LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

SeisSol - High-Performance Computing for Computational Seismology and Earthquake Physics

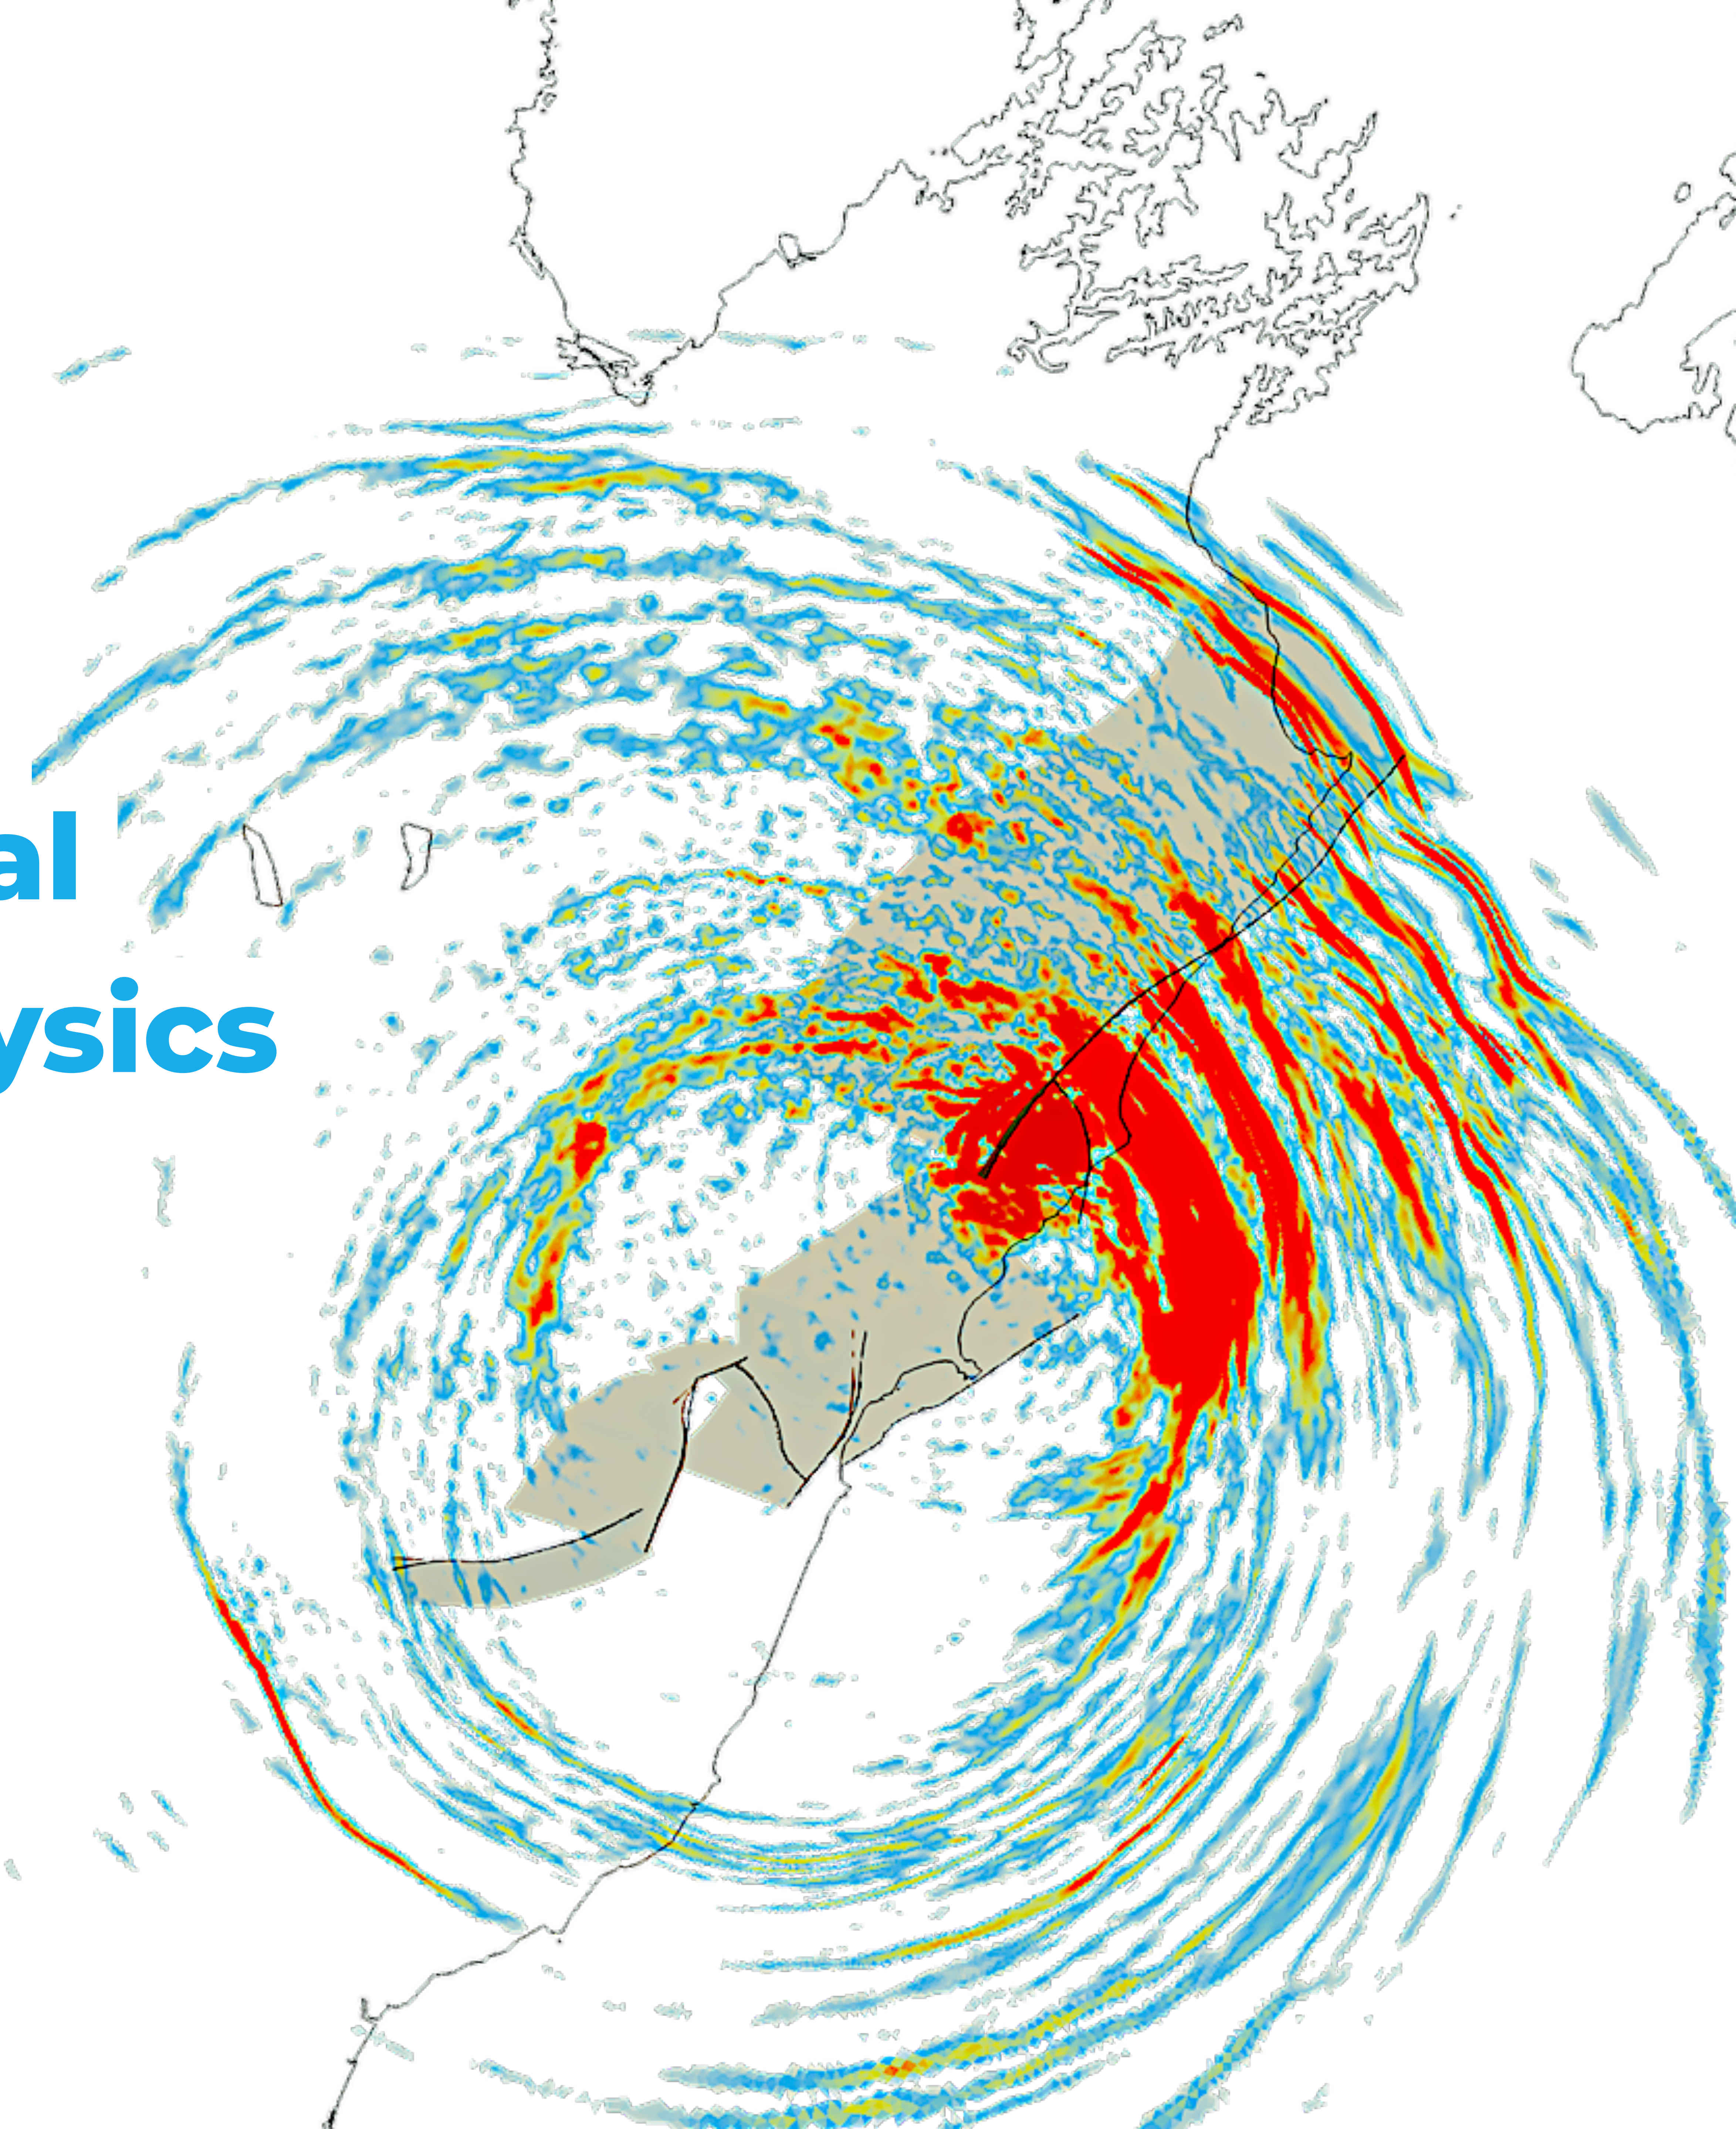

Alice-Agnes Gabriel

gabriel@geophysik.uni-muenchen.de
bit.ly/AAG8LMU
twitter @InSeismoland

ChEESE
Center of Excellence for Exascale in Solid Earth

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

In preparation for this afternoon, please:

—> visit

https://github.com/SeisSol/Training

—> install Docker

https://docs.docker.com/engine/install/

—> install Paraview

https://www.paraview.org/download/

—> run

```
$ docker pull uphoffc/seissol-training
```

ChEESE

Center of Excellence for Exascale in Solid Earth

The SeisSol team

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

TUM TECHNISCHE UNIVERSITÄT MÜNCHEN

lrz Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften

Carsten Uphoff

Duo Li

Taufiqurrahman

Thomas Ulrich

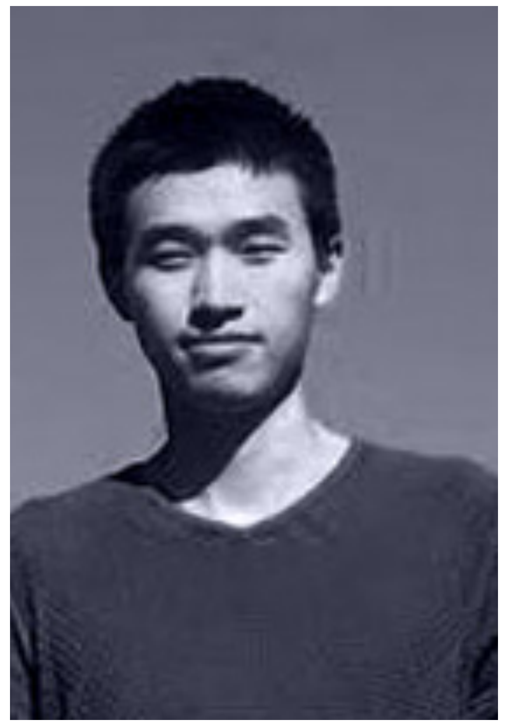

Bo Li

Aniko Wirp

ChEESE

Center of Excellence for Exascale in Solid Earth

Gefördert durch

DFG

erc

European Research Council

Established by the European Commission

www.tear-erc.eu

Michael Bader

Sebastian Wolf

Lukas Krenz

Ravil Dorozhinksii

SuperMUC-NG

Computational seismology

- A pioneering field for **HPC** to image Earth's interior, understand the dynamics of the mantle, track down energy resources
- Seismology is **data-rich** and can often be treated as a **linear** (hyperbolic PDE) system
- **Key activities:** Calculation of synthetic seismograms in 3D Earth and solving seismic inverse problems
- **Common approach:** time-domain solutions of space-dependent seismic wavefield solved by domain decomposition
- **On-going challenges:** 3D Earth structure, computational efficiency, complex geological subsurface
- **Need for open community solutions**

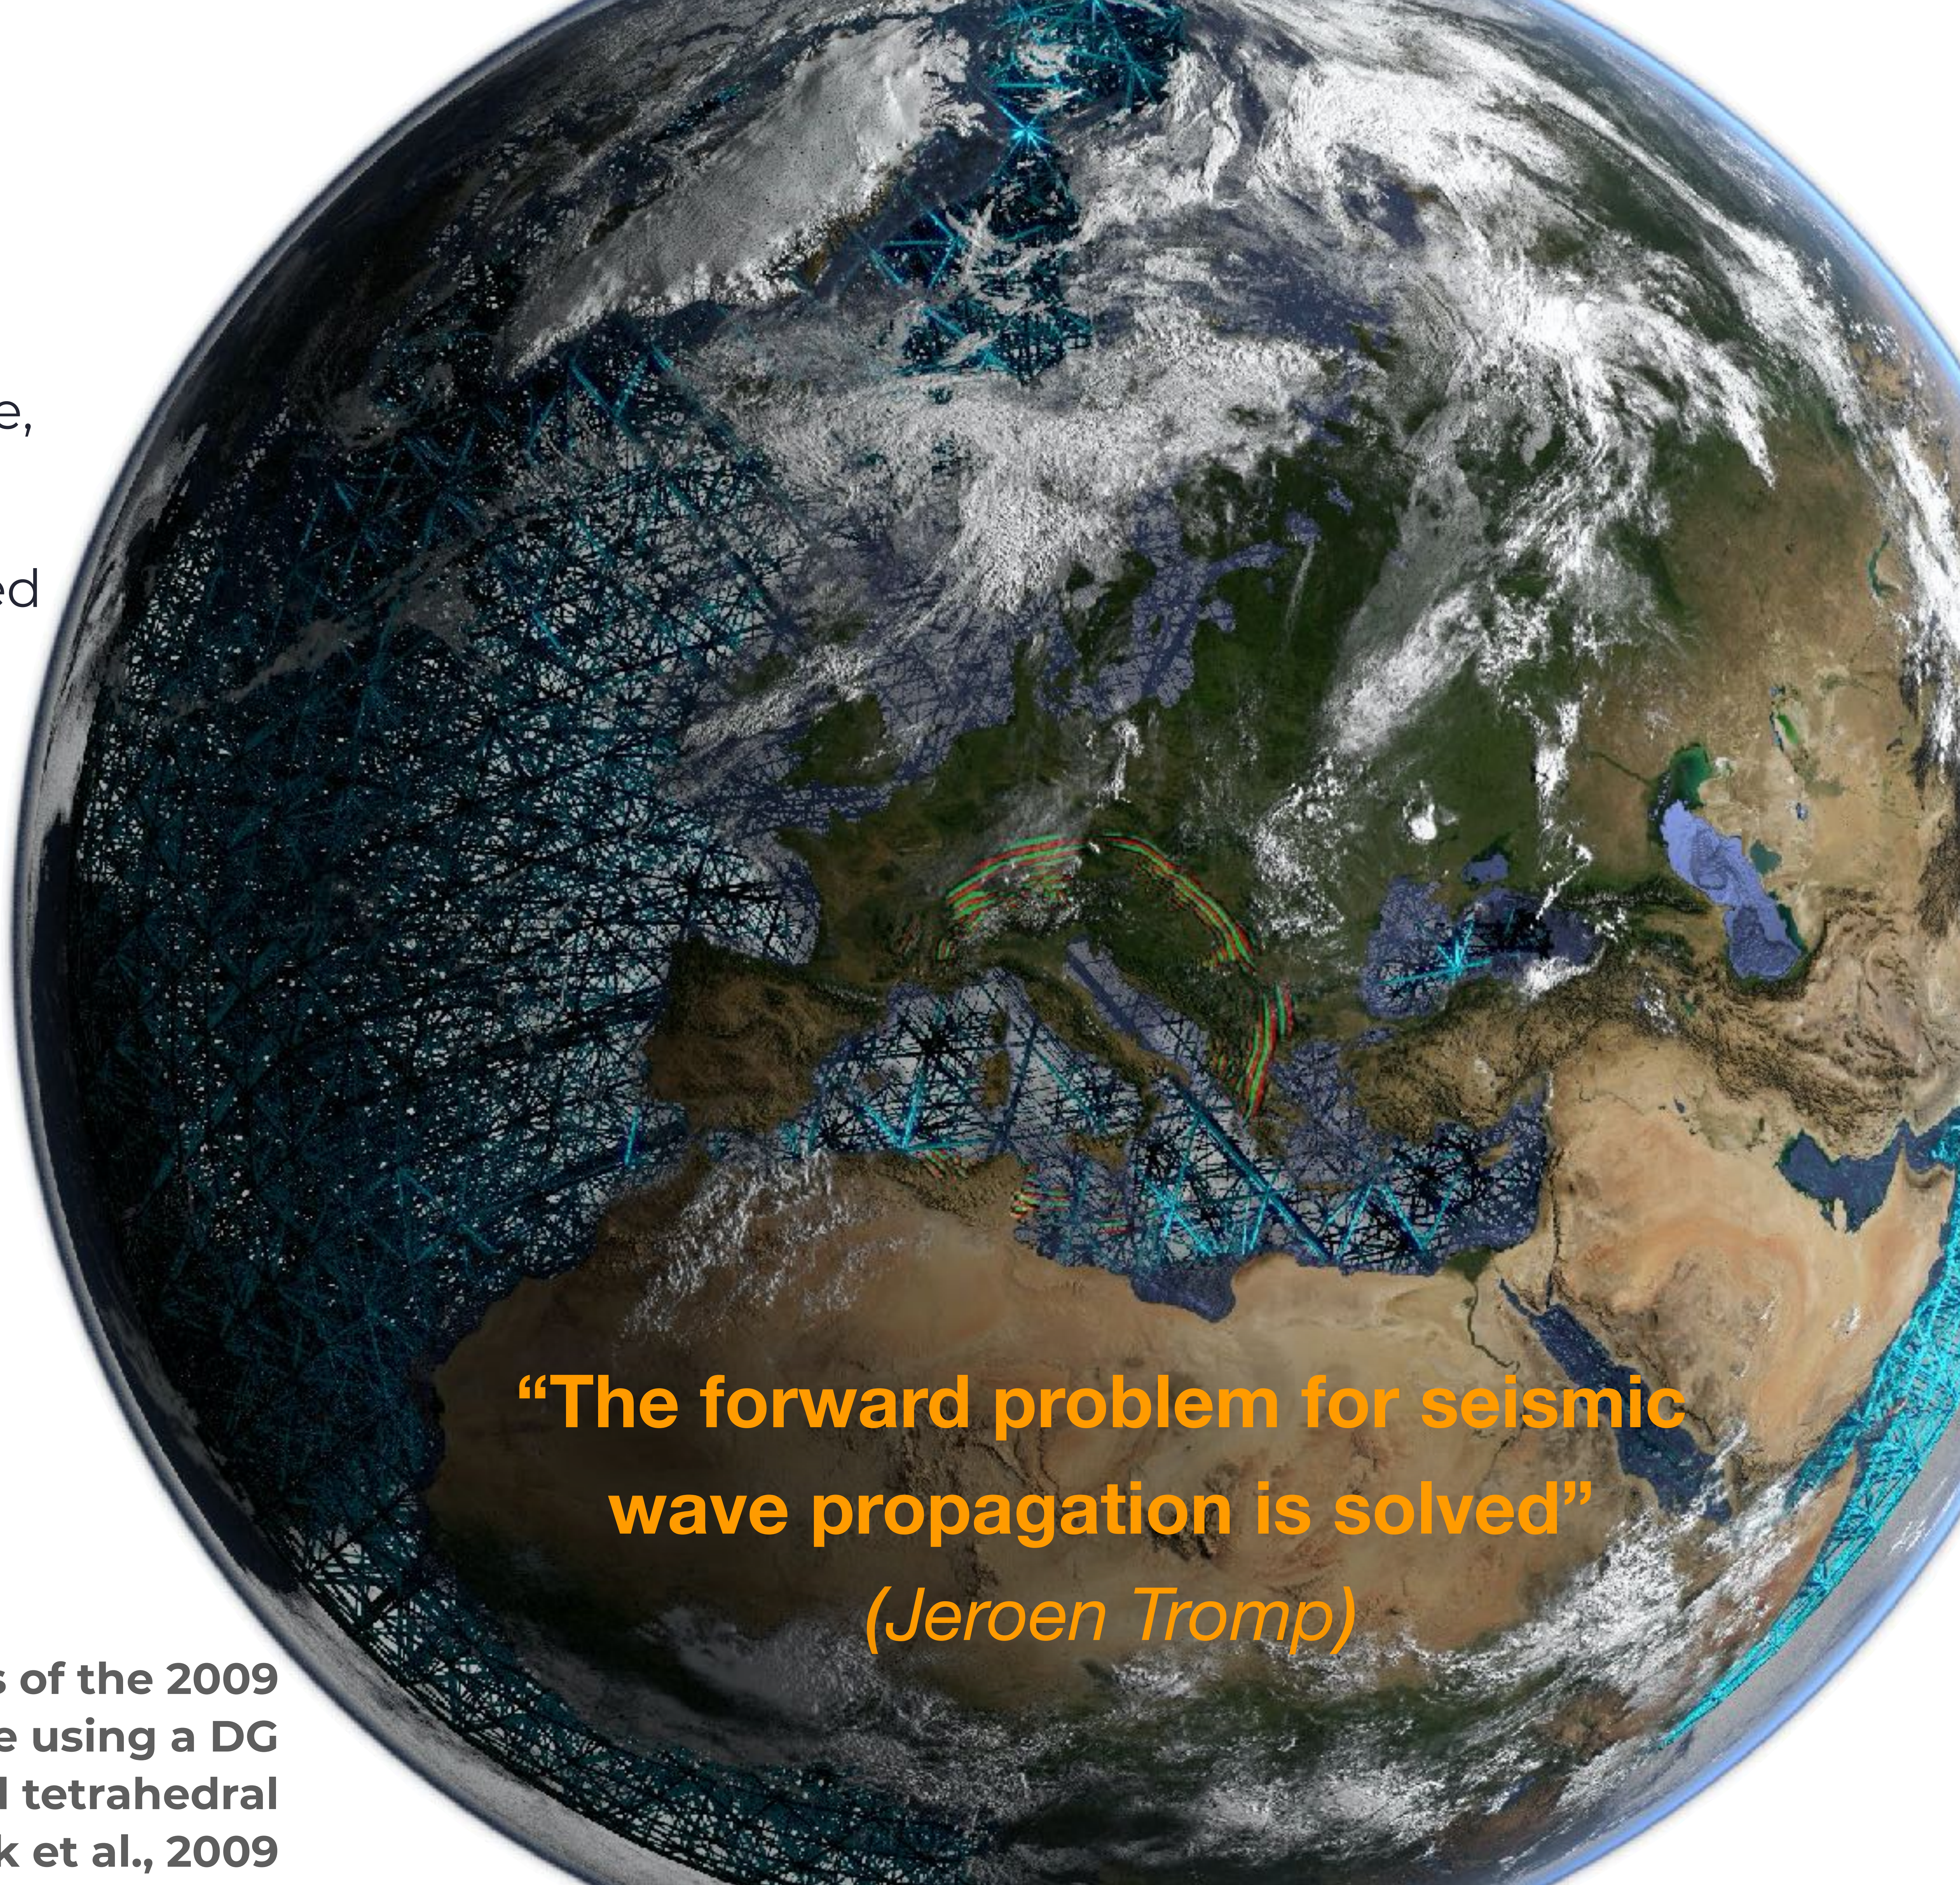

Wave simulations of the 2009 L'Aquila earthquake using a DG method on unstructured tetrahedral meshes, Wenk et al., 2009

ChEESE
Center of Excellence for Exascale in Solid Earth

Computational seismology

- A pioneering field for **HPC** to image Earth's interior, understand the dynamics of the mantle, track down energy resources
- Seismology is **data-rich** and can often be treated as a **linear** (hyperbolic PDE) system
- **Key activities:** Calculation of synthetic seismograms in 3D Earth and solving seismic inverse problems
- **Common approach:** time-domain solutions of space-dependent seismic wavefield solved by domain decomposition
- **On-going challenges:** 3D Earth structure, computational efficiency, complex geological subsurface
- **Need for open community solutions**
- Applications, e.g. from earthquake seismology and engineering, require **multi-scale and multi-physics** capabilities

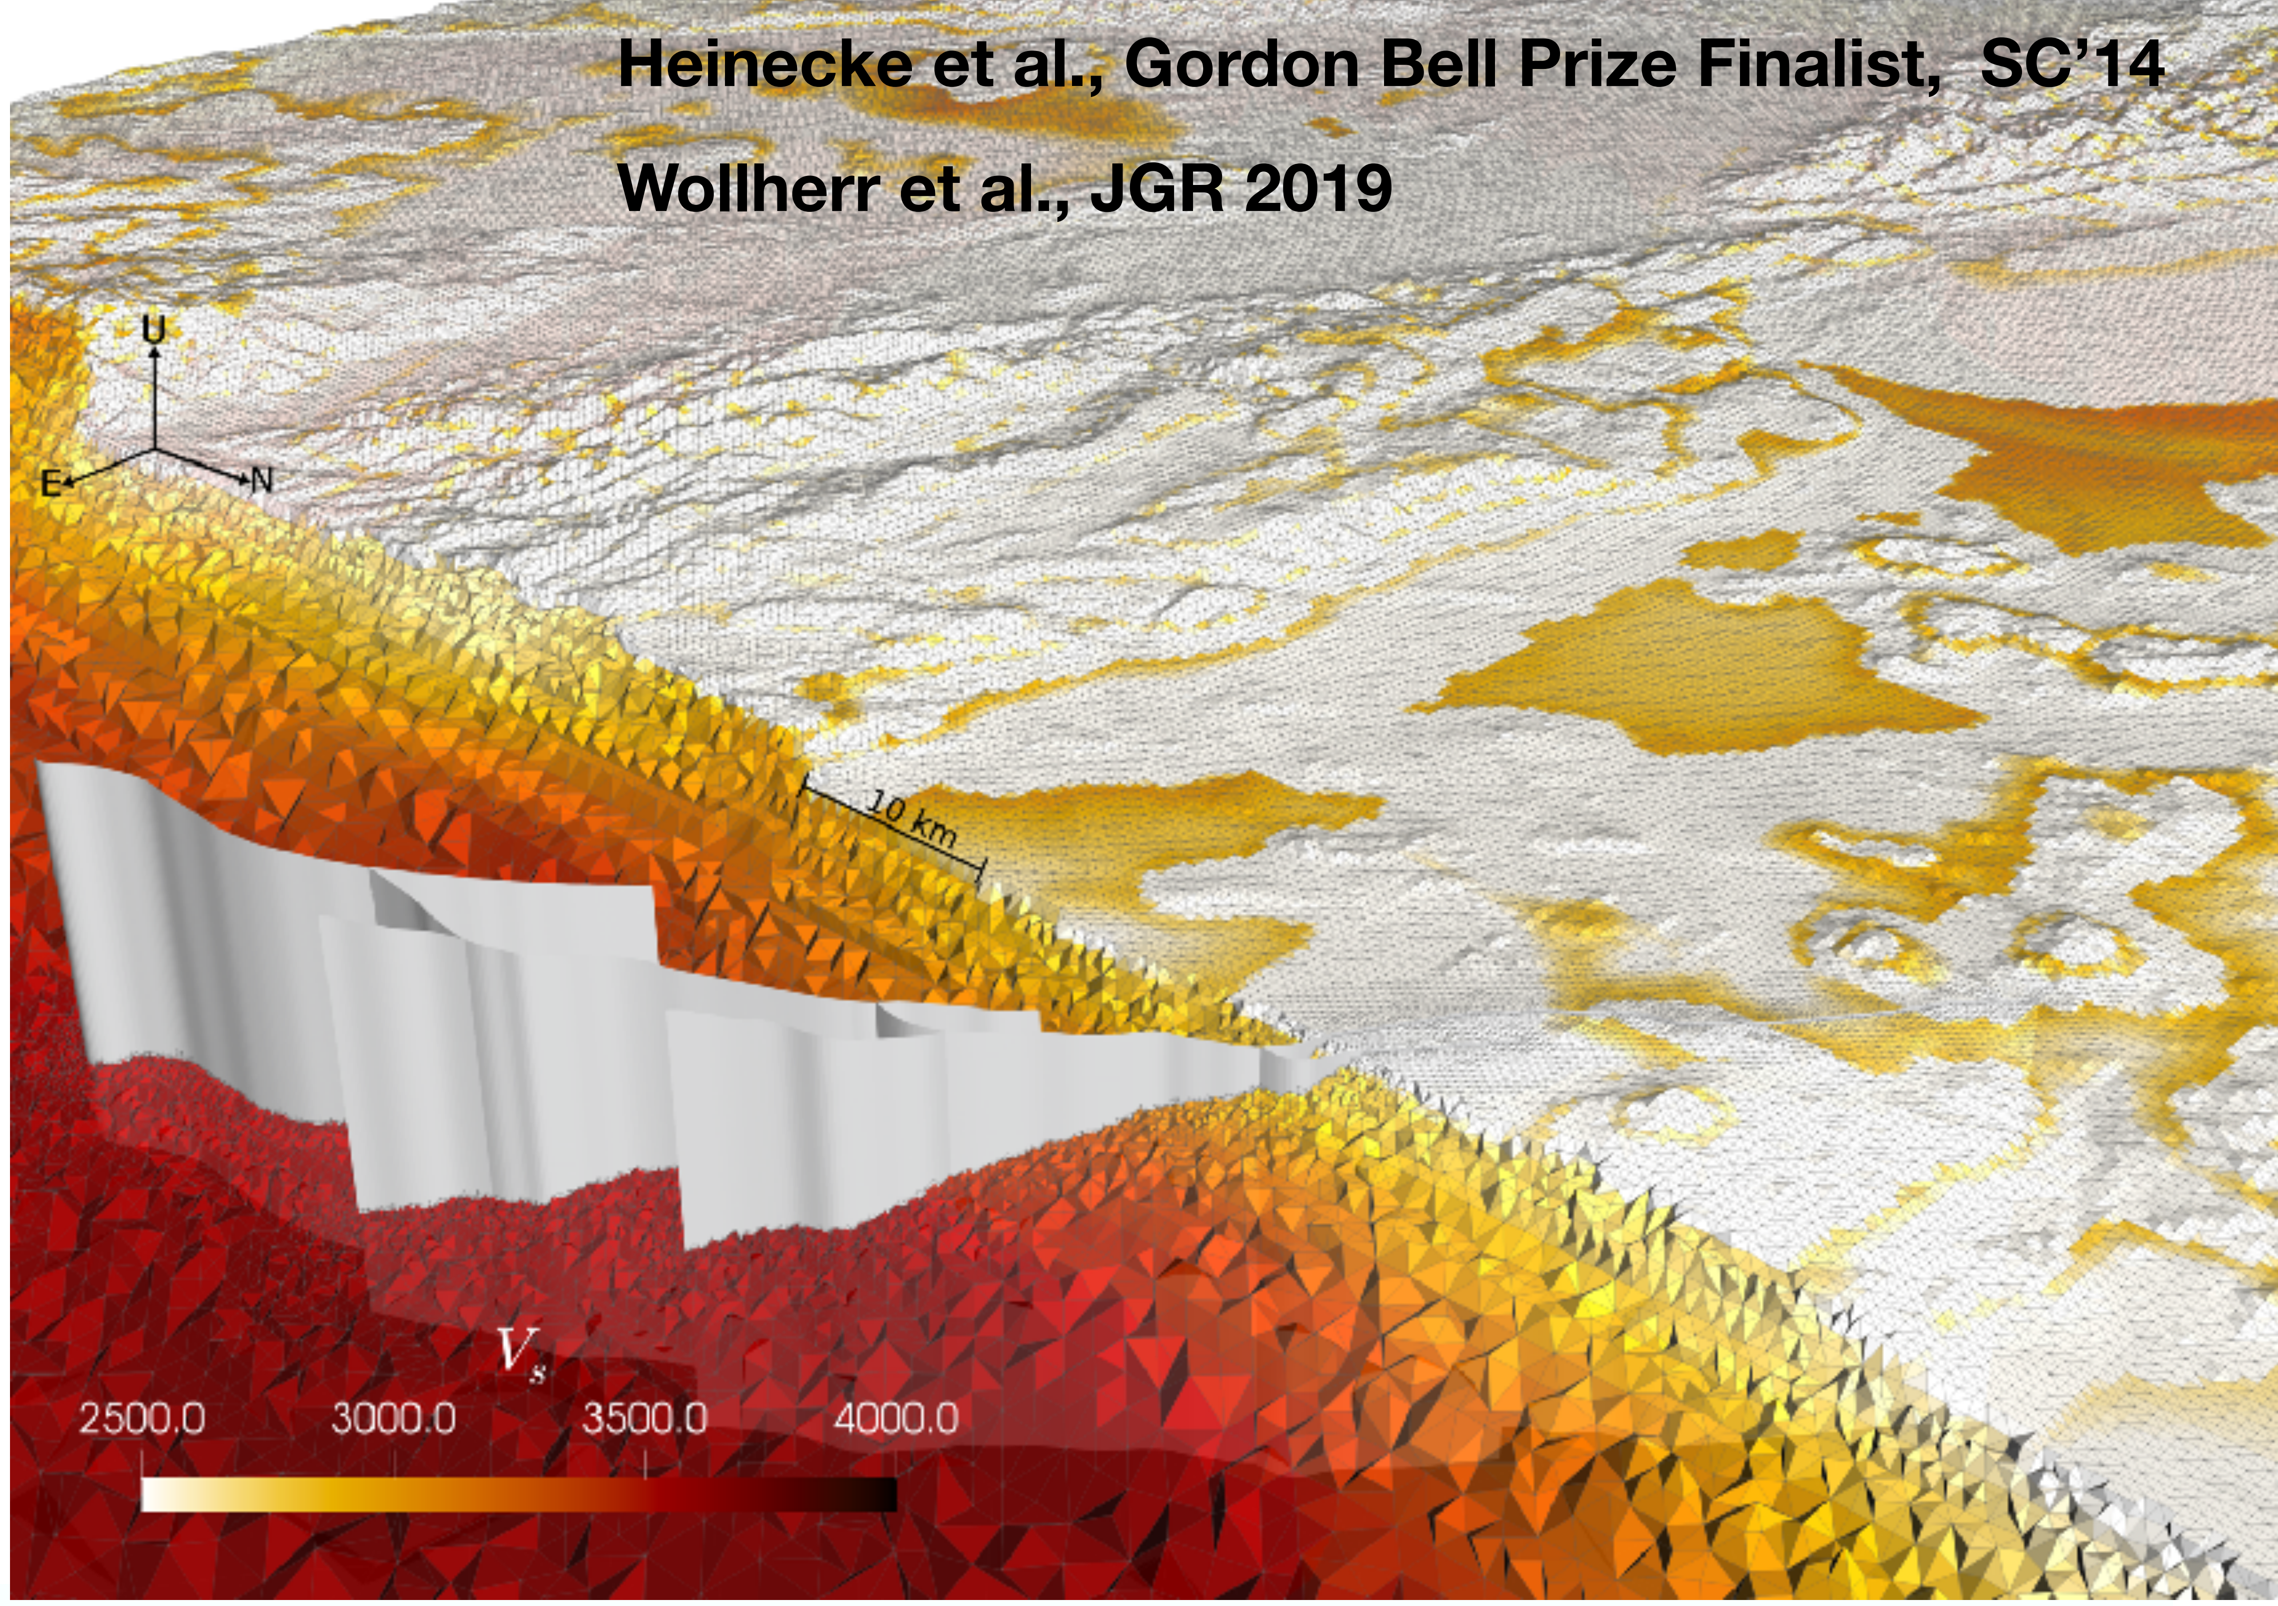

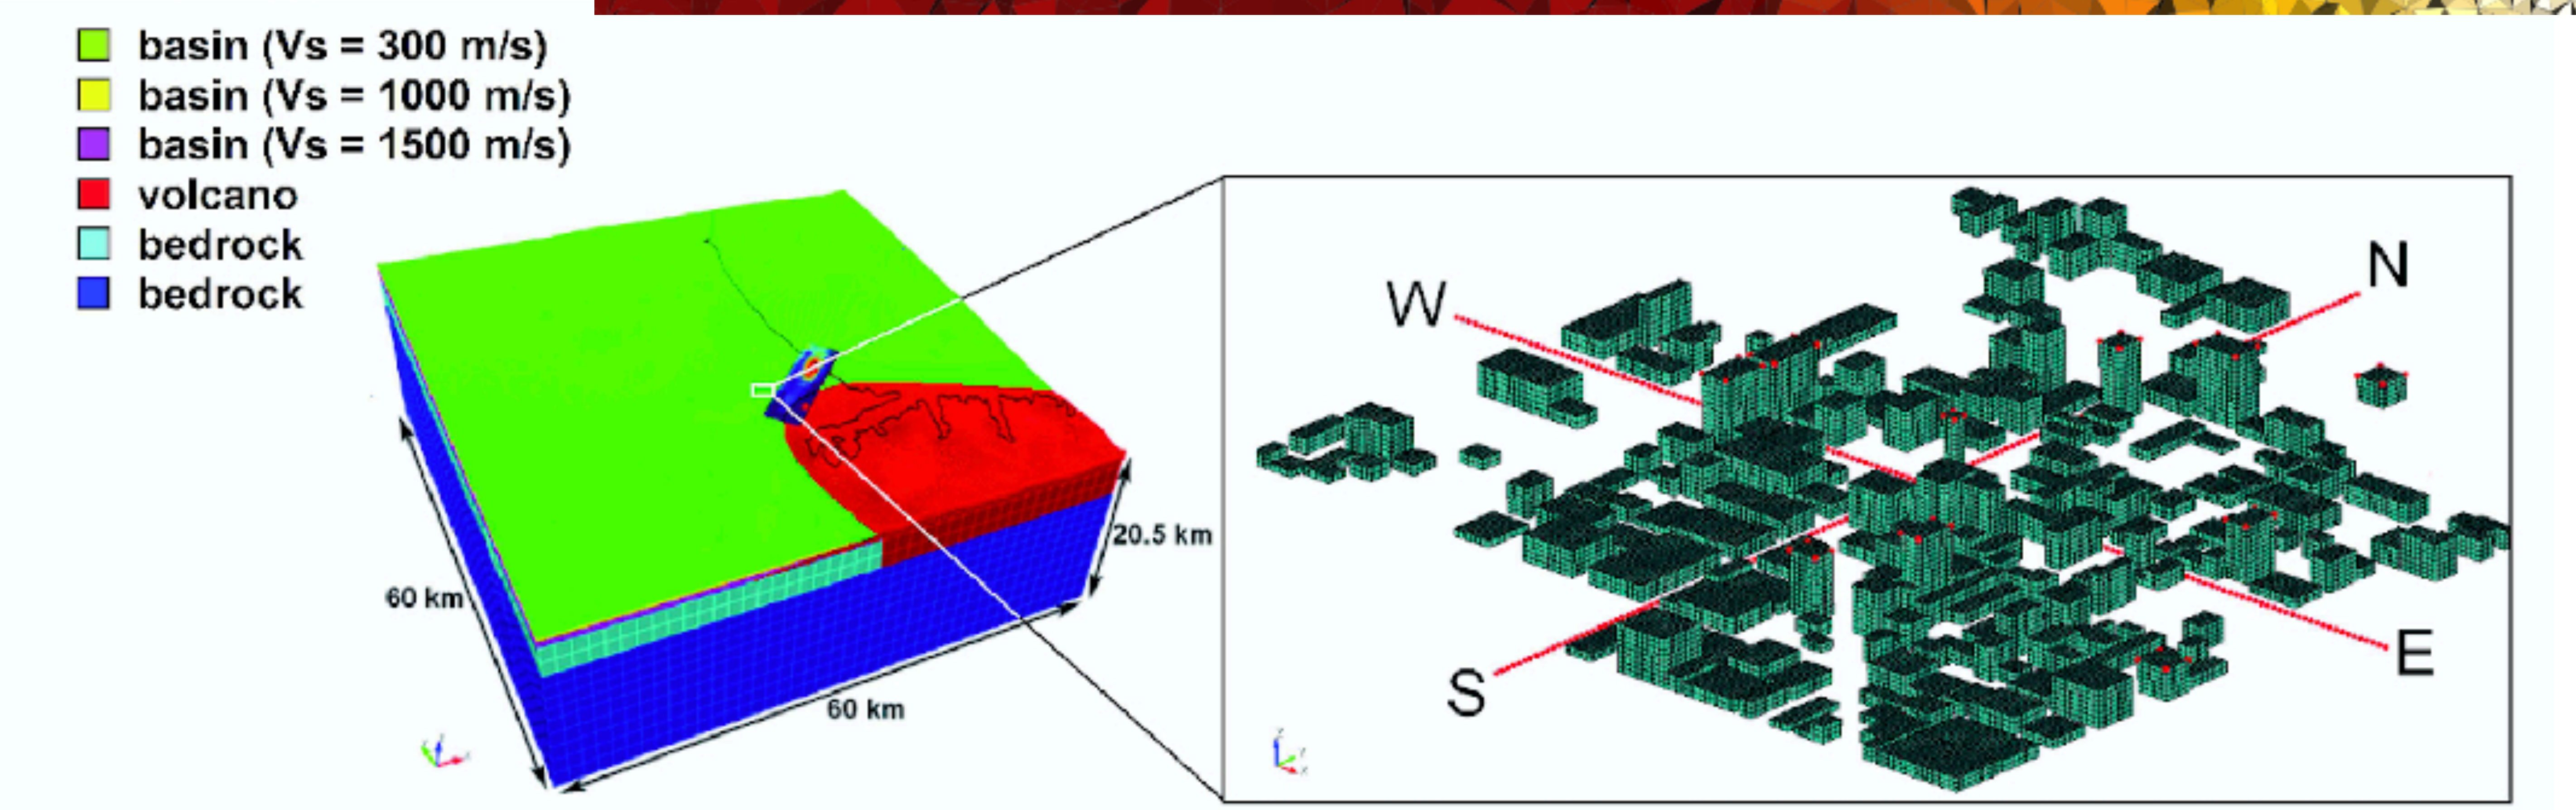

SPEED engineering seismology simulations of the seismic response of the 2011 Canterbury, NZ earthquake strong ground motions and soil city interactions (Mazzieri et al., 2013).

