SCRIPPS INSTITUTION OF OCEANOGRAPHY
UC San Diego

3rd GeoInquire Online Seminar
June 22, 2023

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Digital twins for earthquakes: connecting high-performance computing with observations

Alice Gabriel (and many others!)

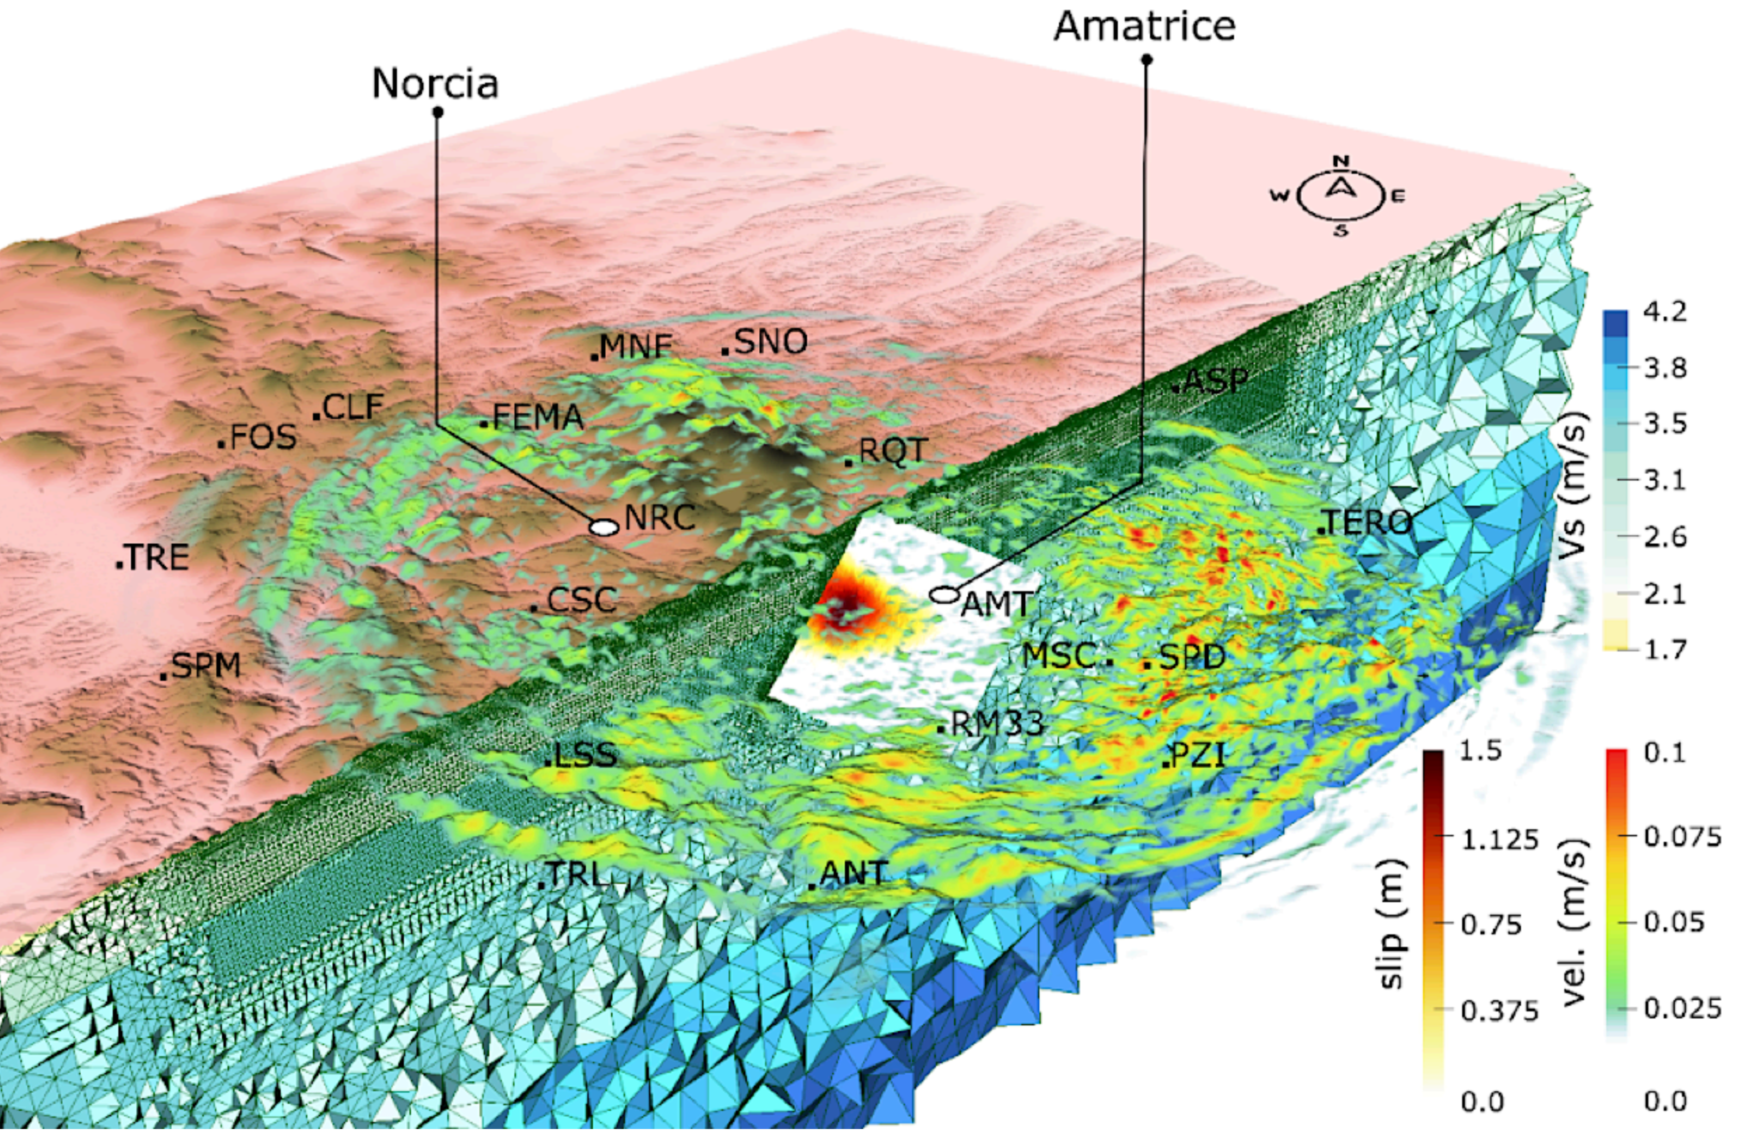

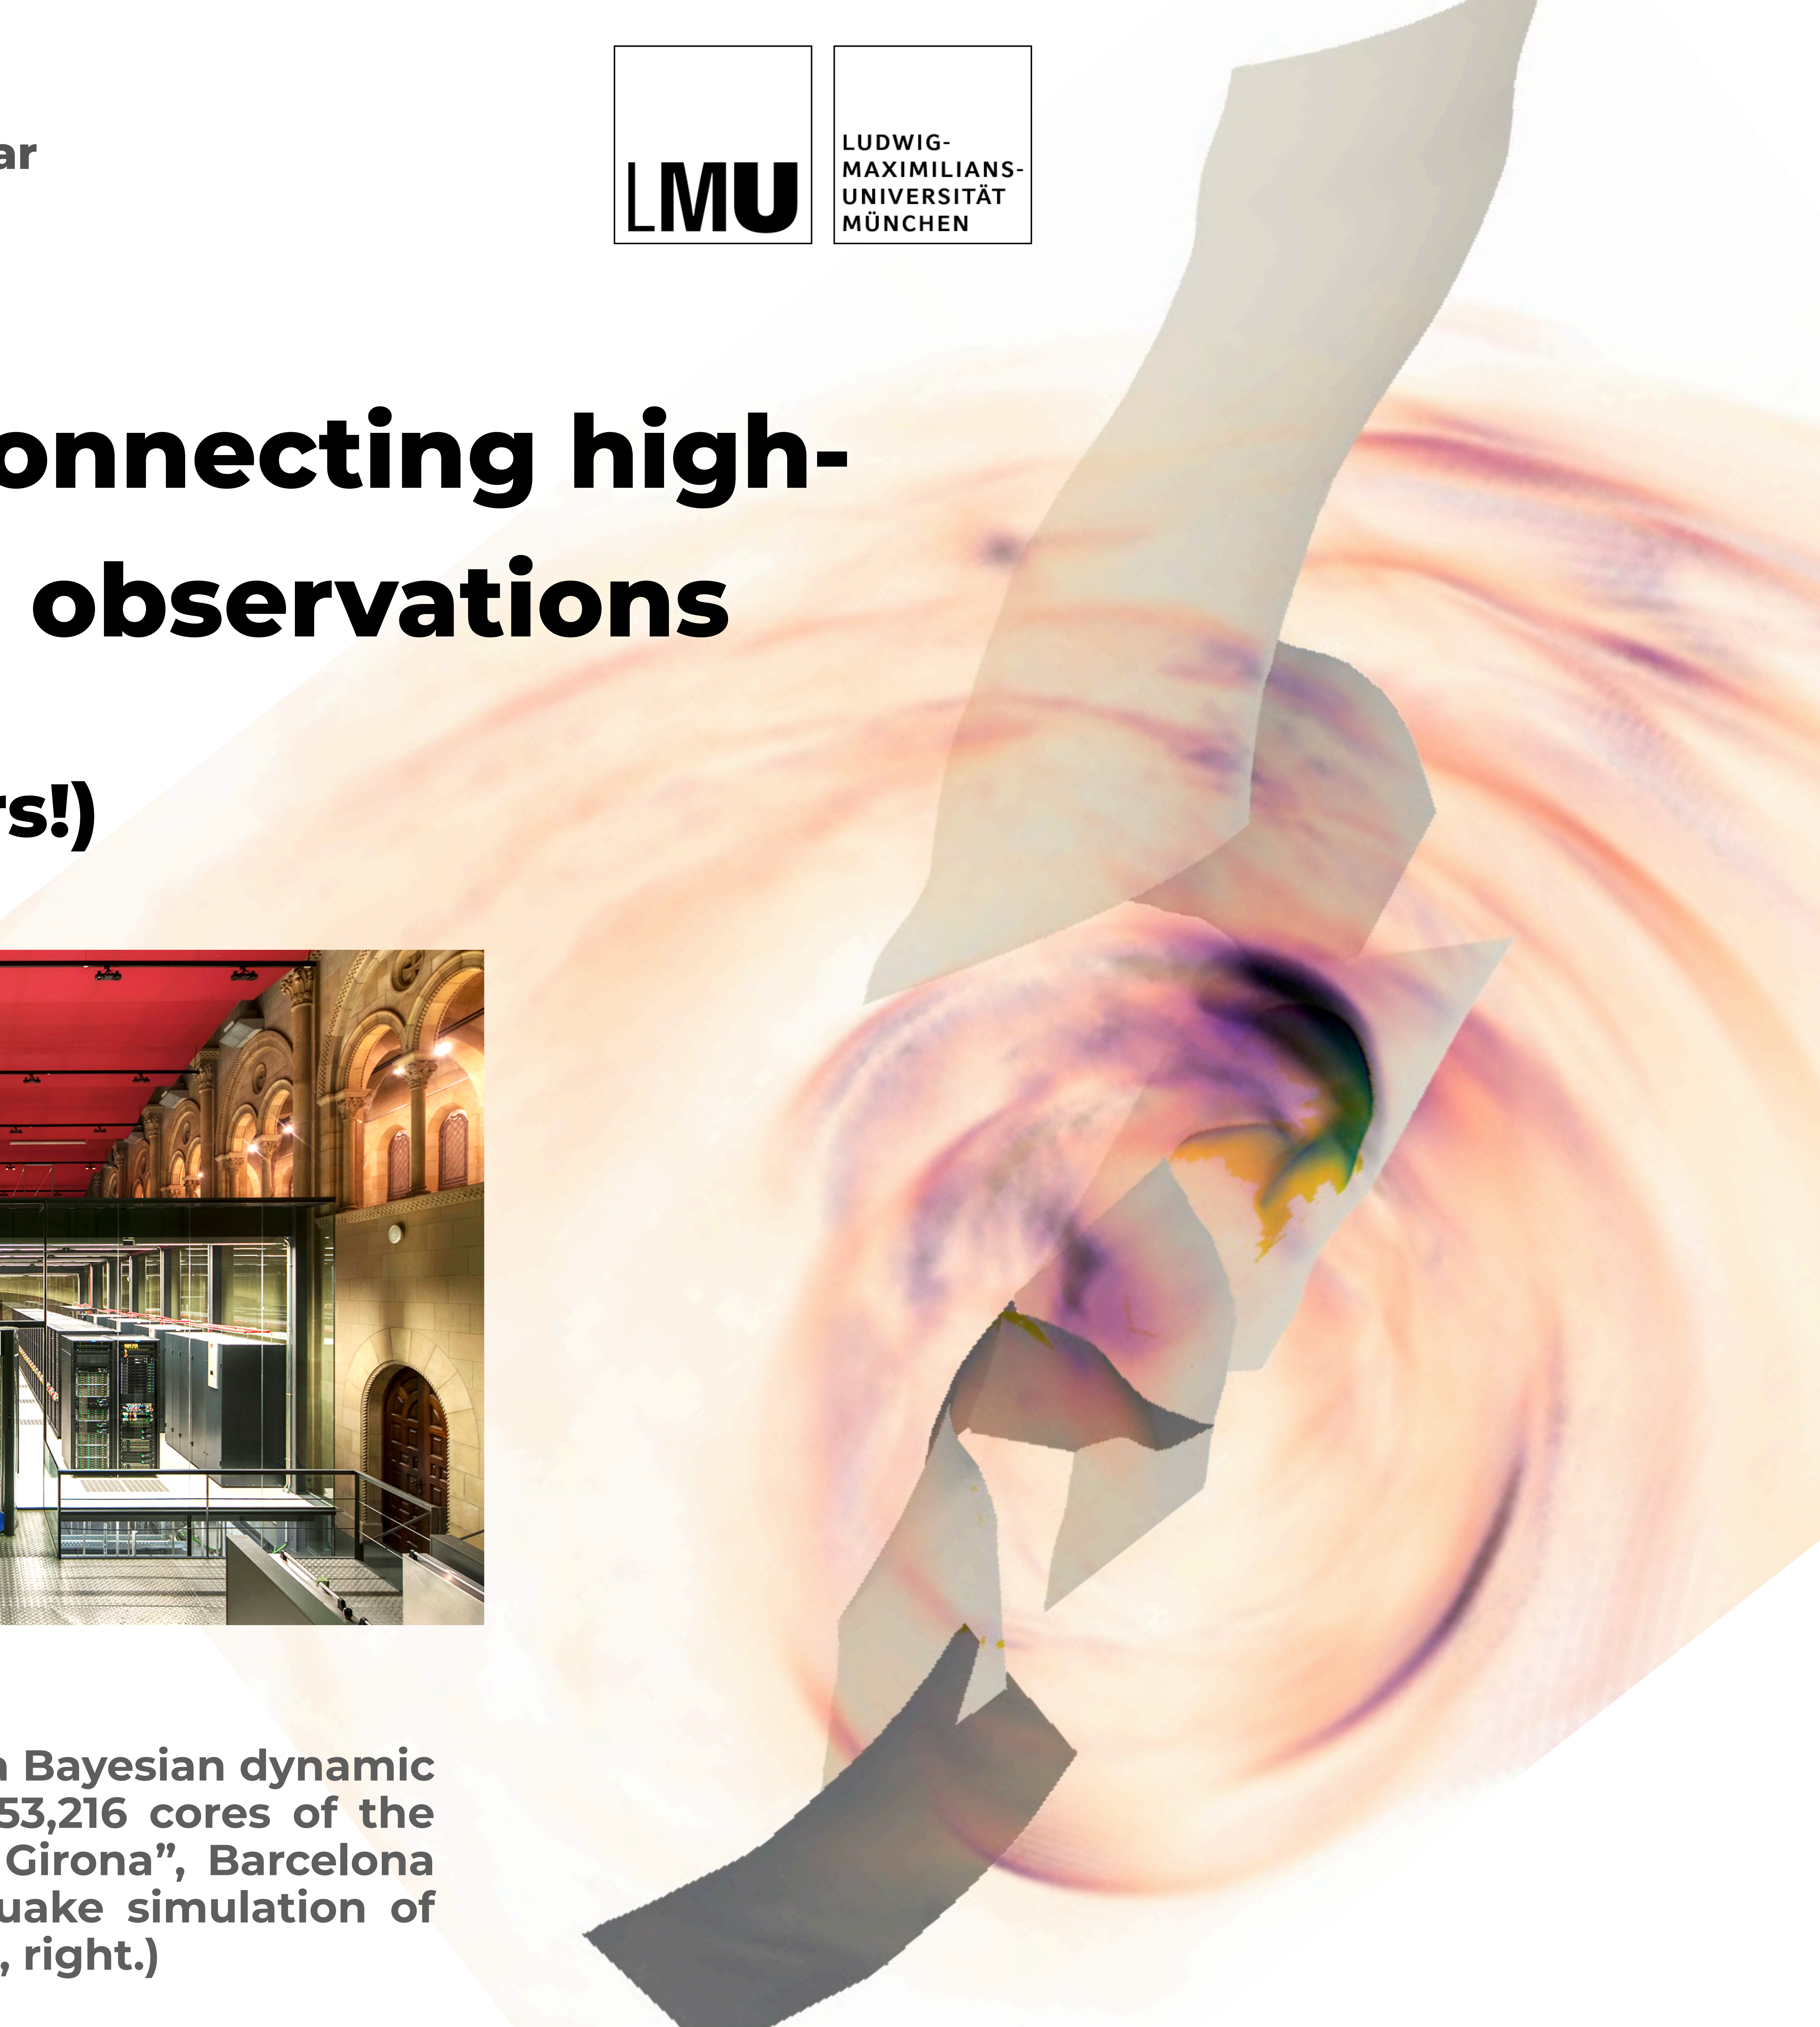

Broadband dynamic rupture modeling based on ~1 million models visited in a Bayesian dynamic strong motion source inversion (Taufiqurrahman et al., GRL'22, left). The 153,216 cores of the MareNostrum supercomputer housed in the deconsecrated Chapel "Torre Girona", Barcelona (middle). Observationally constrained fully dynamic, physics-based earthquake simulation of the 2016 Kaikoura, NZ, multi-fault earthquake (Ulrich et al., Nat. Comms., 2019, right.)

Earthquake seismology is increasingly **data-rich**

- **machine-learning** enhanced seismic catalogues, space **geodesy**, **array-**data processing, Distributed Acoustic Sensing (**DAS**)
- well-recorded earthquakes (and laboratory experiments) reveal striking **variability of earthquake dynamics**

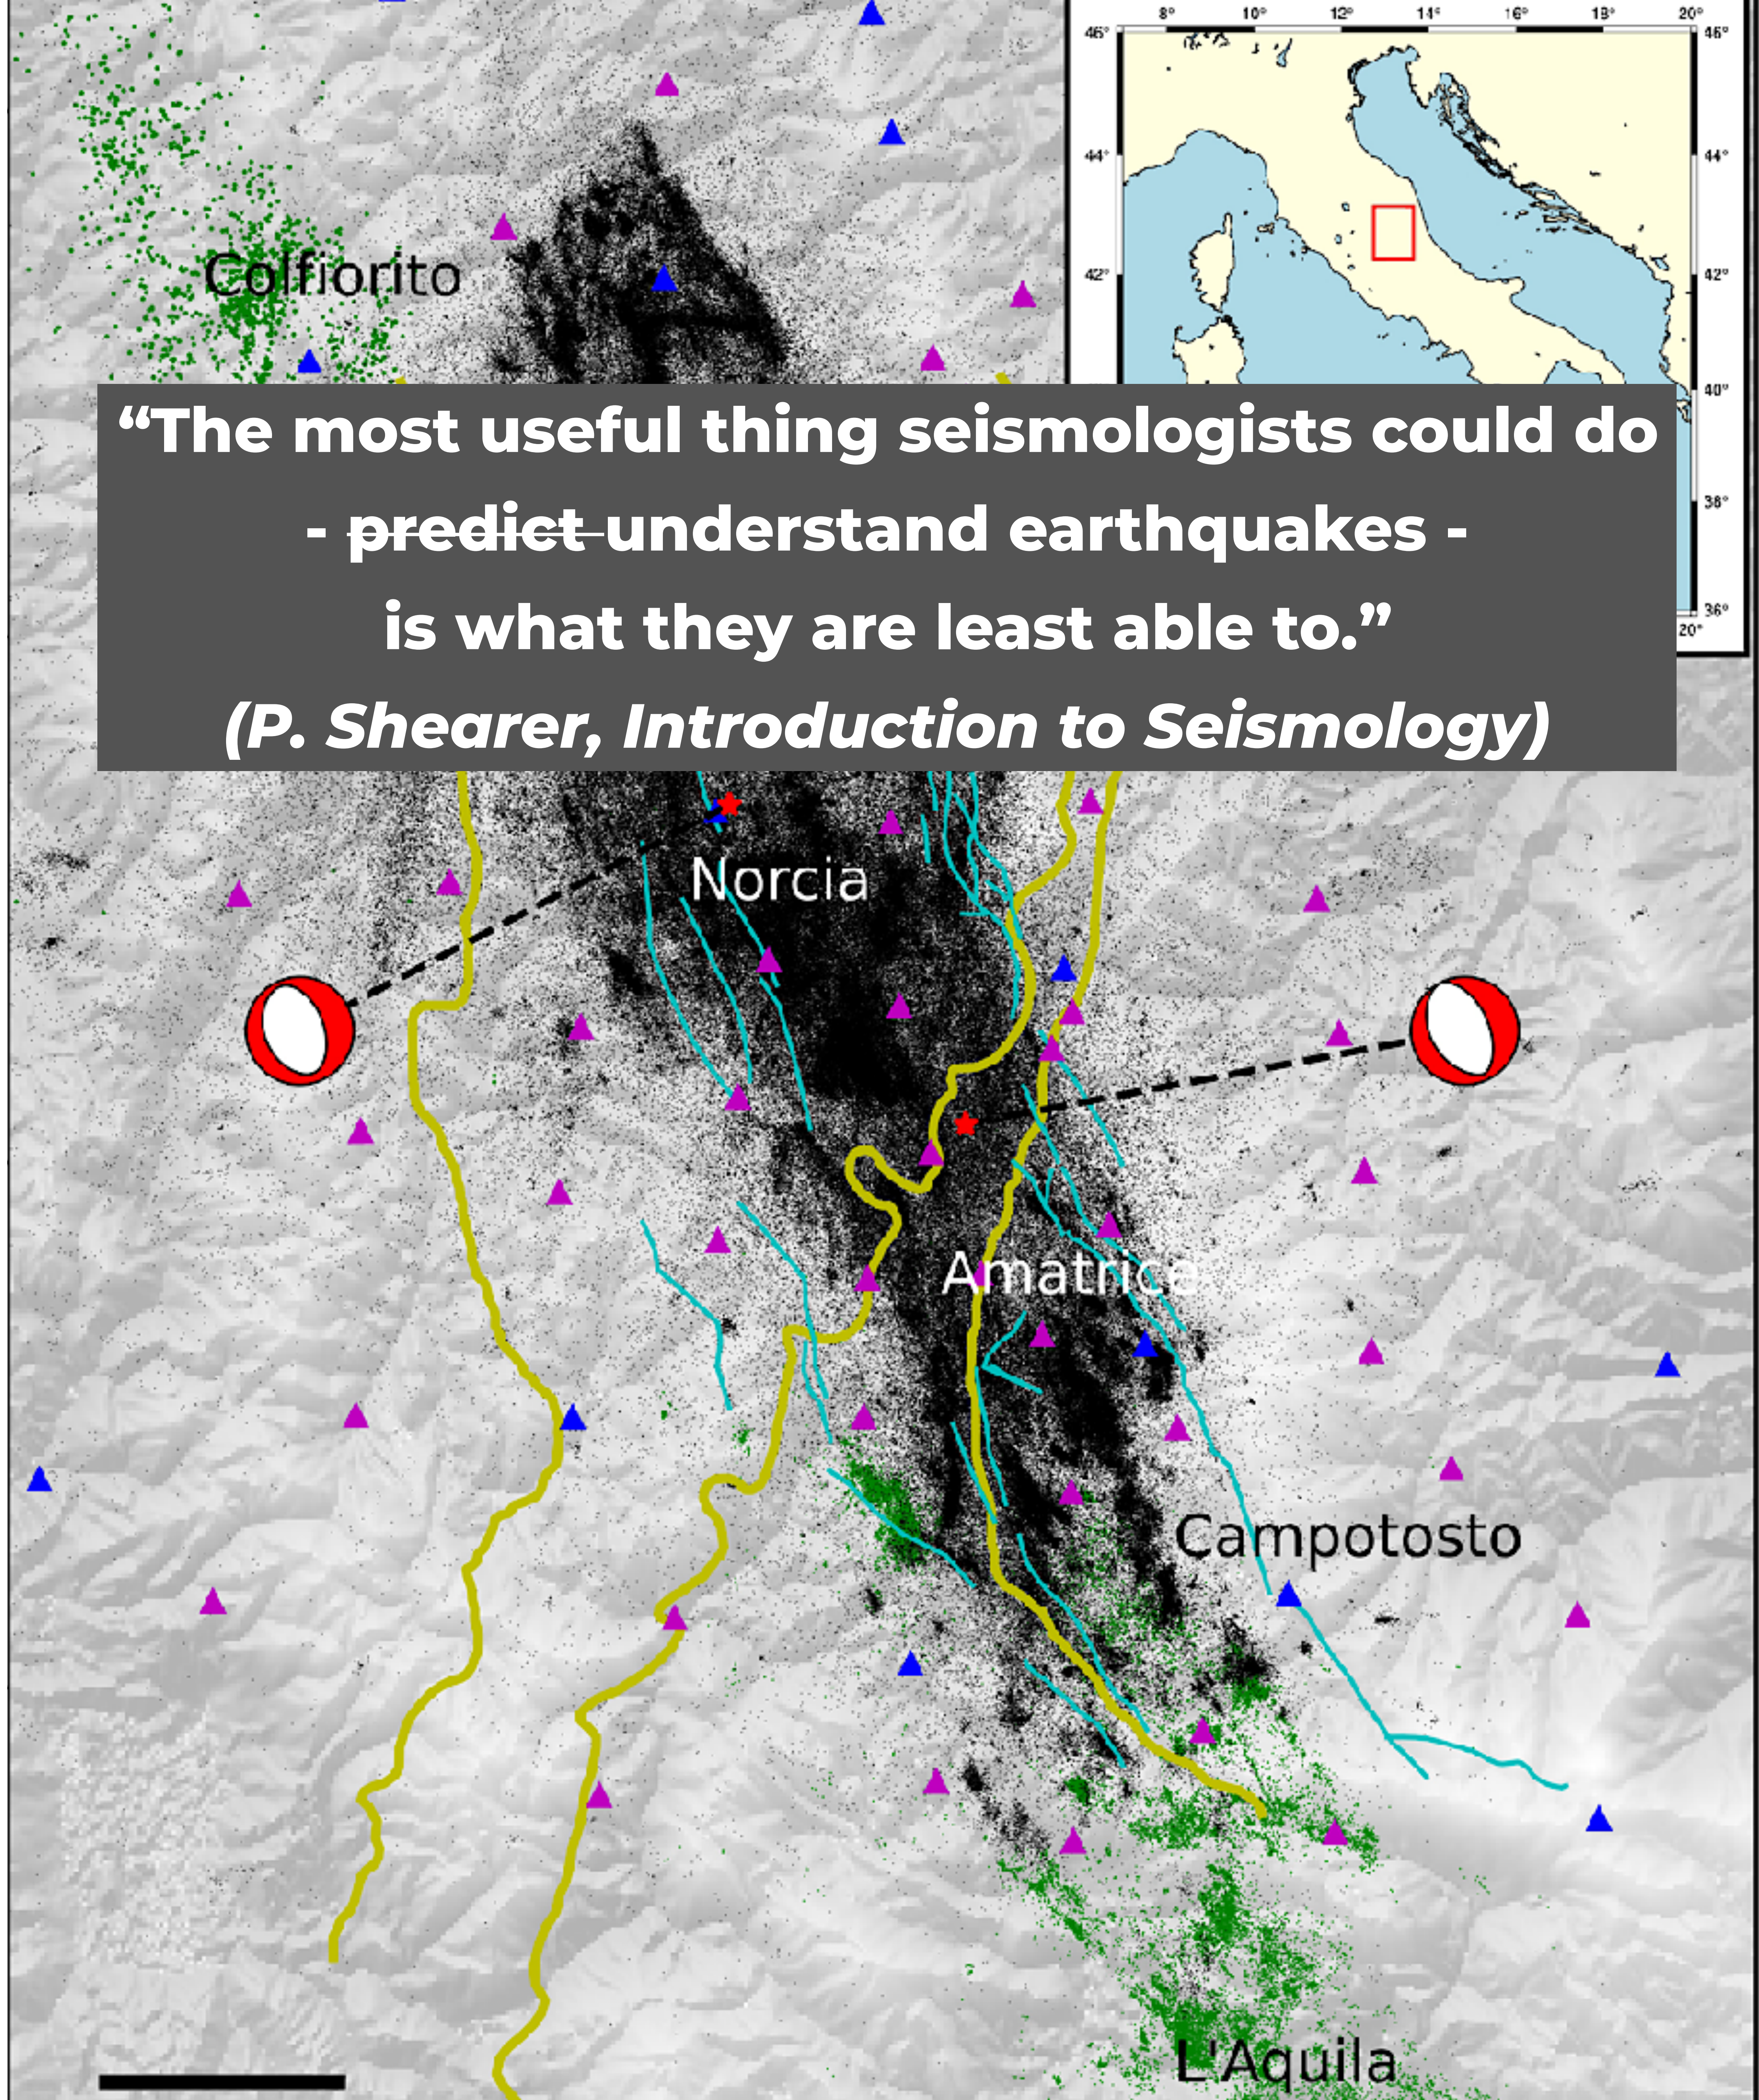

"The most useful thing seismologists could do - ~~predict~~ understand earthquakes - is what they are least able to."
(P. Shearer, Introduction to Seismology)

Machine-Learning-Based High-Resolution Earthquake Catalog Reveals How Complex Fault Structures Were Activated during the 2016–2017 Central Italy Sequence. Tan et al., TSR, 2021

Source inversion model of Tohoku-Oki event (Japan) 2011, from combined local ground motion, teleseismics, GPS & multiple time window parametrization of slip rate. (Lee & Wang, 2011)

Earthquake seismology is increasingly data-rich, yet, remains **model-poor**

- Faults and fractures that host dynamic slip vary over more than **six orders of magnitude** in source dimensions

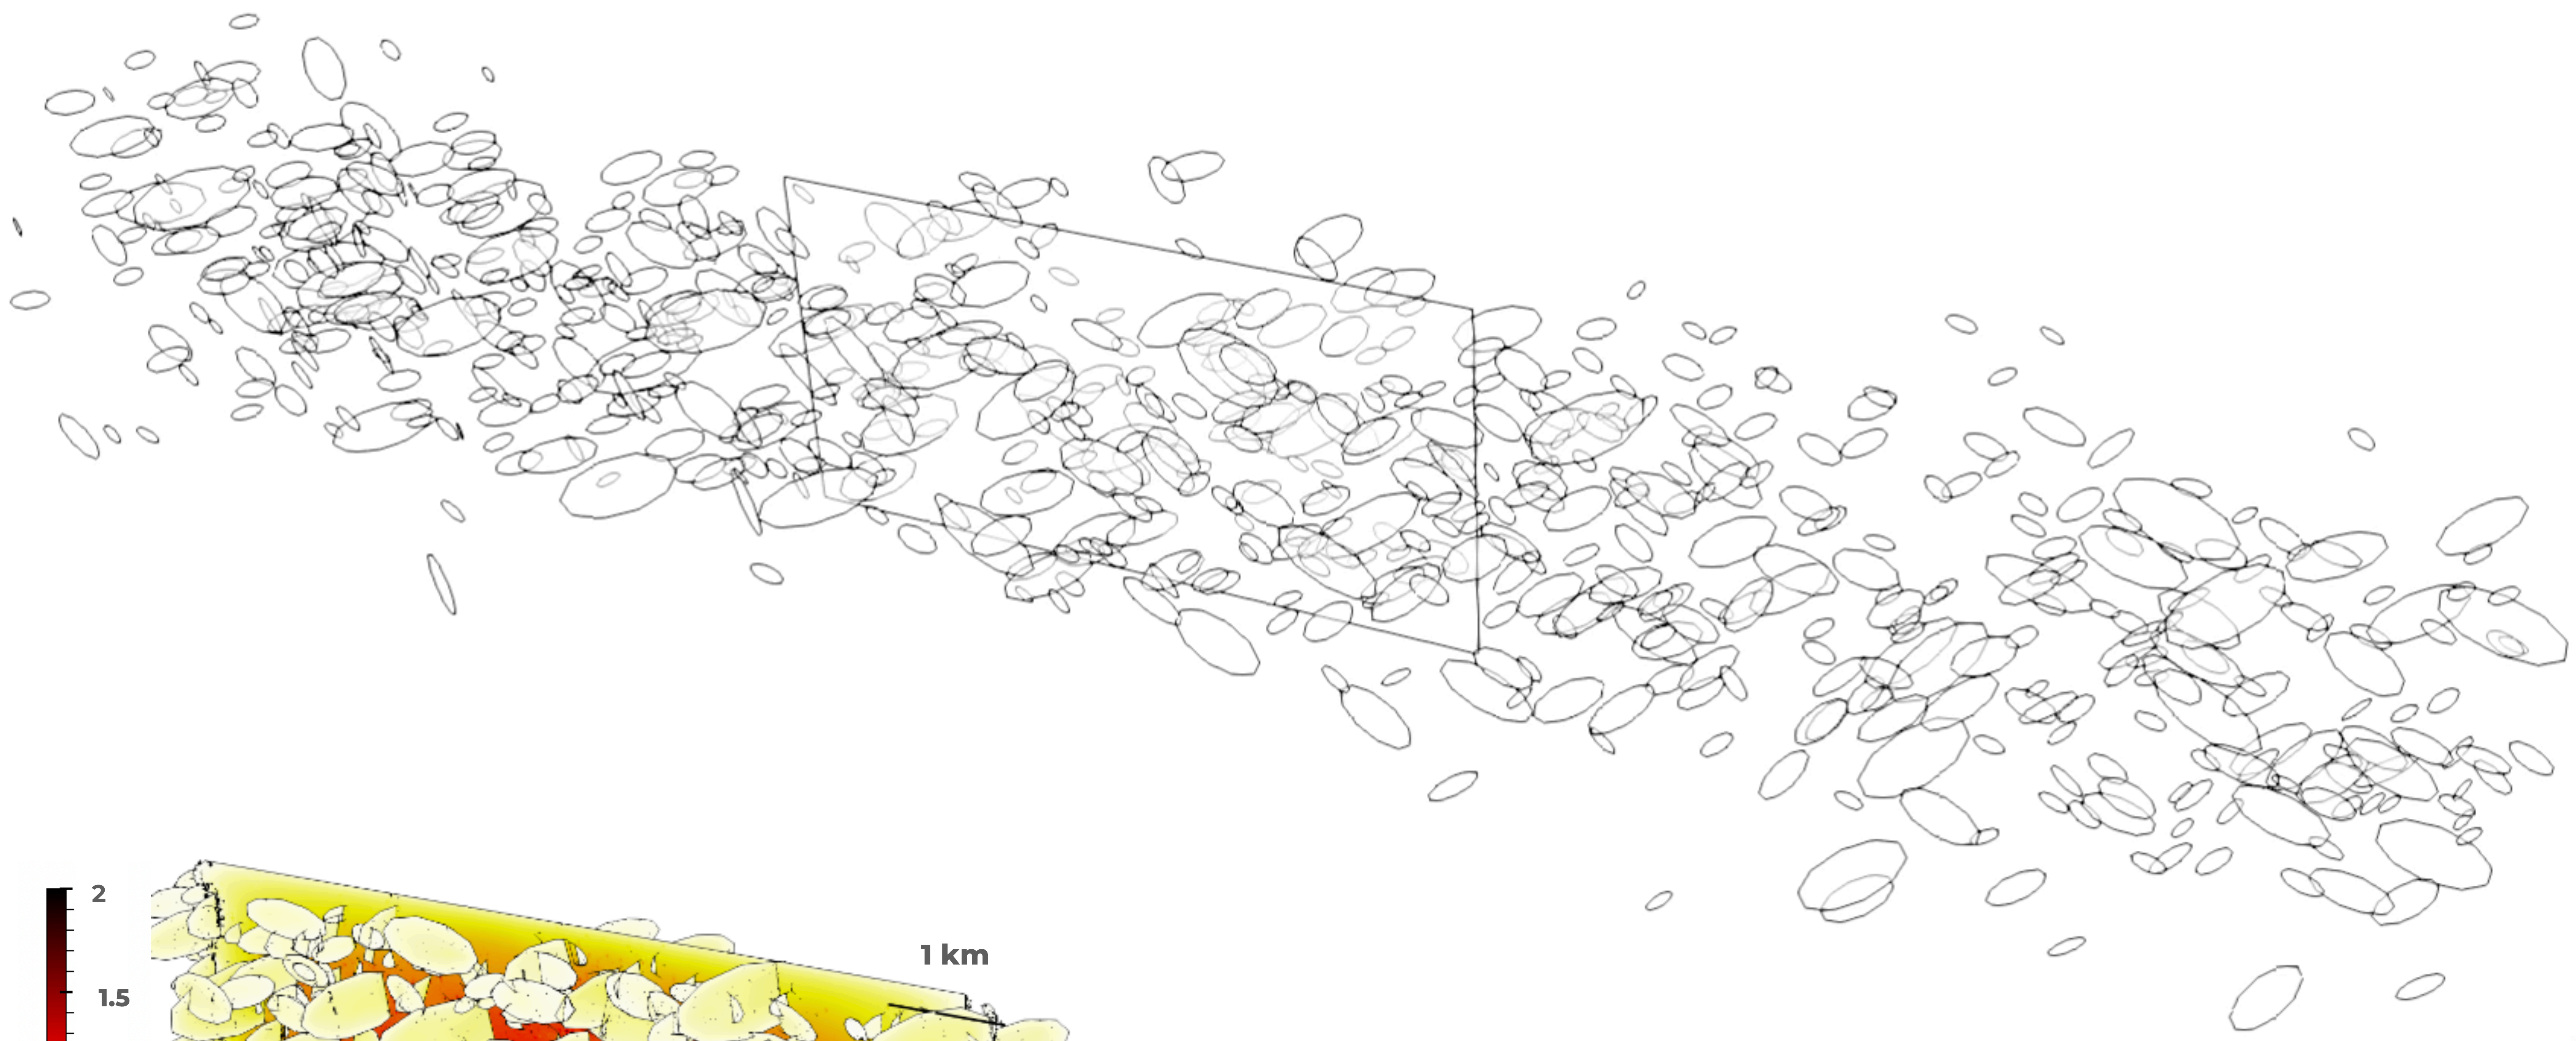

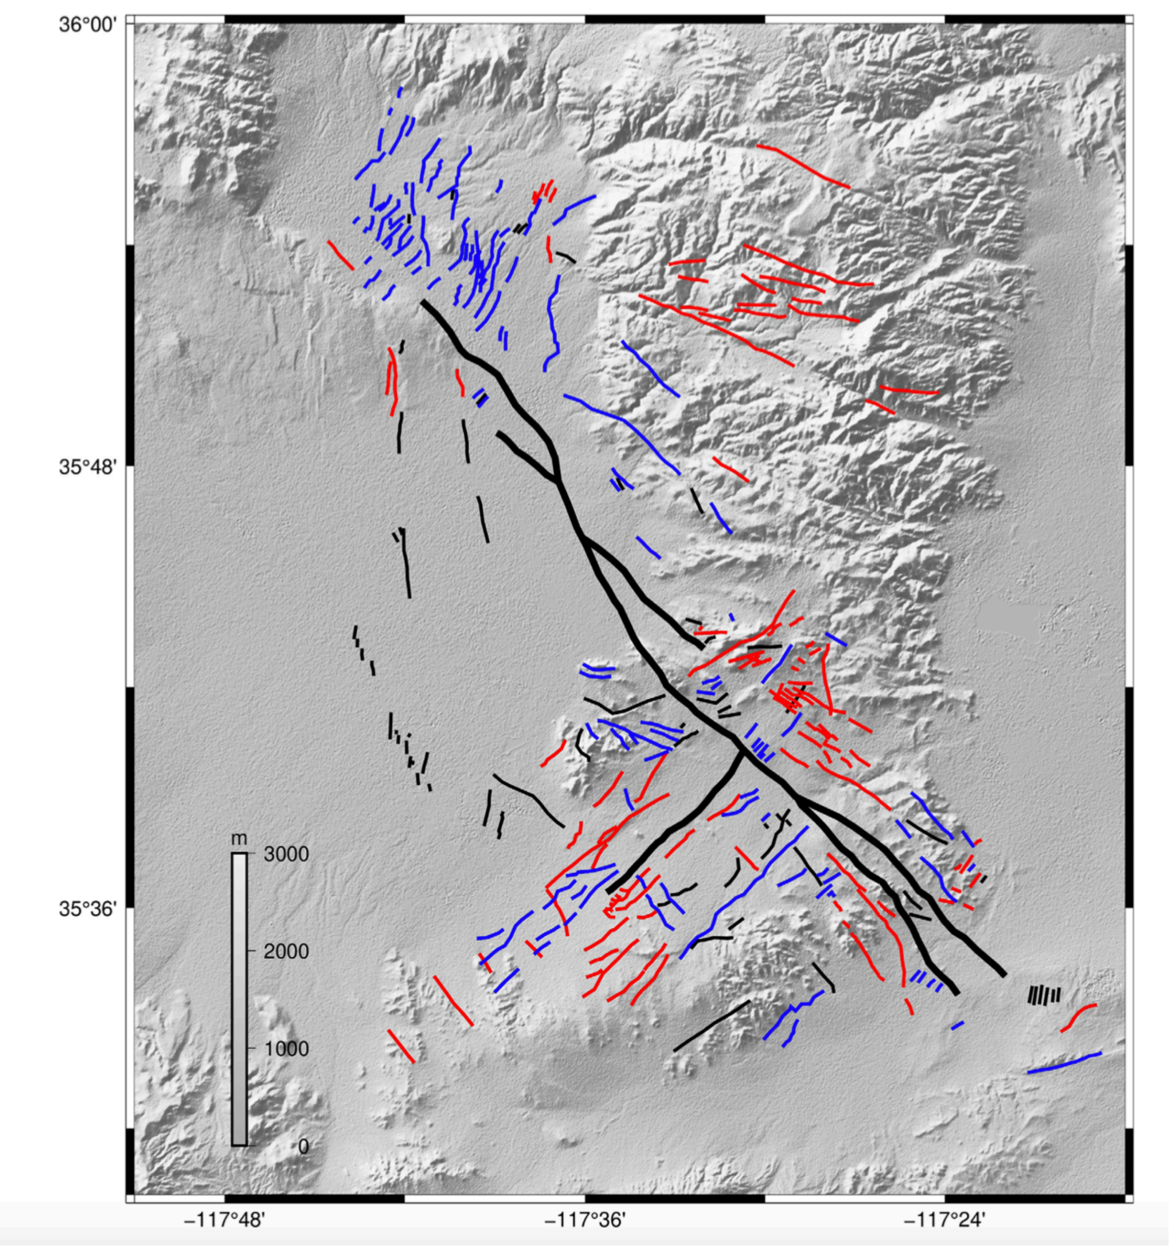

Inferred fracture map from InSAR data around the Ridgecrest earthquake sequence (Xu et al., '19)

Cascading fracture rupture cascade dynamically triggering a large main fault earthquake.

