavelength features in the transfer function between seafloor and sea surface motions [22] or due to conversion between surface gravity, acoustic, and seismic wave modes. Also, nonlinear effects may become important during wave interaction with the coastline. Fully-coupled 3D simulations allow us to explore these hypotheses in great detail and, together with extended numerical analysis, to better understand when one-way linked approaches are justified and when fully coupled modeling is required.

6.3 Strong scaling

The SeisSol software has been shown to scale well for dynamic rupture earthquake simulations on various petascale machines [17, 62]. However, the extensions to the presented coupled elastic-acoustic model impose several new challenges for load balancing and on the cluster-oriented local time-stepping algorithm (due to the smaller

Figure 6: Strong scaling for mesh M on Mahti and SuperMUC-NG (equivalent to parallel efficiency).

timesteps in the acoustic layer and the resulting different distribution of elements to timestep clusters). The main challenge for coupled elastic-acoustic computation is resolving the high frequency range involved and the low wave speeds, e.g., in near-surface sediments, requiring locally high spatial discretization [48]. We therefore present scalability results of the Palu, Sulawesi earthquake-tsunami application scenario on *Mahti* and *SuperMUC-NG*. While SeisSol has previously been optimized towards the use of a single MPI rank per compute node, we observe (see Sec. 5.1) stronger influence of NUMA domains on node-level performance on recent architectures. We therefore experiment with different numbers of MPI ranks in the scalability tests and particularly include *Mahti* as an AMD-based machine, whose CPUs feature 4 NUMA domains. As described in Sec. 5, we expect that a larger number of MPI ranks per node should have a beneficial effect, due to NUMA effects. To evaluate the optimal number of nodes per rank, we ran our scaling study (Fig. 6a) for 1 MPI process per node, per socket (2 ranks per node) and per NUMA domain (8 ranks per node). The situation on *SuperMUC-NG* is similar but less severe, with two NUMA domains per node, stemming from the dual socket configuration. Thus, we ran the scaling tests (Fig. 6b) to compare 1 process per node with 1 process per socket.

We ran all scaling studies until 0.03 s which includes several LTS cycles. In contrast to our production runs, we deactivated output for our scaling study. We report the maximal performance over multiple runs to avoid performance fluctuations. Note that we sacrifice one physical thread of each MPI process for efficient

communication (MPI progression). Despite this, we achieved best results using one process per NUMA domain.

Using this setup, we achieved 1359 GFLOPS per node with 50 nodes on NG and 981 GFLOPS on 1600 nodes. This translates to a parallel efficiency of $\approx 72\%$. On *Mahti*, we achieved 2322 GFLOPS on 50 nodes and 1689 GFLOPS on 700 nodes, resulting in a parallel efficiency of $\approx 73\%$. We achieved up to ≈ 1.57 PFLOPS on *SuperMUC-NG* and up to ≈ 1.18 PFLOPS on *Mahti* in total. We also ran a scaling study on the L mesh on *SuperMUC-NG*. Using 2 nodes per rank, we achieved a performance of 1298 GFLOPS on 768 nodes and of 996 GFLOPS on 3072 nodes. This corresponds to a parallel efficiency of $\approx 76.8\%$. Furthermore, the L mesh (≈ 518 million elements) on 3072 nodes of *SuperMUC-NG* corresponds to a workload of $\approx 1.68 \cdot 10^5$ elements per node, which is equivalent to using roughly 500 nodes for the M mesh. As the latter falls within the region of good parallel efficiency in Fig. 6, we expect that the L mesh scenario will also scale well on the full *SuperMUC-NG* machine.

Finally, we used the strong scaling setup to evaluate the effect of our node weights. When we deactivated this feature, we got a maximum performance (on 700 nodes of *Mahti*) of 1412 GFLOPS, i.e. 84% of the performance with node weights.

7 CONCLUSION AND OUTLOOK

We developed a 3D fully-coupled model for earthquake-tsunami interaction that includes tsunami propagation, and described our novel discretization and implementation in the software SeisSol. Our numerical scheme is high order in both space and time and allows the simulation of multi-physics scenarios that include frictional slip on faults and wave propagation in both elastic and acoustic layers with gravity. We showed results for an earthquake-tsunami interaction benchmark that demonstrate that our method compares well to solutions obtained with “one-way” linking with a shallow water model under certain conditions.

We presented a novel, fully-coupled setup for the 2018 Palu, Sulawesi earthquake and tsunami that captures the dynamics of the entire tsunami-genesis in a single simulation. There are notable differences in the fully coupled and one-way linked models for the Palu event, raising questions regarding the validity of approximations used in standard earthquake-tsunami coupling workflows. Only with a fully coupled solution can these approximations be explored in the complex geometries and material structures representative of real events.

Our production runs achieved petascale performance on both *SuperMUC-NG* and *Shaheen-II*. Our solver scales up to 1600 nodes on *SuperMUC-NG* and up to 700 nodes on *Mahti* for our M production scenario, achieving excellent single node performance (up to 63% for the predictor step on the AMD Rome architecture). We described our heuristic for multi-physics load-balancing and we explained, how we achieve a fairer load-balancing by treating supercomputers as inherently heterogeneous. We discussed limits of the static load balancing, which becomes increasingly difficult when combining local time-stepping and multi-physics and which cannot deal with dynamic load, as posed by the dynamic earthquake rupture friction law. In future work, we want to explore a task-based execution of clustered local time-stepping which enables the use of remote task off-loading for dynamic load balancing [24].

Our work opens up new avenues for research into earthquake-tsunami interaction. Not only is it now possible to capture the entire dynamics of this process in a unified 3D model; it is also efficient. This permits fully-coupled simulations even for large scenarios, instead of relying on approximate 3D-2D linked methods. Our new simulation capabilities open the possibility to fundamentally improve our understanding of earthquake-tsunami interaction in its full complexity, which includes identifying when 3D-2D one-way linked models are justified and when fully coupled modeling is required. Also, our implementation will be instrumental to further elastic-acoustic coupling applications, such as holistic modeling of earthquake-generated acoustic waves [25] in complex geological settings or seismic waves originating from lava lake sloshing [31].