Computational earthquake seismology

- **Recent well-recorded earthquakes and laboratory experiments reveal striking source variability** (pulses/cracks; supershear speed; rupture cascades; slip reactivation)

- **Physics-based:** solving for spontaneous dynamic earthquake rupture as non-linear interaction of frictional failure and seismic wave propagation

- **Earthquakes:** frictional shear failure of brittle solids under compression along preexisting weak interfaces

- “**Bootstrapping**”: on methods originally not developed for earthquake source modelling

ChEESE project video, https://youtu.be/TAue0hEGD-k

Computational earthquake seismology

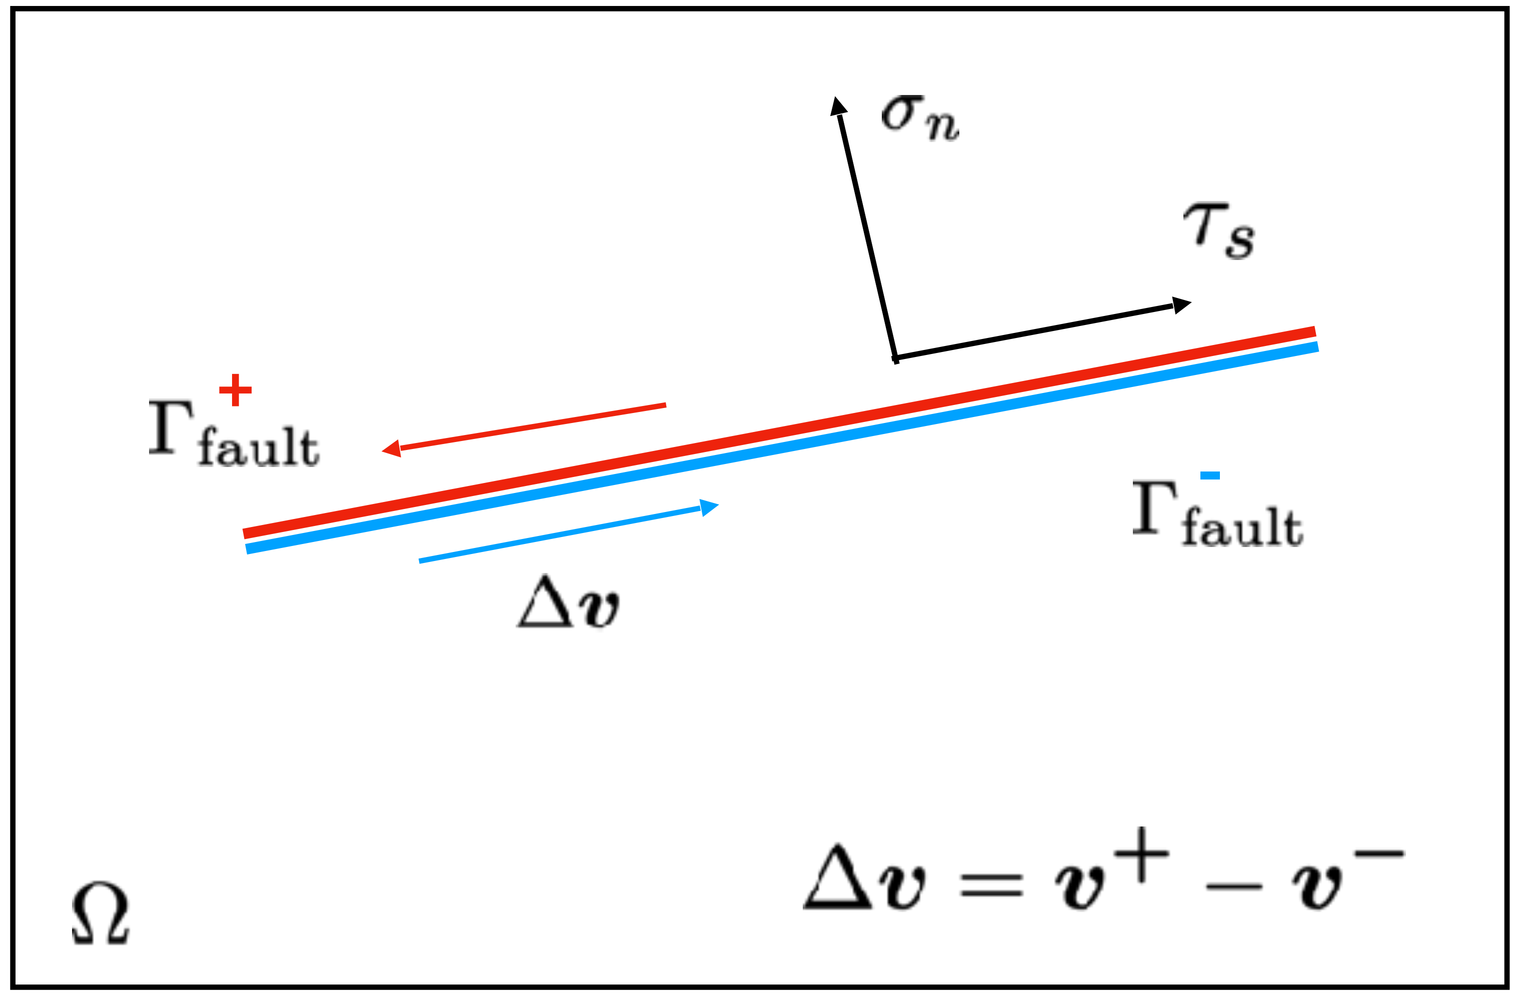

- Earthquake dynamic rupture is treated as a boundary condition in terms of **contact and friction**
- **Thin fault without 'opening'** - two matching fault surfaces are in unilateral contact
- **Displacement discontinuity** across the fault = **slip**
- Much complexity lives in the definition of **friction** (shear traction is bounded by the fault strength), and **fault geometry** and **intersections**

Earthquake dynamics are not predetermined: but evolve as a consequence of the model's initial conditions and the way the fault yields and slides controlled by an assigned friction law relating shear and normal traction on frictional interfaces

ChEESE

Center of Excellence for Exascale in Solid Earth

Computational earthquake seismology

- Earthquake dynamic rupture is treated as a boundary condition in terms of **contact and friction**
- **Thin fault without 'opening'** - two matching fault surfaces are in unilateral contact
- **Displacement discontinuity** across the fault = **slip**
- Much complexity lives in the definition of **friction** (shear traction is bounded by the fault strength), and **fault geometry** and **intersections**
- Can be implemented by **splitting the fault interface**
- **FD, FEM, SEM methods suffer from spurious oscillations - which have to be damped (e.g., by a thin layer of Kelvin-Voigt-Damping cells, Day et al., 2005)**

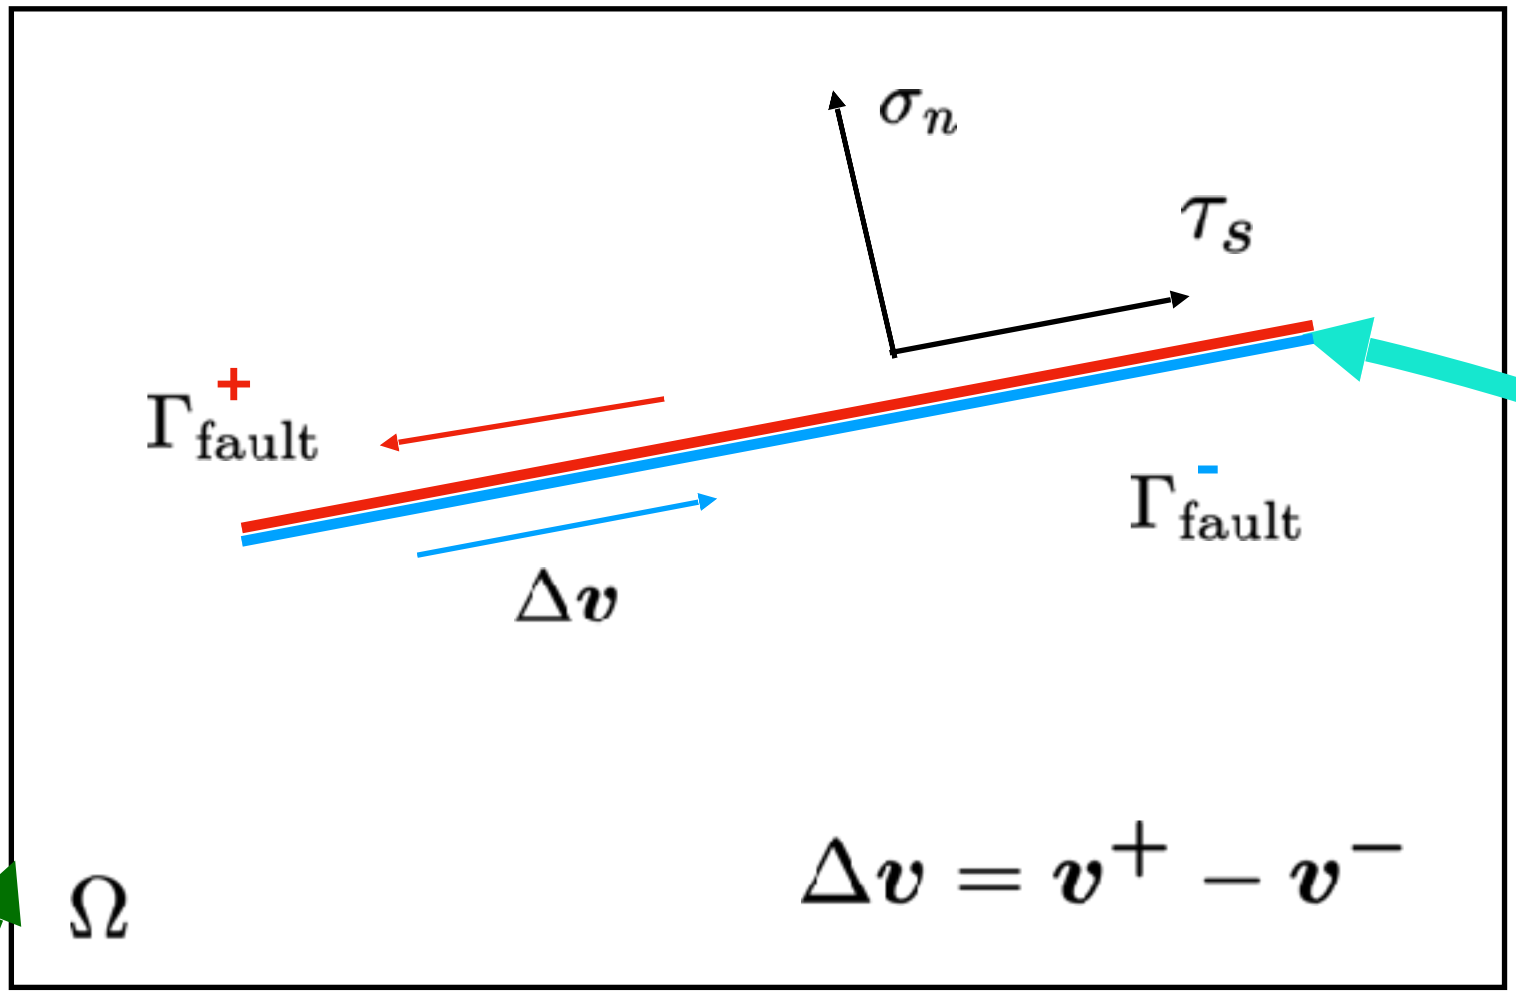

constitutive law (volume)

$$\frac{\partial}{\partial t}\sigma_{xx} - (\lambda + 2\mu)\frac{\partial}{\partial x}u - \lambda\frac{\partial}{\partial y}v - \lambda\frac{\partial}{\partial z}w = 0,$$
$$\frac{\partial}{\partial t}\sigma_{yy} - \lambda\frac{\partial}{\partial x}u - (\lambda + 2\mu)\frac{\partial}{\partial y}v - \lambda\frac{\partial}{\partial z}w = 0,$$
$$\frac{\partial}{\partial t}\sigma_{zz} - \lambda\frac{\partial}{\partial x}u - \lambda\frac{\partial}{\partial y}v - (\lambda + 2\mu)\frac{\partial}{\partial z}w = 0,$$

constitutive law (surfaces)

$$\tau_s = \mu_f \sigma_n.$$
$$|\boldsymbol{\tau}| \leq \tau_s,$$
$$(|\boldsymbol{\tau}| - \tau_s)|\Delta \mathbf{v}| = 0,$$
$$\Delta \mathbf{v}|\boldsymbol{\tau}| + |\Delta \mathbf{v}|\boldsymbol{\tau} = 0.$$

Adapted from Harris et al., SRL 2011, 2018

Computational earthquake seismology

- **Integration** and **interpretation** of full range of observations
- **Tightly links** seismology, geodesy, geology, tectonophysics, hydrology with numerical computing, data science, machine learning, applied mathematics, continuum mechanics, tribology, rock mechanics, materials science, and engineering

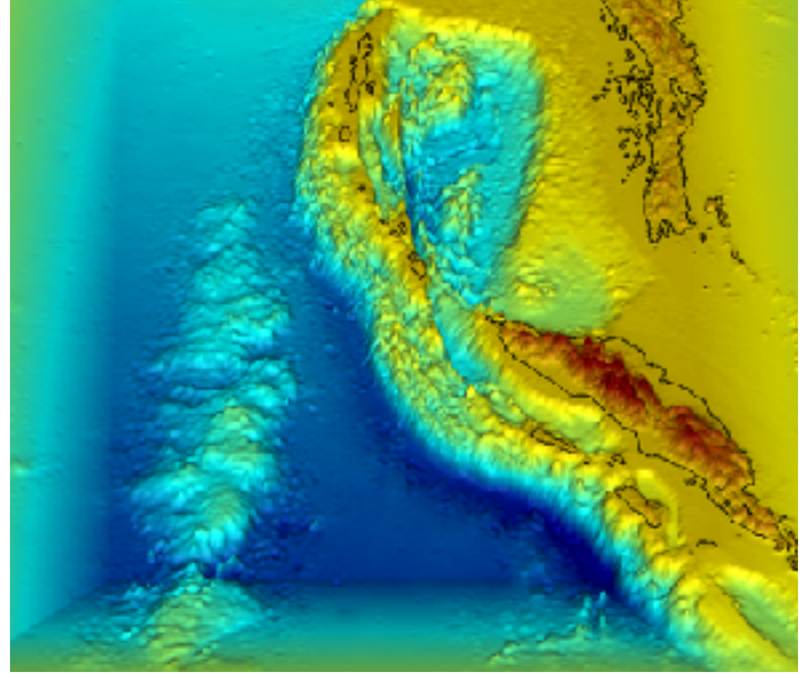

Geology

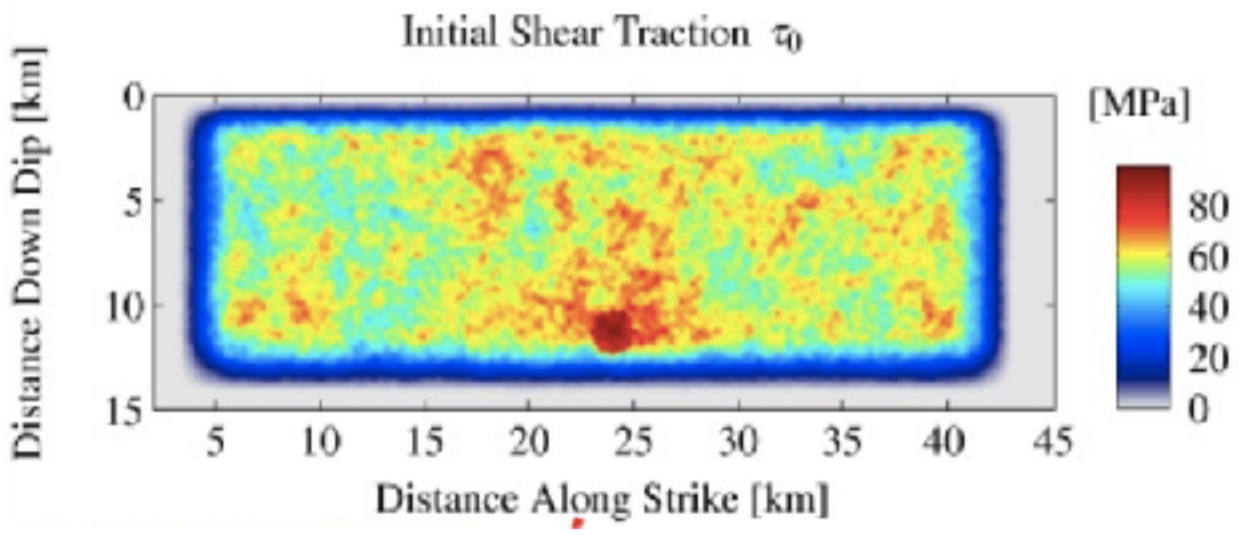

Initial fault stresses

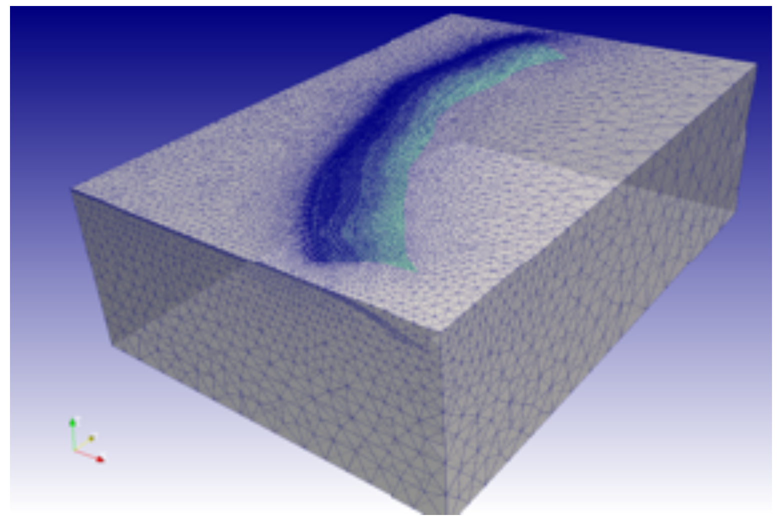

Mesh generation

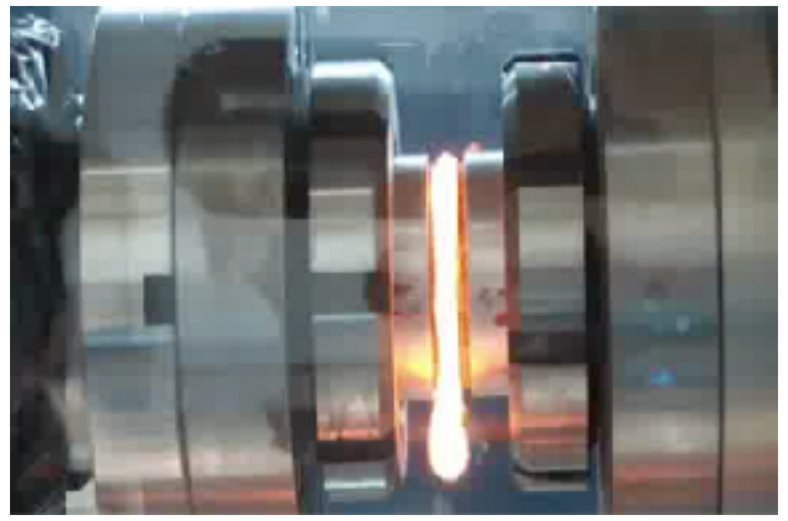

Friction experiments

"Input"

SOLVER

"Output"

Rupture dynamics

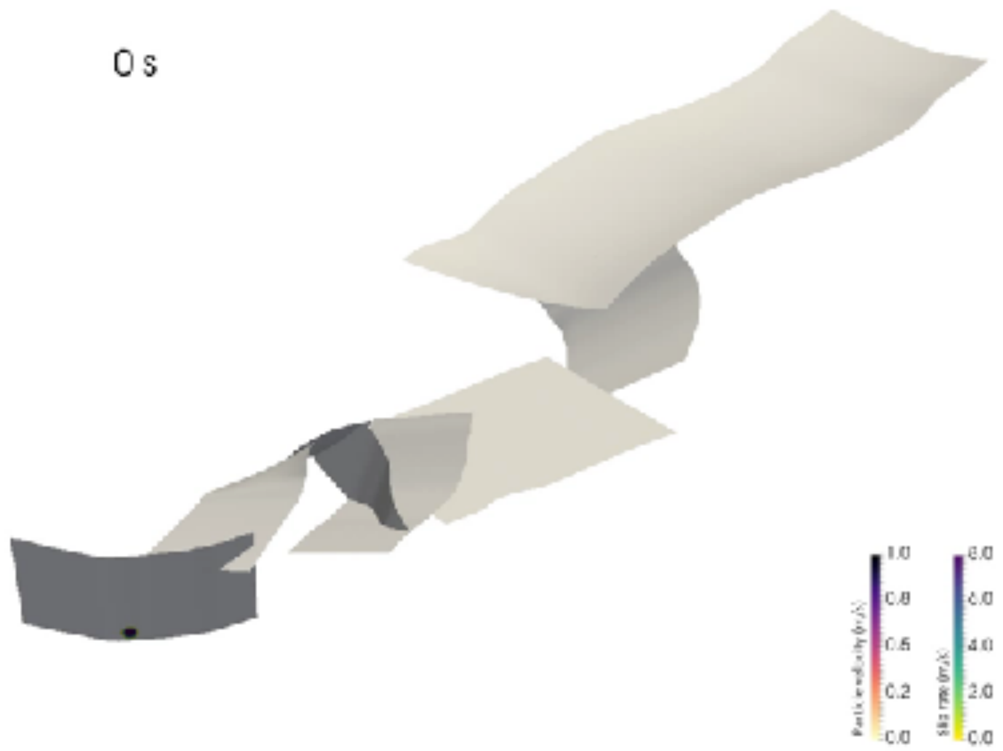

Ground deformation

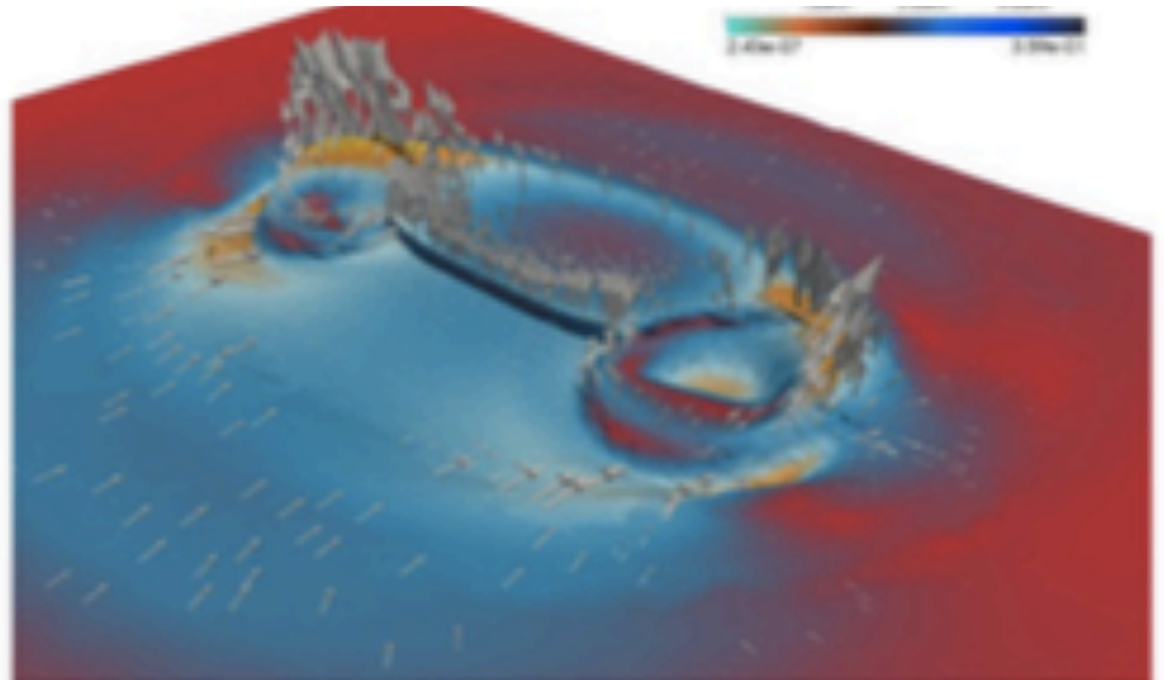

Synthetic observables

ChEESE
Center of Excellence for Exascale in Solid Earth

Computational earthquake seismology

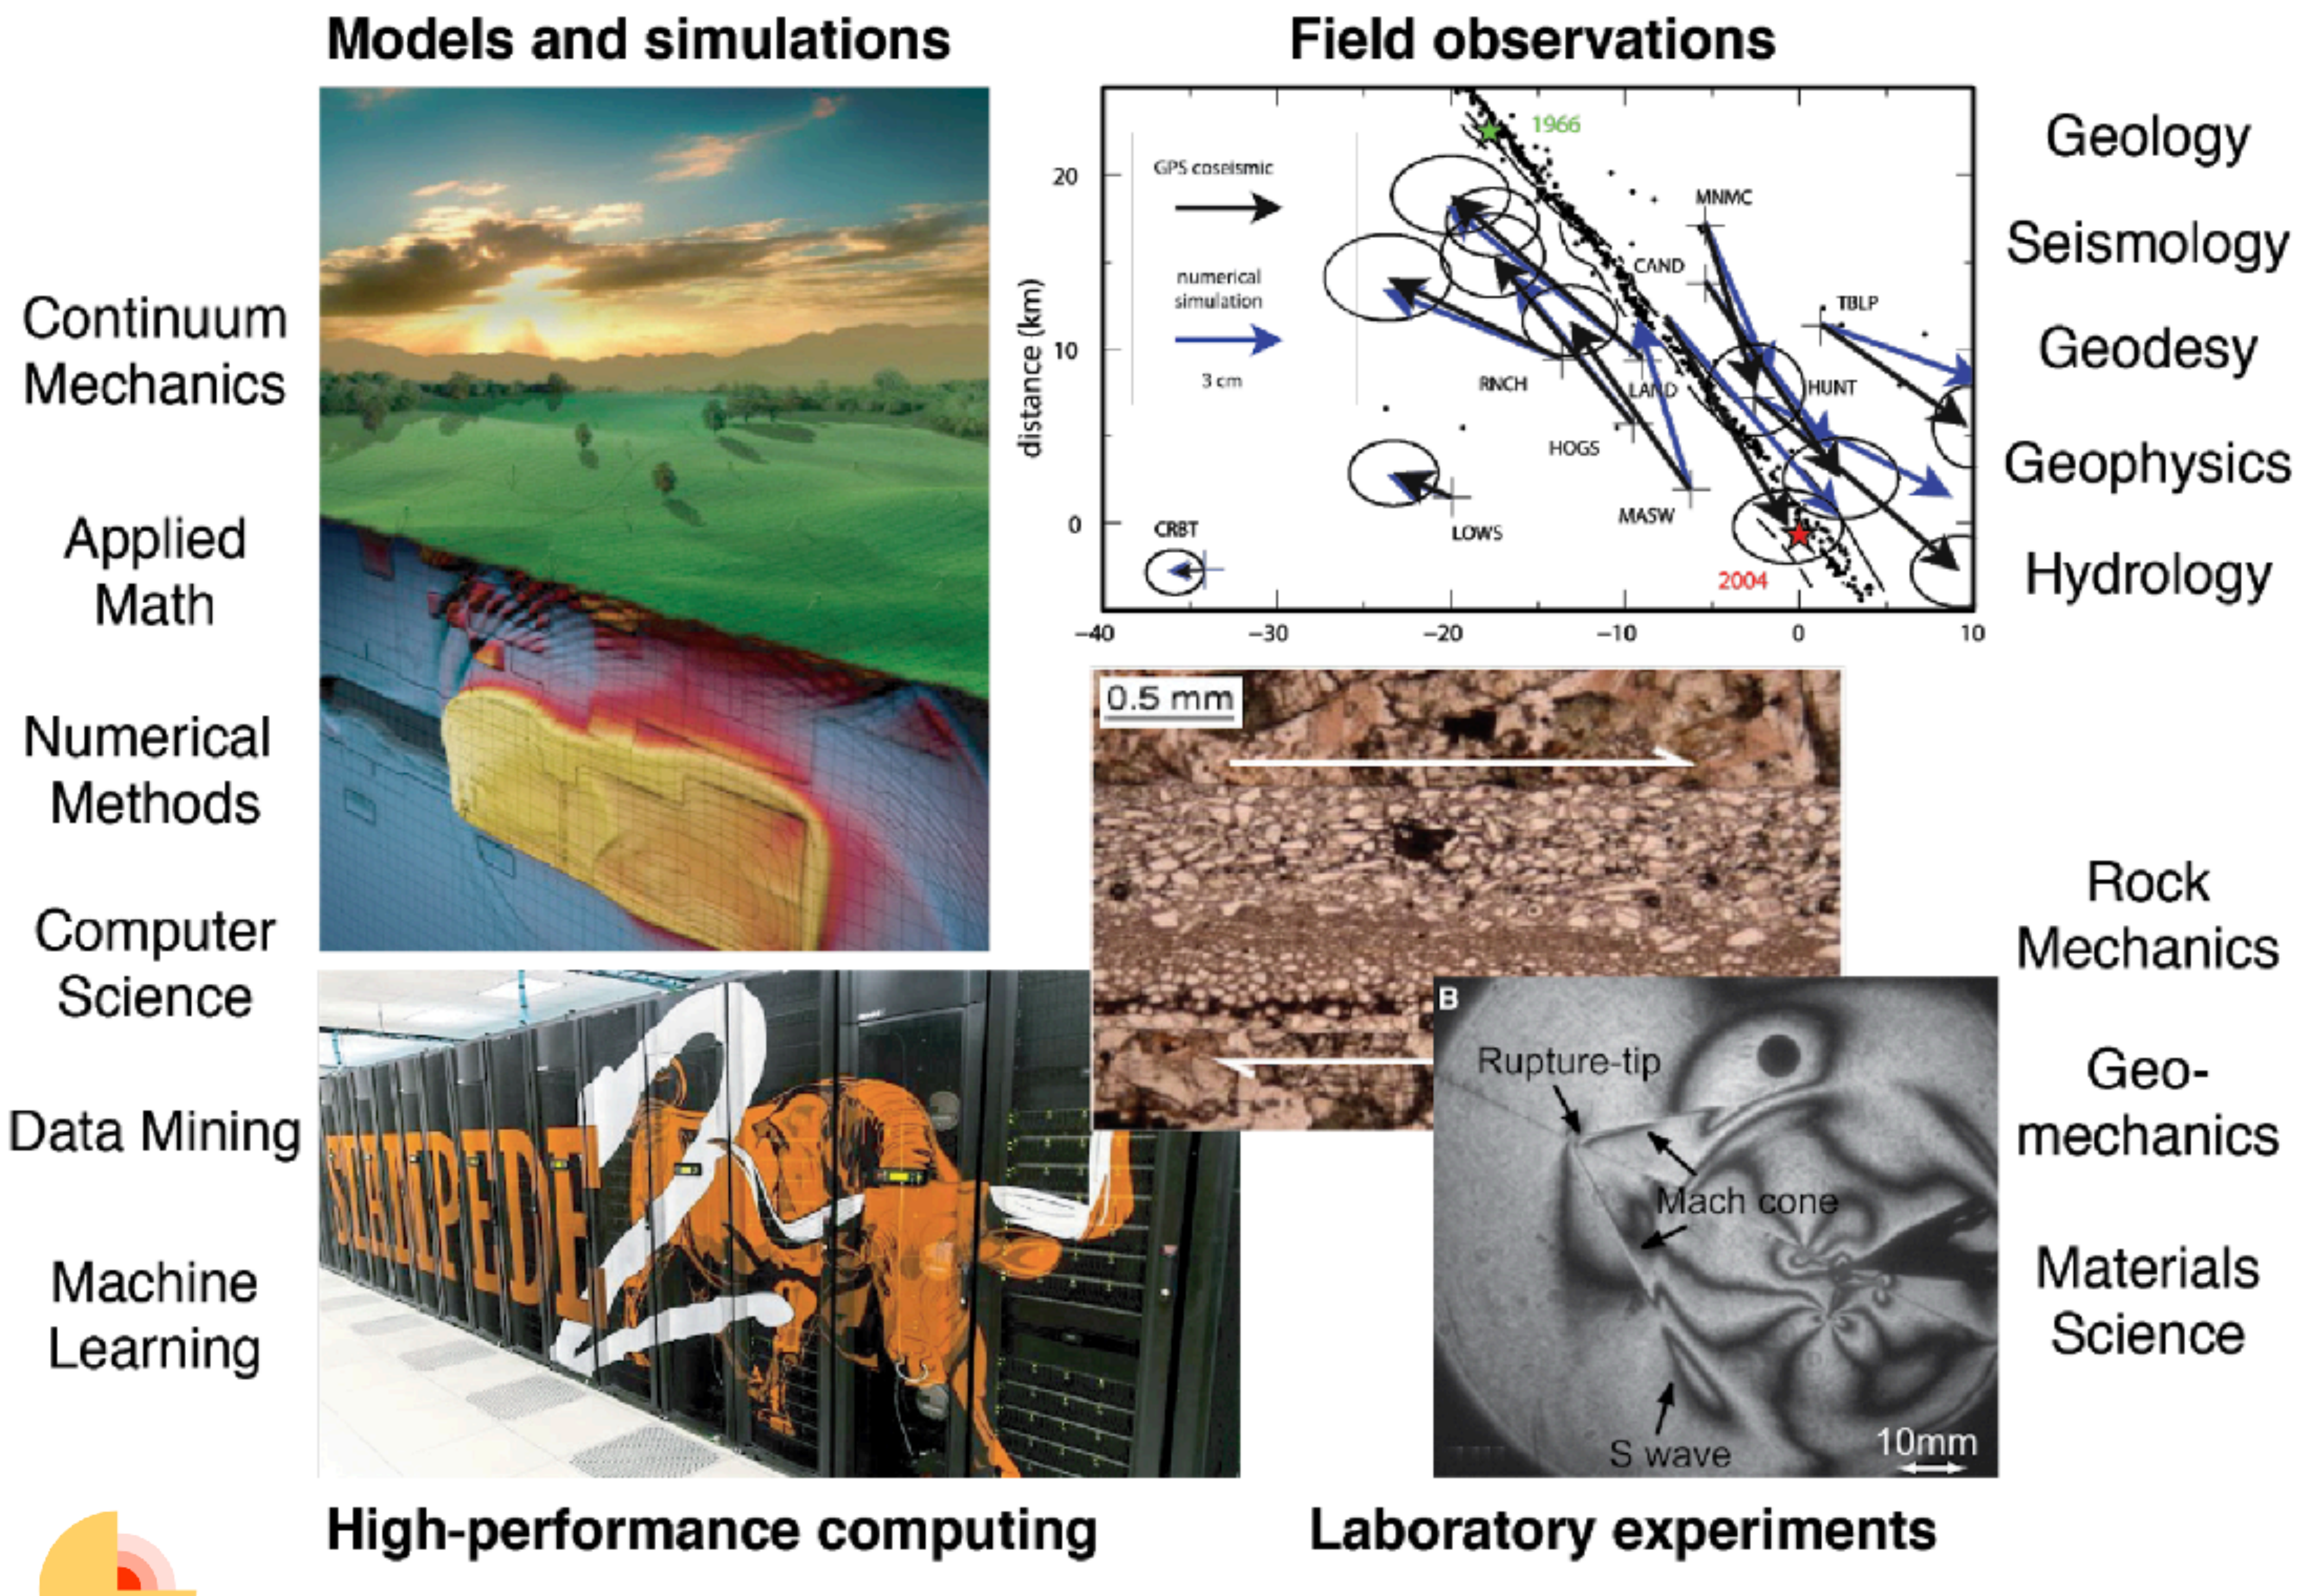

Challenge 1: Earthquake source processes are (very) **ill-constrained** and **highly non-linear.**

Challenge 2: Which **physical processes** are **dominant and relevant** at a given spatio-temporal scale (and in real earthquakes)? Can we justify the "cost" of their inclusion?

Challenge 3: How to **assimilate all available knowledge** in a suitable manner for **software** (numerical discretisation, solvers, equations solved) and **hardware** (heterogeneous HPC systems, energy concerns)?

CHEESE
Center of Excellence for Exascale in Solid Earth

white paper "Modeling earthquake source processes: from tectonics to dynamic rupture", Lapusta et al., 2019

SeisSol - ADER-DG modelling framework

SeisSol solves the seismic wave equations using the ADER-DG method on unstructured tetrahedral meshes.

The method, by design, permits:

- representing **complex geometries** - by discretising the volume via a tetrahedral mesh
- modelling **heterogenous media** - elastic, **viscoelastic**, **viscoplastic**, **anisotropic**, **poroelastic**
- **multi-physics coupling** - flux based formulation is natural for representing physics defined on interfaces
- **high accuracy** - modal flux based formulation allows us to suppress spurious (unresolved) high frequencies
- **high resolution** - suitable for parallel computing environments

DG for elastodynamic wave propagation problems:
Duru et al., 2021; Reinarz et al., 2020; Peyrusse et al., 2014; Mazzieri et al., 2013; Antonietti et al., 2012; Etienne et al., 2010; Wilcox et al., 2010; Grote & Diaz, 2009; de Basabe et al., 2008; Riviere et al., 2007; Chung & Enquist, 2006

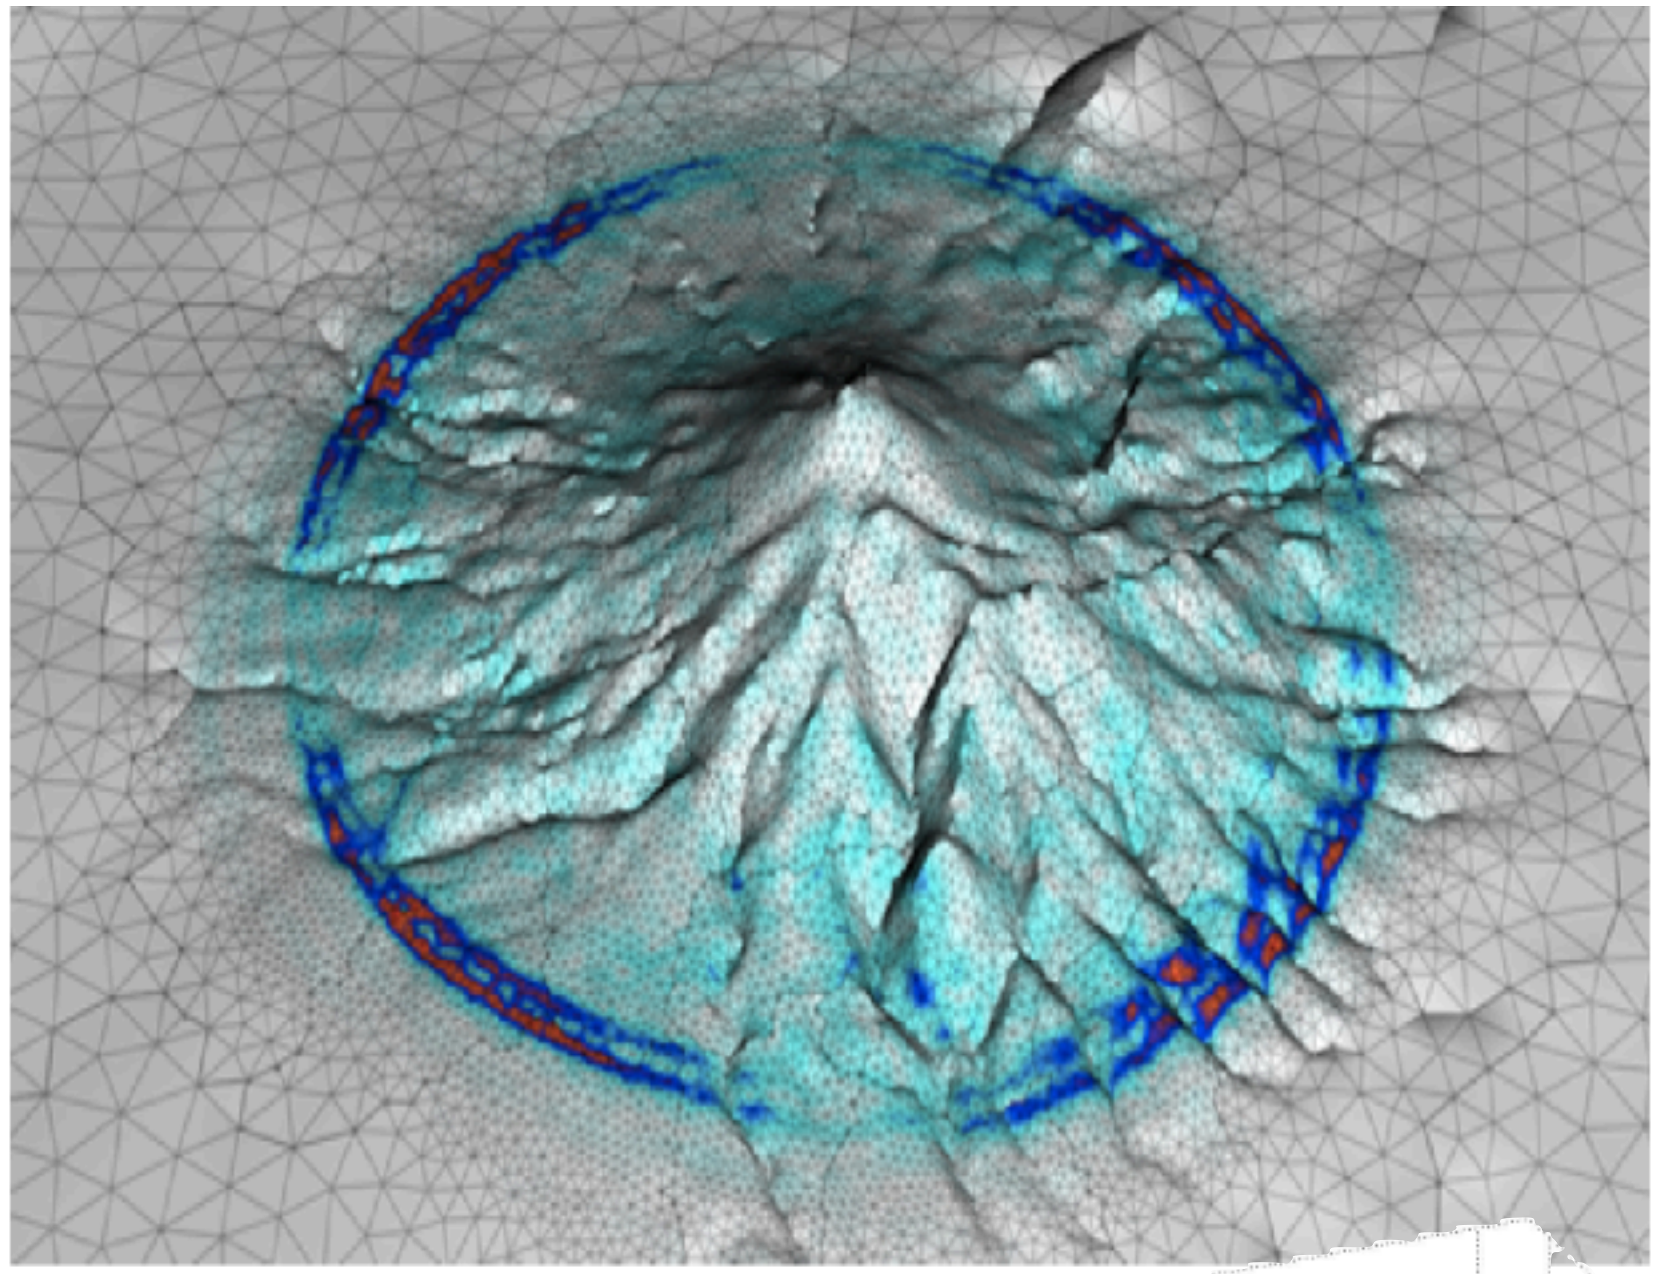

Wave field of a point source interacting with the topography of Mount Merapi Volcano.

PRACE ISC Award for producing the first simulations that obtained the “magical" performance milestone of 1 Peta-flop/s (10^{15} floating point operations per second) at the Munich Supercomputing Centre.

Käser and Dumbser, 2006; de la Puente et al., 2008; Pelties et al., 2014; Breuer et al., ISC’14

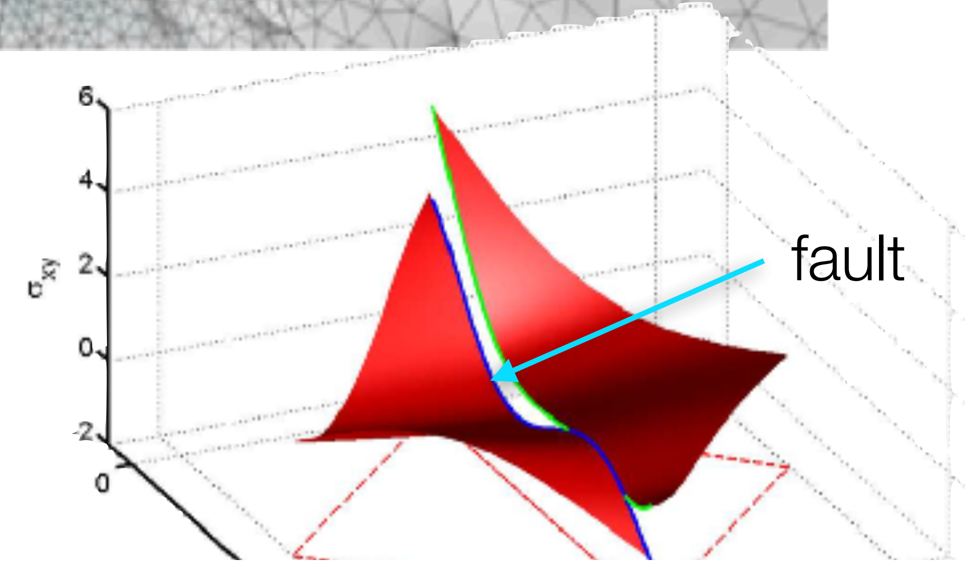

Representation of the shear stress discontinuity across the fault interface. Spontaneous rupture = internal boundary condition of flux term.