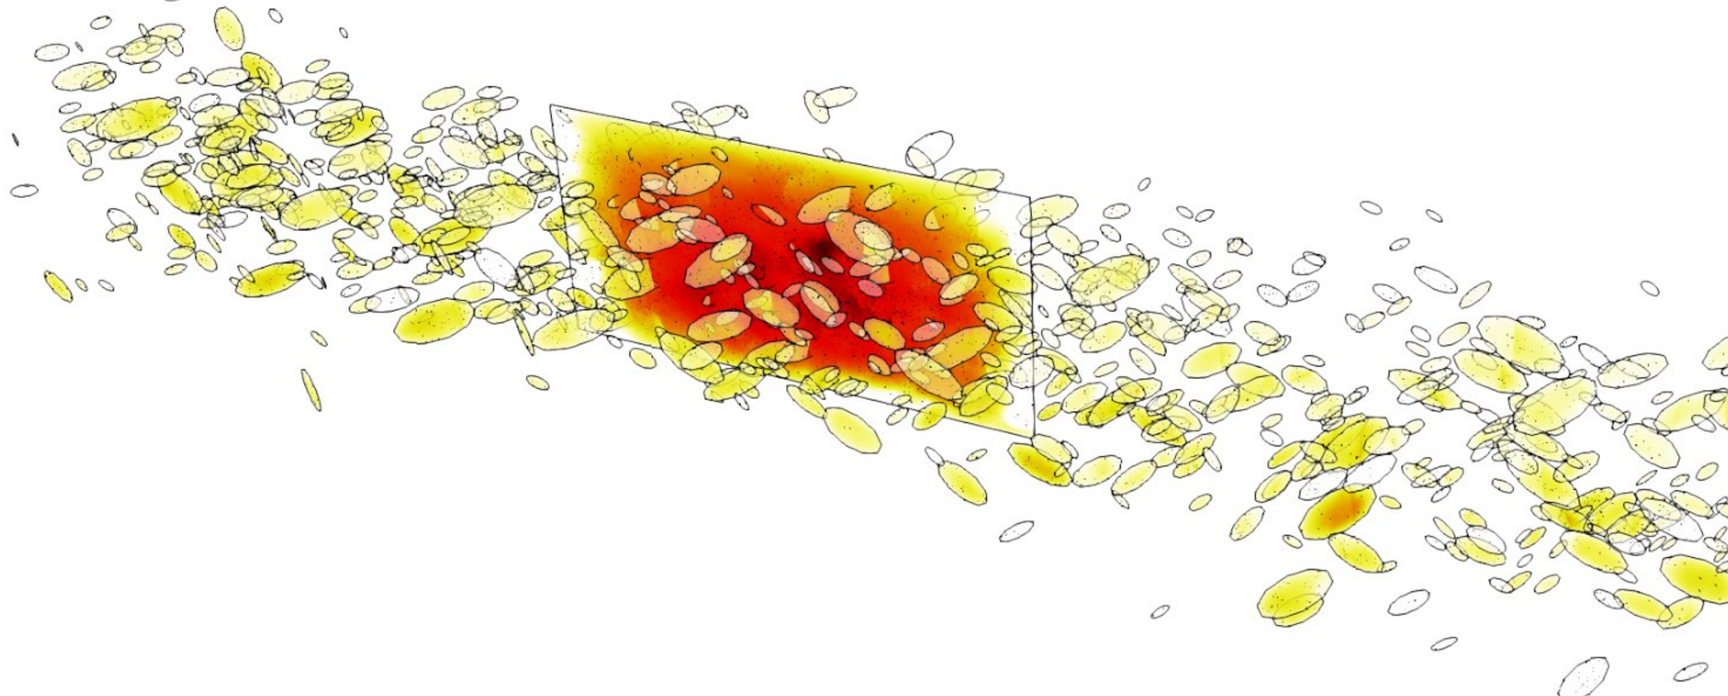

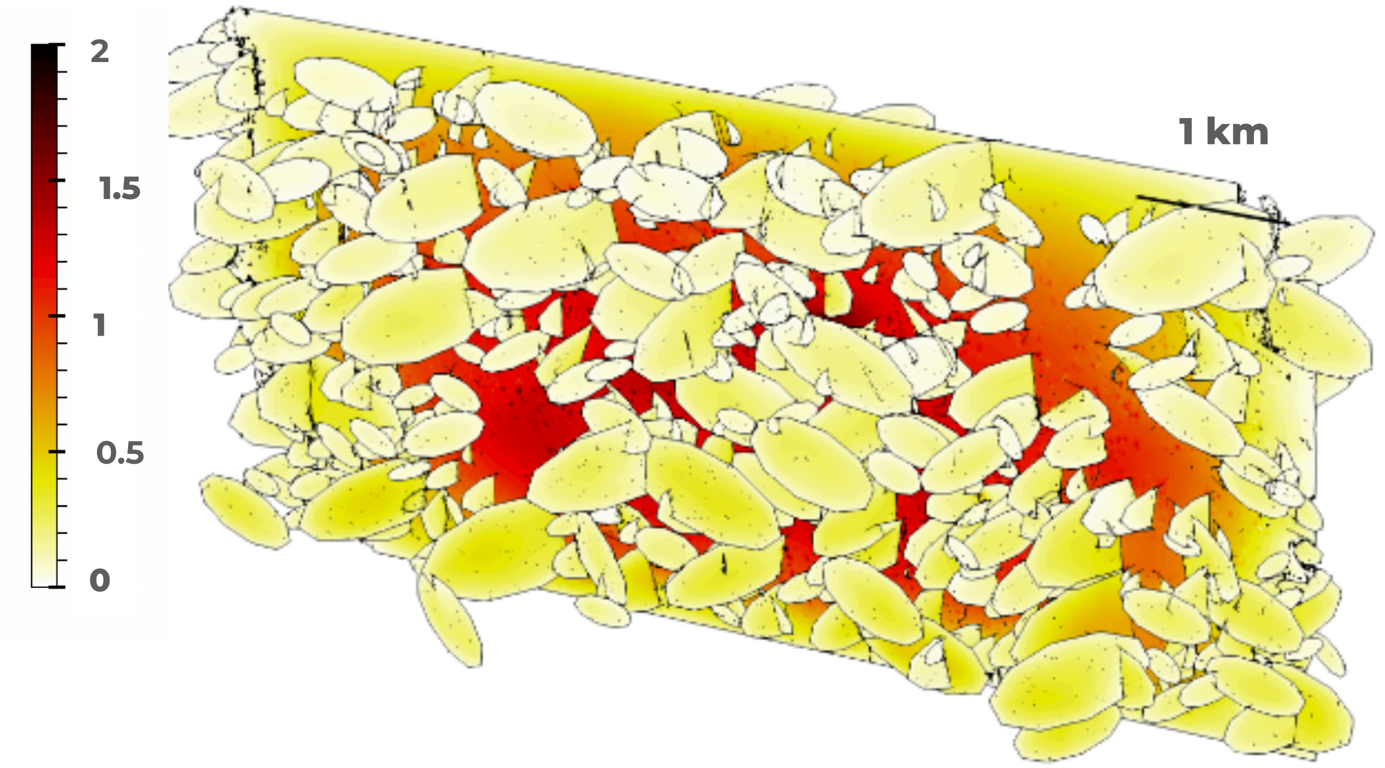

Multi-scale cascading dynamic earthquake rupture across a fault damage zone of ~900 fractures eventually activating the main fault, enabled by corrected scale-dependence of fracture energy (Gabriel et al., SCEC'22). Resolving seismic wave propagation up to 13 Hz requires 18h on 512 nodes (295,000 CPUh, Shaheen II)

Physics-based Earthquake Modeling

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

One of a suite of dynamic rupture simulations informing physics-based PSHA in North Iceland, Li et al., 2023

Physics-based Earthquake Modeling

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

- **"Bootstrapping":** on decade long developments in computational seismology and fracture mechanics

- Much complexity lives in the definition of **friction** (shear traction is bounded by fault strength) and **fault geometry**

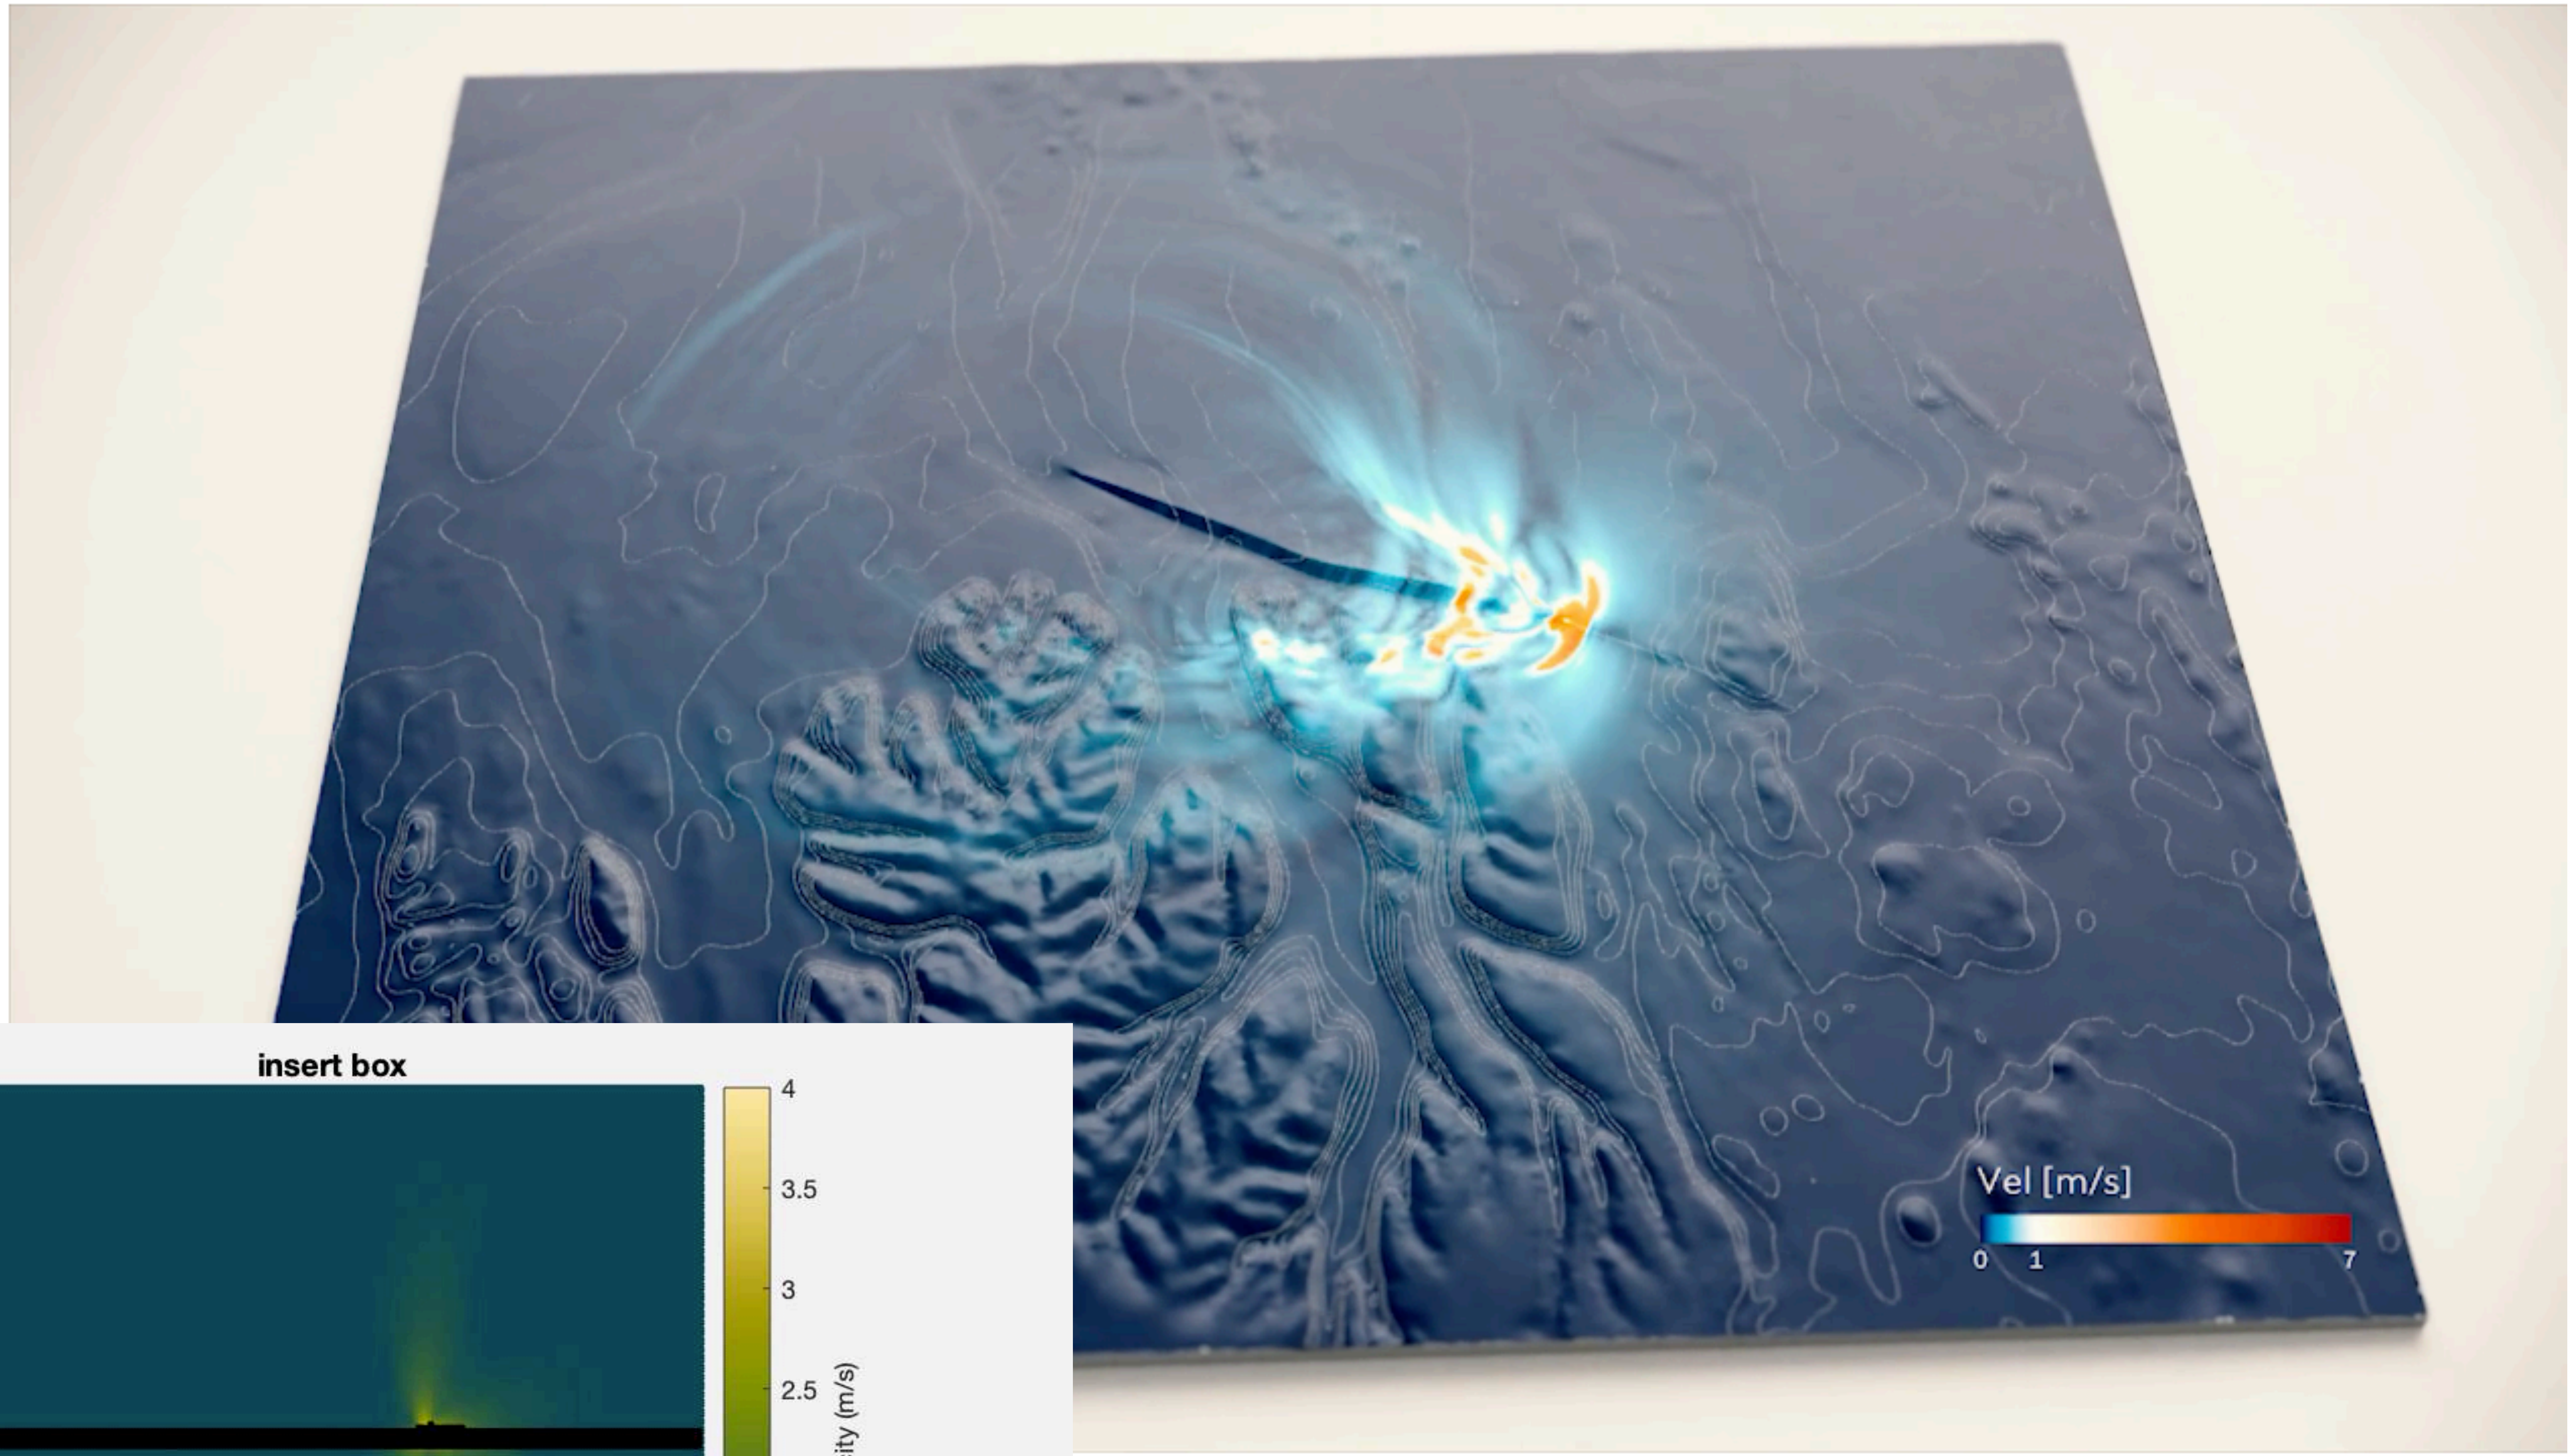

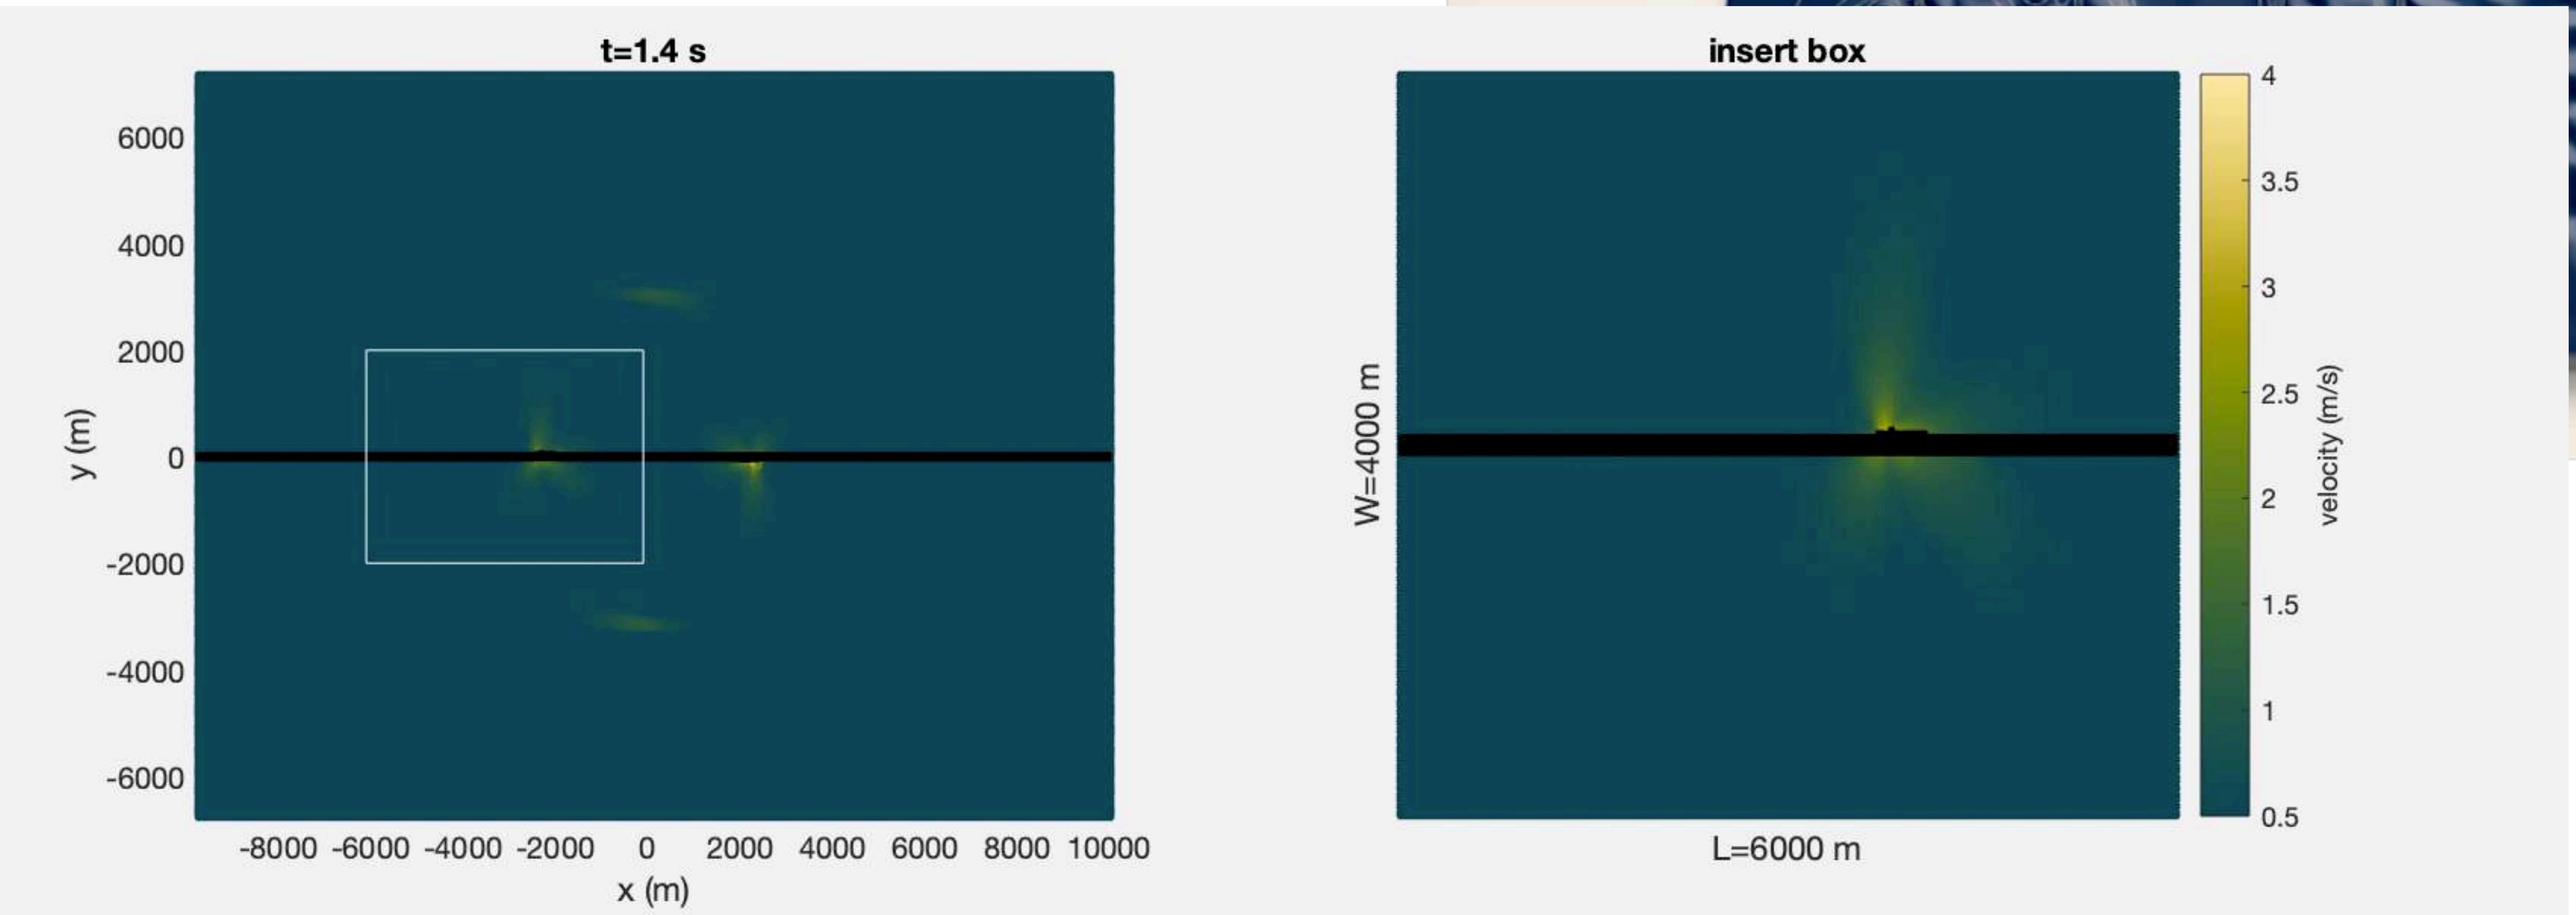

A unified first-order hyperbolic phase-field model for nonlinear dynamic rupture processes in diffuse fracture zones (Gabriel et al., 2021)

3D Dynamic Earthquake Modeling

Adapted from Harris et al., SRL 2011, 2018

- **3D models** strive for the **integration and interpretation of the full breath of observations**

Working with Dynamic Earthquake Rupture Models: A Practical Guide

Marlon D. Ramos[*1,2], Prithvi Thakur[1], Yihe Huang[1], Ruth A. Harris[3], and Kenny J. Ryan[2]

Geology

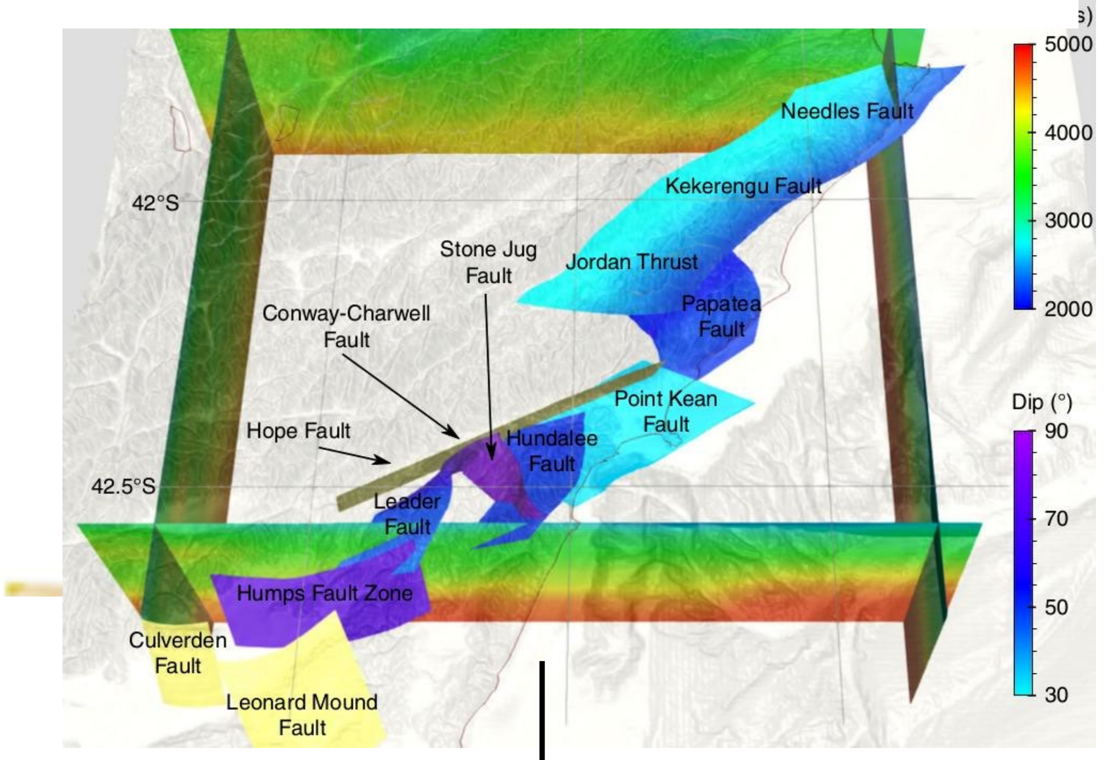

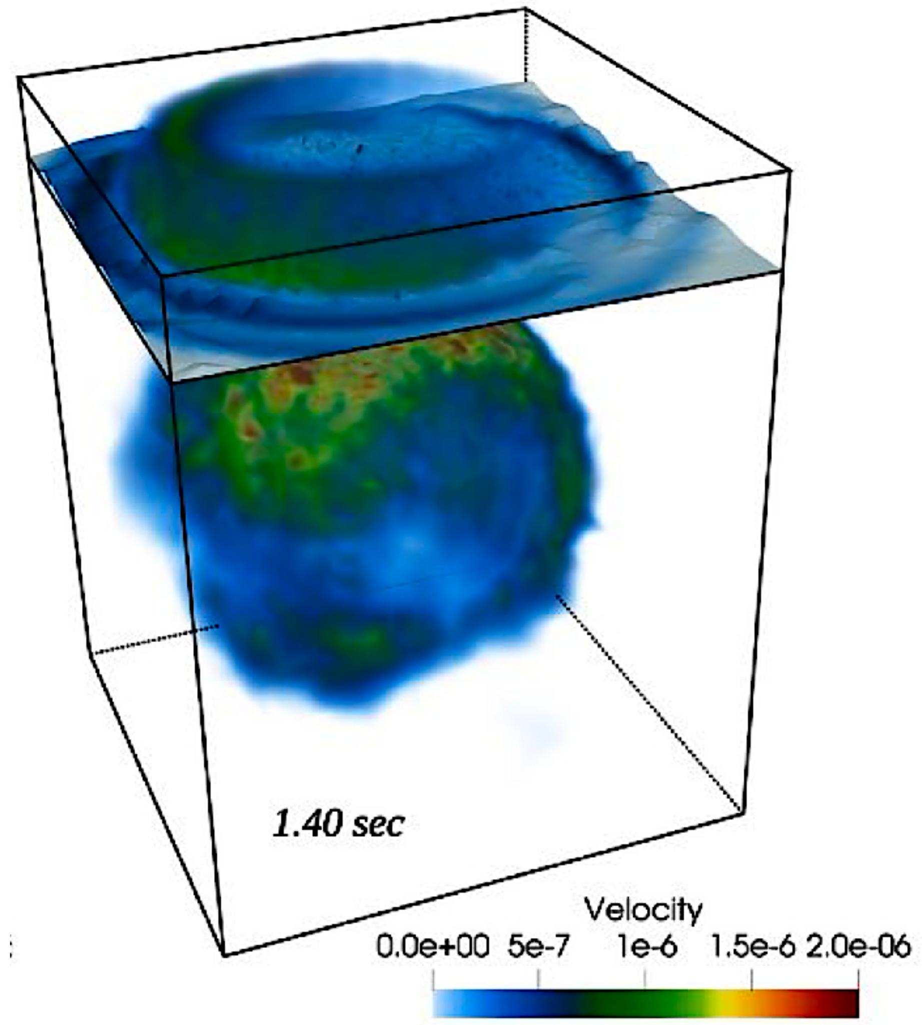

Interdisciplinarity

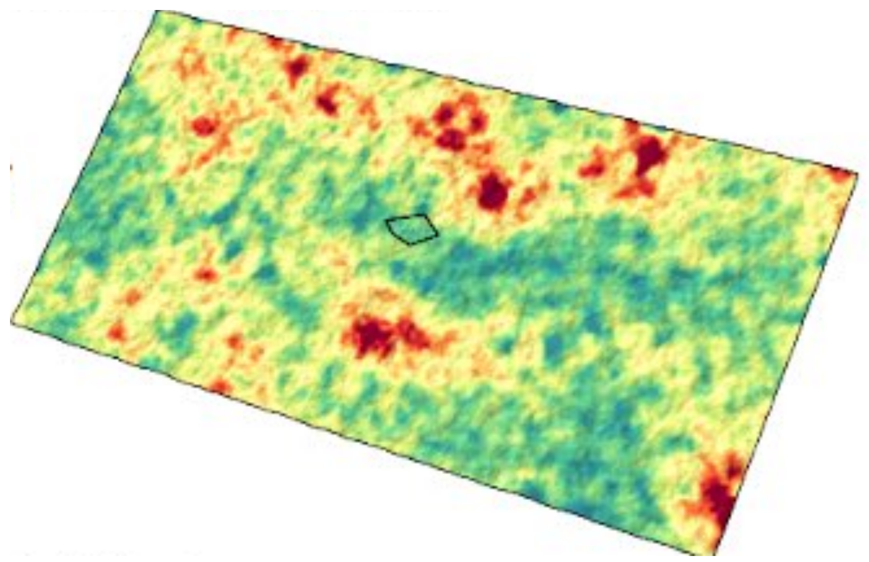

Initial fault loading

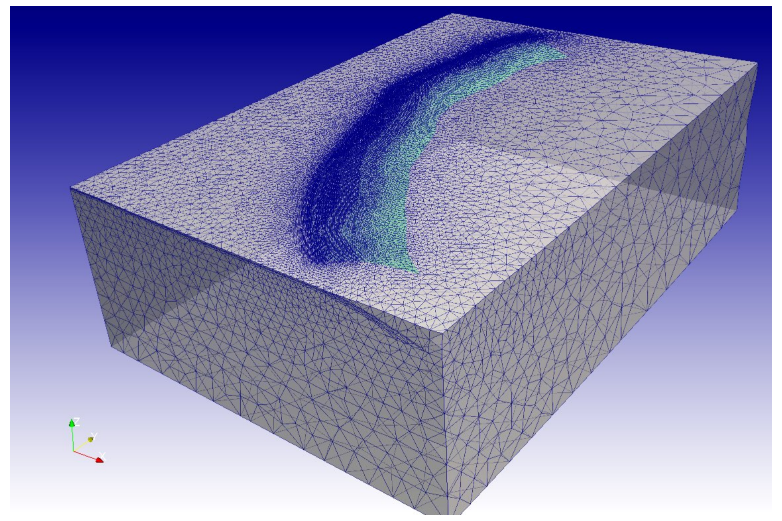

Mesh generation

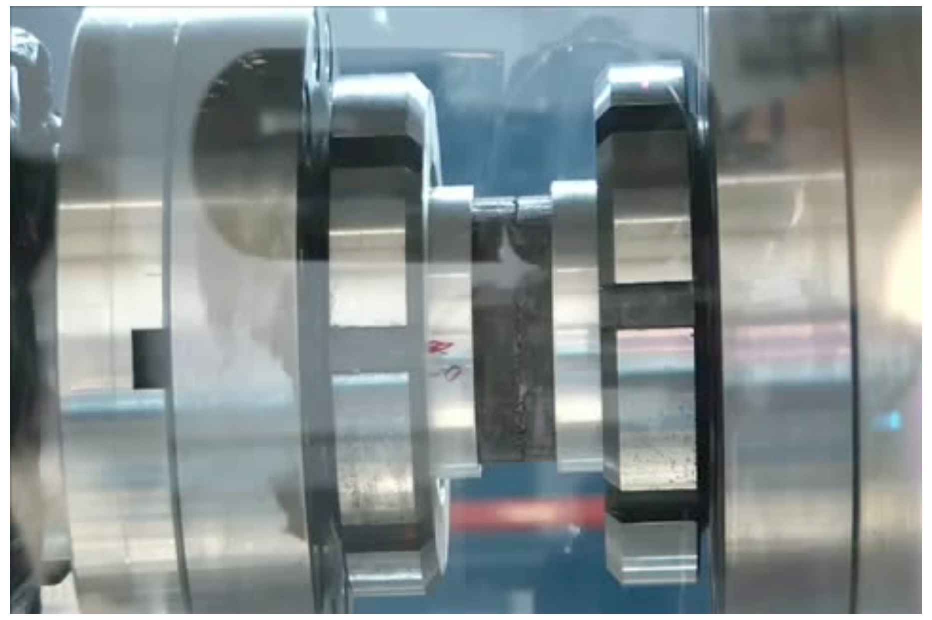

Friction experiments

"Input"

SOLVER

"Output"

