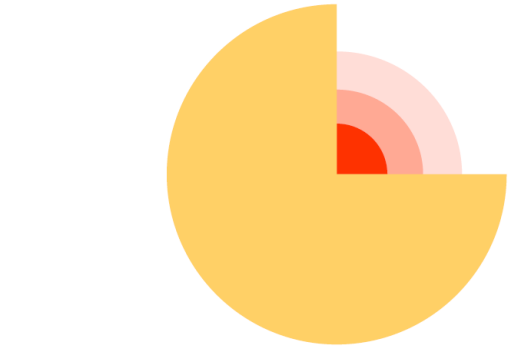

ChEESE

# D3.6 REPORT ON EXASCALE TESTBED DEMONSTRATION

## Version 1.0

## Document Information

| | |
|---|---|
| Contract Number | 823844 |
| Project Website | https://cheese-coe.eu/ |
| Contractual Deadline | 31/03/2022 |
| Dissemination Level | PU |
| Nature | Report |
| Author | Marta Pieńkowska (ETH) |

| | |
|---|---|
| Contributors | Michael Bader (TUM), Andreas Fichtner (ETH), Alice-Agnes Gabriel (LMU), Vadim Monteiller (LMA), Soline Laforêt (BULL) |
| Reviewers | Arnau Folch (BSC), Josep de la Puente (BSC) |

*The ChEESE project has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement No 823844*

## Change Log

| Version | Description of Change |
|---|---|
| V0.1 | Document skeleton and template for case studies. 18/02/2022 |
| V1.0 | First version for internal review. 22/03/2022 |
| V1.1 | Reviewed by J. de la Puente and A. Folch. 26/03/2022 |
| V1.2 | Address contributors' and reviewers' suggestions. 29/03/2022 |
| V1.3 | Final version. 30/03/2022 |

# Index

# 1. Introduction

The ChEESE Center of Excellence prepares 10 flagship codes for exascale computing and develops a series of Pilot Demonstrators (PDs) as candidates for services executed on upcoming exascale machines. This deliverable, D3.6, reports on achievements in Task 3.5 (T3.5) "Pre-Exascale testbed execution" that aims to demonstrate the suitability and performance of selected flagship codes and respective workflows on relevant test cases on the largest scales.

## 1.1 Role of this Deliverable

The objective of T3.5 is to run real scientific workloads on the largest available machines, and D3.6 reports on tested scenarios in the framework of:

- Pilot Demonstrator 1 (PD1) "Urgent seismic simulations",
- Pilot Demonstrator 4 (PD4) "Physics-based tsunami-earthquake interaction",
- Pilot Demonstrator 5 (PD5) "Physics-based Probabilistic seismic hazard assessment".

D3.6 completes the ChEESE milestone MS3 "Exascale testbed simulation" in WP3. It addresses one of the primary objectives of WP3: the demonstration of the suitability and performance of selected codes on the latest available large scale systems including the expected exascale systems. D3.6 therefore incorporates single runs of selected codes (capability runs) as well as pre-exascale workflow implementations as testbed runs (capacity runs). We remark that the activity in this deliverable addresses two Specific Objectives (SO) of the project, namely:

| **Specific objective** | **Target (as in the proposal)** | **Indicator** |
|---|---|---|
| Exascale software engineering, SO1 | Orient to exascale 10 Flagship SE codes | Increase by one order of magnitude the level of scalability of at least half of the codes |
| Scientific-technical challenge, SO3 | Once a European pre-exascale machine is deployed, apply to PRACE resources to run a pre-exascale testbed on $10^5$-$10^6$ cores | Obtain a minimum of two PRACE projects during the 2nd and 3rd year of the project |

Furthermore, the deliverable showcases the impact of developments at the code optimization level (carried out in WP2) on extreme-scale, real workloads.

## 1.2 Targeted Hardware

Task T3.5 aims at running "real scientific workloads on the largest European available HPC systems deployed over the course of the project". The three European pre-exascale

supercomputers (Leonardo, Mare Nostrum 5, and LUMI) to be installed in the framework of the EuroHPC Joint Undertaking were the prime targets for T3.5, with two of these systems at sites participating in ChEESE (Leonardo at CINECA and Mare Nostrum 5 at the BSC).

All three machines were scheduled to be installed towards the end of 2020 and become operational in early 2021, i.e. expected to be available during the final project year (M25-36) of ChEESE. However, due to significant deployment delays, even access during the hardware testing phases - prior to PRACE calls - was out of the question. This exacerbated the risk identified in D3.4 of "insufficient access to computational resources", as our pre-exascale target machines remained unavailable. Our test-case scenarios therefore were executed on the following (available) machines:

1. **Frontera** (TACC, USA, www.tacc.utexas.edu/systems/frontera), based on compute nodes with two Intel Xeon Platinum 8280 CPUs (28 cores per CPU).
2. **Jean-Zay** (IDRIS, CNRS, France,www.idris.fr/eng/jean-zay/index.html) based on compute nodes with two Intel Cascade Lake 6248 CPUs (20 cores per CPU) and four NVIDIA Tesla V100 GPUs.
3. **JURECA-DC** (FZJ, Germany, apps.fz-juelich.de/jsc/hps/jureca/configuration.html) based on compute nodes with two AMD EPYC 7742 CPUs (64 cores per CPU) and for NVIDIA A100 GPUs.
4. **Marconi100** (CINECA, Italy, www.hpc.cineca.it/hardware/marconi100) based on compute nodes with two IBM POWER9 AC922 CPUs (16 cores per CPU) and four NVIDIA Volta V100 GPUs per node.
5. **Mare Nostrum 4** (BSC, Spain, www.bsc.es/marenostrum/marenostrum) based on compute nodes with two Intel Xeon Platinum 8160 CPUs (24 cores per CPU).
6. **Piz Daint** (CSCS, Switzerland, www.cscs.ch/computers/piz-daint) based on compute nodes with an Intel Xeon E5-2690 CPU (12 cores) and a NVIDIA Tesla P100 GPU.
7. **SuperMUC-NG** (LRZ, Germany, doku.lrz.de/display/PUBLIC/SuperMUC-NG) based on compute nodes with two Intel Xeon Platinum 8174 CPUs (24 cores each).

## 1.3 Selected Pilot Demonstrators

### *1.3.1 Pilot Demonstrator PD1*

PD1 (urgent seismic computing, led by ETHZ) has been chosen for the pre-exascale tests within ChEESE, as seismic wave propagation is currently computationally prohibitive at high frequencies relevant for earthquake engineering or for civil protection purposes. With

the large-scale runs with Salvus (see Section 2) and SPECFEM3D Cartesian (see Section 3) we hope to demonstrate that deterministic modelling of ground motions can in the future contribute to the assessment of seismic hazard, both earthquake- and explosion-generated. Developments of computational infrastructures will render high-frequency simulations affordable and thus enable new approaches, including urgent seismic simulations.

### *1.3.2 Pilot Demonstrator PD4*

PD4 (physics-based earthquake-tsunami interaction, led by TUM & LMU) sets the stage for the exascale testbed demonstration with SeisSol. In WP4, we developed a fully-coupled elastic-acoustic model that simulates dynamic rupture, seismic waves and ground motion, acoustic waves in the ocean and tsunami genesis (and onset of propagation) in a single scenario. Such simulations, conducted for the first time in 3D, are computationally extremely challenging, since waves in the acoustic layer need to be resolved simultaneously to the non-linear problem of dynamic rupture and seismic wave propagation of the earthquake simulation. We tackled two extreme-scale earthquake-tsunami scenarios - one to simulate the 2018 Palu-Sulawesi earthquake-tsunami event and one European scenario around the Hellenic Arc, Greece.

### *1.3.3 Pilot Demonstrator PD5*

PD5 (physics-based PSHA, led by LMU & TUM), explores workflows for physics-based Probabilistic Seismic Hazard Analysis (PSHA) empowered by HPC. PD5 provides workflows of data-integrated and data-verified multiple forward simulations of varying complexity - ranging from simple logic-tree approaches to kinematic source modelling to multi-physics dynamic rupture scenarios - towards physics-based PSHA. In PD5, a range of physics-based forward and ensemble simulations have been developed constrained by regional observations yielding realistic earthquake magnitudes. All use cases provide unique PSHA challenges, complexities and physics-based implications, such as fault interaction, site effects, and secondary hazards. For this deliverable, we focus on North Iceland. The Húsavík–Flatey fault system in north Iceland features complex geometries crossing from off-shore to on-shore and consisting of multiple right-lateral strike slip segments distributed across ~100 km. It provides a unique study region to assess fault interaction, dynamic and static stress transfers and rupture jumping across a complex fault network. A range of physics-based dynamic rupture simulations have been informed by regional observations in PD5 yielding magnitudes comparable to historic events. These scenarios lend themselves for analysis of uncertainties in earthquake parameterization using ExaHyPE.

## 1.4 Summary of Achievements

With 3 Pilot Demonstrators and 4 codes, a range of testbed scenarios have been proposed and executed in T3.5. Although the European pre-exascale machines remained unavailable during the ChEESE project, we have demonstrated the suitability of the selected flagship codes to run real scientific workloads on the largest scales, showcasing our progress and continued efforts towards running on the upcoming exascale systems. Many challenges and bottlenecks for complex realistic set-ups have been identified and addressed for each code, bringing the ChEESE pilot demonstrators closer to routine executions on exascale systems in the future. In T3.5 we have successfully performed:

(1) A single Salvus simulation on 3,000 MN4 nodes, on a mesh with ~215 million elements and accurate up to 20 Hz. With this capability run - the largest regional Salvus simulation to date - we have tested Salvus on a realistic scenario that is relevant for the specific seismic urgent computing problem set-up in PD1 and that goes beyond simplistic scaling problems.

(2) Two sets of 66 simultaneous Salvus simulations up to 5 Hz to test a suite of input parameters generated by the PD1 UCIS4EQ workflow (capacity run on 3,000 MN4 nodes). This was the first test of the workflow that has been performed in an urgent-like environment, where jobs entered the queue immediately and completed simultaneously.

(3) A first complete SPECFEM3D Cartesian simulation accurate up to 35 Hz with 2 million time steps which allows us to simulate 25 seconds of signal on 1024 V100 GPUs in 5 hours and 45 minutes. The goal of the run is to pave the way for reliably assessing the risks of damage to buildings on the shore, induced by the detonation of unexploded historical ordnance of large weights in variable shallow water environments with a water depth of less than 50 m.

(4) A full SeisSol production simulation of the first 30s of the Palu-Sulawesi earthquake and tsunami generation on 3,072 nodes of SuperMUC-NG, on a mesh with ~518 million elements (resulting in 261 billion degrees of freedom). The 30s allow to capture the entirety of the earthquake dynamics and the tsunamigenesis. This is the largest simulation of the SeisSol PD4 fully-coupled elastic-acoustic model (and also the largest SeisSol simulation) to date and incorporates dynamic rupture, seismic waves and ground motion, acoustic waves in the ocean and tsunami genesis (and onset of propagation) in a single scenario.

(5) A thorough performance study of the SeisSol Palu-Sulawesi scenario on Frontera in collaboration with TACC. SeisSol achieved more than 9 PFlop/s on 7,000 Frontera nodes. The parallel efficiency from 250 to 7,000 nodes in this experiment was more than 98% for the largest of the tested meshes.

(6) An ExaHyPE simulation on 731 compute nodes of SuperMUC-NG on a 14-million-element grid (896 million integration points, which corresponds to using ~8 billion degrees of freedom for the seismic wave equation). The results obtained with ExaSeis have been compared with simulations that use SeisSol with the same physical parameters, showing excellent agreement of the two codes. The simulation demonstrates fully automated meshing with curvilinear approximation of topography and fault geometry, thus preparing UQ workflows that vary fault location and geometry.

The remainder of this document describes the testbed scenarios in detail and their execution for each of the selected codes. A brief summary is provided in Table 1.1.

**Table 1.1.** Summary of T3.5 testbed scenarios.

| PD | code | testbed type | resources | case definition |
|---|---|---|---|---|
| 1 | Salvus | single run and workflow | Mare Nostrum 4 (3000 nodes), Piz Daint (90 GPUs) | Eastern Mediterranean Urgent Seismic Computing |
| 1 | SPECFEM3D Cartesian | single run | Marconi100 (1024 GPUs), Jean-Zay (512 GPUs) and Jureca-DC (512 GPUs) | Coupled acoustic-elastic simulations of misfires in the Mediterranean sea |
| 4 | SeisSol | single run | SuperMUC-NG (3,072 nodes),<br>Frontera (7000 nodes) | Palu-Sulawesi Tsunami-Earthquake Interaction |
| 5 | ExaHyPE | single run | SuperMUC-NG (731 nodes) | North Iceland Regional PSHA |

# 2. Salvus: Eastern Mediterranean Urgent Seismic Computing on Mare Nostrum 4 and Piz Daint

Salvus (https://mondaic.com/) is under continuous development, with GPU and parallel unstructured higher-order mesh inputs implemented and validated in February 2020 (Hapla et al., 2020) and relevant scaling tests performed on Piz Daint at CSCS. T3.5 provided an excellent opportunity to demonstrate the suitability of Salvus for exascale machines and for complex and realistic use cases, beyond simplistic scaling problems, and in particular for the PD1 urgent computing use case. In the Salvus T3.5 executions we have realised two testbed type scenarios:

1. A large single run that allowed us to test for PD1-specific bottlenecks when upscaling the problem (a capability run), and
2. A workflow execution where a full set of a suite of simulations has been performed (a capacity run).

## 2.1 Goals of the demonstration run

The primary goals of the pre-exascale demonstration runs with Salvus in the context of PD1 were twofold: to test the solver at extreme scales with a single large-scale capability run, as well as to test the entire seismic urgent computing workflow (UCIS4EQ integrated with Salvus) at scale with a suite of runs.

### *2.1.1 Capability: a single run up to 20 Hz*

The aim of the capability run was to test the performance of Salvus on a realistic scenario that is relevant for seismic urgent computing and that goes beyond simplistic scaling problems. Such a large production-like run includes a range of complexities that have previously not been tested at scale, such as the topography, the bathymetry, as well as a kinematic representation of the fault. Thus, the goal was both to test the scalability of Salvus for a complete and realistic set-up, as well as to identify bottlenecks specific to the seismic urgent computing use case that need to be addressed to ensure that the proposed service can be scaled-up to larger problems.

### *2.1.2 Capacity: 132 runs up to 5 Hz*

The aim of the capacity run was to test the entire PD1 workflow, the UCIS4EQ integrated with Salvus, that generates a suite of runs for a given alert. This was the first test of the workflow that has been performed in an urgent-like environment, where jobs entered the queue immediately and completed simultaneously. It is also the first high-frequency suite execution that included the effects of the topography, of the bathymetry and of the water

layer. Such simulations at regional scales with frequencies relevant for structural engineering (up to 10 Hz) are crucial to reproduce the right orders of magnitude of recorded ground motion proxies (such as peak ground acceleration, peak ground velocity, or spectral accelerations, among others).

## 2.2 The scenario

The Salvus PD1 scenario focused on simulating the October 30th 2020 Mw 7.0 Samos-Izmir earthquake. The earthquake resulted in 118 fatalities and over 1000 injuries, as well as building damage and collapse. Wide-ranging effects have been observed, including local zones of high intensity shaking and a moderate tsunami. A range of meshes have been prepared to execute the scenario on 3000 MN4 nodes, with the aim of maintaining about 1500 elements per core (see Table 2.1). All meshes have been generated with elements of polynomial degree 7 (512 grid points per element), and given the offshore location of the event they included both the topography and the bathymetry, as we accounted for the water layer in the simulations. The goal was to:

(1) Run a single scenario on the largest mesh, with nearly 215 million elements that resolves the frequency up to 20 Hz (capability run).

(2) Run two sets of 66 simultaneous simulations up to 5 Hz to test a suite of input parameters generated by the workflow (capacity run).

**Table 2.1.** The summary of the meshes that have been prepared for the range of Samos-Izmir scenarios with different maximum resolved frequency.

| max resolved frequency | # of mesh elements (polynomial degree 7, i.e. 512 grid points per element) | planned # of simultaneous runs | planned nodes per run | planned cores per run | planned elements per core |
|---|---|---|---|---|---|
| 20 | 214,244,147 | 1 | 3000 | 144000 | 1487 |
| 15 | 91,952,202 | 2 | 1500 | 72000 | 1277 |
| 12 | 48,217,095 | 4 | 750 | 36000 | 1339 |
| 10 | 28,524,450 | 8 | 375 | 18000 | 1584 |
| 5 | 4,043,069 | 66 | 45 | 2160 | 1871 |

The meshes for 15, 12 and 10 Hz have been generated to provide fallback scenarios in case of insurmountable issues for the largest capability run at 20 Hz - we had very limited time to test very large-scale problems and thus we have designed multiple scenarios to mitigate the risks of failure (see Section 2.3.3 for further details on the challenges). Those scenarios would have resulted in 2, 4 and 8 simultaneous simulations, respectively, that would have

occupied the 3000 nodes of MN4. Those fallback meshes for lower-frequency capability scenarios have ultimately not been used.