ACKNOWLEDGMENTS

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No. 823844 (ChEESE – Centre of Excellence in Solid Earth). Compute resources were provided by the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) on SuperMUC-NG at the Leibniz Supercomputing Centre (www.lrz.de, project pn68fi), by the CSC – IT Center for Science, Finland (project 2003841, on Mahti) and by the Supercomputing Laboratory at King Abdullah University of Science & Technology (KAUST, project k1488, on Shaheen-II) in Thuwal, Saudi Arabia. We thank all colleagues at LRZ, CSC and KAUST for their excellent support. C.U., T.U. and A.-A.G. acknowledge additional support by the European Union's Horizon 2020 Research and Innovation Programme under ERC StG TEAR, no. 852992, the German Research Foundation (DFG) (grants no. GA 2465/2-1, GA 2465/3-1) and by KAUST-CRG (grant no. ORS-2017-CRG6 3389.02). L.S.A. was supported by National Science Foundation grant DGE-1656518.

REFERENCES

- [1] Lauren S. Abrahams, Eric M. Dunham, Lukas Krenz, Tatsuhiko Saito, and Alice-Agnes Gabriel. 2021. Comparison of techniques for coupled earthquake and tsunami modeling. *Earth and Space Sci. Open Arch.* (Feb. 2021), 49. <https://doi.org/10.1002/essoar.10506178.3> AGU 2020 Fall Meeting.
- [2] Lauren S. Abrahams, Lukas Krenz, Eric M. Dunham, and Alice-Agnes Gabriel. 2019. Verification of a 3D fully-coupled earthquake and tsunami model. In *AGU Fall Meet. Abstr.*, Vol. 2019. NH43F–1000. <https://agu.confex.com/agu/fm19/meetingapp.cgi/Paper/547532>
- [3] D. J. Andrews. 1976. Rupture velocity of plane strain shear cracks. *J. Geophys. Res.* 81, 32 (nov 1976), 5679–5687. <https://doi.org/10.1029/JB081i032p05679>
- [4] Sara Aniko Wirp, Alice-Agnes Gabriel, Maximilian Schmeller, Elizabeth H. Maden, Iris van Zelst, Lukas Krenz, Ylona van Dinther, and Leonhard Rannabauer. 2021. 3D Linked Subduction, Dynamic Rupture, Tsunami, and Inundation Modeling: Dynamic Effects of Supershear and Tsunami Earthquakes, Hypocenter Location, and Shallow Fault Slip. *Front. in Earth Sci.* 9 (2021), 177.
- [5] Marsha J. Berger, David L. George, Randall J. LeVeque, and Kyle T. Mandli. 2011. The GeoClaw software for depth-averaged flows with adaptive refinement. *Adv. Water. Res.* 34, 9 (2011), 1195–1206.
- [6] Alexander Breuer, Alexander Heinecke, and Michael Bader. 2016. Petascale Local Time Stepping for the ADER-DG Finite Element Method. In *2016 IEEE Int. Parallel Distrib. Process. Symp. (IPDPS)*. 854–863. <https://doi.org/10.1109/IPDPS.2016.109>
- [7] Bingwei Chen, Haohuan Fu, Yanwen Wei, Conghui He, Wenqiang Zhang, Yuxuan Li, Wubin Wan, Wei Zhang, Lin Gan, Wei Zhang, Zhenguo Zhang, Guangwen Yang, and Xiaofei Chen. 2018. Simulating the Wenchuan Earthquake with Accurate Surface Topography on Sunway TaihuLight. In *Proc. of the Int. Conf. High Perform. Comput., Networking, Storage, and Analysis* (Dallas, Texas) (SC '18). Article 40. <https://doi.org/10.1109/SC.2018.00043>
- [8] Richard Courant, Kurt Friedrichs, and Hans Lewy. 1928. Über die partiellen Differenzengleichungen der mathematischen Physik. *Math. Ann.* 100, 1 (1928), 32–74.
- [9] DEMNAS 2018. DEMNAS – Seamless Digital Elevation Model (DEM) dan Batimetri Nasional. Badan Informasi Geospasial. <http://tides.big.go.id/DEMNAS>
- [10] Michael Dumbser, Dinshaw S. Balsara, Eleuterio F. Toro, and Claus-Dieter Munz. 2008. A unified framework for the construction of one-step finite volume and discontinuous Galerkin schemes on unstructured meshes. *J. Comput. Phys.* 227, 18 (Sept. 2008), 8209–8253. <https://doi.org/10.1016/j.jcp.2008.05.025>
- [11] Michael Dumbser and Martin Käser. 2006. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – II. The three-dimensional isotropic case. *Geophys. J. Int.* 167, 1 (Oct. 2006), 319–336. <https://doi.org/10.1111/j.1365-246X.2006.03120.x>
- [12] Michael Dumbser, Martin Käser, and Eleuterio F. Toro. 2007. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – V. Local time stepping and p-adaptivity. *Geophys. J. Int.* 171, 2 (Nov. 2007), 695–717. <https://doi.org/10.1111/j.1365-246X.2007.03427.x>
- [13] M. Dumbser, T. Schwartzkopff, and C.-D. Munz. 2006. Arbitrary High Order Finite Volume Schemes for Linear Wave Propagation. In *Comput. Sci. and High Perform. Comput. II (Notes on Numer. Fluid Mech. and Multidisc. Design)*. Springer, Berlin, Heidelberg, 129–144. https://doi.org/10.1007/3-540-31768-6_11