Fundamental Physics

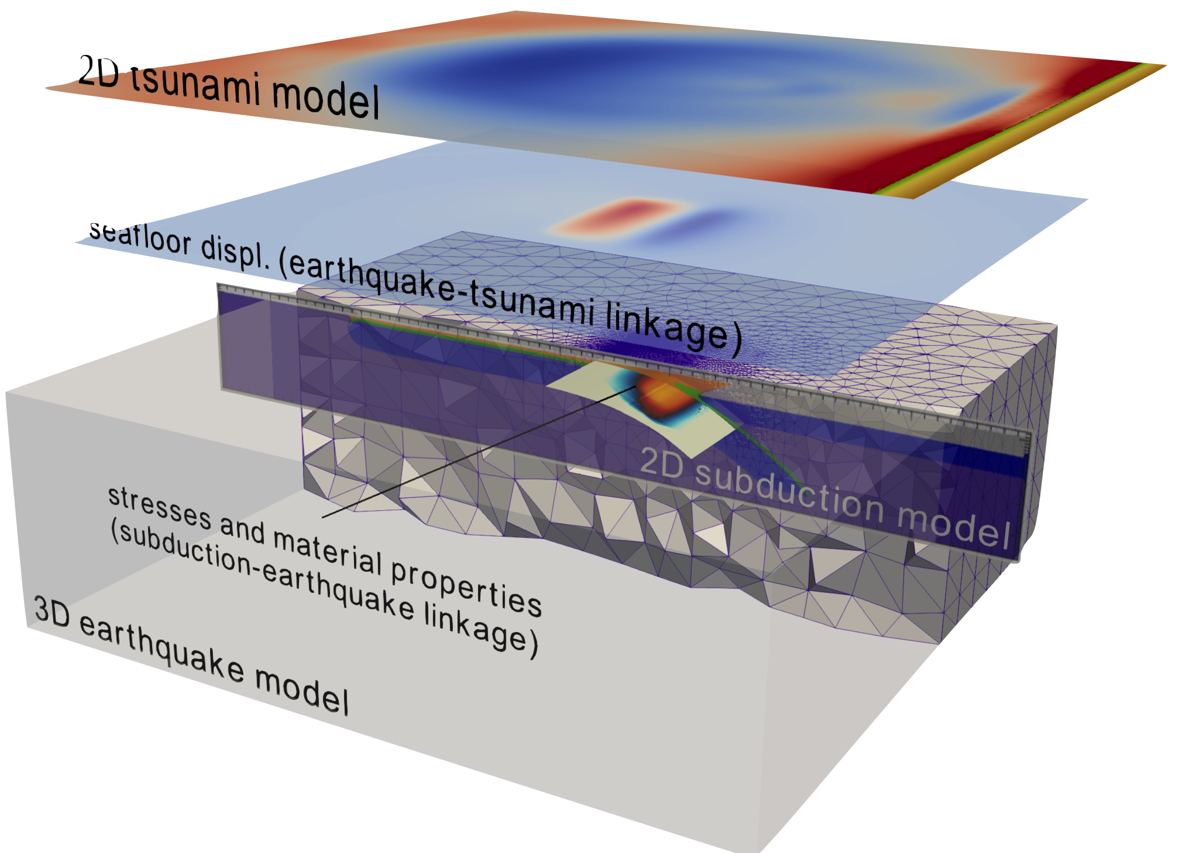

Ground deformation

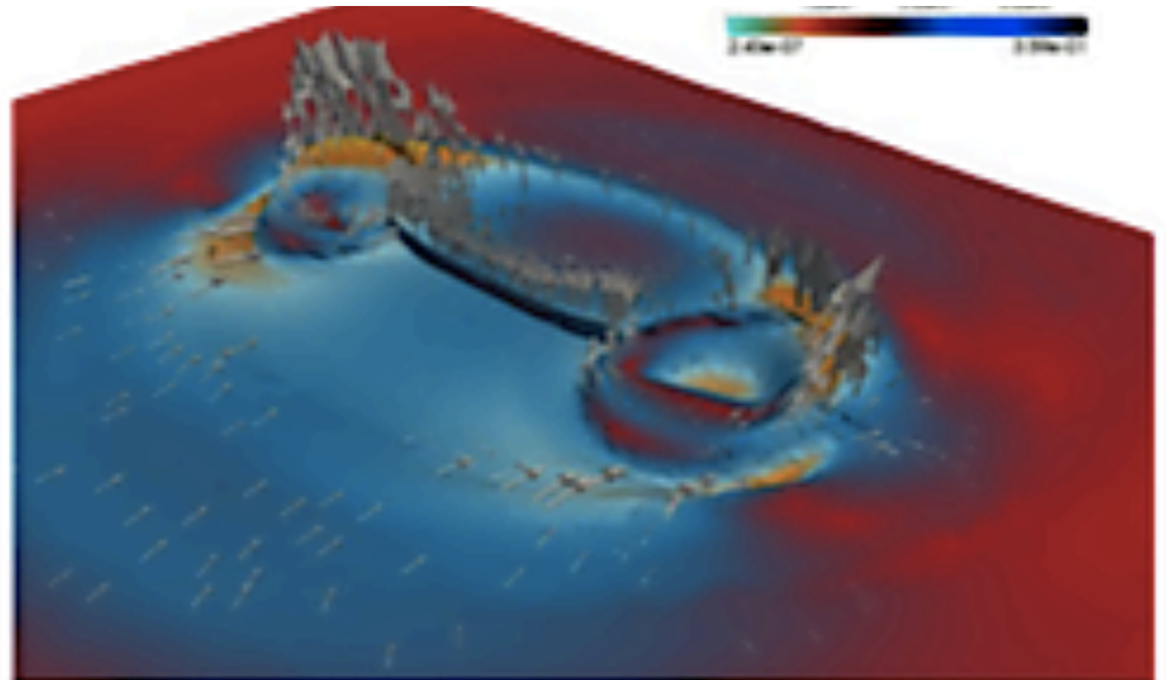

Synthetic observables

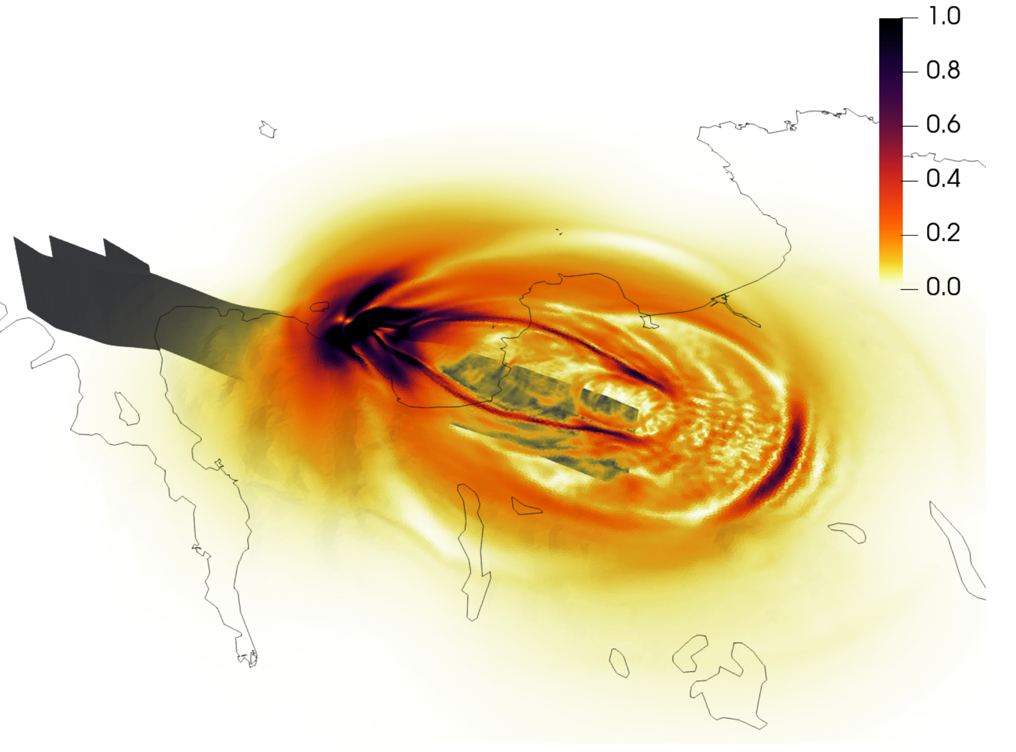

Rupture dynamics of the 2004 Sumatra mega thrust earthquake, Ulrich et al., 2022

Complex models, at scale of real-world observations

Challenge 1: In-situ Earthquake source processes are often ill-constrained and highly non-linear. How can models help to better understand and to better observe earthquakes?

Challenge 2: How to constrain and verify physics-based models? Which physical processes and scales are dominant and relevant for understanding the dynamics of real earthquakes?

Challenge 3: How to assimilate all available (community) knowledge? And how to do so in a suitable manner for software (numerical discretisation, solvers, equations solved), hardware (heterogeneous HPC systems, energy concerns) and the community (accessibility, reproducibility, lowering barriers)?

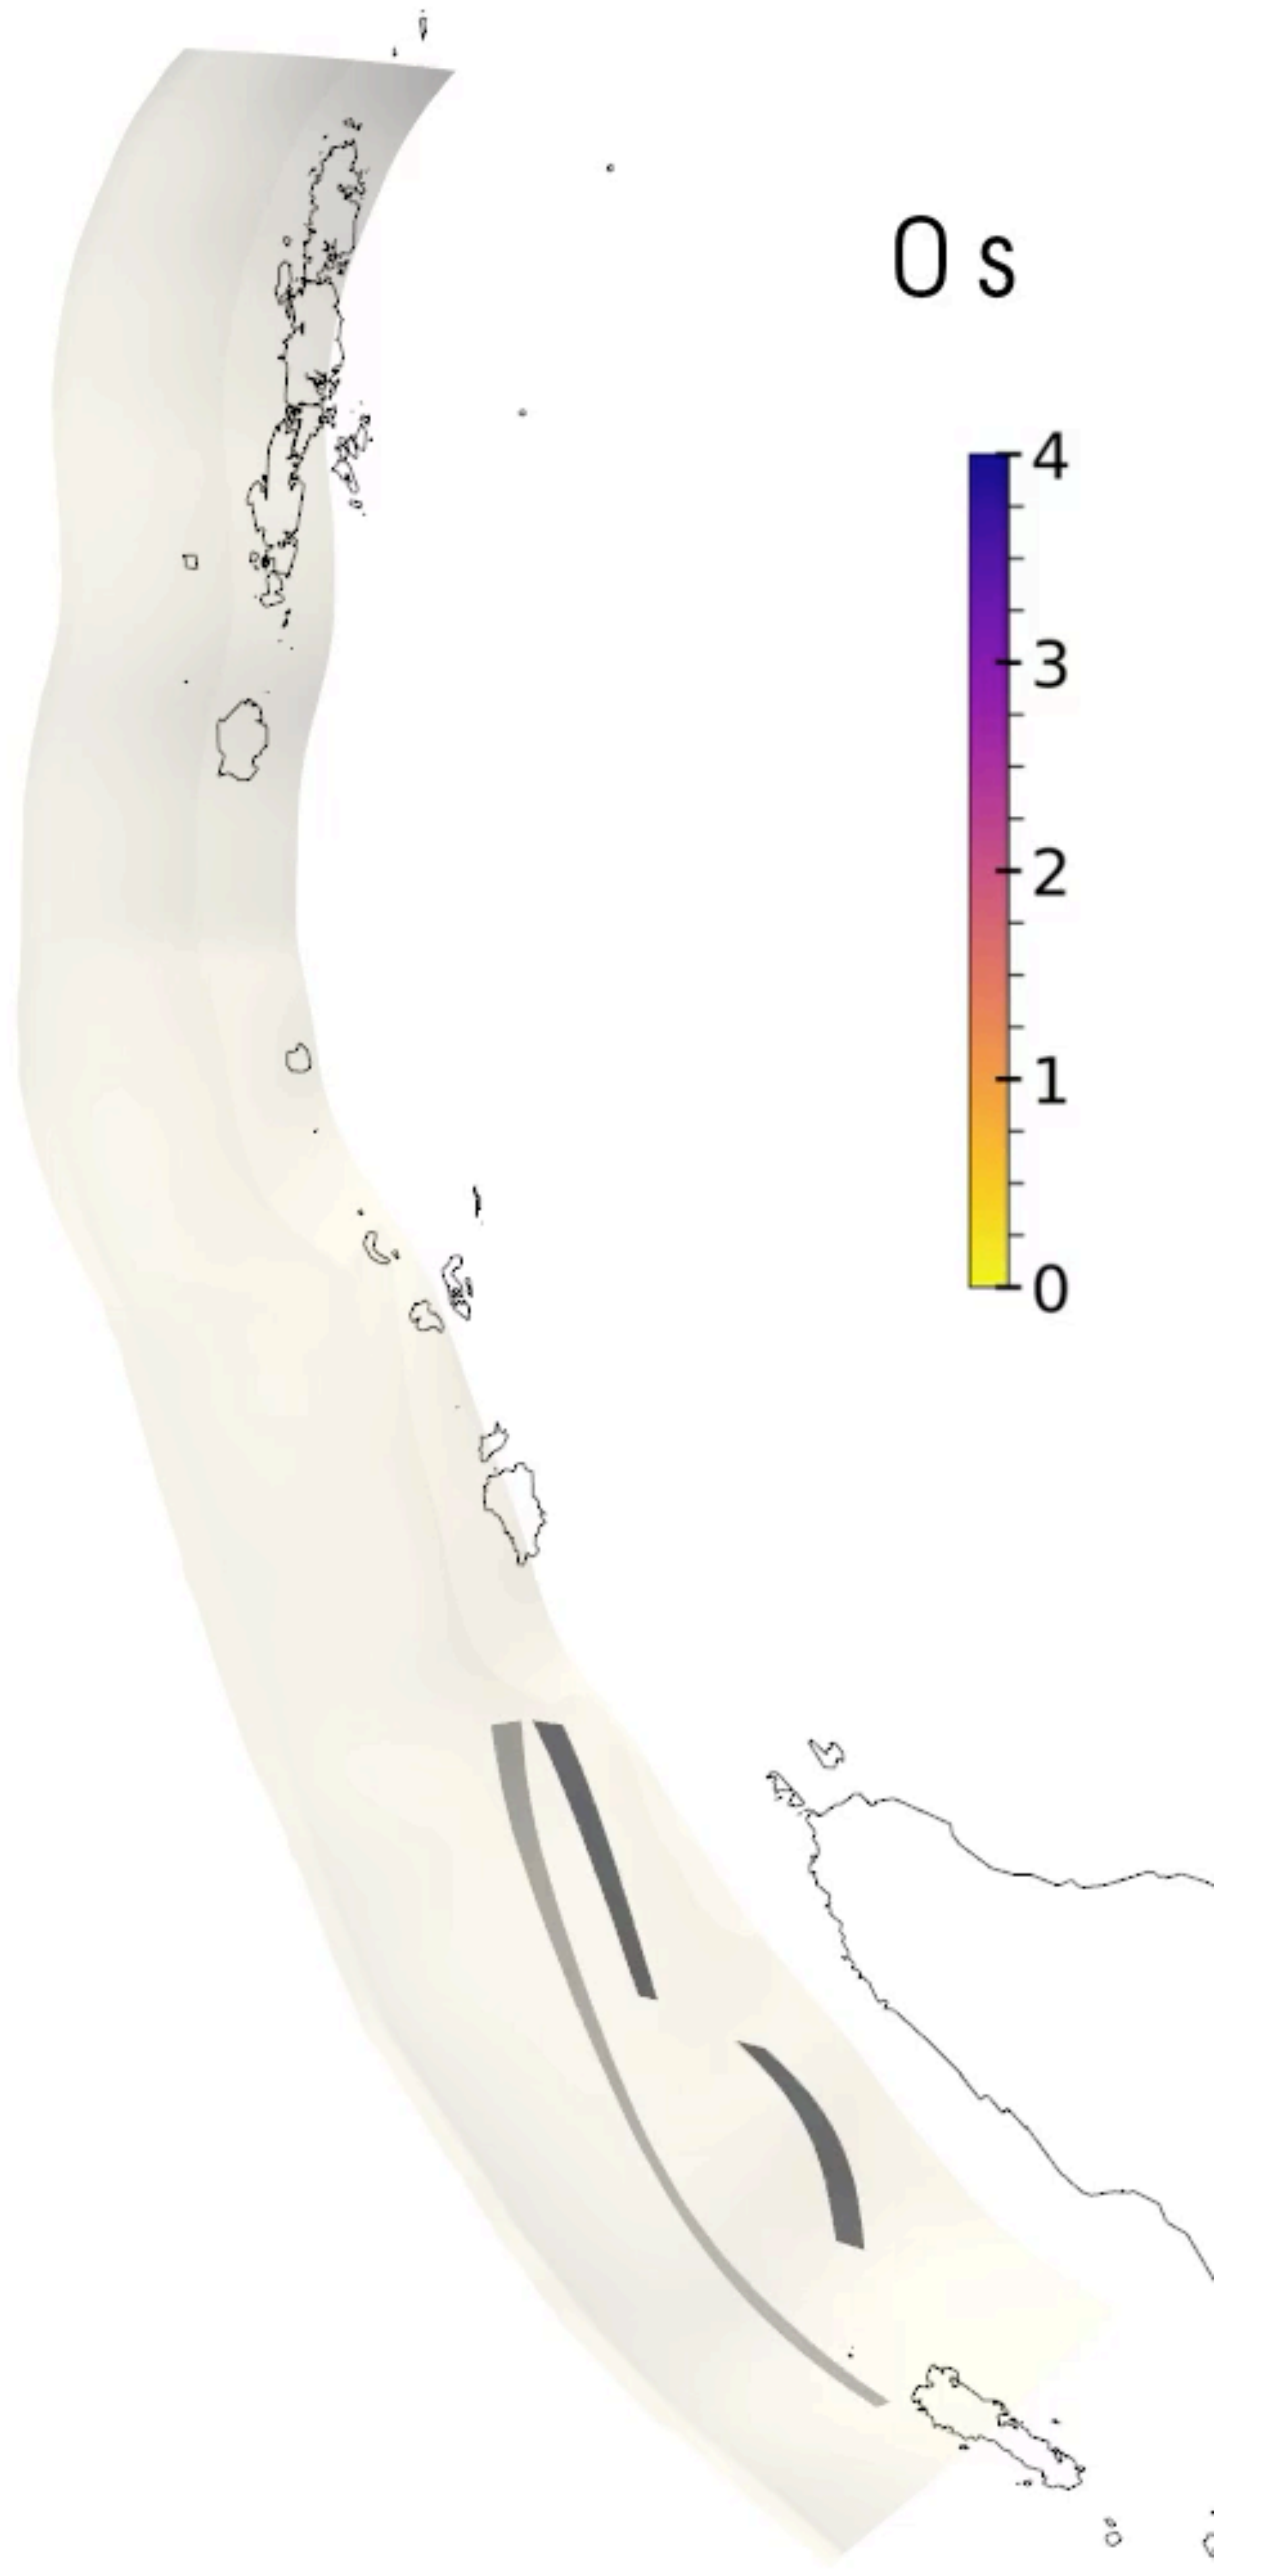

Reaching megathrust scales

- **Large space-time scales, cascading multi-physics and geometric complexity.** Physics-based models of frictional failure & seismic waves & dynamic seafloor displacements & tsunami modelling require **very large and long dynamic rupture simulations.**

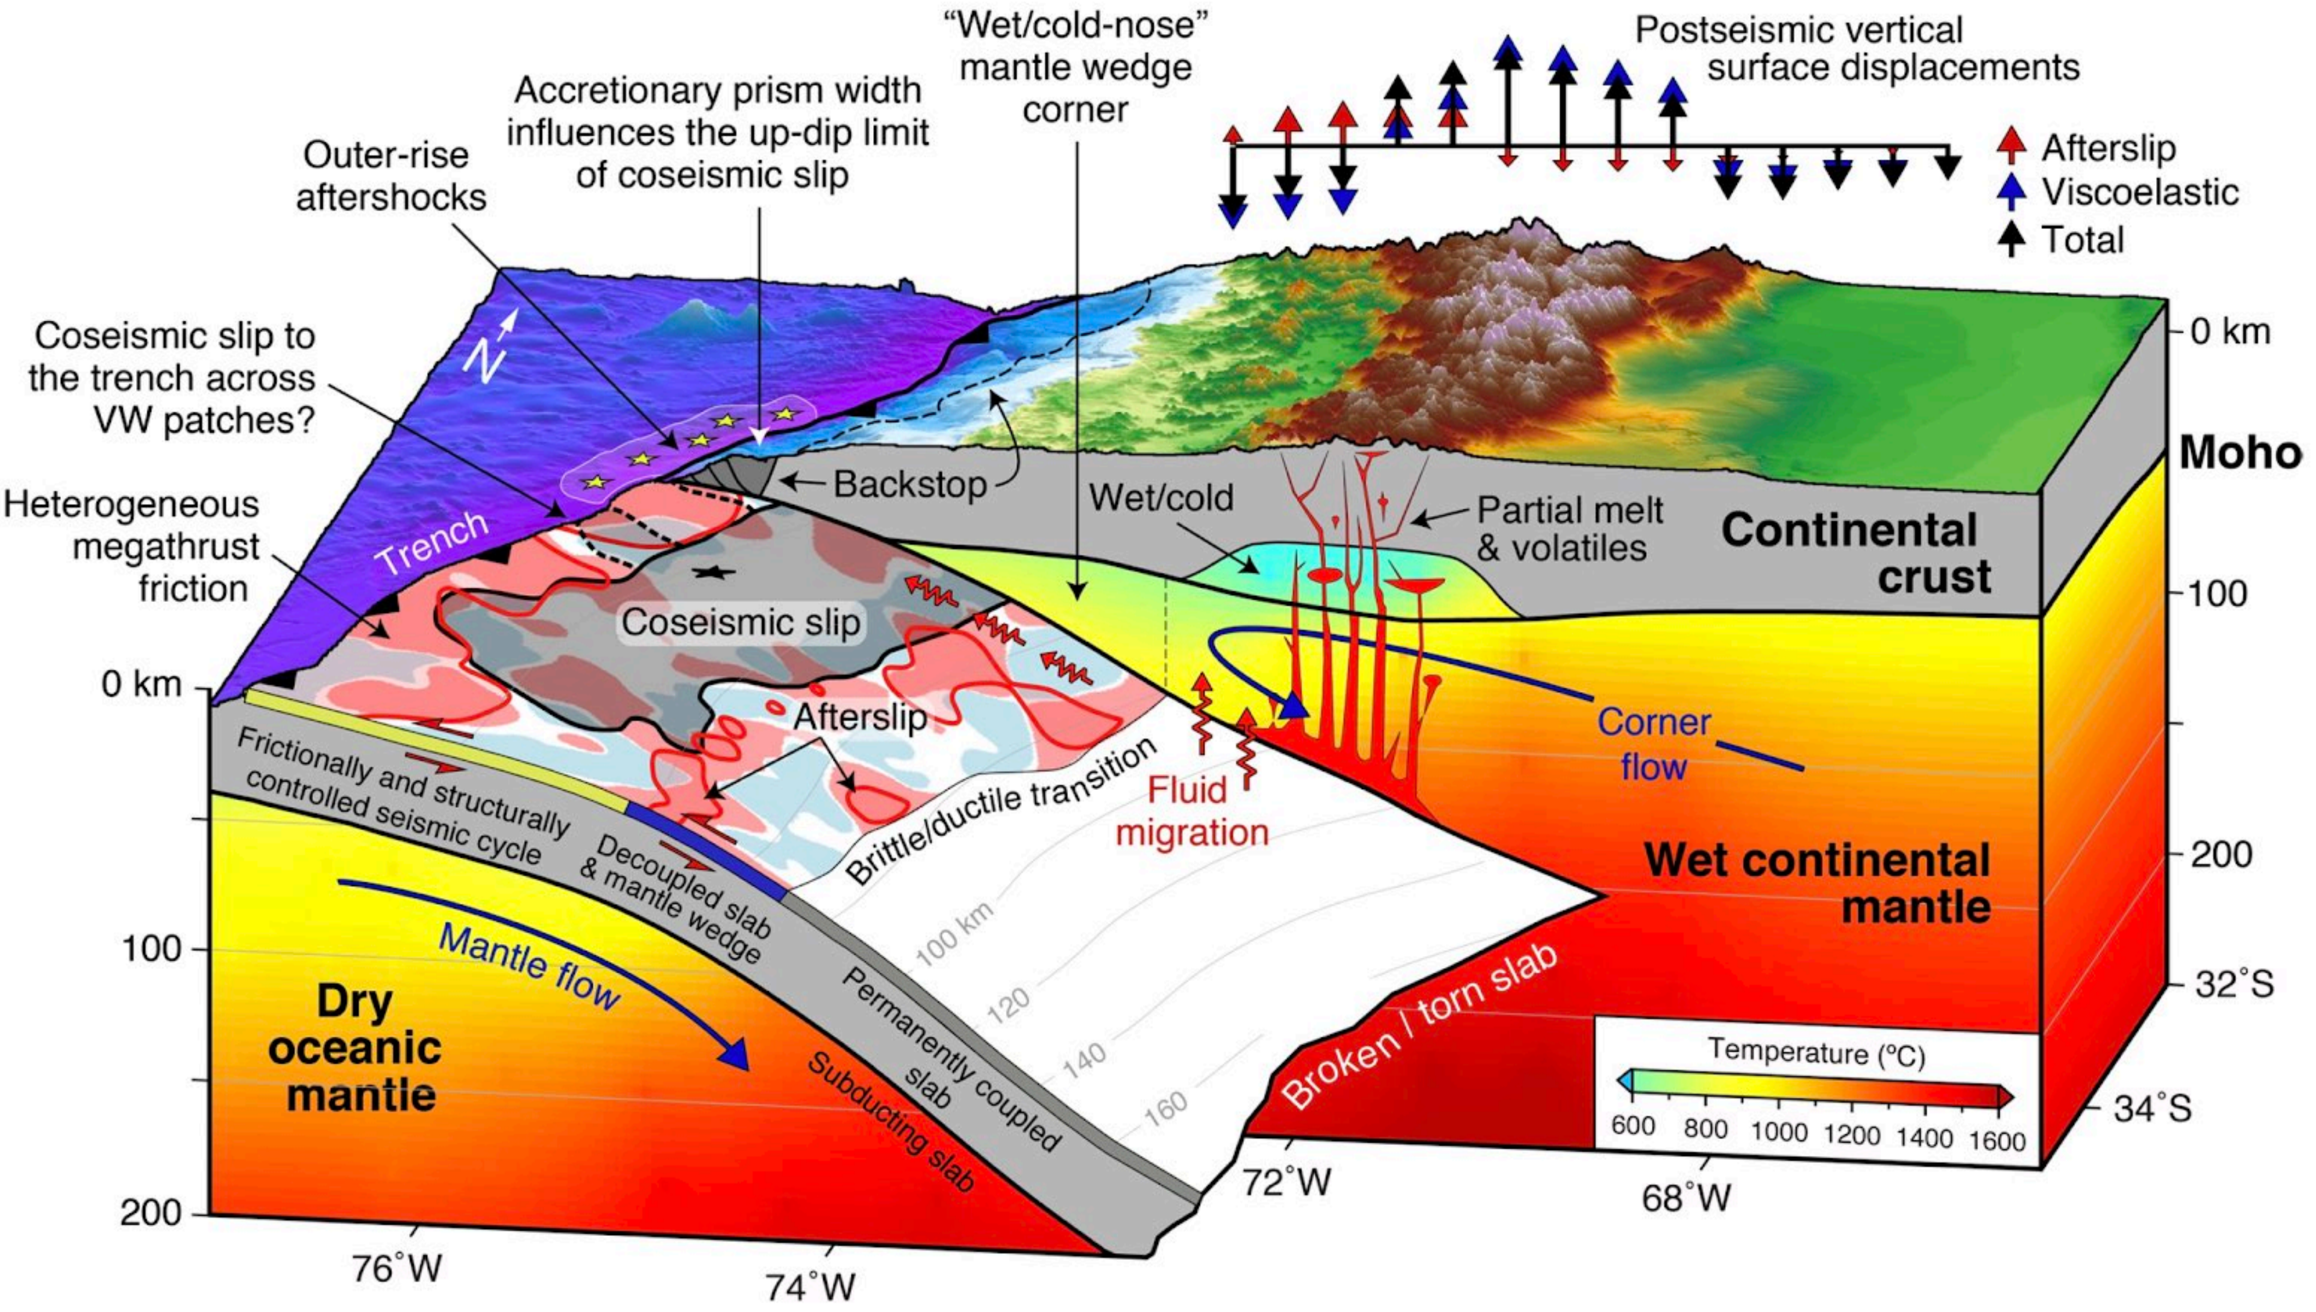

Weiss et al., 2019 Sci. Adv.

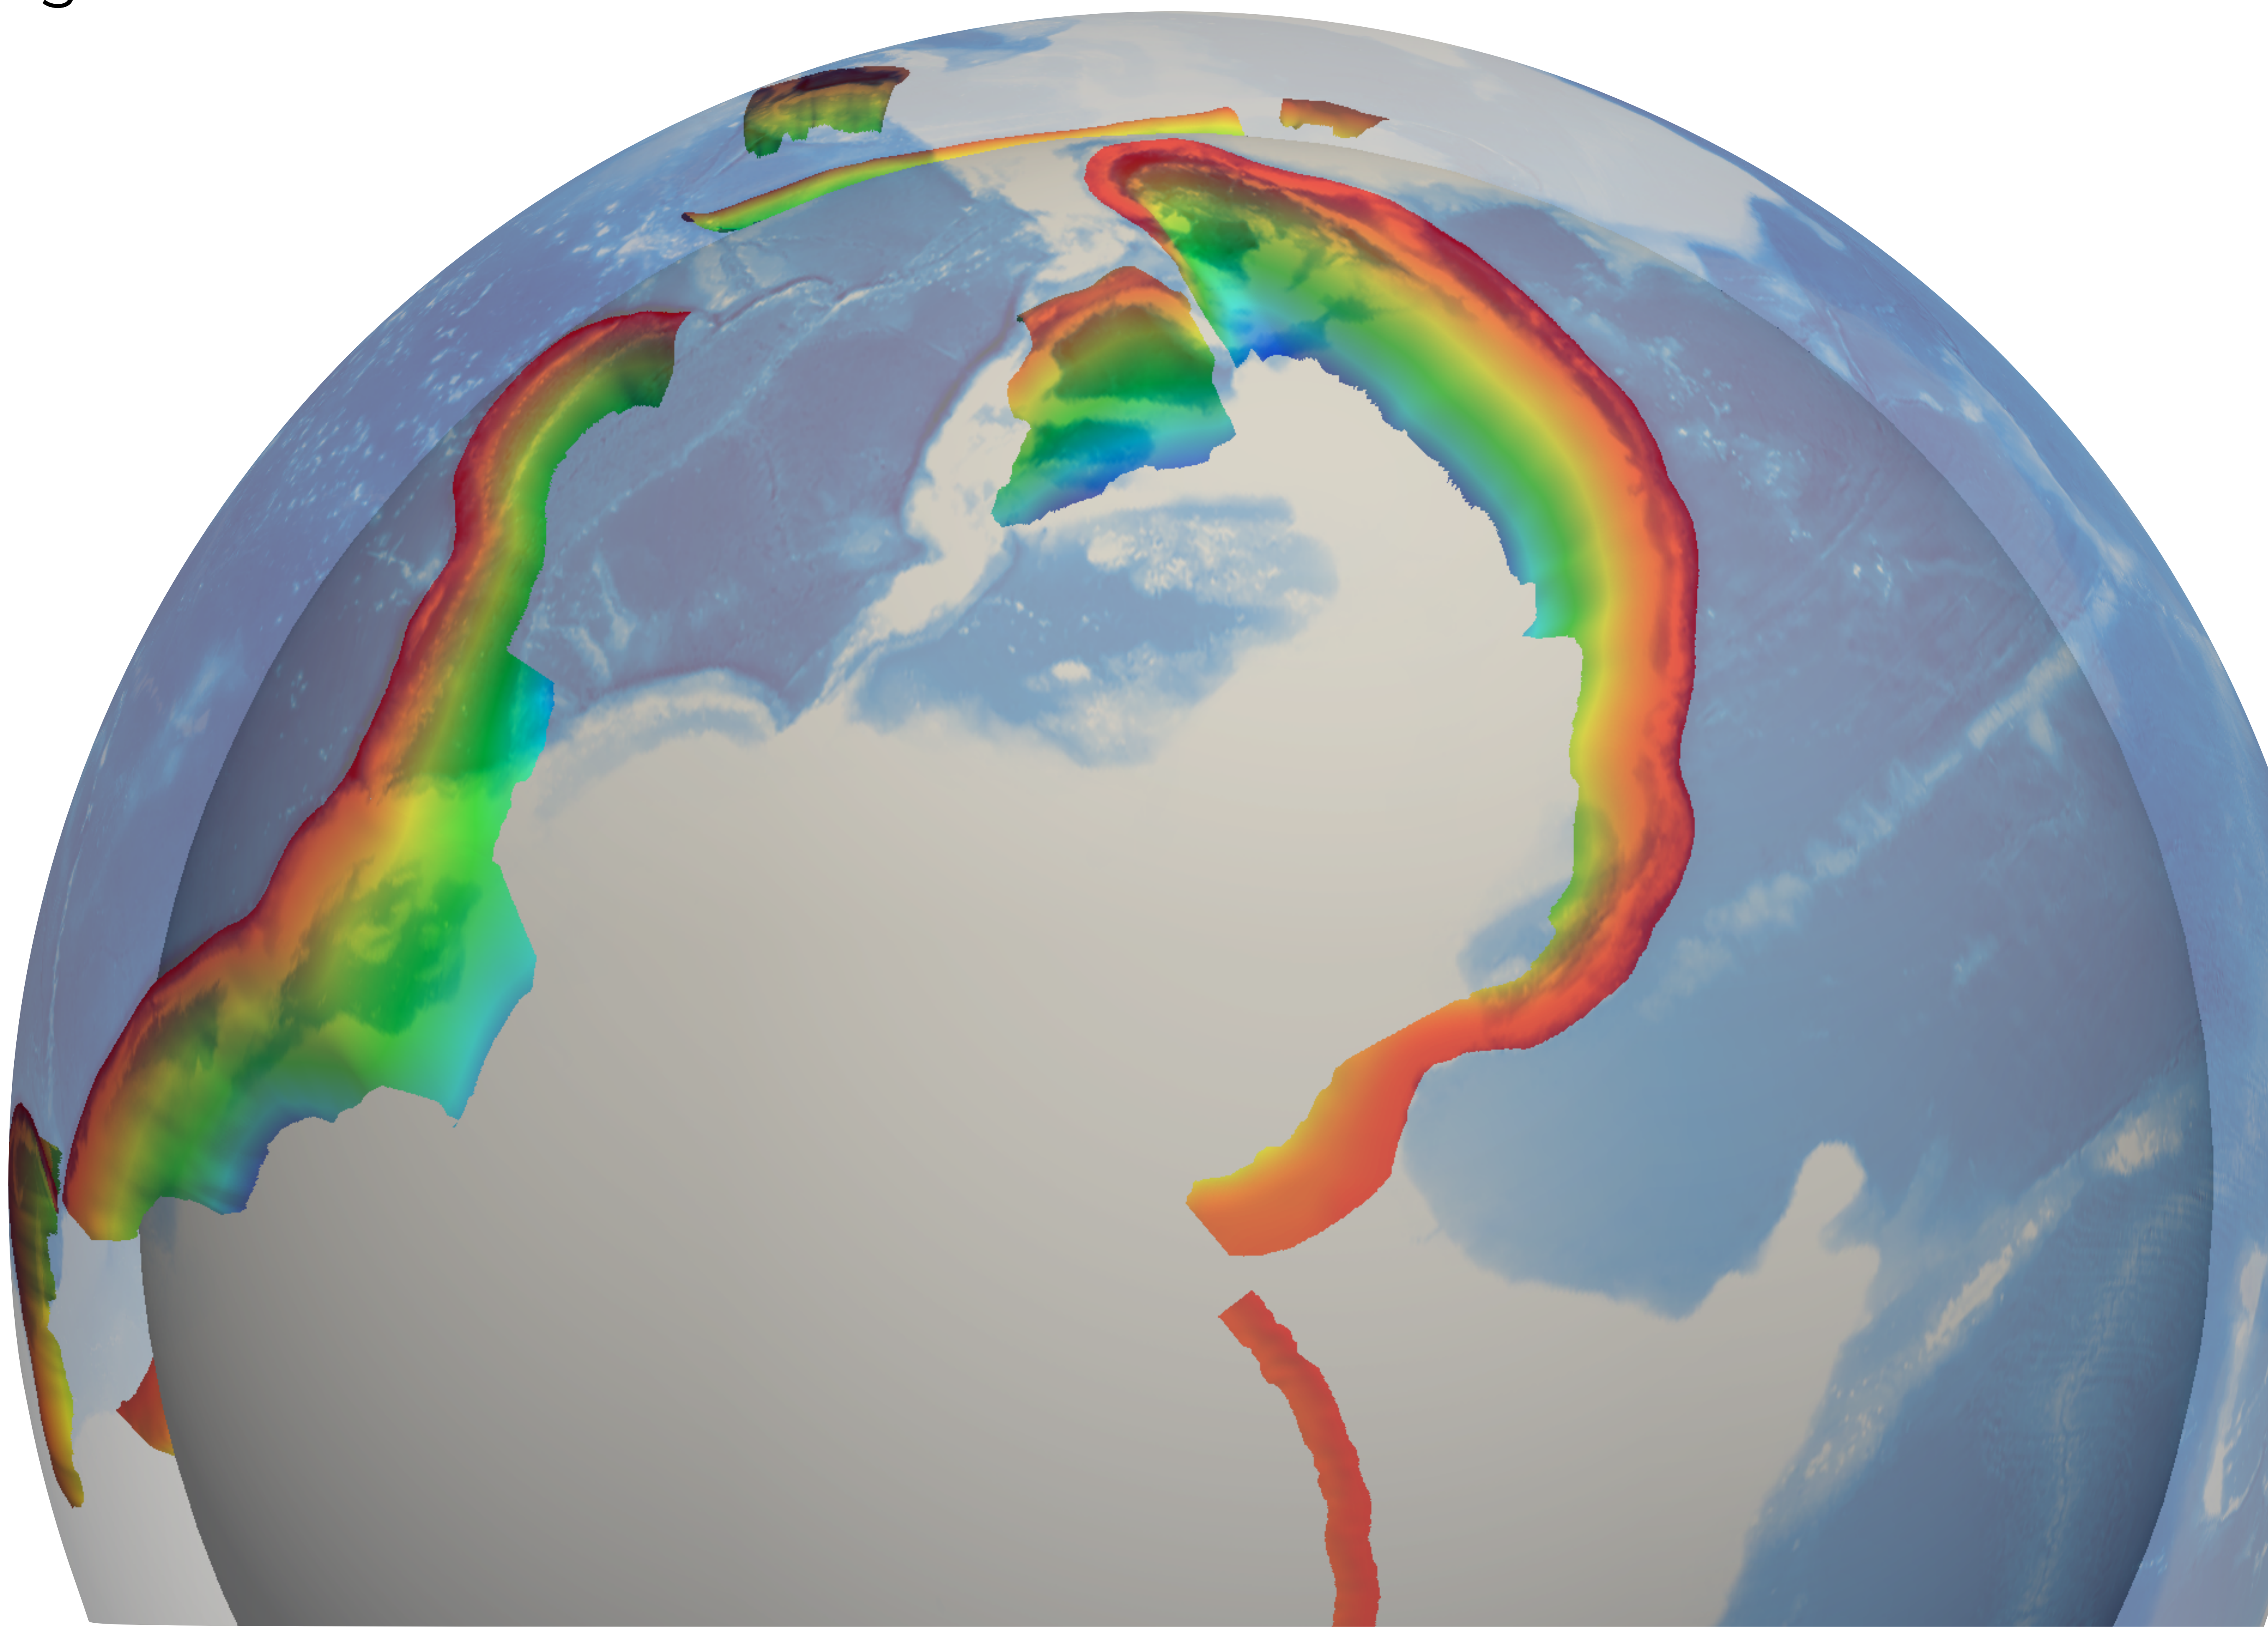

Visualisation of Slab 2.0 (Hayes et al.,2018)

Reaching megathrust scales

2004 Sumatra megathrust and splay dynamic rupture model geometry (Uphoff et al., 2017; Ulrich et al., 2022; Madden et al., 2022)