## 2.3 The execution

For the purpose of PD1 T3.5 we have submitted the proposal AECT-2021-1-0030 “Urgent Earthquake Simulation Demonstrator” via the Spanish Supercomputing Network (Red Espanola de Supercomputacion, RES) and have been granted access to the entire Mare Nostrum 4 (MN4) supercomputer at the BSC for a period of 30 hours starting at 6 PM on Sunday April 11th 2021. The full machine was made available for the PD1 testbed scenarios just before MN4 scheduled maintenance, and we have received continuous assistance from the head of user support at the BSC to perform all large tests and resolve issues. The RES project accounted for 17% core-hours in excess of the allocated 5.1 Mcore-hours for the proposal.

### *2.3.1 Computational resources*

We had 30-hour access to over 3000 MN4 nodes, with each node composed of two Intel Xeon Platinum chips with 24 processors per chip (a total of 144,000 processors). We have used 3000 nodes, with the remaining ~100 as back up if nodes were to fail. We have performed:

- Up to 20 Hz: a single simulation on 3000 MN4 nodes and a total of 84,000 cores, using 28 out of the 48 available processors due to RAM issues with mesh decomposition at start-up (see “Mesh decomposition” in Section 2.3.3 for further information). With a mesh of 214M elements, 512 grid points per element, and 14 fields per grid point, the number of unknowns amounted to over 1.5 trillion. The simulation ran for 6 hours, completing 107,224 time steps, and terminated on the final night of the reservation due to a node failure (see “Checkpointing” in Section 2.3.3). Accurate up to 20 Hz, this is the largest regional Salvus simulation to date despite the interrupted execution.
- Up to 5 Hz: 132 simulations on 45 MN4 nodes each (using all 48 cores, so 2160 cores per run), with a set of 66 running simultaneously on 3000 MN4 nodes. A single set completed in 3.5 hours. An example CPU usage per node, with 48 tasks per node, can be seen in Figure 2.1.n example CPU usage per node, with 48 tasks per node, can be seen in Figure 2.1.

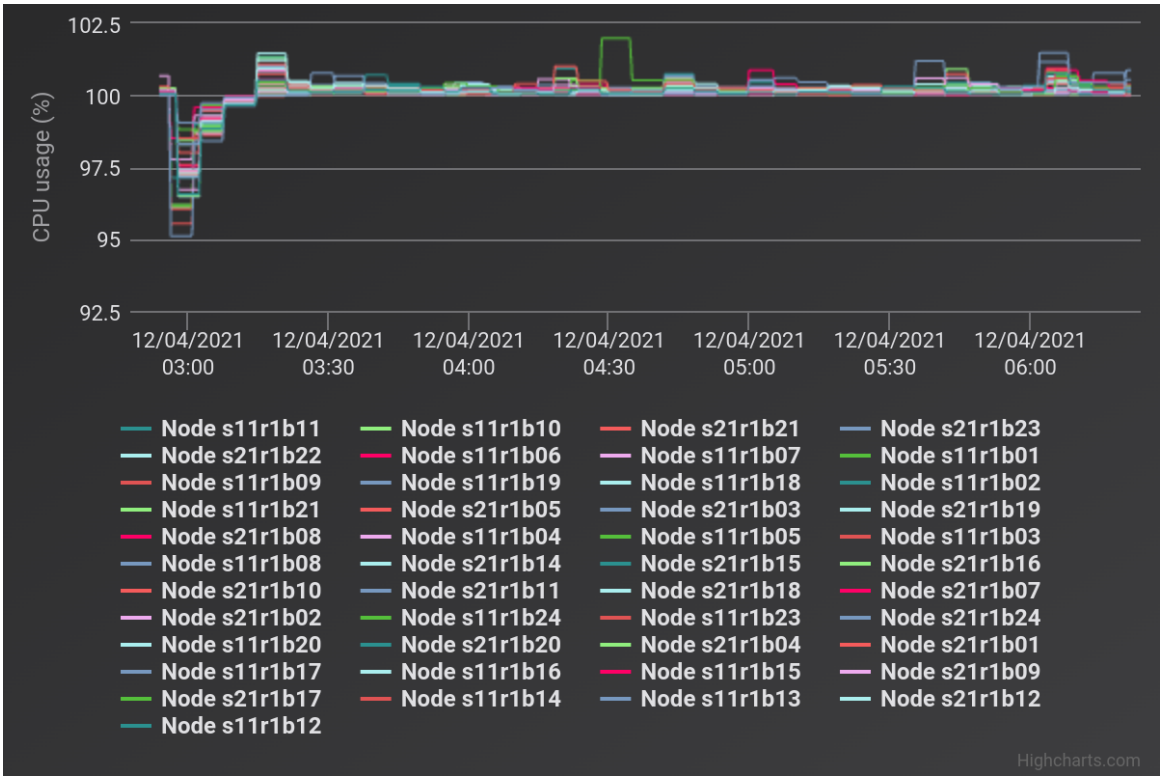


**Figure 2.1.** A sample CPU usage per MN4 node for a single 45-node Salvus simulation with 48 tasks per node.

Further 5 Hz production runs have been executed on Piz Daint in order to test the complex high-frequency set up for PD1 on GPUs and to explore the density of raw outputs. One single 5 Hz simulation takes less than 1.5h on 90 Piz Daint GPUs (with denser output sampling than the runs on MN4, thus increasing the I/O and the entire runtime), allowing us to reach time-scales that are relevant for an urgent computing service. With upcoming exascale infrastructures, running problems on such scales should become more routine, making the computational cost of the service more affordable.

### *2.3.2 Results*

We have shown that the set-up required for PD1 simulations can be scaled up with Salvus to regional simulations with frequency resolution up to 20 Hz. Figure 2.2 shows simulation snapshots for the Samos-Izmir case study, where the multitude of point sources that provide a kinematic description of the earthquake fault are clearly visible. We have also tackled and overcome challenges that now helped us to define further developments in Salvus that are required for the PD1 service prototype to become a pre-exascale service. These are mostly related to the pre-processing and the set-up of the simulations, as well as to the PD1-specific output quantities. The scaling of the time-loop of Salvus simulations on GPU architectures has been demonstrated in WP2, but the pre- and post-processing bottlenecks that are specific

to the workflow parameters remain to be solved for PD1-style production runs on pre-exascale machines. For further details on the challenges see Section 2.3.3.

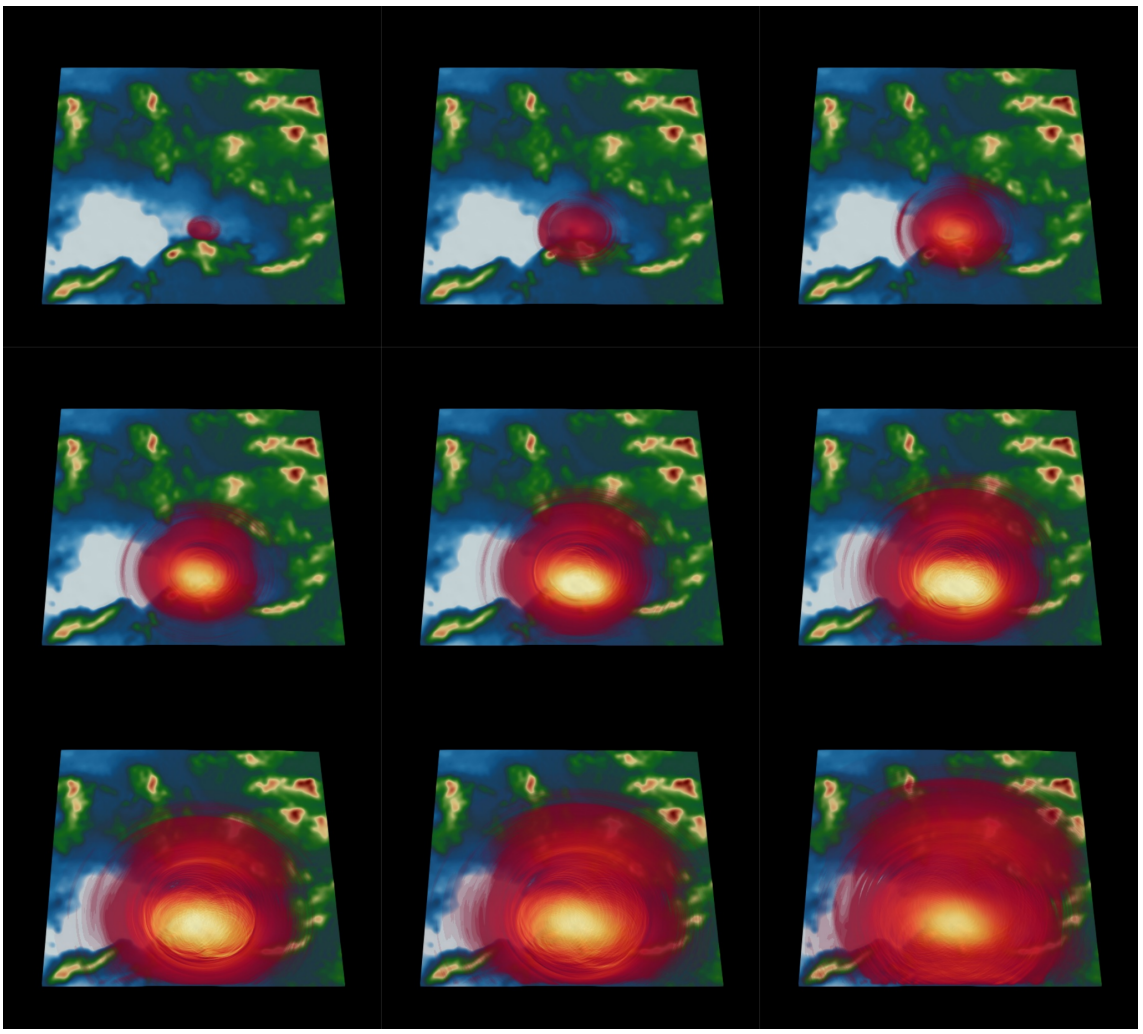

**Figure 2.2** Snapshots of wave propagation for the capability run up to 20 Hz for the 30th of October 2020 Samos-Izmir earthquake.

We have also successfully performed the first high-frequency capacity test with UCIS4EQ up to 5 Hz, generating high resolution spatial maps of synthetic ground motion parameters (see Figure 2.3). This is the highest-frequency run of a full suite of simulations within PD1 to date, and is an important milestone for the development of the prototype. Although a lot of work remains towards calibrating the runs, better defining the ensemble and the corresponding simulation parameters, and estimating uncertainties, this first high-frequency suite of runs demonstrated that the workflow can produce ground motion proxies on the right order of magnitude as seen in the data (see Figure 2.4).

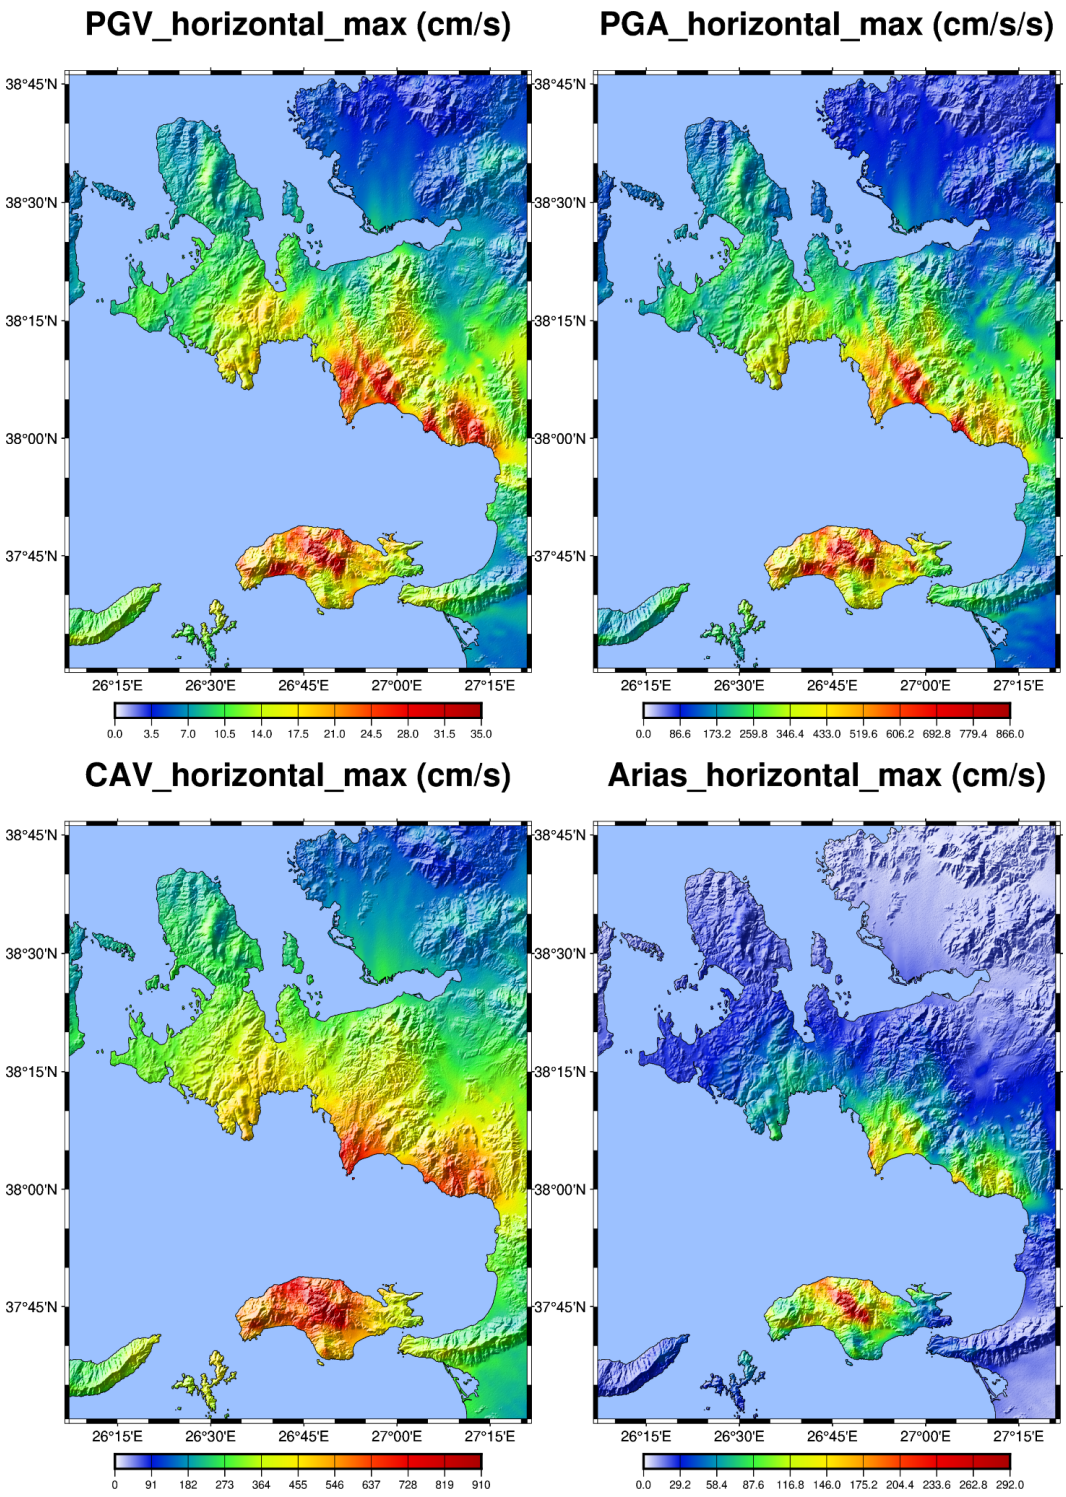


**Figure 2.3** The maximum of the two horizontal components from the simulations from the capability exercise. From top left: peak ground velocity (PGV), peak ground acceleration (PGA), cumulative absolute velocity (CAV), and Arias intensity. Given that full time-series are available on a dense grid, other ground motion parameters can be extracted at the post-processing stage.