- [14] Ahmed Elbanna, Mohamed Abdelmeguid, Xiao Ma, Faisal Amlani, Harsha S. Bhat, Costas Synolakis, and Ares J. Rosakis. 2021. Anatomy of strike-slip fault tsunami genesis. *Proc. of the Natl. Acad. Sci.* 118, 19 (2021). <https://doi.org/10.1073/pnas.2025632118>
- [15] Haohuan Fu, Conghui He, Bingwei Chen, Zekun Yin, Zhenguo Zhang, Wenqiang Zhang, Tingjian Zhang, Wei Xue, Weiguo Liu, Wanwang Yin, Guangwen Yang, and Xiaofei Chen. 2017. 18.9-pflops nonlinear earthquake simulation on Sunway TaihuLight: enabling depiction of 18-Hz and 8-meter scenarios. In *Proc. Int. Conf. High Perform. Comput., Networking, Storage and Analysis* (Denver, Colorado) (SC '17). Association for Computing Machinery, New York, NY, USA, Article 2, 12 pages. <https://doi.org/10.1145/3126908.3126910>
- [16] Ruth A. Harris, Michael Barall, Brad Aagaard, Shuo Ma, Daniel Roten, Kim Olsen, Benchun Duan, Dunyu Liu, Bin Luo, Kangchen Bai, Jean-Paul Ampuero, Yoshihiro Kaneko, Alice-Agnes Gabriel, Kenneth Duru, Thomas Ulrich, Stephanie Wollherr, Zheqiang Shi, Eric Dunham, Sam Bydron, Zhenguo Zhang, Xiaofei Chen, Surendra Nadh Somala, Christian Pelties, Josué Tago, Victor Manuel Cruz-Atienza, Jeremy Kozdon, Eric Daub, Khurram Aslam, Yuko Kase, Kyle Withers, and Luis Dalguer. 2018. A Suite of Exercises for Verifying Dynamic Earthquake Rupture Codes. *Seismol. Res. Lett.* 89, 3 (May 2018), 1146–1162. <https://doi.org/10.1785/0220170222>
- [17] Alexander Heinecke, Alexander Breuer, Sebastian Rettenberger, Michael Bader, Alice-Agnes Gabriel, Christian Pelties, Arndt Bode, William Barth, Xiang-Ke Liao, Karthikeyan Vaidyanathan, Mikhail Smelyanskiy, and Pradeep Dubey. 2014. Petascale High Order Dynamic Rupture Earthquake Simulations on Heterogeneous Supercomputers. In *SC '14: Proc. Int. Conf. High Perform. Comp., Networking, Storage and Analysis*. 3–14. <https://doi.org/10.1109/SC.2014.6> ISSN: 2167-4337.
- [18] Alexander Heinecke, Greg Henry, Maxwell Hutchinson, and Hans Pabst. 2016. LIBXSM: accelerating small matrix multiplications by runtime code generation. In *SC '16: Proc. Int. Conf. High Perform. Comput., Networking, Storage and Analysis*. IEEE, 981–991.
- [19] Torsten Hoefler and Andrew Lumsdaine. 2008. Message progression in parallel computing – to thread or not to thread?. In *2008 IEEE Int. Conf. Cluster Comput.* 213–222. <https://doi.org/10.1109/CLUSTER.2008.4663774>
- [20] Tsuyoshi Ichimura, Kohei Fujita, Seizo Tanaka, Muneo Hori, Madgededara Lalith, Yoshihisa Shizawa, and Hiroshi Kobayashi. 2014. Physics-Based Urban Earthquake Simulation Enhanced by 10.7 BlnDOF × 30 K Time-Step Unstructured FE Non-Linear Seismic Wave Simulation. In *Proc. Int. Conf. High Perform. Comput., Networking, Storage and Analysis* (SC '14). 15–26. <https://doi.org/10.1109/SC.2014.7>
- [21] Tsuyoshi Ichimura, Kohei Fujita, Takuma Yamaguchi, Akira Naruse, Jack C. Wells, Thomas C. Schulthess, Tjerk P. Straatsma, Christopher J. Zimmer, Maxime Martinasso, Kengo Nakajima, Muneo Hori, and Lalith Madgededara. 2018. A fast scalable implicit solver for nonlinear time-evolution earthquake city problem on low-ordered unstructured finite elements with artificial intelligence and transprecision computing. In *SC18: Int. Conf. High Perform. Comput., Networking, Storage and Analysis*. IEEE, 627–637.
- [22] Kinjiro Kajiuura. 1963. The leading wave of a tsunami. *Bull. Earthquake Res. Inst., University of Tokyo* 41, 3 (1963), 535–571.
- [23] Martin Käser and Michael Dumbser. 2006. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – I. The two-dimensional isotropic case with external source terms. *Geophysical J. Int.* 166, 2 (Aug 2006), 855–877. <https://doi.org/10.1111/j.1365-246X.2006.03051.x>
- [24] Jannis Klinkenberg, Philipp Samfass, Michael Bader, Christian Terboven, and Matthias S. Müller. 2020. Reactive task migration for hybrid MPI + OpenMP applications. In *Parallel Proc. and Appl. Math.* Springer International Publishing, 59–71. https://doi.org/10.1007/978-3-030-43222-5_6
- [25] Dimitri Komatitsch. 2011. Fluid–solid coupling on a cluster of GPU graphics cards for seismic wave propagation. *Comptes Rendus Mécanique* 339, 2–3 (2011), 125–135.