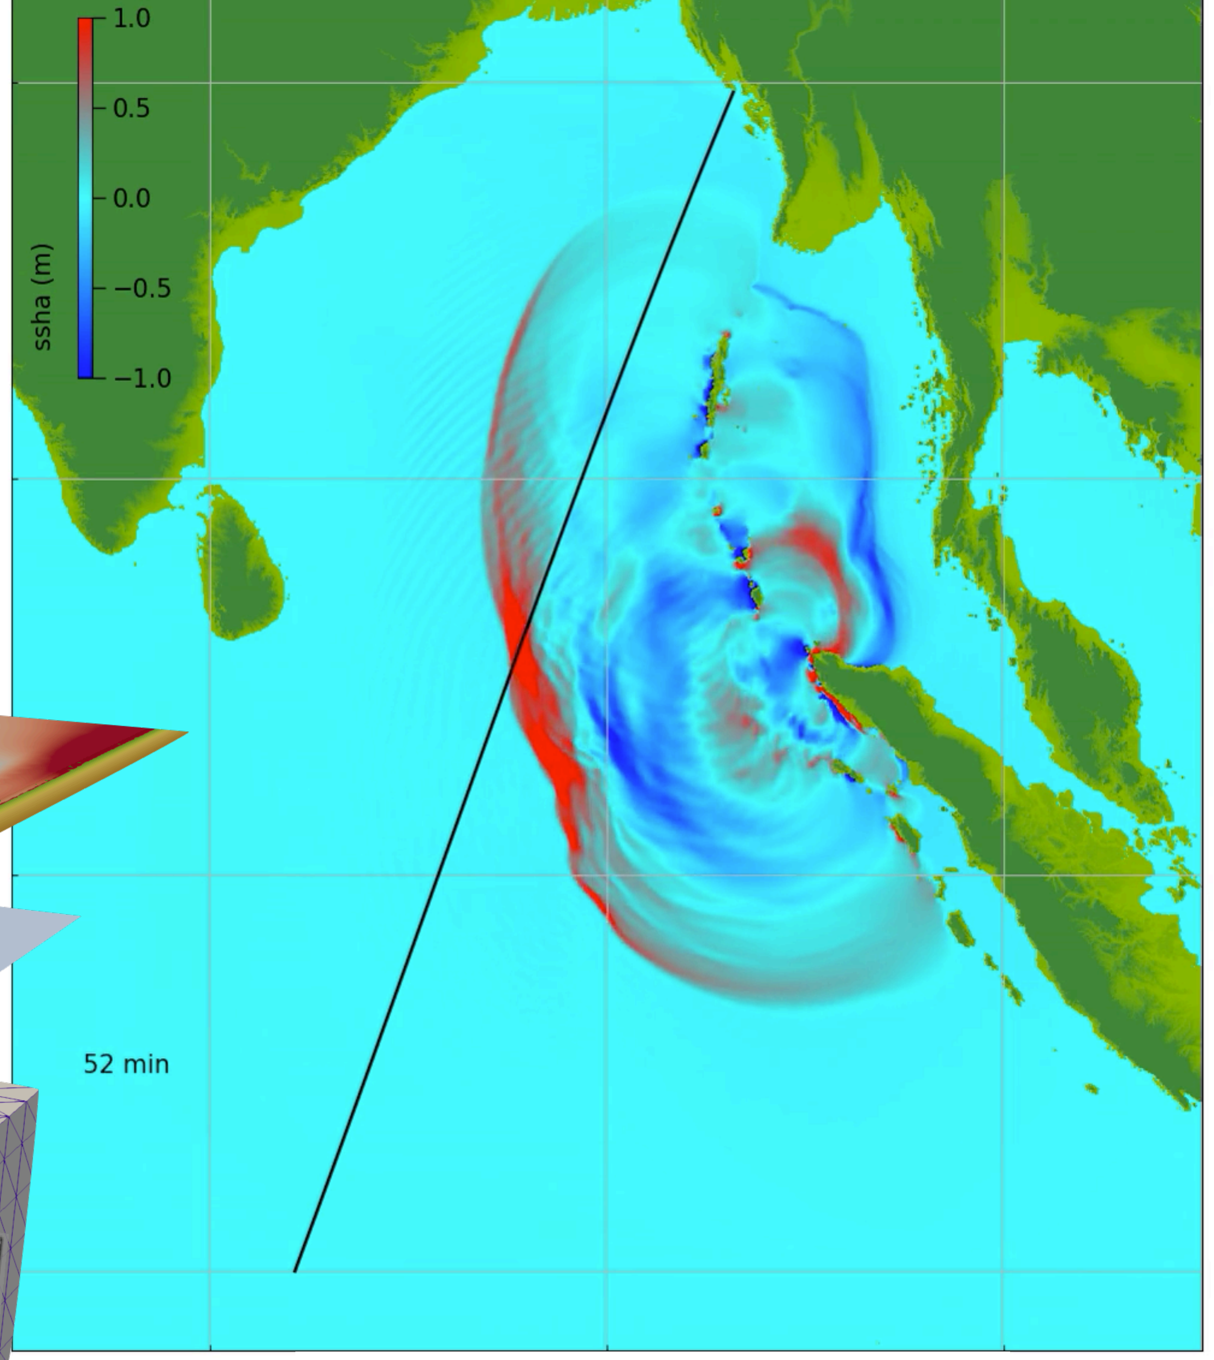

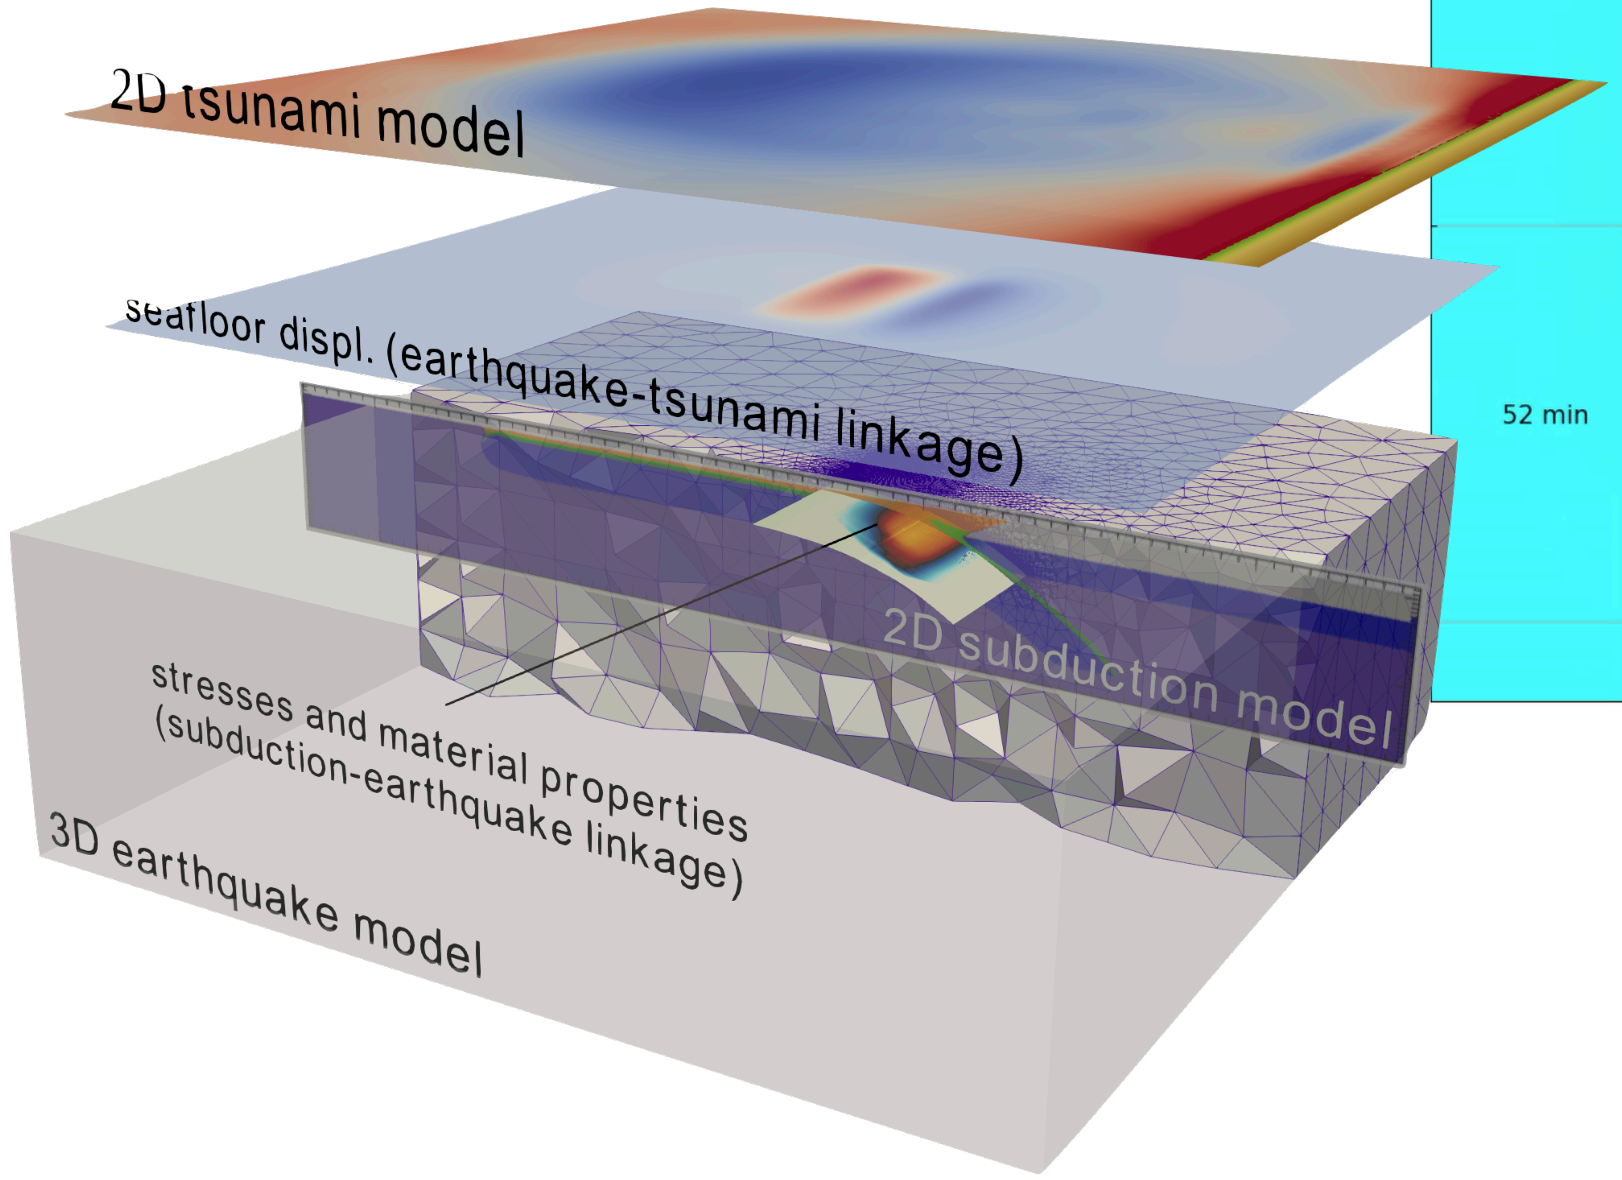

- **reconciling near- and far-field earthquake and tsunami observations of the 2004 Sumatra-Andaman events**
- **Using regional-scale observations of stress, rigidity and sediment strength to decipher megathrust hazards in a physics-based manner**
- **Requires numerical methods handling geometric complexity and highly varying resolution requirements** (SeisSol)
- **End-to-end computational optimization** including auto-generated assembler-level DG kernels, hybrid MPI/OpenMP parallelization, a geoinformation server (ASAGI) for fast and asynchronous input/output and clustered local time stepping
- **"Hero runs", e.g.** with 220 million finite elements (~111 billion degrees of freedom) and 3.3M time steps took 14h (in 2017)
- Todays typical subduction earthquake simulations (20 million elements) now require **5h30 minutes on 16 nodes** (4k CPUh)

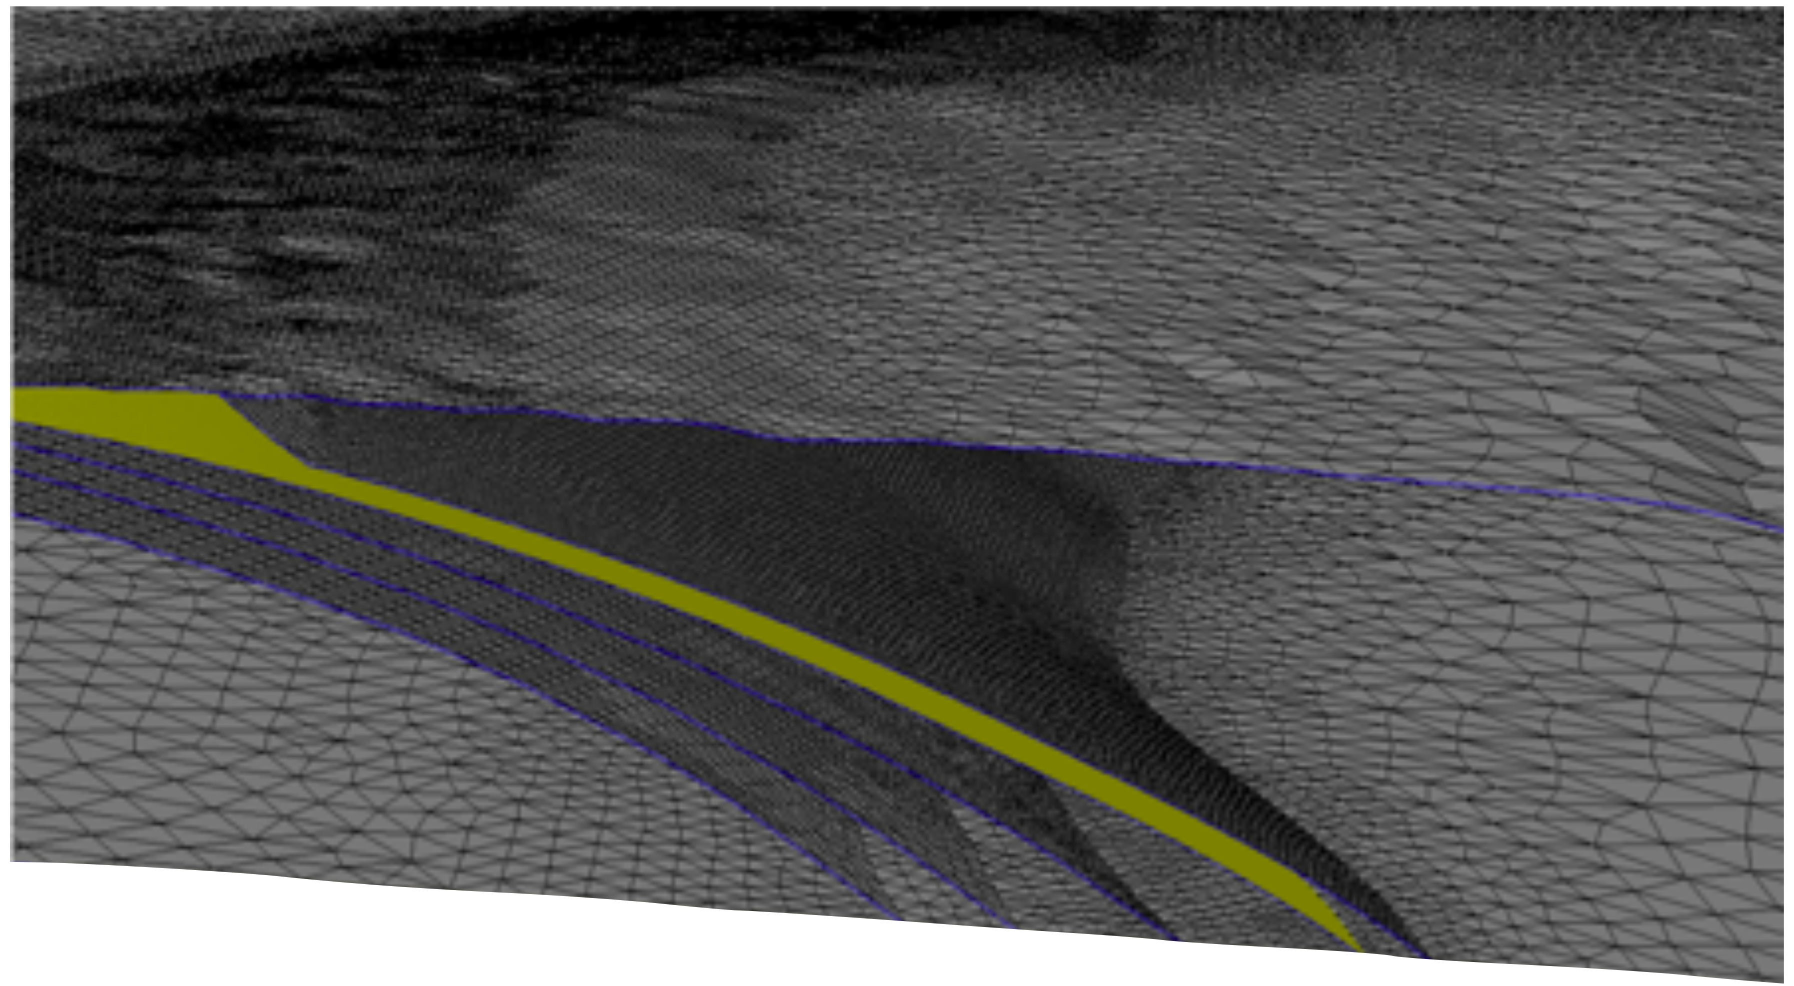

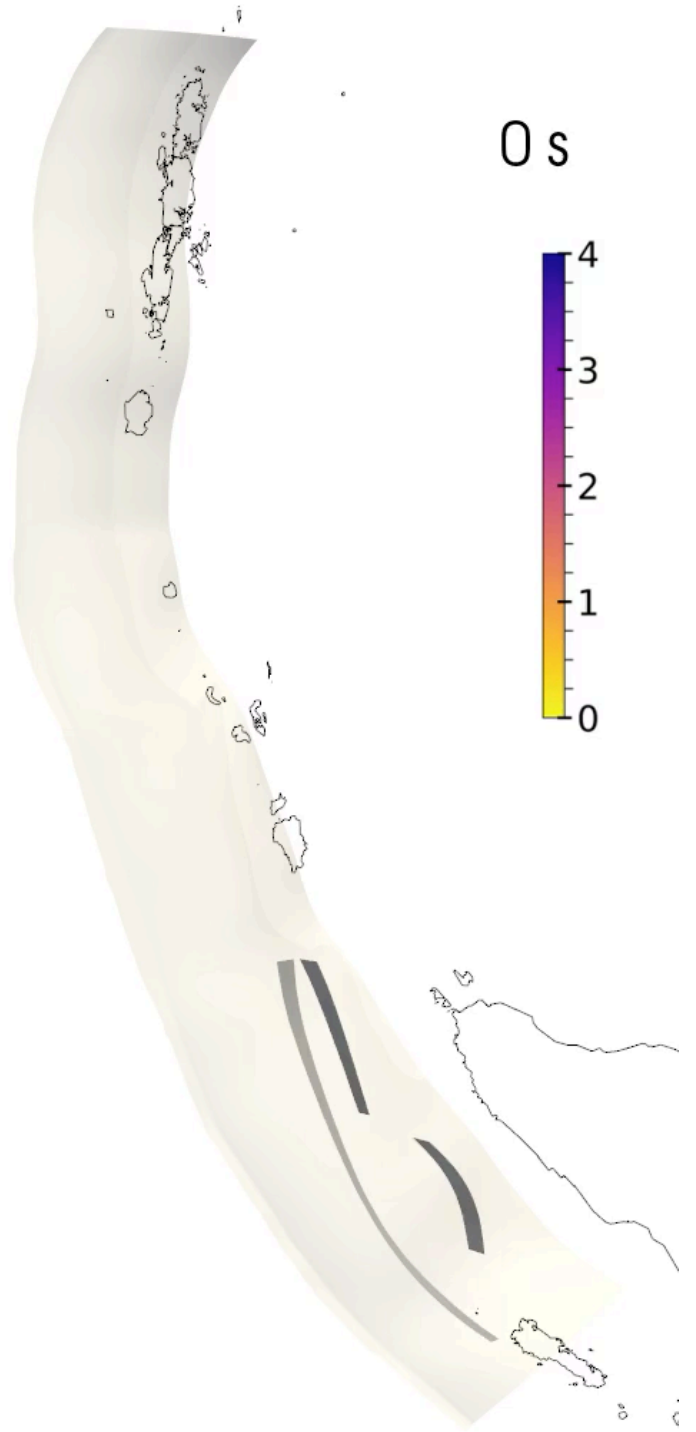

A conversation of earthquakes - The 2019 Ridgecrest, CA, sequence

- Interdisciplinary observations are **interesting** (**difficult**!), hence often studied in isolation, sometimes leading to **non-unique or opposing results** similar to data-driven modeling

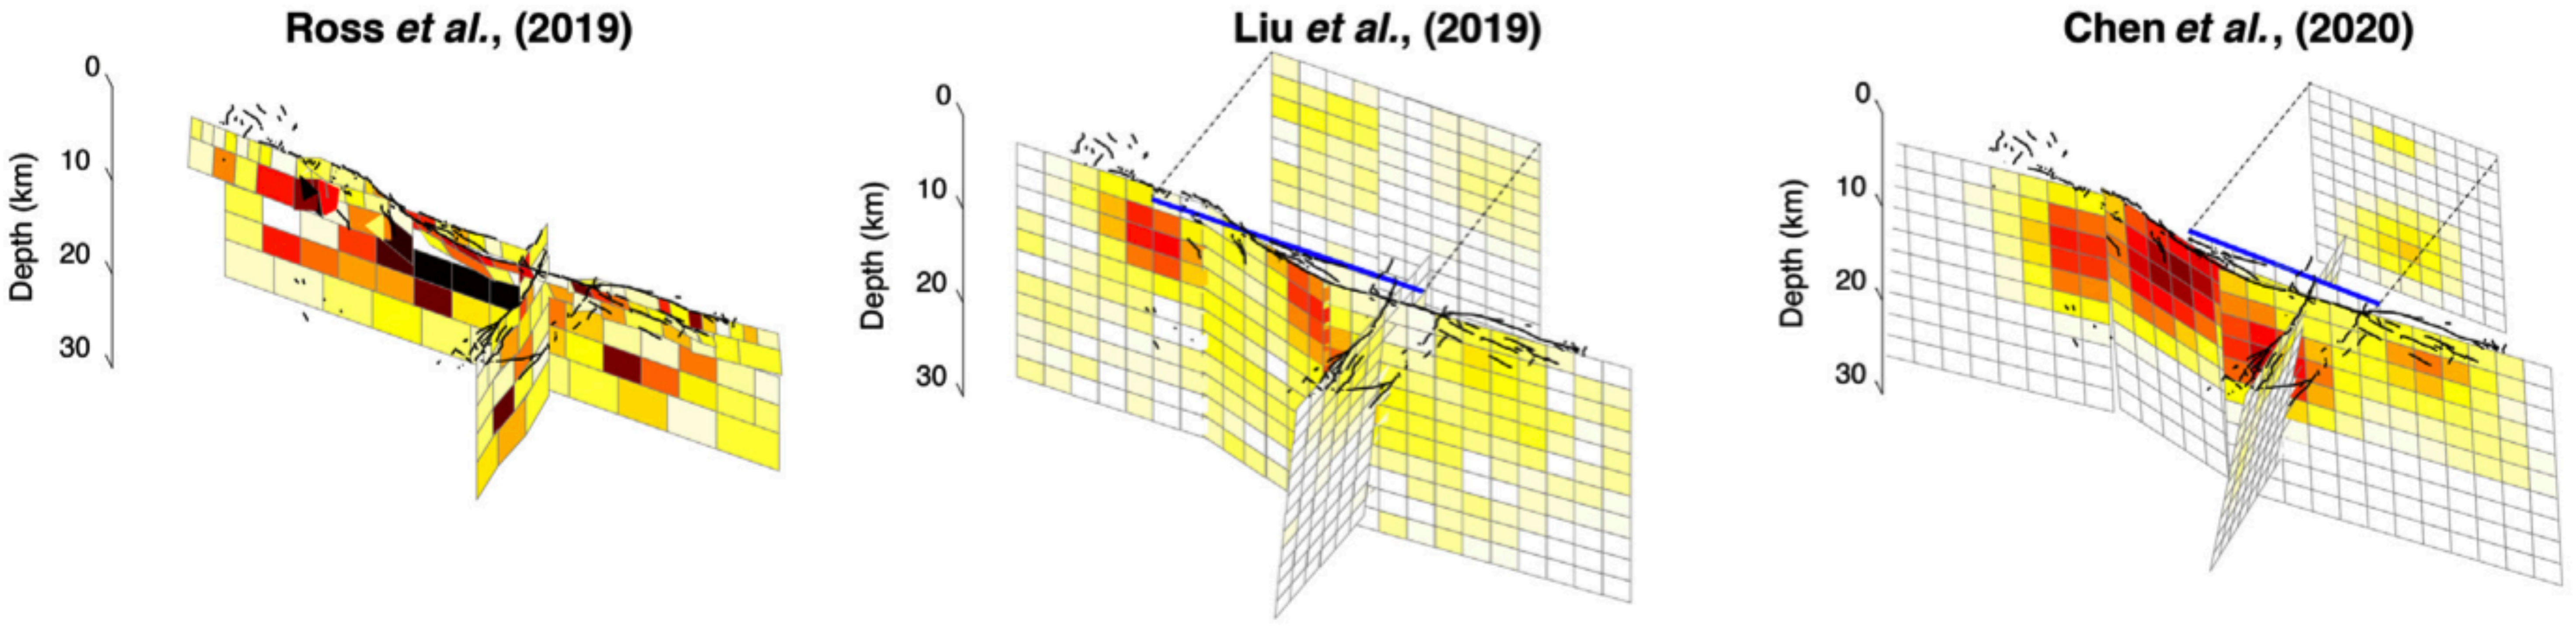

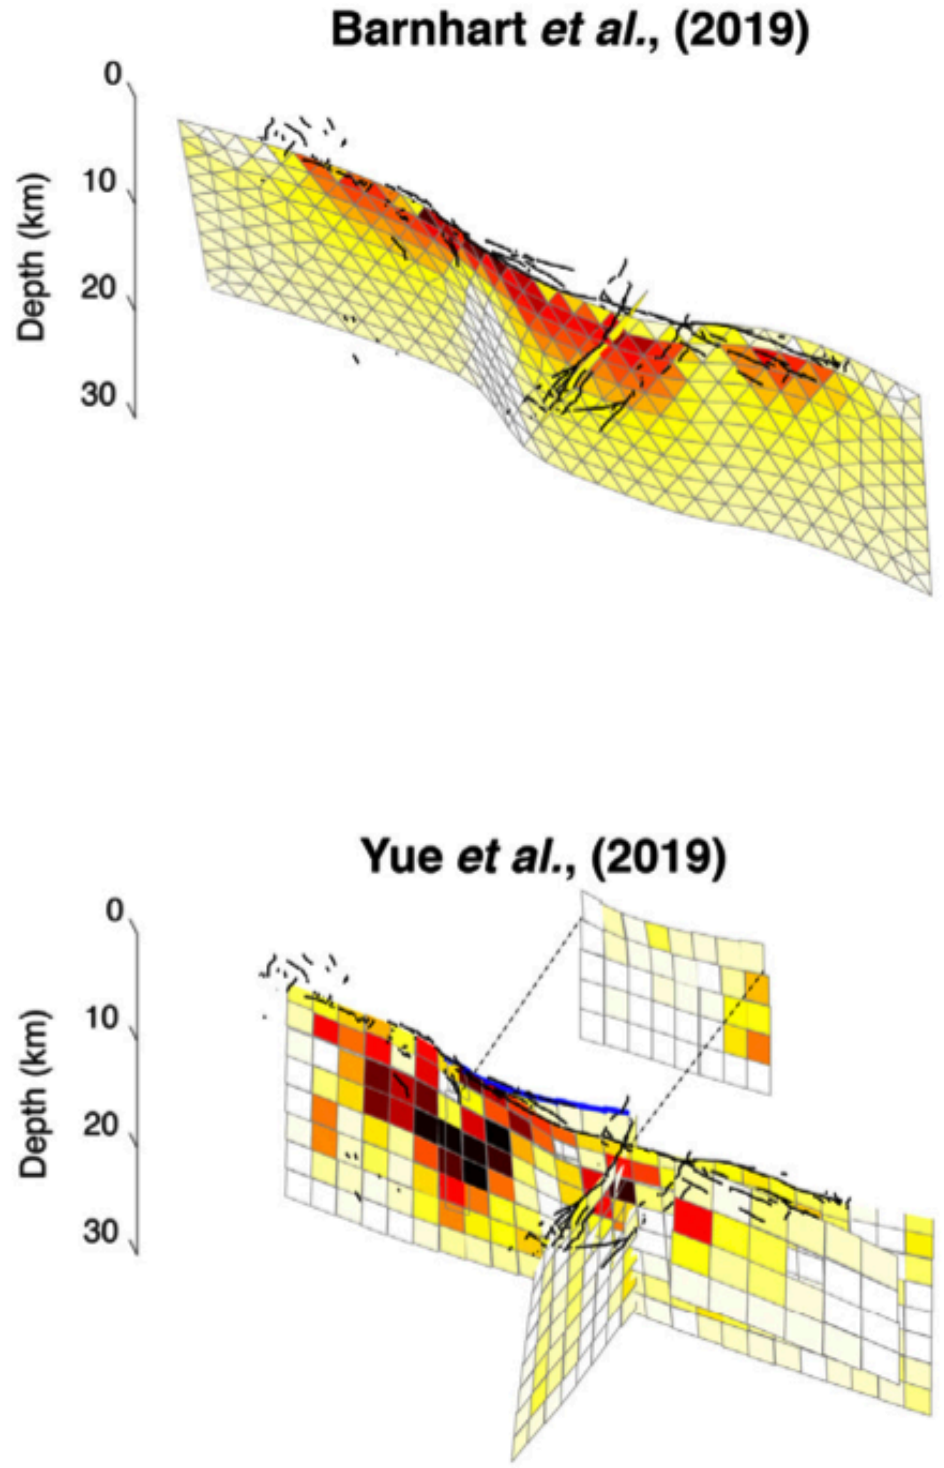

Surface rupture from the M7.1 Ridgecrest earthquake in 2019. Photo: Ben Brooks/USGS

May 24, 2023

“SEGMENT-JUMPING” RIDGECREST EARTHQUAKES EXPLORED IN NEW STUDY

Seismologists use supercomputer to reveal complex dynamics of multi-fault earthquake systems

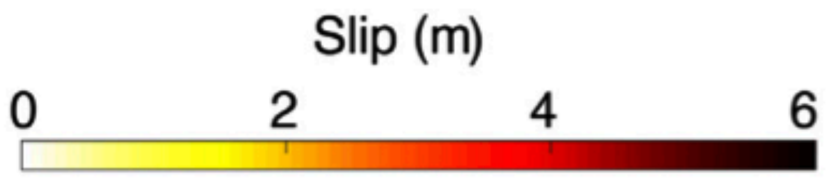

Wang et al., 2020, Comparison of co-seismic models for the 2019 Ridgecrest earthquake sequence.

Integrating multi-scale observations to study dynamics & delays of the 2019 Ridgecrest sequence

Taufiqurrahman et al., Nature, 2023

- **Coupled dynamic rupture models** using supercomputing to find the link between California's biggest earthquakes in 20 years breaking multiple segments of the same fault system
- regional **structure**, ambient long- and short-term **stresses**, as well as **co-seismic fault system interaction**, are crucial to understand the dynamics and delays of the sequence