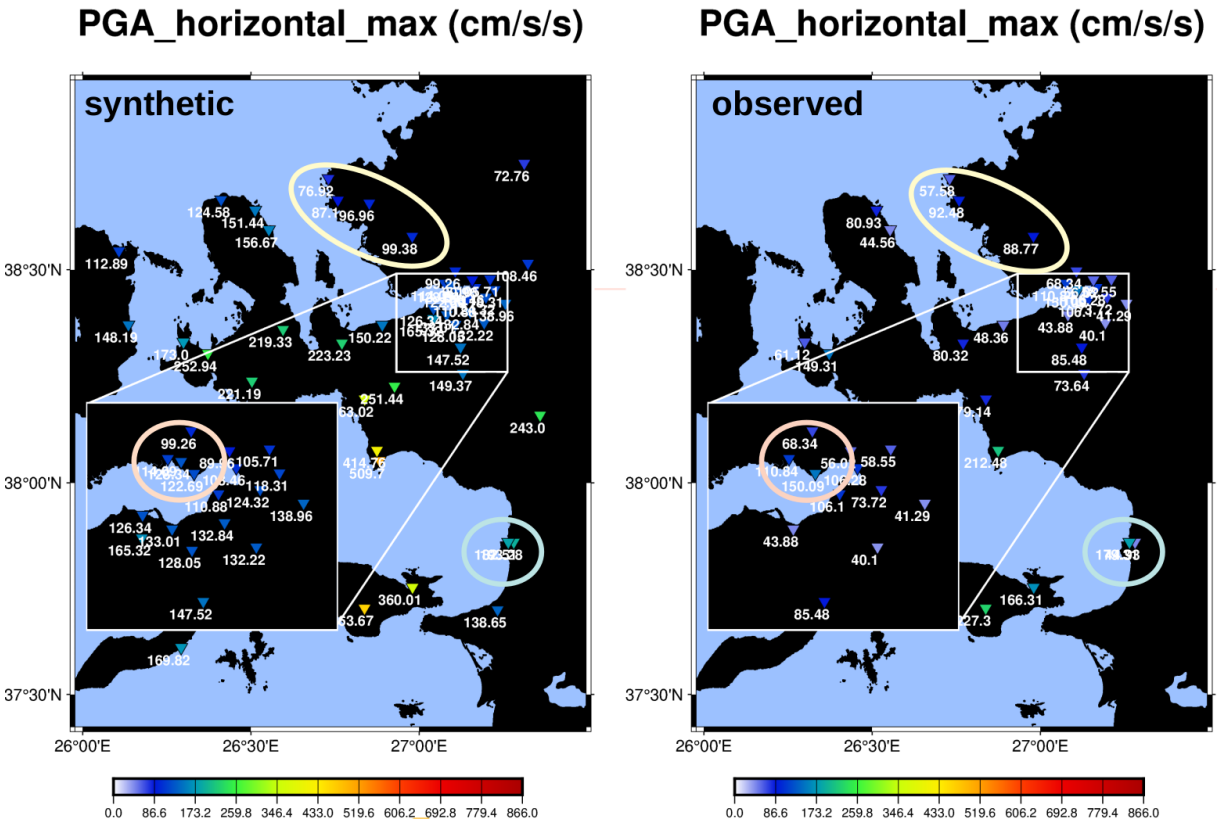


**Figure 2.4** The maximum PGA of the two horizontal components from the simulations from the capability exercise (left) compared with the PGA observed after the 30th of October 2020 Samos-Izmir earthquake (right). The synthetics at 5 Hz reproduce the right order of magnitude.

### *2.3.3 Challenges*

We have been scaling up the problem size and running large tests in preparation for the full-machine run in the run-up to the 11th of April. The regular queue allows access to a maximum of 250 MN4 nodes. From April 1st we could access up to 640 MN4 nodes, while in the last 48 hours, when the queue was being progressively drained for our exercise and the subsequent maintenance, we could start testing on over 1000 MN4 nodes. This restricted node access posed a significant challenge, as a few-fold increase in problem size exposes issues that cannot be otherwise identified on smaller-scale problems. We therefore were ready to run the capacity exercise in due time, but for the large capability run we were solving a range of issues on-the-fly. We therefore also had a range of smaller meshes ready as a fallback plan should the largest simulation prove impossible (see Section 2.2). In our attempts to run on the larges mesh we addressed the following:

- **Mesh generation:** We have uncovered a bug in chunk-wise mesh generation for non-cubic (i.e. realistic and more complex) mesh types. This meant that the meshes had to be generated on a machine with enough RAM to hold the entire mesh in memory, as the chunk-wise saving of the mesh to disk produced non-physical meshes. This poses a limitation on the mesh size: it can be as big as the RAM of the available machine. The chunk-wise mesh generation is thus crucial for further large-scale tests and overall for the PD1 service.

- **Mesh decomposition:** We experienced issues with PTScotch above 9600 cores and moved to a version of Salvus compiled with ParMETIS. Decomposing the mesh with ParMETIS resulted in issues above 90,000 cores. Specifically, ParMETIS threw an error that the weights did not sum up to 1 and we expect this to be due to single precision. ParMETIS is trying to solve a graph-cutting problem to decompose the mesh, and our guess is that with so many ranks and one weight per rank, single precision resulted in accumulating too many errors. We did not try double precision, because (with PETSc and PetscReal in Salvus) we would have had to use double precision for the simulation as well, making it too computationally expensive. This mesh decomposition stage of the start-up phase prevented us from using all 48 cores per node for the largest run, and thus also increased the runtime, as about 40% of the cores remained idle.

- **Placing sources on mesh:** The kinematic representation of the fault that is specific to our PD1 application means that a large number of point sources need to be placed on the mesh - the higher the frequency, the finer the point spacing is on the fault, and so the more points there are. Given the complexity of the mesh that includes both topography and bathymetry (crucial for the fidelity of regional simulations of strong ground motions in PD1), placing those sources correctly is not trivial. As we were scaling up in our tests, it transpired that the combination of more sources to place and a larger mesh to load in memory resulted in RAM issues and out-of-memory (OOM) errors. PD1 is not a ‘typical’ Salvus use case, and this has never been optimised for - it does not pose problems for smaller scale problems, and so efficient memory management for this specific pre-processing step has never been addressed. For the largest run we therefore had to increase the point spacing on the fault to decrease the number of point sources to successfully complete the set-up stage and enter the time-loop of the simulation. This issue of memory allocation in Salvus needs to be addressed in the next workflow releases.

- **Finite fault vs load balancing:** Another issue related to the PD1-specific kinematic representation of the fault is the load balancing. All sources are clustered in just a

small subset of elements in the simulation, as they are all on a fault close to the surface. This means that load balancing is not ensured, as the current mesh decomposition does not take this into account. It means that the finely discretized fault significantly influenced the time loop of the simulation, as only a relatively small subset of elements was treating those sources (the number of elements-updates per core per second decreased significantly).

- **Saving output:** We have uncovered a bug in saving a particular netCDF output file in parallel - for a smaller number of cores it did not result in a crash. Moreover, specific optimisation modules had to be loaded for reading from and saving to HDF5 files. The effect was negligible for smaller runs, but it has proven crucial for the larger execution. As it is not in MN4 documentation, we have only been able to manage this thanks to the continuous support of the head of user support at the BSC. It also has to be noted that for such large runs one has to carefully manage the runtime versus the desired outputs, as saving dense point output to disk increases the runtime. In the future our PD1 service may require developing specific quantities that are processed and output on-the-fly on the mesh points of the simulation (instead of densely sampled time series). This would ensure the most efficient output for our application.

- **Checkpointing:** Salvus does not have checkpointing implemented yet. This is a feature that is crucial to the robustness of our urgent PD1 application, an issue that has been exposed during the capability run that terminated halfway through due to the node failure. At this point of the reservation we did not have time remaining for another full run, and we could not re-run the simulation from a checkpointed state. From the point of view of the service as a whole, checkpointing is a necessary feature for seamless management of simulation ensembles.

### *2.3.4 Deviations from the work plan*

- **Scenario:** We have shifted our focus from the South Iceland case study proposed initially in D3.4 to the Mediterranean. We have decided to do so for two reasons: we have concluded that our largest PD1 test cases should be in a region where recent data is available for large and potentially destructive events, and we had a well-resolved 3D velocity model for the Mediterranean that we believe is more suitable for such large-scale runs than the velocity model for Iceland that was made available for us in the scope of PD1. This became apparent as we were progressing with our initial tests.

- **Machine:** As stated in Section 1.2, our initial plan to secure resources and run on the upcoming European pre-exascale machines was not feasible due to delays. We have therefore decided to apply for resources at the BSC, a key PD1 partner and developer, where we were able to secure the full machine and where all previous tests of UCIS4EQ have taken place. Although this meant that we were running on CPUs, it ensured continuity with our previous tests and allowed us to execute very large production runs with the complexity required for PD1. We focused on identifying case-specific bottlenecks related to meshing, I/O, data transfer, problem set-up and post-processing, and resolving such case-specific bottlenecks for PD1 is also going to be relevant for large-scale production runs on GPUs. It should be noted here that Salvus has also been shown to have good strong and weak scaling on Piz Daint GPUs for large benchmark case studies (for Salvus scaling on GPUs see WP2 deliverables, including the most recent deliverable D2.5).
- **Testbed type:** We have included the workflow-type test on top of the single large-scale simulation. The reservation of the full machine provided us with a unique opportunity to test the PD1 workflow at scale in an urgent-like setting, running a large set of high-frequency simulations simultaneously.

### 2.4 References and Publications

De la Puente, J., Rodríguez, J.E., Monterrubio-Velasco, M., Rojas, O., & Folch, A. (2020). Urgent Supercomputing of Earthquakes: Use Case for Civil Protection. PASC20, https://doi.org/10.1145/3394277.3401853.

Hapla, V., Knepley, M.G., Afanasiev, M., Boehm, C., van Driel, M., Krischer, L. and Fichtner, A. (2021). Fully Parallel Mesh I/O Using PETSc DMPlex with an Application to Waveform Modeling, SIAM Journal on Scientific Computing, 43:2, C127-C153, https://doi.org/10.1137/20M1332748.

# 3. SPECFEM3D Cartesian: Coupled Acoustic-Elastic Simulations of Misfires in the Mediterranean Sea on Marconi100, Jean-Zay and JURECA-DC

## 3.1 Goals of the demonstration run

The upstream assessment of noise pollution generated by anthropogenic activities in the coastal zone (e.g., offshore civil engineering works, disposal of historical explosive devices...) is a crucial issue. However, this accurate assessment is currently a scientific challenge, because it is generally done by using numerical models of wave propagation that suffer from many limitations (eg, one-way model, phase of the propagated signal poorly taken into account, 2D model, etc). In addition, the seabed is often considered as fluid and the study area is limited to the water layer, neglecting the seismo-acoustic exchanges in the vicinity and far from the source. The numerical methods classically used in underwater acoustics are ray or beam tracing, the normal modal method, and more particularly the parabolic equation method. Despite their own advantages, however, none of these methods is capable of solving the 3D propagation problem in the coastal zone, as they all rely on approximations that limit their validity domain.

On the contrary, full-wave numerical models can overcome all these limitations. However, their use is still very marginal in underwater acoustics, because modelling wave propagation in a wide range of frequencies (from a few Hz to a few kilohertz) and over large distances (from a few km to a few tens or even hundreds of km), while realistically taking into account the complexity of the marine environment (in particular the low shear velocity of the sediments), is currently out of reach, due to the exorbitant computational cost.

Nevertheless, it is now possible to start considering realistic 3D configurations in underwater acoustics using modern HPC resources and the objective is to design and run a model capable of simulating a realistic configuration.

## 3.2 The scenario

The goal of the run is to pave the way for reliably assessing the risks of damage to buildings on the shore, induced by the detonation of unexploded historical ordnance of large weights in variable shallow water environments with a water depth less than 50 m. We have chosen to consider a configuration of a counter-mining campaign carried out in the Mediterranean Sea south of France in the Rade d'Hyères. The aim of the run is to model a counter-mining explosion located in the coastal zone a few kilometres offshore. The 3D computation with SPECFEM3D Cartesian allows one to take into account the topography, bathymetry and

sedimentary layers. We have also included attenuation which can be very strong in the sediments.

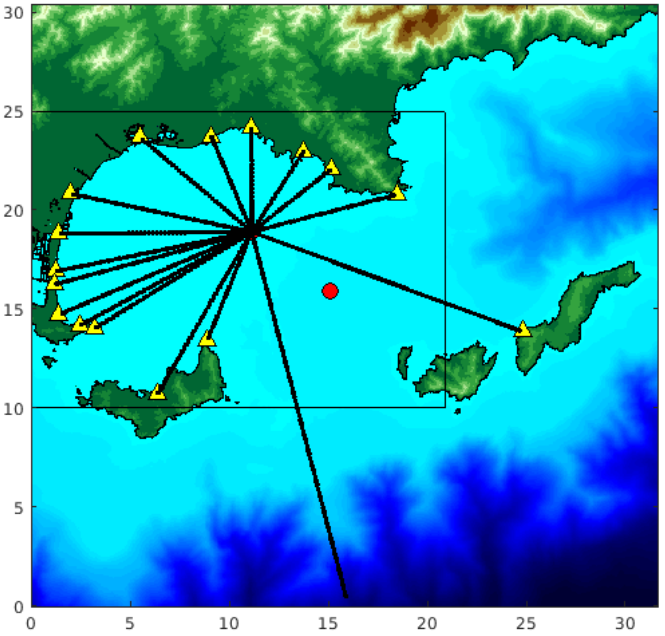


**Figure 3.1.** Acquisition geometry: the red dots are the shots, the yellow triangles are the stations. The black rectangle shows the mesh area.

In the south of France, the coastal area of the Rade d'Hyères was instrumented with seismometers that recorded the seismic waves produced by a counter-mining campaign. Figure 3.1 shows the acquisition geometry in this area and the computational domain. The aim is to simulate up to 35Hz. The main challenge is to correctly take into account the sediments in the simulation. Indeed, the sedimentary layers vary in depth from 1 m to 30m. The S-wave speed in these layers is about 200 m/s with an attenuation factor Qs about 30. This requires meshing with size elements less than 12m at the surface. In order to respect the geometry of the topography and the bathymetry, we used unstructured mesh. We design two meshes, one for strong scaling and another one for weak scaling. The first one uses a 10x10x6.5 m size element on the surface down to the base of the bedrock. At this layer, the mesh size has been multiplied by 3. The second mesh was built with 8x8x6.5m element size at the surface. The water layer is meshed from the bathymetry to the sea surface. A major difficulty is to mesh the transition between water and land on the beach. This would require too much deformation of the elements, so we chose to make a clean transition with 2 m high elements. Below this height we estimate that at 35Hz the water layers below 2 metres will

not affect the wave modelling. Indeed the wavelength at 35Hz in water is 42 metres, we can consider that a layer of water lower than 2 metres will not be significant.