- [26] Jeremy E. Kozdon and Eric M. Dunham. 2014. Constraining shallow slip and tsunami excitation in megathrust ruptures using seismic and ocean acoustic waves recorded on ocean-bottom sensor networks. *Earth Planet. Sci. Lett.* 396 (2014), 56–65. <https://doi.org/10.1016/j.epsl.2014.04.001>
- [27] Lukas Krenz, Carsten Uphoff, Lauren S Abrahams, Alice-Agnes Gabriel, Eric M Dunham, and Michael Bader. 2019. Elastic-acoustic coupling for 3D tsunamigenic earthquake simulations with ADER-DG on unstructured tetrahedral meshes. In *AGU Fall Meet. Abstr.*, Vol. 2019. T52C–09.
- [28] Martin Käser, Michael Dumbser, Josep De La Puente, and Heiner Igel. 2007. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – III. Viscoelastic attenuation. *Geophys. J. Int.* 168, 1 (01 2007), 224–242. <https://doi.org/10.1111/j.1365-246X.2006.03193.x>
- [29] Martin Käser, Verena Hermann, and Josep de la Puente. 2008. Quantitative accuracy analysis of the discontinuous Galerkin method for seismic wave propagation. *Geophys. J. Int.* 173, 3 (June 2008), 990–999. <https://doi.org/10.1111/j.1365-246X.2008.03781.x>
- [30] Randall J. LeVeque. 2002. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press. <https://doi.org/10.1017/cbo9780511791253>
- [31] Chao Liang and Eric M Dunham. 2020. Lava lake sloshing modes during the 2018 Kilauea Volcano eruption probe magma reservoir storativity. *Earth Planet. Sci. Lett.* 535 (2020), 116110.
- [32] Gabriel C. Lotto and Eric M. Dunham. 2015. High-order finite difference modeling of tsunami generation in a compressible ocean from offshore earthquakes. *Comput. Geosci.* 19, 2 (2015), 327–340. <https://doi.org/10.1007/s10596-015-9472-0>
- [33] Gabriel C. Lotto, Tamara N. Jeppson, and Eric M. Dunham. 2019. Fully coupled simulations of megathrust earthquakes and tsunamis in the Japan Trench, Nankai Trough, and Cascadia Subduction Zone. *Pure Appl. Geophys.* 176, 9 (2019), 4009–4041.
- [34] Gabriel C. Lotto, Gabriel Nava, and Eric M. Dunham. 2017. Should tsunami simulations include a nonzero initial horizontal velocity? *Earth Planets Space* 69, 1 (Dec. 2017), 117. <https://doi.org/10.1186/s40623-017-0701-8>
- [35] E. H. Madden, M. Bader, J. Behrens, Y. van Dinther, A.-A. Gabriel, L. Rannabauer, T. Ulrich, C. Uphoff, S. Vater, and I. van Zelst. 2021. Linked 3-D modelling of megathrust earthquake-tsunami events: from subduction to tsunami run up. *Geophys. J. Int.* 224, 1 (2021), 487–516. <https://doi.org/10.1093/gji/ggaa484>
- [36] Takuto Maeda, Takashi Furumura, Shinako Noguchi, Shunsuke Takemura, Shin'ichi Sakai, Masanao Shinohara, Kazuhisa Iwai, and Shiann Jong Lee. 2013. Seismic- and Tsunami-wave propagation of the 2011 Off the Pacific Coast of Tohoku earthquake as inferred from the Tsunami-coupled finite-difference simulation. *Bull. Seismol. Soc. Am.* 103, 2 B (may 2013), 1456–1472. <https://doi.org/10.1785/BSSA0750041135>
- [37] P. Martin Mai. 2019. Supershear tsunami disaster. *Nat. Geosci.* 12, 3 (March 2019), 150–151. <https://doi.org/10.1038/s41561-019-0308-8>
- [38] John D. McCalpin. 2018. HPL and DGEMM performance variability on the Xeon Platinum 8160 Processor. In *Proc. of the Int. Conf. High Perform. Comput., Networking, Storage, and Analysis (SC '18)*. <https://doi.org/10.1109/SC.2018.00021>
- [39] Kevin R. Milner, Bruce E. Shaw, Christine A. Goulet, Keith B. Richards-Dinger, Scott Callaghan, Thomas H. Jordan, James H. Dieterich, and Edward H. Field. 2021. Toward Physics-Based Nonergodic PSHA: A Prototype Fully Deterministic Seismic Hazard Model for Southern California. *Bull. Seismol. Soc. Am.* 111, 2 (01 2021), 898–915.
- [40] Yoshimitsu Okada. 1985. Surface deformation due to shear and tensile faults in a half-space. *Bull. Seismol. Soc. Am.* 75, 4 (Aug. 1985), 1135–1154. <https://doi.org/10.1785/BSSA0750041135>
- [41] C. Pelties, A.-A. Gabriel, and J.-P. Ampuero. 2014. Verification of an ADER-DG method for complex dynamic rupture problems. *Geosci. Model Dev.* 7, 3 (2014), 847–866. <https://doi.org/10.5194/gmd-7-847-2014>
- [42] Sebastian Rettenberger and Michael Bader. 2015. Optimizing I/O for petascale seismic simulations on unstructured meshes. In *2015 IEEE Int. Conf. Cluster Comp. IEEE*, 314–317. <https://doi.org/10.1109/CLUSTER.2015.51>
- [43] Max Rietmann, Marcus Grote, Daniel Peter, and Olaf Schenk. 2017. Newmark local time stepping on high-performance computing architectures. *J. Comput. Phys.* 334 (2017), 308–326.
- [44] Max Rietmann, Daniel Peter, Olaf Schenk, Bora Uçar, and Marcus Grote. 2015. Load-balanced local time stepping for large-scale wave propagation. In *2015 IEEE Int. Parallel and Distrib. Process. Symp. IEEE*, 925–935.
- [45] Arthur J. Rodgers, Arben Pitarka, Ramesh Pankajakshan, Björn Sjögreen, and N. Anders Petersson. 2020. Regional-Scale 3D Ground-Motion Simulations of M_w 7 Earthquakes on the Hayward Fault, Northern California Resolving Frequencies 0–10 Hz and Including Site-Response Corrections. *Bull. Seismol. Soc. Am.* 110, 6 (2020), 2862–2881.
- [46] Arthur J. Rodgers, Arben Pitarka, N. Anders Petersson, Björn Sjögreen, and David B. McCallen. 2018. Broadband (0–4 Hz) Ground Motions for a Magnitude 7.0 Hayward Fault Earthquake With Three-Dimensional Structure and Topography. *Geophys. Res. Lett.* 45, 2 (2018), 739–747. <https://doi.org/10.1002/2017GL076505>
- [47] Julien Diaz Ronan Madec, Dimitri Komatitsch. 2009. Energy-conserving local time stepping based on high-order finite elements for seismic wave propagation across a fluid-solid interface. *Comp. Model. in Eng. & Sci.* 49, 2 (2009), 163–190. <https://doi.org/10.3970/cmescs.2009.049.163>