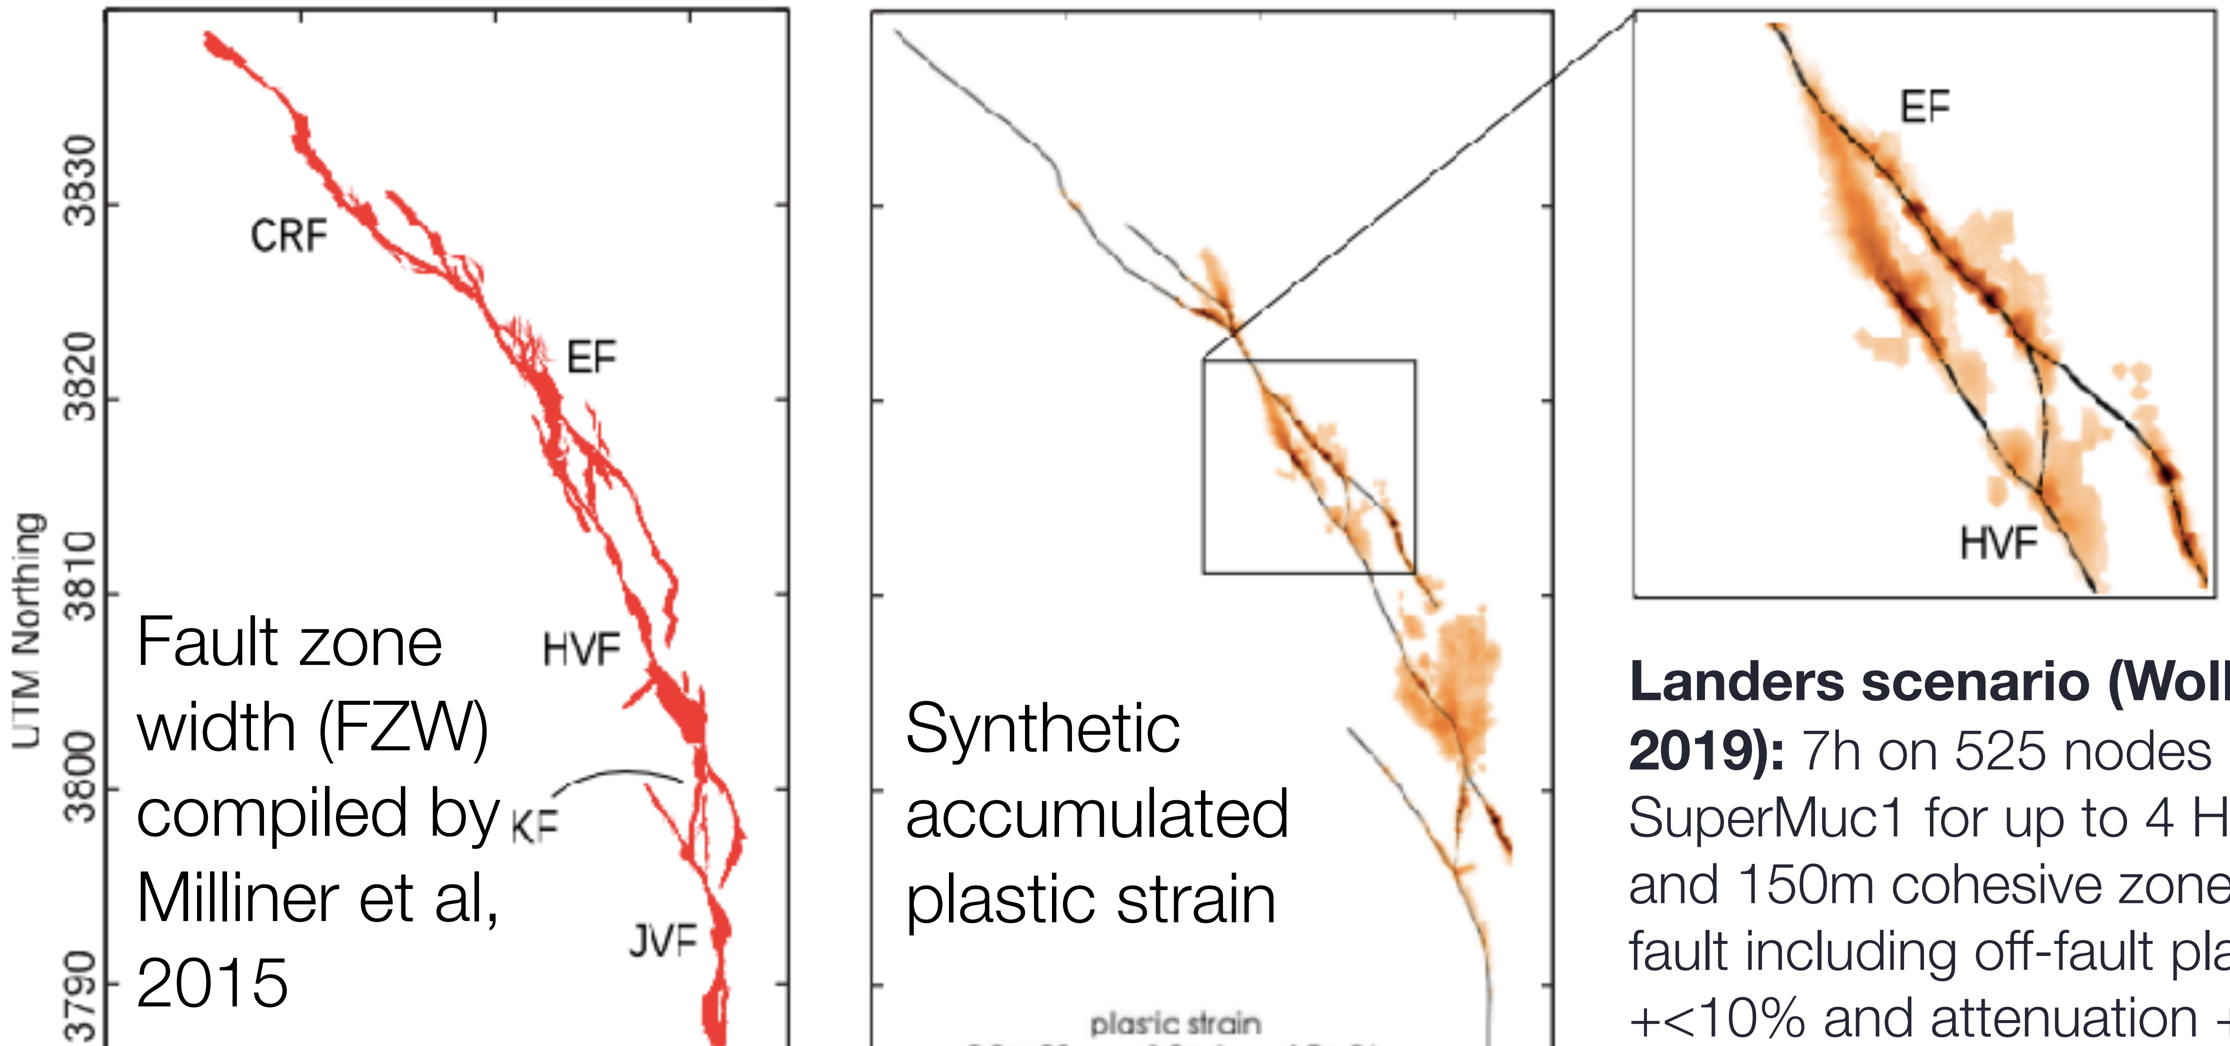

Landers scenario (Wollherr et al., 2019): 7h on 525 nodes of SuperMuc1 for up to 4 Hz Wavefield and 150m cohesive zone size on-fault including off-fault plasticity +<10% and attenuation +80%)

SeisSol - ADER-DG modelling framework

Why DG? Low numerical dispersion, minor changes for dynamic rupture, suitable for intersecting and branching faults/structure, favourable numerical dissipation of the Godunov flux (Hu et al. 1999; Kaeser et al. 2008; Hesthaven & Warburton 2010)

Why ADER? Equivalent high-order accuracy as in space using a single explicit time integration step. Increasing order of accuracy can be ‘cheap’ if hardware is exploited)

Why tets? Complex realities of geological subsurface, non-planar fault surfaces, intersecting undulating surfaces, static mesh refinement and coarsening

Why modal formulation? easy to build arbitrary high-order basis functions for tetrahedra, block-structured sparsity patterns with ADER

Why orthogonal basis functions? Dubiner’s basis functions (Cockburn et al. 2000), leads to well-conditioned diagonal mass matrix, all matrices can be pre-calculated analytically leading to a quadrature-free scheme (e.g., Atkins & Shu 1996)

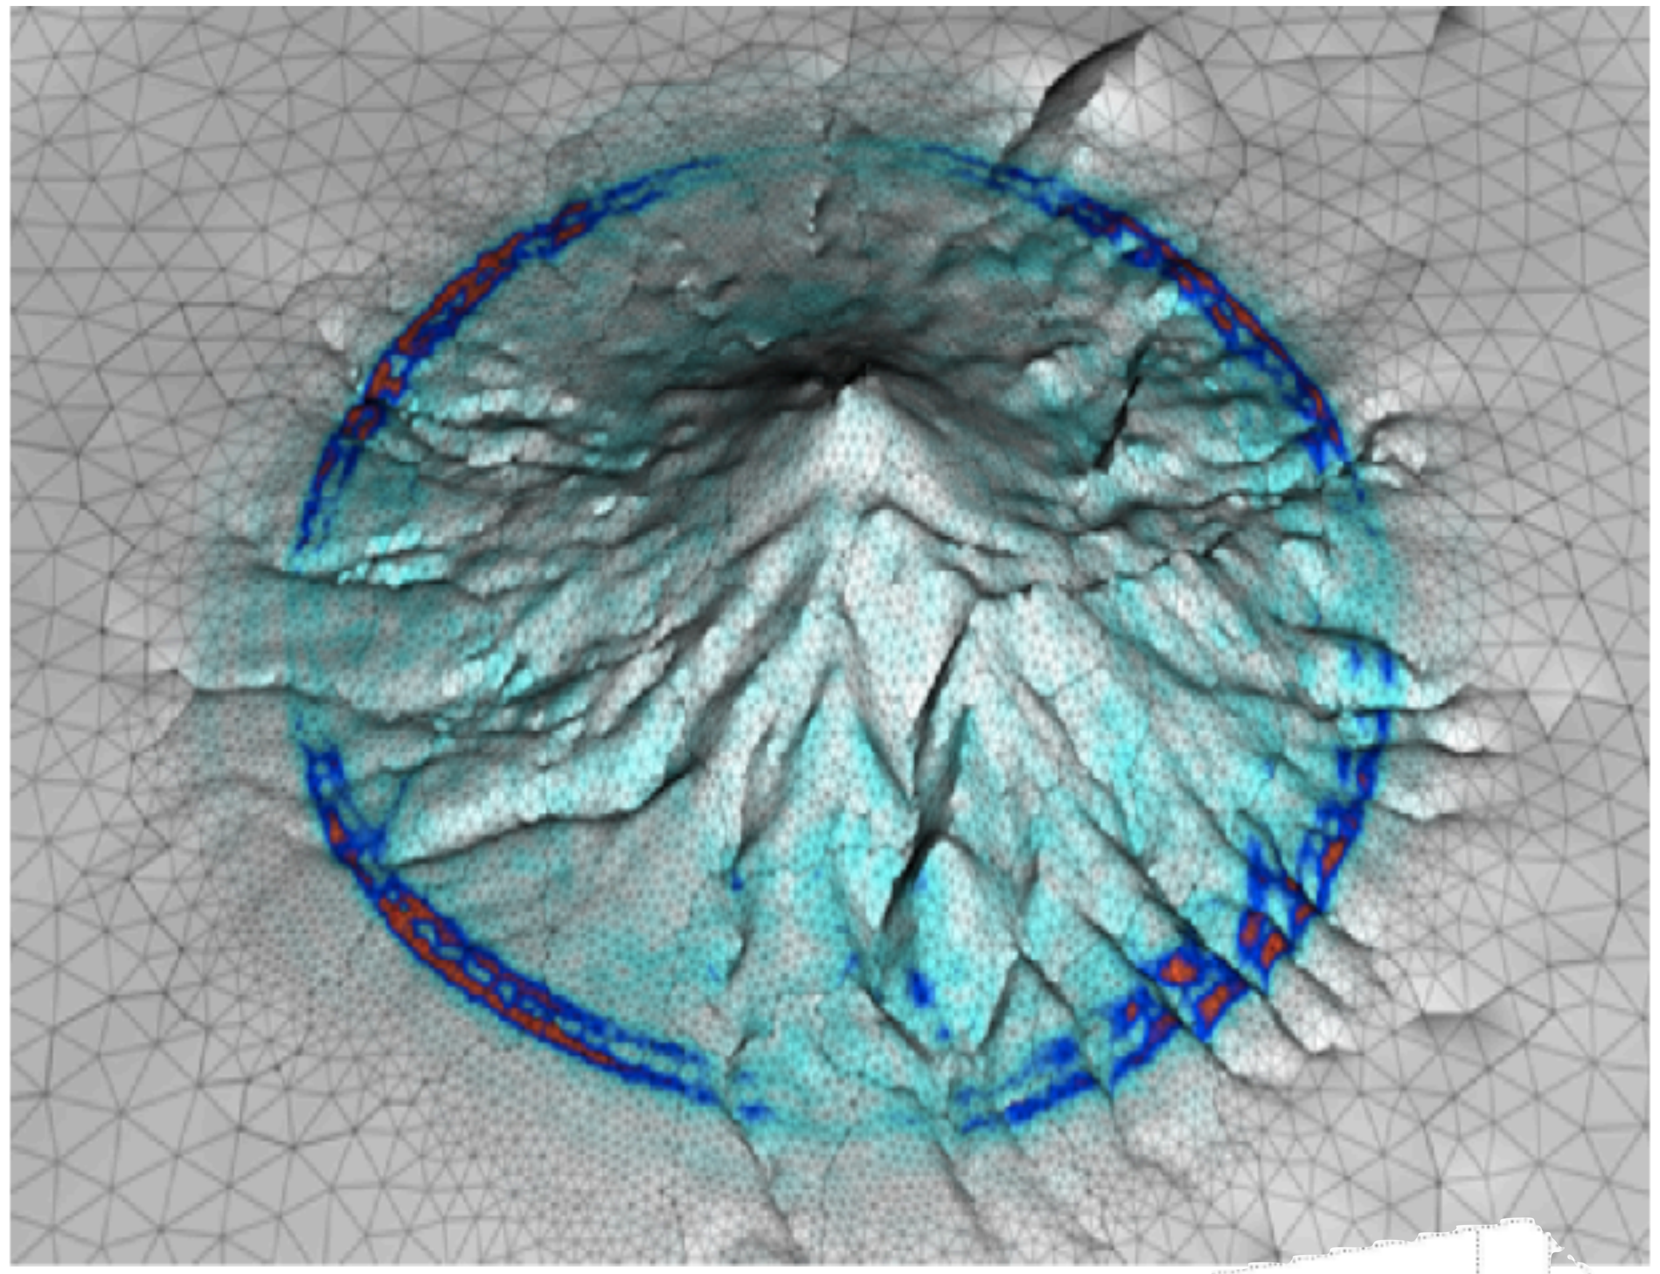

Wave field of a point source interacting with the topography of Mount Merapi Volcano.

PRACE ISC Award for producing the first simulations that obtained the “magical" performance milestone of 1 Peta-flop/s (10^{15} floating point operations per second) at the Munich Supercomputing Centre.

Käser and Dumbser, 2006; de la Puente et al., 2008; Pelties et al., 2014; Breuer et al., ISC’14

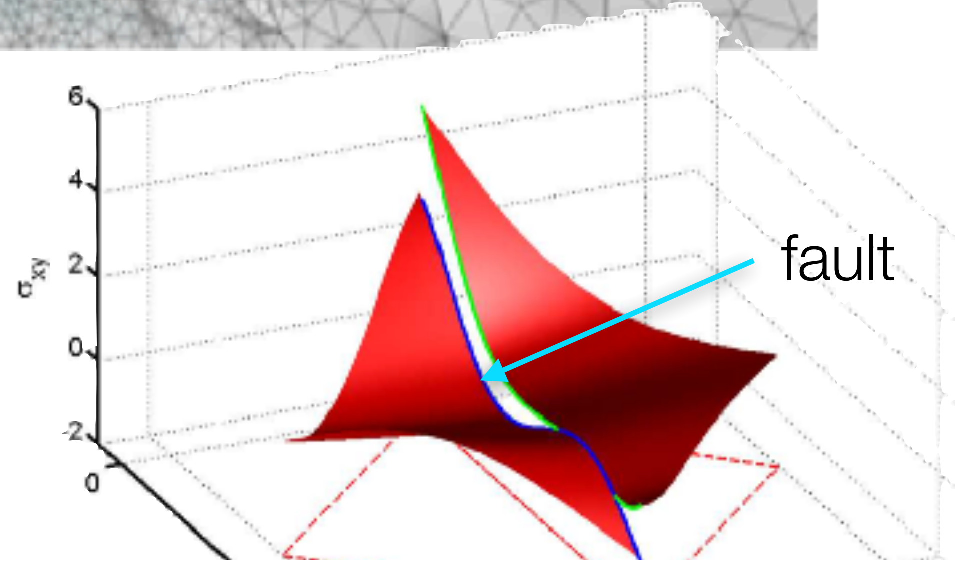

Representation of the shear stress discontinuity across the fault interface. Spontaneous rupture = internal boundary condition of flux term.

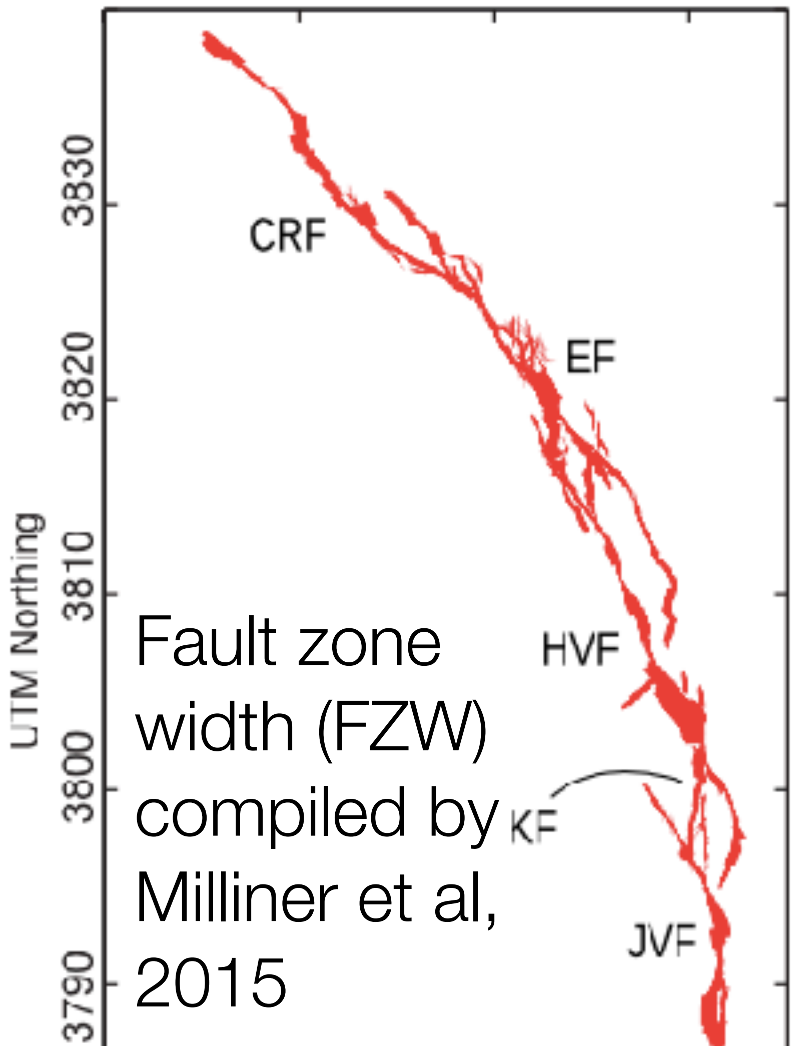

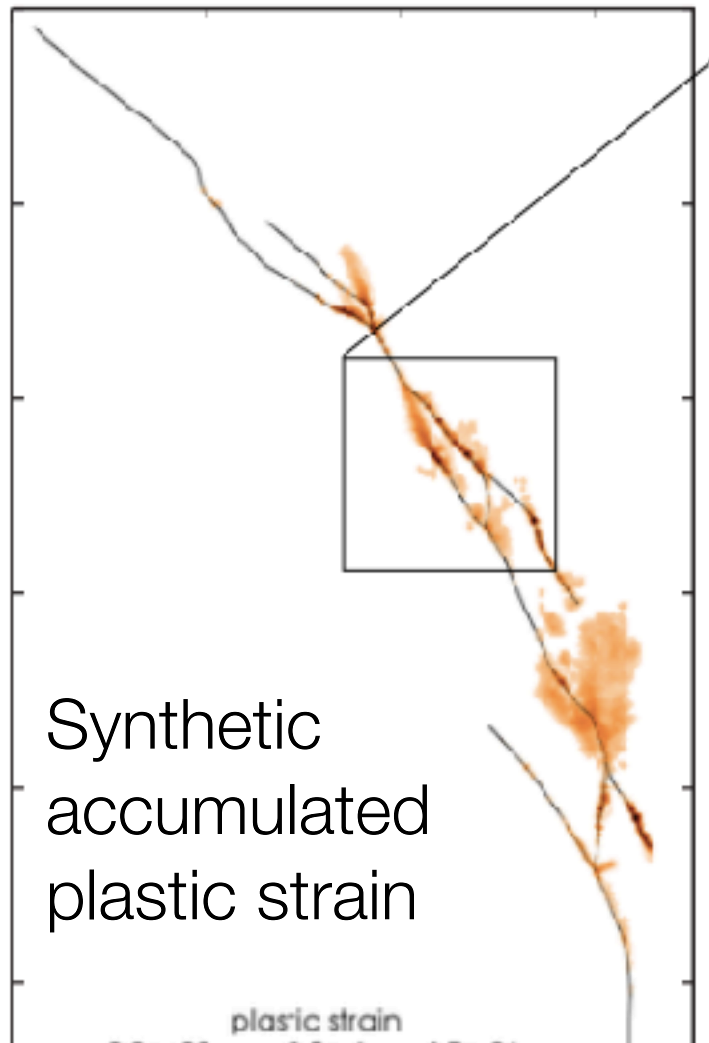

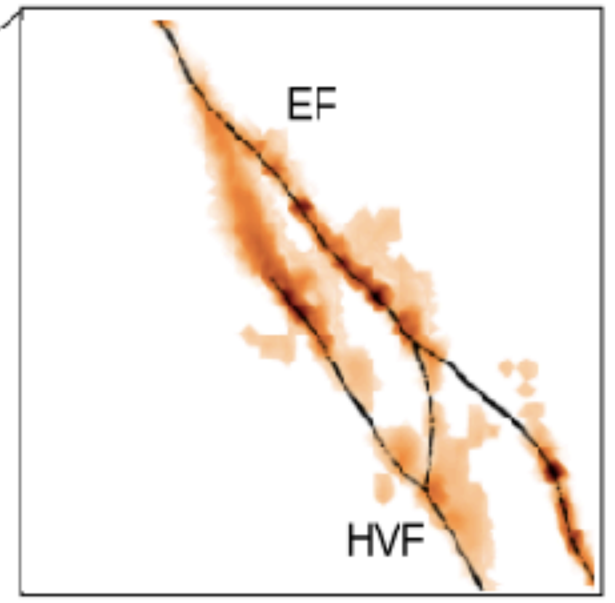

Landers scenario (Wollherr et al., 2019): 7h on 525 nodes of SuperMuc1 for up to 4 Hz Wavefield and 150m cohesive zone size on-fault including off-fault plasticity +<10% and attenuation +80%)

SeisSol - ADER-DG modelling framework

Why DG? Low numerical dispersion, minor changes for dynamic rupture, suitable for intersecting and branching faults/structure, favourable numerical dissipation of the Godunov flux (Hu et al. 1999; Kaeser et al. 2008; Hesthaven & Warburton 2010)

Why ADER? Equivalent high-order accuracy as in space using a single explicit time integration step. Increasing order of accuracy can be 'cheap' if hardware is exploited)

Why tets? Complex realities of geological subsurface, non-
plana
mesh

Why
basis ... erns
with

A software that allows for rapid setup of models with realistic non-planar and intersecting fault systems while exploiting the accuracy of a high-order numerical method

Why orthogonal basis functions? Dubiner's basis functions (Cockburn et al. 2000), leads to well-conditioned diagonal mass matrix, all matrices can be pre-calculated analytically leading to a quadrature-free scheme (e.g., Atkins & Shu 1996)

Good news: DG's "extra" flops (storage, time to solution) compared to e.g. FEM can be performed fast using Computational Science

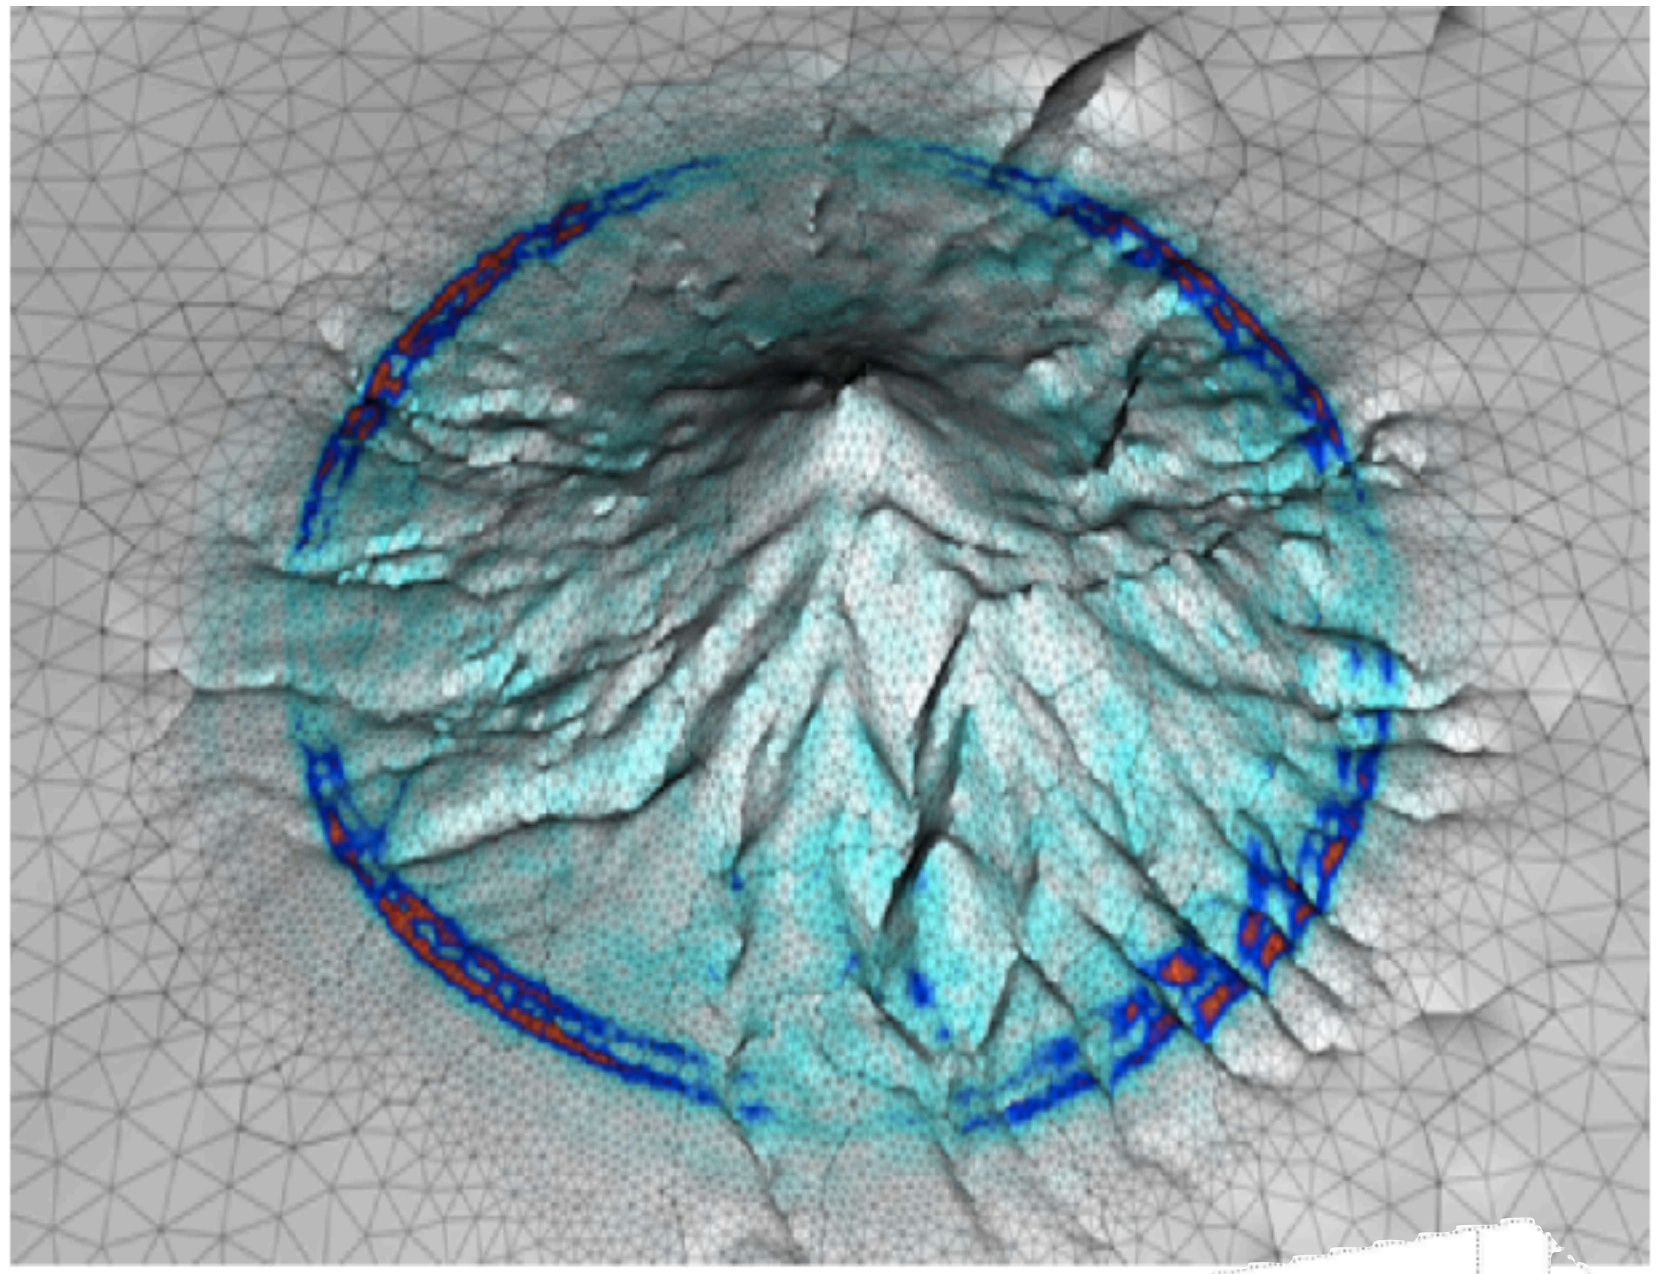

Wave field of a point source interacting with the topography of Mount Merapi Volcano.

PRACE ISC Award for producing the first simulations that obtained the "magical" performance milestone of 1 Peta-flop/s (10^{15} floating point operations per second) at the Munich Supercomputing Centre.

Käser and Dumbser, 2006; de la Puente et al., 2008; Pelties et al., 2014; Breuer et al., ISC'14

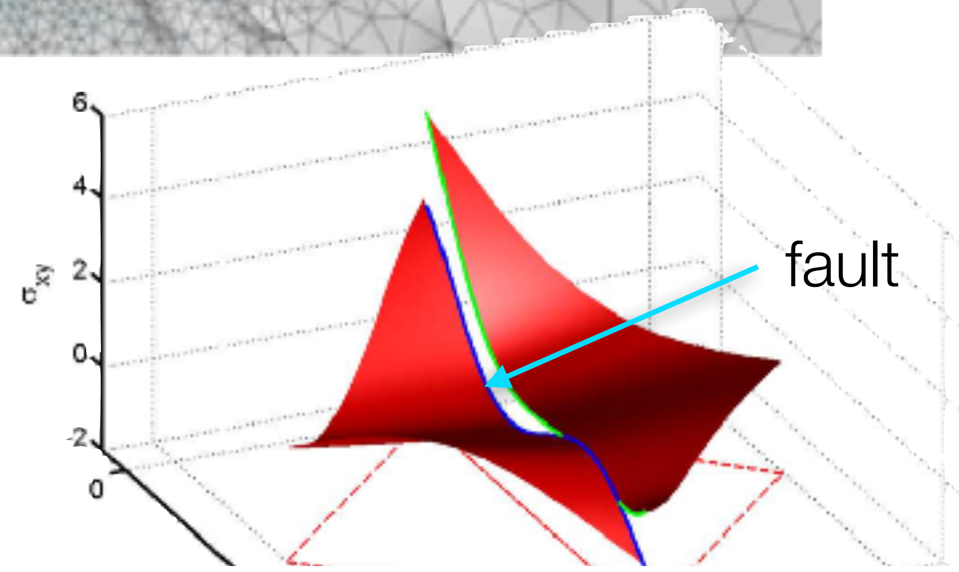

Representation of the shear stress discontinuity across the fault interface. Spontaneous rupture = internal boundary condition of flux term.

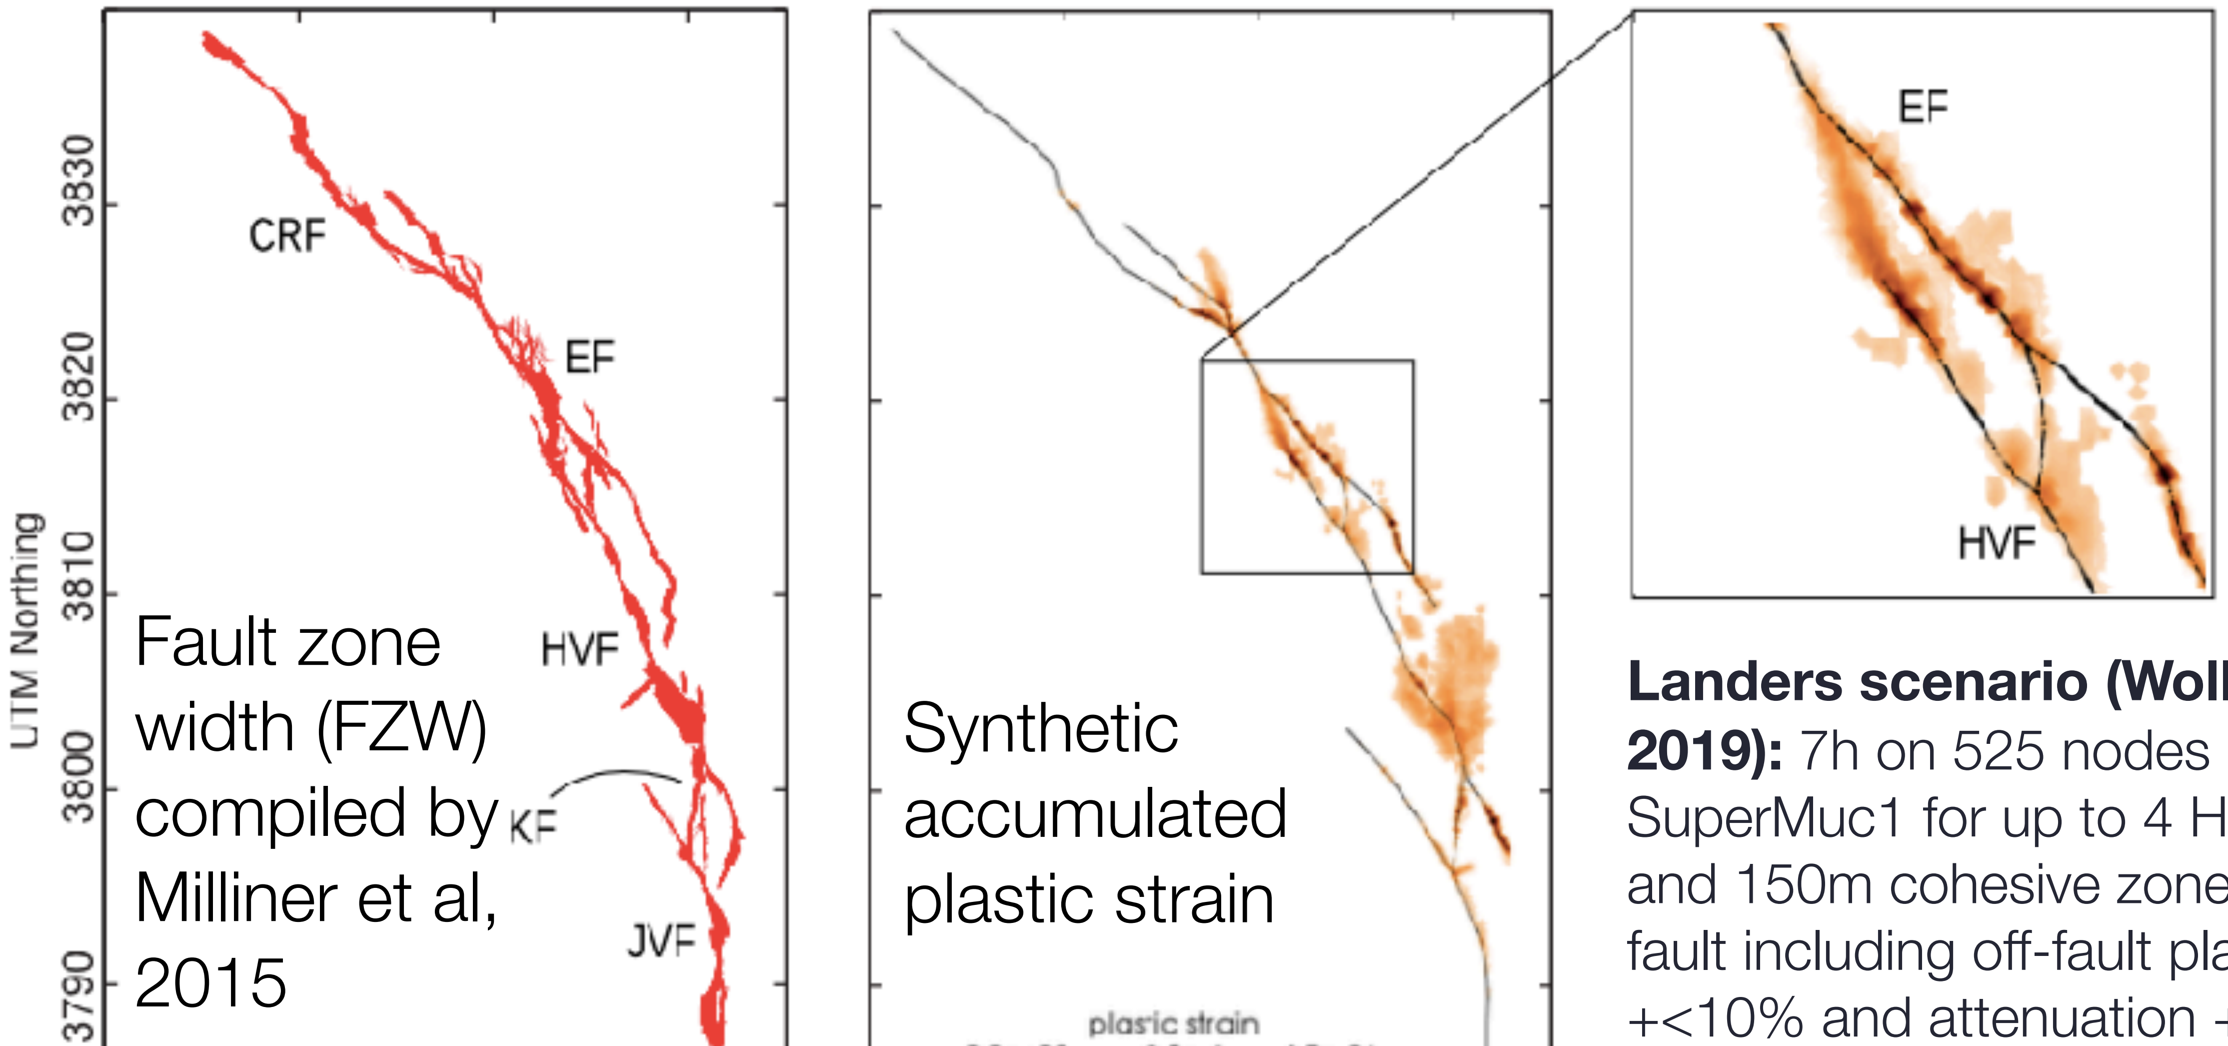

Landers scenario (Wollherr et al., 2019): 7h on 525 nodes of SuperMuc1 for up to 4 Hz Wavefield and 150m cohesive zone size on-fault including off-fault plasticity +<10% and attenuation +80%)

Computational earthquake seismology

What is the cost (and why is optimisation worthwhile to do) ?
Two linked 100 second scenarios (3D subsurface, topography, attenuation, off-fault plasticity, resolving up to **5 Hz**) of the **2018 Ridgecrest, CA, earthquake sequence.**

- Computing
8.8 hours on 400 nodes (48 Skylake cores) of the SuperMUC-NG supercomputer **~170k CPUh**

- Money
Energy charged at $0.10 per kWh —> $320
Cloud service such as AWS ~ $6500

- Energy*
~3.2 MWh
~ 2 barrel of oil equivalent (BOE)

- Carbon*
~2968 pounds of CO2 (~flying from London to Los Angeles)

*numbers for Pangea (TOTAL) 220k Xeon E5 cores
e.g. Shaheen-2 uses 30% less energy = flying from MUC -> ORD