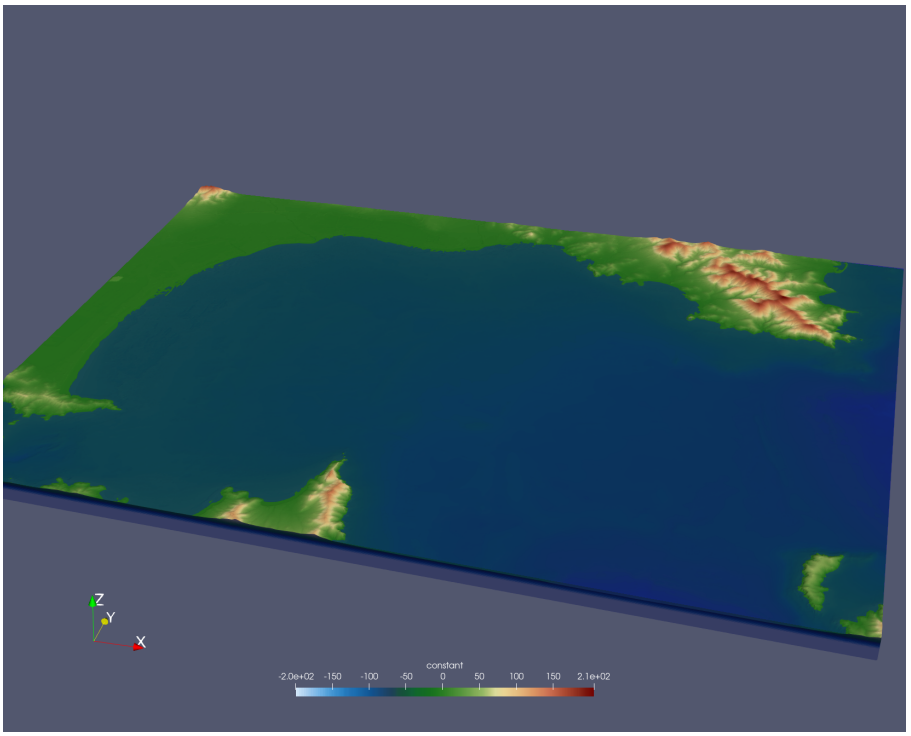


**Figure 3.2.** Mesh size of the area is 15x20x0.5km with 27 million hexahedral elements.

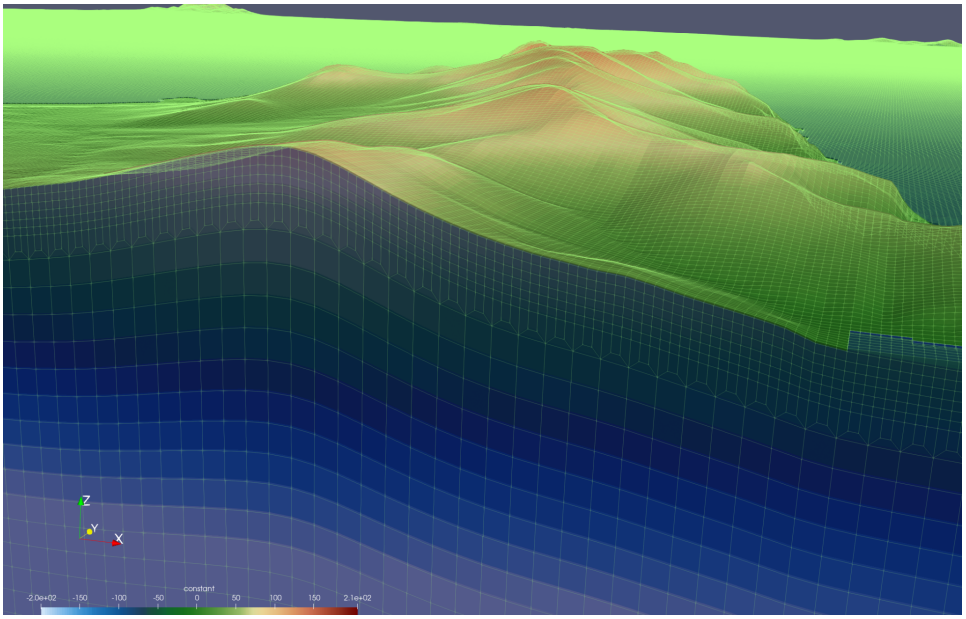


**Figure 3.3.** Detail of the mesh showing the tripling layer which allows the element size to be multiplied by 3 in a consistent way.

## 3.3 The execution

We used the mesh with 27M elements to make strong scaling runs and did weak scaling tests using the 100, 50 and 25M elements mesh (Table 3.1).

**Table 3.1.** The mesh 1 is used for strong scaling and the others are used for weak scaling performance tests.

| | Domain (km) | Element size at surface (m) | Number of elements | Number of degree of freedom | Required memory (TB) |
|---|---|---|---|---|---|
| **Mesh 1** | 20 x 15 x 0.5 | 10 x 10 x 6.5 | 27E+06 | 9E+09 | 3.5 |
| **Mesh 2** | 20 x 15 x 1 | 8 x 8 x 6.5 | 100E+06 | 3.3E+10 | 12.8 |
| **Mesh 3** | 10 x 15 x 1 | 8 x 8 x 6.5 | 50E+06 | 1.5E+10 | 6.4 |
| **Mesh 4** | 10 x 7.5 x 1 | 8 x 8 x 6.5 | 25E+06 | 8E+09 | 3.2 |

### *3.3.1 Computational resources*

We used NVIDIA V100 GPUs to do our tests on 3 different machines. The node configuration of the 3 computers used are :

- MARCONI 100 CINECA (https://www.hpc.cineca.it/hardware/marconi100):
  2x16 cores IBM POWER9 AC922 at 3.1 GHz
  4 x NVIDIA Volta V100 GPUs, Nvlink 2.0, 16GB
- JEAN-ZAY IDRIS (http://www.idris.fr/eng/jean-zay/index.html):
  2 x Intel Cascade Lake 6248
  4 GPU Nvidia Tesla V100 SXM2 32 Go
- JURECA (https://apps.fz-juelich.de/jsc/hps/jureca/configuration.html)
  2xAMD EPYC 7742, 2× 64 cores, 2.25 GHz
  4× NVIDIA A100 GPU, 4× 40 GB HBM2e

The resources we were able to use on these computers were 1024 GPUs on MARCONI 100 (the maximum GPUs allowed to be used for a job) and 512 GPUs on JEAN-ZAY and JURECA.

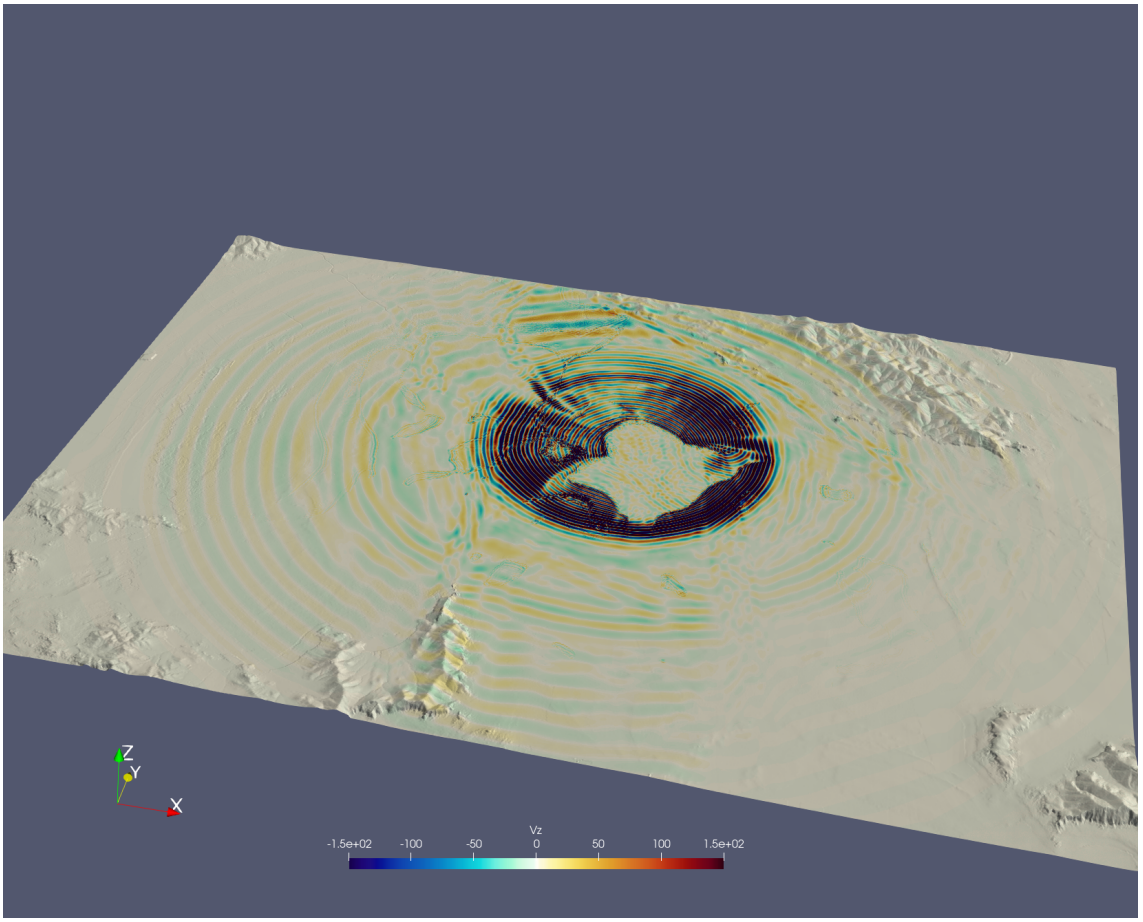


**Figure 3.4.** snapshot of wave propagation on bathymetry/topography after the shot.

### *3.3.2 Results*

We have done a first complete simulation with 2 million time steps which allows us to simulate 25 seconds of signal on 1024 V100 GPUs in 5 hours and 45 minutes. In this way we were able to model seismograms at different locations. We also developed a new feature that allows us to view snapshots (or movie) on the topography/bathymetry (Figure 3.4), since making a movie on the whole volume is too big because it would require several hundreds of TB to be stored on the disk.

We performed strong scaling tests with the 27M element mesh using 256, 512 and 1024 V100 GPUs on Marconi 100. Figure 3.5 shows the performance obtained. We see that we get 90% parallel efficiency by using 1024 GPUs compared to 256.

We performed weak scaling tests with 3 different meshes (table 3.2) on the Jean-Zay and JURECA machines. The average number of elements is 185000 with a load imbalance of +/- 2.5% due to the difficulty of decomposing such large meshes. Note that a speedup of 3 is

obtained by going from V100 to A100. This suggests that the complete run made for the strong scaling would take around 2h00 on 1024 A100.

**Table 3.2.** Weak scaling using V100 and A100.

| | Nb GPUs | Memory per GPU | Elapsed time (s.) V100 | Elapsed time (s.) A100 | Speedup V100/A100 | Number of element per GPU |
|---|---|---|---|---|---|---|
| **Mesh 2** | 512 | 25 GB | 2158 | 716 | 3 | 185000 +/- 2.5% |
| **Mesh 3** | 256 | 25 GB | 2054 | 698 | 2.9 | 185000 +/- 2.5% |
| **Mesh 4** | 128 | 25 GB | 2199 | 698 | 3.1 | 185000 +/- 2.5% |

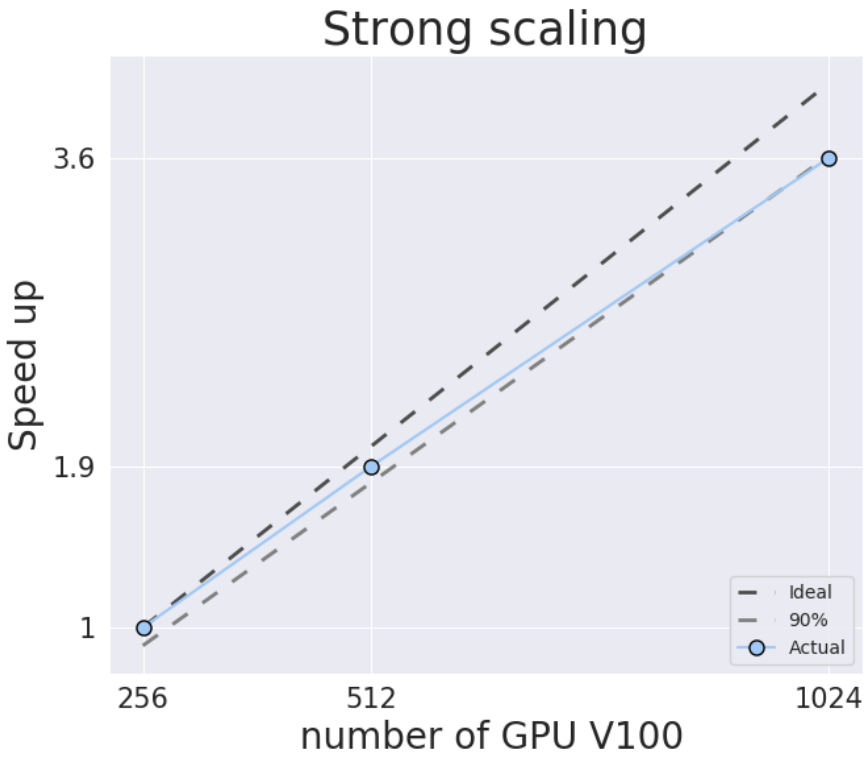


**Figure 3.5.** Strong scaling for the mesh 1 (Table 3.1) on Marconi 100.

### *3.3.3 Resolved challenges*

The main challenge was to implement a new mesher that correctly represents a coastal area that contains a complex topography and bathymetry as well as different sedimentary layers.

This mesher can handle several hundred million elements but we do not currently have an efficient parallel decomposer. We have developed a heuristic parallel decomposer, but it does not provide perfect load balancing. Indeed we have a load imbalance of about 2.5% which suggests a possible improvement of the strong scaling curves. Once the mesh was built and decomposed the code worked perfectly on the 3 machines used without having to make any specific modifications.

### *3.3.4 Deviations from the work plan*

The main difficulty we got is the access to pre-exascale machines, which was not possible during the project. The machines, on which we were able to compute, have usage rules that prevented us from accessing a large part of the resources. We were limited to 1024 V100 GPUs on MARCONI 100 and 512 GPUs on the other two machines used. This prevented us from testing larger meshes and thus testing the limits of SPECFEM3D cartesian as we had no difficulty running the code on the cases described and with the resources we were granted.

## 3.4 References and Publications

Favretto-Cristini N. et al. (2022). Assessment of risks induced by counter-mining unexploded large-charge historical ordnance in a shallow water environment : Part 1. Real case study. Part 2. Modeling of seismo-acoustic wave propagation. IEEE Journal of Oceanic Engineering. In Press.

# 4. SeisSol: 3D Fully Coupled Earthquake-Tsunami Simulation (Palu, Sulawesi, 2018 and Hellenic Arc) on SuperMUC-NG and Frontera

As part of pilot demonstrator PD4, “physics-based tsunami-earthquake interaction”, we developed simulation pipelines for physics-based modelling of complex earthquake tsunami events. Starting out from linked simulations of earthquakes (3D dynamic rupture simulations) that provide ocean floor displacement for subsequent tsunami simulations (2D shallow water models), we developed a novel 3D fully-coupled elastic-acoustic model that combines the entire process of tsunami generation in a single scenario: from the earthquake’s dynamics rupture process via seismic wave propagation and resulting displacement of the ocean floor up to acoustic wave propagation and displacement of the sea surface (incl. the onset of tsunami propagation). Fully coupled multi-physics modelling of earthquakes and tsunami has the potential to identify the key factors and physical links controlling earthquake and tsunami dynamics. By integrating a wide range of observations, from local to megathrust scales, such models bridge scientific disciplines and space-time scales.

As PD4 was not envisaged to become a “service” in the sense of WP5 in ChEESE, developing a grand challenge scenario for the exascale testbed demonstrator also served to establish a mature simulation workflow for coupled earthquake-tsunami events.

## 4.1 Goals of the demonstration run

Large tsunamis may be caused by various earthquake mechanisms. Yet, the specific dynamics of tsunamigenic earthquakes and tsunami generation are often unclear. The role of PD4 is to illuminate the entire process of earthquake tsunami generation in a self-consistent, fully coupled, physics-based approach relying on fewer modelling assumptions than commonly used. Only with a fully coupled solution can these approximations be explored in the complex geometries and material structures representative of real events.

In Sections 4.2 and 4.3 (and in Krenz et al., SC '21) we present 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. This tsunami was as surprising to scientists as it was devastating to communities in Sulawesi. It occurred near an active plate boundary, where earthquakes are common. Surprisingly, the earthquake caused a major tsunami, although it primarily offset the ground horizontally—normally, large-scale tsunamis are typically caused by vertical motions. The aim of the Palu, Sulawesi scenario was to test the scaling of the fully-coupled elastic-acoustic model developed in PD4, including our novel gravitational boundary condition.

In Section 4.4 we discuss our progress towards modelling the Hellenic Arc, Greece scenario that has been proposed in D3.4 as the fully-coupled 3-dimensional subduction pre-exascale earthquake-tsunami test case. Producing a mesh for a fully-coupled elastic-acoustic model has proven to be the key challenge to realise this extreme-scale scenario, due to small features in the topography close to the water layer, such as sharp coastal regions or small islands. These need to be resolved for a grand challenge simulation in this region. We finally also shortly report on the realisation of two smaller use-cases to illustrate that the workflows developed in PD4 are applicable in flexible tectonic settings, beyond the strike-slip Palu, Sulawesi testbed.

## 4.2 Scenario: 2018 Palu, Sulawesi Earthquake and Tsunami

Tsunamis occur due to abrupt perturbations to the water column, for example caused by the seafloor deforming during earthquakes. Devastating tsunamis associated with submarine *strike-slip* earthquakes are therefore rare: While such events may trigger landslides that in turn trigger tsunamis, the associated ground displacements are predominantly horizontal, not vertical, which does not favour tsunami genesis. However, strike-slip fault systems in complex tectonic regions, such as the Palu-Koro fault zone cutting across the island of Sulawesi, may host significant vertical deformation.