- [48] Eric Rosenkrantz, Alexis Bottero, Dimitri Komatitsch, and Vadim Monteiller. 2019. A flexible numerical approach for non-destructive ultrasonic testing based on a time-domain spectral-element method: Ultrasonic modeling of Lamb waves in immersed defective structures and of bulk waves in damaged anisotropic materials. *NDT & E International* 101 (2019), 72–86. <https://doi.org/10.1016/j.ndteint.2018.10.002>
- [49] Daniel Roten, Yifeng Cui, Kim B Olsen, Steven M Day, Kyle Withers, William H Savran, Peng Wang, and Dawei Mu. 2016. High-frequency nonlinear earthquake simulations on petascale heterogeneous supercomputers. In *SC'16: Proc. Int. Conf. High Perform. Comput., Networking, Storage and Analysis*. IEEE, 957–968.
- [50] Tatsuhiko Saito, Toshitaka Baba, Daisuke Inazu, Shunsuke Takemura, and Eiichi Fukuyama. 2019. Synthesizing sea surface height change including seismic waves and tsunami using a dynamic rupture scenario of anticipated Nankai trough earthquakes. *Tectonophysics* (2019), 228166. <https://doi.org/10.1016/j.tecto.2019.228166>
- [51] Kirk Schloegel, George Karypis, and Vipin Kumar. 2002. Parallel static and dynamic multi-constraint graph partitioning. *Concurrency Comp.: Pract. and Exp.* 14, 3 (2002), 219–240.
- [52] Fengyan Shi, James T. Kirby, Jeffrey C. Harris, Joseph D. Geiman, and Stephan T. Grilli. 2012. A high-order adaptive time-stepping TVD solver for Boussinesq modeling of breaking waves and coastal inundation. *Ocean Modelling* 43–44 (Jan 2012), 36–51. <https://doi.org/10.1016/j.ocemod.2011.12.004>
- [53] Anthony Sladen, Diane Rivet, Jean-Paul Ampuero, Louis De Barros, Yann Hello, Gaëtan Calbris, and Patrick Lamare. 2019. Distributed sensing of earthquakes and ocean-solid Earth interactions on seafloor telecom cables. *Nat. Commun.* 10, 1 (2019), 1–8.
- [54] J. S. Sochacki, J. H. George, R. E. Ewing, and S. B. Smithson. 1991. Interface conditions for acoustic and elastic wave propagation. *GEOPHYSICS* 56, 2 (1991), 168–181. <https://doi.org/10.1190/1.1443029>
- [55] Y. Tony Song, L.-L. Fu, Victor Zlotnicki, Chen Ji, Vala Hjorleifsdottir, C.K. Shum, and Yuchan Yi. 2008. The role of horizontal impulses of the faulting continental slope in generating the 26 December 2004 tsunami. *Ocean Modell.* 20, 4 (2008), 362–379. <https://doi.org/10.1016/j.ocemod.2007.10.007>
- [56] W.J. Stephenson, N.G. Reitman, and S.J. Angster. 2017. *U.S. Geological Survey Open-File Report 2017–1152: P- and S-wave velocity models incorporating the Cascadia subduction zone for 3D earthquake ground motion simulations, version 1.6—Update for Open-File Report 2007–1348*. Technical Report. U.S. Geol. Surv. <https://doi.org/10.3133/ofr20171152>
- [57] Vladimir A. Titarev and Eleuterio F. Toro. 2002. ADER: Arbitrary high order Godunov approach. *J. Sci. Comput.* 17, 1 (2002), 609–618.
- [58] Eleuterio F. Toro. 2009. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer Berlin Heidelberg. <https://doi.org/10.1007/b79761>
- [59] Thomas Ulrich, Alice-Agnes Gabriel, and Elizabeth Madden. 2021. Stress, rigidity and sediment strength control megathrust earthquake and tsunami dynamics. *Nat. Geosci.* (2021). <https://doi.org/10.31219/osf.io/9kdhb>
- [60] Thomas Ulrich, Stefan Vater, Elizabeth H. Madden, Jörn Behrens, Ylona van Dinther, Iris Van Zelst, Eric J. Fielding, Cunren Liang, and A.-A. Gabriel. 2019. Coupled, physics-based modeling reveals earthquake displacements are critical to the 2018 Palu, Sulawesi tsunami. *Pure Appl. Geophys.* 176, 10 (2019), 4069–4109.
- [61] Carsten Uphoff and Michael Bader. 2020. Yet Another Tensor Toolbox for Discontinuous Galerkin Methods and Other Applications. *ACM Trans. Math. Softw.* 46, 4, Article 34 (Oct. 2020), 40 pages. <https://doi.org/10.1145/3406835>
- [62] Carsten Uphoff, Sebastian Rettenberger, Michael Bader, Elizabeth H. Madden, Thomas Ulrich, Stephanie Wollherr, and Alice-Agnes Gabriel. 2017. Extreme scale multi-physics simulations of the tsunamigenic 2004 Sumatra megathrust earthquake. In *Proc. Int. Conf. High Perf. Comput., Networking, Storage and Analysis (SC '17)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3126908.3126948>
- [63] Jim Verner. 2013. Jim Verner's Refuge for Runge-Kutta Pairs. <http://people.math.sfu.ca/~jverner/>. Accessed: 2020-04-01.
- [64] Lucas C. Wilcox, Georg Stadler, Carsten Burstedde, and Omar Ghattas. 2010. A high-order discontinuous Galerkin method for wave propagation through coupled elastic-acoustic media. *J. Comput. Phys.* 229, 24 (Dec. 2010), 9373–9396. <https://doi.org/10.1016/j.jcp.2010.09.008>
- [65] Stephanie Wollherr, Alice-Agnes Gabriel, and Carsten Uphoff. 2018. Off-fault plasticity in three-dimensional dynamic rupture simulations using a modal Discontinuous Galerkin method on unstructured meshes: implementation, verification and application. *Geophys. J. Int.* 214, 3 (sep 2018), 1556–1584. <https://doi.org/10.1093/gji/ggy213>
- [66] Brian J. N. Wylie. 2020. Exascale potholes for HPC: Execution performance and variability analysis of the flagship application code HemeLB. In *2020 IEEE/ACM Int. Workshop HPC User Support Tools (HUST) & Workshop Program. & Perform. Visual Tools (ProTools)*. IEEE. <https://doi.org/10.1109/hustprotocols51951.2020.00014>
- [67] Zhongwen Zhan, Mattia Cantono, Valey Kamalov, Antonio Mecozzi, Rafael Müller, Shuang Yin, and Jorge C. Castellanos. 2021. Optical polarization-based seismic and water wave sensing on transoceanic cables. *Science* 371, 6532 (2021), 931–936. <https://doi.org/10.1126/science.abe6648>