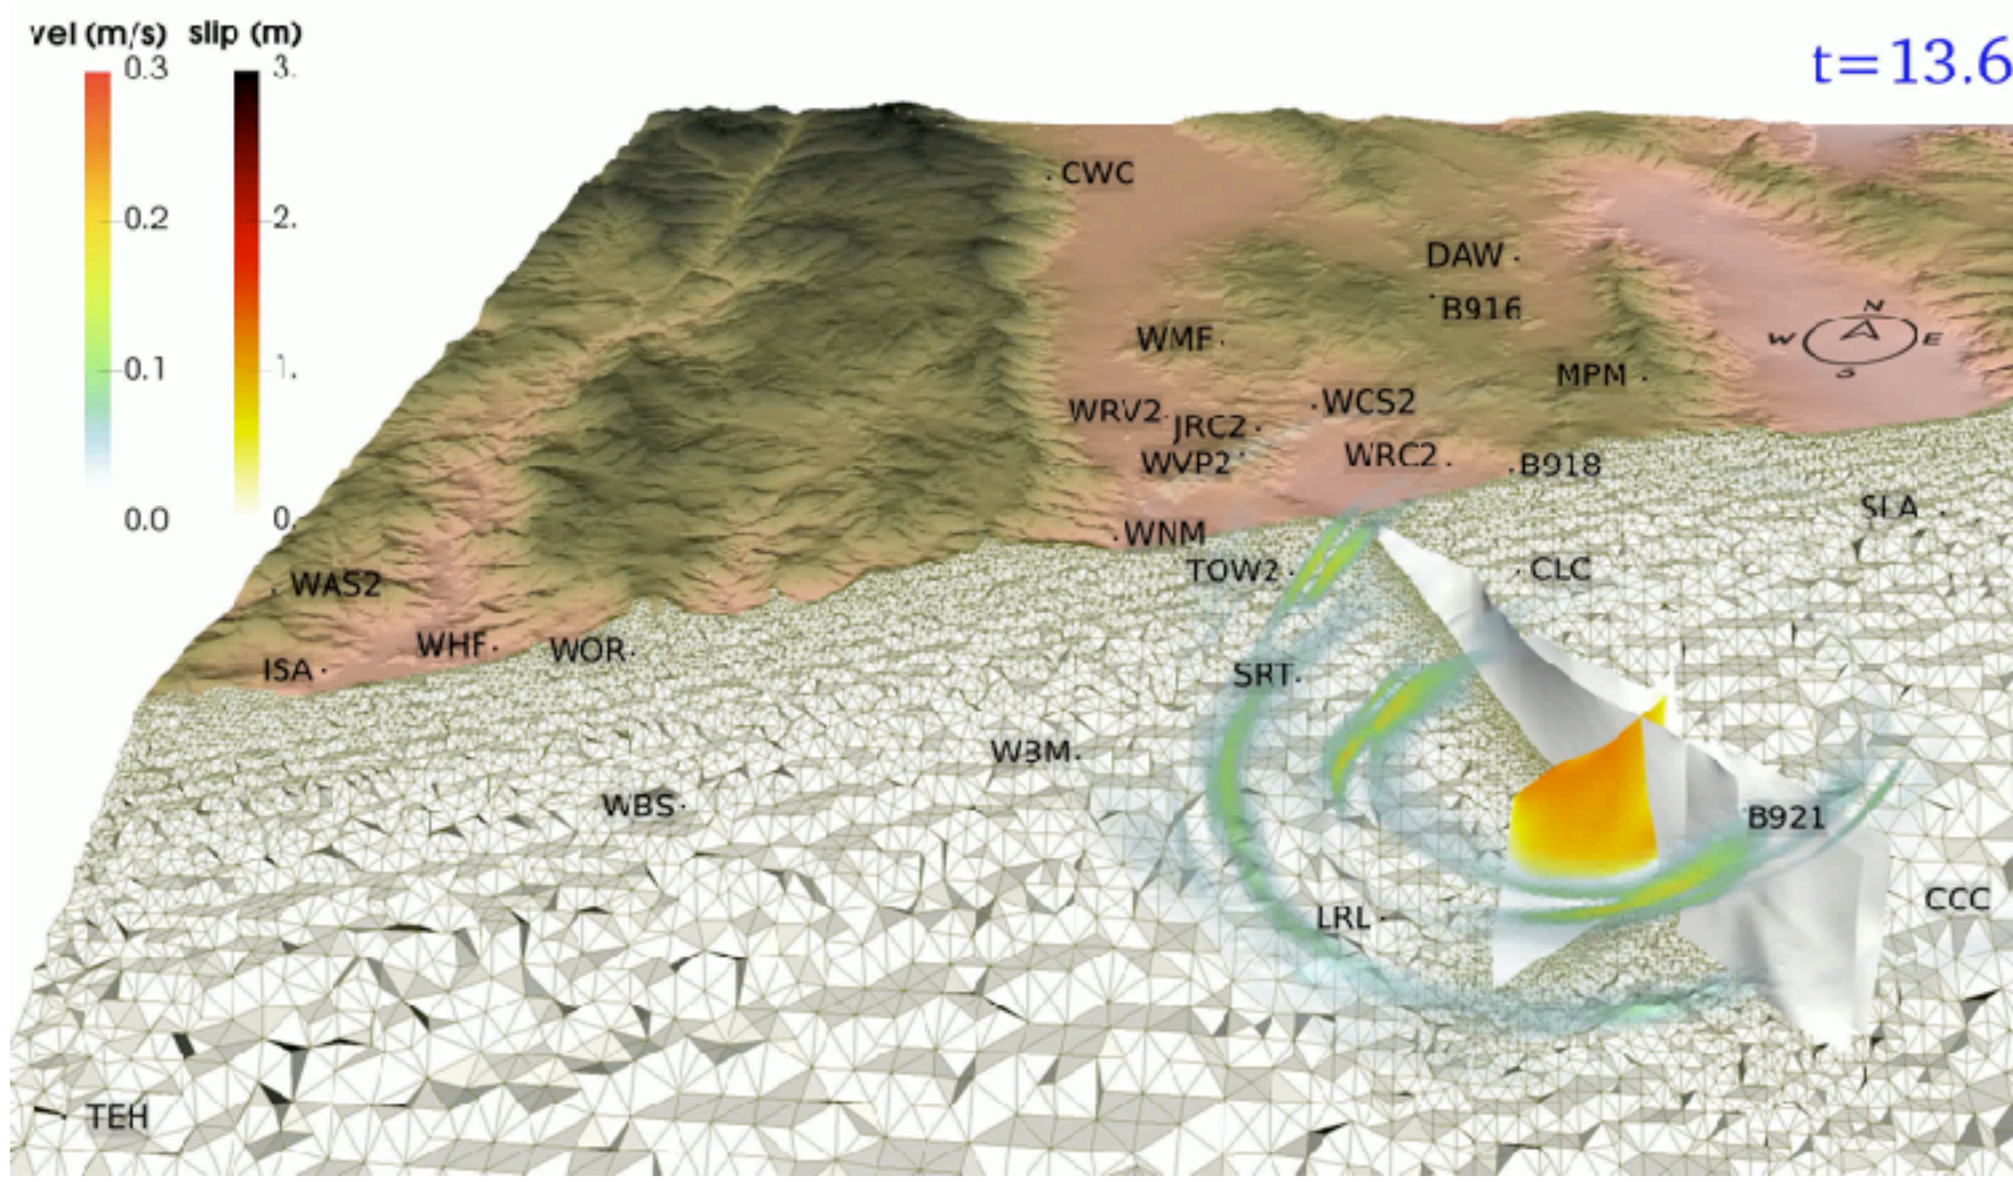

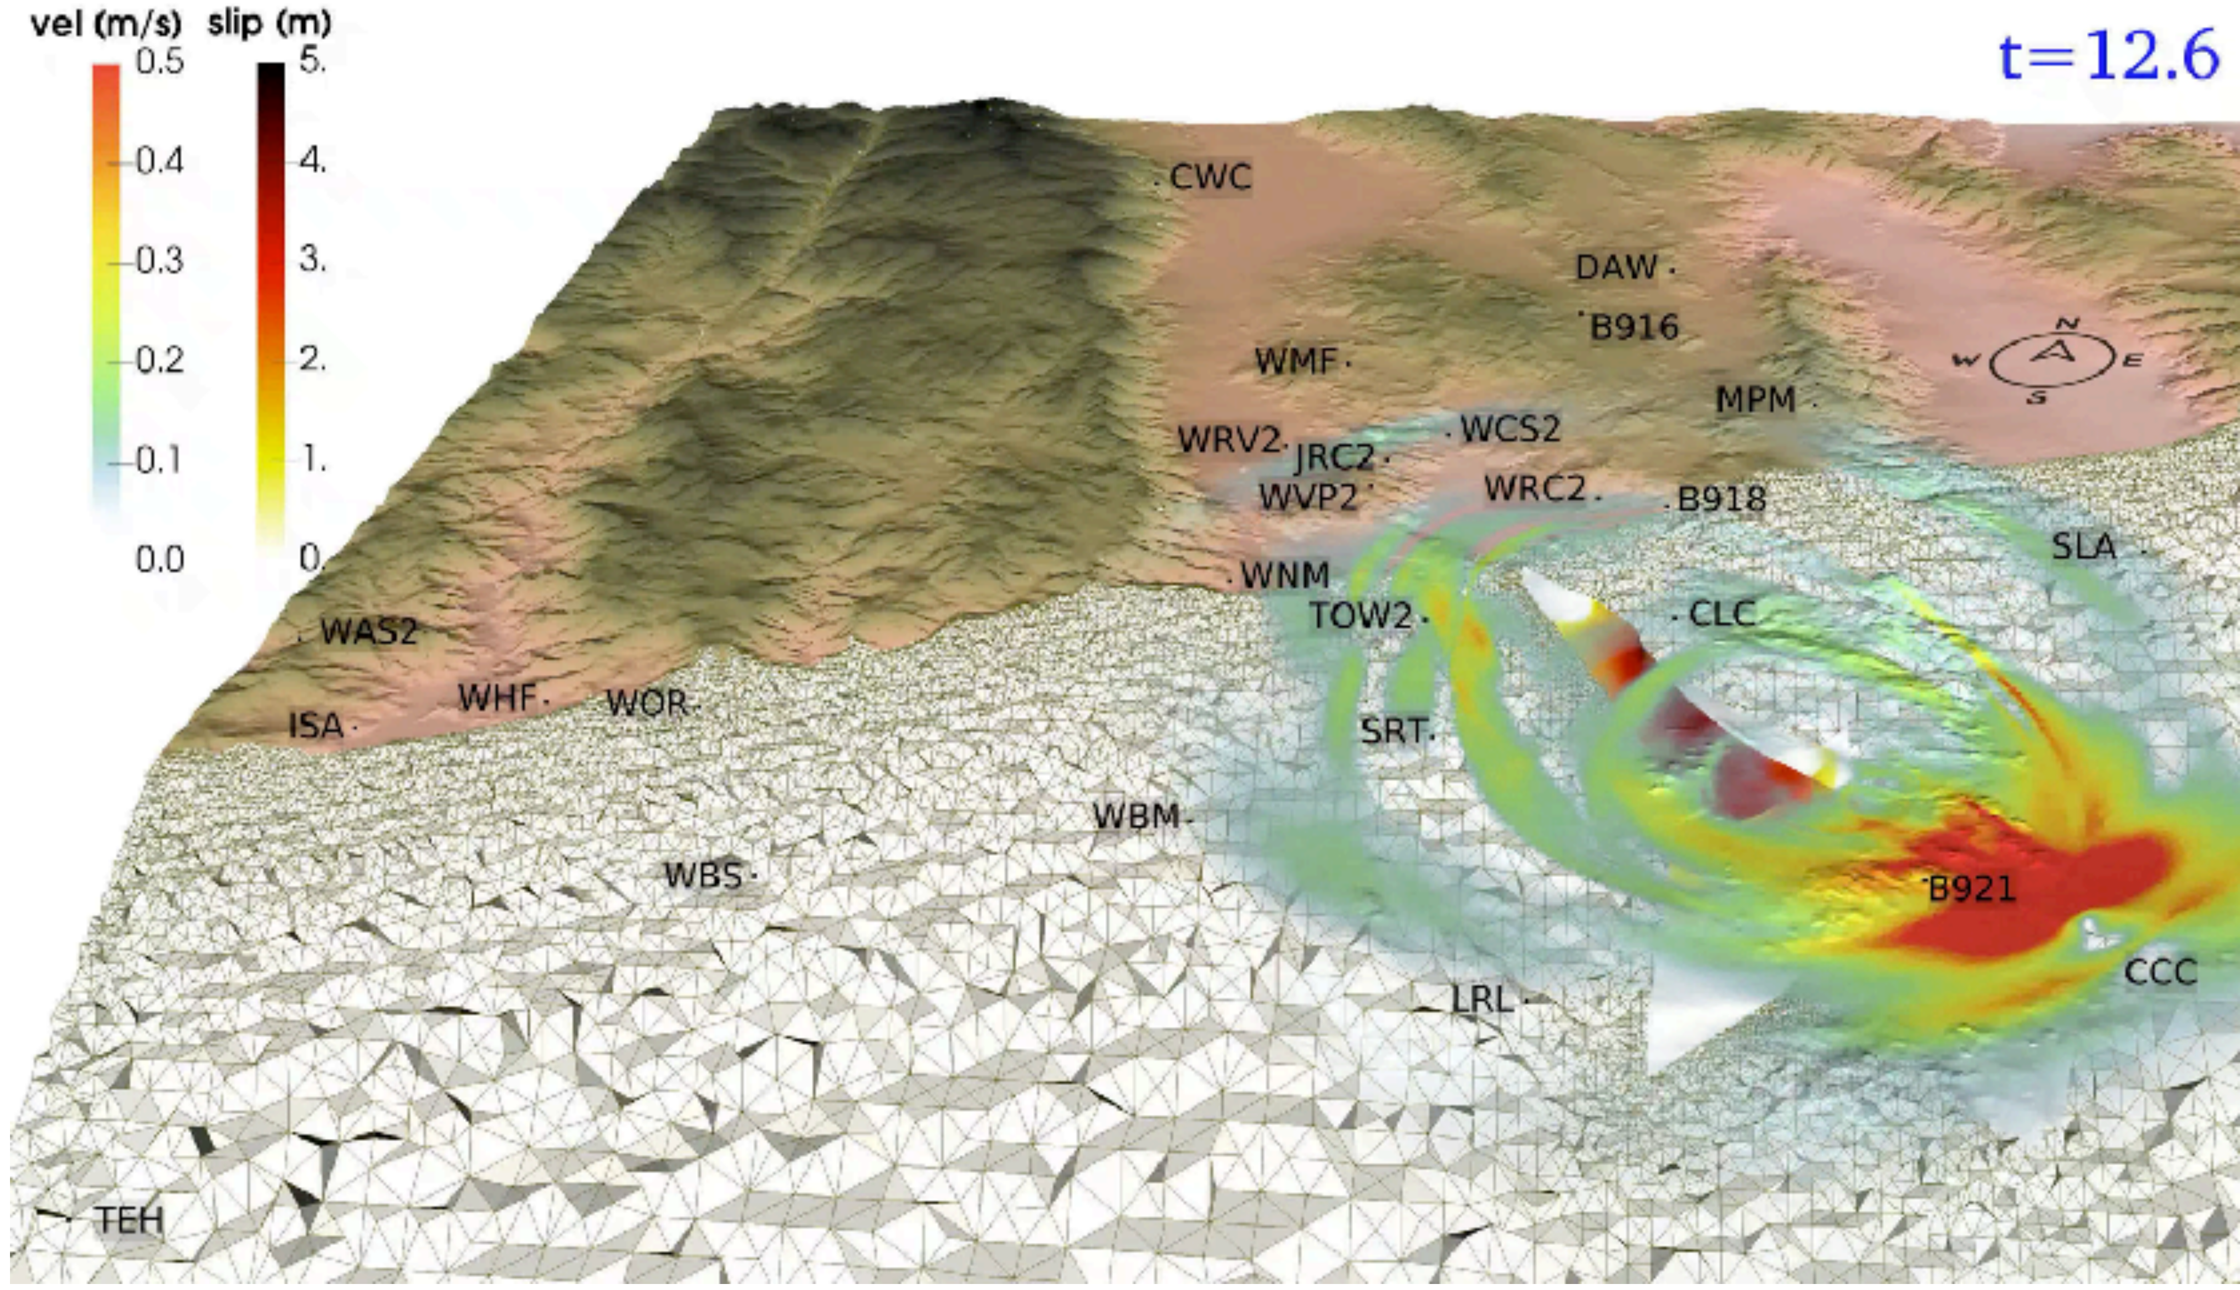

Observationally constrained 3D dynamic rupture scenarios of the 2018 Ridgecrest Earthquake Sequence (Taufiqurrahman et al, to appear)

Balancing HPC and geophysics

Breuer et al.,ISC14, Heinecke et al.,SC14
Breuer et al.,IEEE16, Heinecke et al.,SC16
Rettenberger et al., EASC16
Uphoff & Bader, HPCS'16
Uphoff et al., SC17
Wolf et al., ICCS'20
Uphoff & Bader, TOMS'20
Dorozhinskii & Bader, HPC Asia'21

"Geophysics" Version

- Fortran 90
- MPI parallelised
- Ascii based, serial I/O

ChEESE
Center of Excellence for Exascale in Solid Earth

Balancing HPC and geophysics

Breuer et al.,ISC14, Heinecke et al.,SC14
Breuer et al.,IEEE16, Heinecke et al.,SC16
Rettenberger et al., EASC16
Uphoff & Bader, HPCS'16
Uphoff et al., SC17
Wolf et al., ICCS'20
Uphoff & Bader, TOMS'20
Dorozhinskii & Bader, HPC Asia'21

Gordon Bell Prize Finalist, SC14

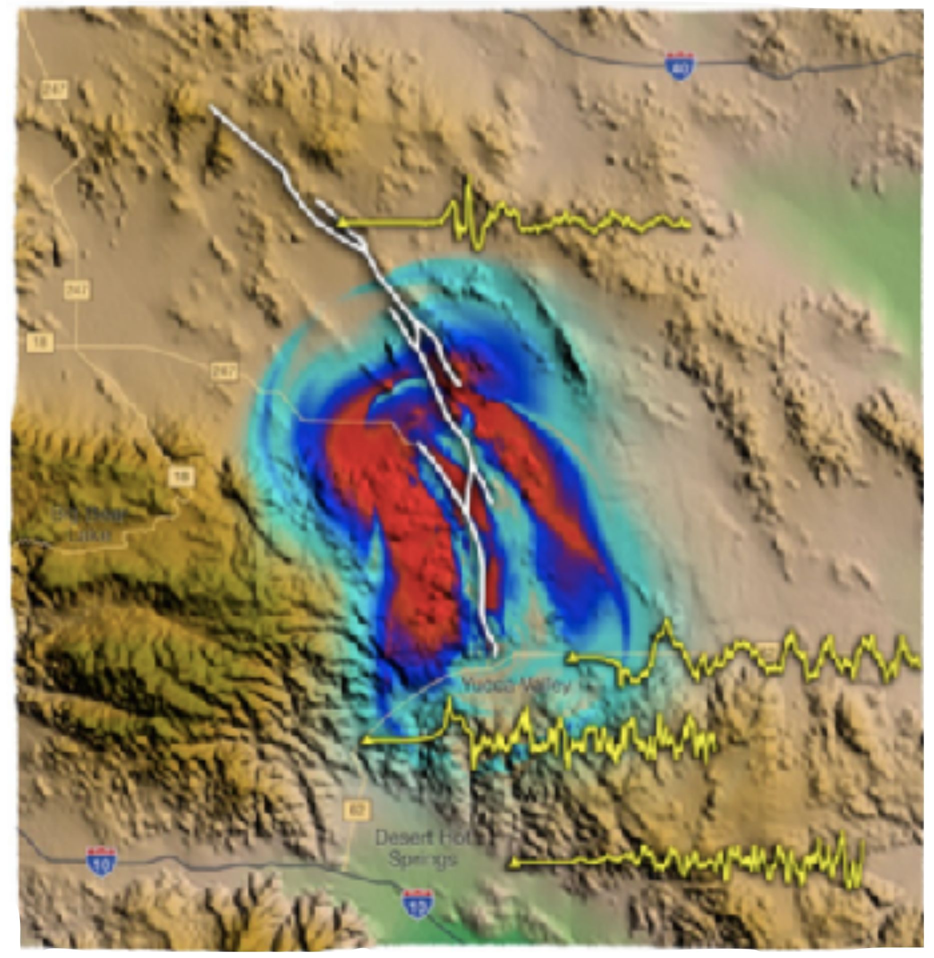

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 billion DoF, 200,000 time steps)	• Hybrid MPI+OpenMP parallelisation • Parallel I/O (HDF5, inc. mesh init.) • **Assembler-level DG kernels** • multi-physics off-load scheme for many-core architectures

- **> 1 PFlop/s performance**
- **90% parallel efficiency**
- **45% of peak performance**
- **5x-10x faster time-to-solution**
- **10x-100x bigger problems**

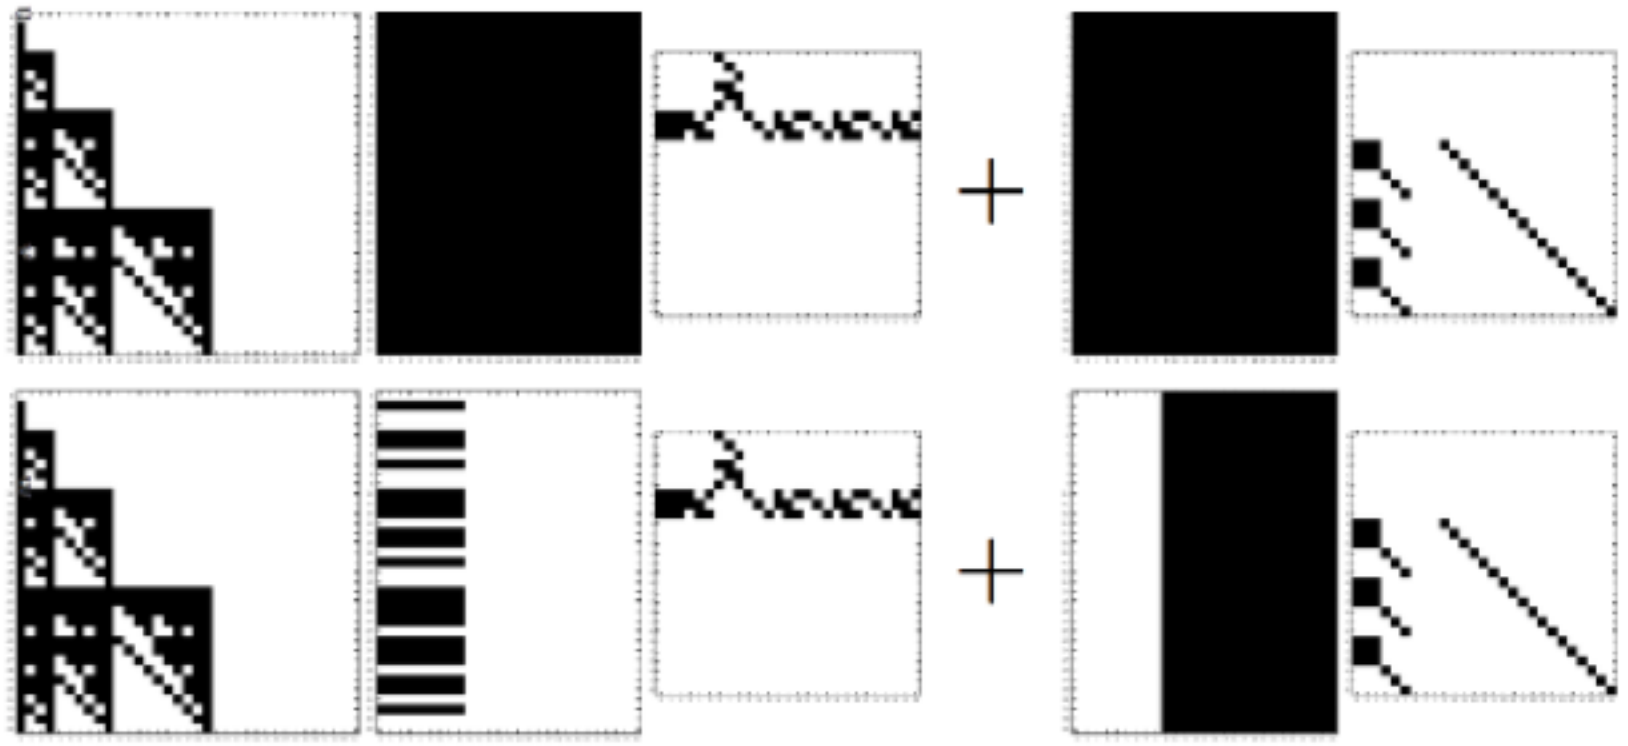

Partial kernel before (top) and after (bottom) removing irrelevant entries in matrix chain products

➡ **A code generator automatically detects and exploits sparse block patterns**
➡ **Hardware specific full "unrolling" and vectorization of all element operations**
➡ **Customised code for each matrix-matrix multiplication via the libxsmm back-end**

➡ **Efficiently exploits as of 2014 available hardware (AVX, MIC), reaching unto 8.6 PFLOPS on Tianhe-2 supercomputer**

ChEESE
Center of Excellence for Exascale in Solid Earth

Balancing HPC and geophysics

Breuer et al.,ISC14, Heinecke et al.,SC14
Breuer et al.,IEEE16, Heinecke et al.,SC16
Rettenberger et al., EASC16
Uphoff & Bader, HPCS'16
Uphoff et al., SC17
Wolf et al., ICCS'20
Uphoff & Bader, TOMS'20
Dorozhinskii & Bader, HPC Asia'21

Gordon Bell Prize Finalist, SC14

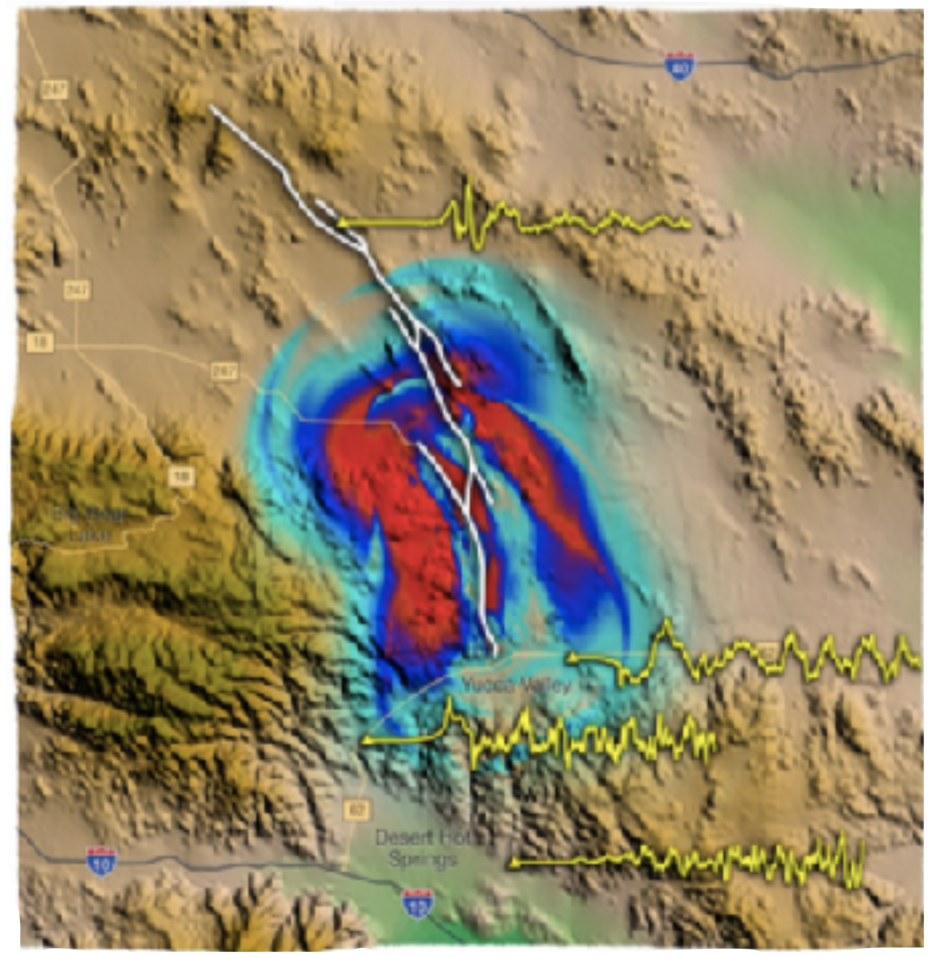

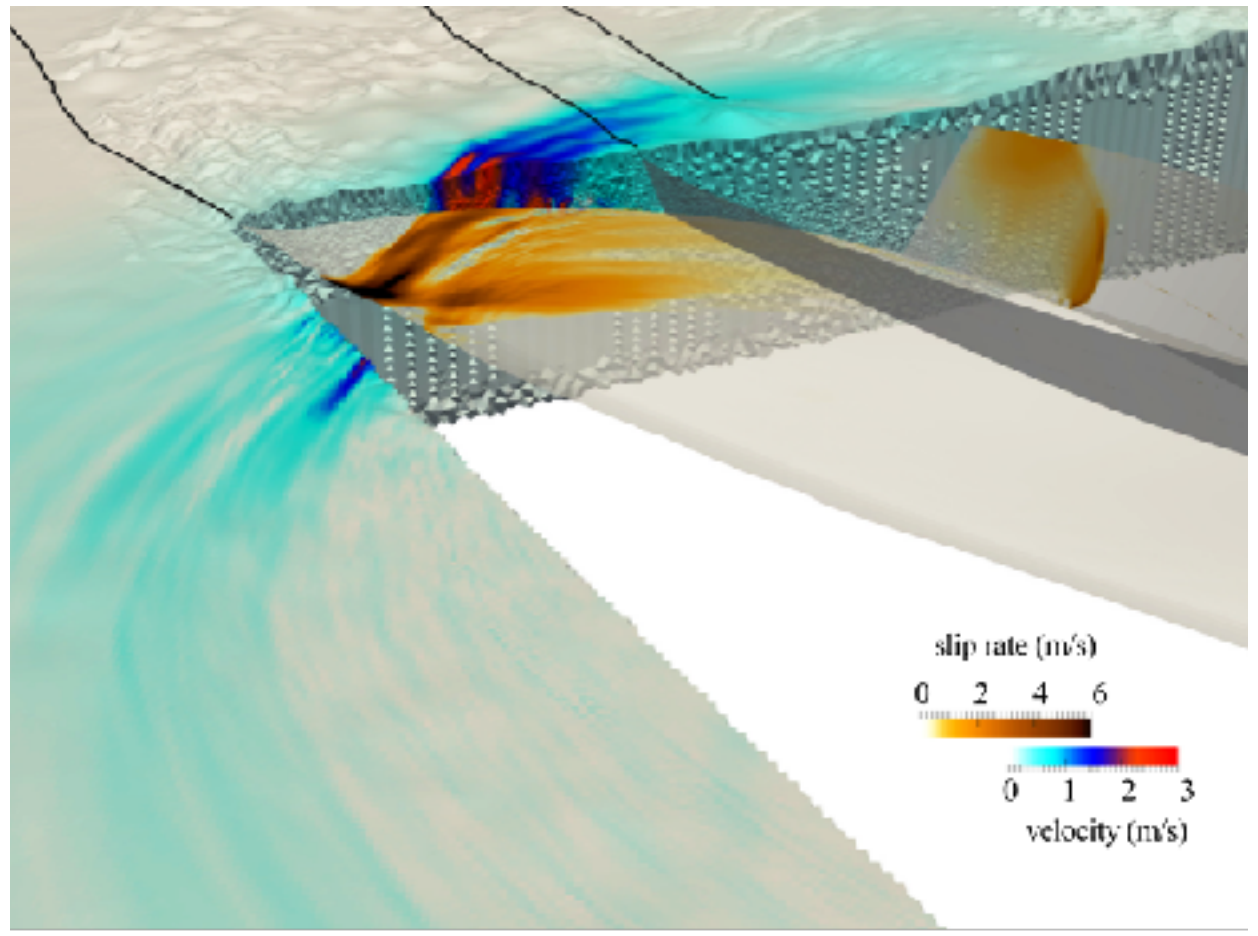

Best Paper Award, SC17

"Geophysics" Version

- Fortran 90
- MPI parallelised
- Ascii based, serial I/O

Landers scenario
(96 bi
200,000

- Hybrid MPI+OpenMP parallelisation
- Parallel I/O (HDF5, inc. mesh init.)
- G kernels
- oad scheme for
- ctures

- > 1 PFlop/s performance
- 90% parallel efficiency
- 45% of peak performance
- 5x-10x faster time-to-solution
- 10x-100x bigger problems

➡ **Multi-scale and multi-physics modelling is routinely feasible**
(few kCPUh per high resolution forward simulation)

Sumatra scenario
(111 billion DoF,
3,300,000 time steps)

- l time stepping
- Code generator also for advanced PDE's as viscoelastic attunation
- Asagi (XDMF)-geoinformation server
- Asynchronous input/output
- Overlaping computation and communication

- Optimized for Intel KNL
- Speed up of 14x
- 14 hours compared to almost 8 days for Sumatra scenario on SuperMuc2

The “grand challenge” of meshing

- Community standard 1) **Hexahedral** meshes - may consume weeks in mesh generation **,** is **limited** for complex geometries (external / internal boundary conditions), common tool: TRELIS
- Community standard 2) **Unstructured tetrahedral** meshes - allows automatised meshing and complex internal/external boundary conditions, however are numerically challenging (sliver elements), common tools: GMSH, SIMMODELER

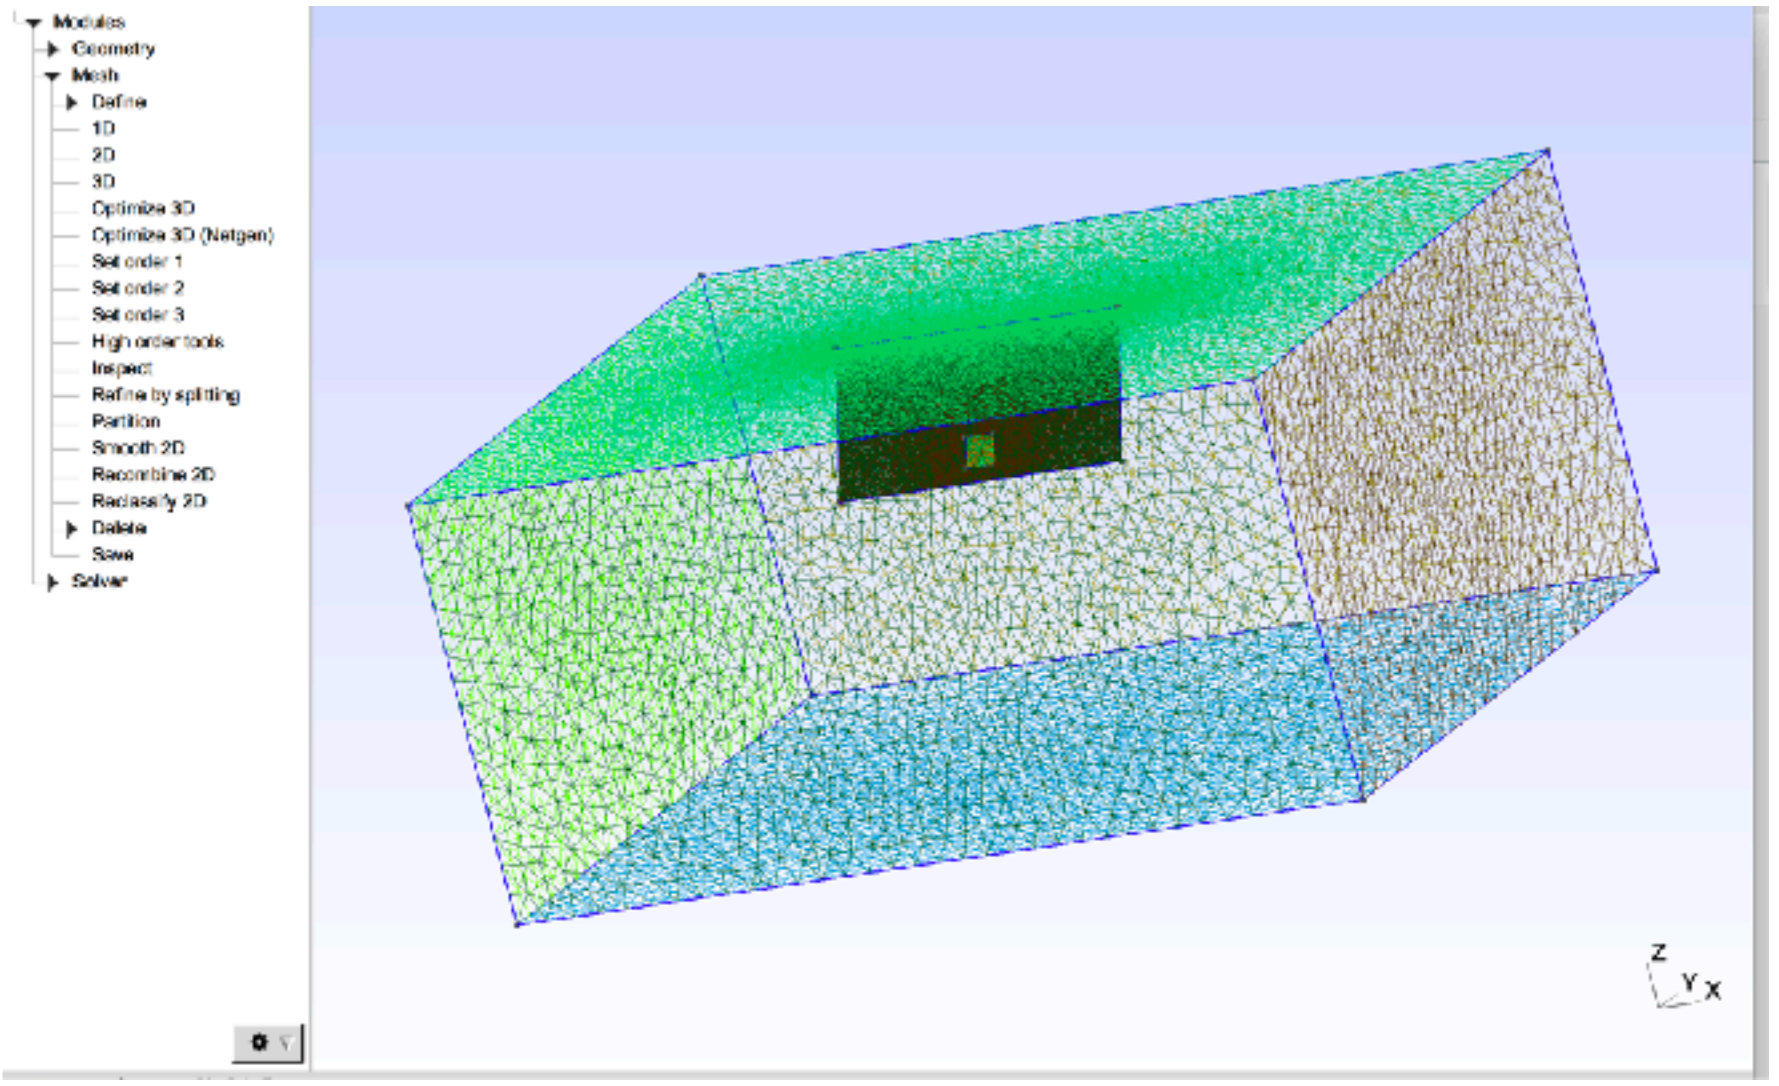

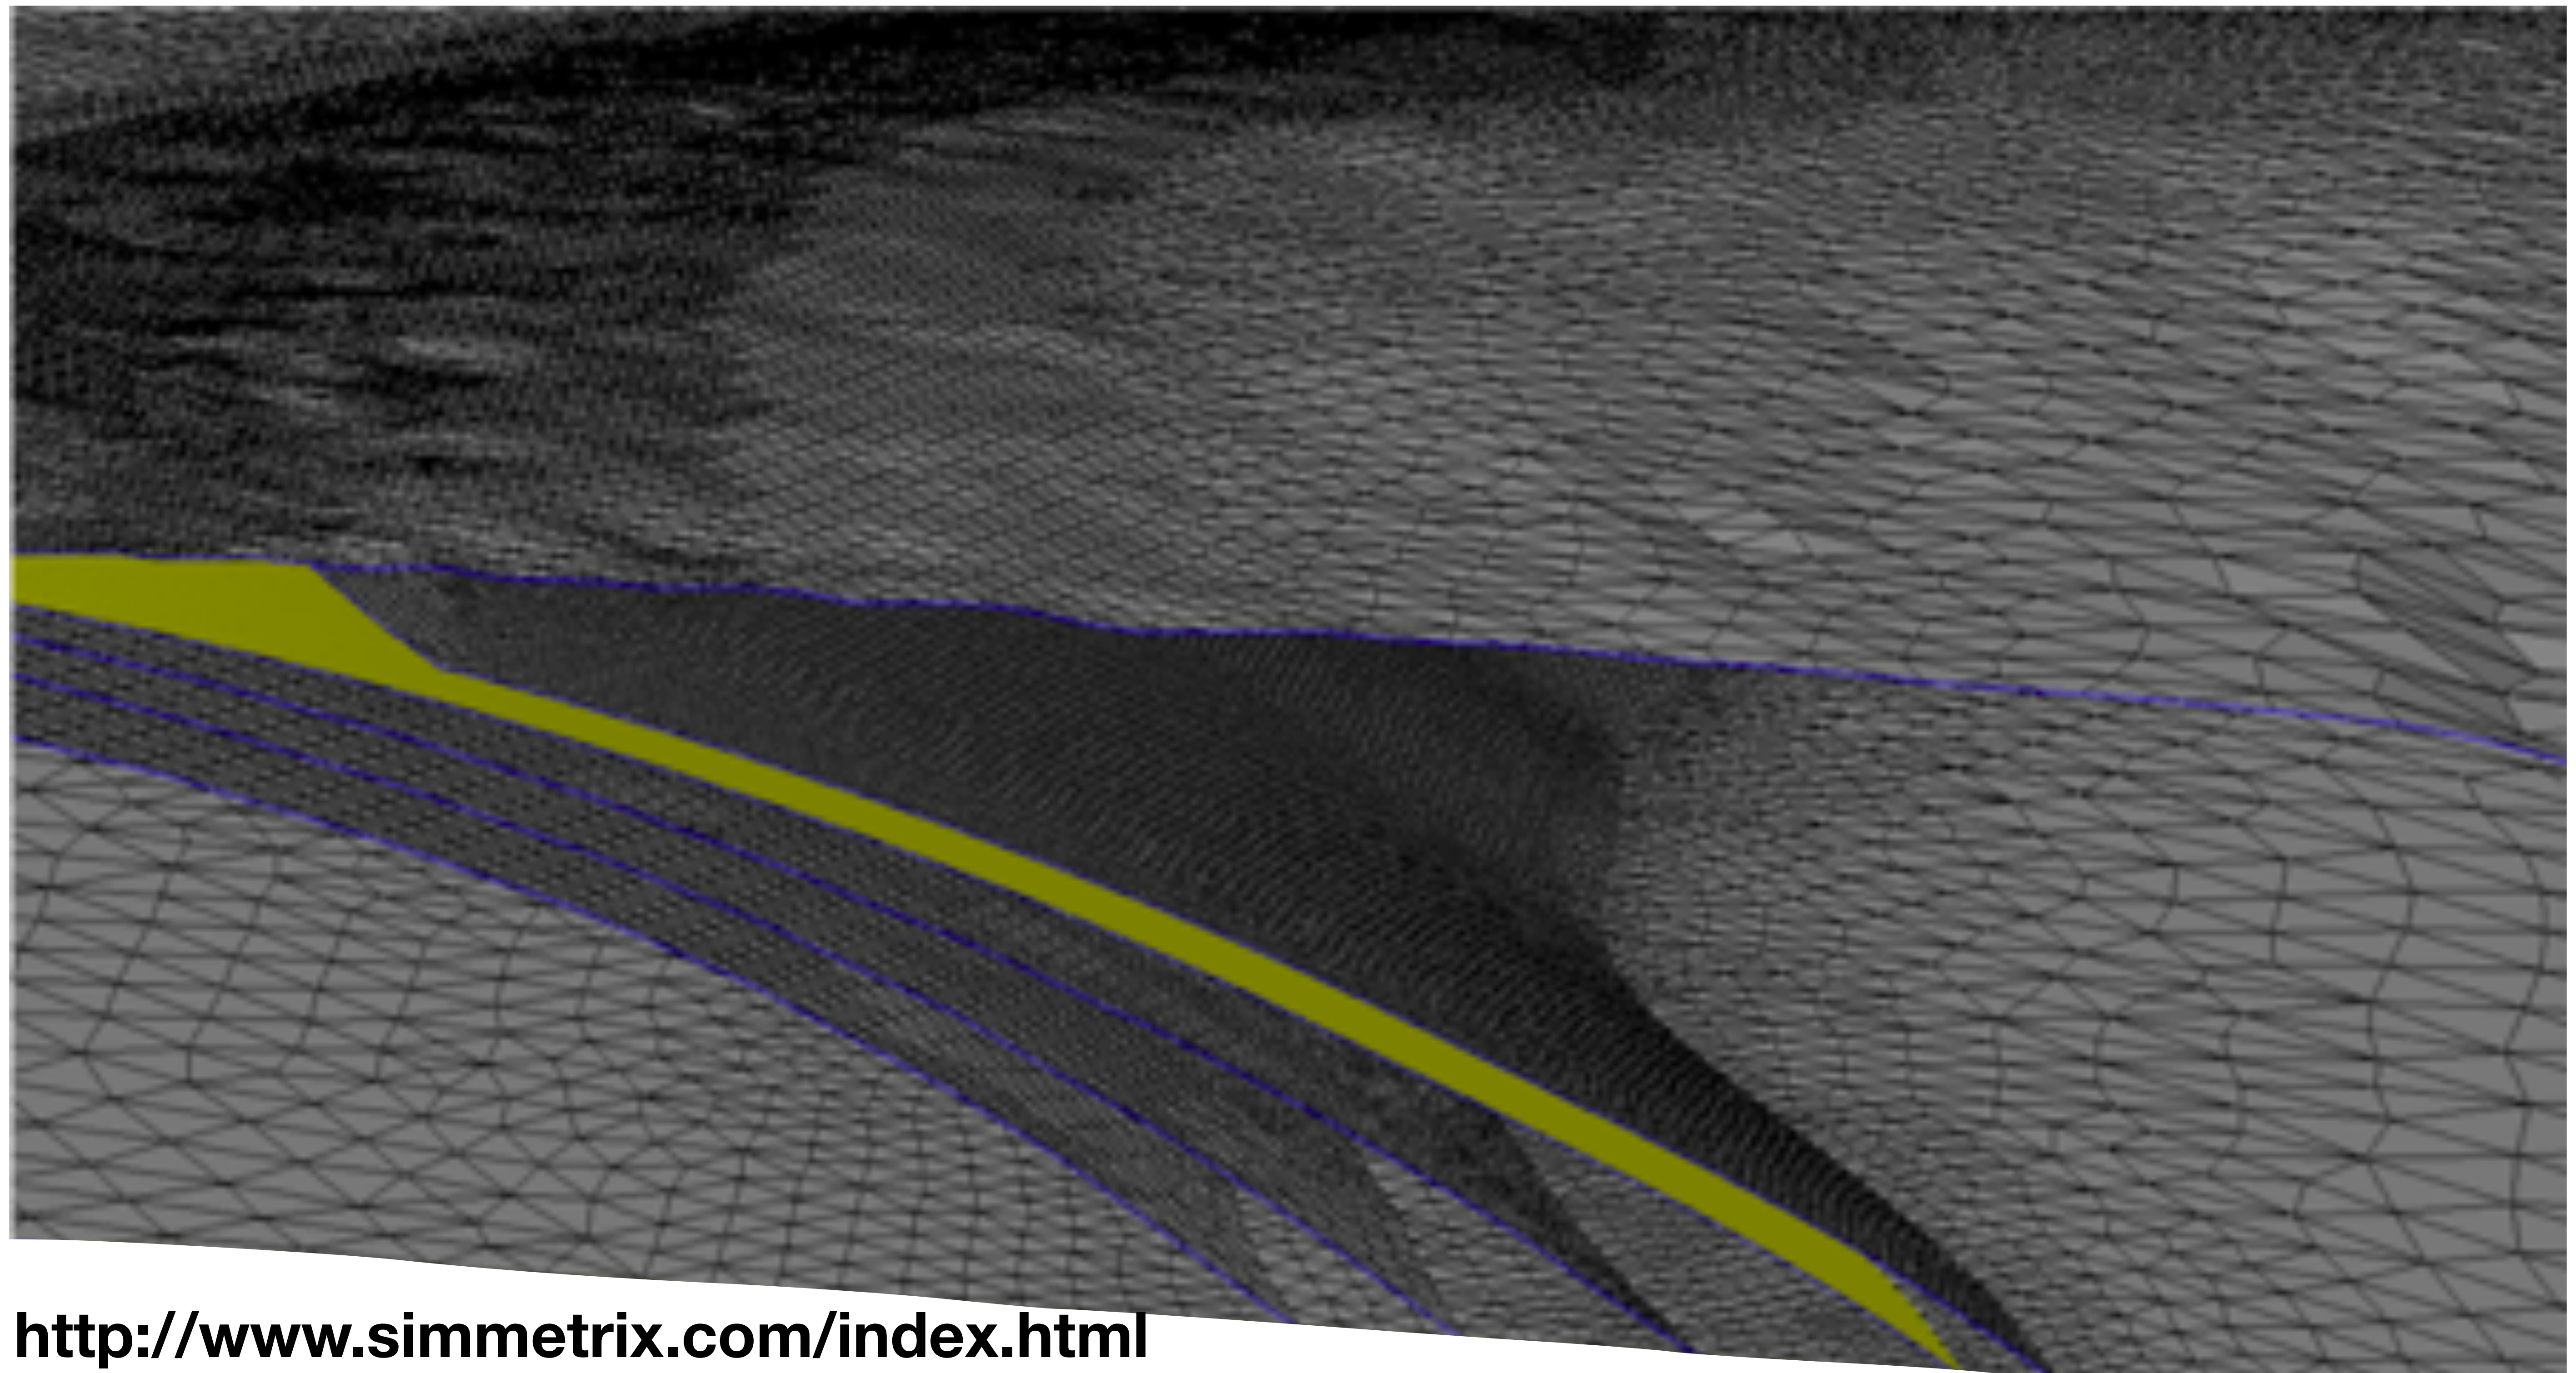

http://www.simmetrix.com/index.html

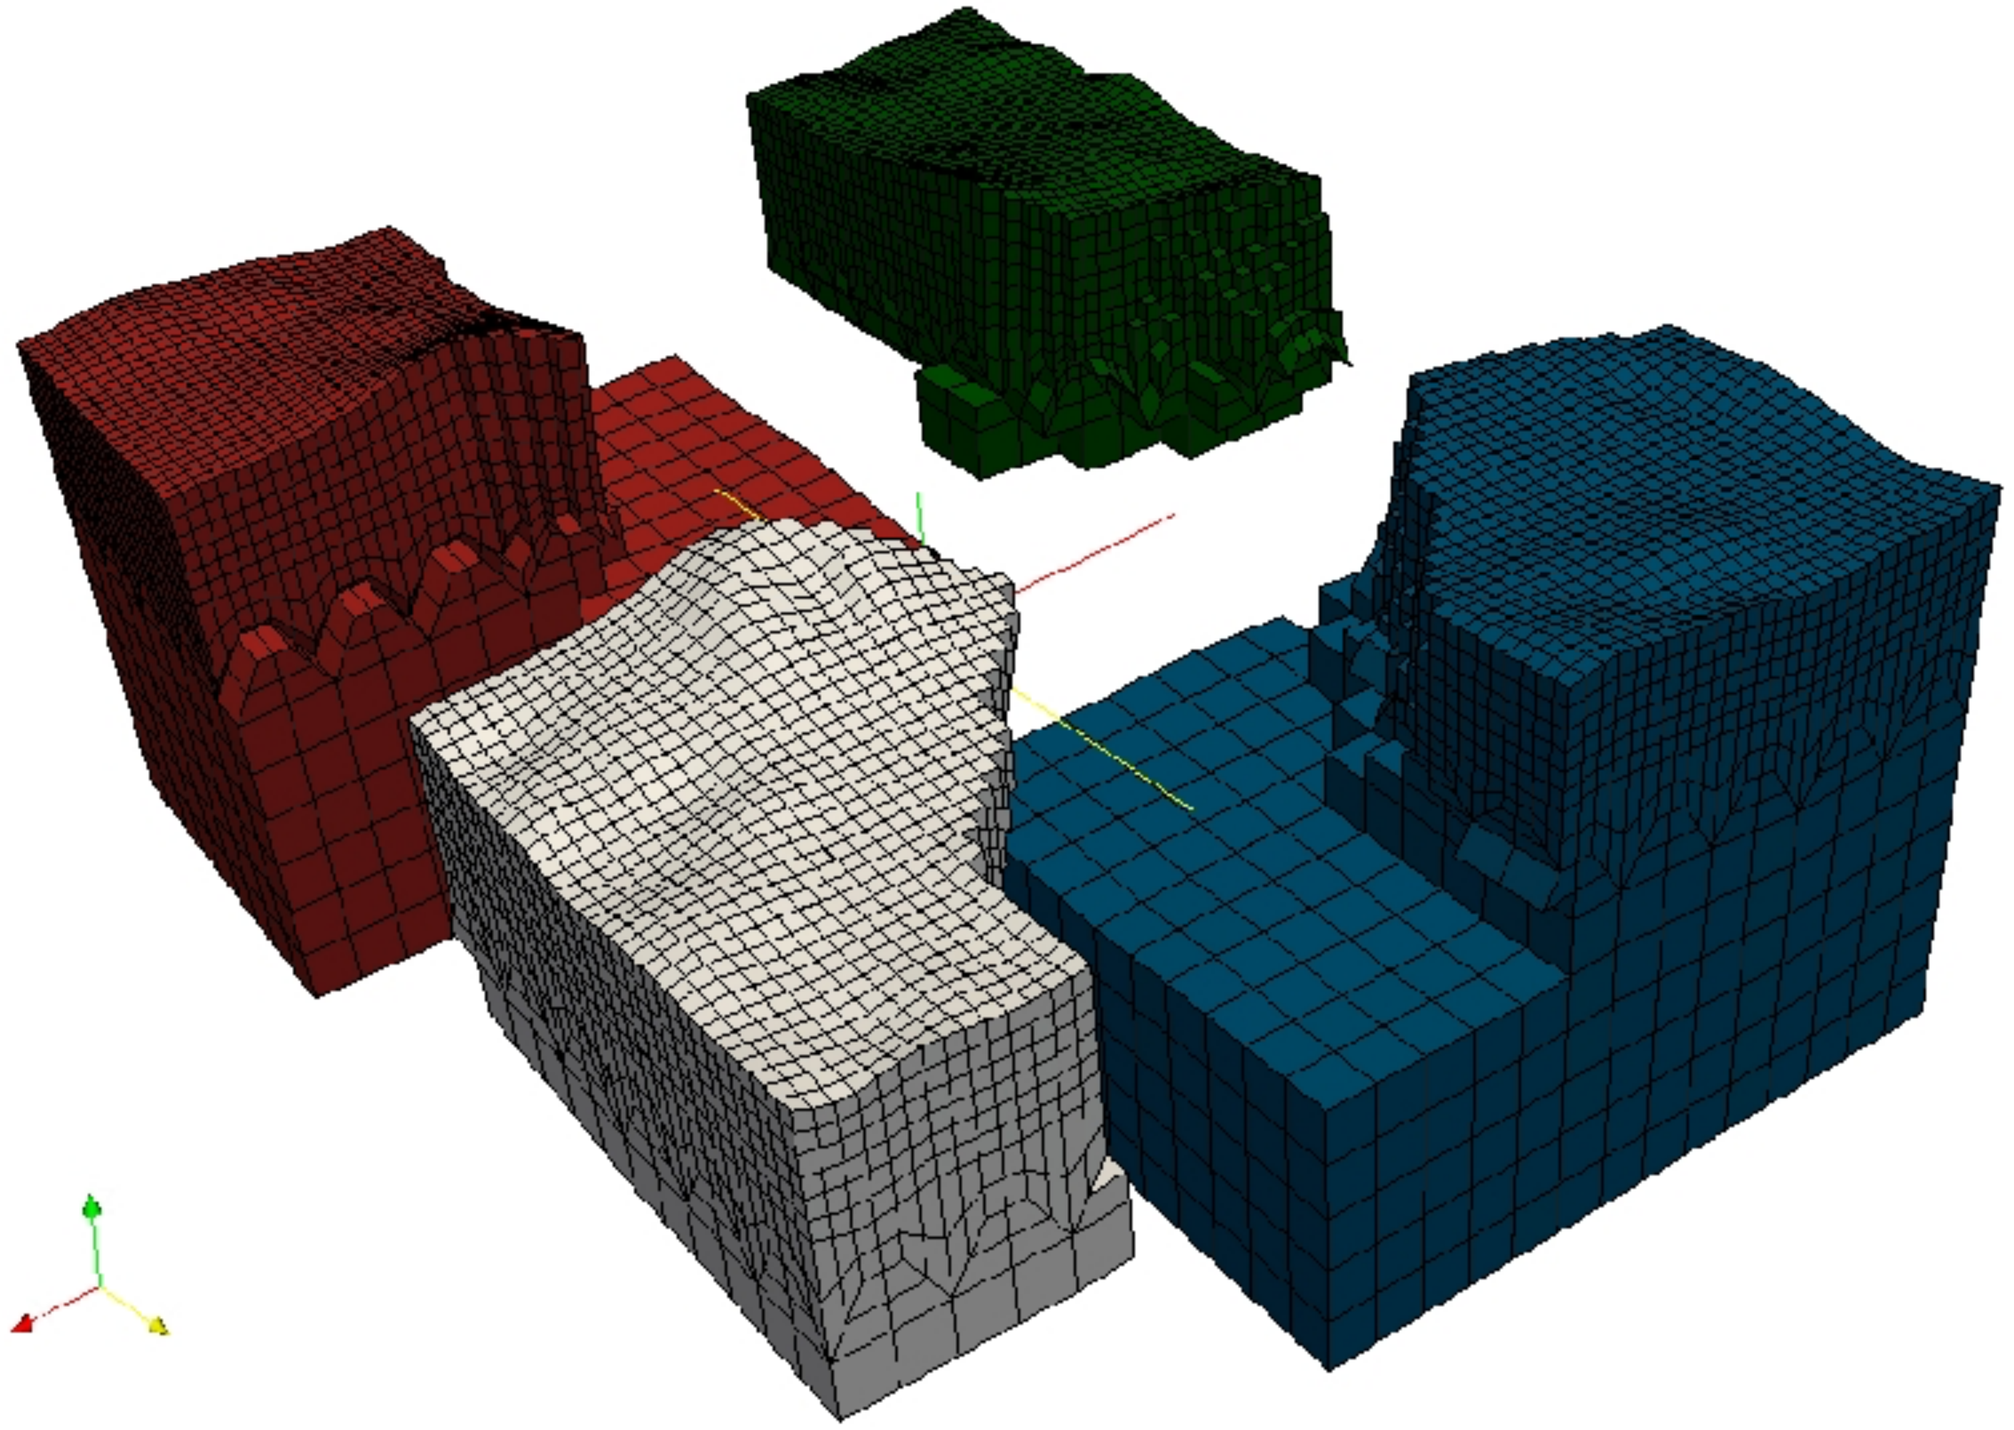

The “grand challenge” of meshing

- **Gmsh** (http://gmsh.info, open source) for most simple geometries and every-day mesh sizes, many tutorials, limited in terms of geometry
- **Simmodeler** (Simmetrix, **free for academics**) for large meshes / complex geometries: customised GUI for SeisSol, pumgen library for parallel meshing on Clusters
- Mesh is provided in parallel data format - code does internal **partitioning**