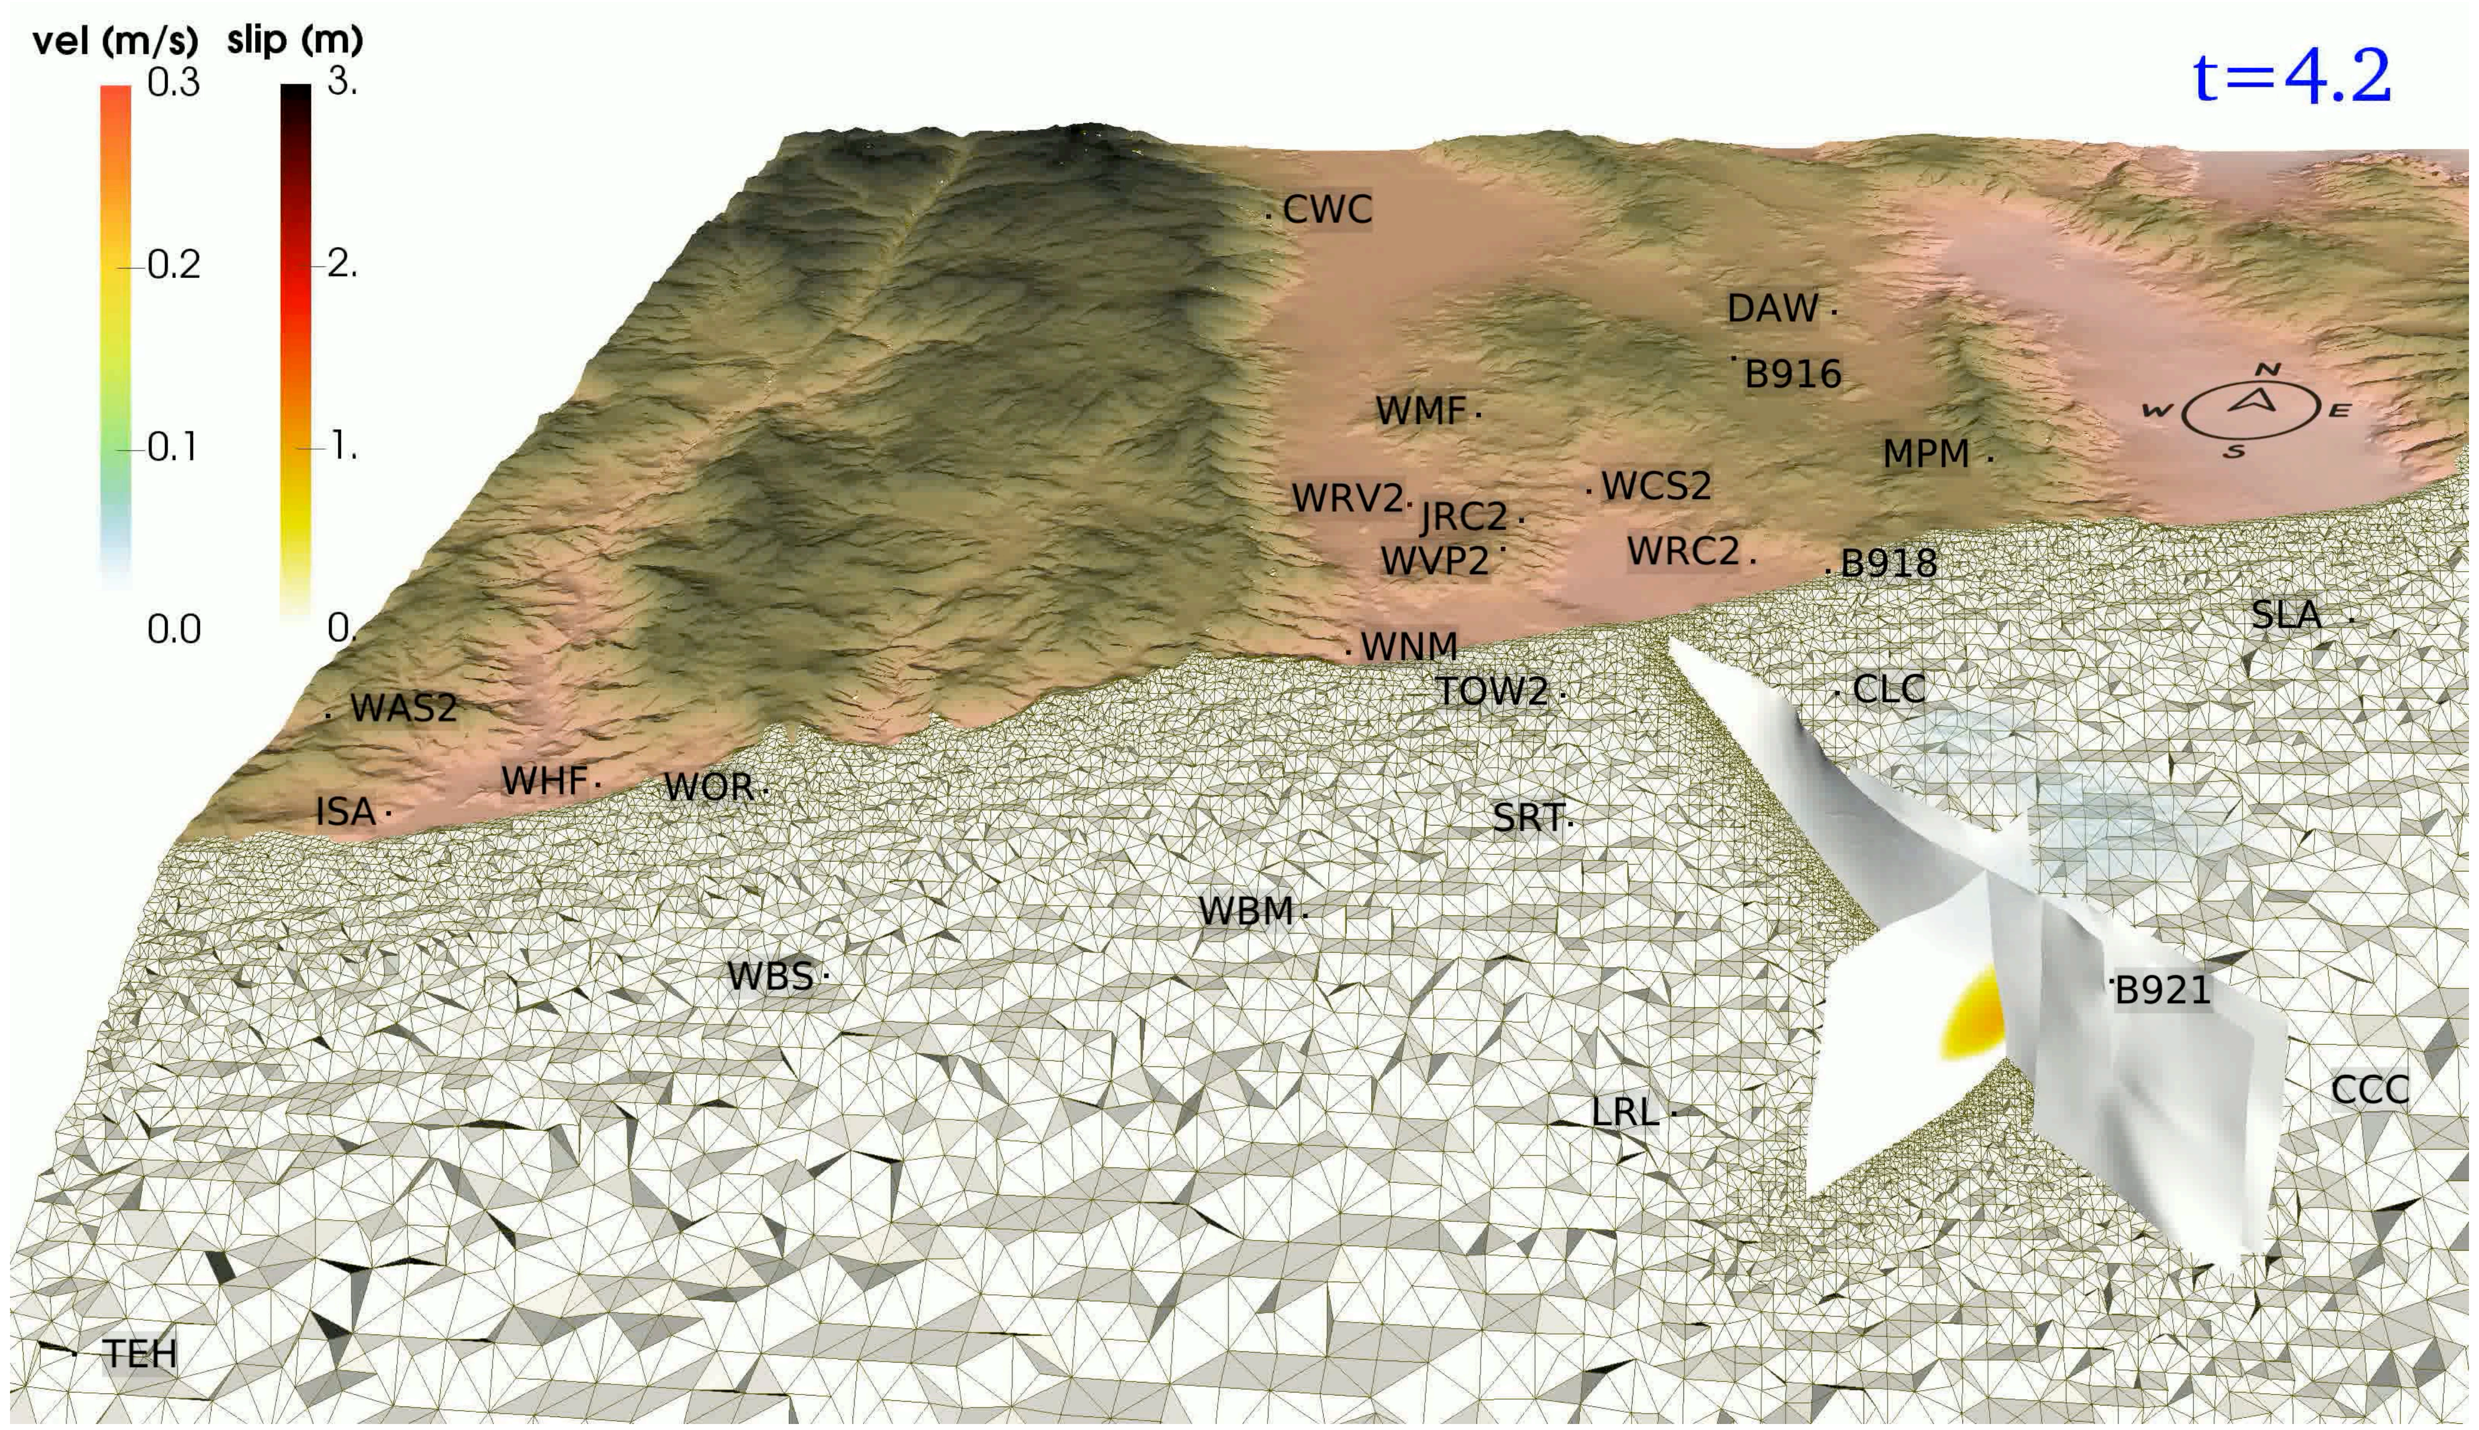

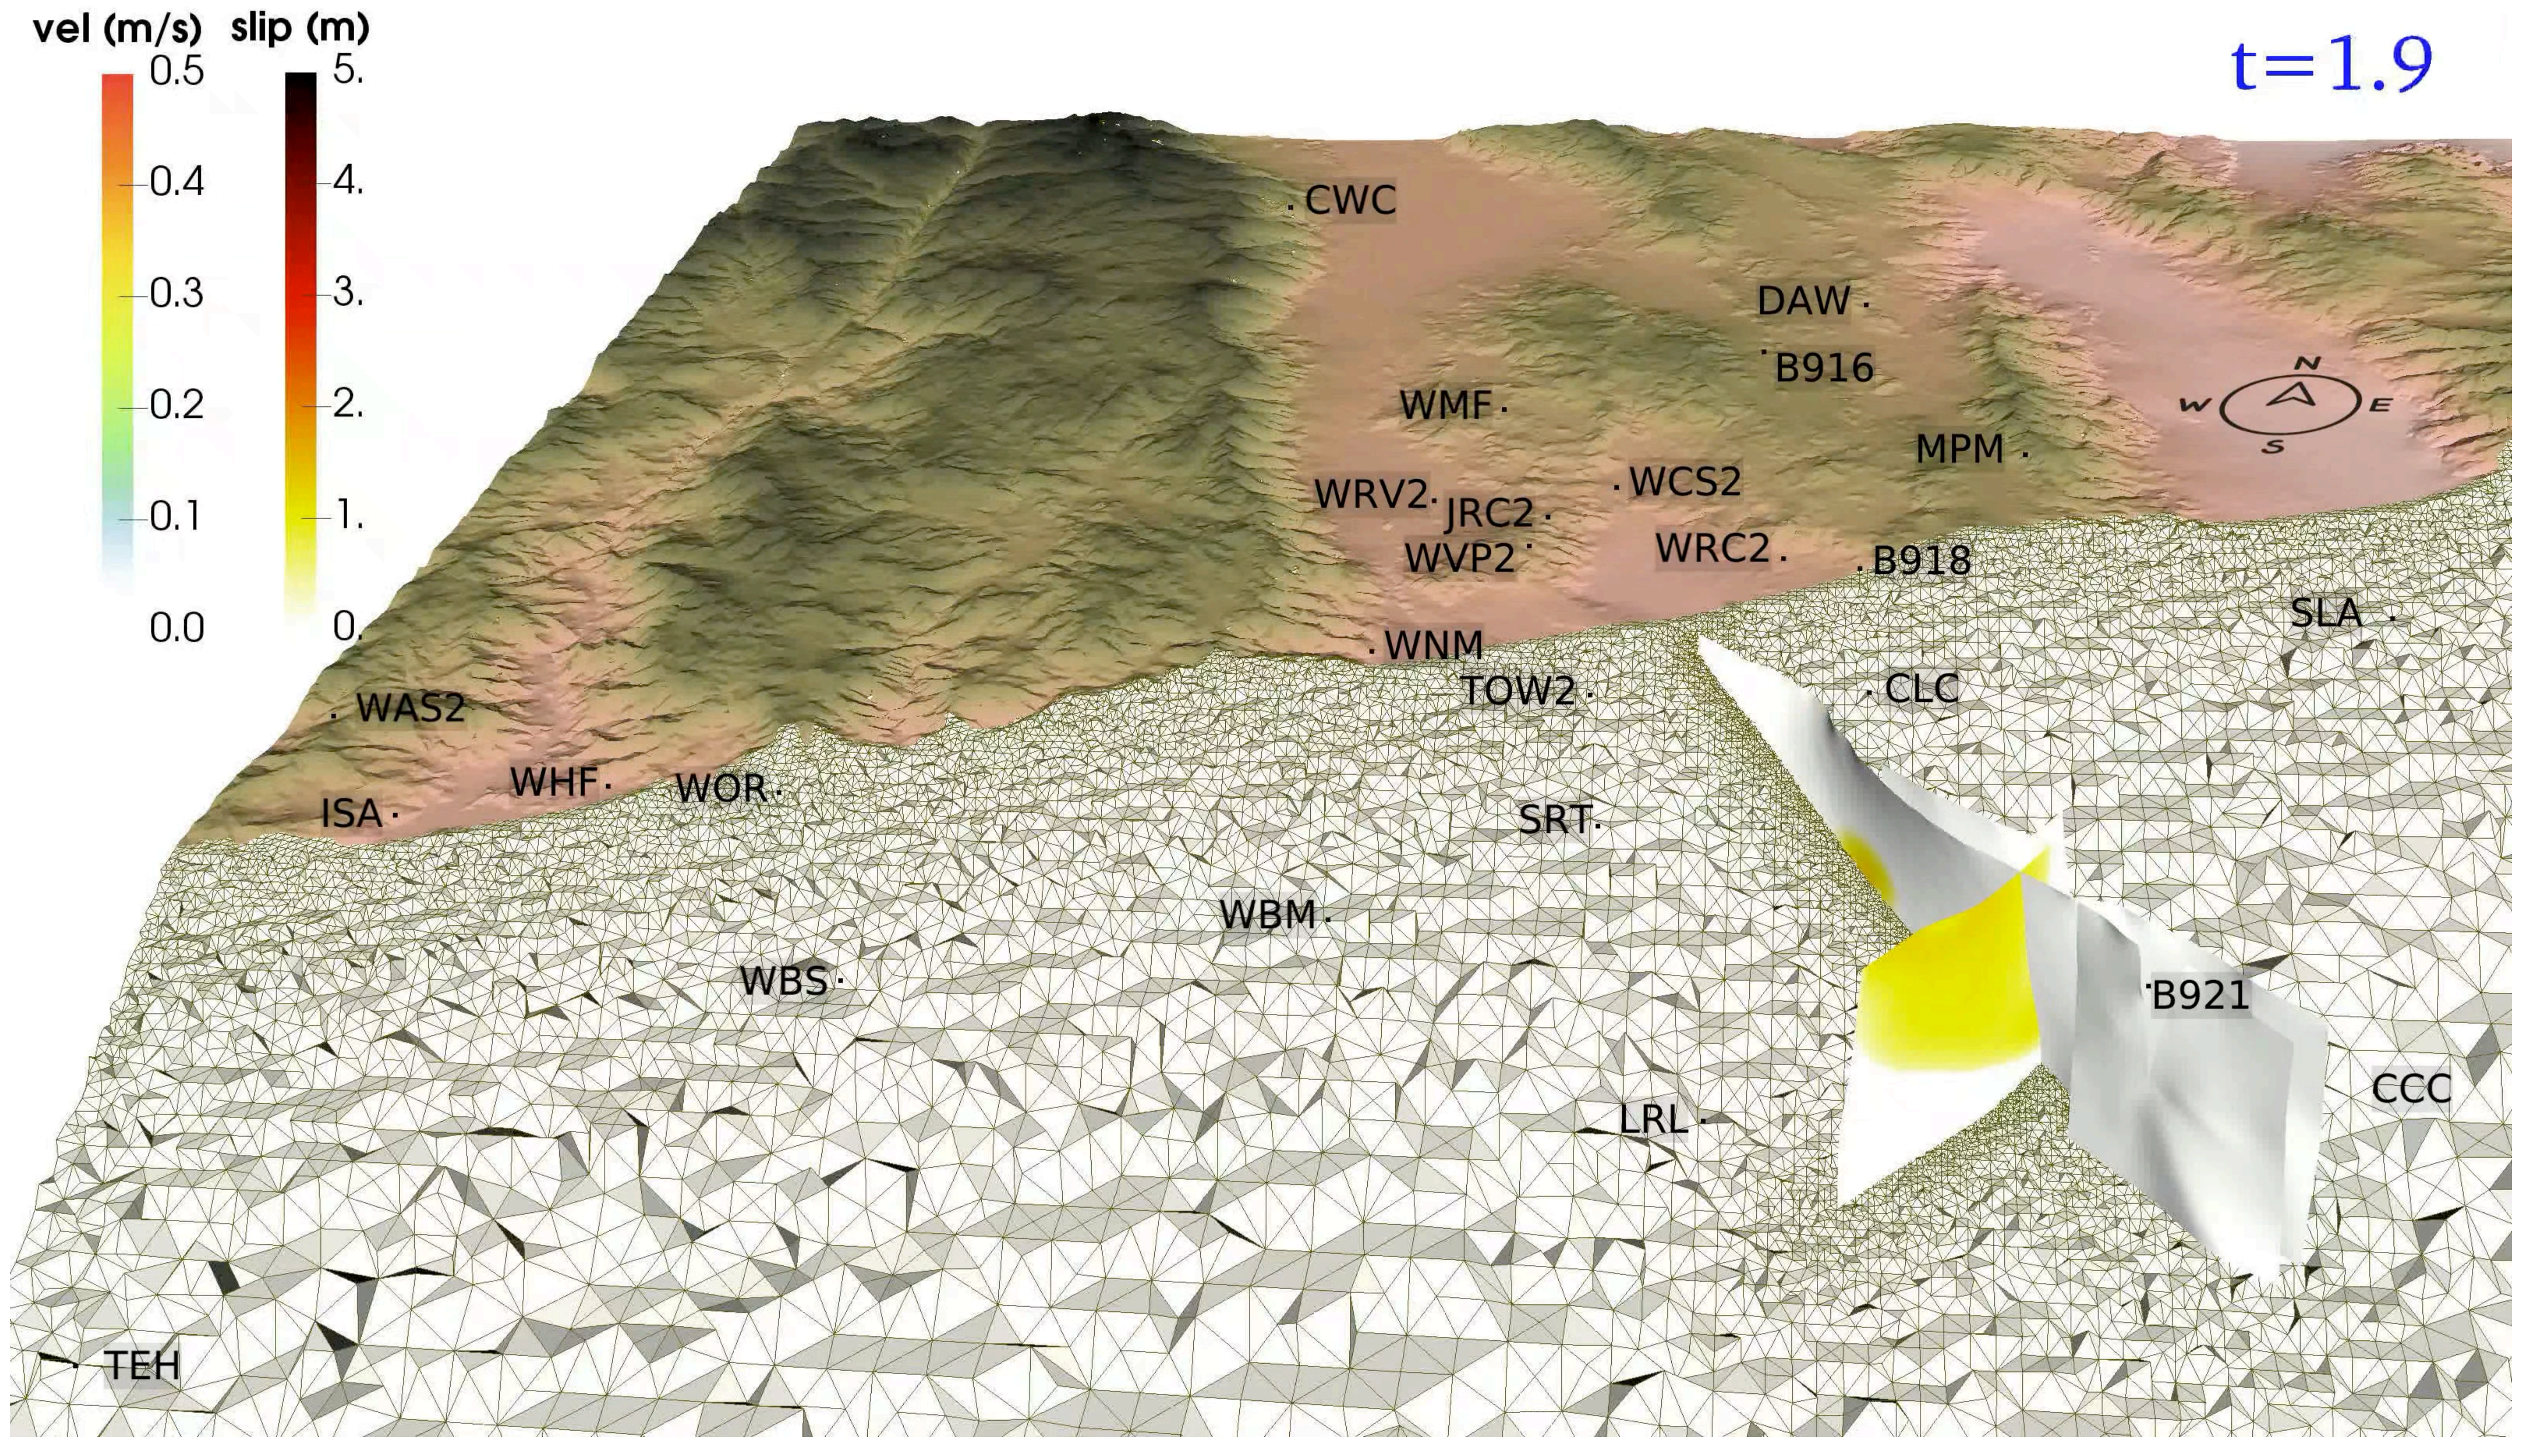

Dynamic rupture models of the Mw 6.4 Searles Valley and Mw 7.1 Ridgecrest earthquakes, Taufiqurrahman et al., 2023

A geometrically complex 3D fault network

- Both earthquakes were **highly complex**, including likely fault reactivation and being set apart in time by 34 hr while driving aftershocks, shallow aseismic creep, and swarm activity

- A **non-vertical quasi-orthogonal crosscutting 3D fault geometry** intersecting with **topography** from integrating geological field mapping of rupture traces, geodetical InSAR data, relocated seismicity of Ross et al., 2019 and selected focal mechanisms from SCEDC catalog (Carena and Suppe, 2002)

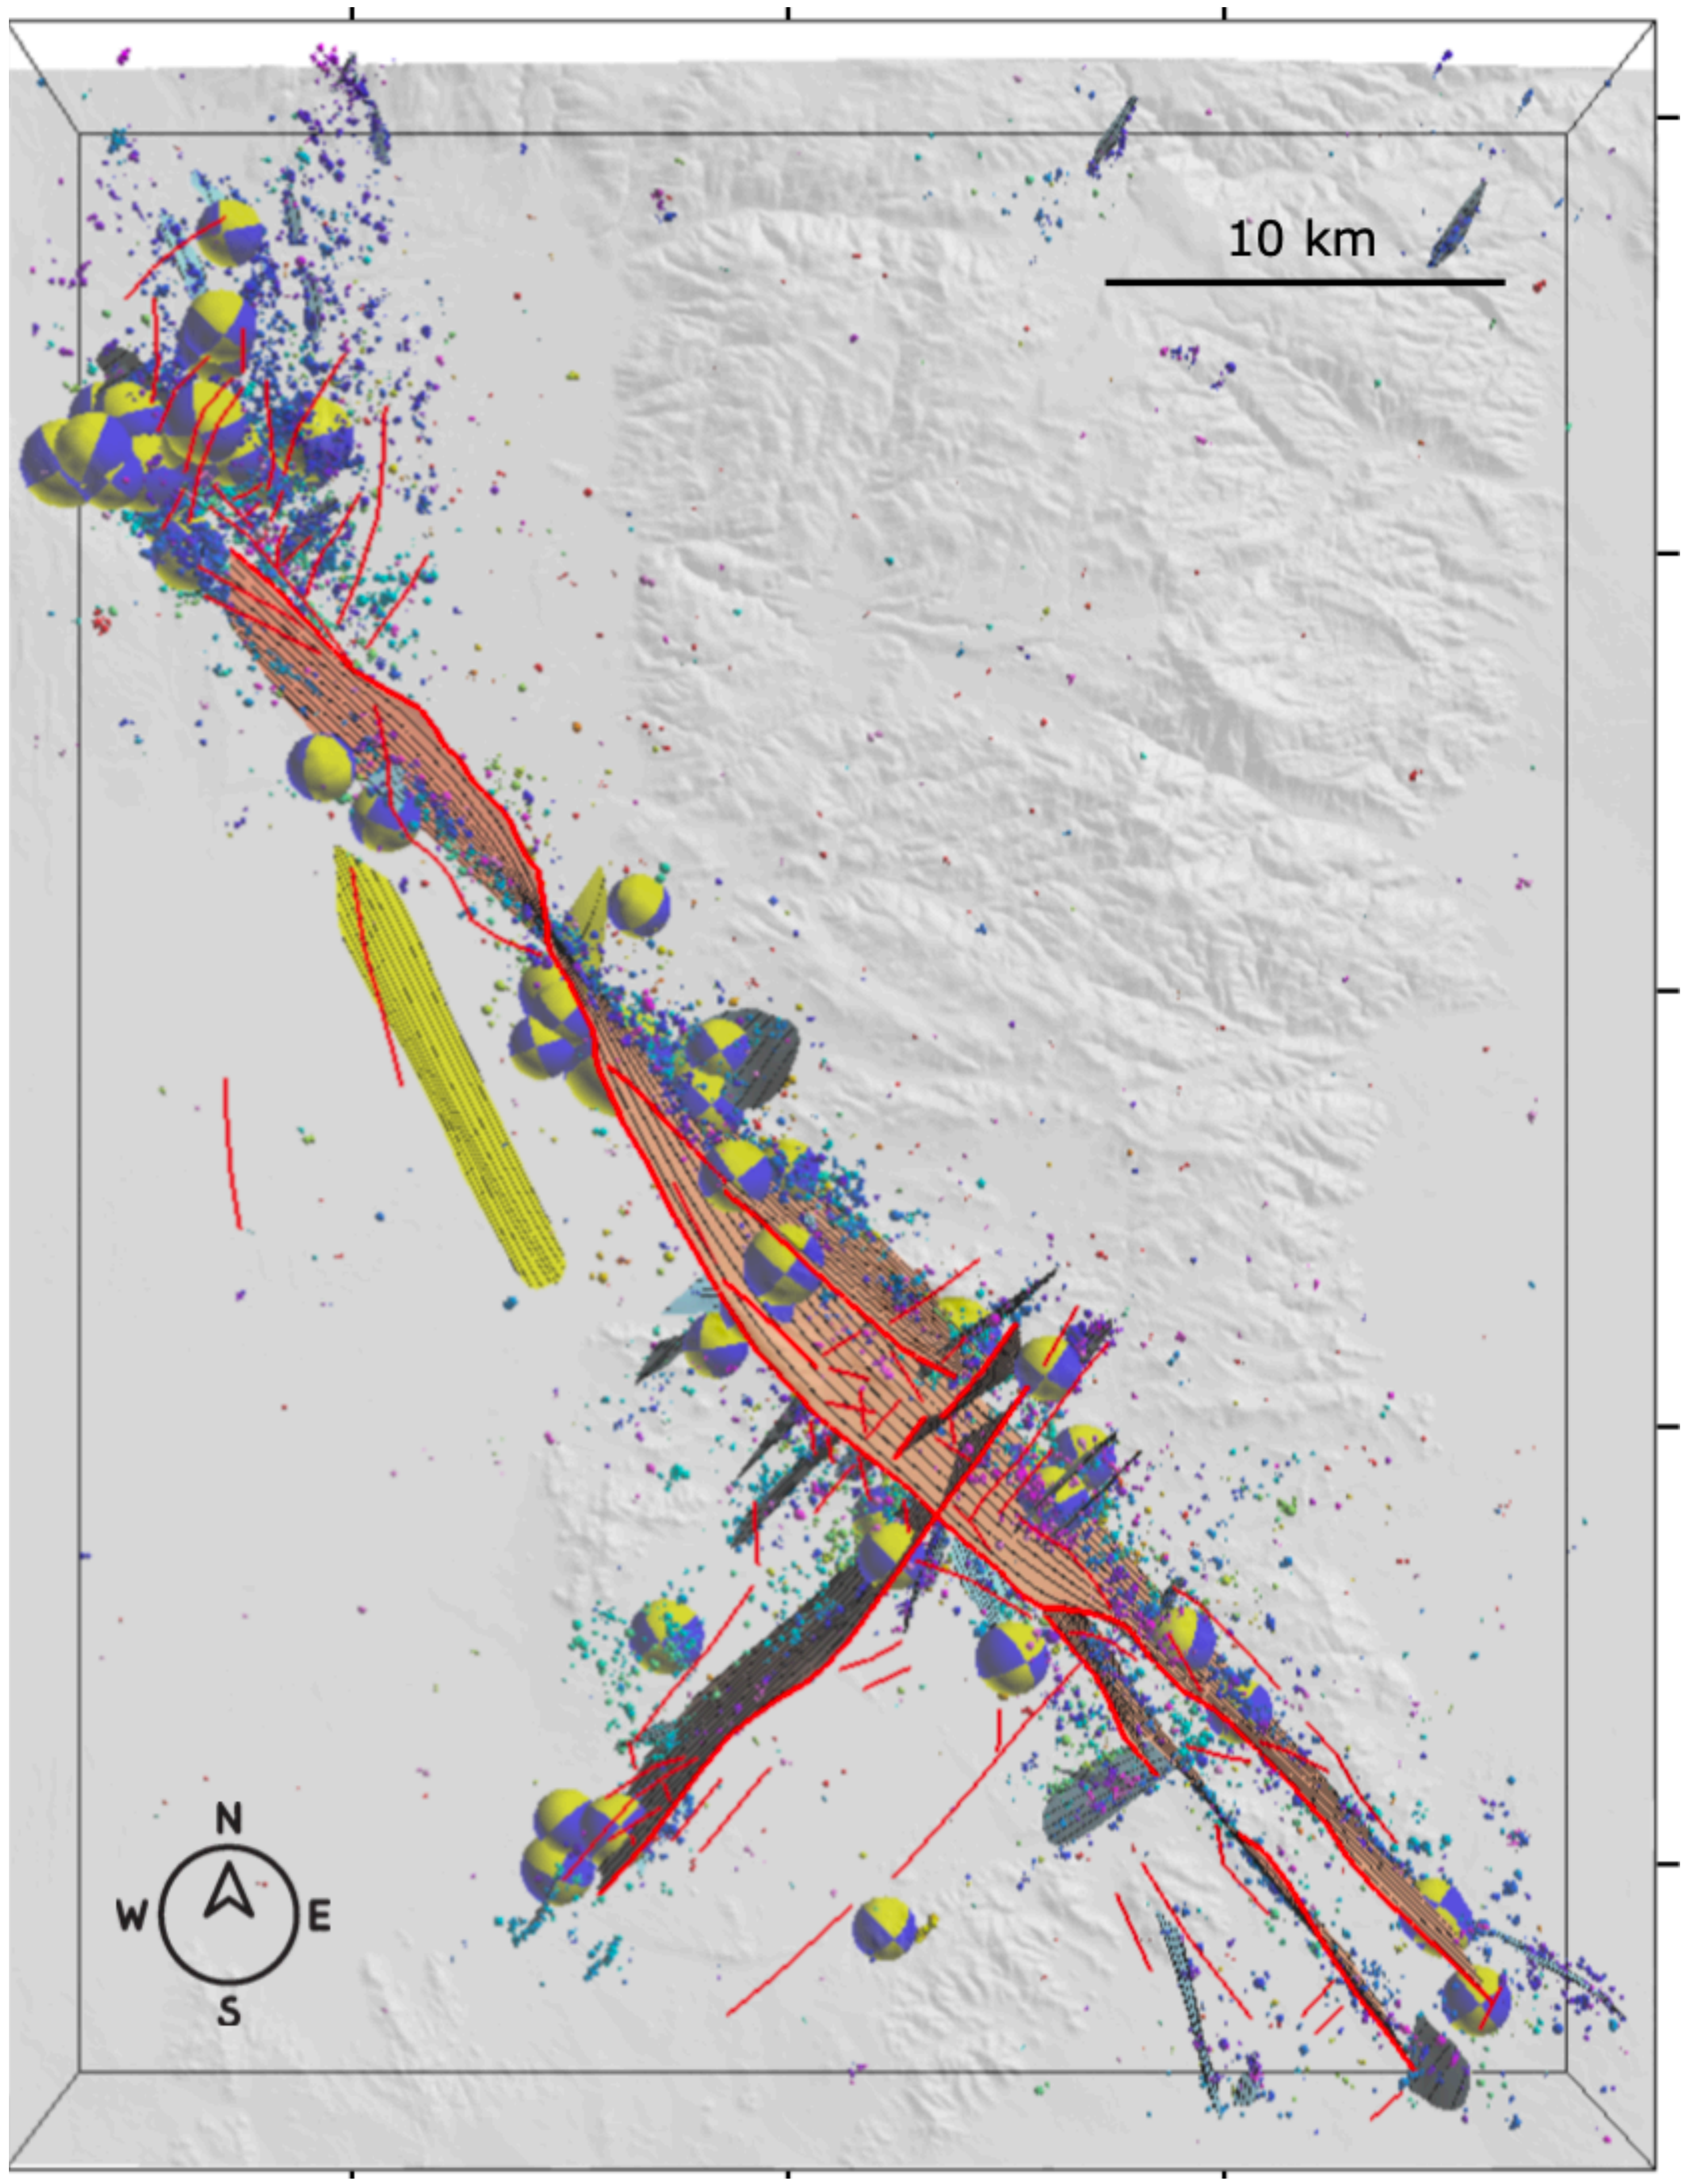

Taufiqurrahman et al., Nature, 2023

3D structure and loading stresses

- All faults are embedded in the **3D CMV-S** velocity model and exposed to **ambient tectonic stress state** (SCEC CSM YHSM-2013, Yang & Hauksson 2012)

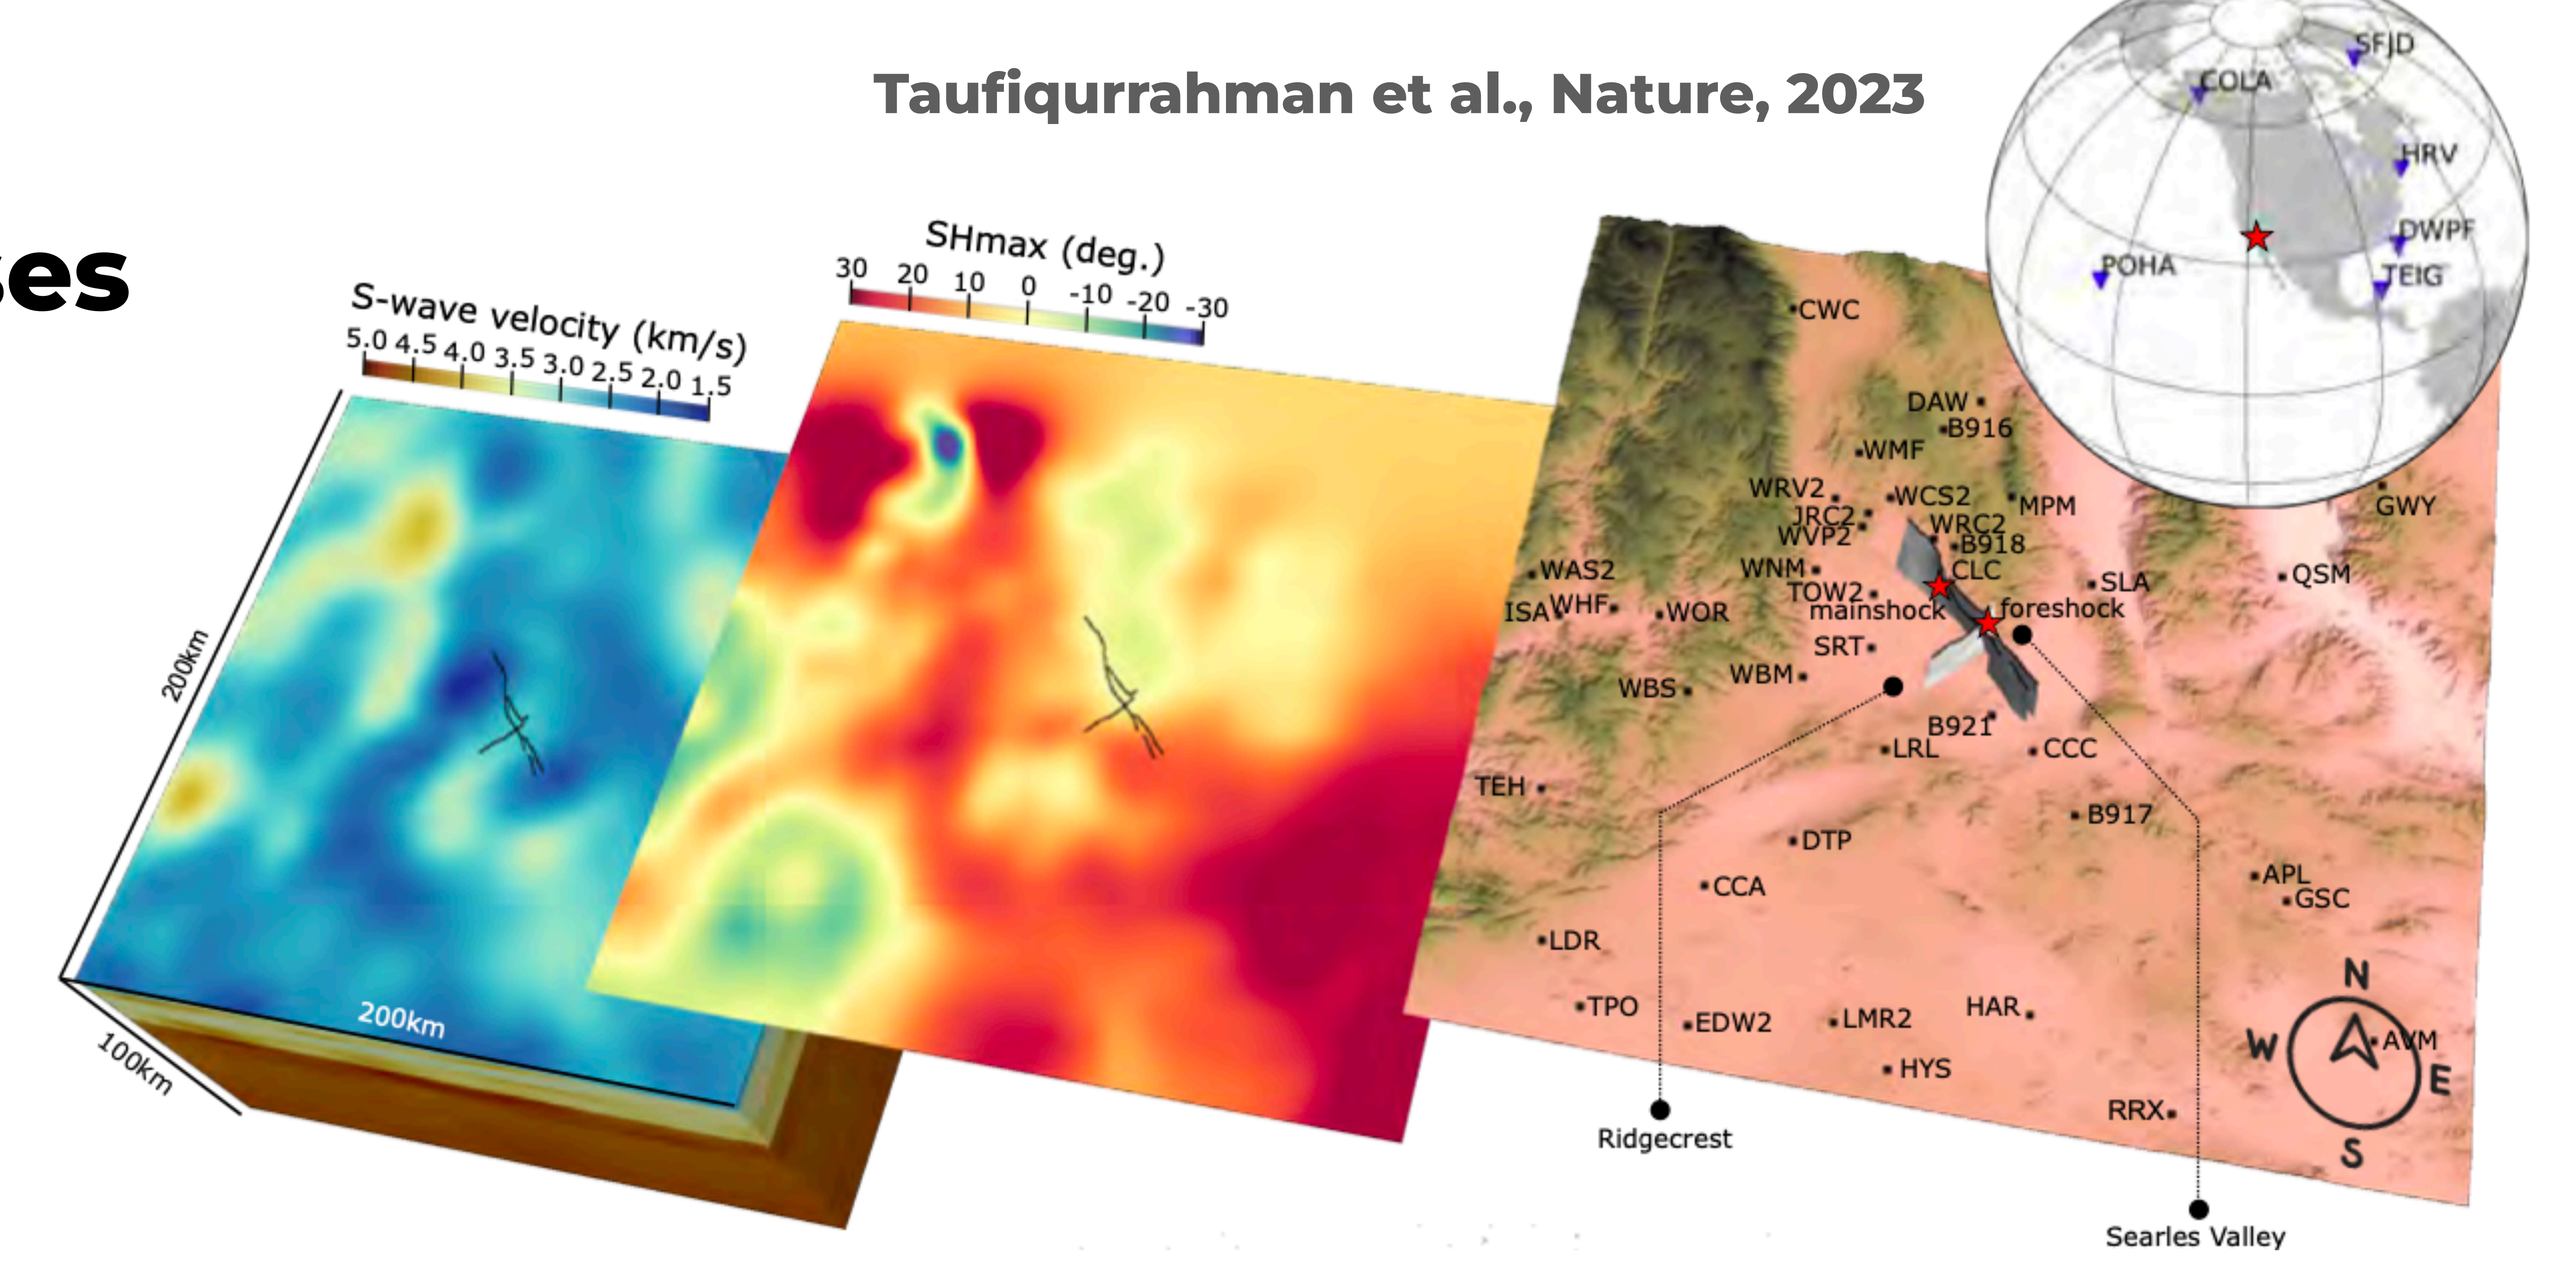

Taufiqurrahman et al., Nature, 2023

3D structure and loading stresses

- All faults are embedded in the **3D CMV-S** velocity model and exposed to **ambient tectonic stress state** (SCEC CSM YHSM-2013, Yang & Hauksson 2012)

- pre-Ridgecrest **co-seismic and post-seismic Coulomb failure stress changes** (dCFS) due to previous major earthquakes occurring in the region in the last **~1400 years** (Verdecchia & Carena, 2016; Friedrich et al., 2019) yields positive stress redistribution additionally loading the source region