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

In this appendix, we describe technical details for the simulations in our paper, which are required to reproduce the respective results. All earthquake simulations were performed by the code SeisSol (www.seissol.org), using the extensions towards elastic-acoustic coupling and gravity boundary conditions featured in the manuscript. The respective source code is available in a persistent artifact. For shallow water simulations, we use a solver implemented in sam(oa)², which is available as open source on <https://gitlab.lrz.de/samoa/samoa>.

Additionally, all used input files together with documentation are available in the persistent artifact.

In the following, we first describe shared environment variables, followed by job scripts for the individual experiments.

Briefly, we compiled SeisSol with the following commands

```
CC=mpicc
CXX=mpicpc
FC=mpiifort
cmake -DNETCDF=ON \
  -DMETIS=ON \
  -DCOMMTHREAD=ON \
  -DHDF5=ON \
  -DCMAKE_BUILD_TYPE=Release \
  -DTESTING=OFF \
  -DLOG_LEVEL=warning \
  -DLOG_LEVEL_MASTER=info \
  -DHOST_ARCH={{ arch }} \
  -DPRECISION=double \
  -DNUMA_AWARE_PINNING=ON \
  -DASGI=ON \
  -DORDER=6 \
  -DPLASTICITY={{ plasticity }} \
  -DMINI_SEISSOL={{ mini_seissol }}
-- {{seissol_source_dir}}
make
```

The placeholder {{ arch }} takes the value skx, hsw, rome for experiments on SuperMUC-NG, Shaheen and Mahti respectively. For the Palu scenario, the placeholder {{ plasticity }} is set to ON, for all other scenarios it is set to OFF. We always set the placeholder {{ mini_seissol }} to ON for all experiments except when we wish to deactivate the node level weights for load balancing. A detailed documentation on how SeisSol can be build can be found at <https://seissol.readthedocs.io/>. This documentation is also contained in the artifact.

We always used the following environment flags when running SeisSol:

```
export XDMFWRITE_ALIGNMENT=8388608
export XDMFWRITE_BLOCK_SIZE=8388608
export SC_CHECKPOINT_ALIGNMENT=8388608
```

```
export SEISSOL_CHECKPOINT_ALIGNMENT=8388608
export SEISSOL_CHECKPOINT_DIRECT=1
```

```
export ASYNC_MODE=THREAD
export ASYNC_BUFFER_ALIGNMENT=8388608
```

```
export MP_SINGLE_THREAD=no
unset KMP_AFFINITY
```

These have to be added to all SeisSol jobscripsts.

For the sake of brevity, we only include slurm jobfiles for our largest production scenario and for the strong scaling here in this appendix. The jobfiles for other SeisSol jobs can be obtained by simply changing the number of nodes and the runtime.

We used the following job script for the L mesh on Supermuc-NG:

```
#!/bin/bash
#SBATCH --account={{ project_id }}
#SBATCH --partition=large
#SBATCH --array=0-0

#SBATCH --nodes=3072
#SBATCH --ntasks-per-node=2
#SBATCH --no-requeue

#SBATCH -J ss_xxl

#SBATCH -o ./logs/%A_%a.out
#SBATCH -e ./logs/%A_%a.err

#SBATCH -D .

#SBATCH --mail-type=BEGIN,END,ARRAY_TASKS
#SBATCH --mail-user={{ email }}

#SBATCH --time=08:00:00
#SBATCH --export=NONE
#SBATCH --get-user-env

source /etc/profile.d/modules.sh
module load slurm_setup
export LD_LIBRARY_PATH=$HOME/lib:$LD_LIBRARY_PATH

ulimit -Ss 2097152

# Include environment definitions from above!
export OMP_NUM_THREADS=46
export OMP_PLACES="cores(23)"

mkdir output

mpiexec -n $SLURM_NTASKS ./SeisSol large_parameters.par

We used the the following jobscrip on Shaheen:

#!/bin/bash
#SBATCH --account={{ project_id }}
#SBATCH --partition=workq
#SBATCH --array=0-0
```



```

#SBATCH --nodes=6144
#SBATCH --ntasks-per-node=1
#SBATCH --cpus-per-task=64
#SBATCH --no-requeue

#SBATCH -J ss_xxl

#SBATCH -o ./logs/%A_%a.out
#SBATCH -e ./logs/%A_%a.err

#SBATCH -D .

#SBATCH --mail-type=BEGIN,END,ARRAY_TASKS
#SBATCH --mail-user={{ email }}

#SBATCH --time=04:00:00

source /etc/profile.d/modules.sh

# Include environment definitions from above!
export OMP_NUM_THREADS=62
export OMP_PLACES="cores(31)"

source /etc/profile.d/modules.sh

mkdir output

lfs setstripe -c 144 output

ulimit -Ss unlimited
srun --mem-bind=v,local --cpu-bind=threads \
    ./SeisSol large_parameters.par

We ran both SeisSol and Samoa-flash for the Earthquake-tsunami
benchmark on SuperMUC-NG, as described in the paper.

We ran our M Palu mesh on SuperMUC-NG, as described in the
paper. We ran the L mesh on SuperMUC-NG and on Shaheen, as
described in the paper. For the latter, we used all available nodes
exclusively. We ran the strong scaling studies on both Mahti and
SuperMUC-NG.

For this, we used the following script for Mahti:
#!/bin/bash
#SBATCH --account={{ project_id }}
#SBATCH --partition={{ partition }}
#SBATCH --array=0-0

#SBATCH --nodes={{ nodes }}
#SBATCH --ntasks-per-node={{ ranks_per_node }}
#SBATCH --hint=multithread
#SBATCH --cpus-per-task={{ 256 // ranks_per_node }}

#SBATCH -J {{ job_name }}

#SBATCH -o {{ workdir }}/logs/%A_%a.out
#SBATCH -e {{ workdir }}/logs/%A_%a.err