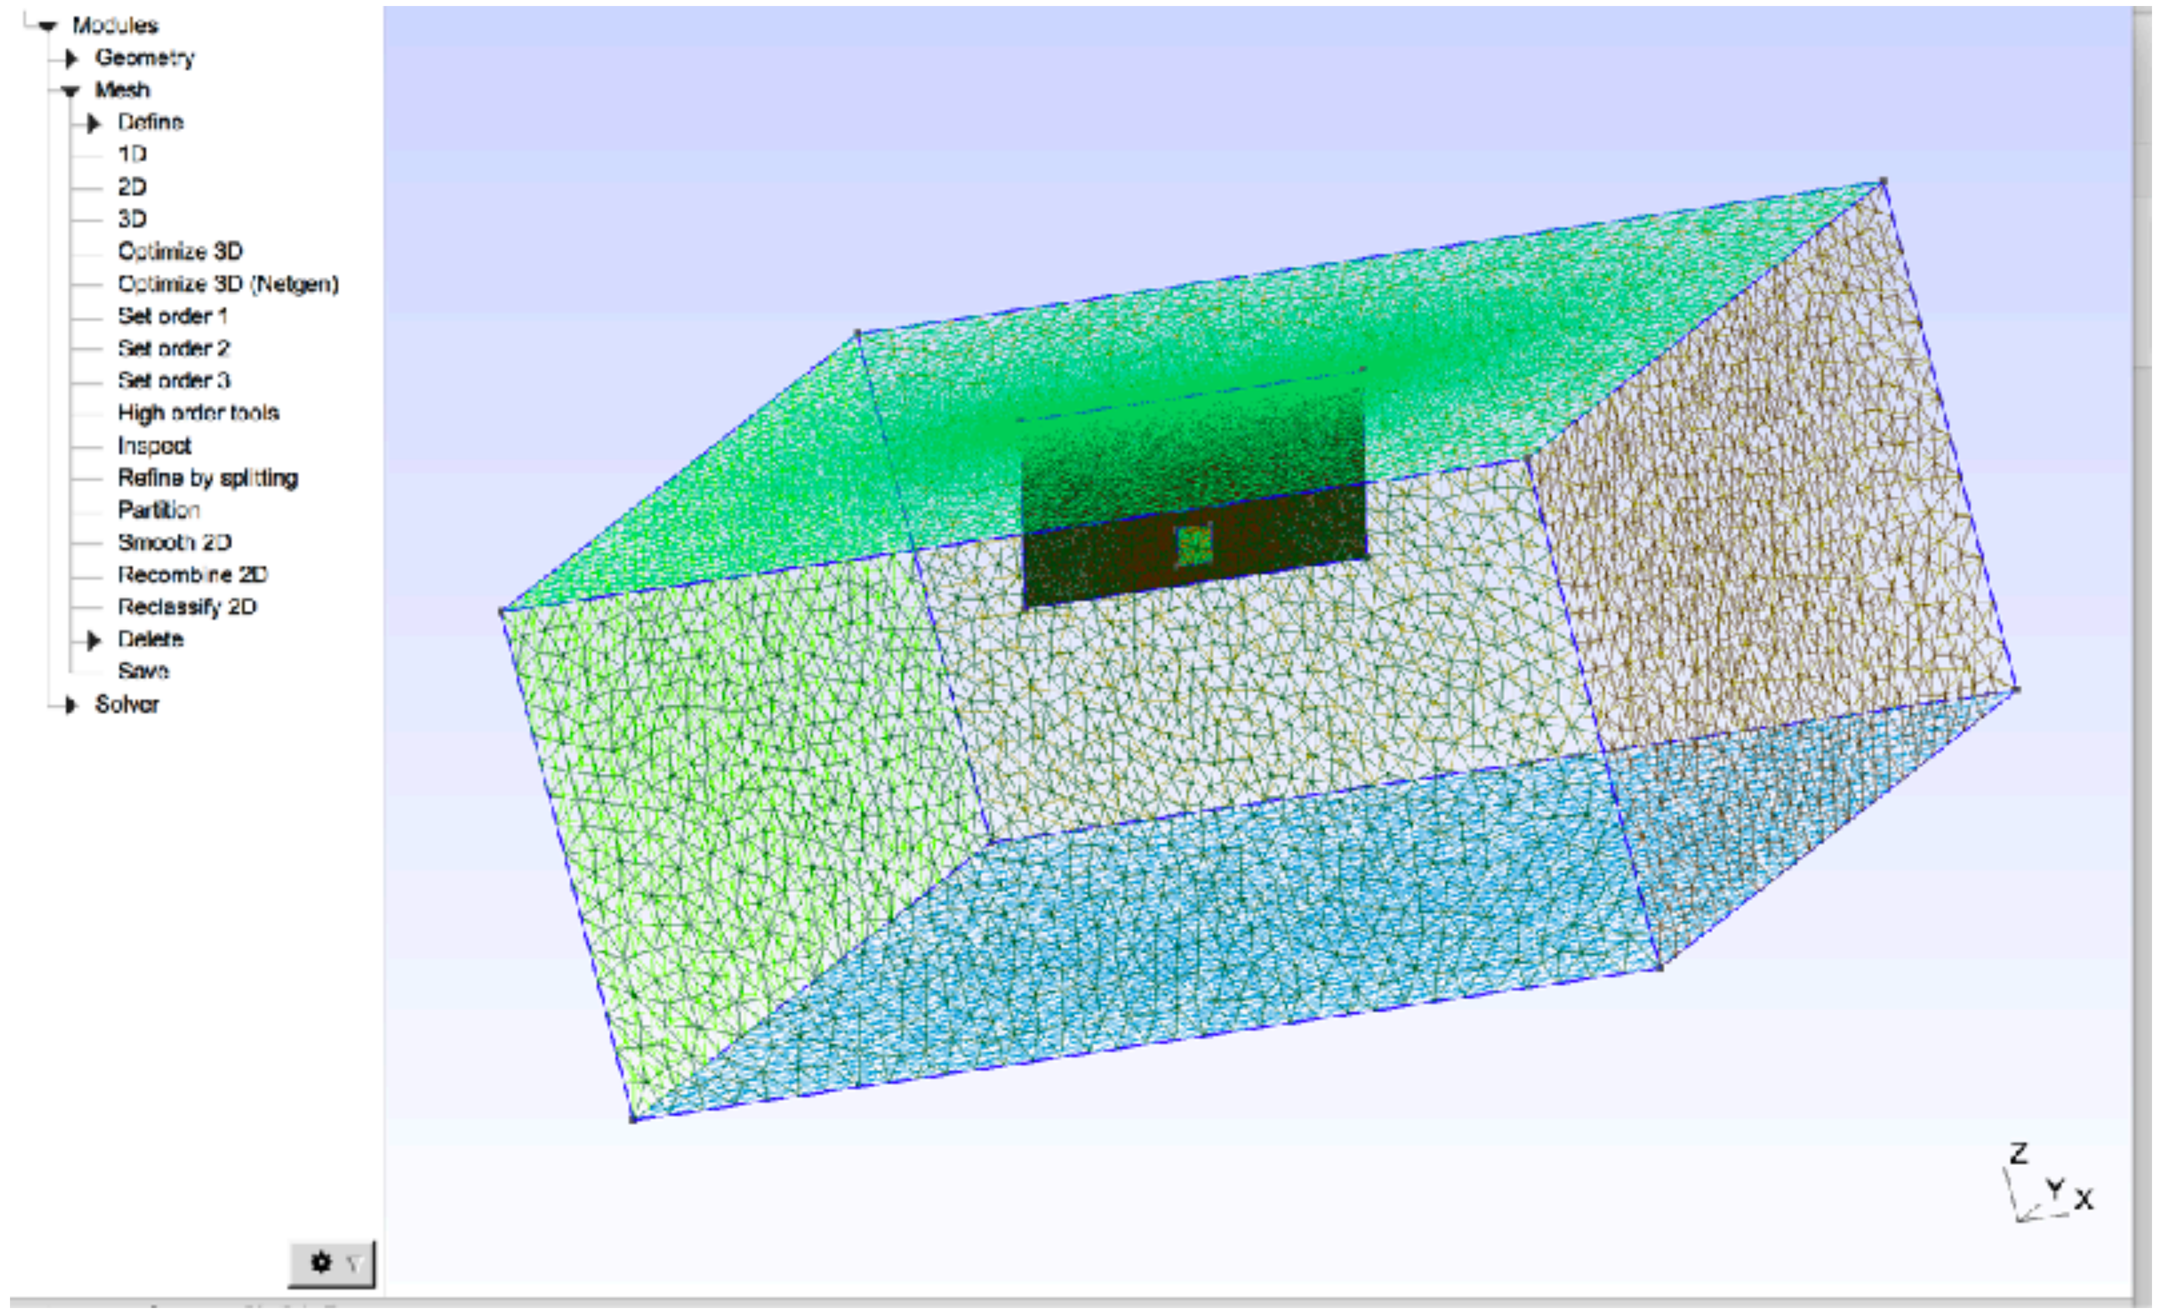

Gmsh interface for example geometry and 2D mesh

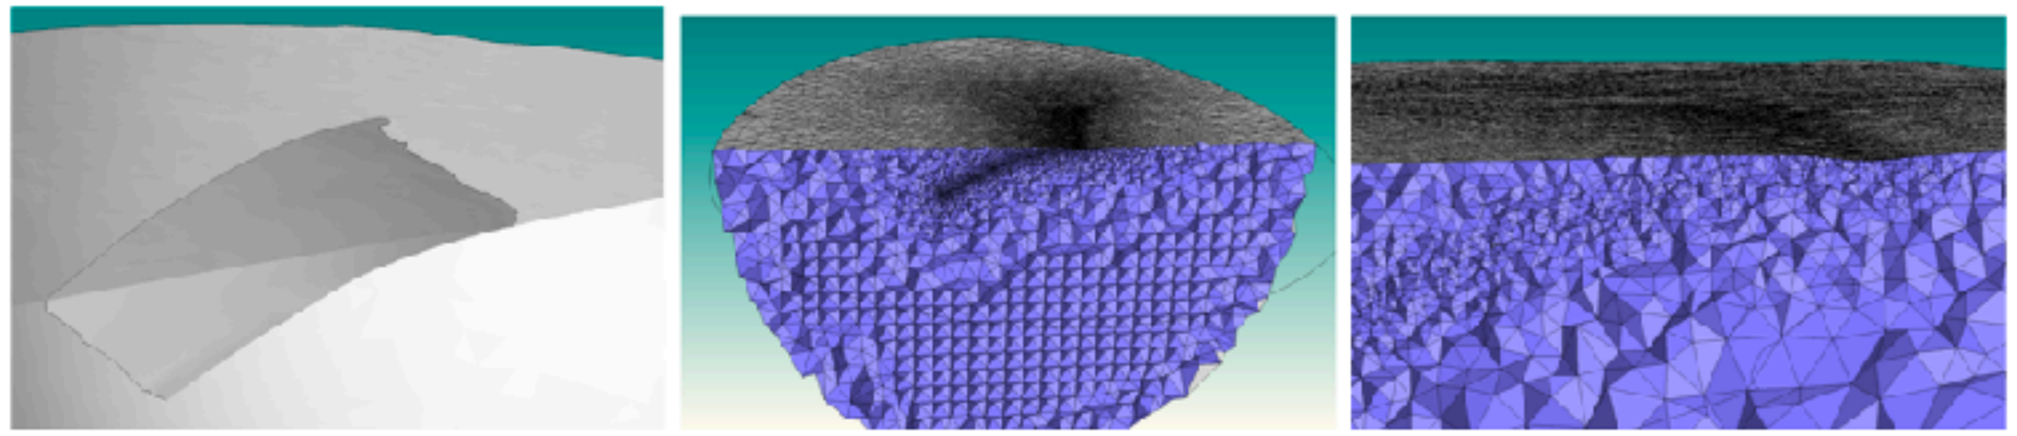

http://www.simmetrix.com/index.html

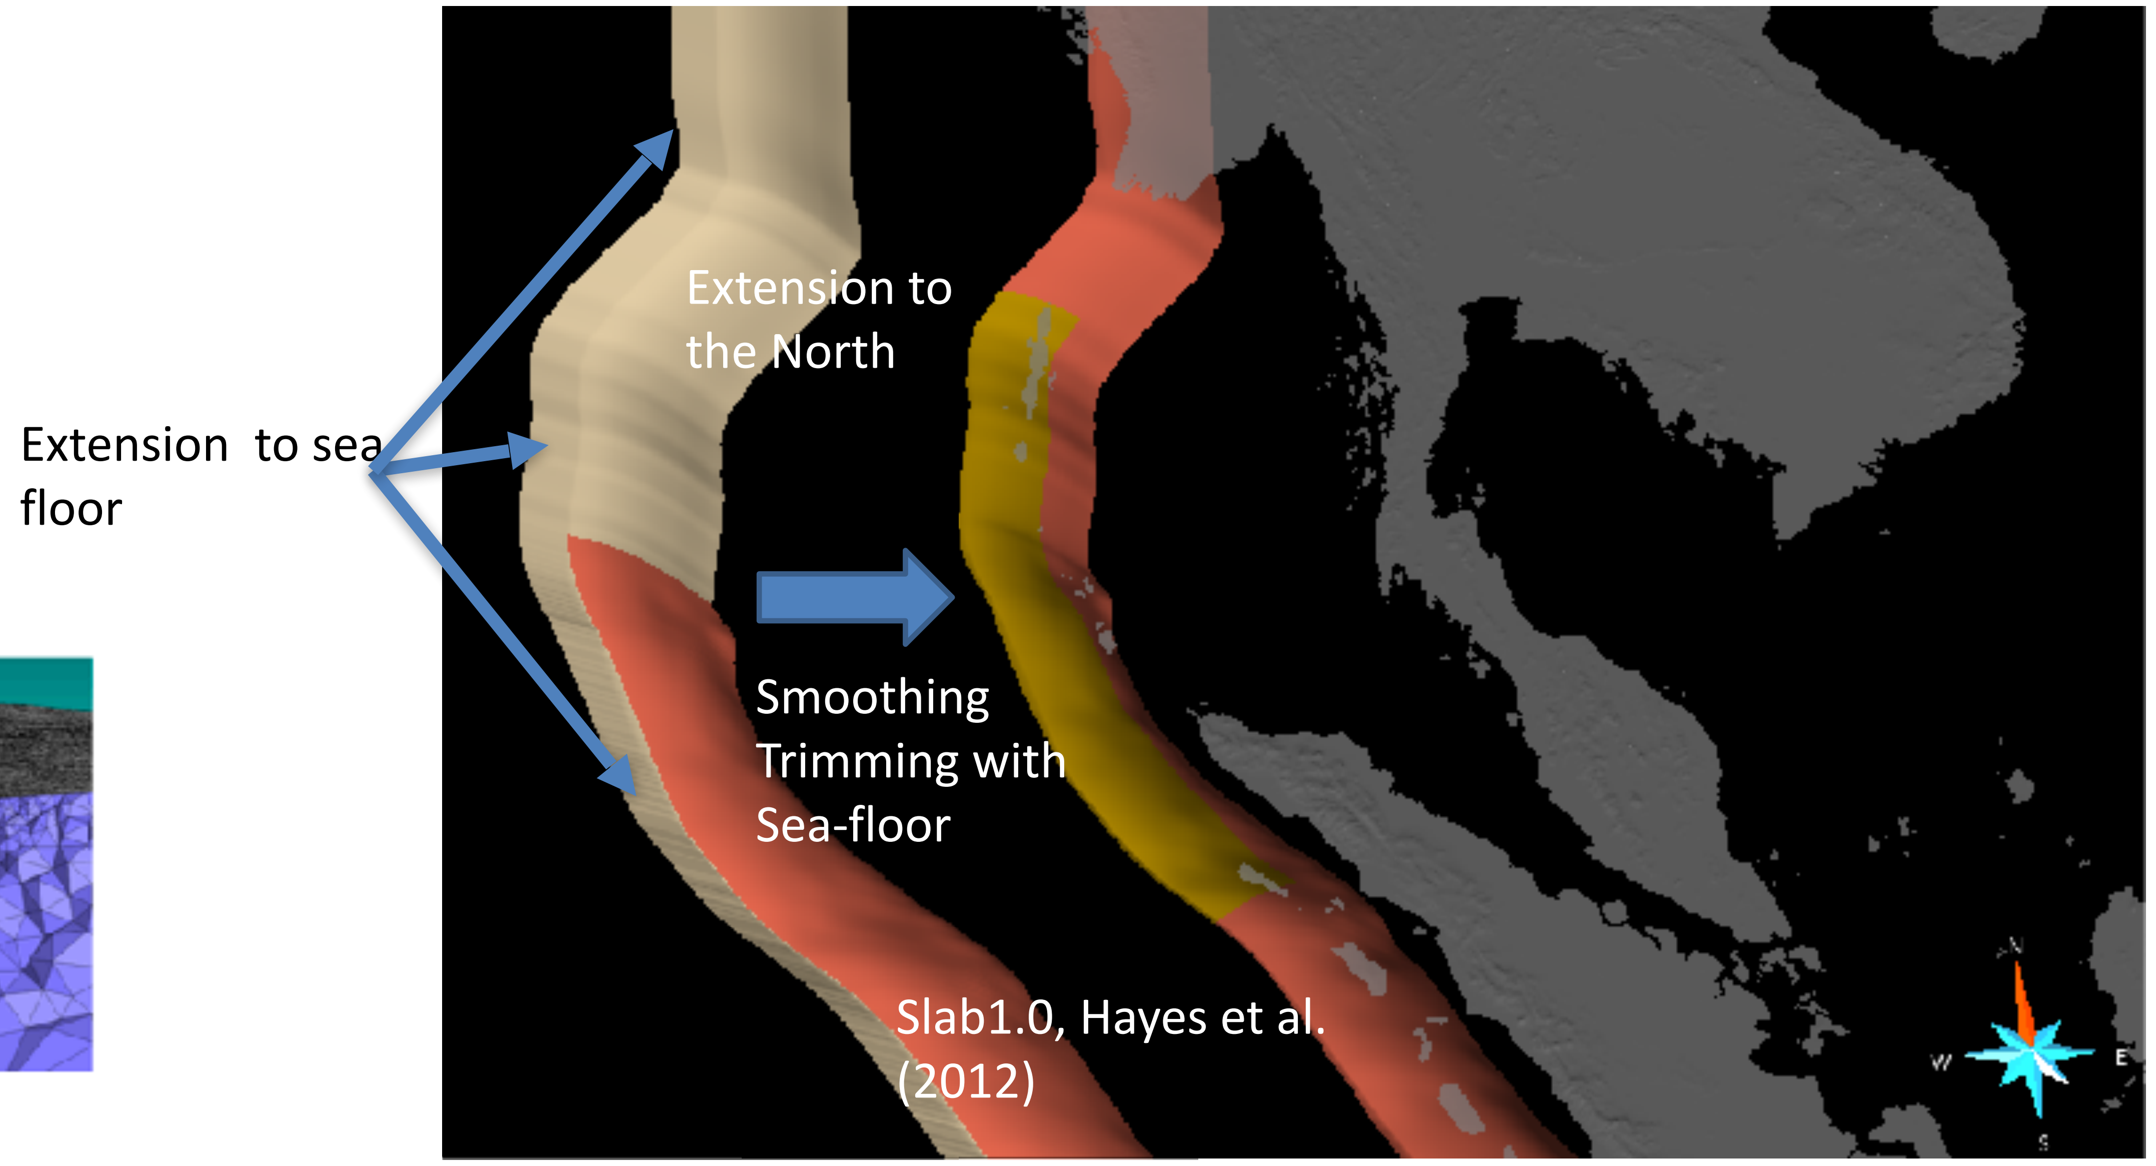

GoCAD interface for complex geometries

ChEESE
Center of Excellence for Exascale in Solid Earth

The “grand challenge” of meshing

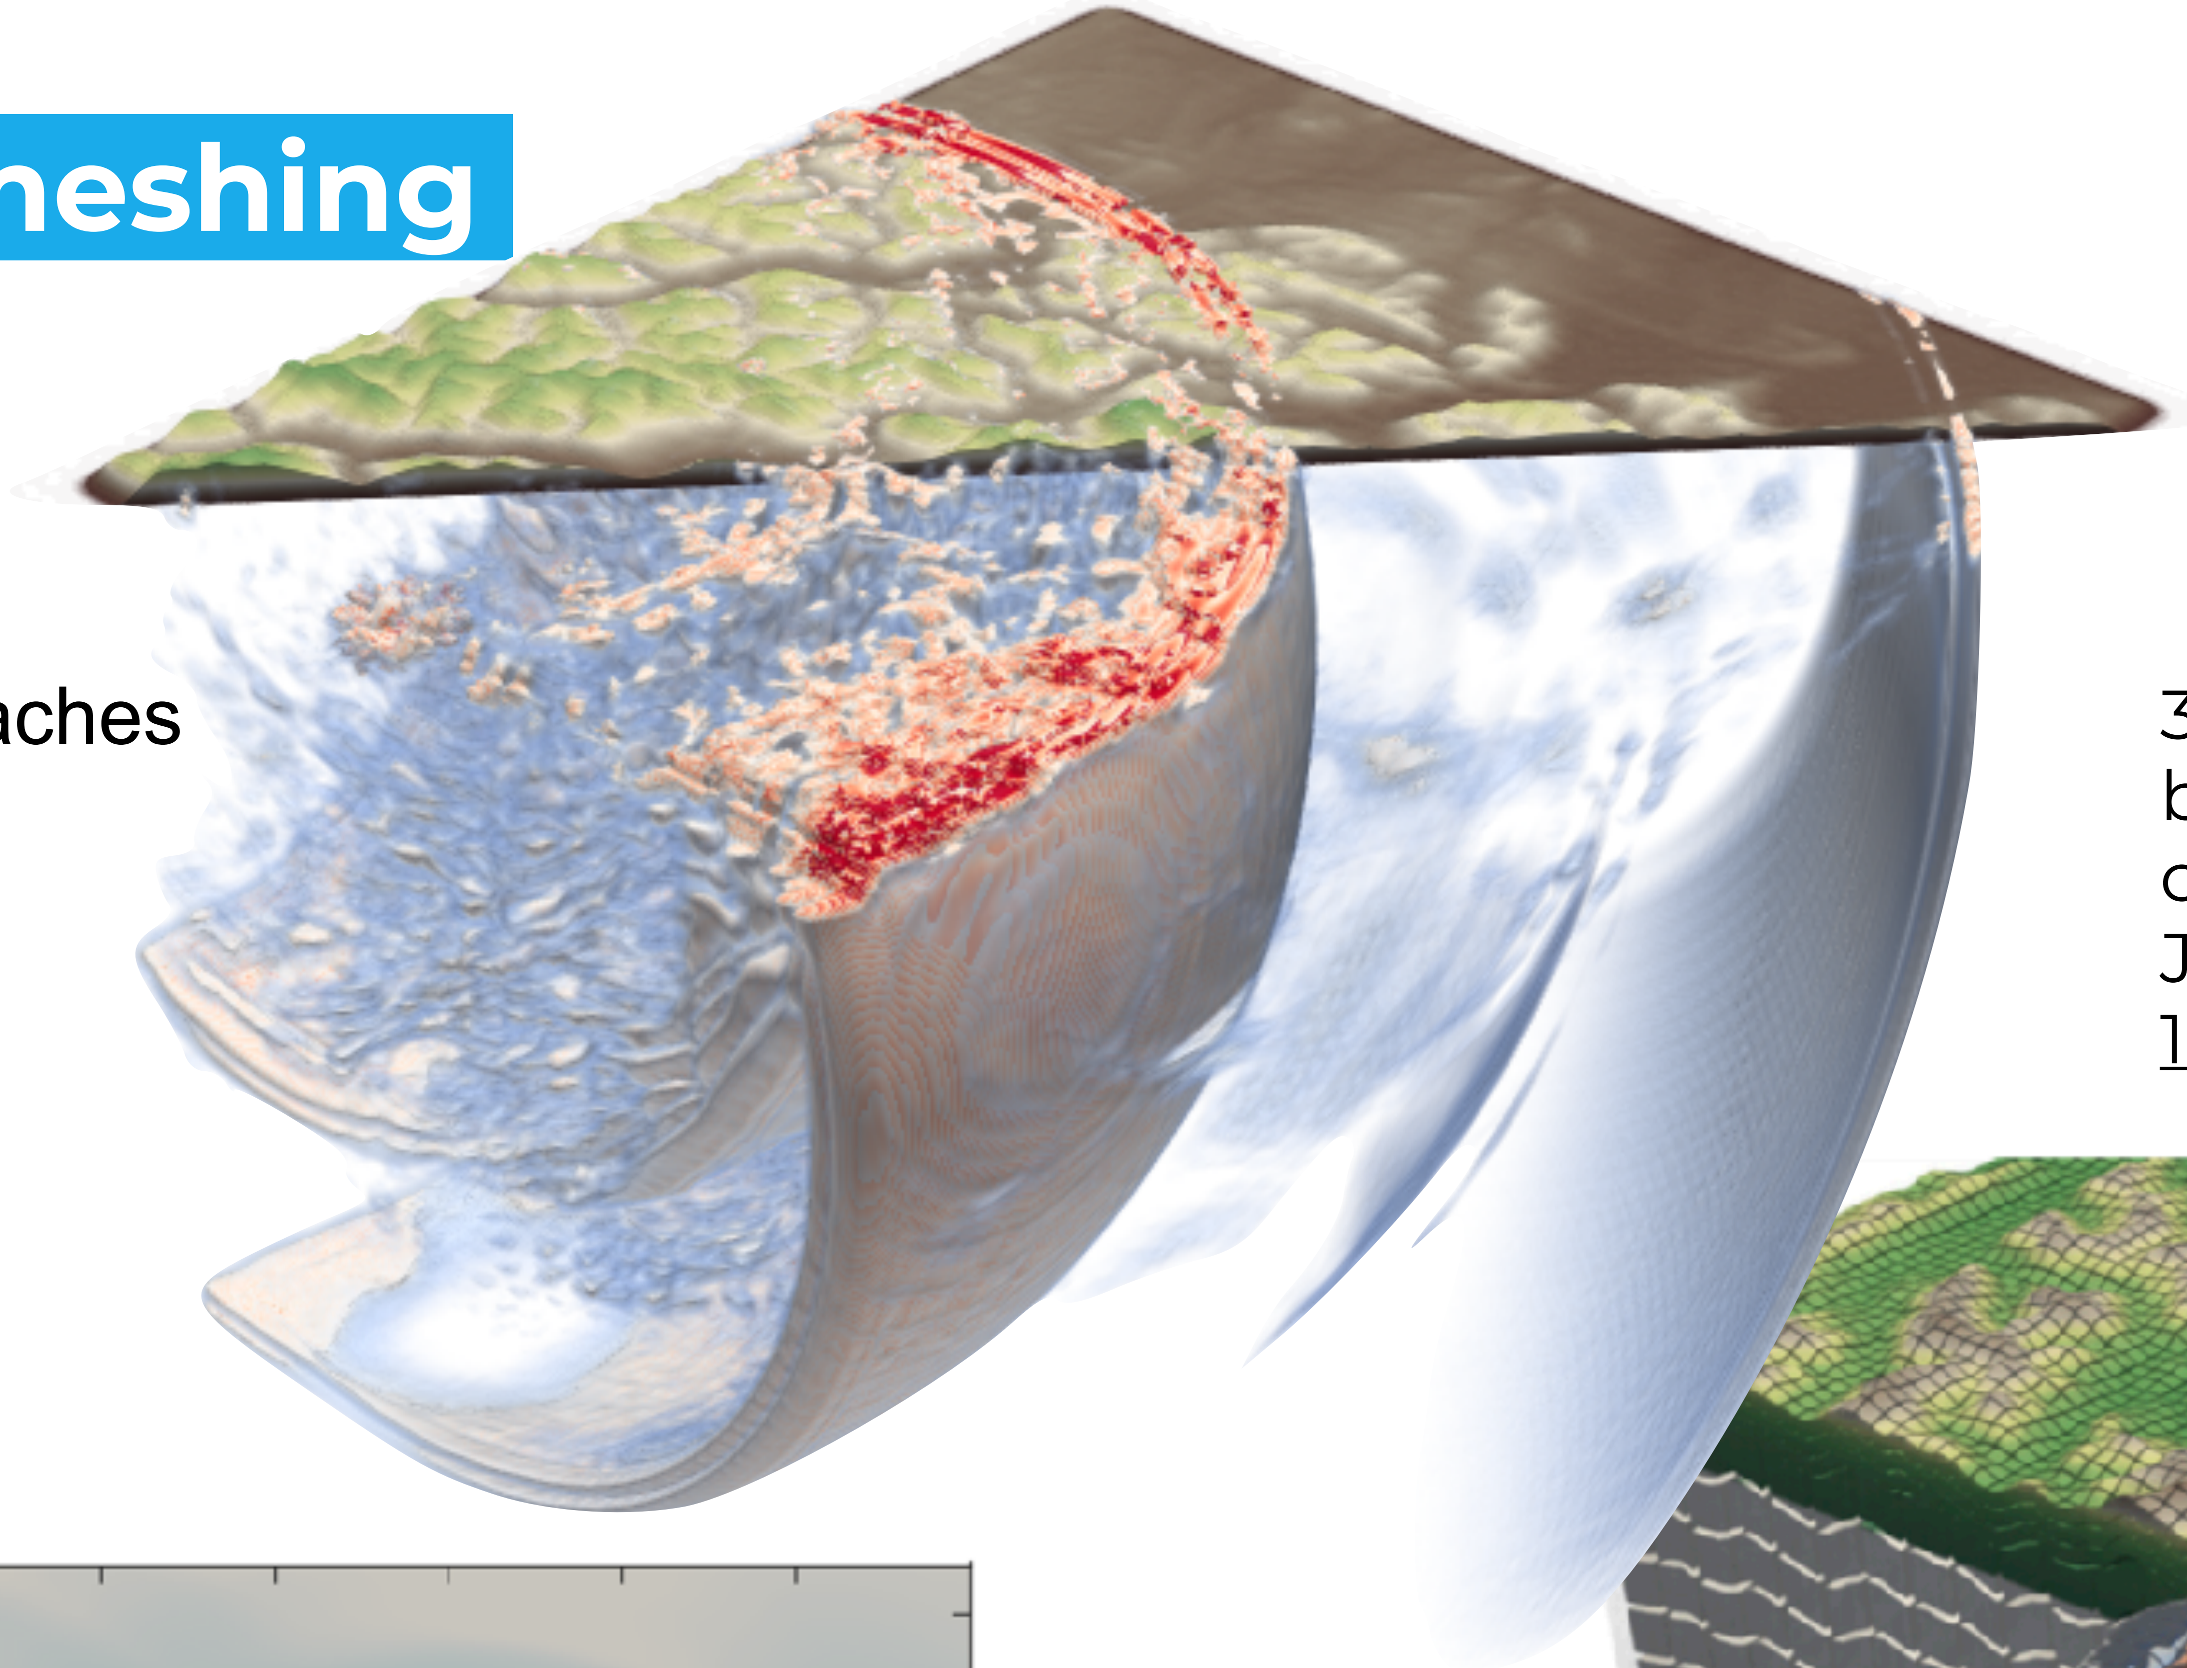

Reinarz et al., CPC, 2020

- Emerging:

Diffuse interface and **curvilinear mesh** approaches

3D curvilinear meshes via multi-block boundary conforming curvilinear meshes (Duru et al., JSC’21 & https://arxiv.org/abs/1907.02658

Propagation of an out-of-plane brittle crack using the diffuse interface GPR model and ExaHyPE [Tavelli et al., JCP’20, Gabriel et al. Proc. R. Soc. A, 2021].