We investigated the magnitude Mw 7.5 Sulawesi earthquake and tsunami that occurred on September 28, 2018 across Palu Bay, Sulawesi, Indonesia. This combined event led to over two thousand fatalities. From a geophysical standpoint, this event is interesting for two main reasons: Firstly, the rupture propagated with supershear speed. Secondly, it sourced an unexpected tsunami that was localised in the shallow bay. The exact mechanism by which the tsunami was triggered is debated (P.M. Mai, 2019). The earthquake ruptured a 180 km long section of the Palu-Koro fault. It nucleated 70 km north of the city of Palu at shallow depths, with inferred hypocentral depths varying between 10 and 22 km. The rupture propagated predominantly southward, passing under Palu Bay and the city of Palu. The Palu earthquake triggered a local but powerful tsunami that devastated the coastal area of Palu Bay quickly after the earthquake. Inundation depths of over 6 m and run-up heights of over 9 m were recorded at specific locations.

Our coupled earthquake-tsunami simulations show that the time-dependent, 3D seafloor displacements are translated into bathymetry perturbations with a mean vertical offset of 1.5 m across the submarine fault segment. This sources a tsunami with wave amplitudes and periods that match those measured at the Pantoloan wave gauge and inundation that reproduces observations from field surveys. We conclude that a source related to earthquake displacements is probable and that landsliding may not have been the primary source of the

tsunami. This scenario is challenging because it exhibits a complicated fault network and a complex geometry. The inclusion of the shallow bathtub-like shaped bay with realistic bathymetry and topography turned out to be especially challenging.

The Palu, Sulawesi scenario makes use of our novel gravitational boundary condition. This boundary condition allows us to capture a time-dependent changing water height in the ocean without the need of using a moving mesh. We model the ocean and Earth with the acoustic and elastic wave equations respectively. The usage of the exact one-dimensional Riemann solver makes sure that we simulate the behaviour at element interfaces and especially at the elastic-acoustic interface with sufficient detail.

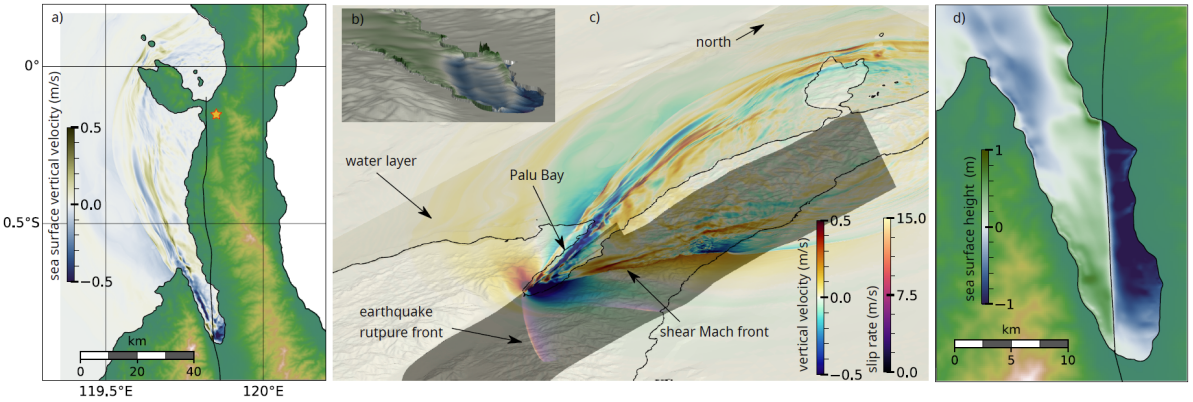


**Figure 4.1.** Summary of main results of the Palu-Sulawesi simulation: a) Map view of the vertical sea surface velocity (at 15 s simulation time, star indicates epicentre, black lines mark the fault); (b) sea surface height at 60 s (metres, same colormap as in d); c) 3D view of earthquake and tsunami model in Palu Bay (with slip rate, vertical sea-surface and Earth ground velocity at 15 s); d) sea surface height in Palu Bay at 15 s (vertical displacements). Image from Krenz et al., 2021.

## 4.3 Grand Challenge simulation and scalability of the Palu, Sulawesi scenario

For the coupled elastic-acoustic simulation of the Palu, Sulawesi scenario, we created two meshes, one medium-sized mesh with ~89M elements and a large mesh with ~518M elements. Using polynomials of order 5 (leading to 6th-order convergence) in SeisSol's ADER-DG discretisation, we obtain problem sizes of 46 billion (medium-sized mesh) and 261 billion (large mesh) degrees of freedom. The smaller mesh has a water layer resolution of 100m and resolves seismic waves with a maximum element size of 1000 m. The large mesh has a water layer resolution of 50 m and resolves the seismic waves with a resolution of 500 m. We used both meshes for production simulations. The smaller mesh was run for 100 s, which is enough to capture important parts of earthquake and dynamics dynamics. We ran the larger mesh for 30 s which captures the entirety of the earthquake dynamics and the

tsunamigenesis. The medium-sized mesh served to evaluate scalability in the extreme-strong-scalability regime. The large mesh was used in the Grand Challenge run.

### *4.3.1 Computational resources*

We ran our simulations primarily on the SuperMUC-NG machine at Leibniz Supercomputing Centre. SuperMUC-NG offers 6,336 nodes with dual socket Intel Skylake Xeon Platinum 8,174 CPUs (24 cores each). The nodes are organised in 8 islands. A fat tree OmniPath network is used, between each island the connection is pruned with a pruning factor of 1:4.

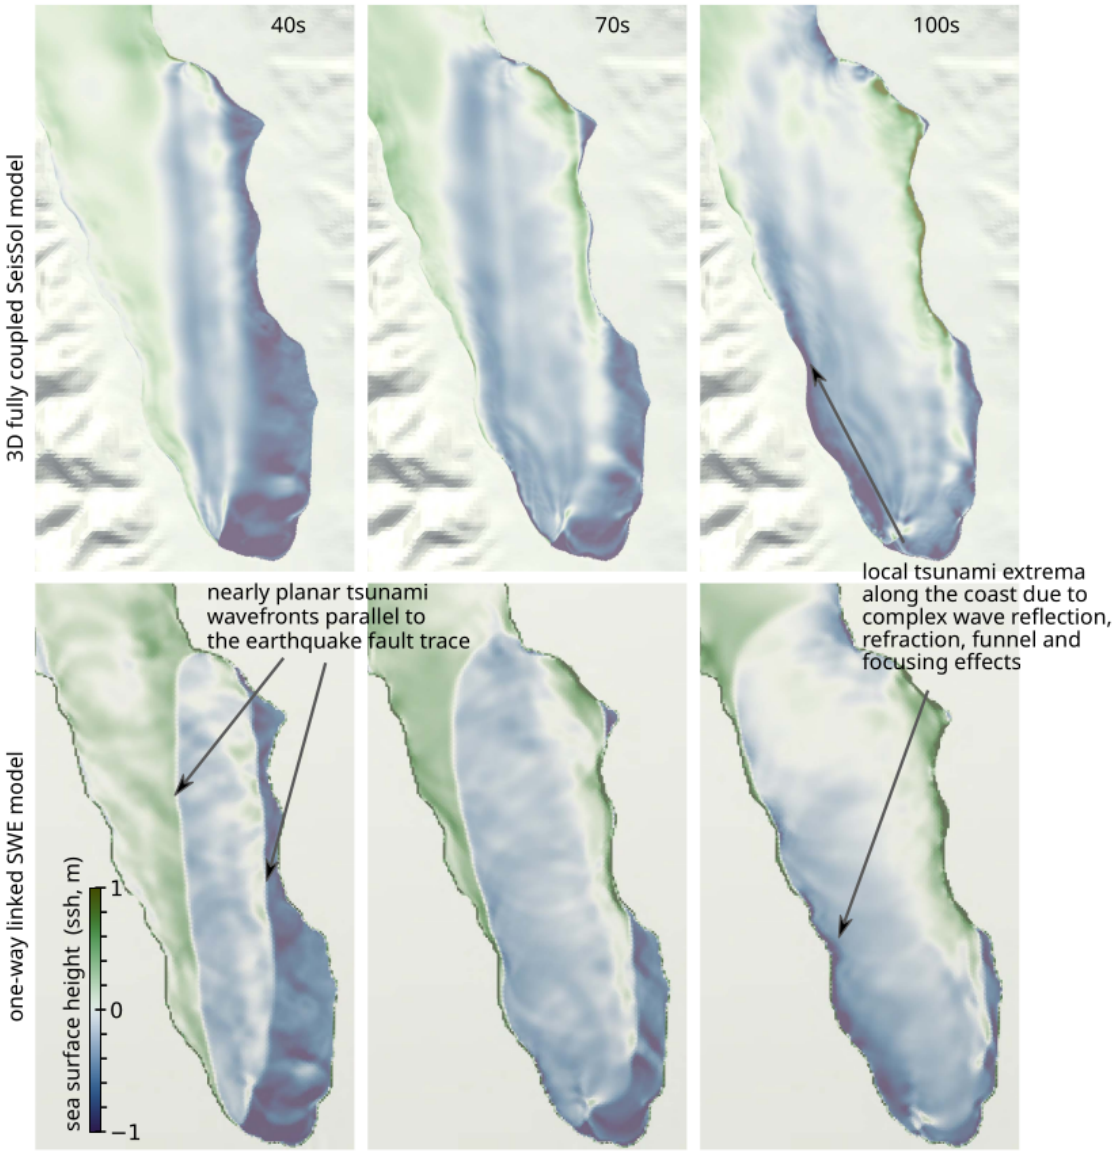


**Figure 4.2.** Results from Palu, Sulawesi Earthquake-Tsunami setup. The top row shows the water height obtained with our fully coupled simulation, the results below show the water height for a one-way linked tsunami simulation. Image from Krenz et al., 2021.

### *4.3.2 Results*

Figure 4.2 shows snapshots of the generated tsunami in the Palu Bay, as simulated by our Grand Challenge simulation. We compare the results to those obtained via a classical linked modelling approach, that simulates the earthquake separately and pipes the simulated ocean floor displacements into a subsequent 2D tsunami simulation. Both simulations manage to reproduce salient features of the Palu, Sulawesi 2018 tsunami. 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. 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 or due to conversion between surface gravity, acoustic, and seismic wave modes. Also, nonlinear effects may become important during wave interaction with the coastline.

### *4.3.3 Performance and scalability on SuperMUC-NG*

We ran a full production simulation of the first 30s of the earthquake and tsunami generation using the large mesh on 3,072 nodes of SuperMUC-NG. The simulation ran for 5h 30min and achieved 3.14 PFlops on average, i.e., a bit more than 1 TFlops per node. This simulation included receiver output (every 0.01 s) and free surface output (every 0.1 s). We also ran a performance test using the L mesh on 768 nodes (and using 2 MPI ranks per node), for which we achieved a performance of 1.3 TFlops per node (equivalent to ~78% parallel efficiency between these two results).

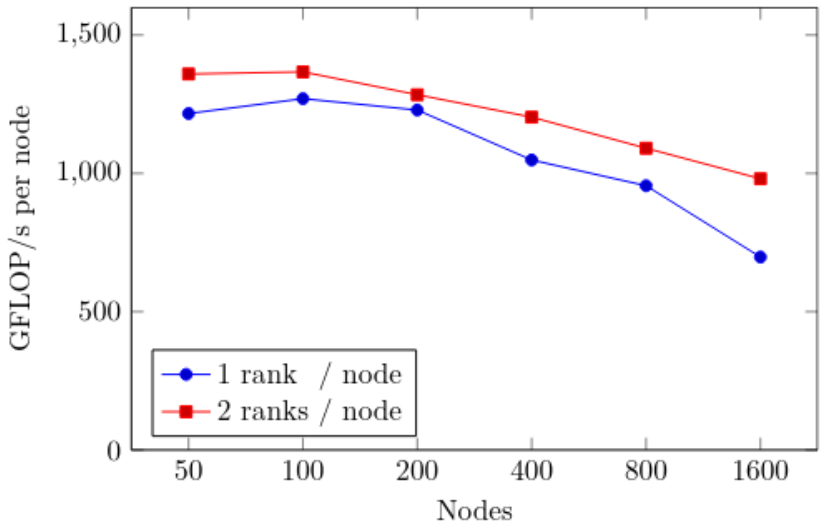


**Figure 4.3.** Strong scaling for the medium sized Palu setup with SeisSol.

Using the medium-sized mesh, we performed a full scaling study on SuperMUC-NG, starting from 50 nodes and using up to 1,600 nodes. Figure 4.1 shows the obtained scalability plot. We achieved 1,359 GFlops per node with 50 nodes and 981 GFLOPS on 1,600 nodes. This translates to a parallel efficiency of ≈72%.

### *4.3.4 Scalability and performance evaluation on Frontera*

In 2021, TUM and LMU submitted a proposal to become a *Leadership-Class Computing Facility Application Partner* of the Texas Advanced Computing Centre (TACC) with a proposal on performing multi-physics earthquake simulations on TACC's next-generation supercomputer (proposal was successful; the respective project will start in later 2022). As part of the proposal process, TACC performed a "deep dive" with SeisSol on its current petascale system Frontera. As one of the benchmarks for this deep dive, TACC used the Palu-Sulawesi scenario described in Sec. 4.2. and the two meshes (medium-sized and large) described above in Sec. 4.3. Table 4.1 lists important performance data (runtime, memory bandwidth, double-precision GFLOP/s per node) for running SeisSol on up to 7,000 nodes of Frontera. SeisSol achieved more than 9 PFlop/s on 7,000 nodes. The parallel efficiency from 250 to 7,000 nodes in this experiment was more than 98% for the large mesh.

**Table 4.1.** Performance data for SeisSol on Frontera for two different problem sizes (medium and large, see description in text).

| Case name | Nodes | Runtime[s] | MBW[GB/s] | DP[GF] |
|---|---|---|---|---|
| Seissol-medium | 100 | 1007.49 | 98.43 | 1499.96 |
| | 200 | 522.78 | 94.64 | 1443.48 |
| | 400 | 276.06 | 89.63 | 1374.49 |
| | 800 | 140.92 | 86.72 | 1348.75 |
| Seissol-large | 250 | 509.85 | 75.74 | 1329.13 |
| | 500 | 255.09 | 75.82 | 1332.90 |
| | 1000 | 109.57 | 87.90 | 1553.32 |
| | 2000 | 61.15 | 79.24 | 1416.10 |
| | 4000 | 33.47 | 71.84 | 1305.29 |
| | 6000 | 22.19 | 76.16 | 1403.36 |
| | 7000 | 18.45 | 71.82 | 1327.34 |

### *4.3.5 Performance on an NVIDIA DGX A100 node with 8 GPUs*

The results of 4.3.3 and 4.3.4 indicate that the PD4 simulation workflow is already well-prepared in terms of scalability to several thousands of compute nodes. On upcoming European pre-exascale supercomputers, such as Leonardo, the additional challenge will be to cope with heterogeneous nodes that provide multiple GPU accelerators per node. We therefore evaluated the performance of the elastic-acoustic model on an NVIDIA DGX A100 node with 8 GPUs (installed at CINECA), a platform well-suited to identify possible problems with node-level performance.

As the porting of some functionality to GPUs is still work in progress (incl. dynamic rupture and the gravity boundary conditions required for tsunami propagation), we evaluated the performance for a simpler elastic-acoustic modelling scenario motivated by assessing audible noise generated by induced seismicity (cmp. Figure 4.4).

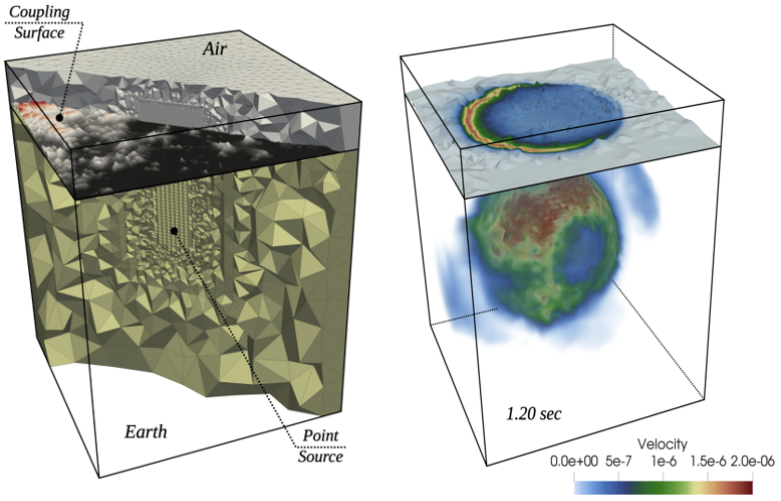


**Figure 4.4.** Coupled elastic-acoustic simulation of a point-source signal created in solid Earth and propagating into air.