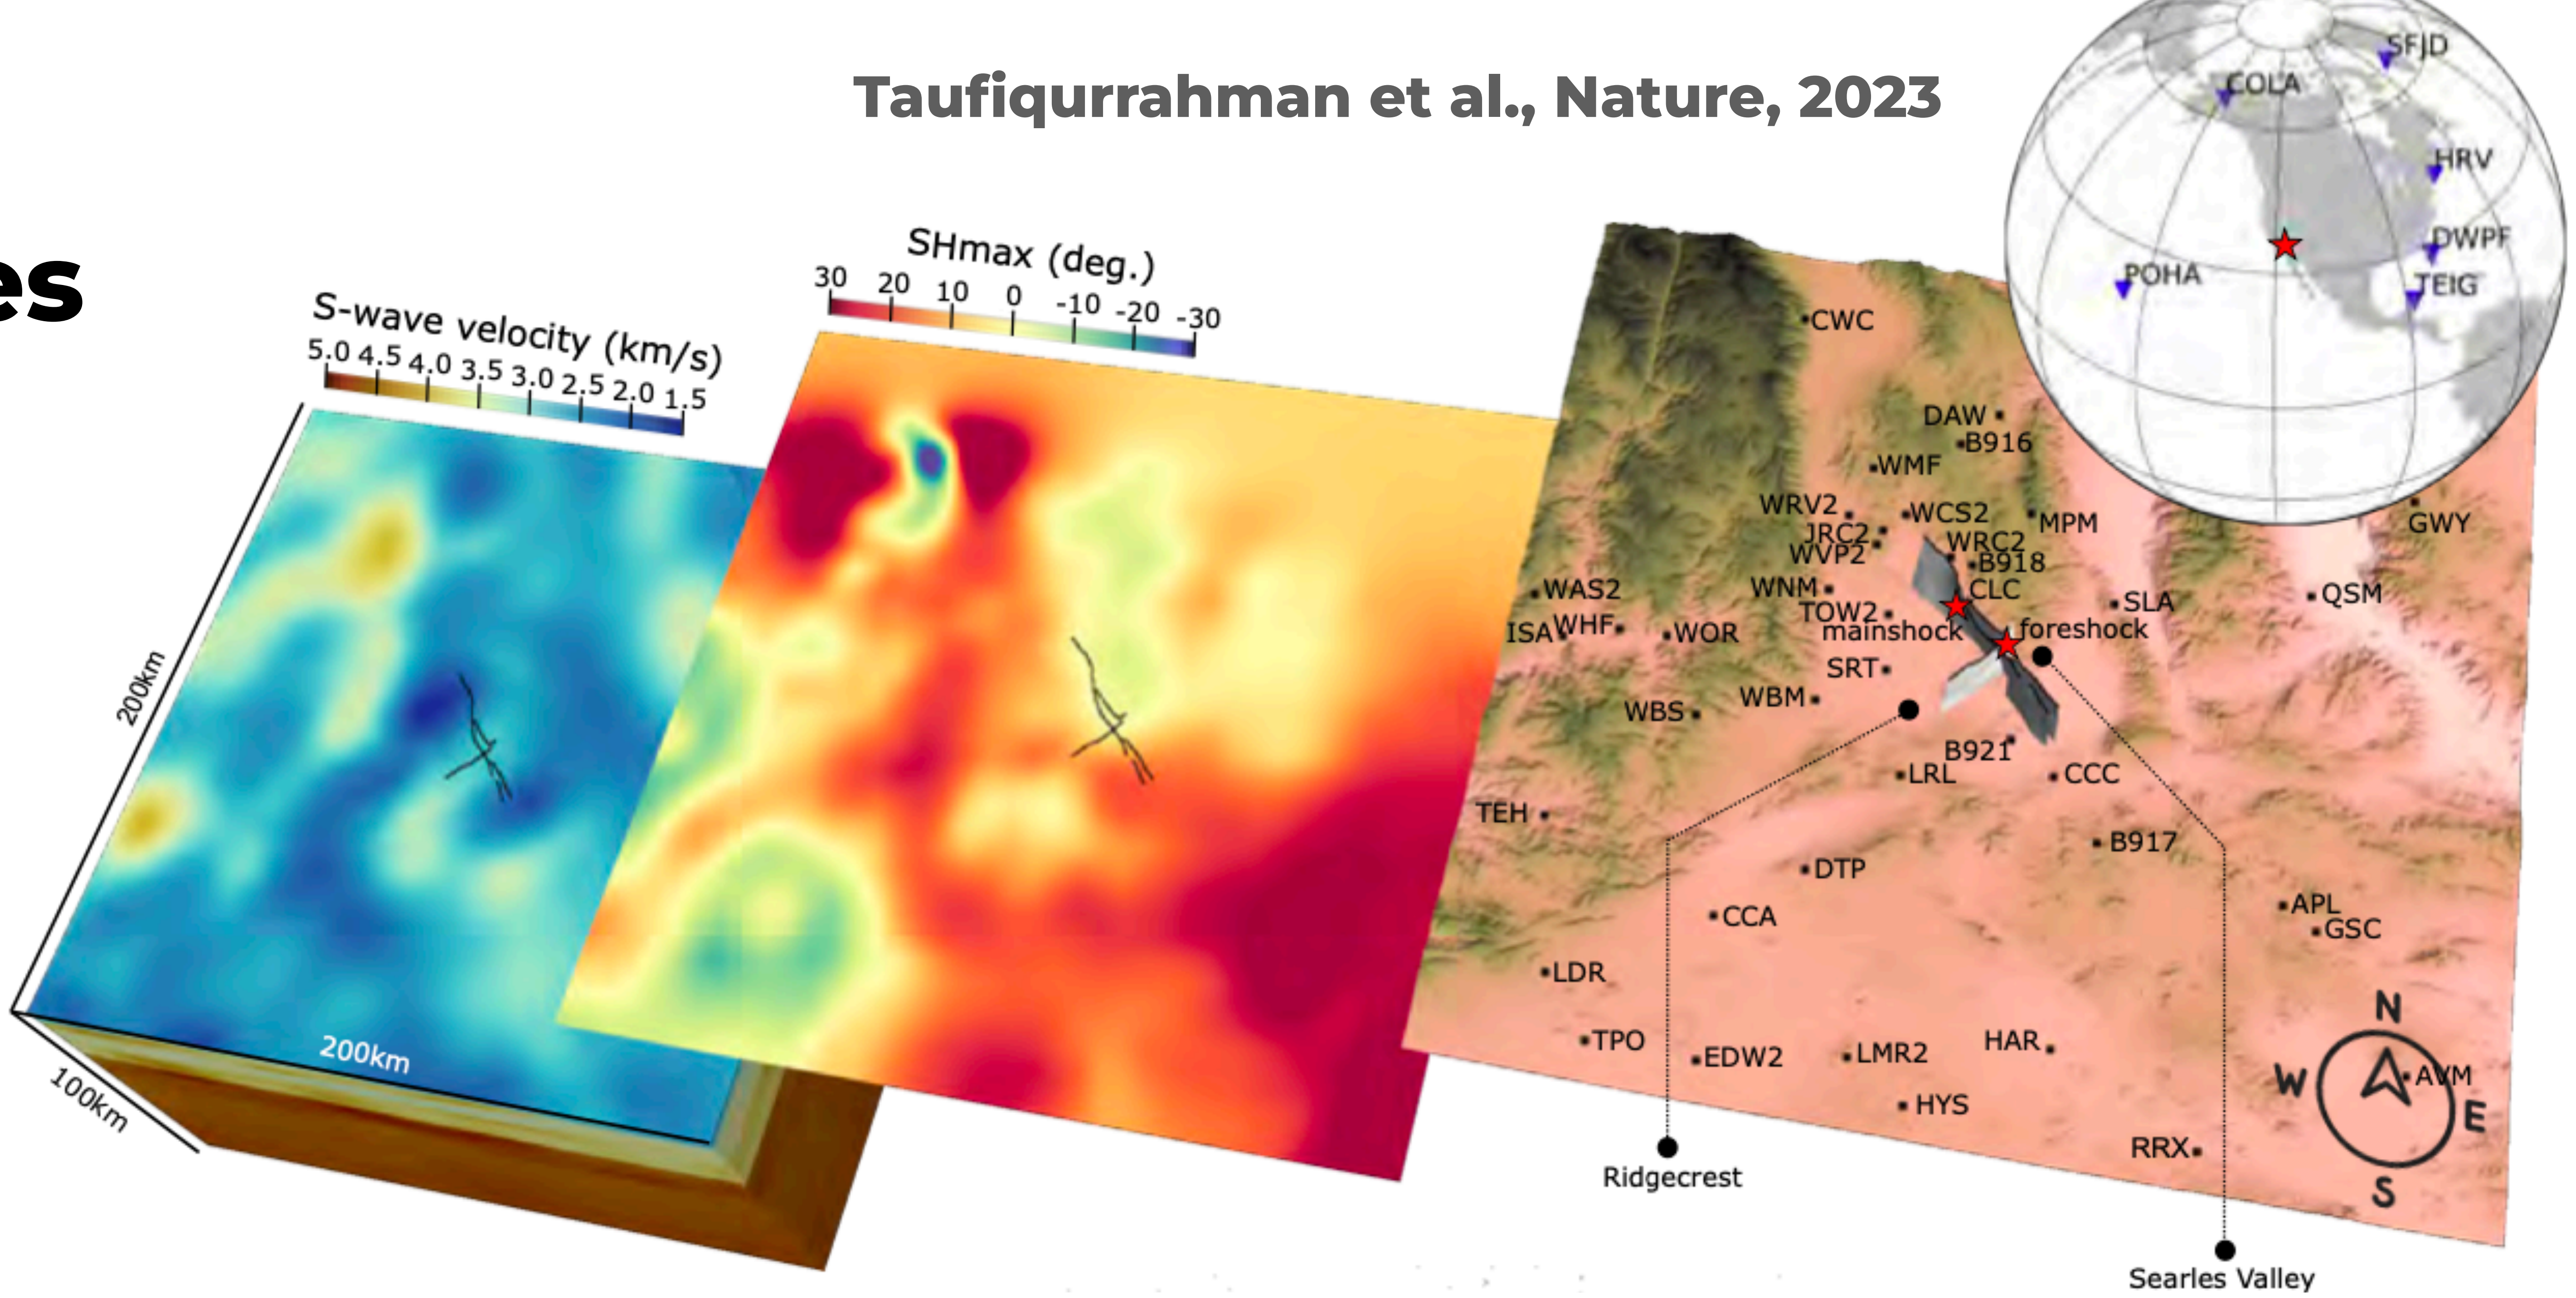

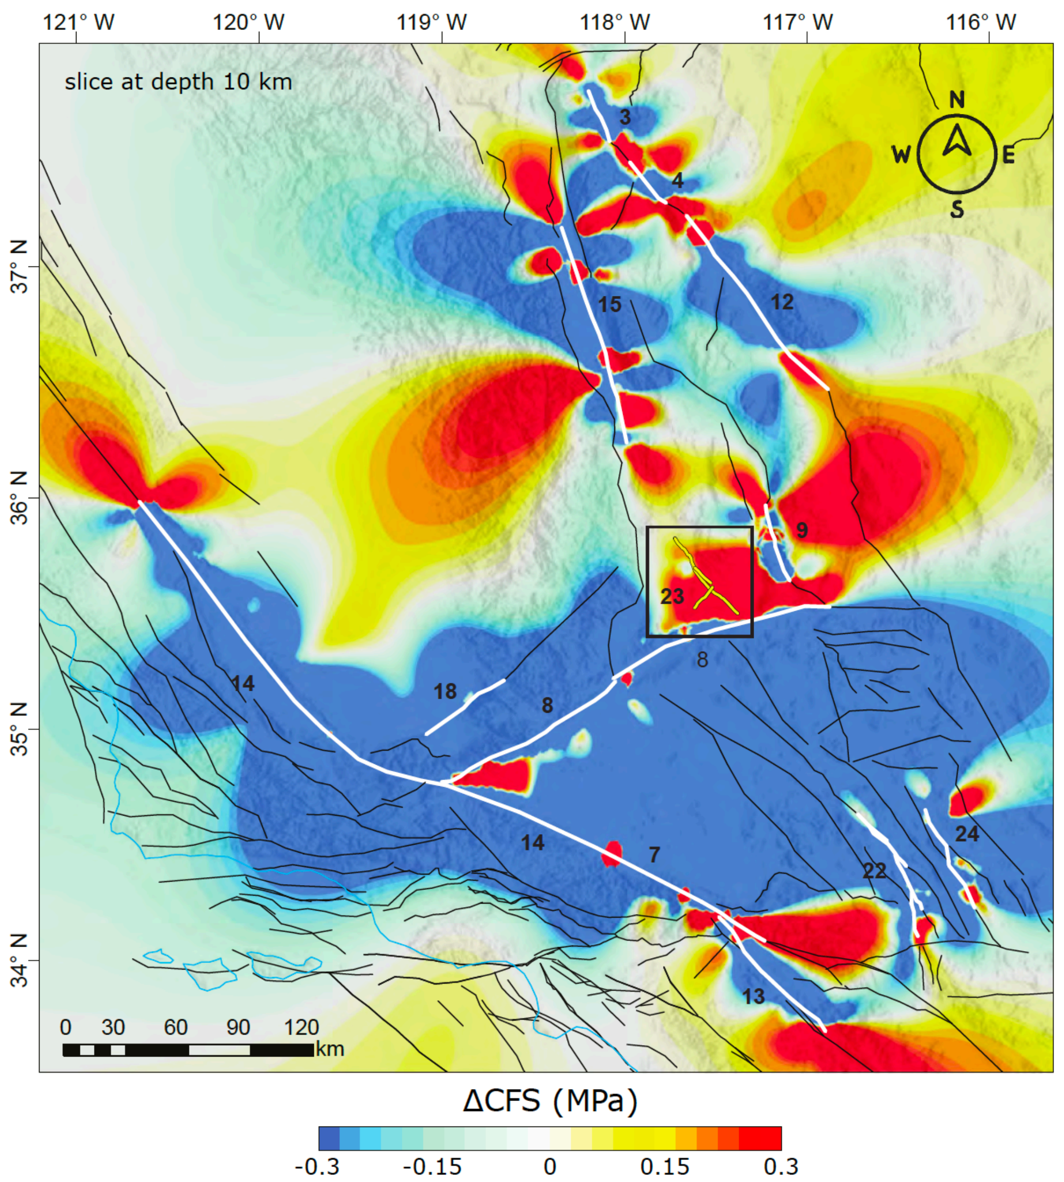

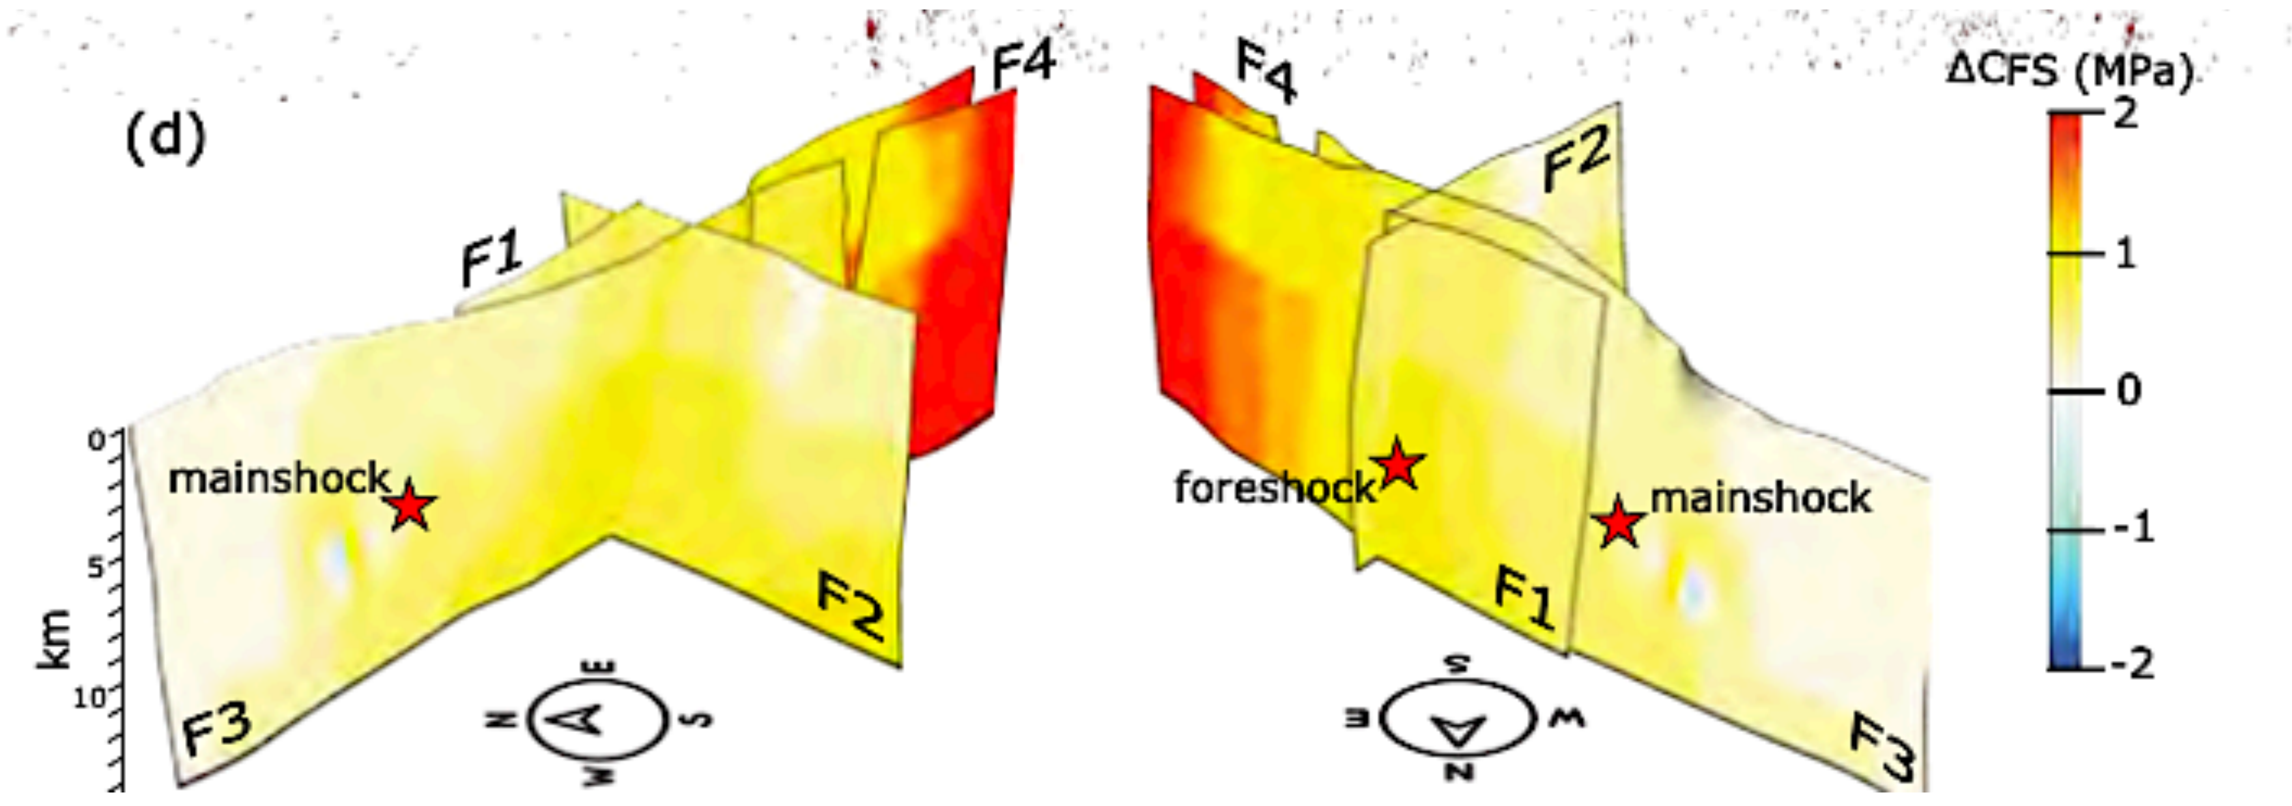

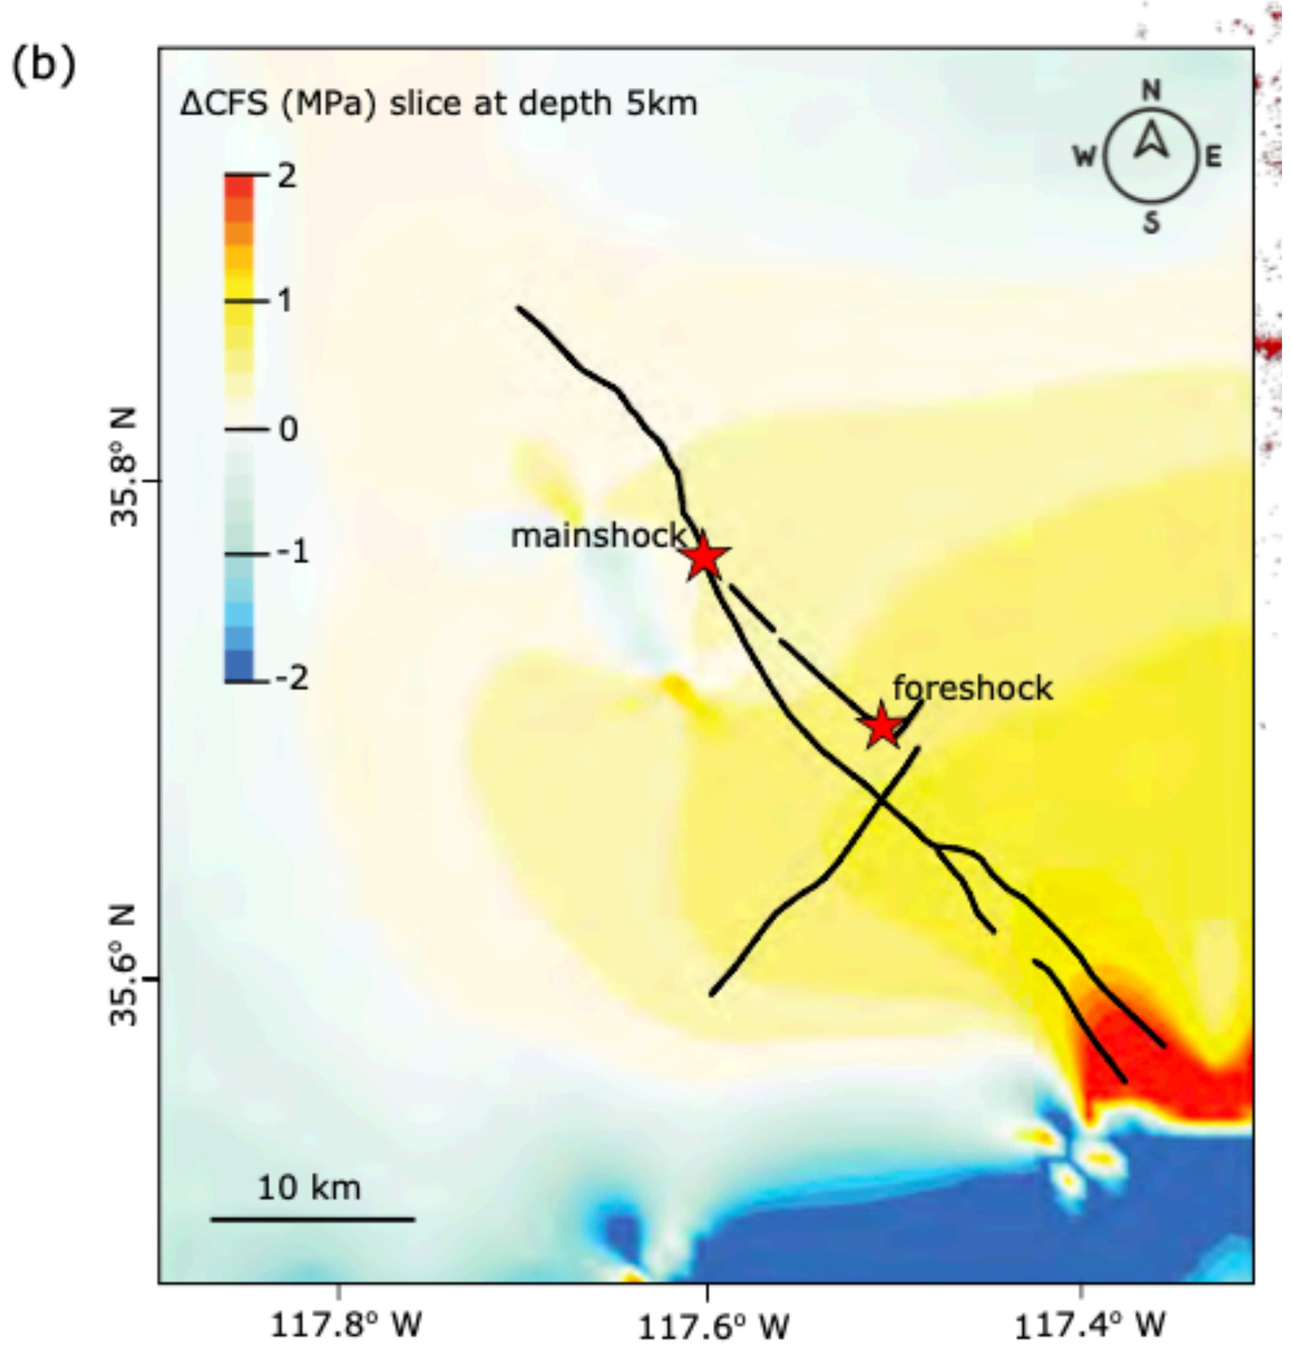

Fault strength

Taufiqurrahman et al., Nature, 2023

- **Apparently weak faults** due to combined effects of **severe velocity-weakening friction, elevated fluid pressure**
- **Relative fault strength can be constrained observationally**