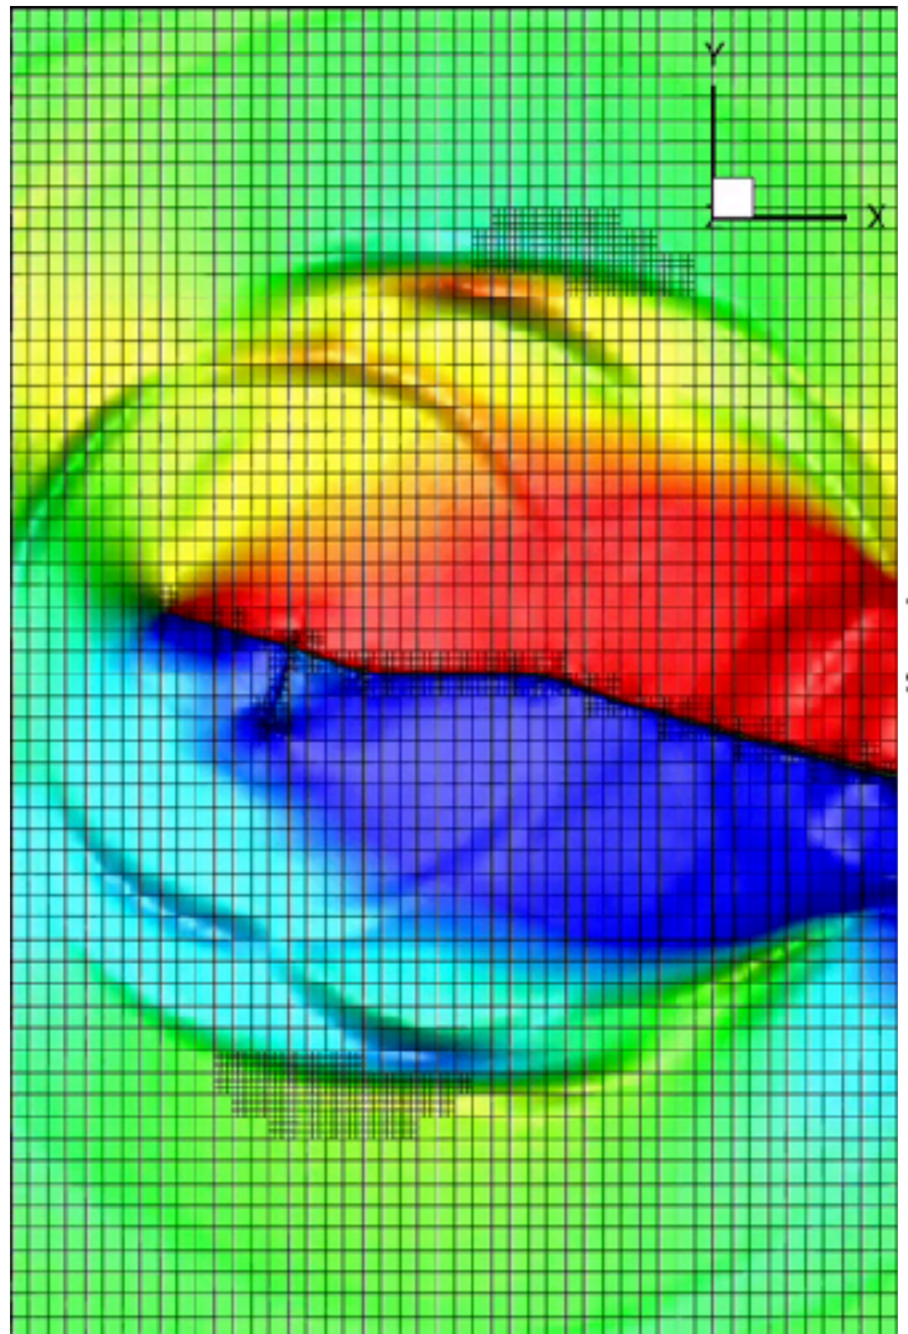

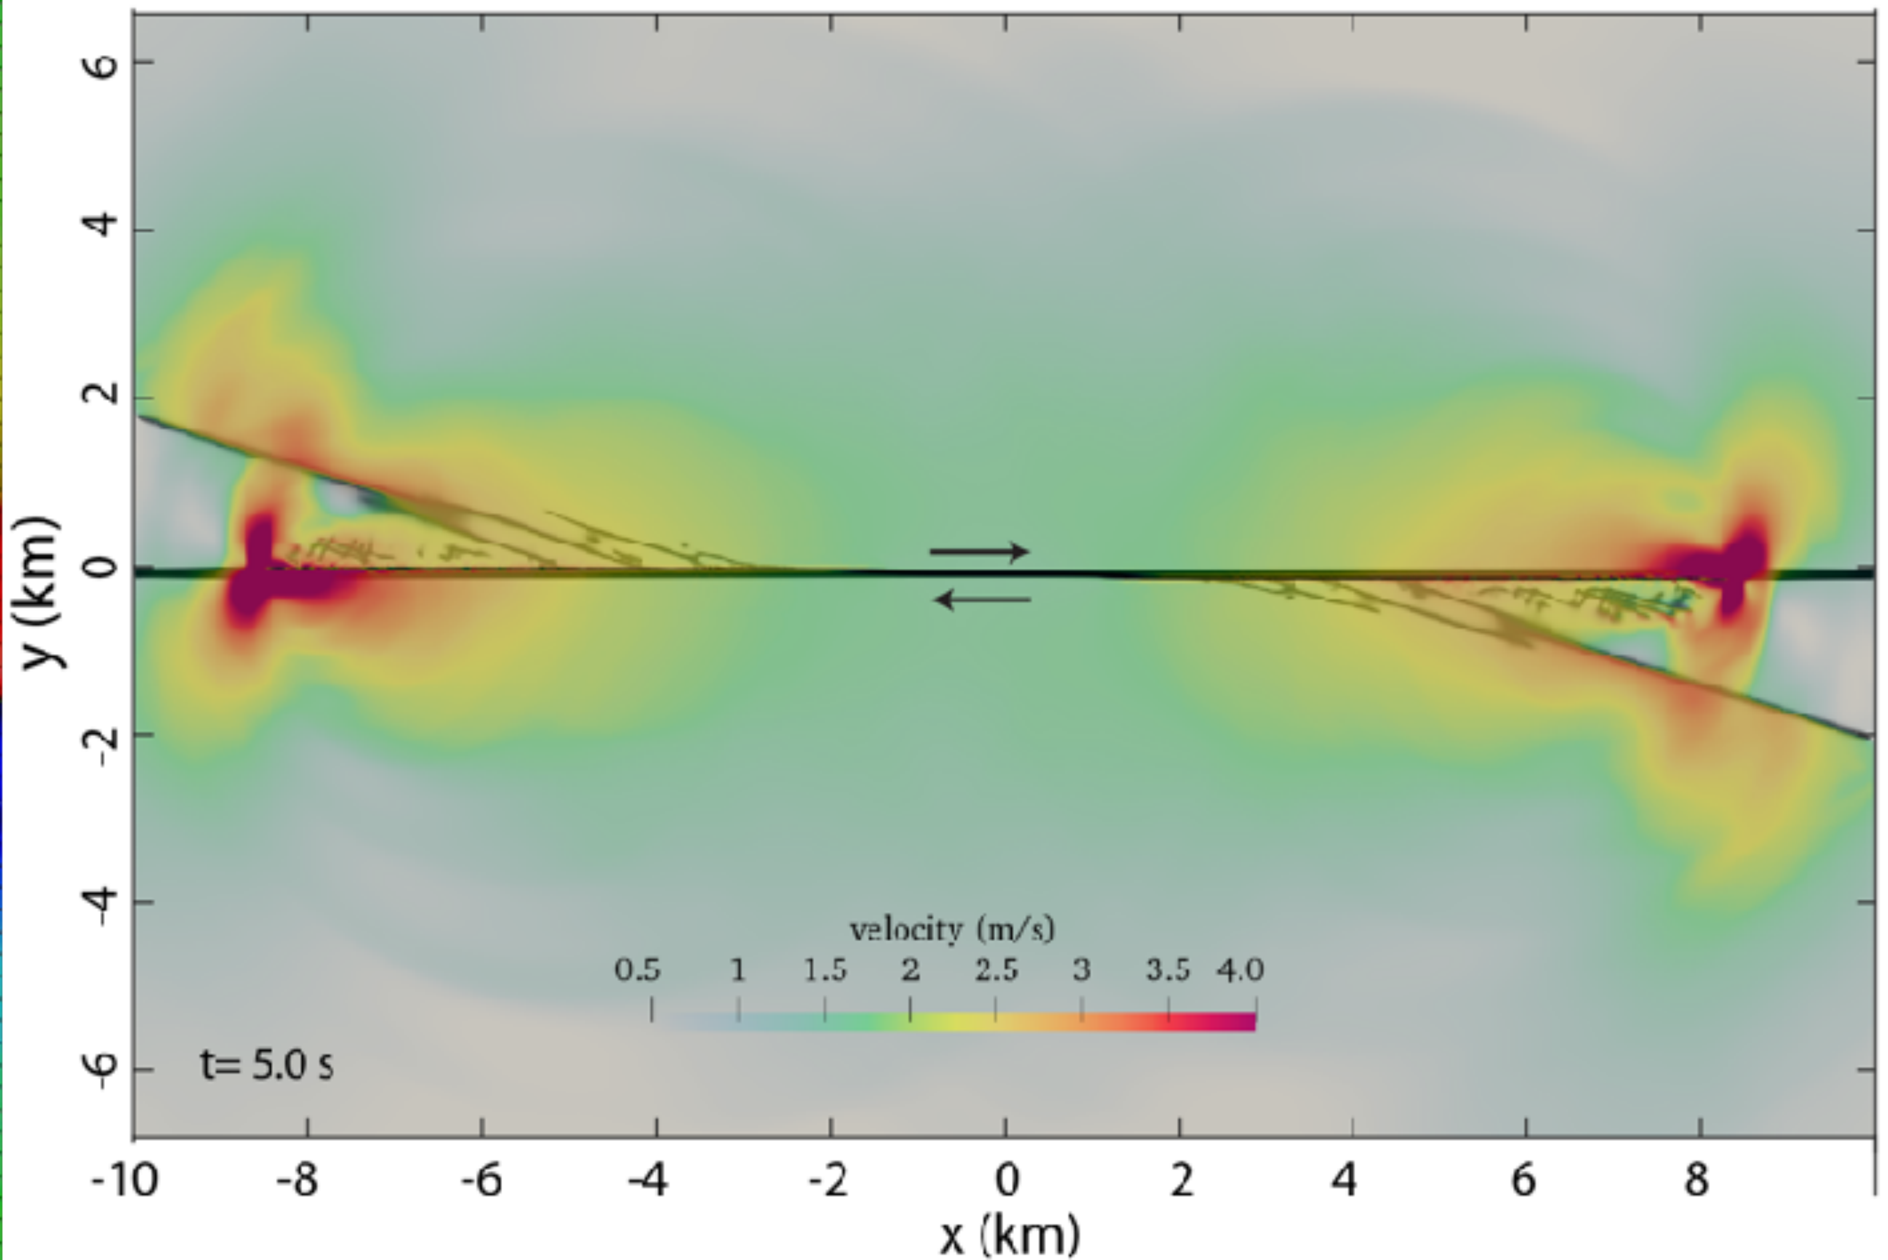

without feature-preserving meshing

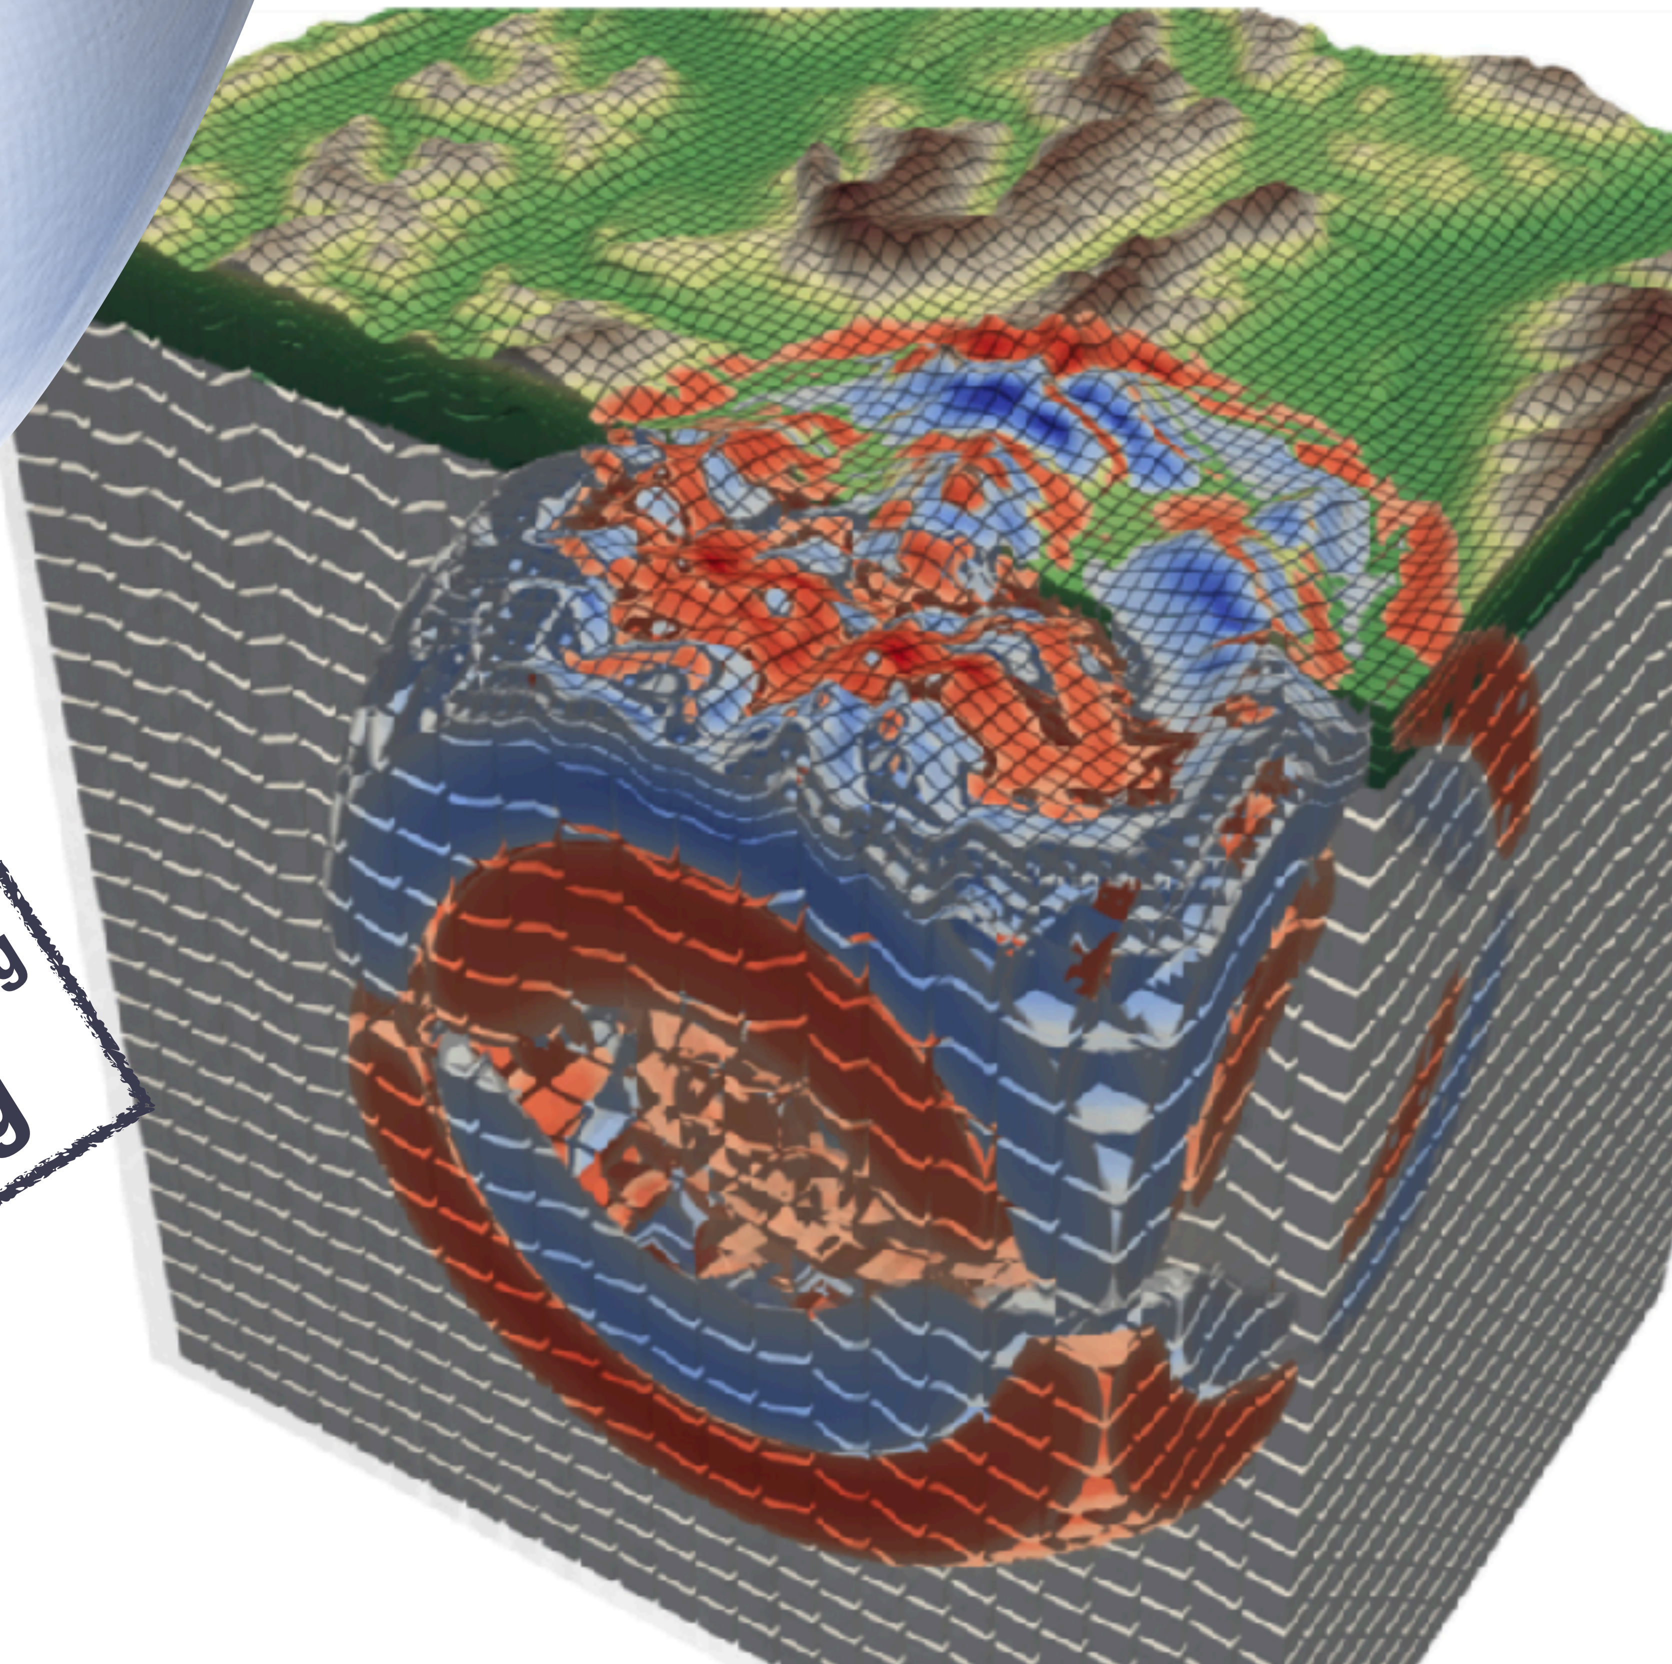

The 2004 Sumatra-Andaman earthquake and tsunami - rise to a modelling challenge

- Extreme-scale runs tackle the large space-time scales of **megathrust earthquakes and tsunami** modeling
- Requires numerical methods handling geometric complexity and **highly varying element sizes**
- parallel automatic mesh generation for tetrahedral meshes (tested up to **≈ 1 billion elements**)
- Local-time stepping mitigates e.g. 'sliver elements'

BEST PAPER

Uphoff et al., SC17

Ulrich et al., Nature Geoscience, in press
Madden et al., EarthArxiv'21

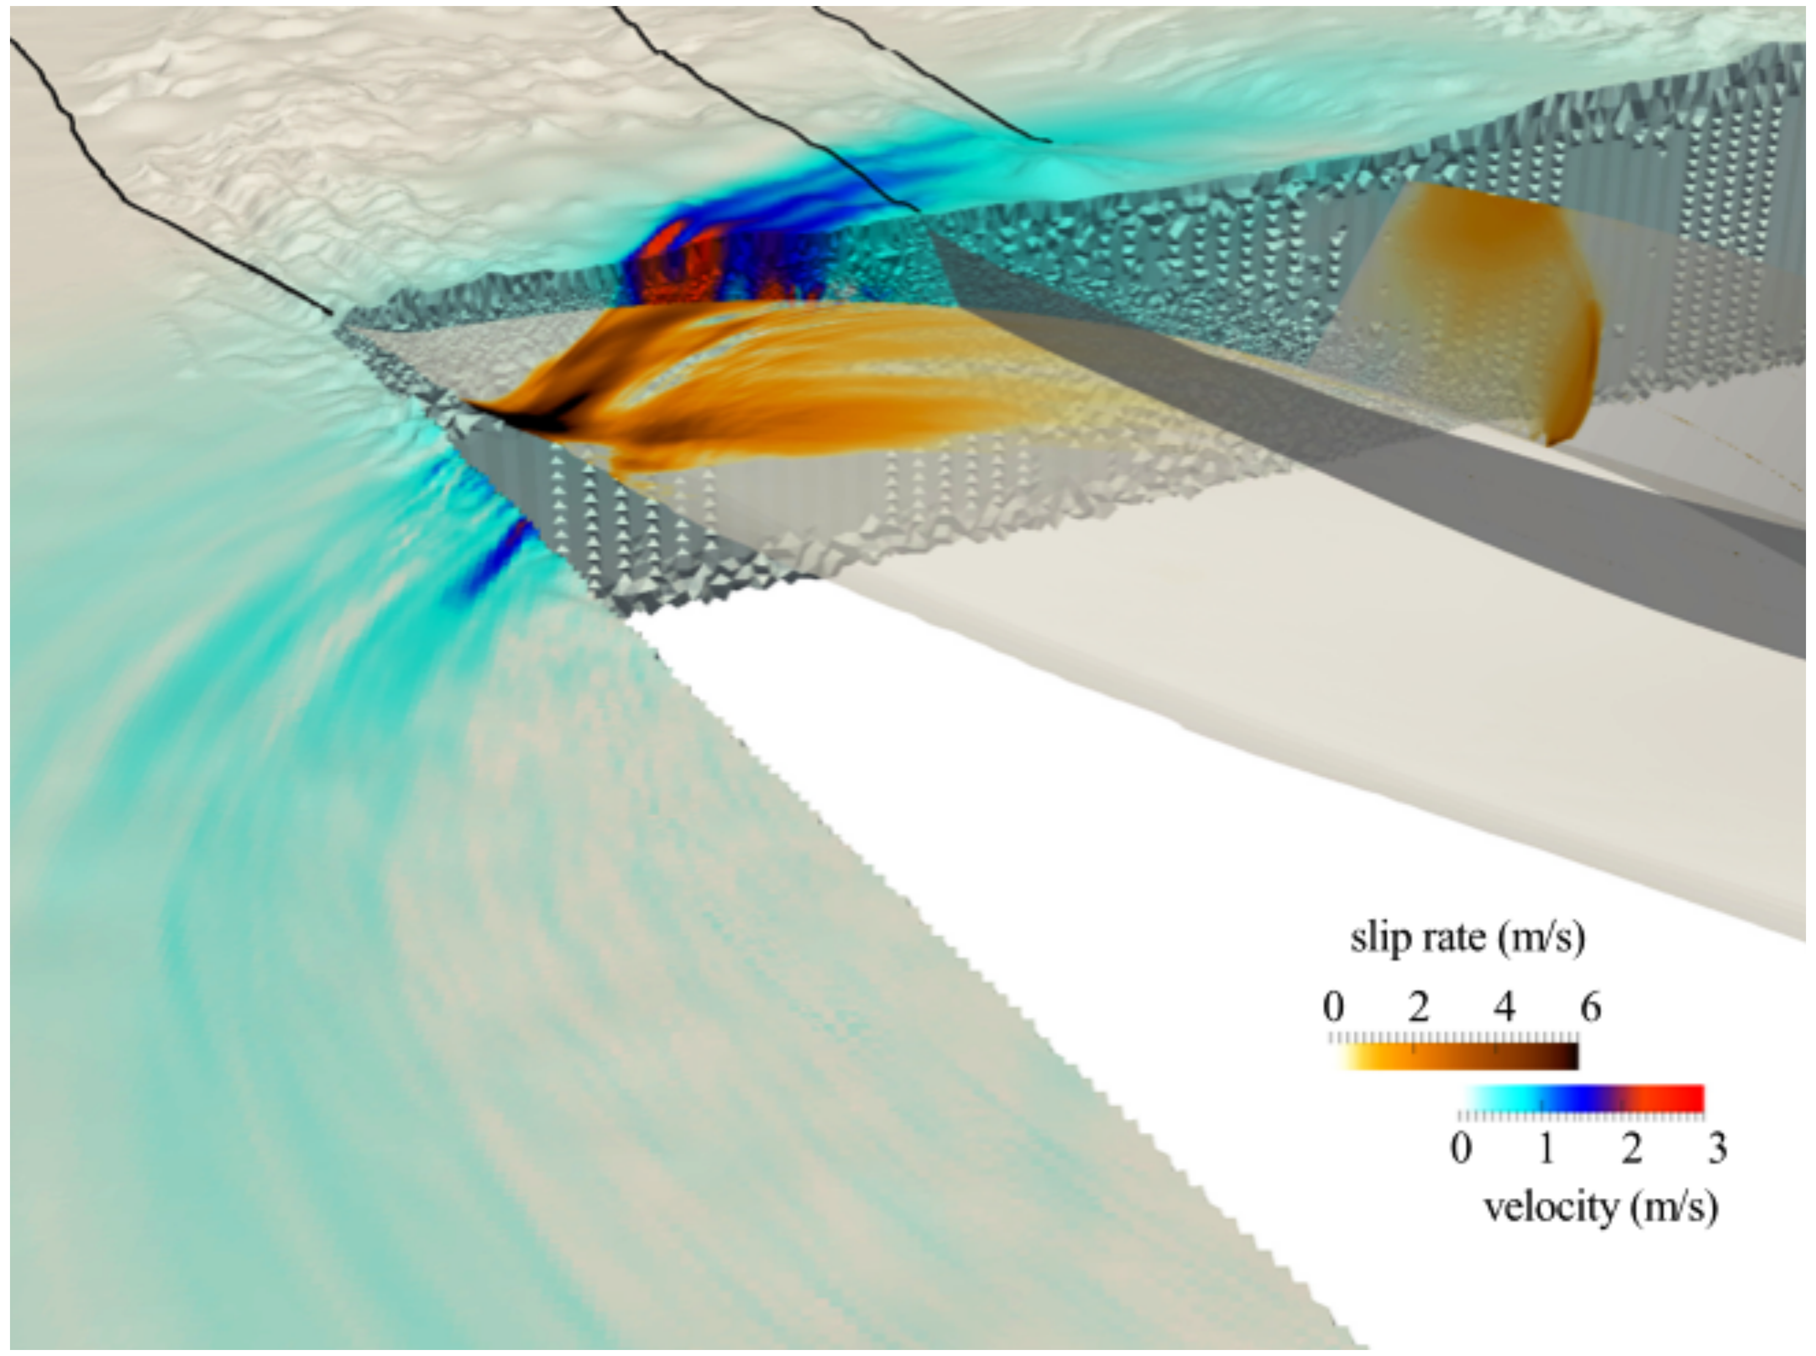

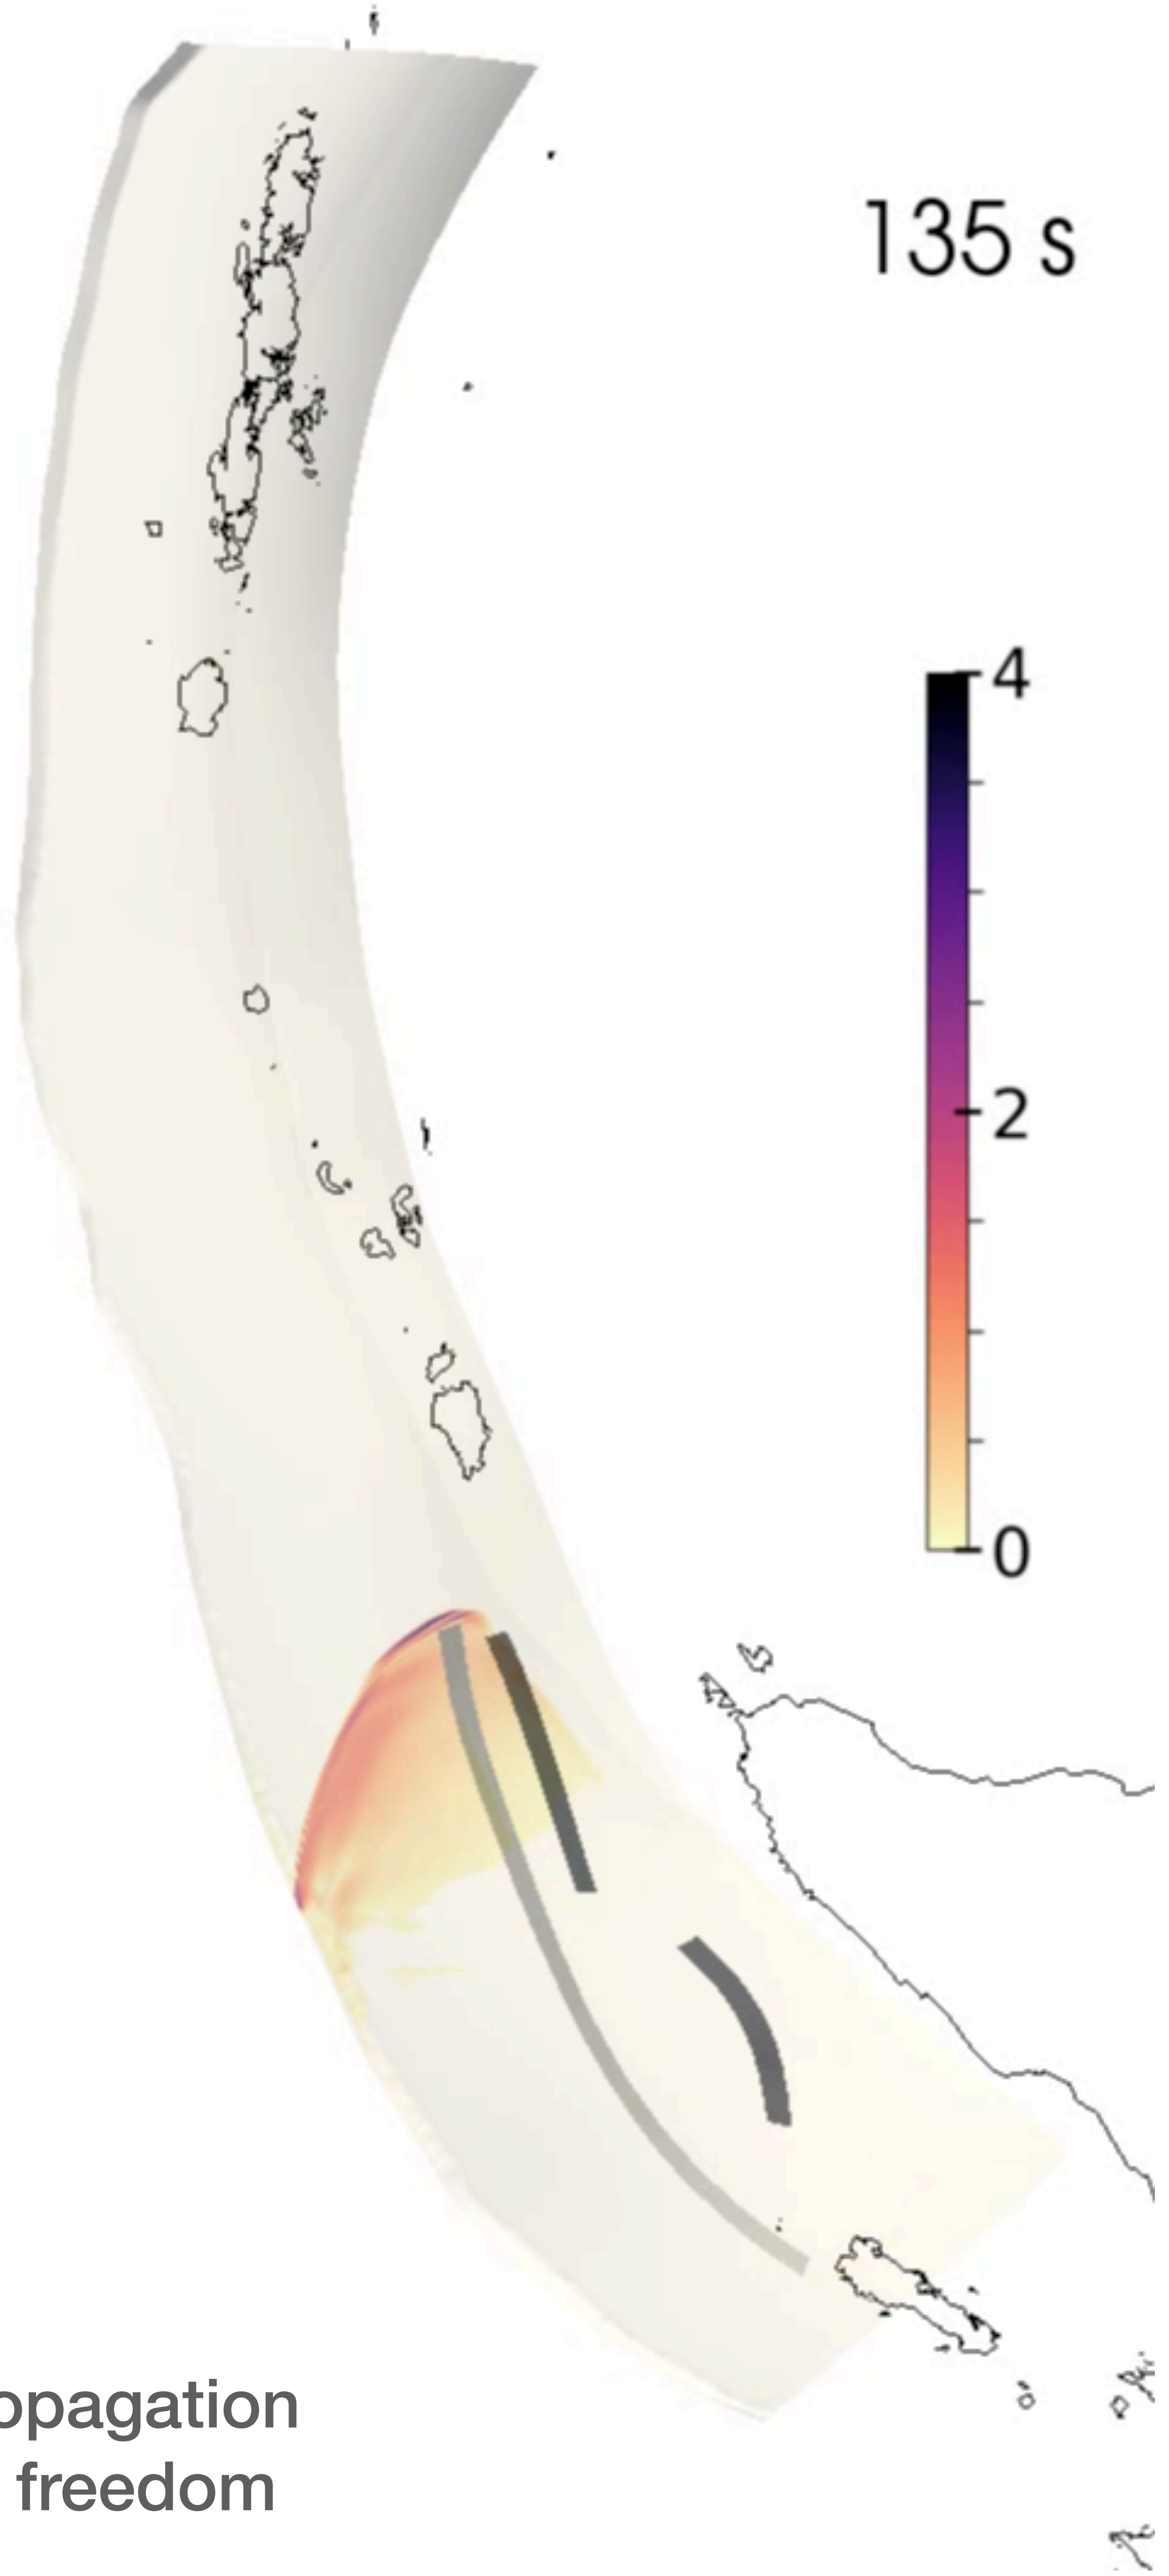

Sumatra earthquake dynamic rupture and tsunami propagation scenario with 220M elements ~111 x 10^9 degrees of freedom

ChEESE
Center of Excellence for Exascale in Solid Earth

The 2004 Sumatra-Andaman earthquake and tsunami - rise to a modelling challenge

BEST PAPER

Uphoff et al., SC17

- While a **"hero run"** took 14h on full SuperMuc2; now typical high-order simulations: >10 Mio elements run 5h30 minutes on 16 nodes **(4k CPUh)**
- Enabled by end-to-end computational optimization including a **geoinformation server** for fast and asynchronous input/output, clustered local time stepping, off-line code generator, flexible boundary conditions (e.g. gravity), GPU optimisation (in progress)

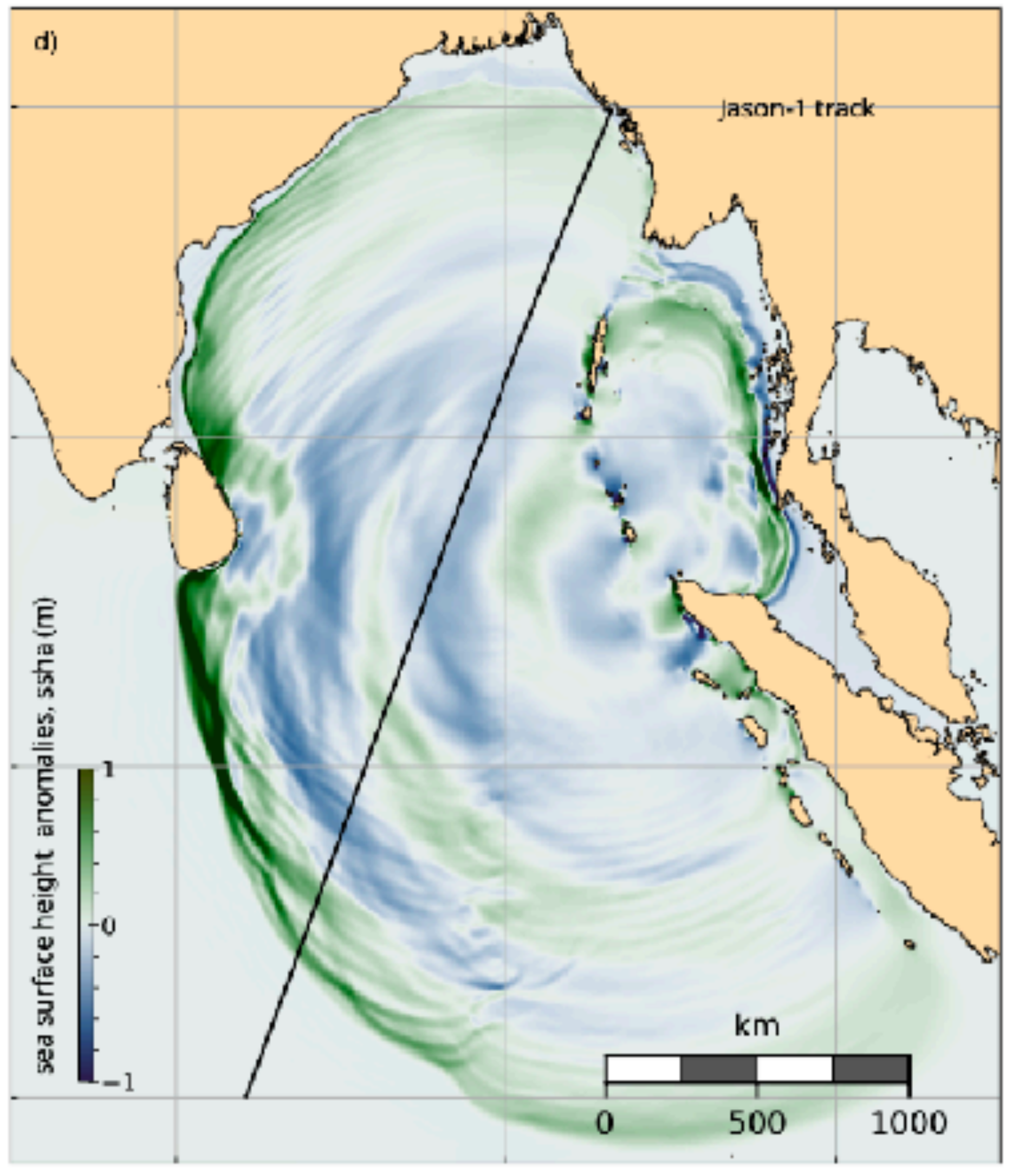

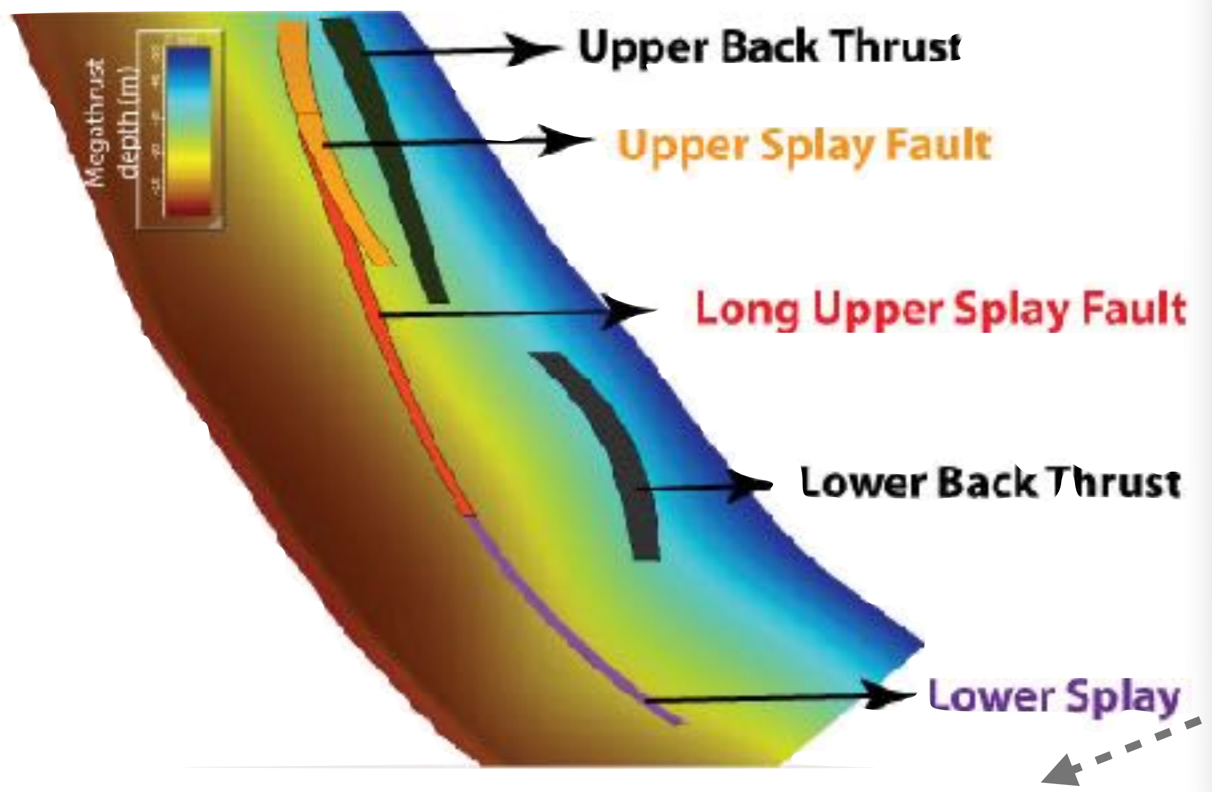

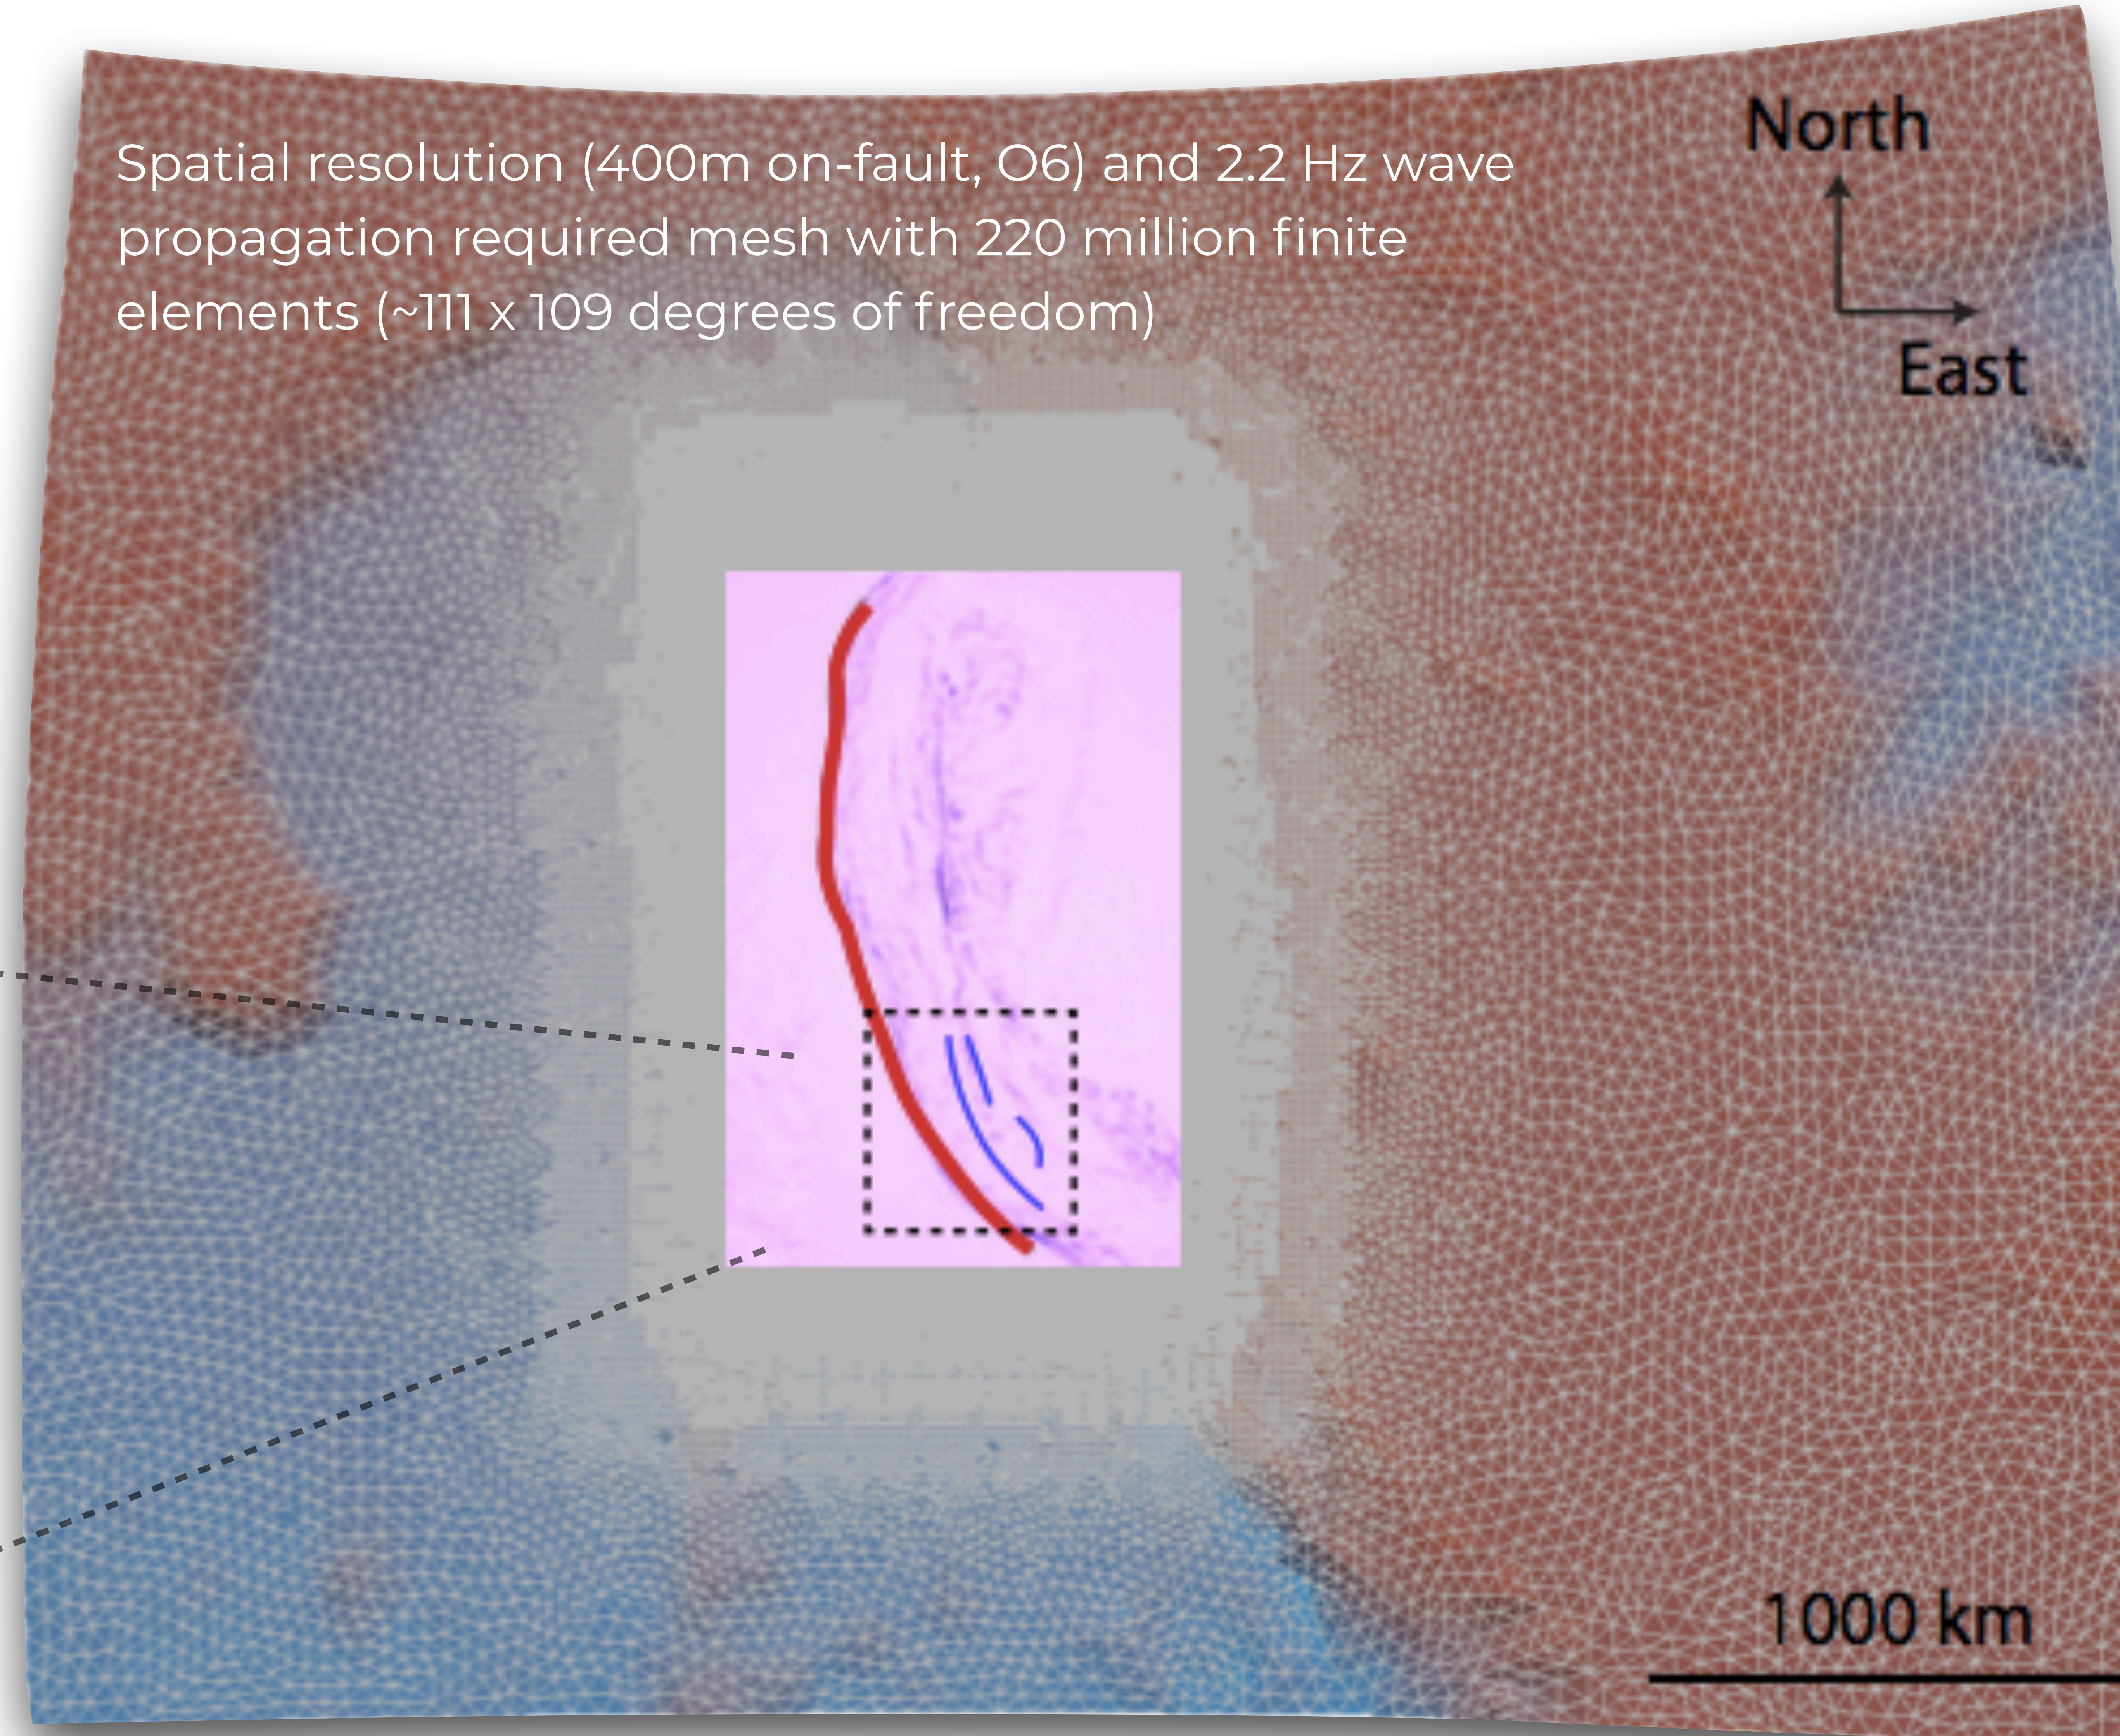

Rapid earthquake/tsunami modelling - the 2018, Palu-Sulawesi Event

- Rapid 3D dynamic rupture setup from **sparse data**
- **Fault system** from Sentinel-2, SAR data, regional seismicity; **Stress and strength** based on World stress map; and assuming a **transtensional** regime; high fluid pressure, mechanic viability across the fault system's geometric complexities, dynamics constrained by teleseismics and moment rate release