```

```

#SBATCH -D {{ workdir }}

#SBATCH --mail-type=BEGIN,END,ARRAY_TASKS
#SBATCH --mail-user={{ email }}

#SBATCH --time=00:10:00

ulimit -Ss unlimited

# Include environment definitions from above!
export OMP_NUM_THREADS={{ threads_per_rank }}
export OMP_PLACES="cores({{ cores_per_rank }})"

mkdir output

srun --cpu-bind=loms \
    --ntasks={{ nodes * ranks_per_node }} \
    --ntasks-per-node={{ ranks_per_node }} \
    --cpus-per-task={{ 256 // ranks_per_node }} \
    ./SeisSol ./parameters.par

We used the following on SuperMUC-NG:

#!/bin/bash
#SBATCH --account={{ project_id }}
#SBATCH --partition={{ partition }}
#SBATCH --array=0-0

#SBATCH --nodes={{ nodes }}
#SBATCH --ntasks-per-node={{ ranks_per_node }}
#SBATCH --no-requeue

#SBATCH -J {{ job_name }}

#SBATCH -o {{ workdir }}/logs/%A_%a.out
#SBATCH -e {{ workdir }}/logs/%A_%a.err

#SBATCH -D {{ workdir }}

#SBATCH --mail-type=BEGIN,END,ARRAY_TASKS
#SBATCH --mail-user={{ email }}

#SBATCH --time=00:10:00
#SBATCH --export=NONE
#SBATCH --get-user-env

source /etc/profile.d/modules.sh
module load slurm_setup
export LD_LIBRARY_PATH=$HOME/lib:$LD_LIBRARY_PATH

ulimit -Ss 2097152

# Include environment definitions from above!
export OMP_NUM_THREADS={{ threads_per_rank }}
export OMP_PLACES="cores({{ cores_per_rank }})"

source /etc/profile.d/modules.sh

```

```
module load slurm_setup
```

```
mkdir output
```

```
mpiexec -n $SLURM_NTASKS ./SeisSol ./parameters.par
```

All quantities in curly brackets are placeholders whose value either depends on your setup (e.g. project id) or on quantities that are described in the paper (e.g. number of nodes).

To estimate the single node performance and the strength of NUMA effects on the AMD Rome architecture, we used the following jobscript on Mahti:

```
#!/bin/bash
#SBATCH --job-name=proxy
#SBATCH --account={{ project_id }}
#SBATCH --partition=test
#SBATCH --time=00:10:00
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=1
#SBATCH --hint=multithread
#SBATCH --cpus-per-task=256

export PROXY_PATH="SeisSol_proxy_Release_drome_6_elastic"
export NUMBER_OF_CELLS=300000
export NUMBER_OF_TIMESTEPS=100

declare -a kernels=("local" "all")
declare -a domains=("numa" "socket" "node")
declare -A d2t=( ["numa"]="16"
  ["socket"]="64"
  ["node"]="128")
declare -A d2tasks=( ["numa"]="8"
  ["socket"]="2"
  ["node"]="1")

output_file="output.csv"
if [ ! -f ${output_file} ]; then
  echo "Initialise output"
  echo "kernel,domain,cores,threads,gflops,\" \
  \"number_of_cells,number_of_timesteps\" \
  > \"$output_file"
fi

for kernel in "${kernels[@]}"
do
  for domain in "${domains[@]}"
  do
    cores="${d2t[$domain]}"
    let threads="2*cores"
    tasks="${d2tasks[$domain]}"
    export OMP_PROC_BIND="close"
    export OMP_PLACES="cores(${cores})"
    export OMP_NUM_THREADS="${threads}"
    GFLOPS=$(srun --label --ntasks=${tasks} \
      --cpus-per-task=${threads} \
      --cpu-bind=ldoms ${PROXY_PATH} \
```

```
    ${NUMBER_OF_CELLS} \
    ${NUMBER_OF_TIMESTEPS} \
    ${kernel} \
    | grep "0: GFLOPS (hardware).*" \
    | grep -P -o "\d*\.\d*")
    echo "${kernel}",\
    "${domain}",\
    "${cores}",\
    "${threads}",\
    "${GFLOPS}",\
    "${NUMBER_OF_CELLS}",\
    "${NUMBER_OF_TIMESTEPS}"\
    >> ${output_file}
  done
done
```

This script generates a csv file that contains all measurements. We ran this for five times and took the maximum resulting number of GFLOPS. The path to the SeisSol-proxy (our performance reproducer) needs to be adjusted depending on the installation directory. We used the same compilation settings and SeisSol version as for the Palu simulations.