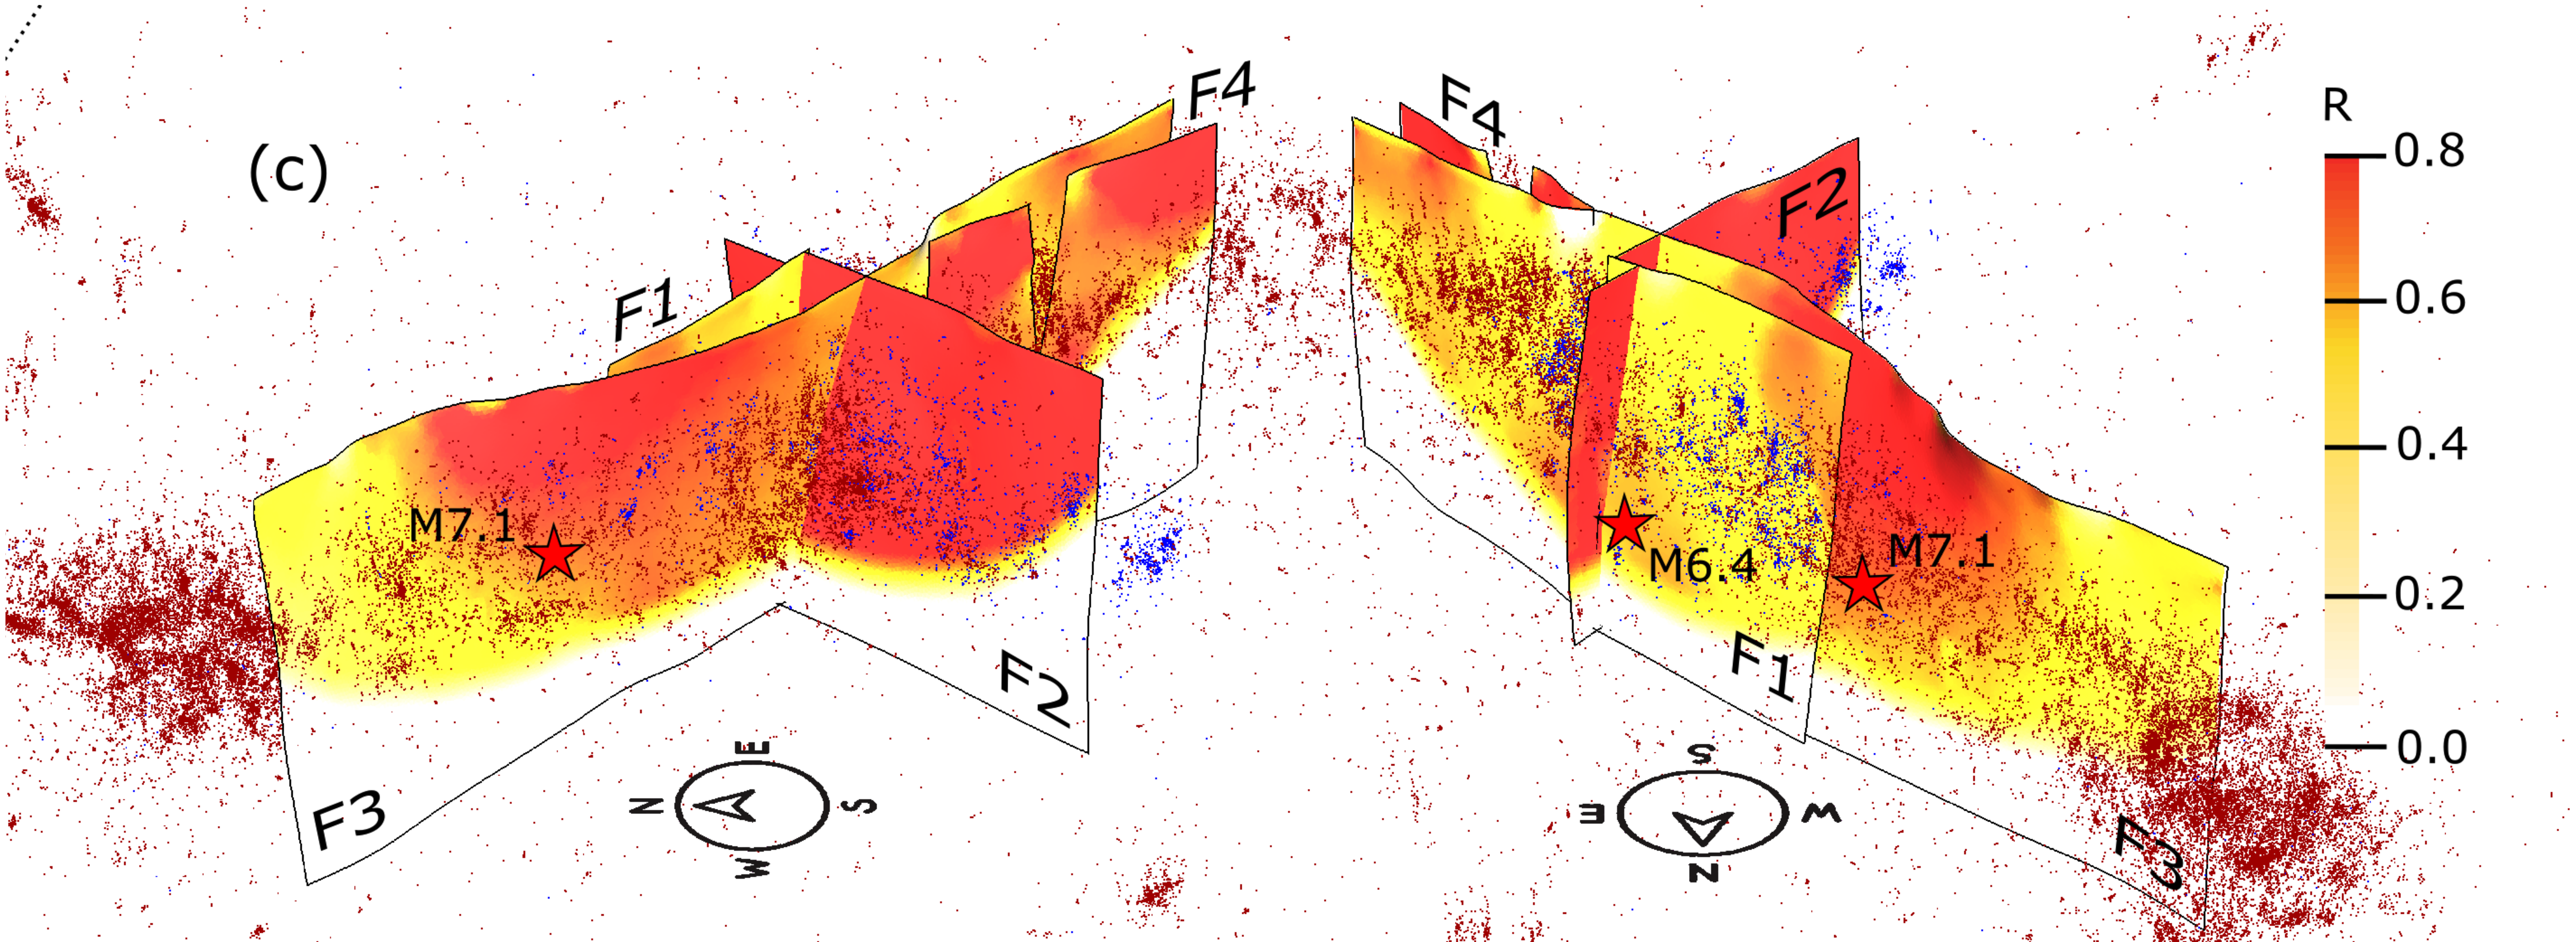

Model Verification - Mw 6.4 Searles Valley fore shock

Taufiqurrahman et al., Nature, 2023

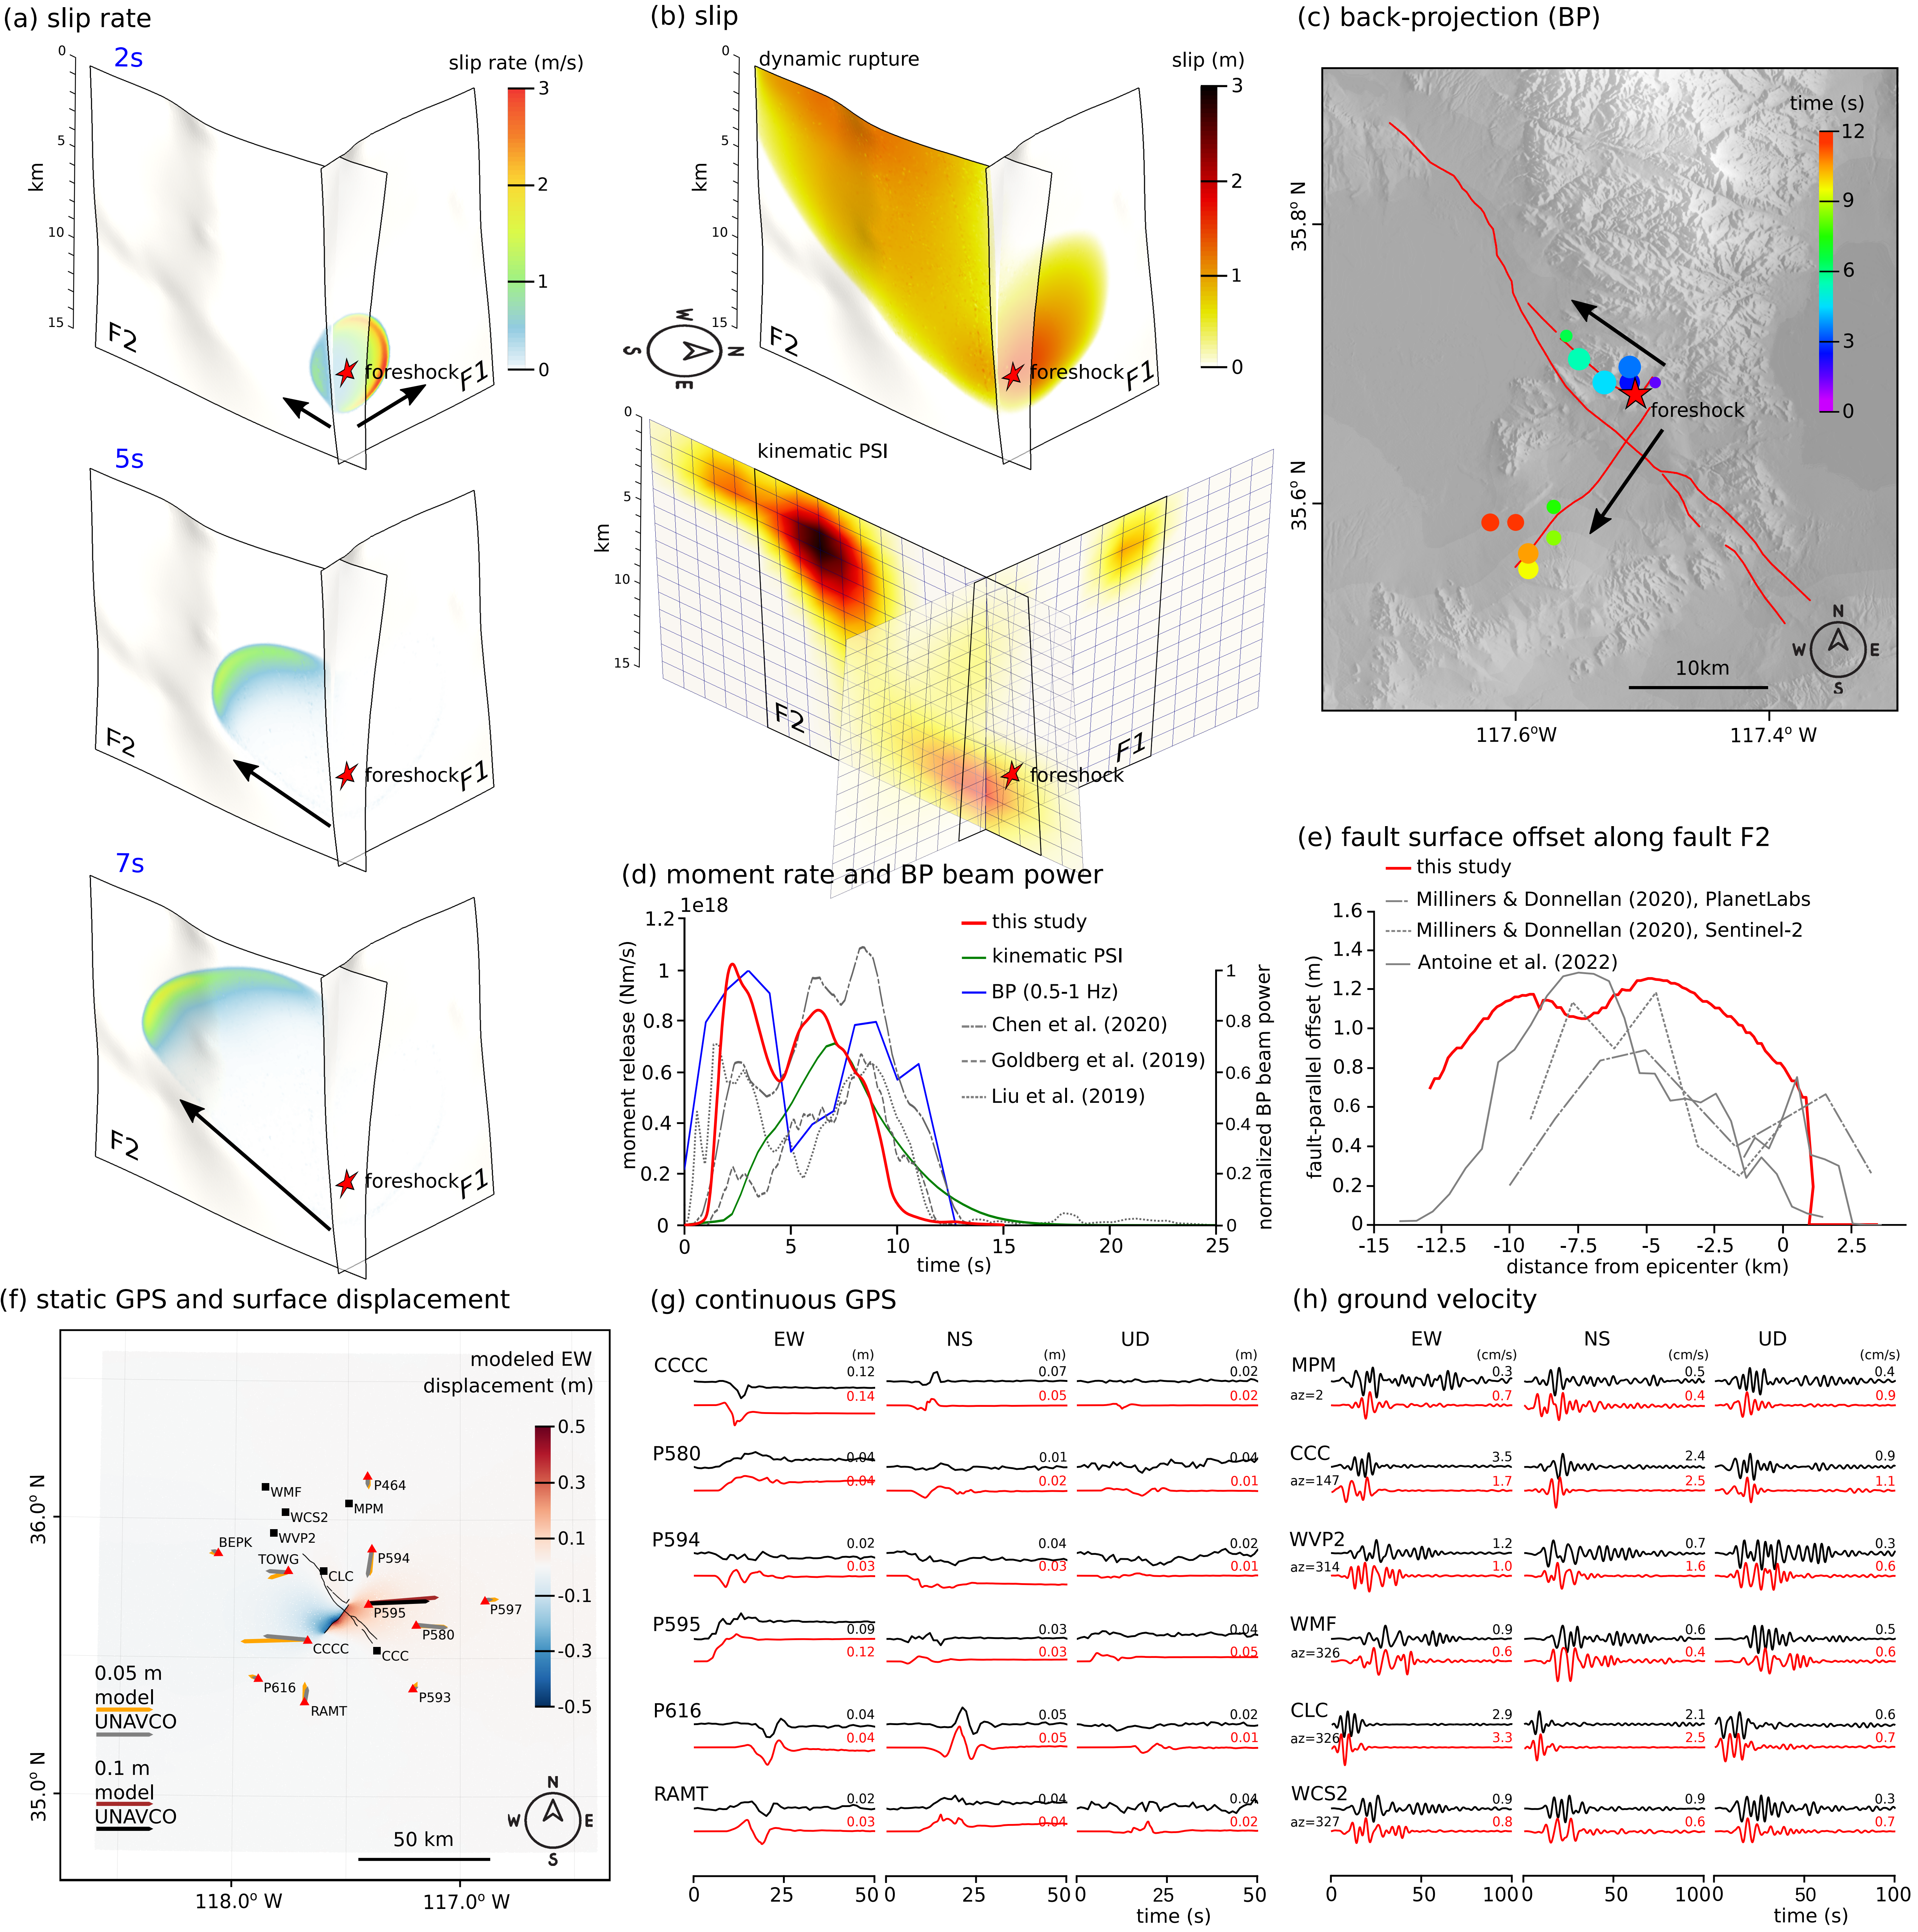

Model Verification - Mw 7.1 Ridgecrest main shock

Taufiqurrahman et al., Nature, 2023

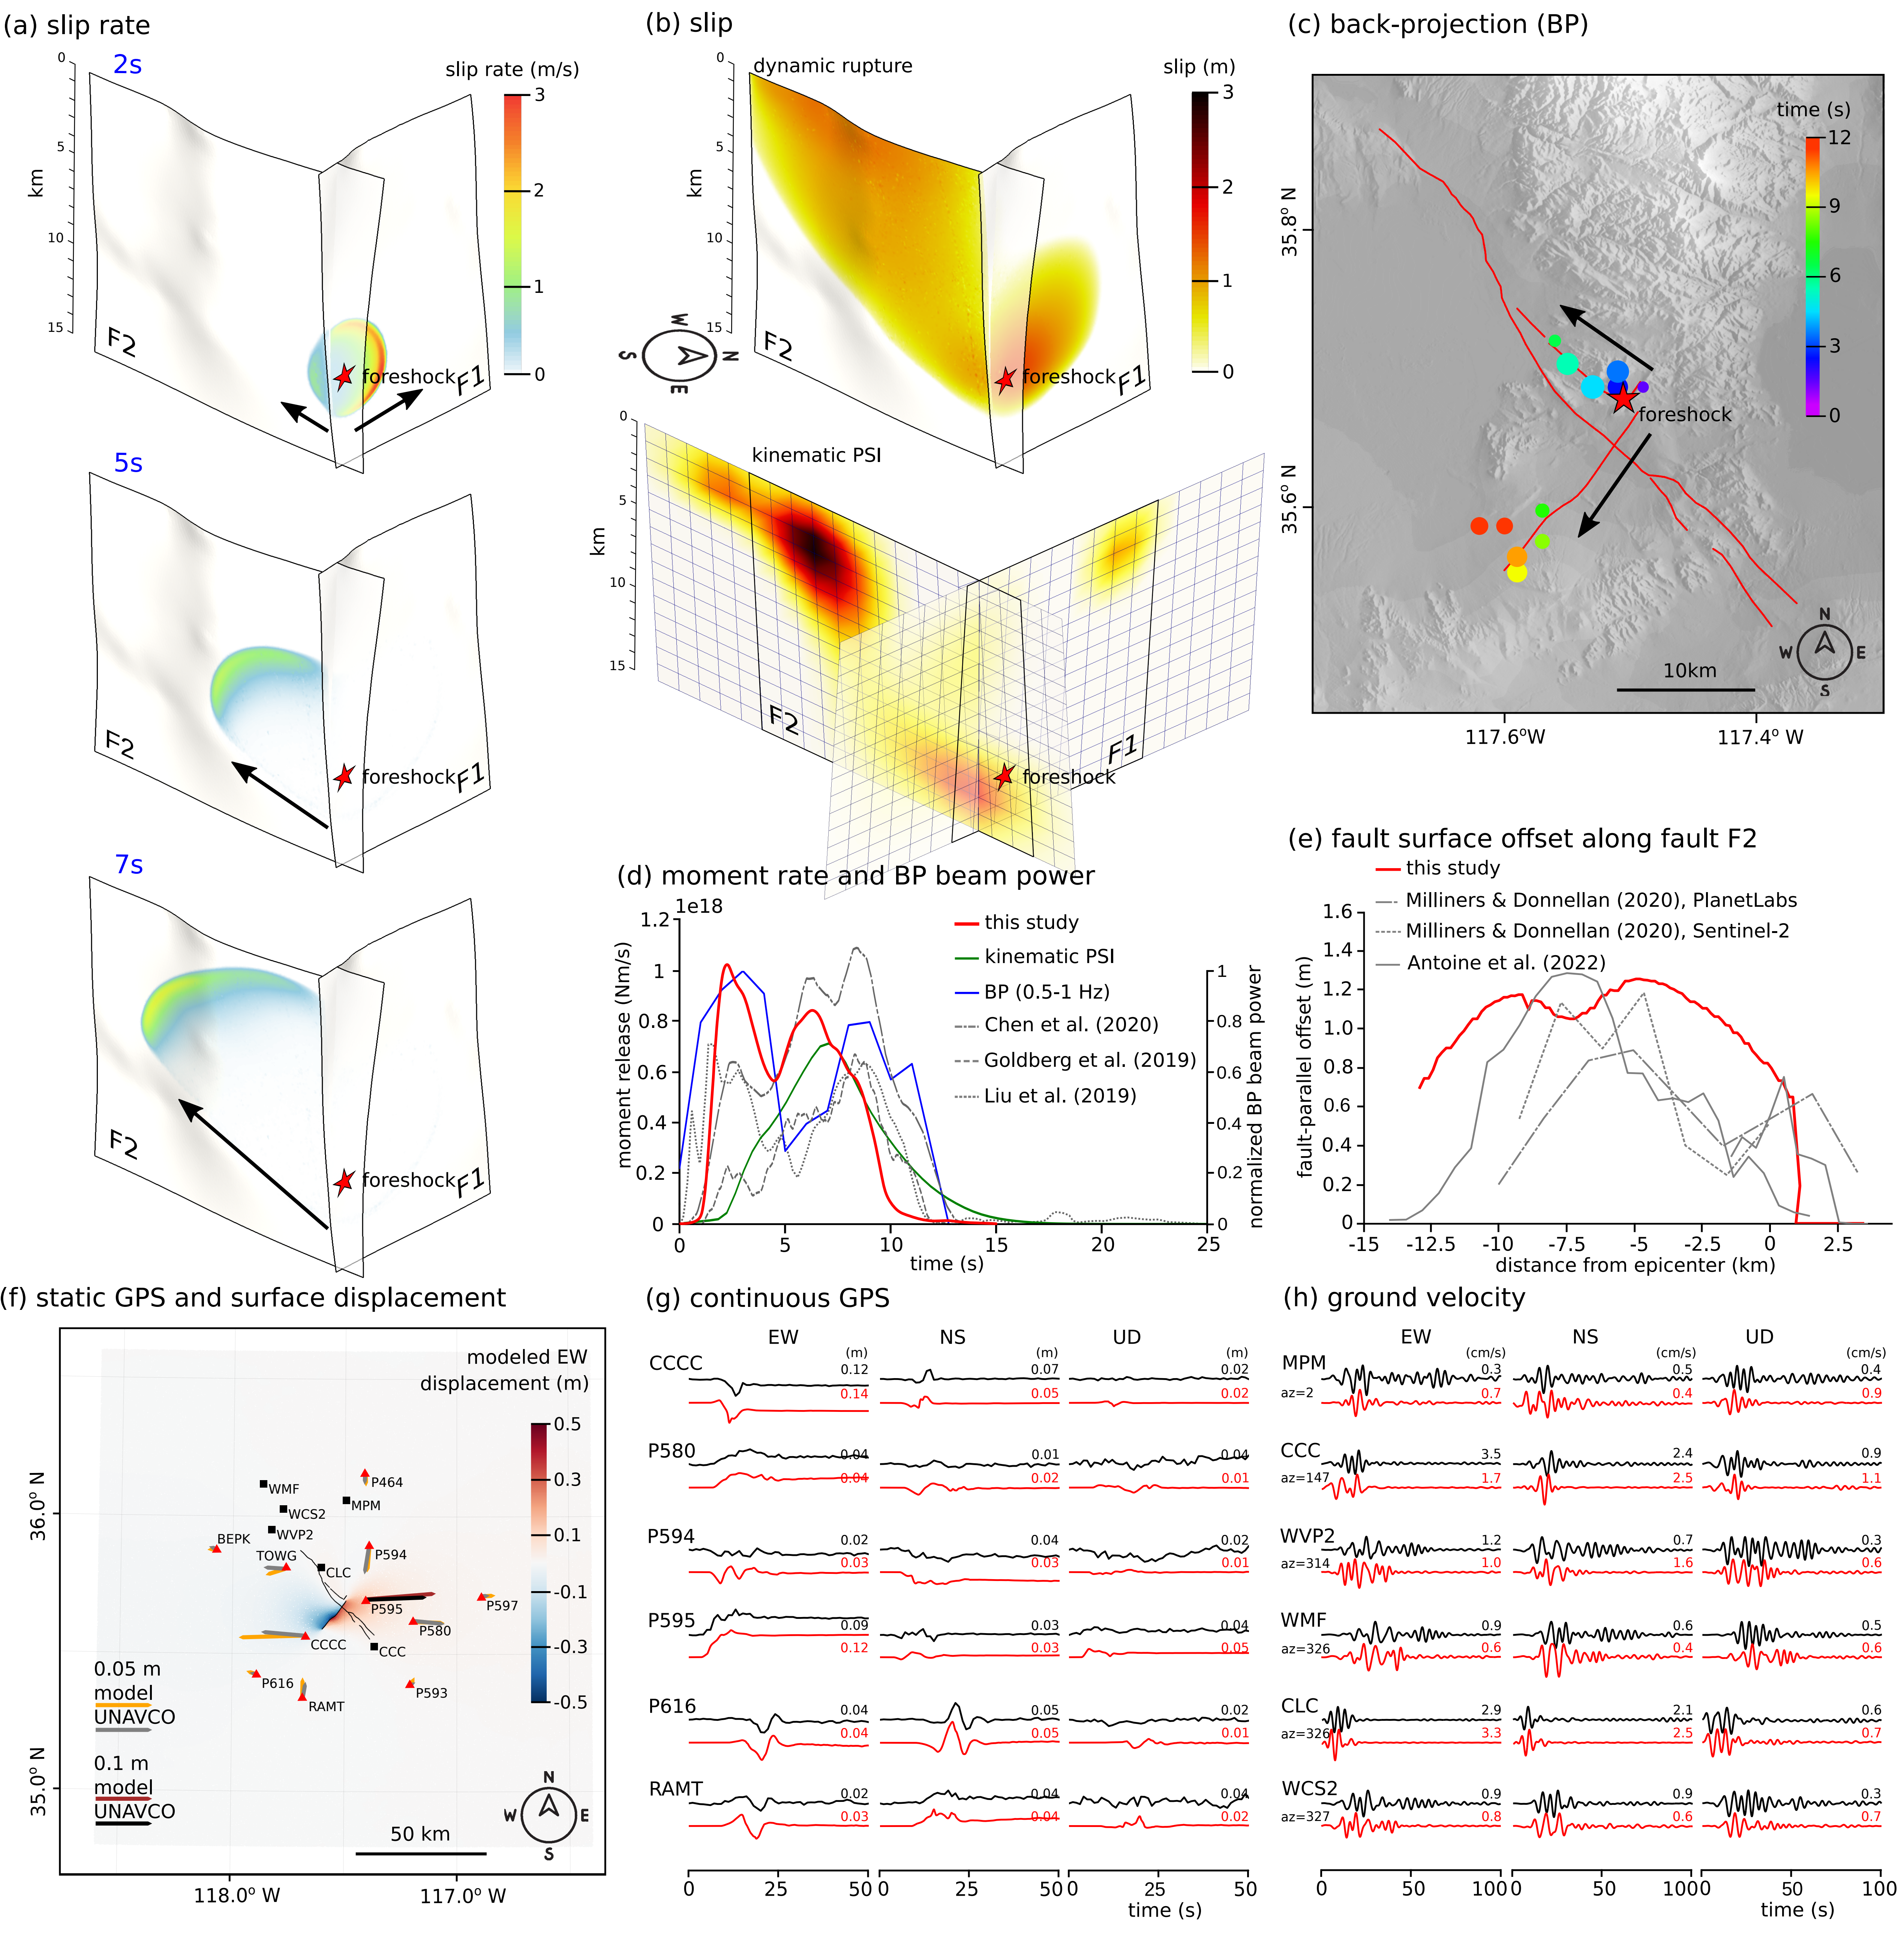

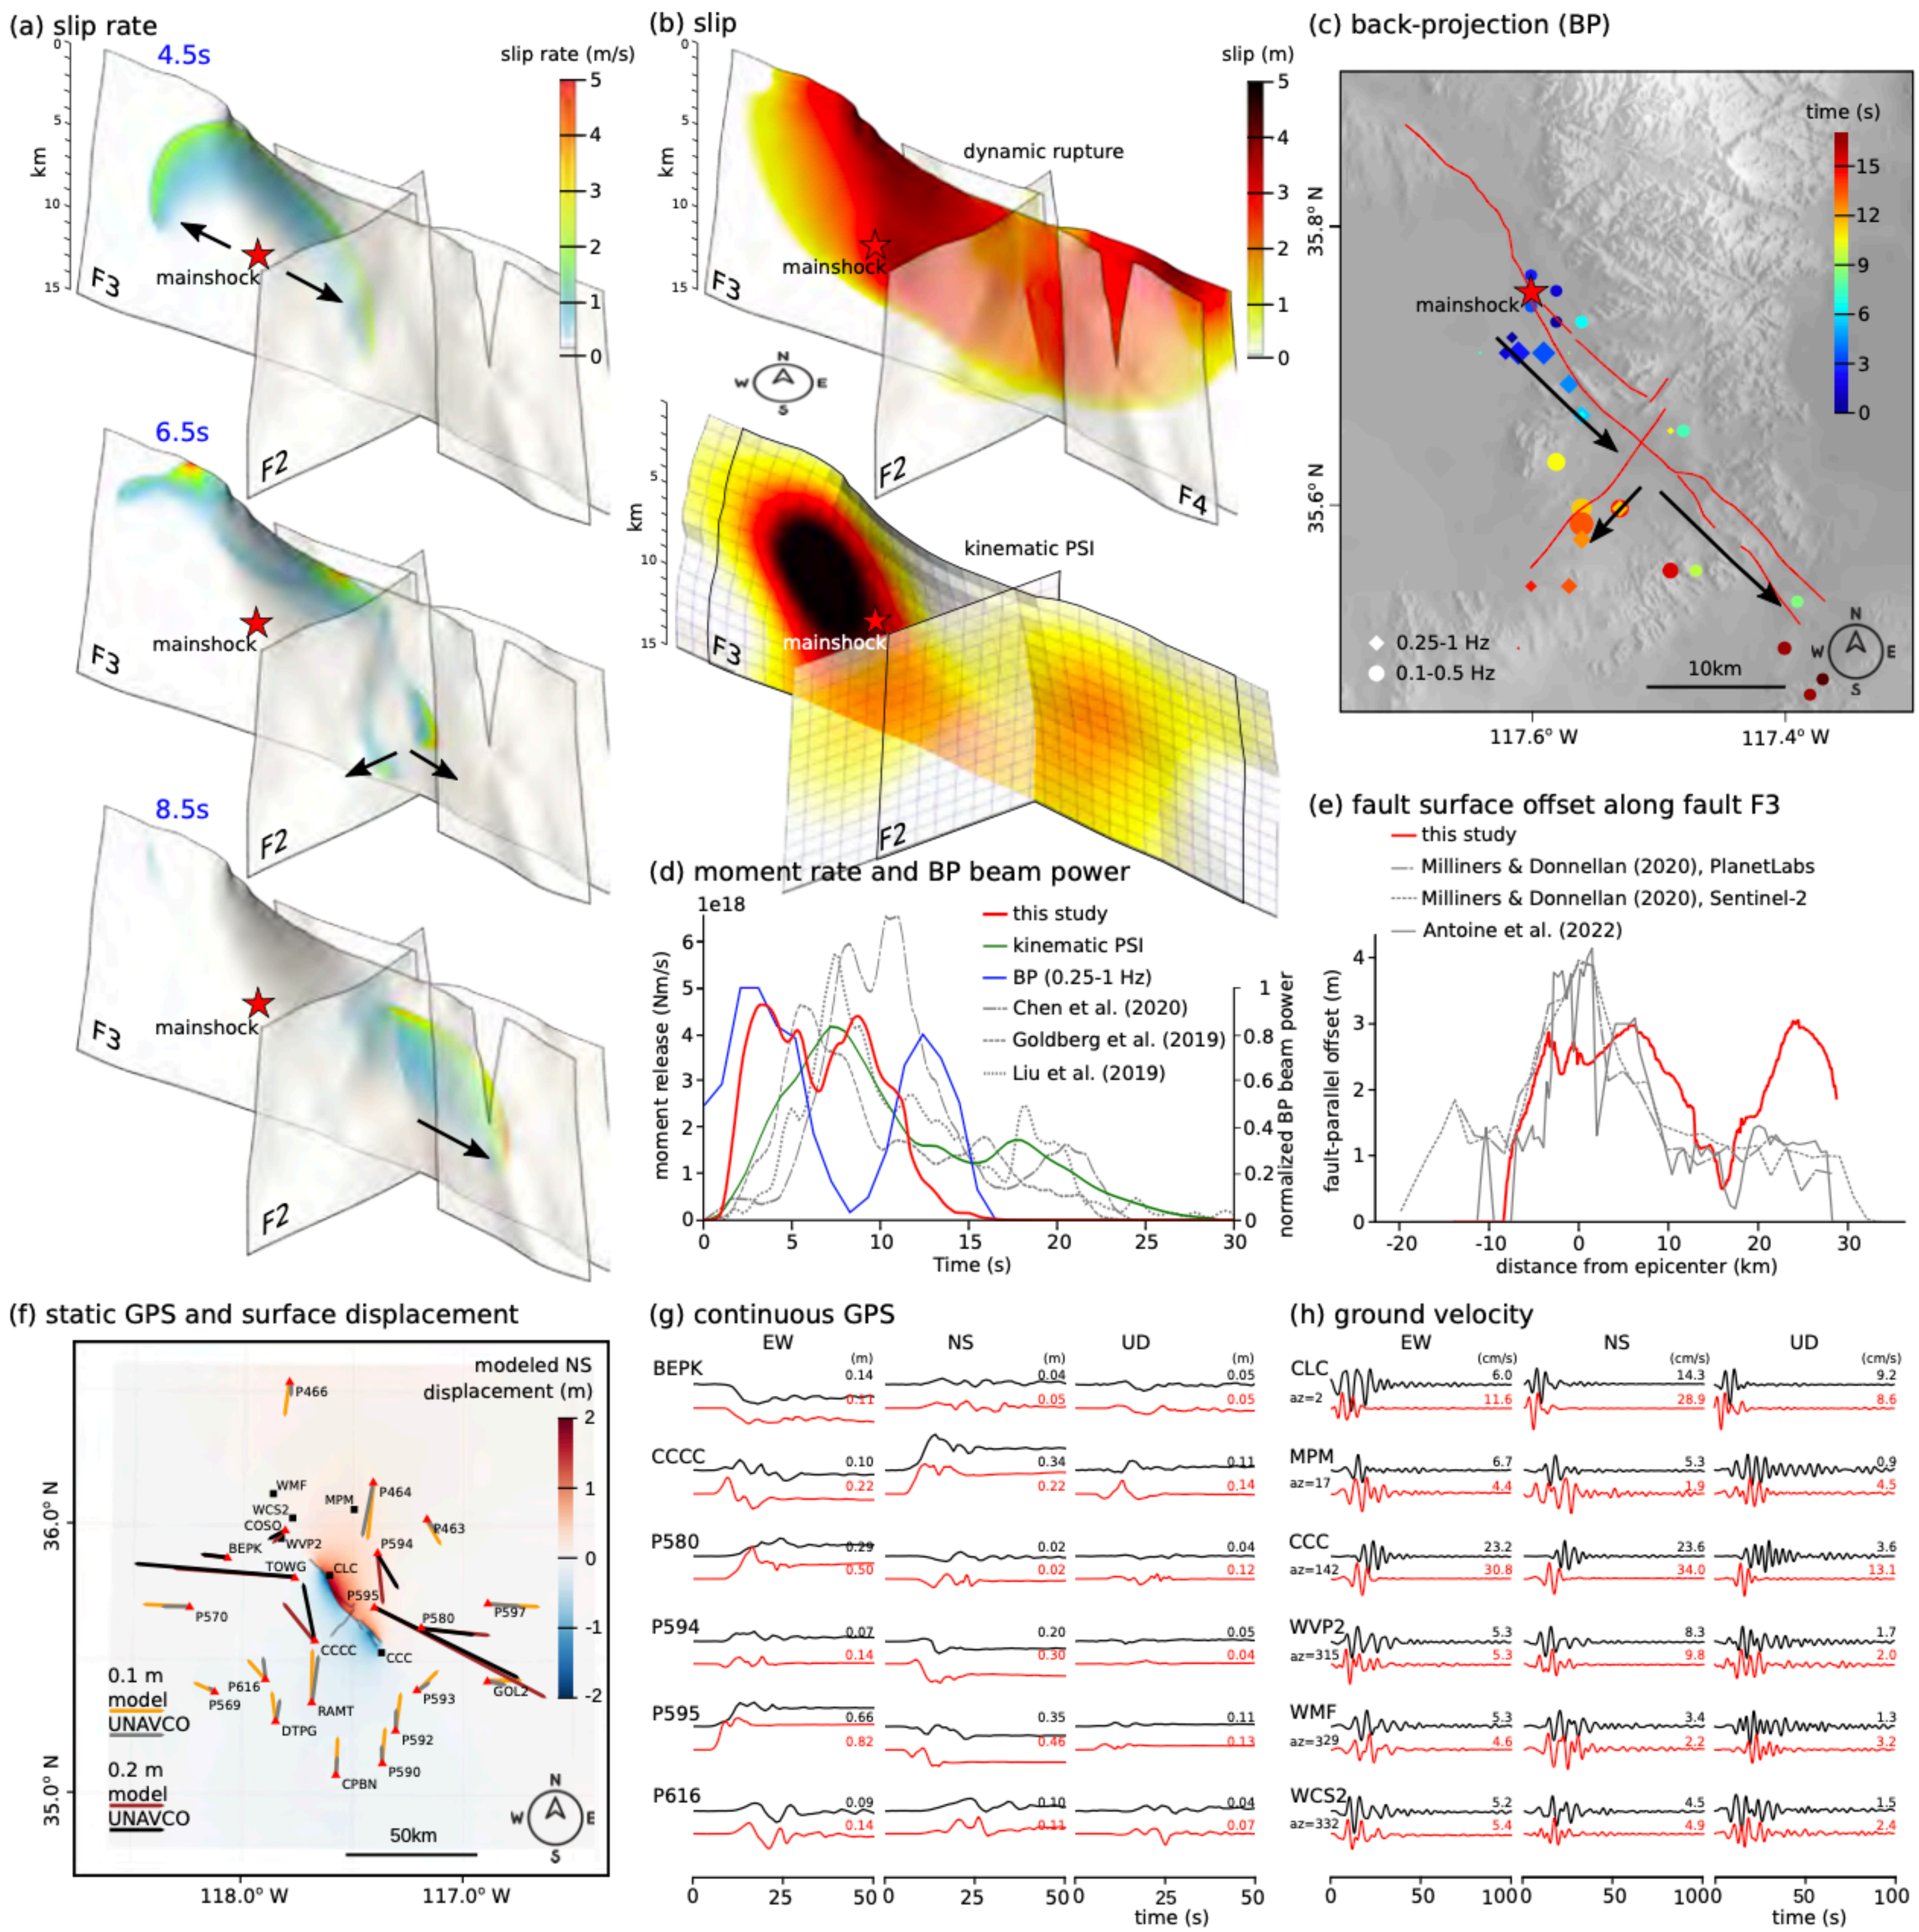

SeisSol - an open source earthquake modelling framework

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 attenuation** (+80%), **fault-zone (visco-)plasticity** (+10%), **anisotropy** (+0%)**, poroelastic** (+360%)
- **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

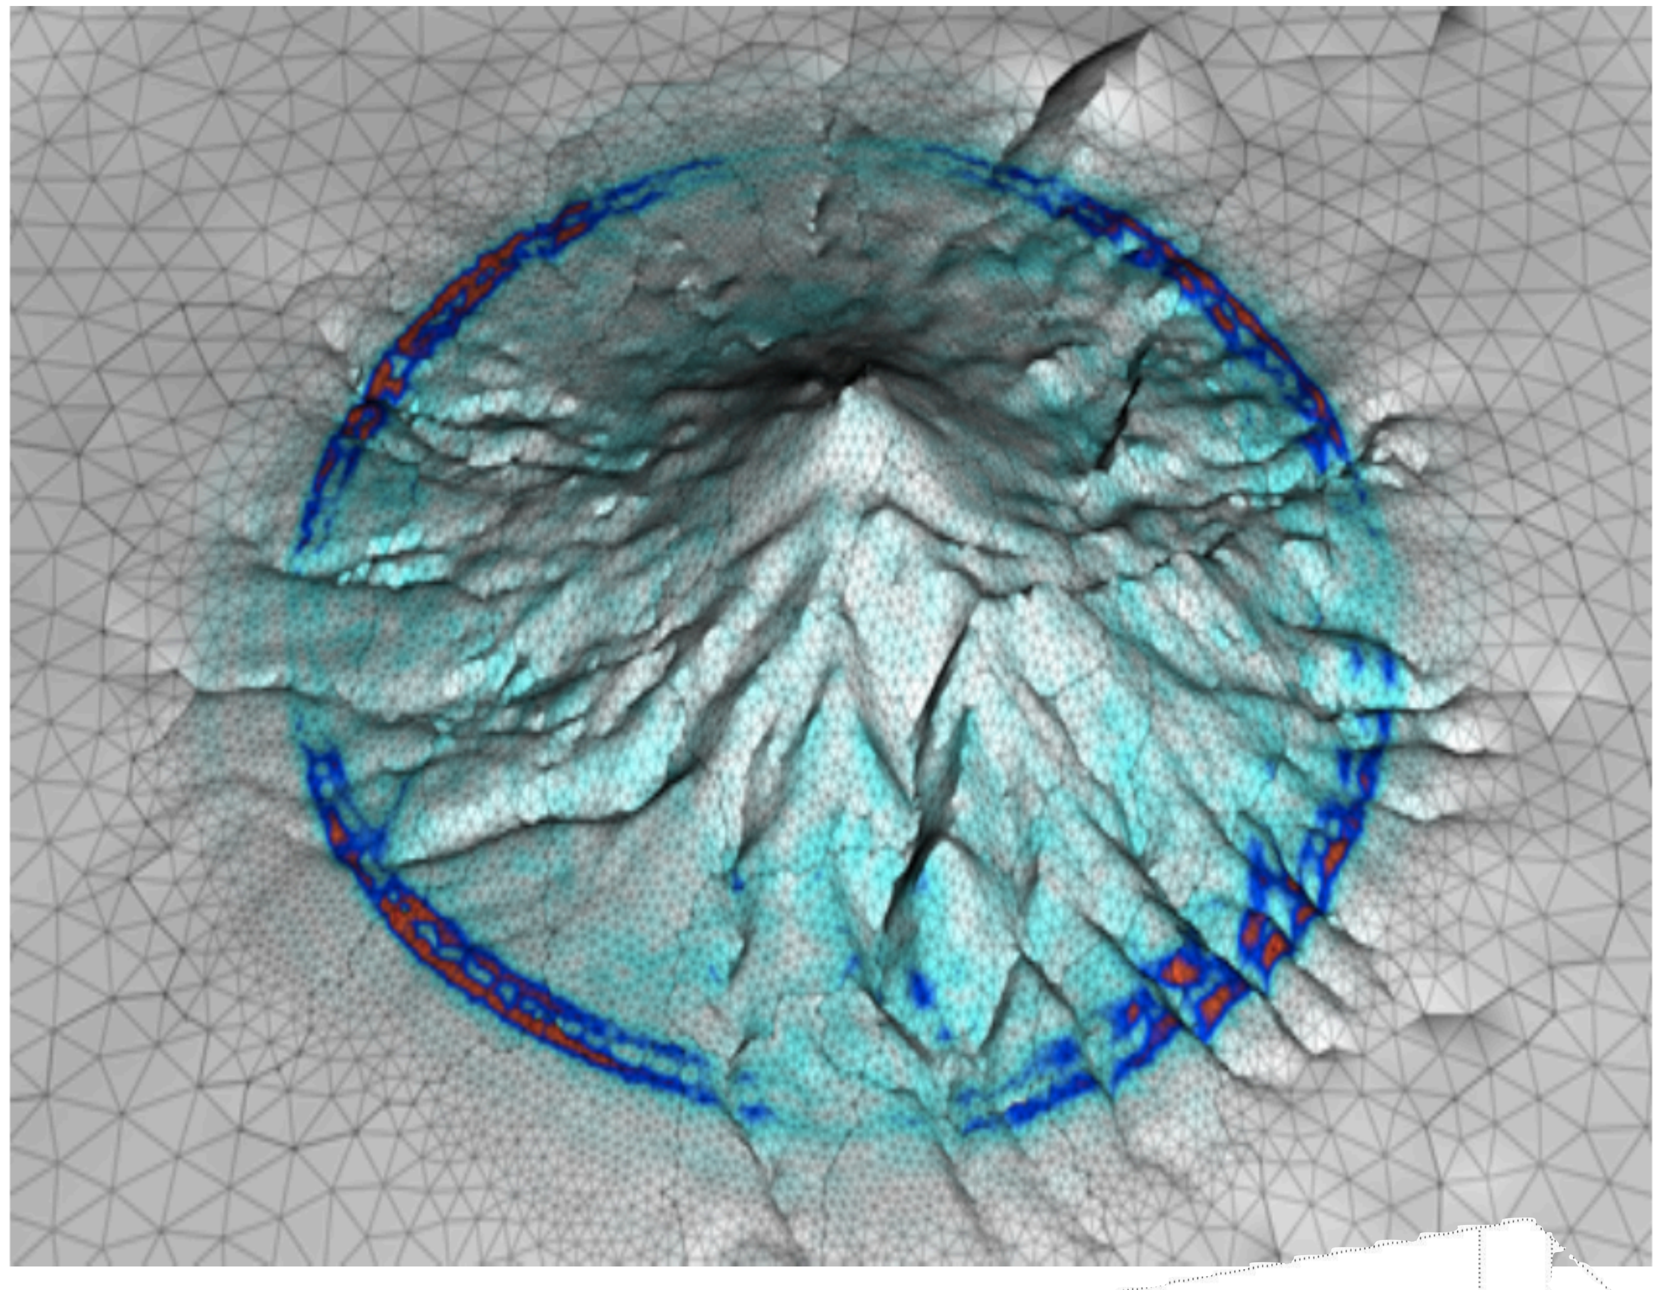

Wave field of a point source interacting with the topography of Mount Merapi Volcano.

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

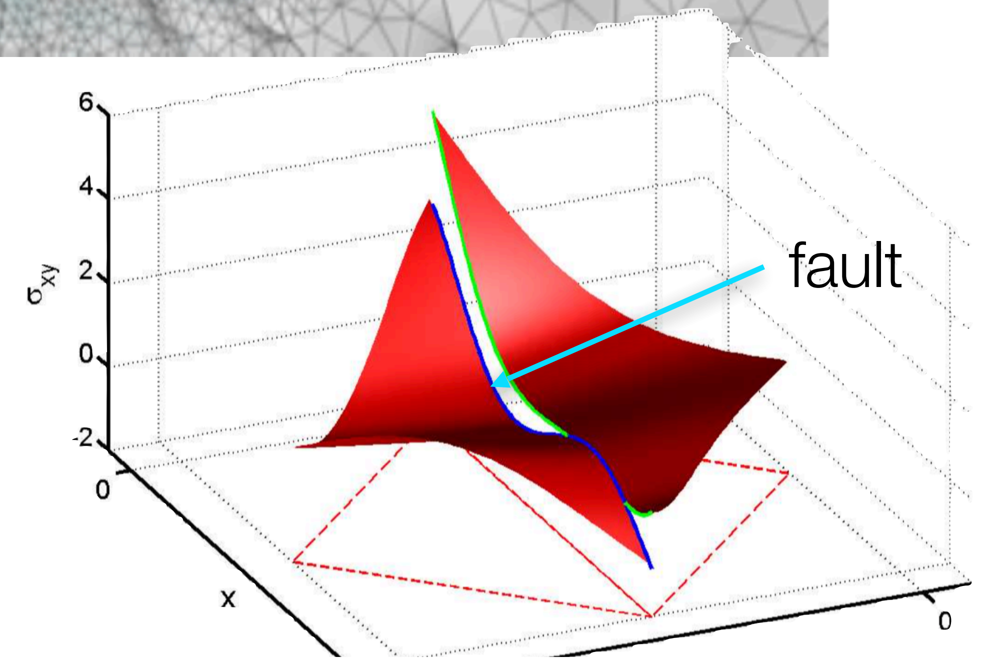

Representation of the shear stress discontinuity across the fault interface. Spontaneous rupture = internal boundary condition of flux term.

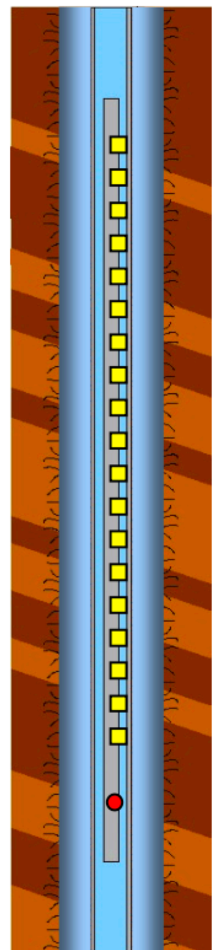

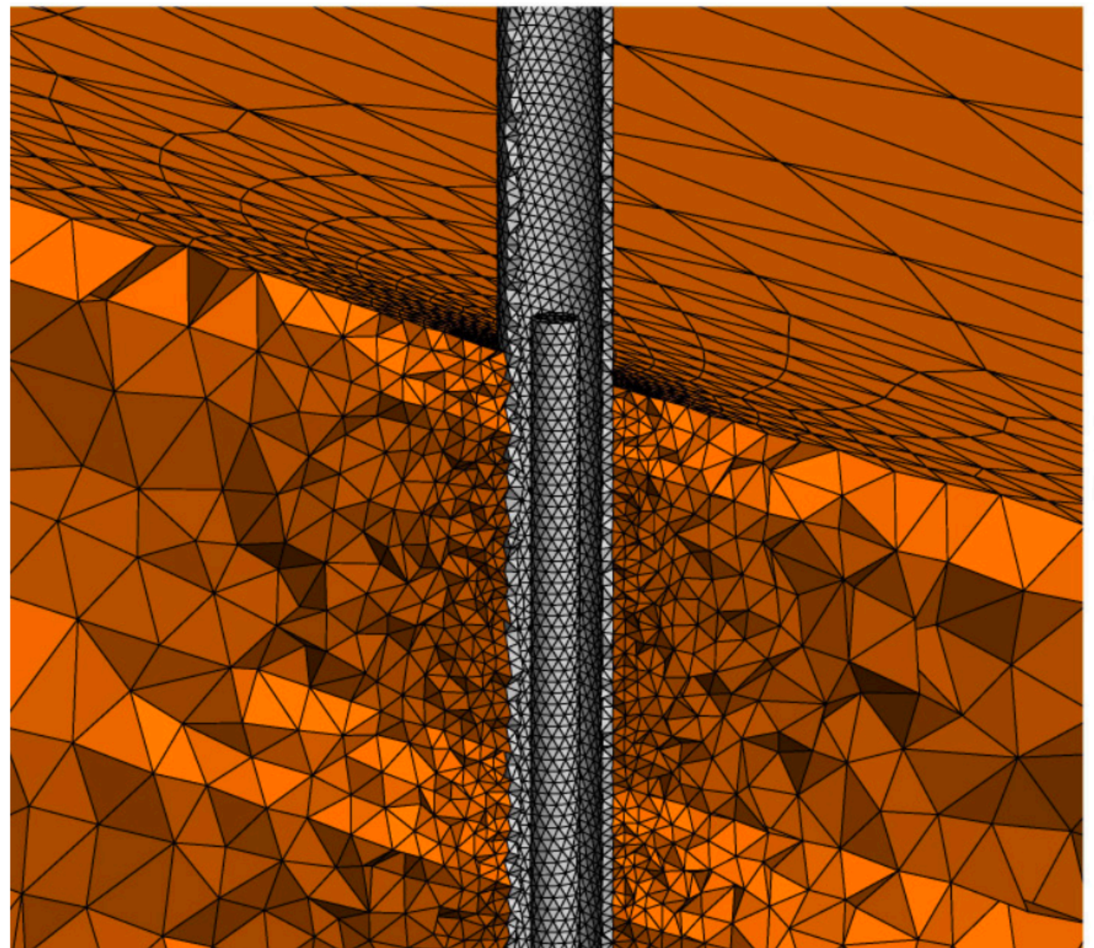

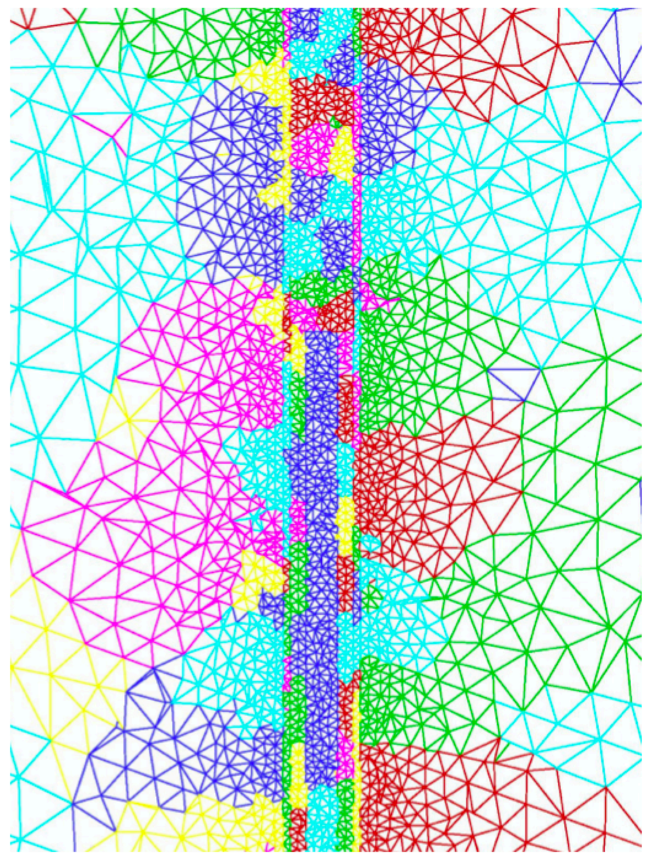

SeisSol - an open source earthquake modelling framework

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

The method, by design, permits:

- re ... ne via
- m at ar
- **m** ... r re

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.

DG's "extra" flops (storage, time to solution) can be performed fast exploiting Computational Science

- **high accuracy** - modal flux based formulation allows us to suppress spurious (unresolved) high frequencies
- **high resolution** - suitable for parallel computing environments

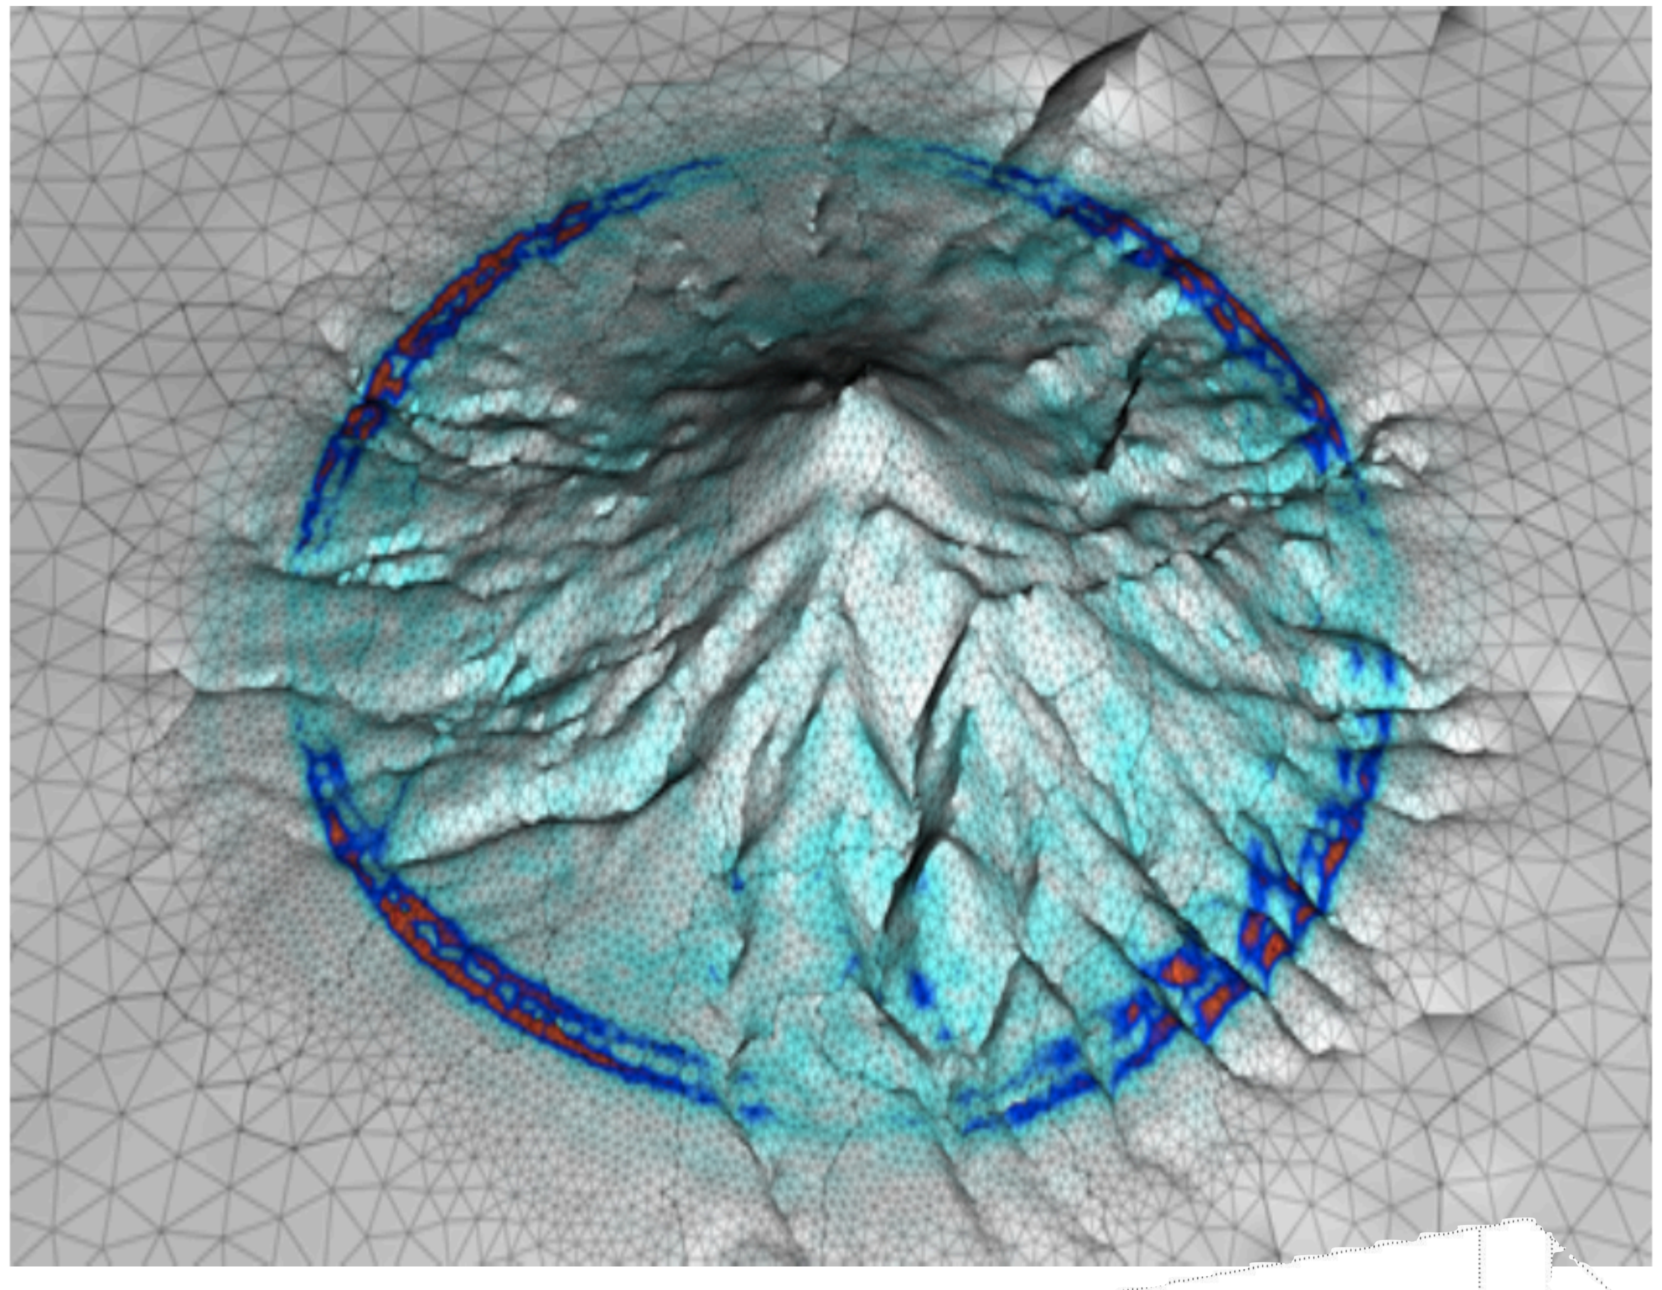

Wave field of a point source interacting with the topography of Mount Merapi Volcano.