Ulrich et al., PAGEOPH 2019

Rapid earthquake/tsunami modelling - the 2018, Palu-Sulawesi Event

Earthquake-induced movement of seafloor beneath Palu Bay itself may played a critical role generating the tsunami

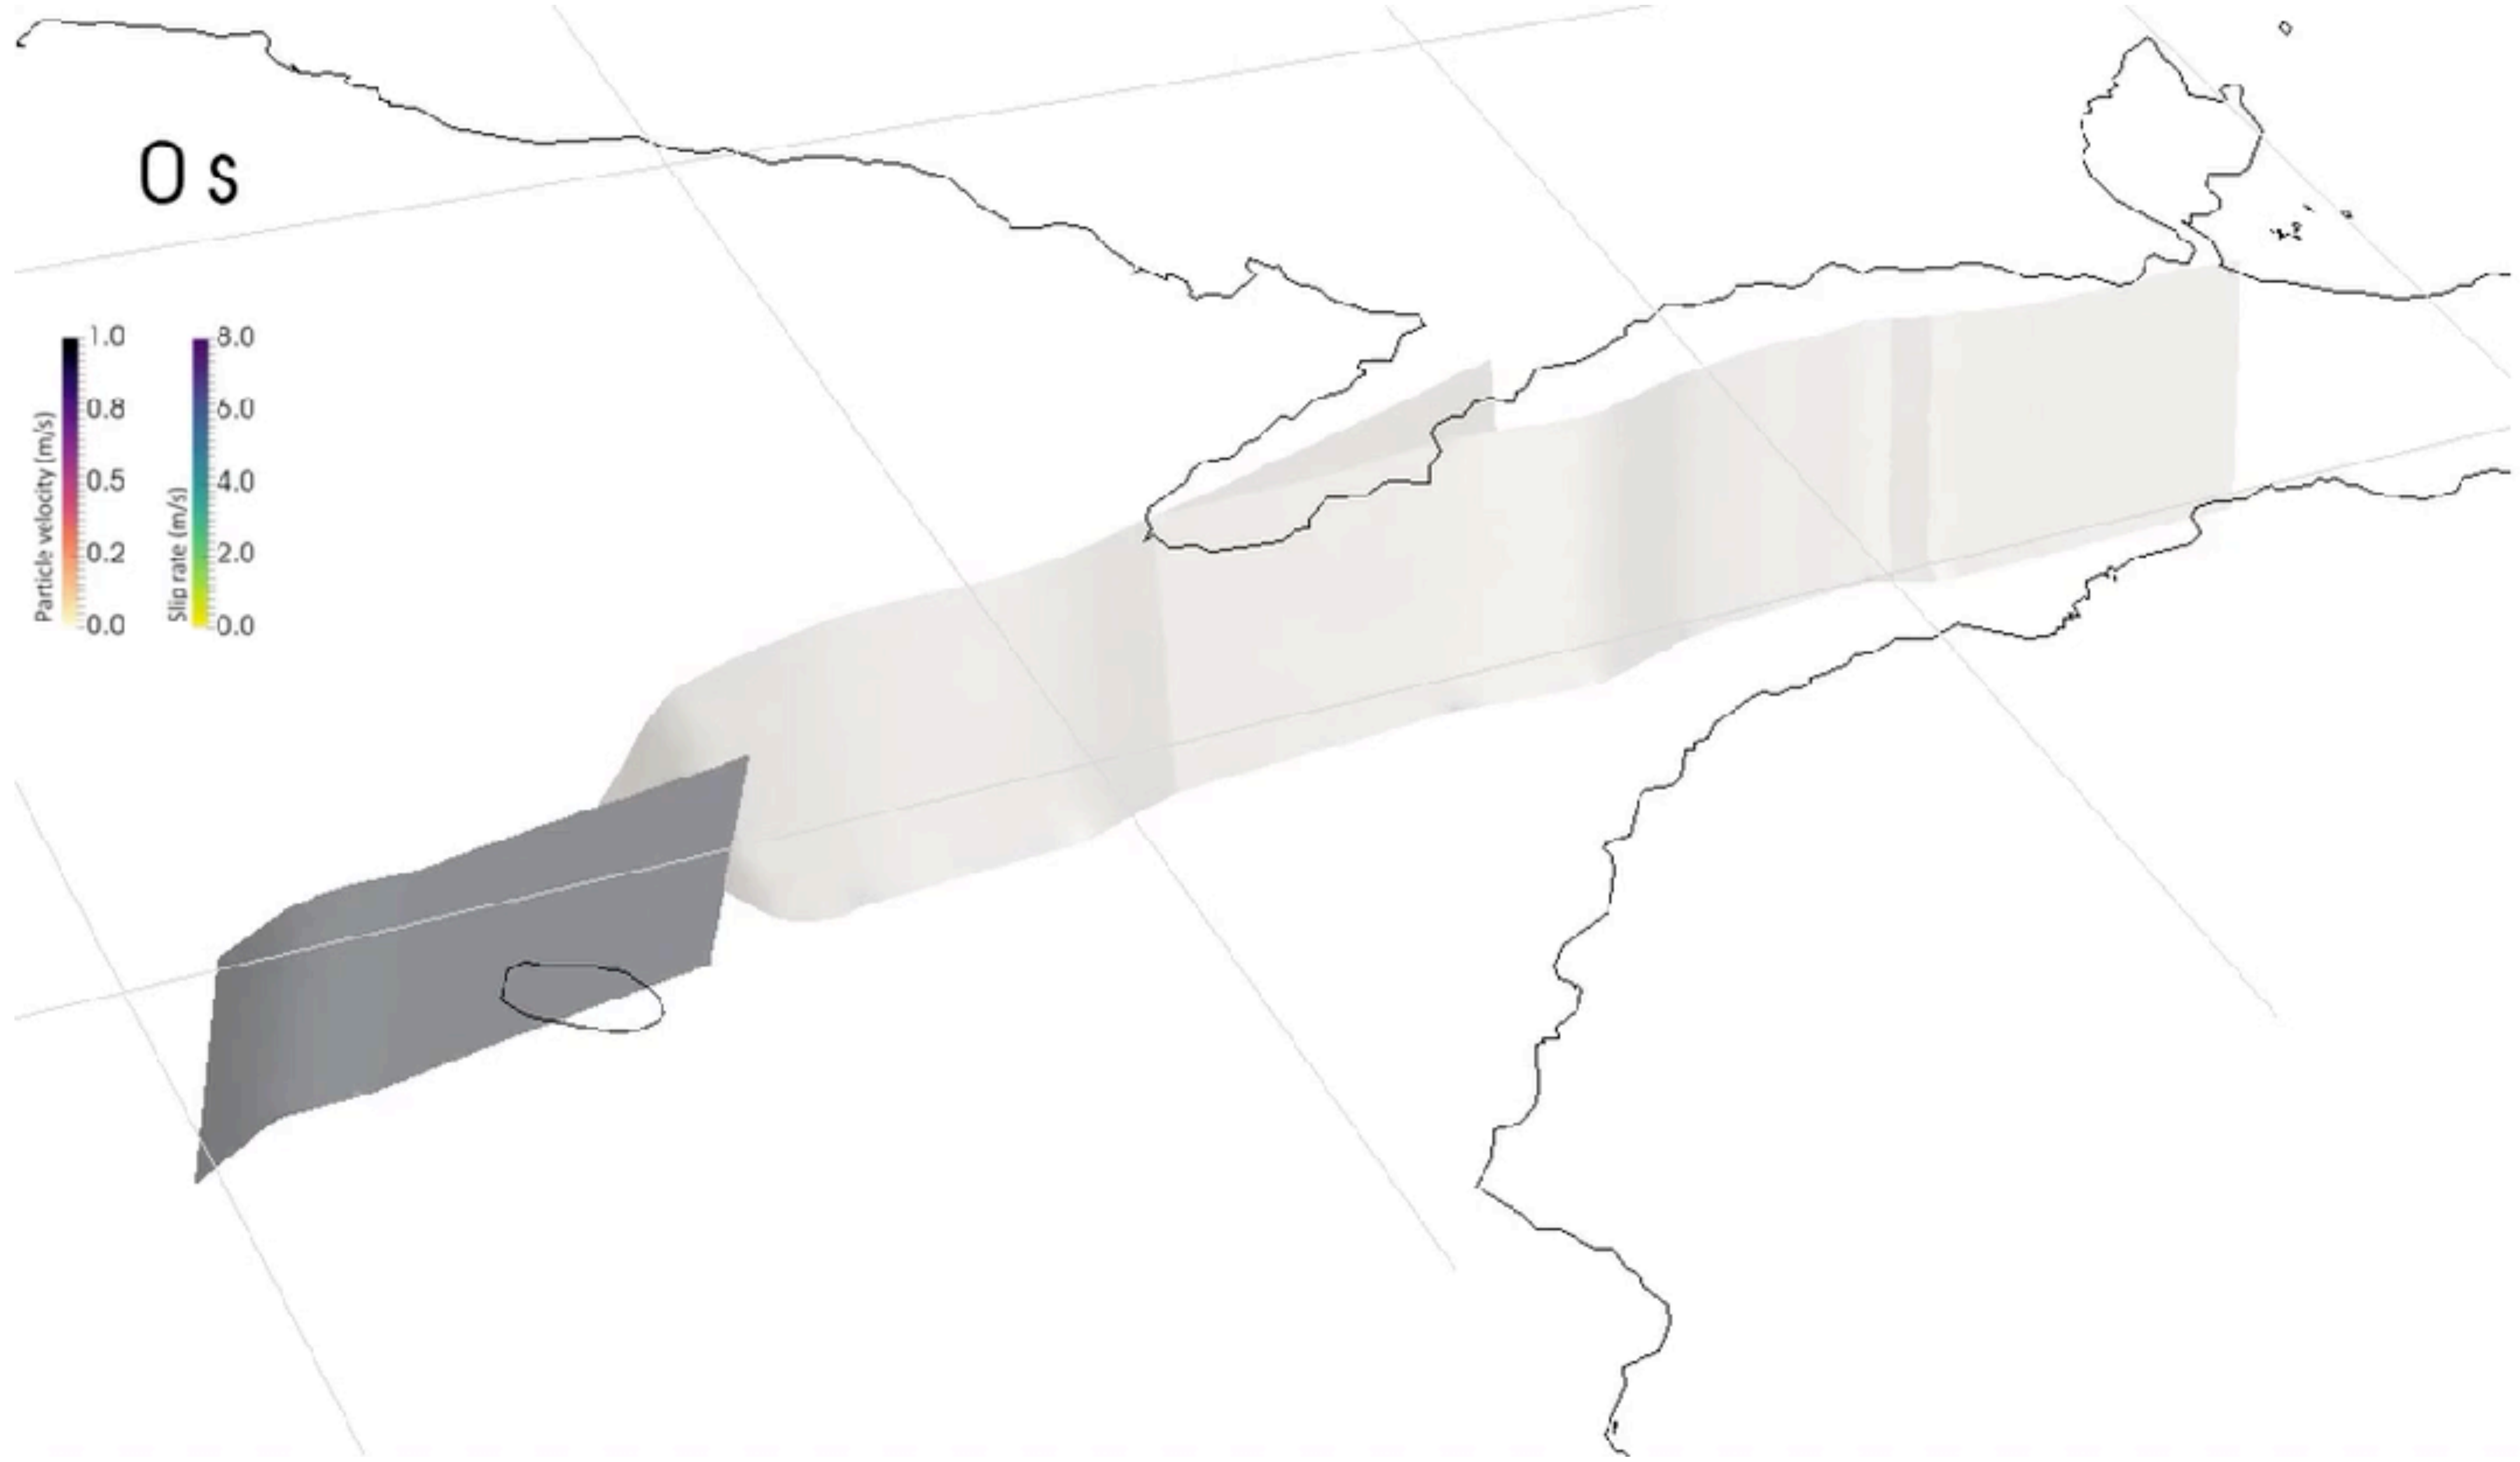

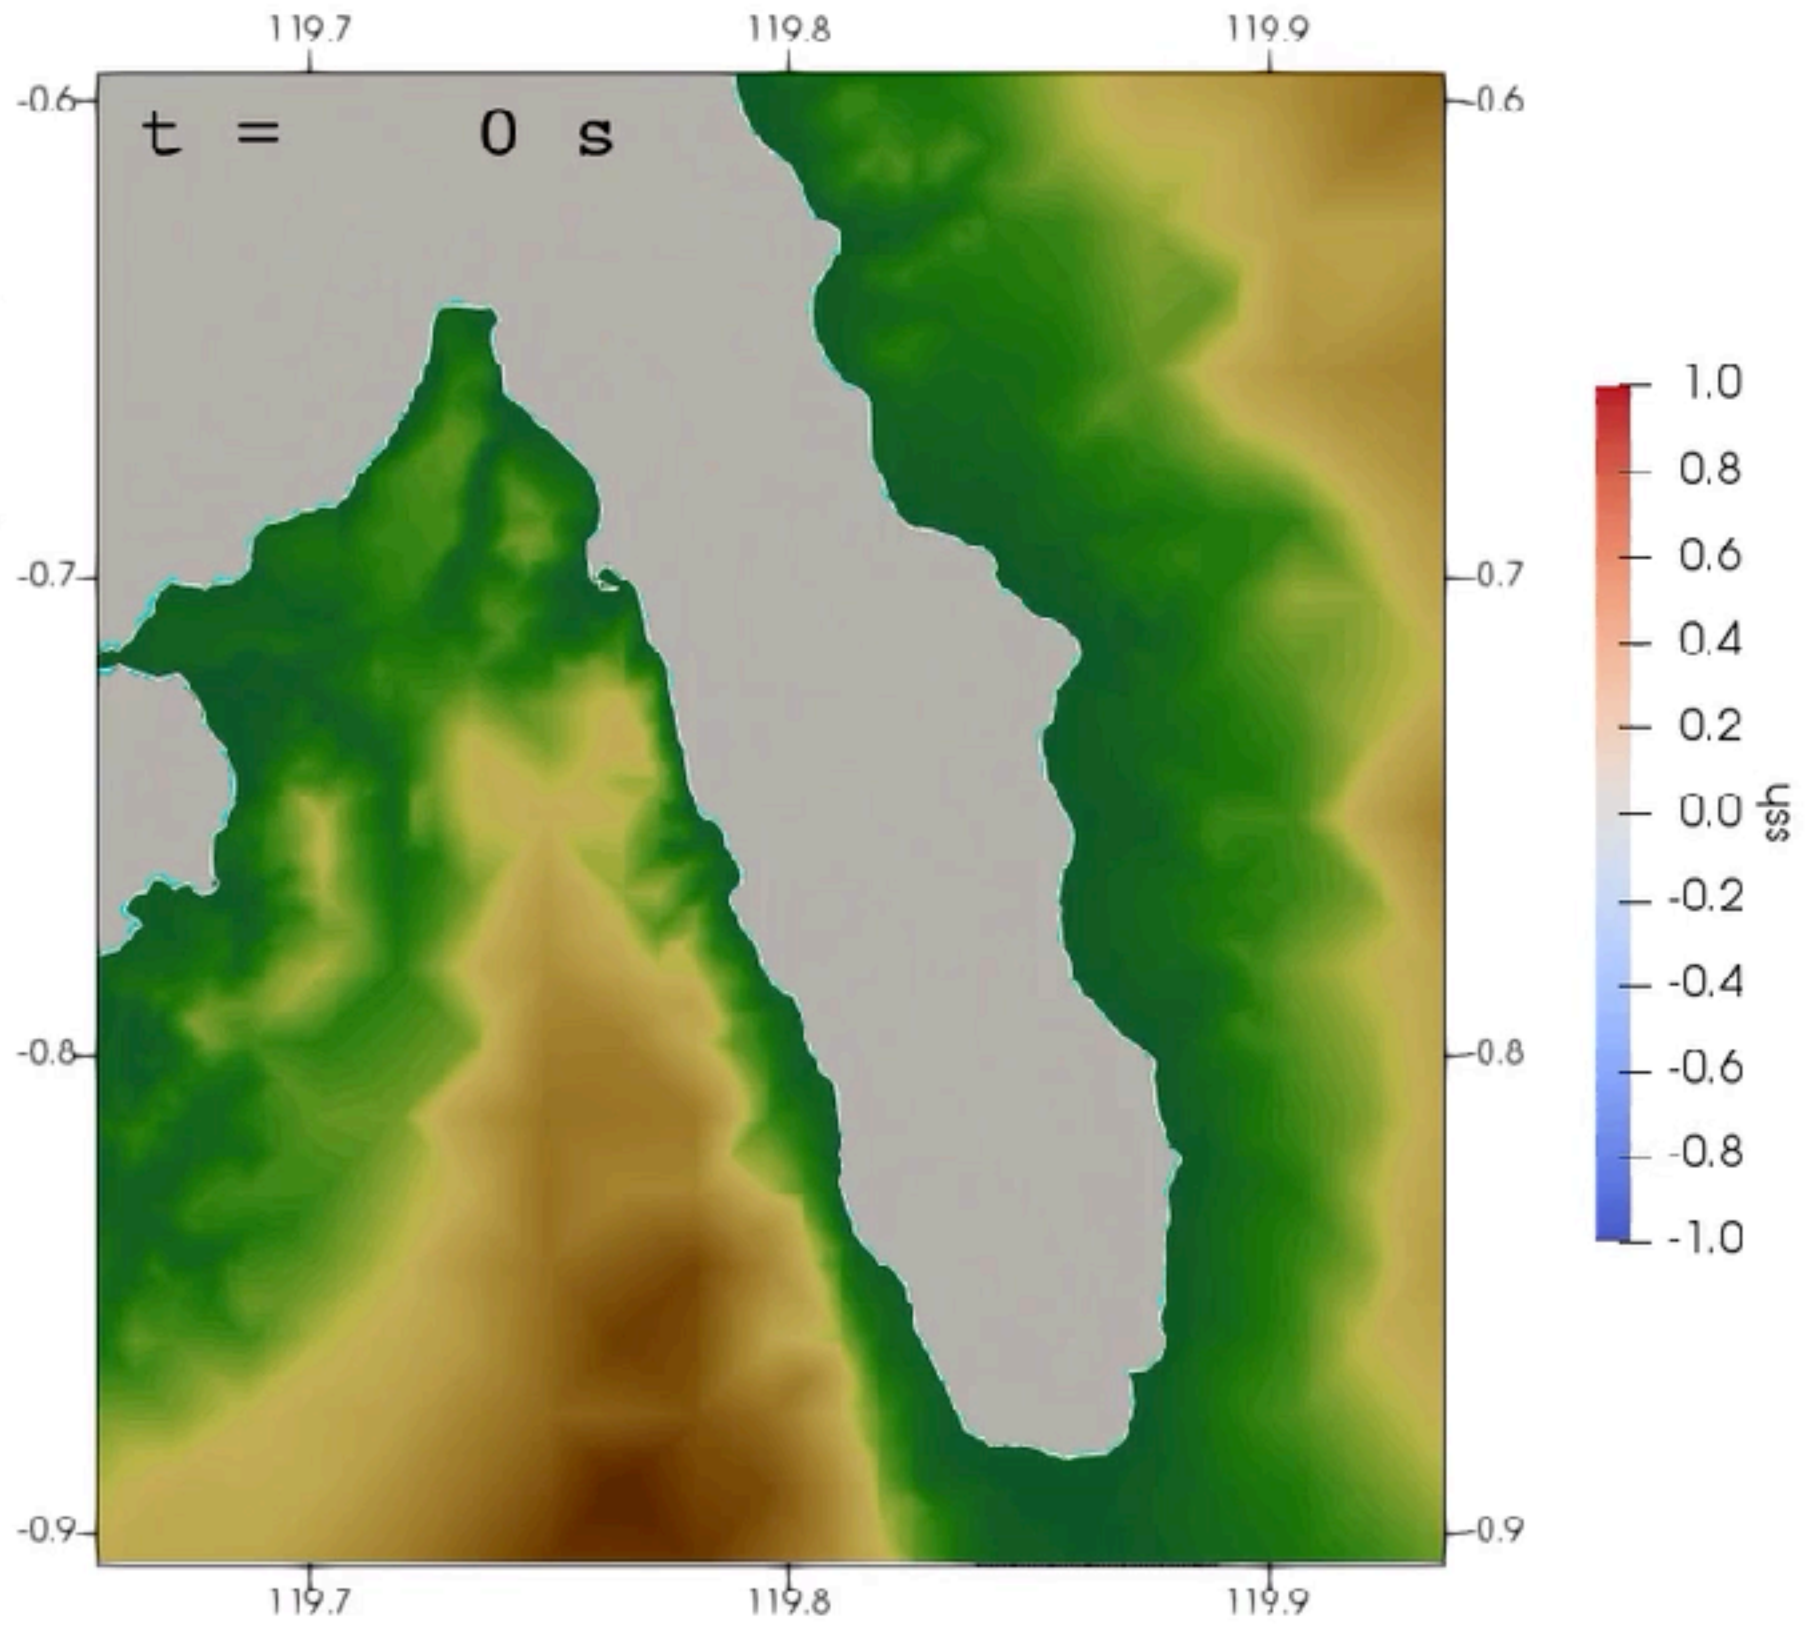

Fully coupled earthquake-tsunami modeling of the 2018 Palu, Sulawesi events

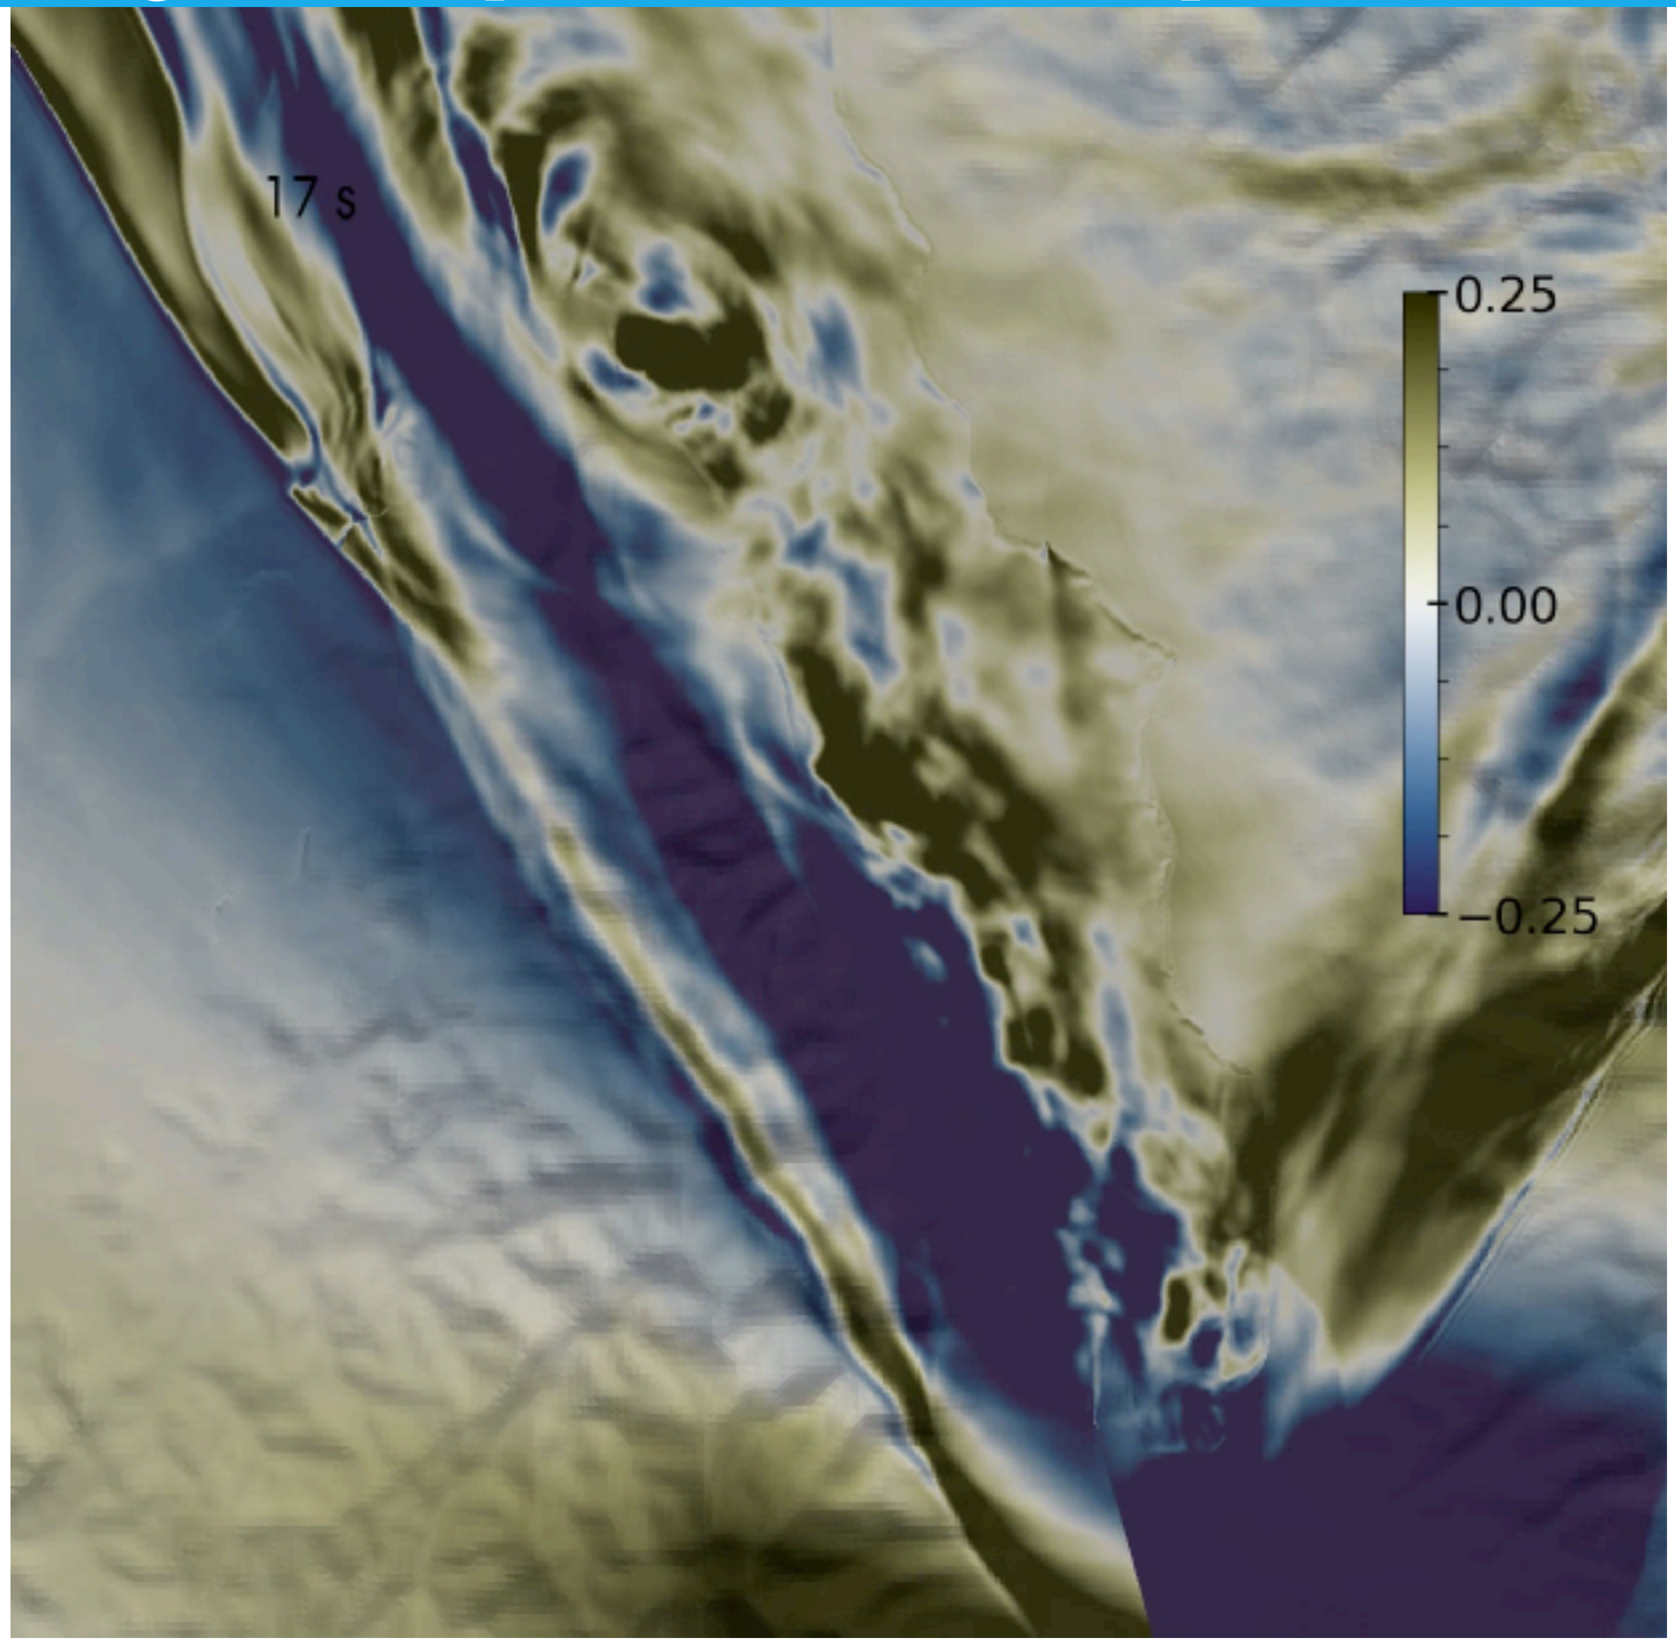

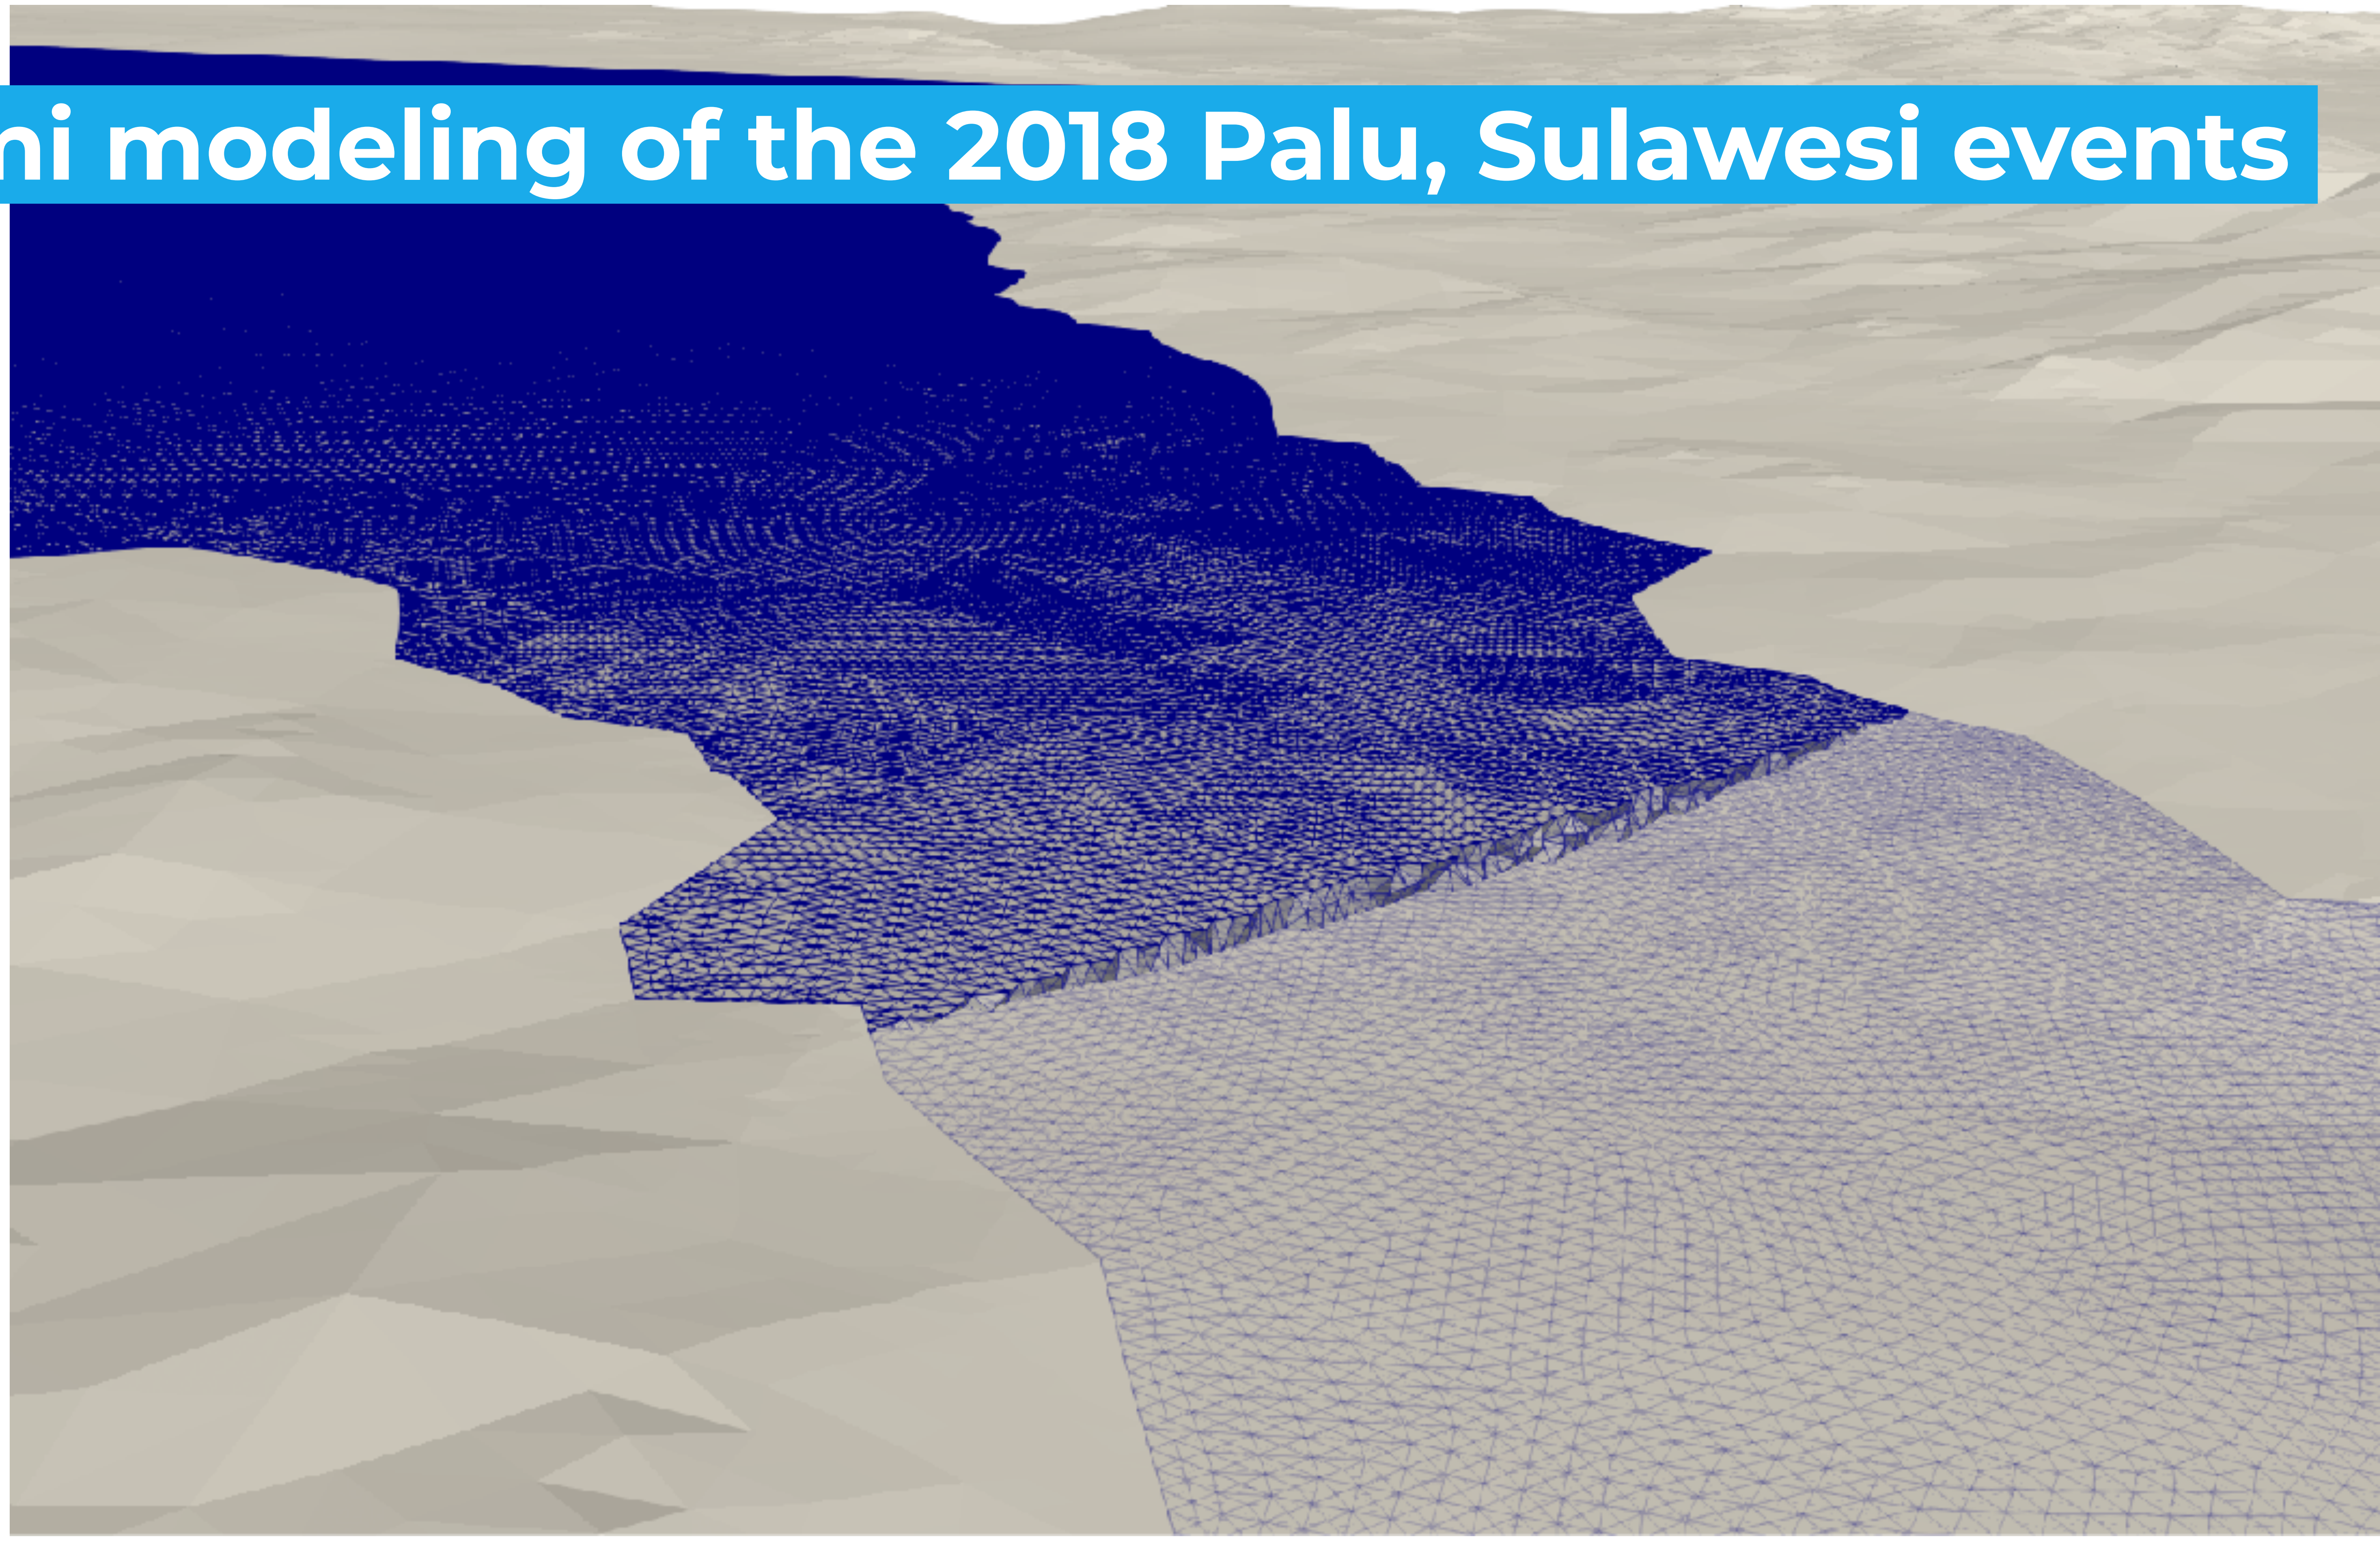

- **fully coupled 3D acoustic-elastic coupling with gravity implemented in SeisSol,** via free surface tracking (gravitational effects) by linearised free surface boundary condition
- **multi-petaflop simlation:** 518 mio. elements = 261 billion degrees of freedom, ran for 5.5 hours on 3,072 nodes SuperMUC-NG, sustained performance of 3.1 PFLOPS
- Computing Facility Application Partner competition for **NSF's future Leadership Class Computing Facility (LCCF)** reaching >95% efficiency on Frontera