We compare the elastic-acoustic performance with that for three scenarios for purely elastic wave propagation models studied in WP2: the LOH.1 benchmark and a production scenario for the Northridge 1994 earthquake. Figure 4.5 shows the performance obtained for three different meshes, two for LOH.1 and one for Northridge, each consisting of approx. 5 million grid elements, but having different choices of adaptive mesh refinement and thus different distribution of allowed time step sizes for elements. The results show a strong influence of the time step size distribution on scalability, which can be explained by

overheads caused by small time step clusters (see also the results and discussion in Deliverable 2.5).

For the elastic-acoustic scenario, despite its substantially higher complexity, we achieved 26.8 double-precision TFlop/s on the 8-GPU node.

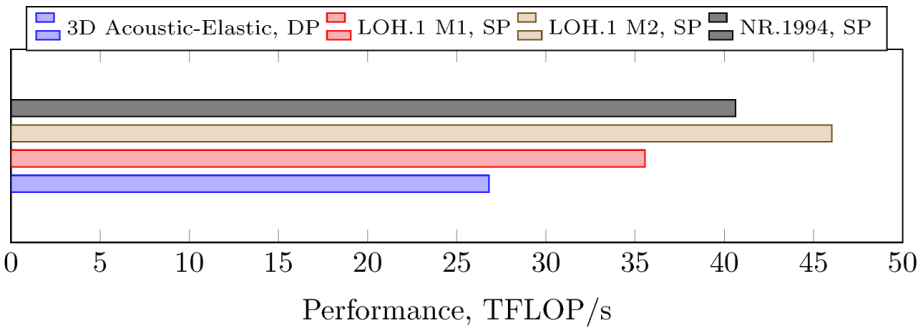


**Figure 4.5.** Performance comparison of the 3D elastic-acoustic scenario (blue) with three elastic wave equation scenarios (2x LOH.1, Northridge) on a node with 8 Nvidia A100 GPUs. Note that the elastic-acoustic scenario used double precision.

## 4.4 Prototyped Grand Challenge scenario: Hellenic Arc subduction zone

In D3.4 we defined the Hellenic Arc, Greece, as the envisioned fully-coupled 3-dimensional subduction pre-exascale earthquake-tsunami testbed scenario. Besides a high-resolution run, we envisaged modelling several multi-physics dynamic rupture earthquake scenarios with varying hypocenter locations at the subducting slab, varying rupture speeds, and varying resulting magnitudes. The key challenge to realise this extreme-scale scenario turned out to be not so much the scalability, but rather the modelling of the respective event - in particular, producing a mesh for a fully-coupled elastic-acoustic model. In this section, we therefore detail on the model development for the Hellenic Arc subduction zone and discuss challenges for mesh generation. We also shortly report on the realisation of two smaller use-cases to illustrate that the developed workflows are already mature enough to tackle state-of-the-art simulation tasks.

### *4.4.1 Modelling of earthquake-tsunami events at the Hellenic Arc subduction zone*

In the Mediterranean, the Aegean Sea Plate, which is part of the Eurasian plate, and the African plate converge, which leads to the large subduction area of the Hellenic Arc. The subducting plate itself extends over an area of 860km x 625km x 90km, including a complex slab geometry. Our SeisSol model uses the slab geometry from the European Database of Seismogenic Faults (EDSF, Basili et al., 2013) (see Figure 4.6). The resulting dynamic rupture model covers great parts of the Eastern Mediterranean Sea. The bathymetry and

topography of the area are highly resolved, and the initial conditions (stresses and strengths) are constrained on the subduction scale, following Ulrich et al. (2022). Optimal stress parameters were extracted from the static analysis (see Figure 4.6), where we varied $SH_{max}$ (azimuth of maximum horizontal compressive stress), as well as the strike, dip and rake of an optimally oriented fault. For a fault strike of 280°, a dip between 10° and 20° and a rake of 85°, we obtain the highest relative prestress ratio R (the fault is closest to failure and optimally prestressed).

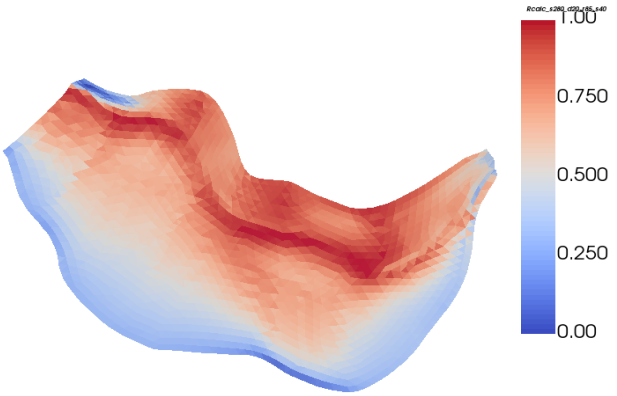


**Figure 4.6.** Static Analysis result for strike = 280°, dip = 20°and rake = 85°. Red coloured parts of the fault are close to failure, while blue coloured parts are non-optimally prestressed.

The Hellenic Arc has a seismogenic zone from 15km depth down to ~40km depth (Scala et al., 2020). However, most of its convergence is accommodated aseismically (Papazachos et al., 2000). We consider this by adopting the static and dynamic friction coefficients accordingly. From 0km to 15km the dynamic friction coefficient $\mu_d$ is 1.2 and greater than the static friction coefficient ($\mu_s$ = 0.6), a zone where we allow for shallow creep. The seismogenic zone is the locked part of the fault ($\mu_s$ = 0.6, $\mu_d$= 0.2), ranging from 15km to 40km depth. From 40km to 50km we assign a slip neutral region ($\mu_s = \mu_d$ = 0.6), followed by a slip-strengthening region below 50km depth ($\mu_s$ = 0.6, $\mu_d$ = 1.2).

These parameters were combined in a worst-case scenario for the Hellenic Arc, featuring a $M_w$ 9.01 earthquake. The hypocenter is located on the eastern Hellenic Arc in 30km depth. During the simulation, the rupture propagates from east to west (see Figure 4.7), until reaching the trench, producing fault slip of up to 14m (see Figure 4.8). Also, the shallow portion of the fault (zone where we allow for shallow creep) is dynamically ruptured. The highest peak slip is observed at the steep-dipping portion of the slab. This worst-case Hellenic Arc scenario produces an uplift of up to 2 metres.

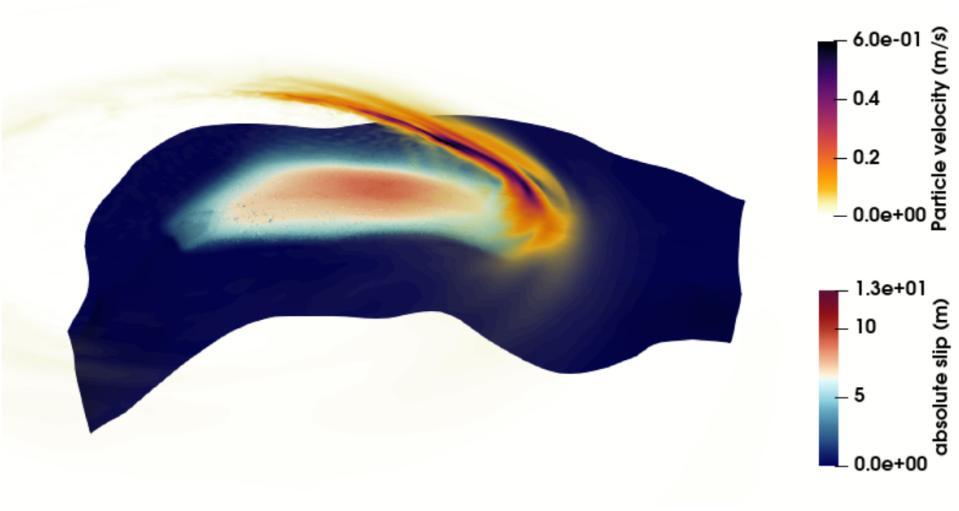


**Figure 4.7.** Particle velocity and fault slip during the rupture propagation (after 20s) for the worst-case Hellenic Arc dynamic rupture scenario (front view). The rupture propagates from East to West.

We vary the hypocenter location along the arc and observe drastic changes in shallow slip, together with an increase of spontaneously modelled magnitude ($M_w$ = 9.26). We observe the further West the hypocenter, the greater the slip amplification on the shallow fault (see Figure 4.9).

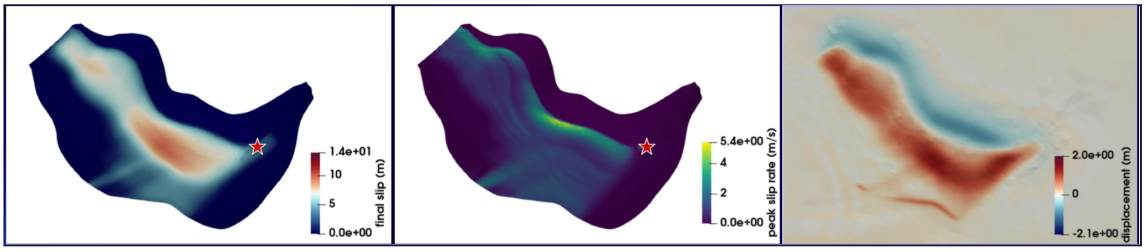


**Figure 4.8.** Worst-case scenario of Mw 9.01 for the Hellenic Arc. Left: Accumulated fault slip in metres. Middle: Peak slip rate in m/s. Right: Final seafloor displacement in metres, caused by the earthquake.

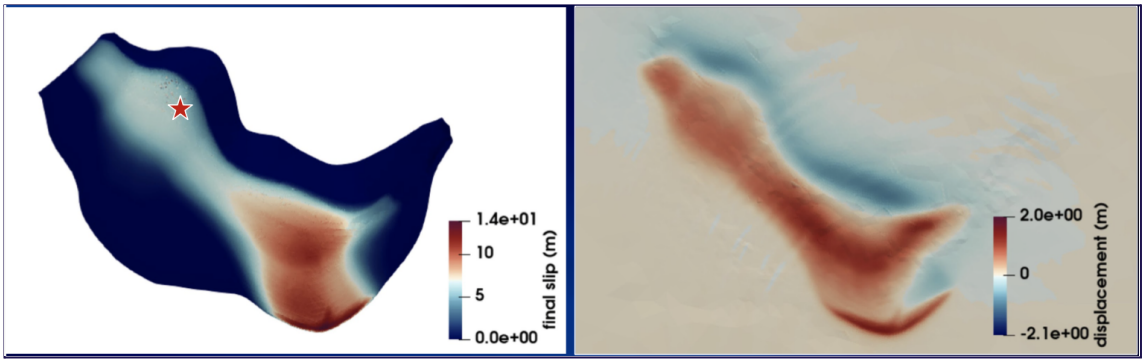


**Figure 4.9.** Left: Accumulated fault slip in metres; Right: final seafloor displacement for a hypocenter location on the eastern Hellenic Arc. Note the shallow slip amplification on the eastern fault.

In the next step, we allow for off-fault plastic yielding. We distinguish between 1) partially consolidated and 2) weaker (less consolidated) sediments, following Ulrich et al. (2022) . In case 1) the maximum final slip is 19,98m which leads to a maximum seafloor displacement of 1.60m ($M_w$ = 9,36). For case 2) we observe a maximum seafloor displacement of 8.30m ($M_w$ = 9,26, see Figure 4.10).

We can use the dynamic rupture models to test the impact of gradual introduction of modelling complexity and compare the slip distributions and possibly seismic effects on tsunami generation in the Mediterranean. Especially shallow fault slip and off-fault plastic yielding may enhance the seafloor uplift drastically and are important parameters for dynamic rupture scenarios that contribute to tsunami studies.

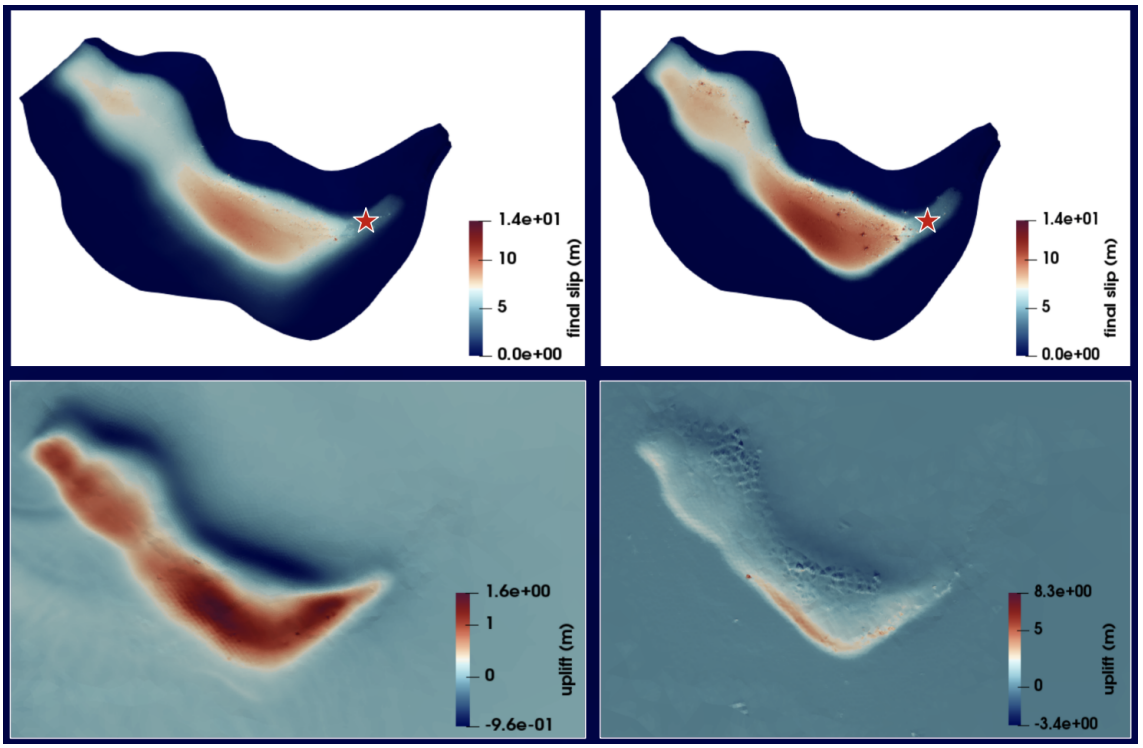


**Figure 4.10.** Accumulated fault slip in meter (top) and final seafloor displacement (bottom) for plastic dynamic rupture models. Left: Model including off-fault plasticity with partially consolidated sediments. Right: Model including off-fault plasticity with "weaker" sediments. Note the difference in scale for the seafloor displacement.

We also generate a smaller ($M_w$ = 8.53) earthquake, based on the "worst-case" scenario. Here, the fault is not optimally prestressed anymore, such that the rupture stops before reaching the trench. The fault geometry limits the rupture and no more shallow slip is observed (see Figure 4.11).

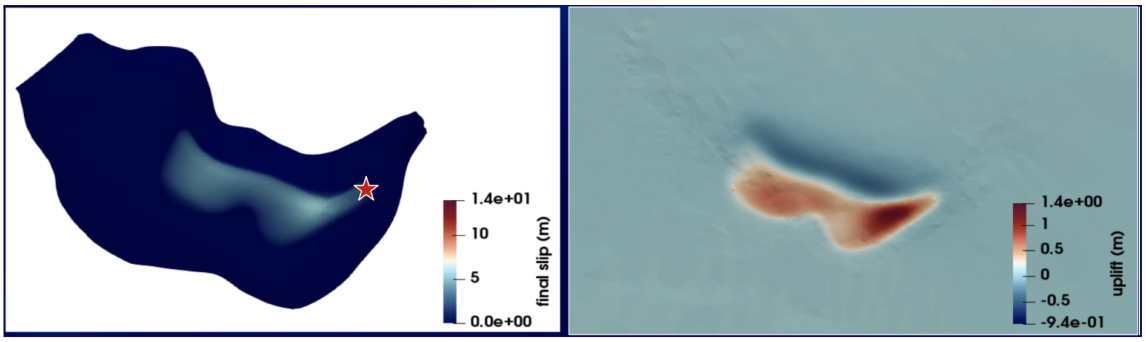


**Figure 4.11.** Accumulated fault slip (left) and final seafloor displacement (right) for a $M_w$ 8.53 earthquake with a hypocenter on the eastern Hellenic Arc.