We ran samoa-flash as a comparison for the tsunami-benchmark scenario. For this, we used the following setup:

```
#!/bin/bash
#SBATCH -J samoa_ascete
#SBATCH -o ./logs/%j.out
#SBATCH -D .
#SBATCH --mail-type=ALL
#SBATCH --mail-user={{ email }}
#SBATCH --no-requeue
#SBATCH --account={{ project_id }}
#SBATCH --time=00:30:00
#SBATCH --nodes=16
#SBATCH --ntasks=16
#SBATCH --cpus-per-task=48
#SBATCH --exclusive
#SBATCH --mem=MaxMemPerNode
#SBATCH --ear=off

source /etc/profile.d/modules.sh

module load slurm_setup
module switch python/3.6_intel

export OMP_NUM_THREADS=48
export LC_NUMERIC="en_US.UTF-8"

ulimit -s unlimited

printf -v date '%(Y-%m-%d_%H_%M_%S)T' -1
output_dir="output/${date}"
mkdir -p "$output_dir"

mpiexec -np ${SLURM_NTASKS} samoa_0221d8b2 \
  -dmax 24 -dmin 16 \
  -courant 0.5 -fdisp disp_1000_withCoh.nc \
```

```
-displacement_resolution 500 -drytolerance 0.1d0 \
-refined_bathymetry 50.0 \
-stestpoints_file receiver.dat \
-tmax 120.0 -xdmfoutput \
-tout 20.0 -output_dir "$outdir"
```

Furthermore, we ran `samoa-flash` (with a different version/branch) as a comparison for a Palu setup. As this simulation is small, we ran it on a workstation. We used the following run script:

```
export OMP_NUM_THREADS=4
export MP_PE_AFFINITY=yes
export MP_TASK_AFFINITY=core:4

dmin=14
dmax=18
domain='-30000 1376000 22000 1476000'
# Initial Displacement height [m]
displacement_height=10
# Initial Displacement resolution [m]
displacement_resolution=100
# Bathymetry Refinement
refined_bathymetry=100
dry_tolerance=$(echo "scale=16; 0.01/9.80616" | bc -l)
#displacement file
disp="Sulawesi_WL_Batv1_CFL04bSul8e6_100s.nc"
#bathymetry file
bath="BATNAS_MERGED_TRIMMED_v1.0_50_smooth_proj.nc"
#output directory
outdir="outputSamoa/"
rm -rf $outdir
mkdir $outdir
```

```
#Point output file
points="pointsAroundBay.txt"
# domain: considered domain
# [-29000 1370000]x[42000 1440000]
# If you want the wavefield,
# add the argument -xmlpointoutput
samoa/bin/samoa_flash -fbath $bath -fdisp $disp \
-tmax 150 -dmin $dmin -dmax $dmax -tout 5.0 \
-max_displacement $displacement_height \
-displacement_resolution $displacement_resolution \
-refined_bathymetry $refined_bathymetry \
-output_dir $outdir -drytolerance $dry_tolerance \
-sections 1 -courant 0.5 -domain $domain -xdmfoutput
```

We computed the input displacement file from the medium sized SeisSol simulation using scripts available in the appendix. These scripts extract the seafloor displacements, apply the Tanioka filter and move it to the coordinate system used by the shallow water simulation.

We also include the input displacement file directly such that the tsunami simulation can be ran without running the SeisSol simulation beforehand.

Author-Created or Modified Artifacts:

Persistent ID: <https://doi.org/10.5281/zenodo.5159333>

Artifact name: Supplementary material and artifacts
↪ (including SeisSol)

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: SuperMUC-NG: 6336 nodes with dual socket Intel Skylake Xeon Platinum 8174 (24 cores each). The nodes are organized in 8 islands. A fat tree OmniPath network is used, between each island the connection is pruned (pruning factor 1:4); Shaheen: 6174 nodes with dual socket Intel Xeon E5-2698 v3. Nodes are connected with an Aries interconnect and Dragonfly topology; Mahti: 1404 nodes with dual socket AMD Rome 7H12 (64 cores each). Nodes are connected by a Mellanox HDR InfiniBand interconnect using a Dragonfly+ topology.

Operating systems and versions: SuperMUC-NG: SUSE Linux Enterprise Server 12 SP3 (Linux 4.4.180-94.100-default); Shaheen: SUSE Linux Enterprise Server 15 (Linux 4.12.14-150.66-default); Mahti: Red Hat Enterprise Linux Server 7.9 (Linux 3.10.0-1160.21.1.el7.x86_64)

Compilers and versions: Intel Compiler, version: SuperMUC-NG: 19.0.5.281 2019081; Shaheen: 19.0.5.281 20190815; Mahti: 19.1.1.217 20200306

Applications and versions: `samoa` (<https://gitlab.lrz.de/samoa/samoa>) commits `ca1e325e572463cc09fc1ff9d531f08a9e7aaa89` (earthquake-tsunami-benchmark) and `4571cac8bf864e20af673fa08e933d2e2e3b90db` for (Palu)

Libraries and versions: NCO (netCDF Operators): version 4.9.3; GDAL: 3.1.4; seissolxdmf: 0.0.8, <https://pypi.org/project/seissolxdmf/>; libxsmm: `db22ad26975adb39409a7803bd48e5f75beb963e`, <https://github.com/hfp/libxsmm>; asagi: `601886131ab3cda7e75e19d10964198093674048`, <https://github.com/TUM-I5/ASAGI>; PSpaMM: `344c06c183854f72224c1e88ad2ced2e092d4efb`, <https://github.com/peterwauligmann/PSpaMM>; fox: `3199302381316608fb11daad9bf93ae2166bfe13` <https://github.com/StarGate01/fox>; eigen: version included as submodule of SeisSol; SuperMUC-NG: `mpi.intel/2019, hdf5/1.8.20-intel-impi-frt-threadsafe, netcdf/4.6.1-intel-impi-hdf5v1.8-parallel, parmetis/4.0.3-intel-impi-i64-r64`; Shaheen: `cray-mpich/7.7.14, hdf5-1.8.21, netcdf-4.6.1, parmetis-4.0.3`; Mahti: `openmpi/4.0.3, hdf5/1.10.6-mpi, netcdf-4.6.1, parmetis-4.0.3`

Key algorithms: ADER-DG, Runge-Kutta (Verner's "most efficient" order 7 without adaptive error estimate and without interpolation), Samoa: RK-DG

Input datasets and versions: ASCETE-Earthquake-Tsunami Scenario

URL to output from scripts that gathers execution environment information.

<https://doi.org/10.5281/zenodo.4899349>