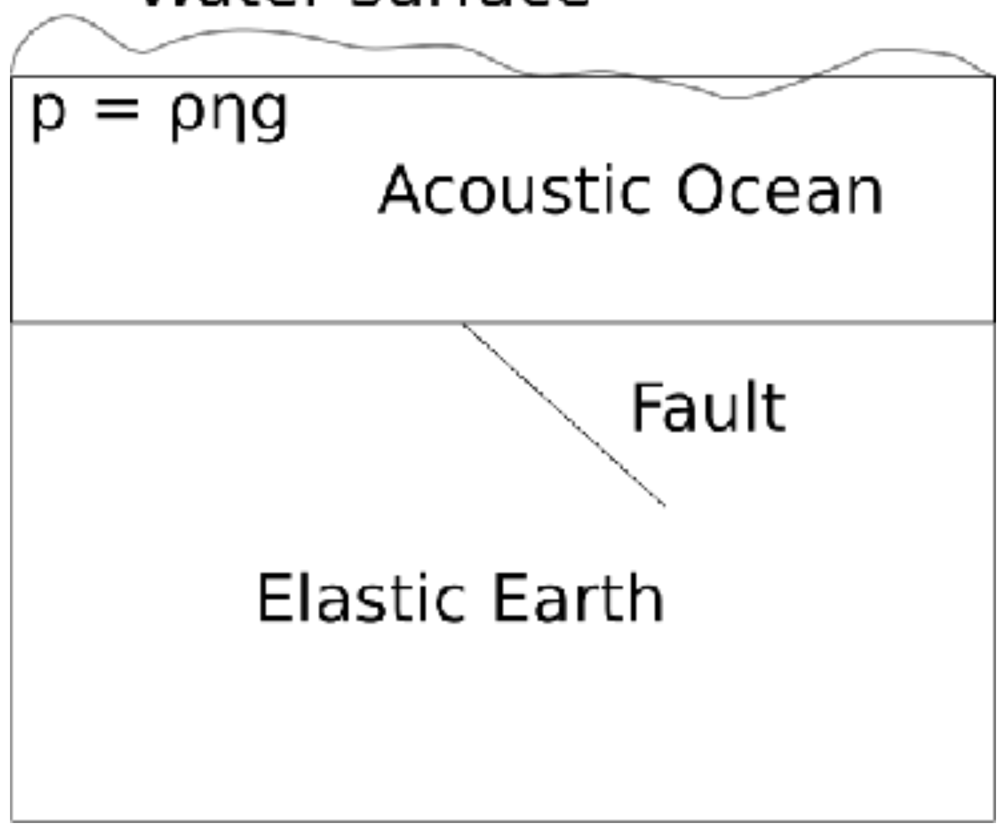

Fully coupled earthquake-tsunami modeling

Krenz et al. SC'2021

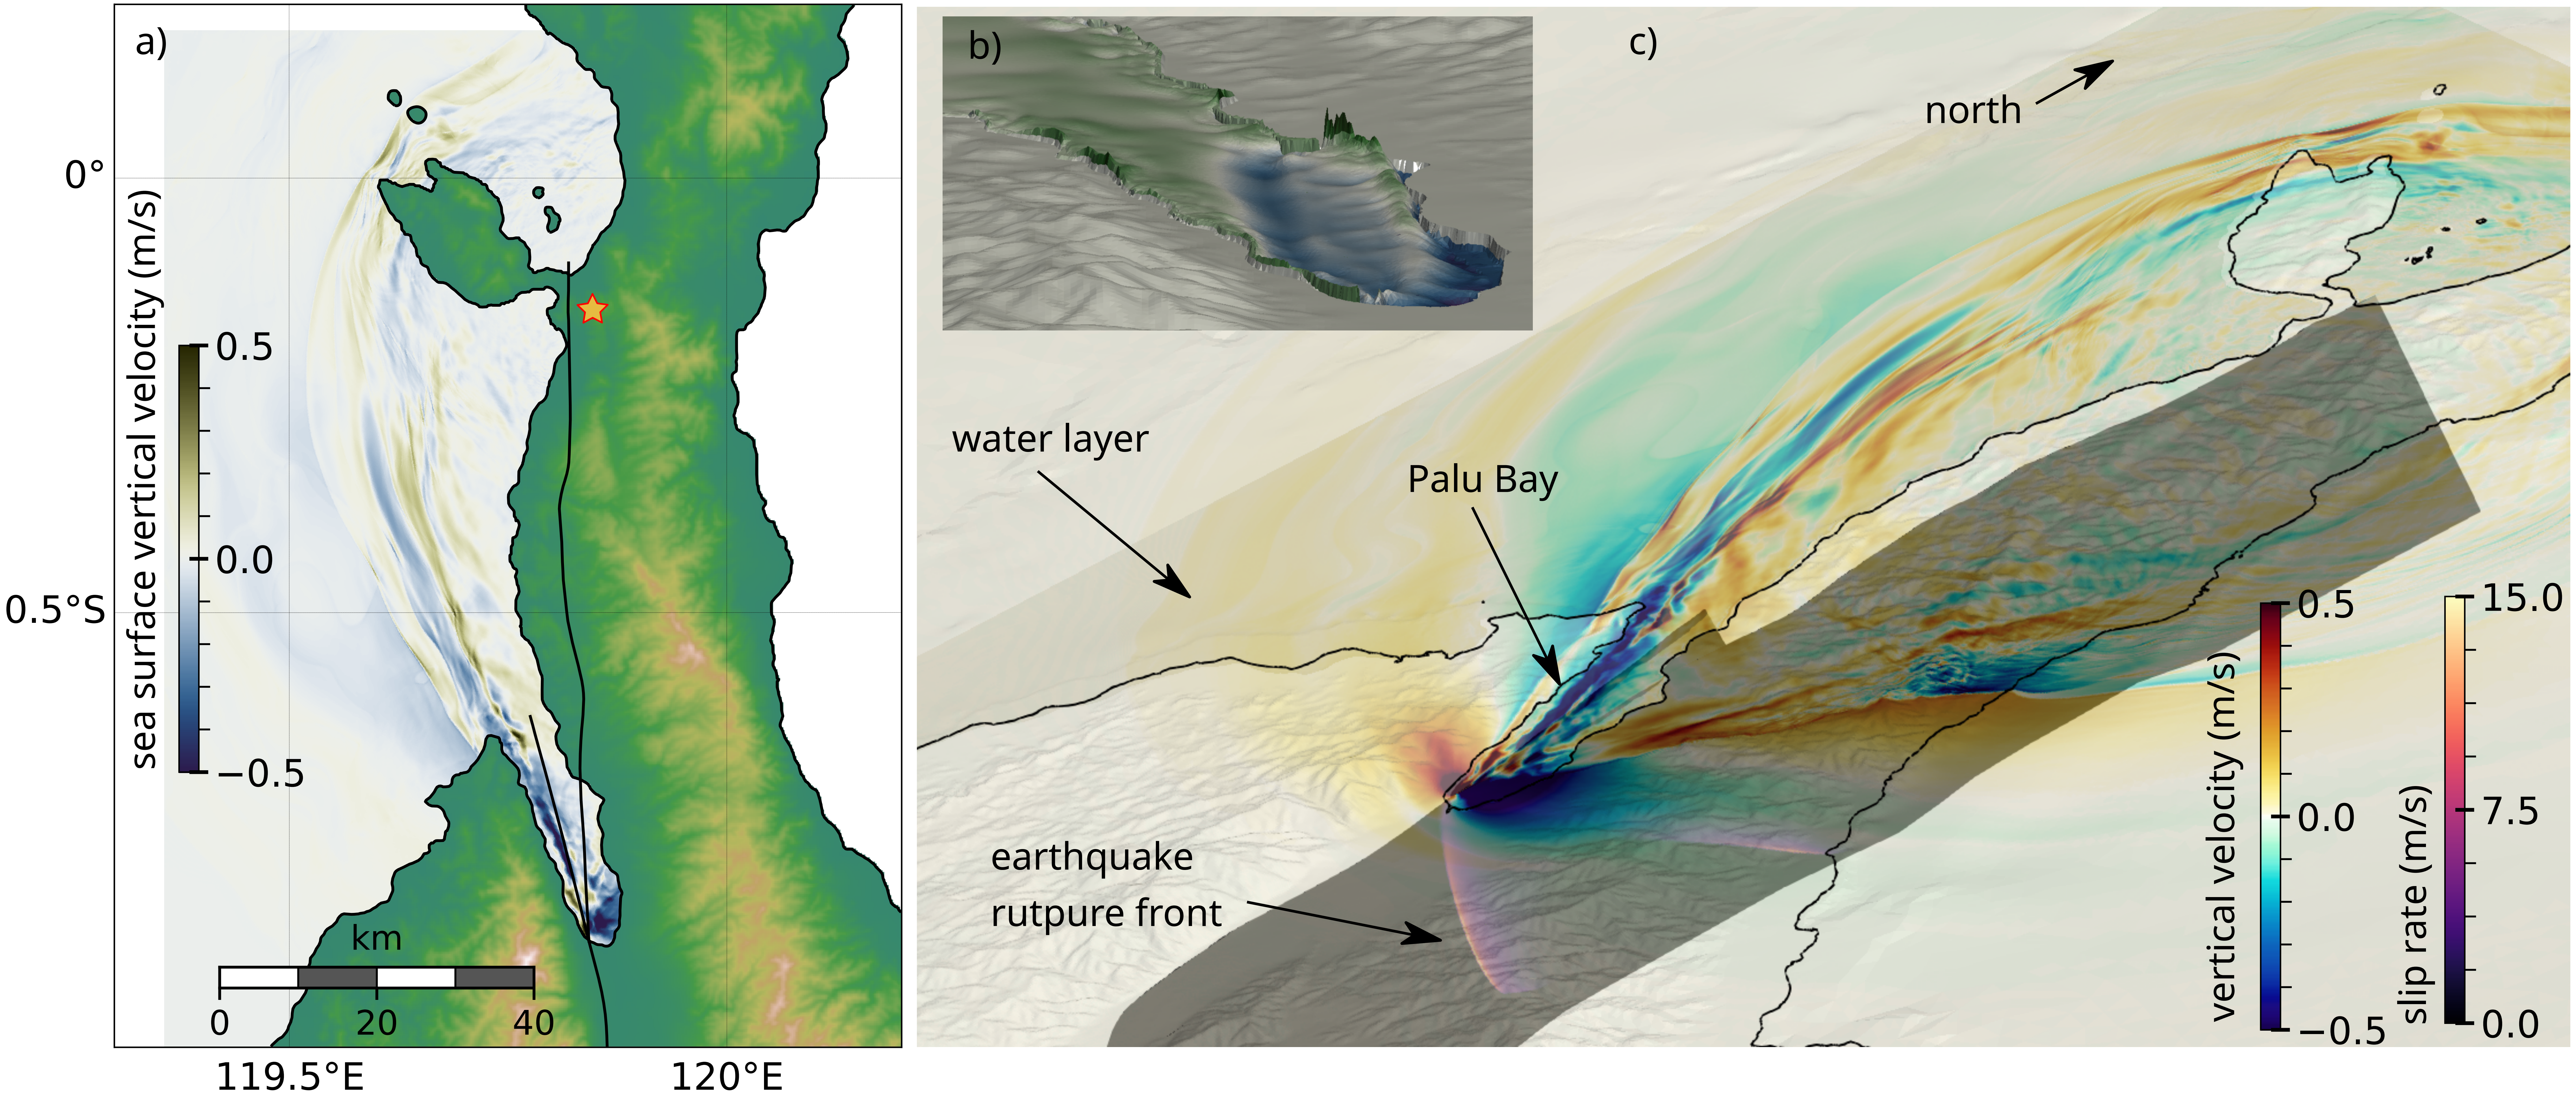

- we solve the **elastic wave equation coupled to non-linear frictional sliding in a complex fault network + the acoustic wave equation, describing perturbations about an equilibrium hydrostatic state in a compressible, inviscid ocean of variable depth + the effects of gravitational restoring forces through a modification of the standard free surface boundary condition**
- resolving wave excitation of up to **30 Hz** in the Fourier spectra of the recorded acoustic waves

Training example 2 - Community benchmark dipping normal fault with off-fault plasticity

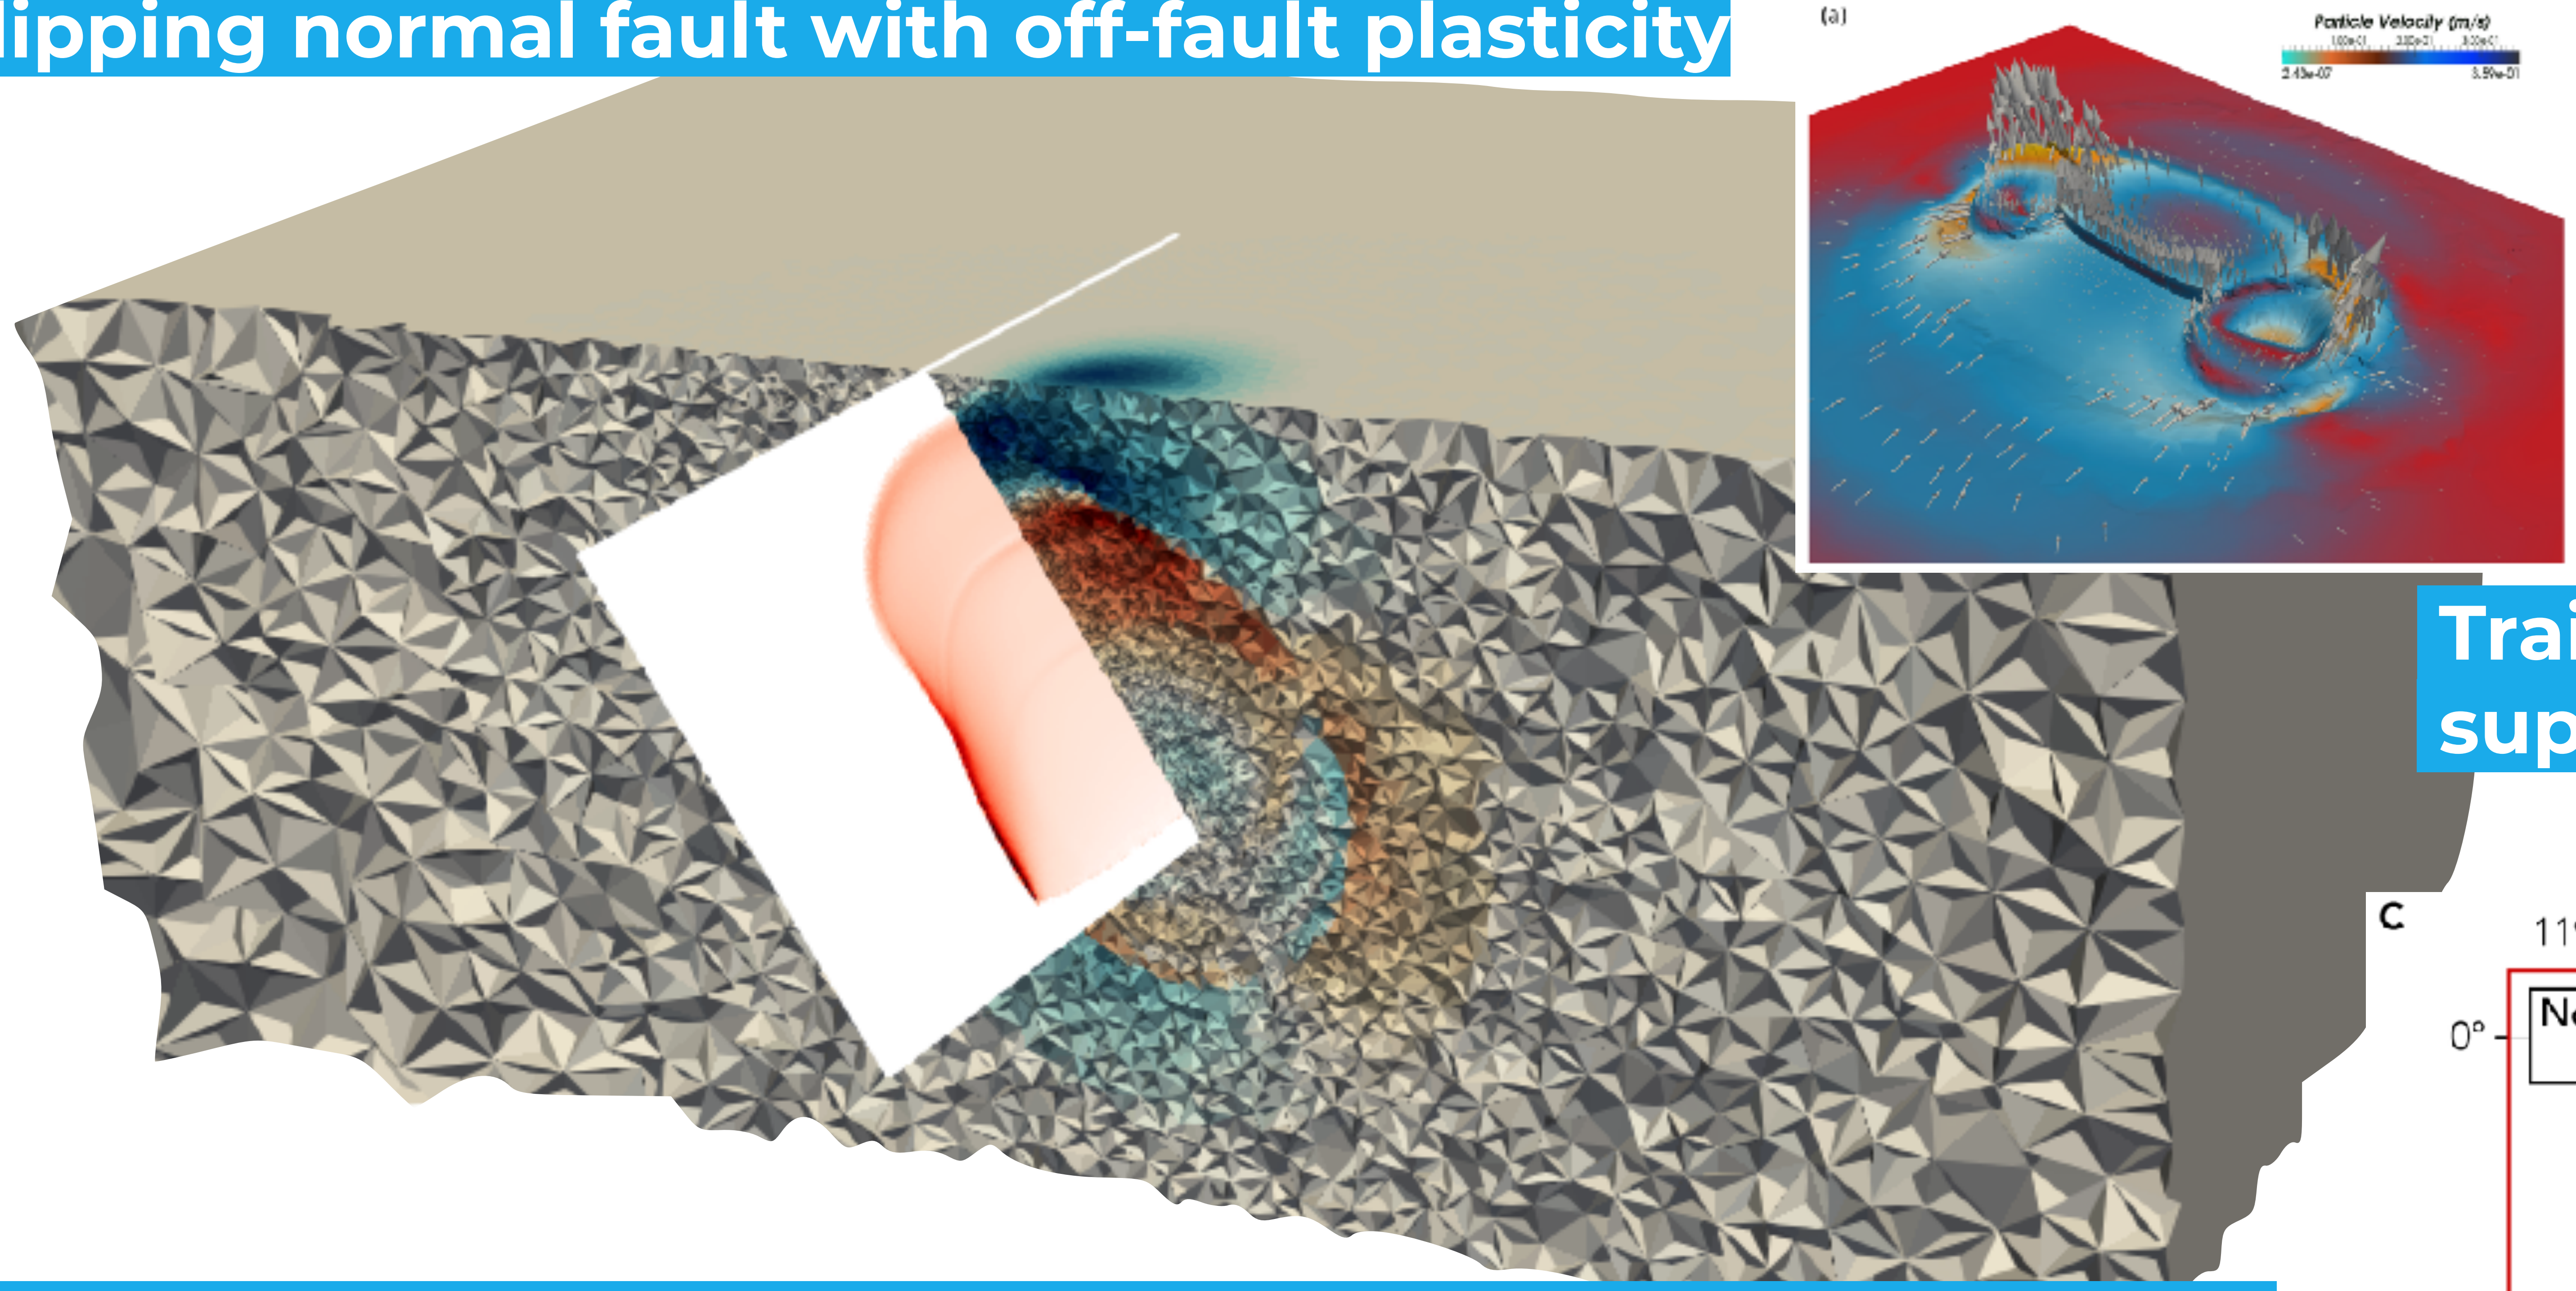

Training example 3 - Palu, Sulawesi strike-slip supershear earthquake

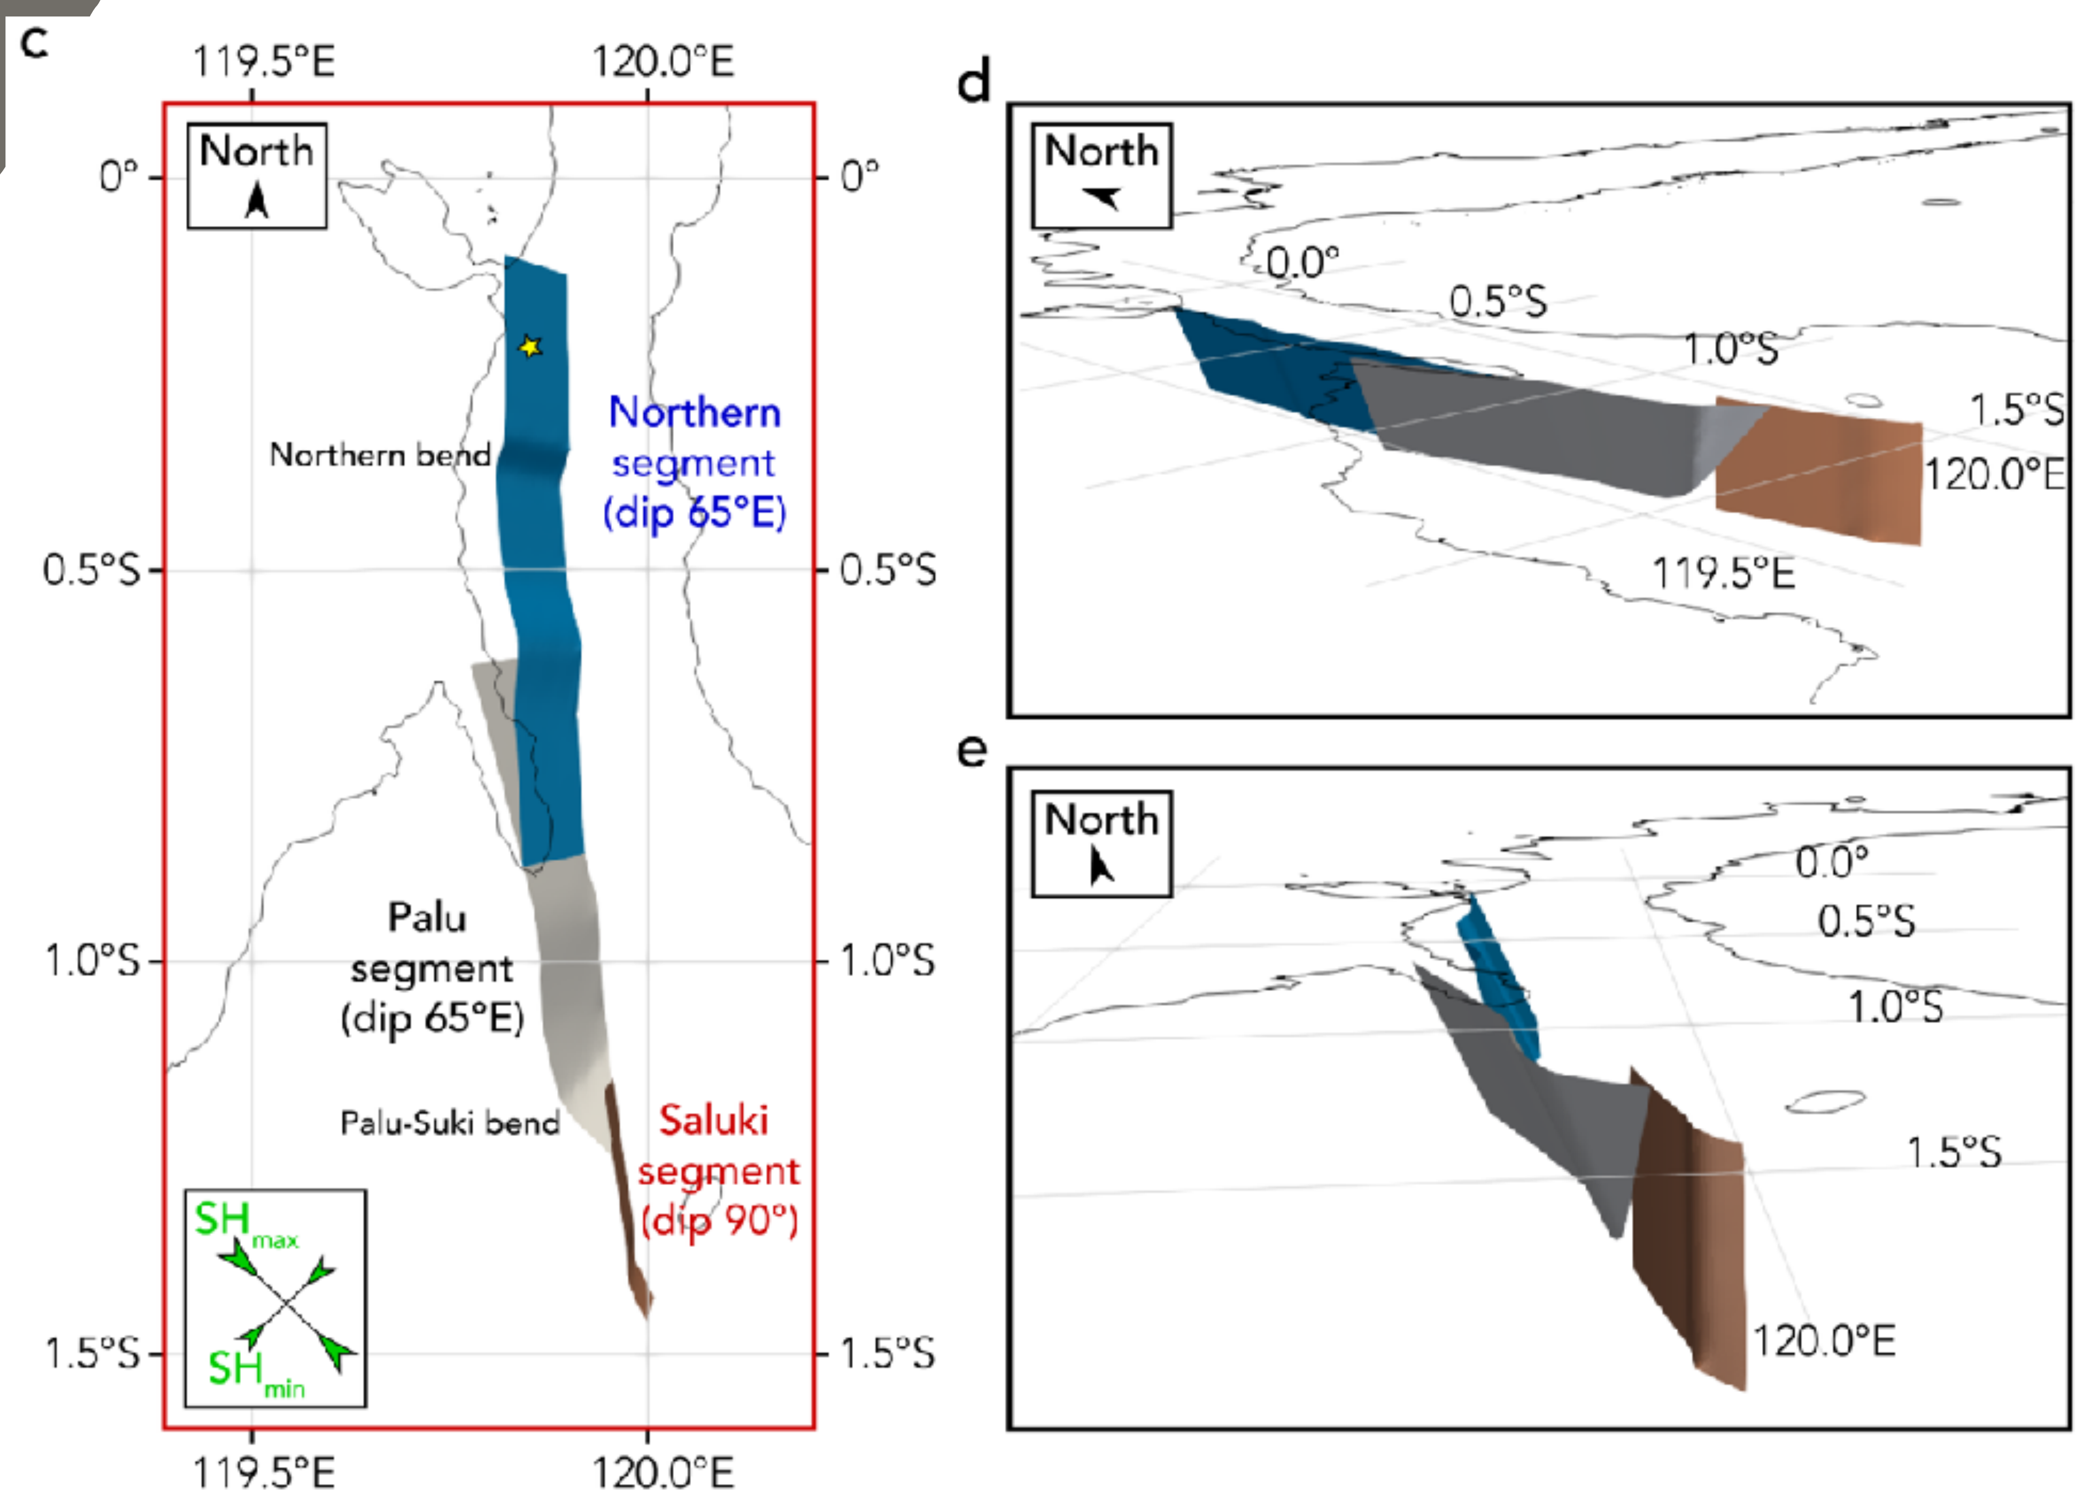

Training example 1 - kinematic finite source model and wave propagation with attenuation and topography

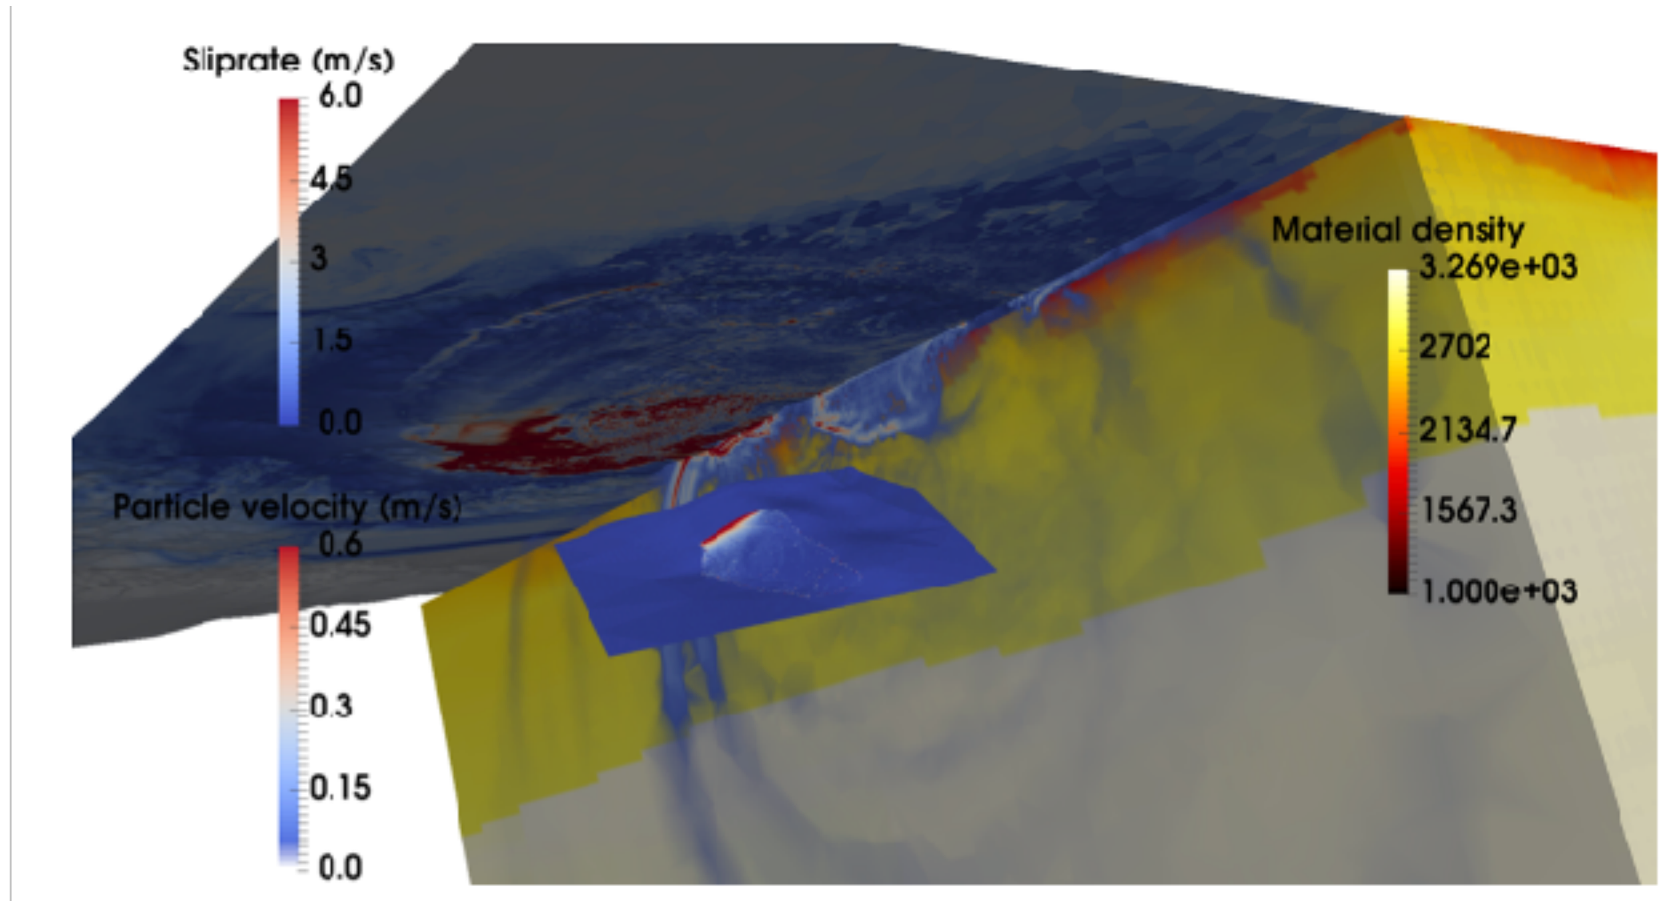

Summary

- **Computational earthquake seismology provides mechanically viable insight into the physical conditions that allow earthquake rupture on complex fault systems and helps constraining competing views on earthquake behaviour**

- **Geo-data, specifically community models, can be routinely included; Observational methods can themselves be constrained**

- **Bridging scales by coupling to tsunami, global seismic wave propagation, engineering, intermediate and long-term geodynamic modelling**

- **The interplay of advances in high-performance computing and dense observations will allow us to go beyond scenario-based analysis, aiming for, e.g., fully non-linear source-path-site effects, urgent response, data-driven dynamic source inversion, (Bayesian) uncertainty quantification, …**

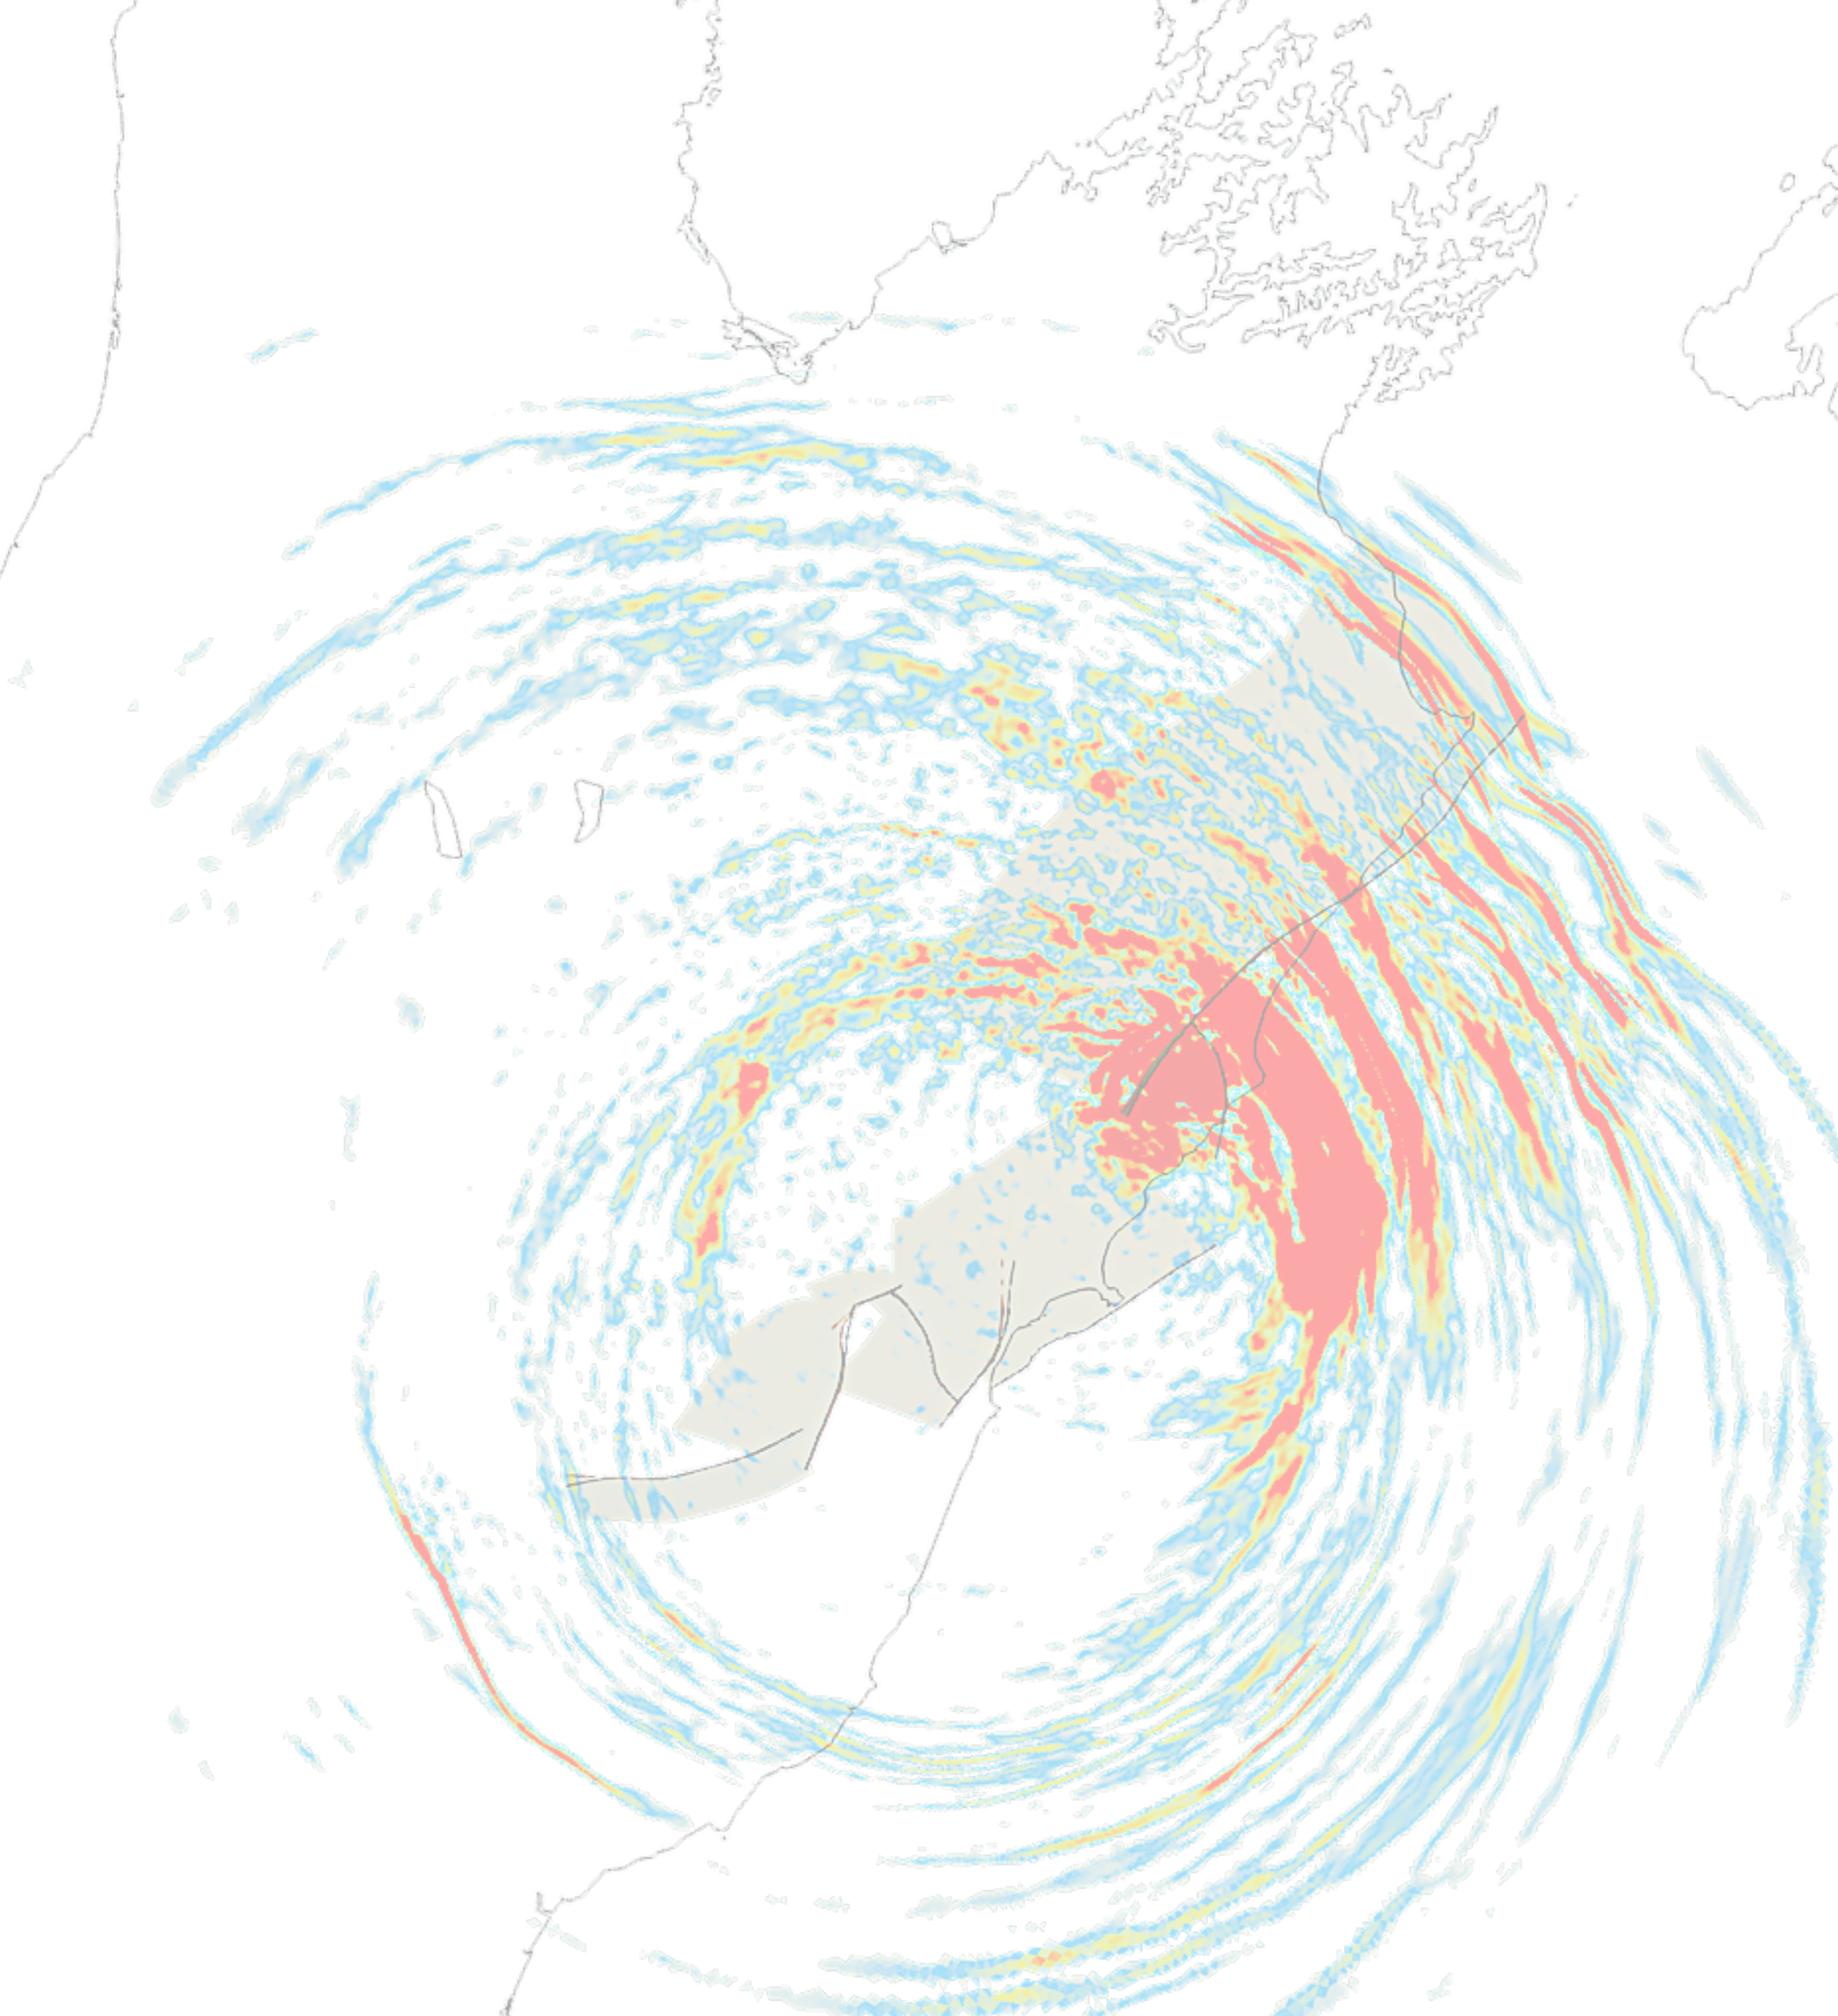

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

This afternoon:

—> After installation (with the docker engine running) type

```
$ docker run -p 53155:53155 uphoffc/seissol-training
```

or run the start.sh script.

After some time you should see

http://127.0.0.1:53155/lab?token=some5cryptic8hash123

Click on that link or enter the link in the address bar of your favourite web browser.

Then use the navigation bar to open the exercises (e.g., tpv13/tpv13.ipynb).