All these are plausible tsunami-genic earthquake scenarios, which can contribute to common tsunami studies, as they reveal differences in shallow fault slip, rupture directivity and final seafloor displacement.

### *4.4.2 Meshing challenges*

To generate high-quality unstructured tetrahedral meshes for our scenarios, we use SimModeler (Simmetrix, http://www.simmetrix.com/index.php), which is suited to approximate complex 3D model geometries and rapid model generation. The workflow of mesh generation for SeisSol is described in detail on https://seissol.readthedocs.io/en/latest/simmodelerCAD-workflow.html. For a fully-coupled model, this workflow has to be adapted accordingly. Here, a waterlayer has to be created with Gmsh (https://gmsh.info/) and intersected with the topography before proceeding with the normal meshing workflow. The generated mesh surfaces of waterlayer and topography must have similar mesh sizes to be intersected without problems. Otherwise, the software is not able to intersect and combine them accordingly. Additionally, the mesh space between two nodes needs to be as small as possible to sufficiently resolve the topography and the intersecting parts. On the other hand, a finer mesh on the single model parts slows down the GUI of SimModeler drastically.

Challenging for the intersection of topography and waterlayer are especially small features in the topography close to the water layer, such as sharp coastal regions, small islands, and topography that lies slightly above or below 0 km. In the intersection process, these features can not be united in the GUI if they are too small, due to the limiting 2D mesh size of each surface. To avoid these problems, we use a preprocessing step that smoothens the topography around 0 km depth, before applying the intersection. Nonetheless, an extensive and complex region such as the Hellenic Arc includes many of such small-scale features:

tiny islands and parts with a shallow seafloor, such that the intersection of waterlayer and topography is difficult. Figure 4.12 shows the successful intersection for this region. Currently, we struggle in adding sides and a bottom to the model, as the topography and waterlayer run parallel, due to a shallow coastal water depth. The intersection process produces a general error in SimModeler, when trying to generate a whole modelling domain. The generality of the error and the size of the modelling domain makes it difficult to solve the issue straight forward. To find the geometry that causes the error is like searching for "the needle in a haystack".

For smaller and less complex regions, as for the Samos-Izmir earthquake or the Husavík Flatey Fault zone, we successfully created the modelling domain, including the intersection between waterlayer and topography (see section 4.3.3).

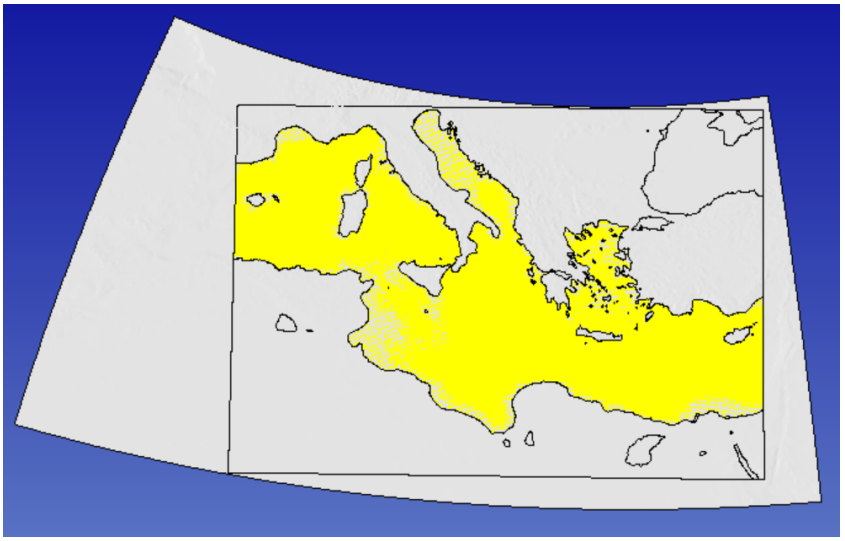

**Figure 4.12.** Intersection of a waterlayer (squared patch) with the topography, including great parts of the Mediterranean Sea. The yellow region represents the water.

### *4.4.3 Addendum: towards workflows for fully coupled scenarios*

As PD4 was not envisaged to become a "service" in ChEESE' WP5, preparing the Hellenic Arc grand challenge scenario was the main driver for increasing the maturity of PD4's simulation workflow for fully-coupled earthquake-tsunami modelling. Hence, even though these do not directly address exascale computing, we present two further fully-coupled (but smaller) scenarios, that demonstrate the usability of the simulation workflow on two ChEESE-related regions: the Samos-Izmir earthquake ($M_w$ 7, October 2020) and the Husavík Flatey Fault in Northern Iceland ($M_w$ 7.34), thus creating strong links to PDs 1 and 5, respectively.

### Samos-Izmir Scenario

We contributed to the PD1 live demonstration exercise (see Deliverable D5.4) and presented a fully-coupled model for the Samos-Izmir earthquake. For this, we adapted the kinematic north-dipping earthquake model of 147 sub-faults provided by Ryo Okuwaki (Heidarzadeh et al., 2021). The model includes a geometric bathymetry resolution of 400m and ran with 256 CPU hours. The displacement discontinuity of the model can be accurately represented in SeisSol's discontinuous finite element space. The model shows a dynamically complex peak ground velocity (see Figures 4.13 and 4.14).

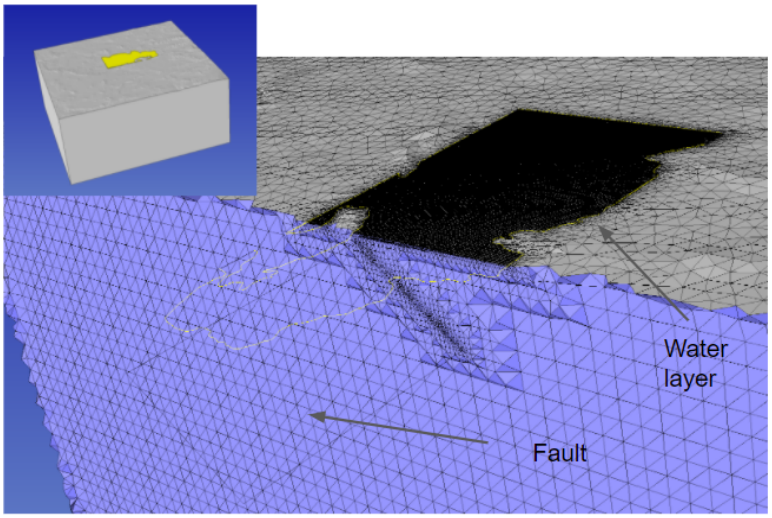


**Figure 4.13.** Waterlayer mesh for the Samos-Izmir fully-coupled kinematic model. The water layer covers the area right above the fault (small figure: yellow patch). Note the smaller mesh size in the water layer and around the fault.

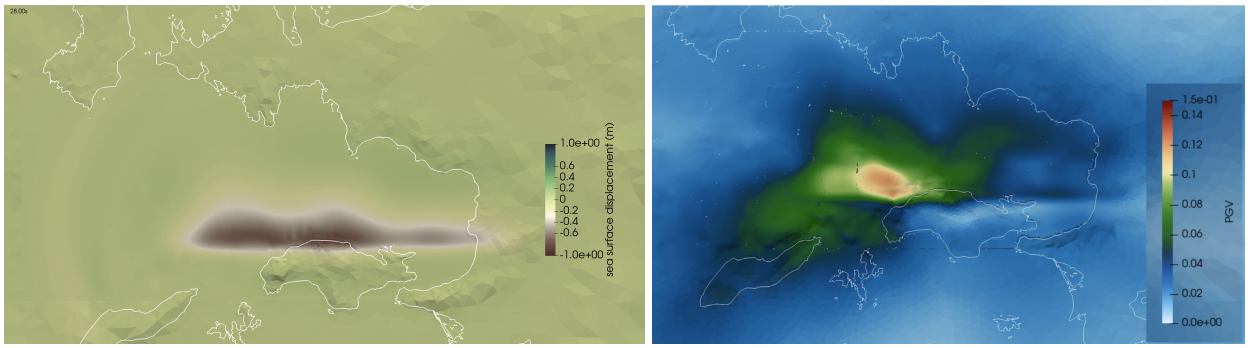


**Figure 4.14.** Left: Sea surface displacement (in metre) caused by the kinematic model of the Samos-Izmir earthquake. Right: Peak ground velocity resulting from the kinematic Samos-Izmir earthquake simulation.

### PD5 Húsavík Flatey Fault Scenario

Picking up the PSHA from PD5, we also tackled a fully coupled 3-dimensional earthquake-tsunami scenario on the Húsavík-Flatey Fault (HFF) in Iceland (see also Sec. 5.2). The knowledge on tsunami genesis and propagation is essential for hazard forecast and potential warning in Icelandic coastal regions in case of a tsunami risk.

The earthquake hypocenter is located on the on-shore portion of the HFF (Figure 4.15), such that the rupture propagates in north-eastern direction on the complex strike-slip right lateral fault system, resulting in a $M_w$ 7.34 earthquake with one metre uplift. We include complex bathymetry and topography, as well as the 3D subsurface structure by Abril et al. (2021). Primary stress orientations and the stress shape ratio are taken from Ziegler et al. (2016). We follow the one-way linked dynamic earthquake rupture and shallow water equations tsunami workflow (Madden et al., 2021), where the dynamic rupture output (seafloor displacement) is used to source a tsunami in the shallow-water solver sam(oa)$^2$ (sam(oa)$^2$, 2020).
For comparison, the same earthquake scenario was used for the fully-coupled earthquake-tsunami approach (see Figure 4.16). Here, the water layer covers the part right above the fault system. The fault and the waterlayer are meshed with a 200m mesh spacing. The maximum grid size of the domain is 5km, such that the final 3D-mesh contains ~4 million elements.

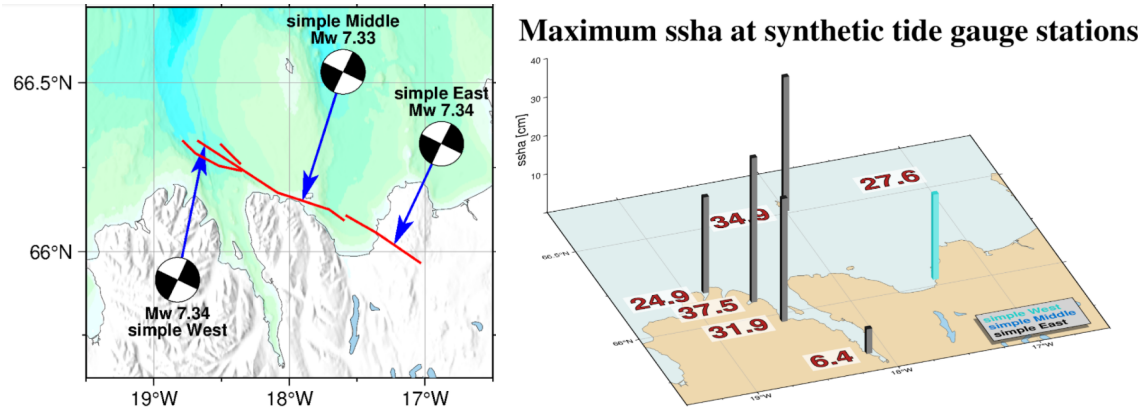


**Figure 4.15. Left:** Fault system of the Húsavík-Flatey Fault, with moment tensors of simulated earthquake rupture scenarios. The scenario "simple East" is used for the fully-coupled rupture-tsunami simulation. **Right:** Maximum sea surface height (ssha) for specific points along the Icelandic coast, as a result of one-way linking of the time-dependent seafloor uplift (generated in SeisSol) to a shallow water tsunami model (modelling with sam(oa)$^2$). Grey bars indicate the sea surface height produced by the "simple East" scenario.

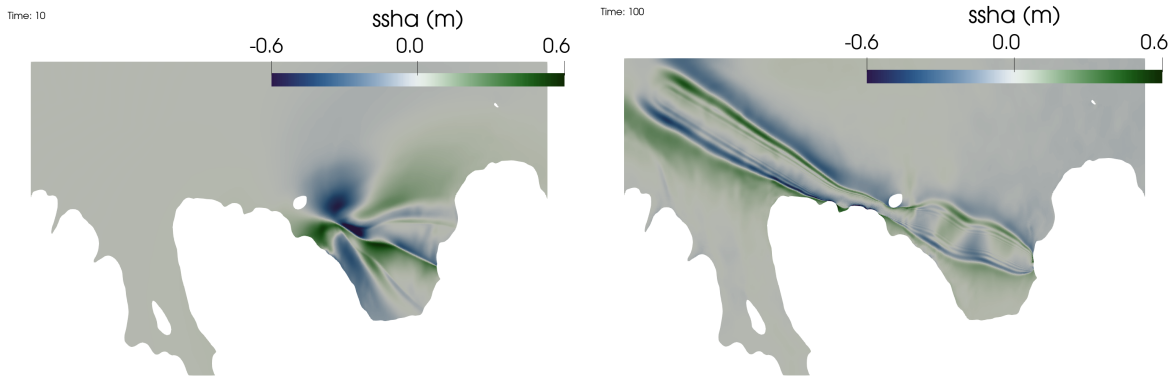


**Figure 4.16.** Tsunami initiation and propagation process of the fully-coupled earthquake-tsunami model for the "simple East" event of Mw 7.34 on the on-shore portion of the HFF. **Left:** Sea surface height after 10s simulation time. **Right:** Sea surface height after 100s simulation time.

## 4.5 References and Publications

Krenz, Lukas; Uphoff, Carsten; Ulrich, Thomas; Gabriel, Alice-Agnes; Abrahams, Lauren S.; Dunham, Eric M.; Bader, Michael: *3D acoustic-elastic coupling with gravity: the dynamics of the 2018 Palu, Sulawesi earthquake and tsunami.* Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, ACM, 2021.

Wirp, S.A., Gabriel, A.-A., Schmeller, M., Madden, E.H., van Zelst, I., Krenz, L., van Dinther, Y. and Rannabauer, L.: *3D Linked Subduction, Dynamic Rupture, Tsunami, and Inundation Modeling: Dynamic Effects of Supershear and Tsunami Earthquakes, Hypocenter Location, and Shallow Fault Slip.* Frontiers in Earth Science, 9, 177, 2021.

# 5. ExaHyPE: Dynamic Rupture Simulation on the Húsavík-Flatey Fault on SuperMUC-NG

The exascale testbed scenario for ExaHyPE had to be changed substantially due to unforeseen problems that manifested in the earthquake simulation approach (particularly, modelling of dynamic rupture) and in difficult-to-overcome performance and scalability limitations of ExaHyPE, but also in unexpected "shortage" in personnel. See section 5.4 for further discussion.

We therefore split our testbed scenario into two aspects: evaluation of the general scalability of MUQ+ExaHyPE and realisation of UQ-enabling dynamic rupture scenarios. We concentrate on the latter in this report. For a large-scale evaluation of MUQ+ExaHyPE for an inverse UQ problem in tsunami simulation (using 2D shallow water models instead of 3D seismic simulation), we refer to the publication by Seelinger et al. (2021), which was presented at the SC21 Supercomputing Conference.

## 5.1 Goals of the demonstration run

The goal for the scenario described in Section 5.2 was to represent a single simulation run in a forthcoming UQ scenario, having originally in mind the Southern Iceland Seismic Zone (SISZ), which requires fully-automatic and efficient generation of the curvilinear mappings to approximate topography and bookshelf fault geometry. The combination of using dynamic rupture on curvilinear-represented faults and topography with PML boundary and using high order discretisation is a genuinely novel contribution of the ChEESE project. For the validation and for comparison of results for complicated fault geometries, we therefore chose to switch to the Húsavík-Flatey fault (in Northern Iceland) as the setup, which is the scenario for the Pilot Demonstrator 5 (physics-based probabilistic seismic hazard analysis) in ChEESE and thus offered the best data and allowed comparison to results obtained with SeisSol.

## 5.2 Dynamic rupture simulation on the Húsavík-Flatey Fault

In Iceland, the Eurasian Plate and the North American Plate diverge with an average rate of 21.8 mm per year (Demets et al., 2010), producing complex fault systems on- and off-shore. The Húsavík-Flatey Fault, an approx. 90 km long, right lateral strike-slip fault extending from the Tjörnes peninsula, poses a high seismic risk to the north Icelandic town Húsavík with potential magnitudes of up to Mw 6.8 (Metzger and Jónsson, 2014). We simulated an earthquake caused on the main branch of the Húsavík-Flatey fault. The setup picks up the fault model, topography data and material properties developed for PD5.

### *5.2.1 Challenges for scalability*

To be useful for MUQ scenarios, individual simulation runs should execute on a medium-size (tens to few hundreds) number of MPI ranks and lead to execution times in the range of hours. It turned out that the smallest setup that could accurately simulate dynamic rupture for the Husavik-Flatey scenario required 731 compute nodes of SuperMUC-NG and barely fit into the 48-hours wall-time limit for single jobs enforced on SuperMUC-NG (cmp. Sec. 5.3.1). A mixture of reasons showed to be responsible for this unexpectedly high resource requirements:

- To avoid unphysical reflections at domain boundaries, either a highly accurate treatment of waves near the boundary (using Perfectly Matched Layers, PML; Duru et al., 2020) or a large extension of the simulation domain (in combination with adaptive coarsening of the mesh away from the region of interest) is necessary. As ExaHyPE's underlying mesh management layer Peano featured problems with aggressive coarsening over multiple AMR layers, we opted for making simulation domains as small as possible and thus using PML.
- Usage of PML is only required in a layer near the boundary. However, as ExaHyPE (actually as almost all PDE environments) still enforces a uniform set of quantities per grid element, the additional degrees of freedom for PML need to be stored in all grid elements, even though PML computations are restricted to boundary layers. PML increases the degrees of freedom from 9 (elastic wave equation) to 36 unknowns per integration point (e.g., leading to 18,432 degrees of freedom per element at discretisation order 7), which is further enlarged by local storage required to evaluate the ADER-DG scheme. Using curvilinear meshes leads to 10 further persistent degrees of freedom needed for each integration point such that ExaSeis with PML becomes a strongly memory-dominant code.
- For dynamic load balancing, ExaHyPE relies on reactively offloading lightweight tasks from overloaded to idle ranks, which proved to be ineffective for PML and thus prevented simulations with strong adaptive mesh refinement (and coarsening).
- In addition, the large working set led to problems with available memory, such that the most effective parallelisation was identified to be using 1 MPI rank per compute node and largely skip shared-memory parallelism (cmp. Sec. 5.3).

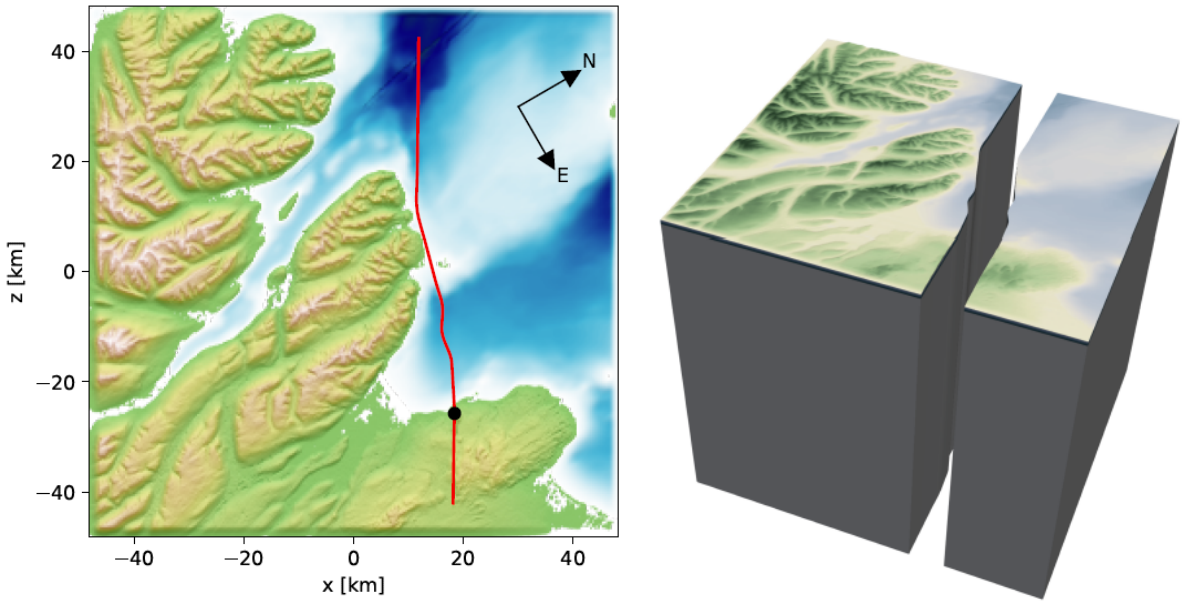


**Figure 5.1.** Setup for the Húsavík-Flatey scenario. Left: topography data and fault structure (red line); right: curvilinear mesh aligning with the fault. The setup is rotated to match the overall orientation of the fault (images from Rannabauer, 2022).

## 5.3 Execution on SuperMUC-NG

Figure 5.1 illustrates the topography, fault structure and resulting curvilinear mesh setup for the scenario. We used a setup that can be modelled by the main fault without branches, reflecting the fault line by a smoothed curve.

We rotate the cuboid computational domain (96.8 km in each direction) by 60 degrees to simplify alignment of the fault with the Cartesian coordinates of the curvilinear mesh (cmp. also Figure 4.15). Thus, the fault separates the computational mesh into two sections. We use a uniform grid with $241^3$ elements and discretisation order 3 for the discontinuous Galerkin method, which results in a subcell resolution of 0.1 km between the degrees of freedom. The chosen resolution led to a time step size of 0.7 ms. We simulated the earthquake and seismic wave propagation over a duration of 16 s, after which the rupture has propagated over the entire fault.

### *5.3.1 Computational resources*

The simulation setup leads to a grid with roughly 14 million elements. At the chosen discretisation order, we thus obtain 896 million integration points, which corresponds to using approx. 8 billion degrees of freedom for the seismic wave equation. However, we require additional storage to also enable PML and to store material parameters and

Jacobians for the curvilinear mesh. The setup therefore requires storage for over 46 billion variables.

The simulation was executed on 731 compute nodes of SuperMUC-NG. Problems with the memory consumption forced us to place only 1 MPI rank per node, even though OpenMP scalability stops to be satisfactory already beyond using 4 cores for using PML. For the pure seismic wave equation, the recommended MPI+TBB setup would be to use at least 6 MPI ranks per node, such that satisfactory performance is currently not achievable. Due to the resulting performance problems, the simulation required nearly 48 hours to finish, which makes UQ scenarios infeasible, until these problems have been solved.

### *5.3.2 Results*

Figure 5.2 illustrates the rupture dynamics obtained for the simulation. We compared the results obtained with ExaSeis with simulations that use SeisSol with the same physical parameters, showing excellent agreement of the two codes.

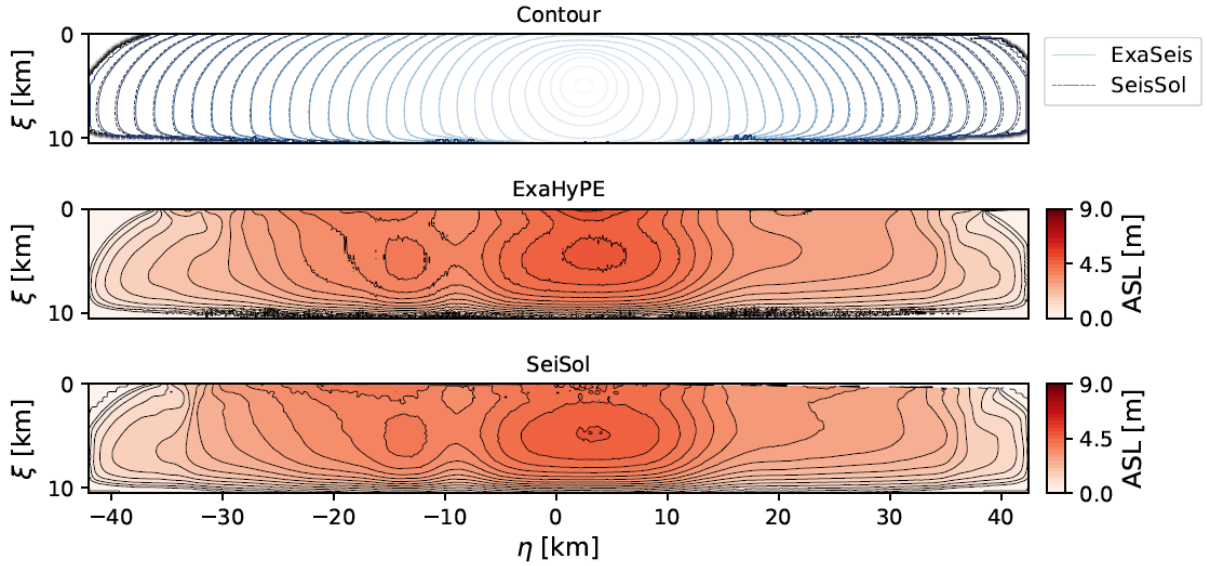


**Figure 5.2.** Comparison of results obtained with ExaSeis and SeisSol for the Húsavík-Flatey scenario. We compare the rupture conture (top) and the absolute slip rates (ASP) obtained with ExaSeis (middle) and SeisSol (bottom), respectively. Near the surface, small discrepancies are observed, which can be explained by differing representations of geometry. (Plots from Rannabauer, 2022)

## 5.4 Deviations from the work plan and resolved challenges

At the time of planning the exascale testbed scenario (Deliverable 3.4, April 2020), we had three researchers working on ExaSeis scenarios: Leonhard Rannabauer (TUM, partly funded via ChEESE; implementation of models, curvilinear meshing), Duo Li (LMU, partly funded via ChEESE; realisation of dynamic rupture scenarios) and Anne Reinarz (TUM, funded via additional resources; uncertainty quantification and respective workflows). In August 2020,

Anne Reinarz became assistant professor at Durham University. We continued our collaboration on uncertainty quantification, but naturally with a reduced extent compared to her time at TUM. In addition, Duo Li went on parental leave from January 2021 until December 2021 and returned working 50% part-time. From the several scenario candidates envisaged in Deliverable 3.4, we therefore had to early-on extract one scenario that we considered most interesting in terms of scientific outcome. We decided to concentrate on dynamic rupture for the two hazard areas envisaged in ChEESE - the Southern Iceland Seismic Zone, SISZ, and the Húsavík-Flatey fault system in Northern Iceland (the latter chosen for this deliverable due to the more complicated fault).

For the exascale testbed scenario, we had envisaged a large-scale UQ forward simulation that studies sensitivity of earthquake hazards to variations in topography, fault geometry and similar parameters. Both Icelandic hazard regions pose two major challenges on earthquake modelling:

- Modelling complex fault systems, which may be curved and/or rough, and for which exact location and orientation is at least to some extent uncertain - multiple, parallel faults for SISZ vs. branched faults for Húsavík-Flatey
- capturing the topography.

For UQ setups, this requires a flexible description of the topography and fault geometry, which can then be parameterised in the UQ workflow. The curvilinear mesh approach was realised in ExaSeis to address these challenges and is a unique feature for such UQ workflows.

During ChEESE, we carefully validated ExaHyPE's novel dynamic rupture implementation on curvilinear grids against several community benchmarks (TPV) and in comparison to other dynamic rupture codes (SeisSol, WaveQLab). The results revealed a strong influence of the grid resolution at the fault on the obtained rupture behaviour, which we could not mitigate via adaptive mesh refinement. Hence, using a multi-resolution approach with MUQ (as in the shallow water demonstrator, Seelinger et al., 2021) was not feasible. The high model resolution necessary for accurate dynamic rupture, however, makes running large UQ workloads unfeasible on current machines. We therefore opted for a two-step approach: the general applicability and feasibility of ExaHyPE within large-scale UQ workflows with MUQ was evaluated in collaboration with Anne Reinarz and Linus Seelinger (Univ. Heidelberg) in a setup that uses a 2D shallow water model in ExaHyPE (Seelinger et al., 2021). The ExaSeis setup was developed in parallel, primarily by Leonhard Rannabauer (and Duo Li, as far as possible besides her parental leave). For the demonstration run reported here, we opted for the Húsavík-Flatey scenario, as its complicated fault geometry

makes the scenario more interesting for a single-run demonstration, and because the results could be directly compared with the PSHA setups simulated with SeisSol.

## 5.5 References and Publications

**Related Publications:**

K. Duru; L. Rannabauer; A.-A. Gabriel; G. Kreiss; M. Bader: *A stable discontinuous Galerkin method for the perfectly matched layer for elastodynamics in first order form.* Numerische Mathematik 146, 2020.

K. Duru; L. Rannabauer; A.-A. Gabriel; O. K. A. Ling; H. Igel; M. Bader:
*A stable discontinuous Galerkin method for linear elastodynamics in 3D geometrically complex elastic solids using physics based numerical fluxes.* Computer Methods in Applied Mechanics and Engineering 389, 2022, 114386.

K. Duru, L. Rannabauer, A.-A. Gabriel, and H. Igel (2021): *A new discontinuous Galerkin spectral element method for elastic waves with physically motivated numerical fluxes*, Journal of Scientific Computing, 88(51), doi:10.1007/s10915-021-01565-1.

L. Rannabauer: *Earthquake and tsunami simulation with high-order ADER-DG methods.* PhD thesis, submitted to the Department of Informatics, Technical University of Munich, in January 2022.

L. Seelinger; A. Reinarz; L. Rannabauer; M. Bader; P. Bastian; R. Scheichl:
*High Performance Uncertainty Quantification with Parallelized Multilevel Markov Chain Monte Carlo.* Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC21).

**Further References:**

Abril, C., Tryggvason, A., Gudmundsson, Ó., & Steffen, R. (2021). Local Earthquake Tomography in the Tjörnes Fracture Zone (North Iceland). Journal of Geophysical Research: Solid Earth, 126(6), e2020JB020212.

Basili, R., Kastelic, V., Demircioglu, M. B., Garcia Moreno, D., Nemser, E. S., & Petricca, P., et al. (2013). The European Database of Seismogenic Faults (EDSF) compiled in the framework of the Project SHARE. URL http://diss.rm.ingv.it/share-edsf/.

DeMets, C., Gordon, R. G., Argus, D. F., & Stein, S. (1990). Current plate motions. Geophysical journal international, 101(2), 425-478.

Heidarzadeh, M., Pranantyo, I. R., Okuwaki, R., Dogan, G. G., & Yalciner, A. C. (2021). Long tsunami oscillations following the 30 October 2020 Mw 7.0 Aegean Sea earthquake: observations and modelling. Pure and Applied Geophysics, 178(5), 1531-1548.

Madden, E. H., Bader, M., Behrens, J., van Dinther, Y., Gabriel, A. A., Rannabauer, L., Ulrich, T., Uphoff, C., Vater, S., & van Zelst, I. (2021). Linked 3-D modelling of megathrust earthquake-tsunami events: from subduction to tsunami run up. Geophysical Journal International, 224(1), 487-516.

Metzger, S., & Jónsson, S. (2014). Plate boundary deformation in North Iceland during 1992–2009 revealed by InSAR time-series analysis and GPS. Tectonophysics, 634, 127-138.

Papazachos, B. C., Karakostas, V. G., Papazachos, C. B., & Scordilis, E. M. (2000). The geometry of the Wadati–Benioff zone and lithospheric kinematics in the Hellenic arc. Tectonophysics, 319(4), 275-300.

sam(oa)$^2$, 2020: GitLab, URL https://gitlab.lrz.de/samoa.

Scala, A., Lorito, S., Romano, F., Murphy, S., Selva, J., Basili, R., Babeyko, A., Herrero, A., Hoechner, A., Løvholt, F., Maesano, F.E., Perfetti, P., Tiberti, M.M., Tonini, R., Velope, M., Davies, G., Festa, G., Power, W., Piatanesi, A., & Cirella, A. (2020). Effect of shallow slip amplification uncertainty on probabilistic tsunami hazard analysis in subduction zones: use of long-term balanced stochastic slip models. Pure and Applied Geophysics, 177(3), 1497-1520.

Ulrich, T., Gabriel, A. A., & Madden, E. H. (2022). Stress, rigidity and sediment strength control megathrust earthquake and tsunami dynamics. Nature Geoscience, 1-7.

Ziegler, M., Rajabi, M., Heidbach, O., Hersir, G. P., Ágústsson, K., Árnadóttir, S., & Zang, A. (2016). The stress pattern of Iceland. Tectonophysics, 674, 101-113.