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

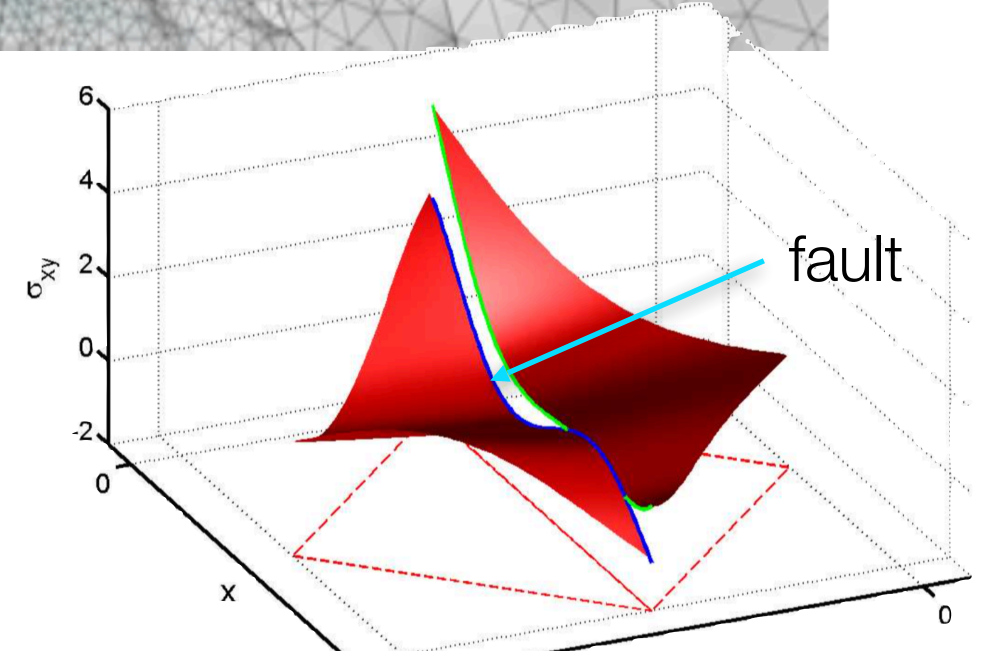

Representation of the shear stress discontinuity across the fault interface. Spontaneous rupture = internal boundary condition of flux term.

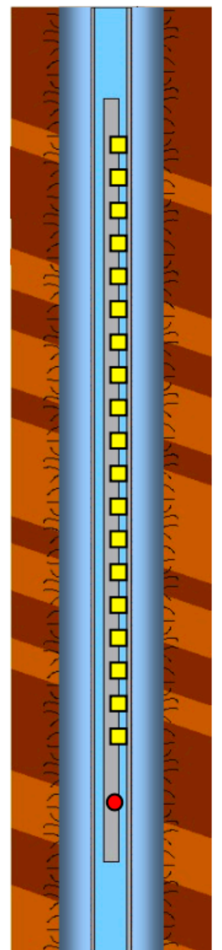

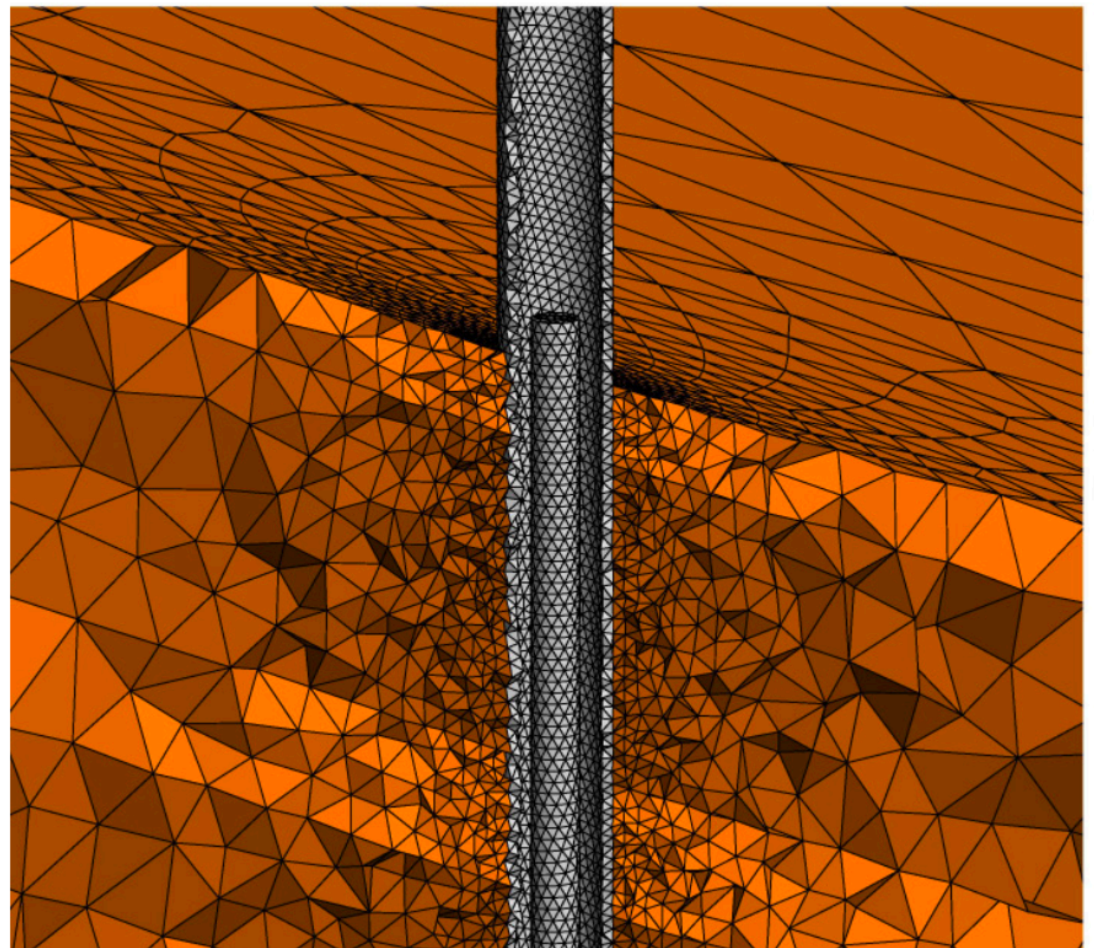

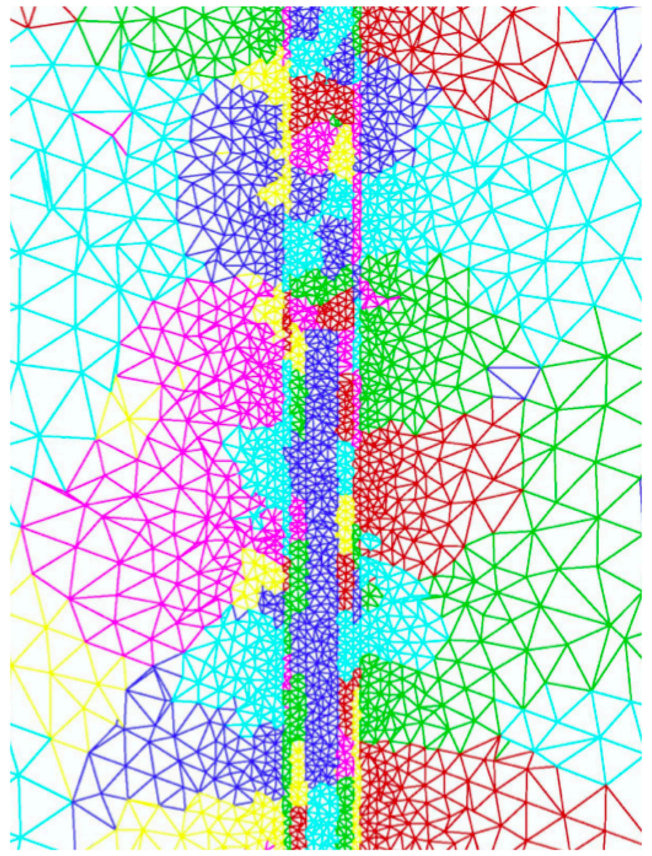

Balancing HPC and geophysics

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

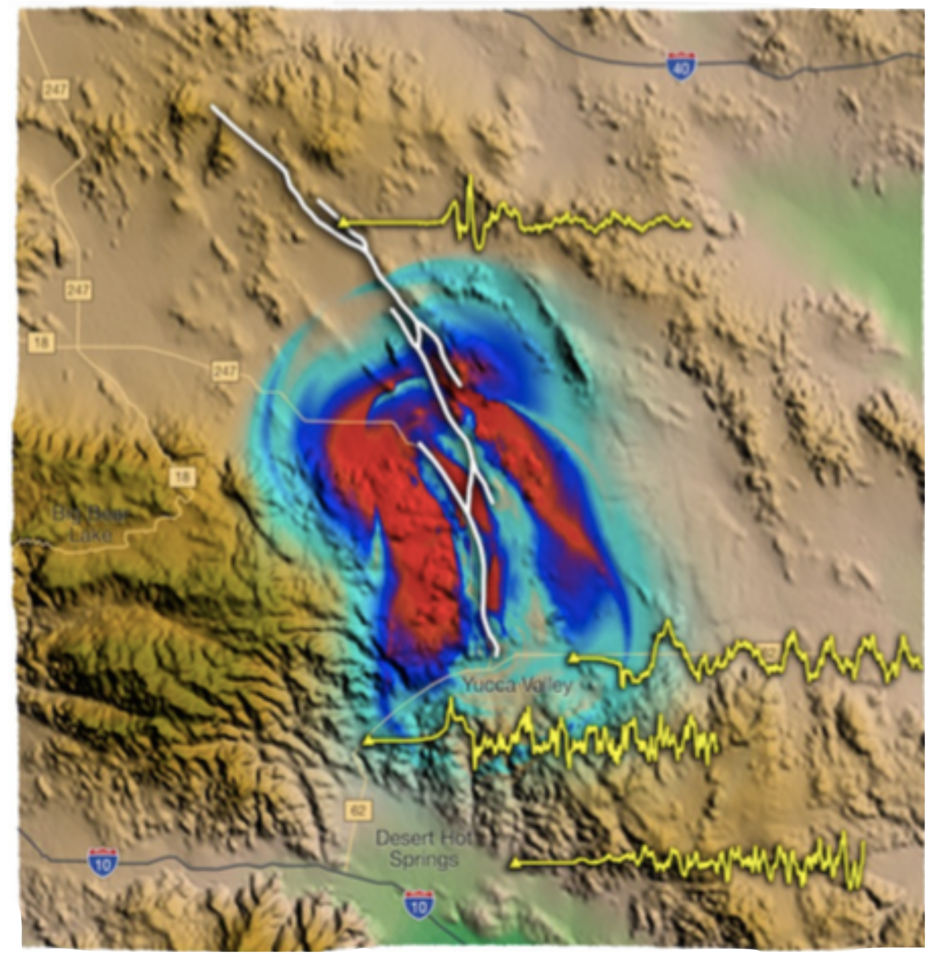

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 billion DoF, 200,000 time steps)	• Hybrid MPI+OpenMP parallelisation • Parallel I/O (HDF5, inc. mesh init.) • **Assembler-level DG kernels** • multi-physics off-load scheme for many-core architectures

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

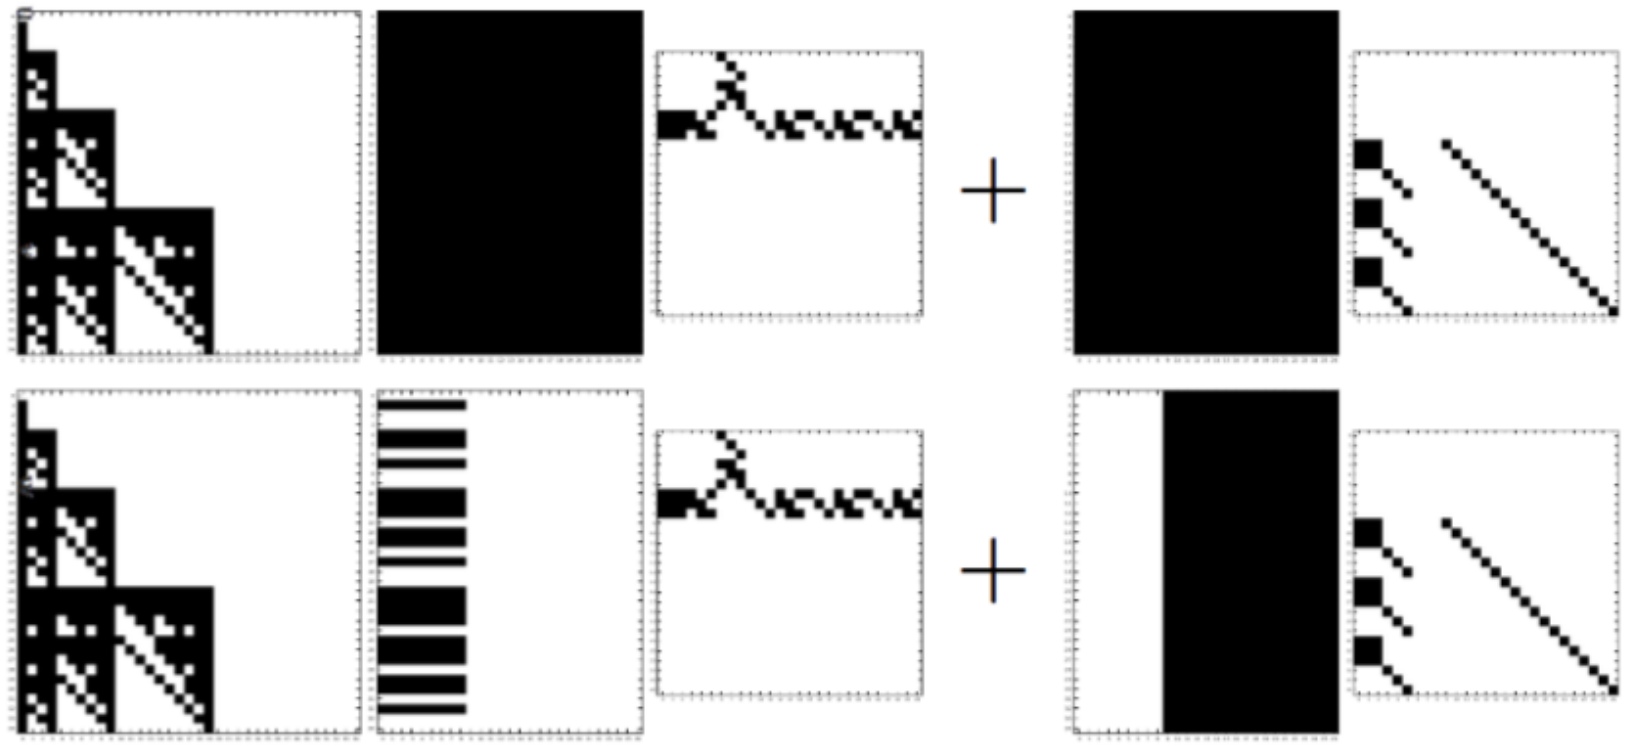

Partial kernel before (top) and after (bottom) removing irrelevant entries in matrix chain products

➡ **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**

Balancing HPC and geophysics

TECHNISCHE UNIVERSITÄT MÜNCHEN

Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften

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

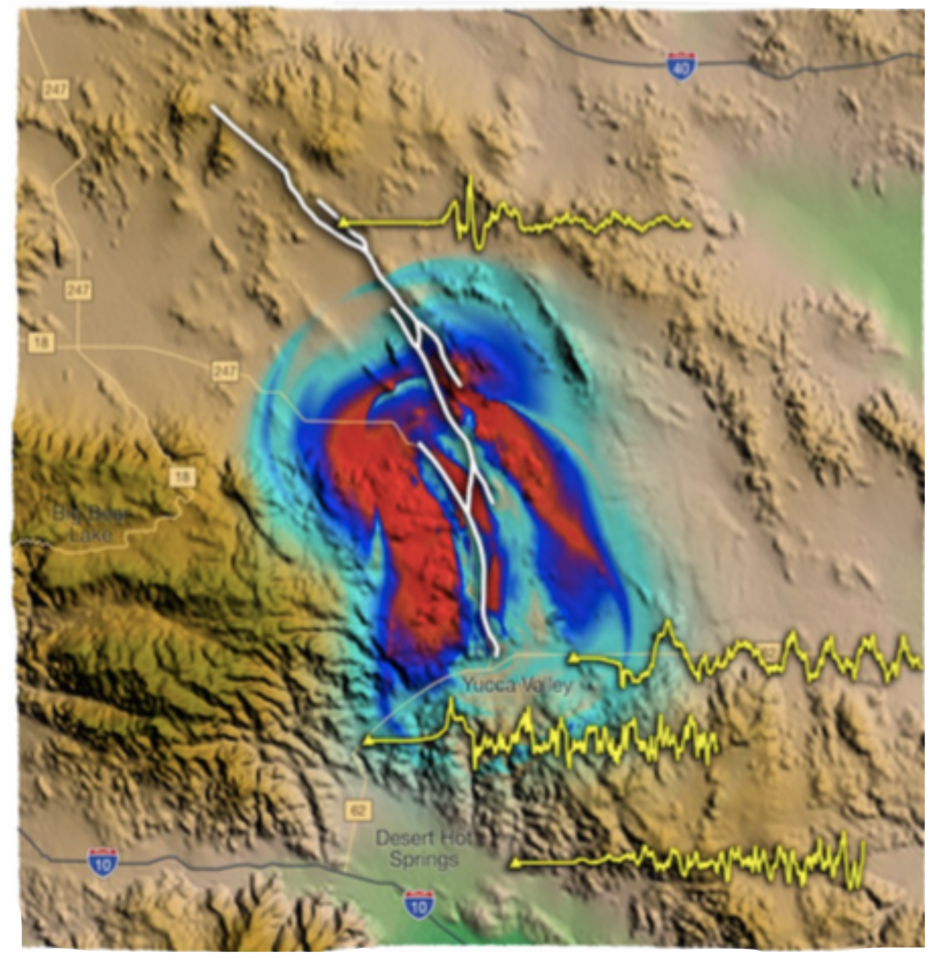

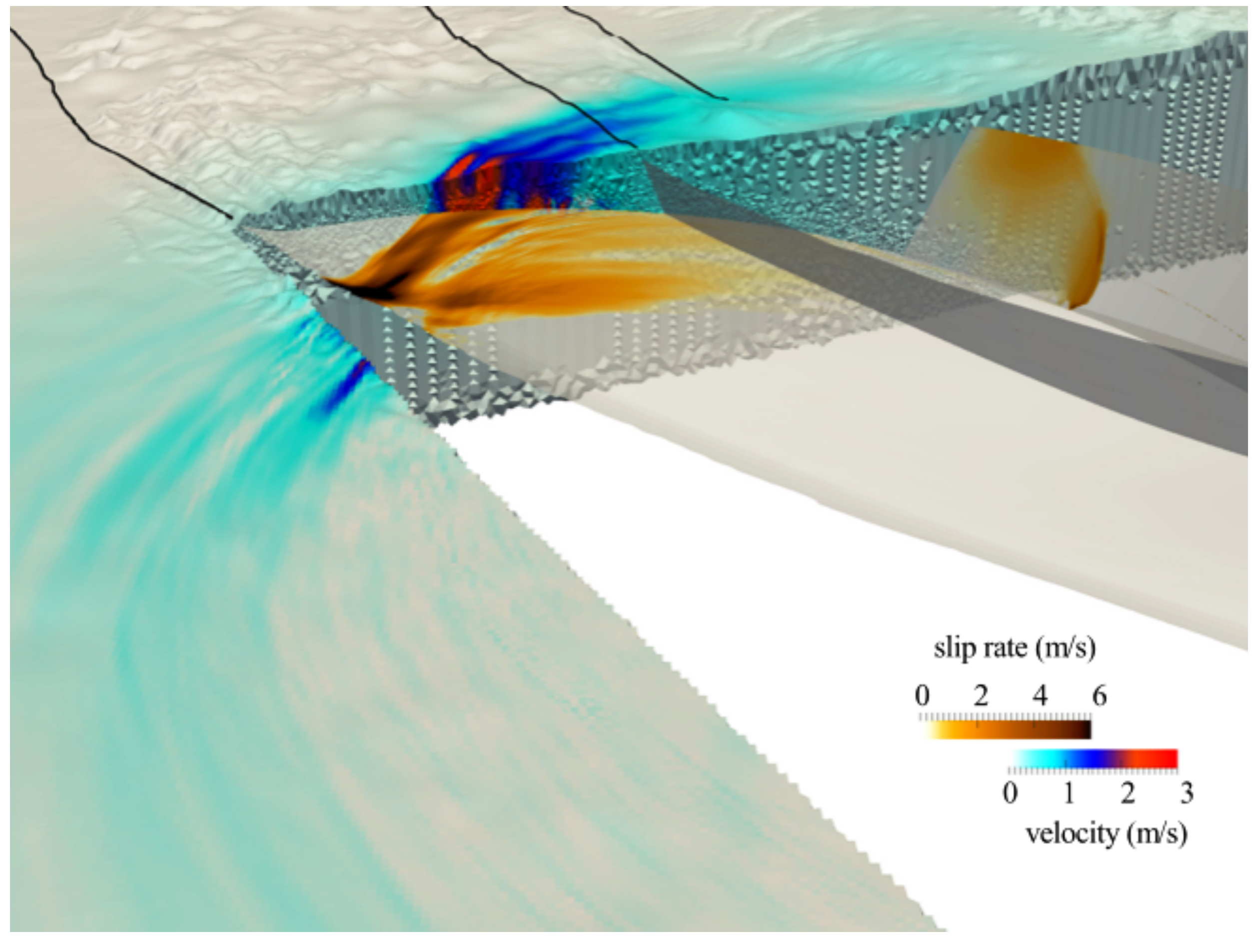

Best Paper Award, SC17

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 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
Sumatra scenario (111 billion DoF, 3,300,000 time steps)	• **Cluster-based local time stepping** • Code generator also for advanced PDE's as viscoelastic attunation • Asagi (XDMF)-geoinformation server • Asynchronous input/output • Overlaping computation and communication

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

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

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

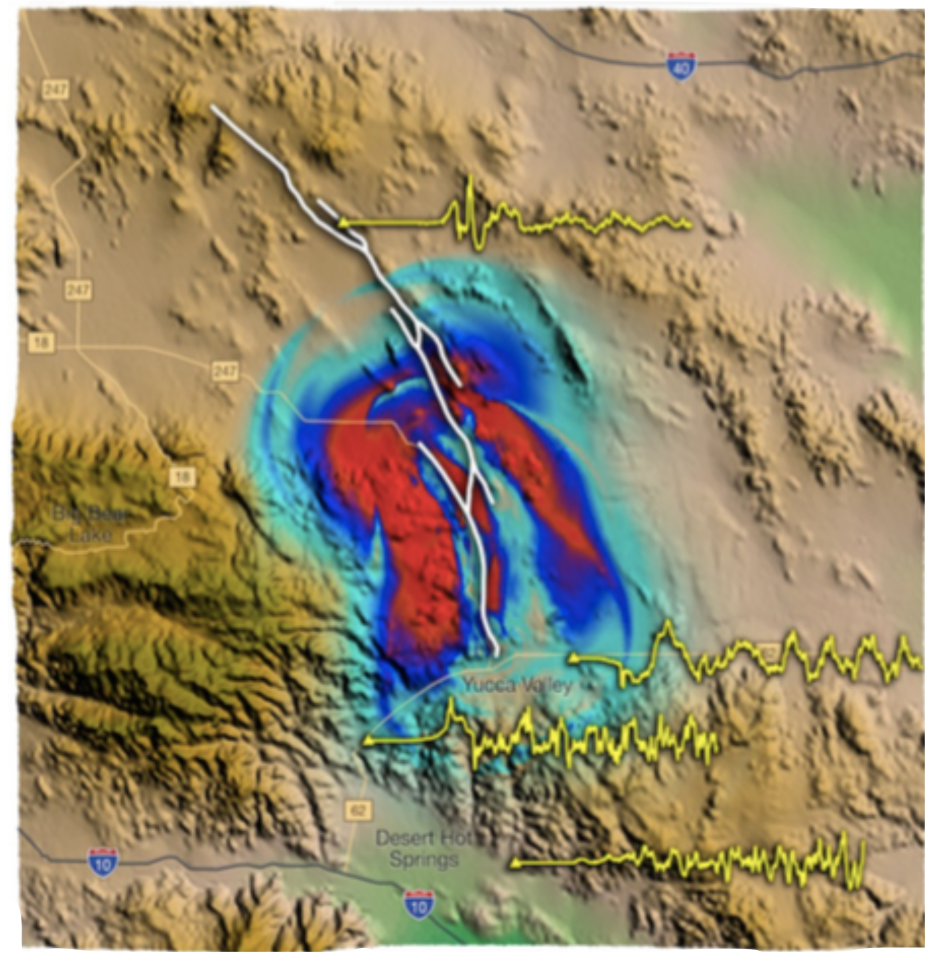

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 billion DoF, 200,000 time steps)	• Hybrid MPI+OpenMP parallelisation • Parallel I/O (HDF5, inc. mesh init.) • **Assembler-level DG kernels** • multi-physics off-load scheme for many-core architectures

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

Elastic Wave Equations: (velocity-stress formulation)

$$q_t + Aq_x + Bq_y + Cq_z = 0$$

with $q = (\sigma_{11}, \sigma_{22}, \sigma_{33}, \sigma_{12}, \sigma_{23}, \sigma_{13}, u, v, w)^T \in \mathbb{R}^9$

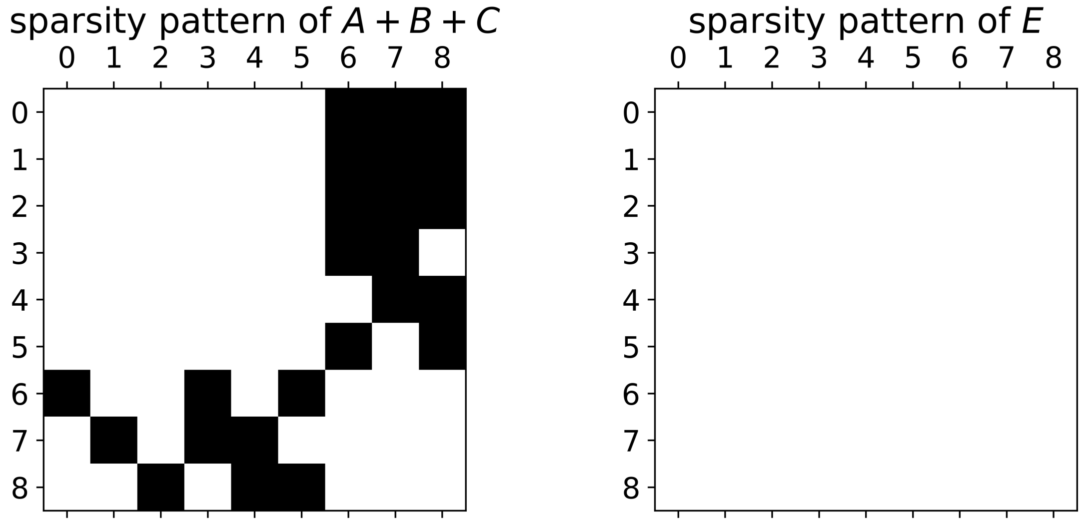

Sparse vs. Dense (Small) Matrix Operations

Determine equivalent sparsity patterns:

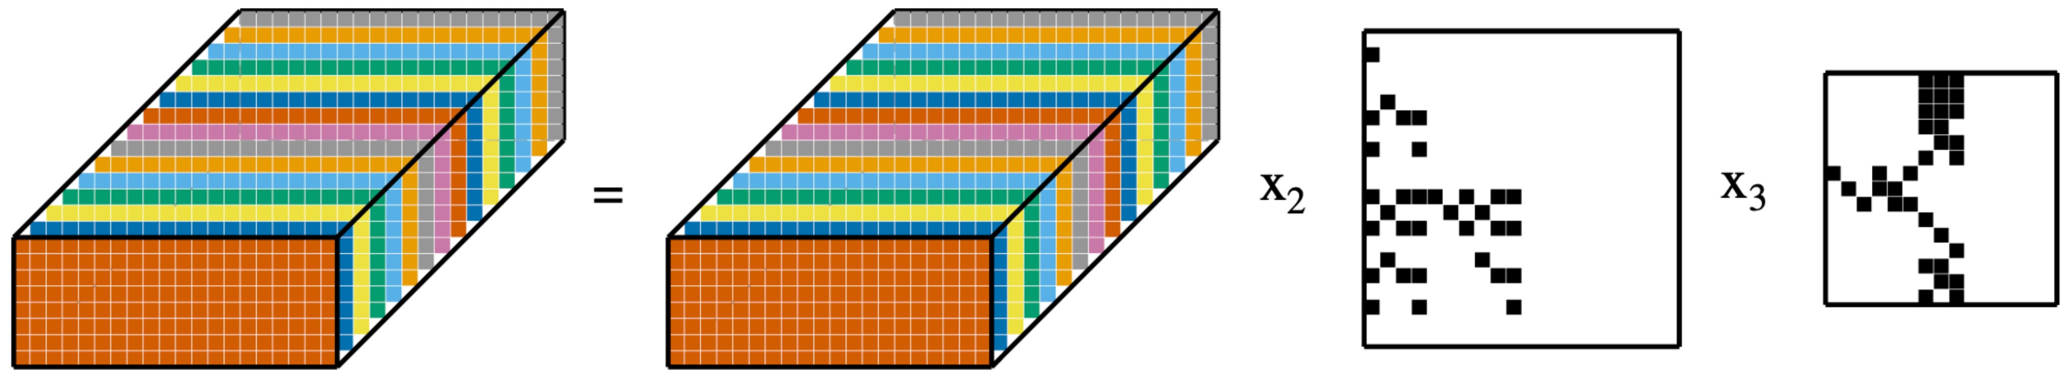

Code generation for chains of matrix/tensor multiplications:

- optimisation of chain products: find best-possible execution order (strength reduction), memory layout, index permutations, etc.

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

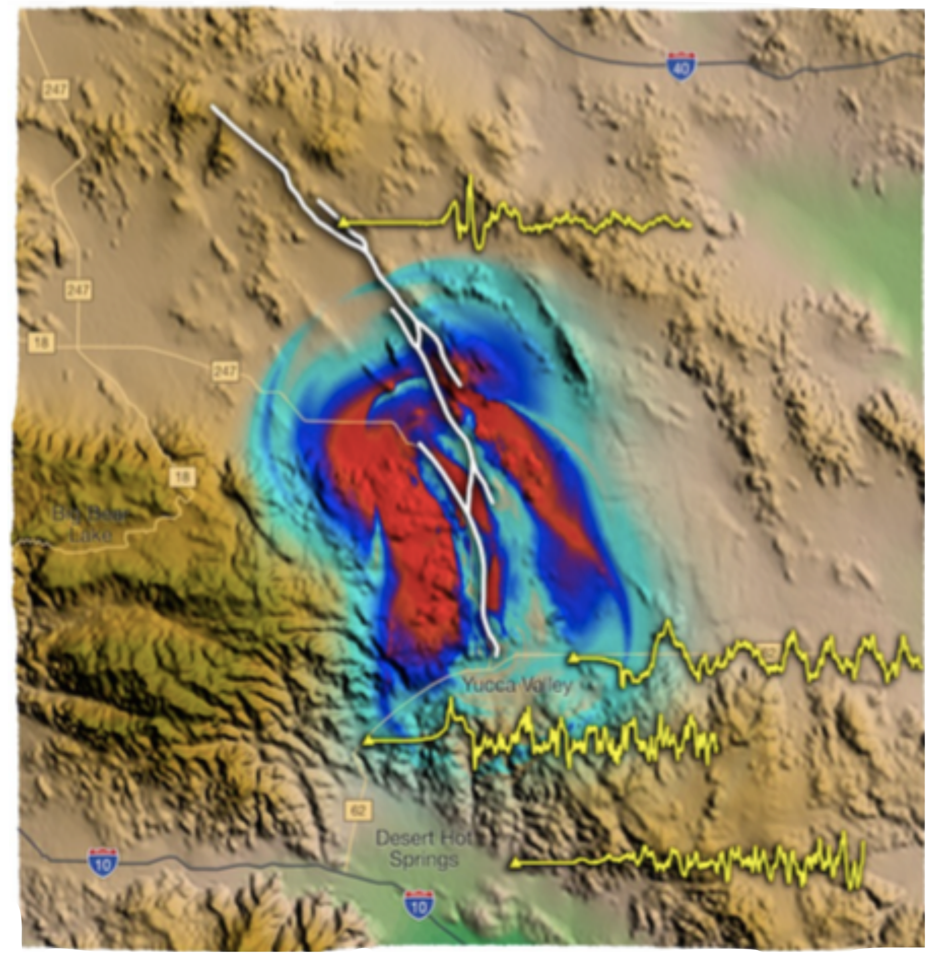

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 billion DoF, 200,000 time steps)	• Hybrid MPI+OpenMP parallelisation • Parallel I/O (HDF5, inc. mesh init.) • **Assembler-level DG kernels** • multi-physics off-load scheme for many-core architectures

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

Elastic Wave Equations: (velocity-stress formulation)

$$q_t + Aq_x + Bq_y + Cq_z = 0$$

with $q = (\sigma_{11}, \sigma_{22}, \sigma_{33}, \sigma_{12}, \sigma_{23}, \sigma_{13}, u, v, w)^T \in \mathbb{R}^9$

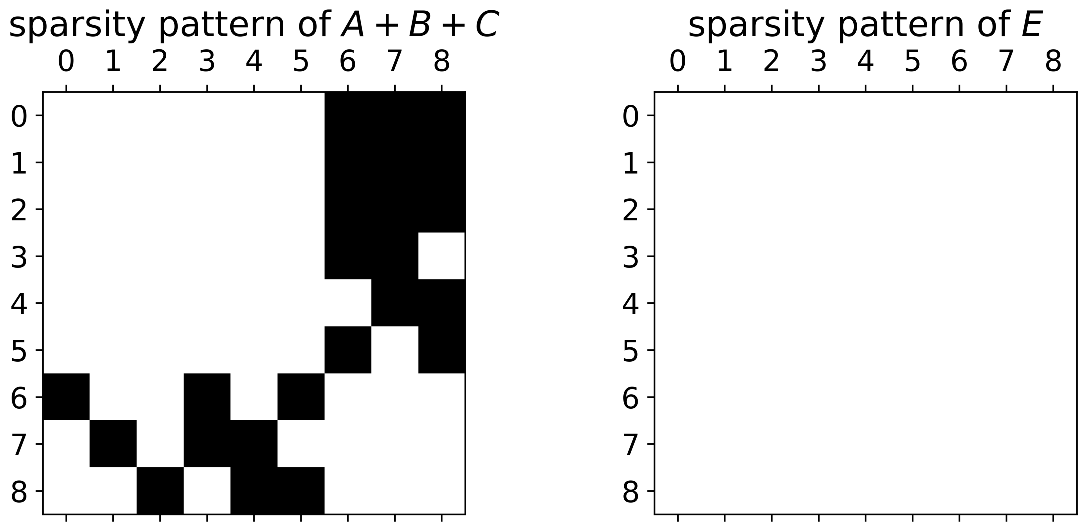

Sparse vs. Dense (Small) Matrix Operations

Determine equivalent sparsity patterns:

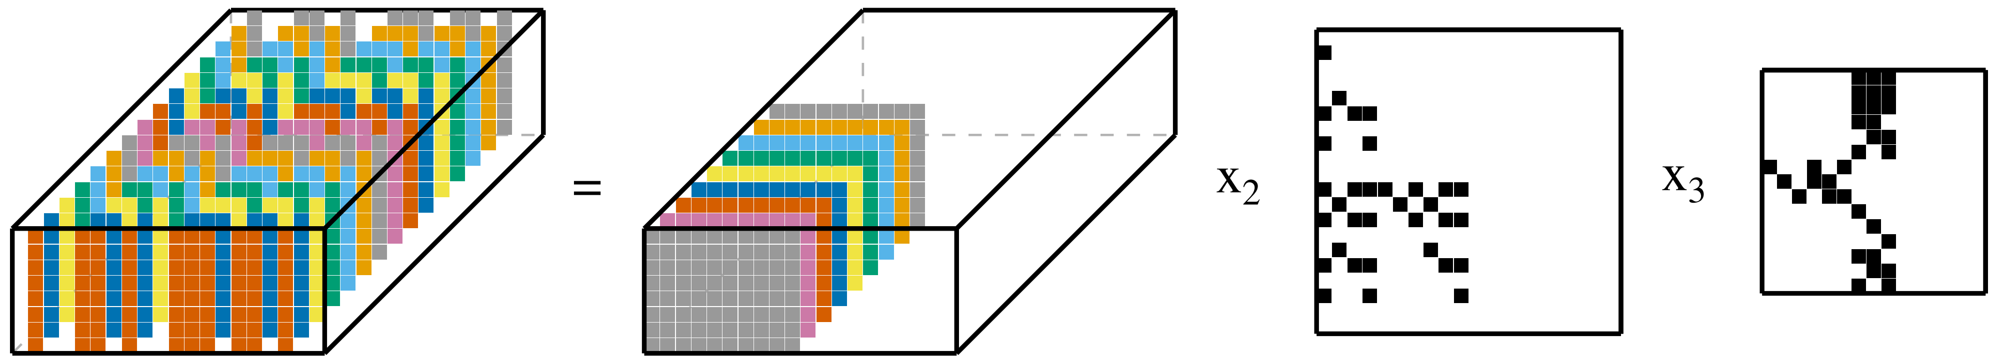

Code generation for chains of matrix/tensor multiplications:

- optimisation of chain products: find best-possible execution order (strength reduction), memory layout, index permutations, etc.

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

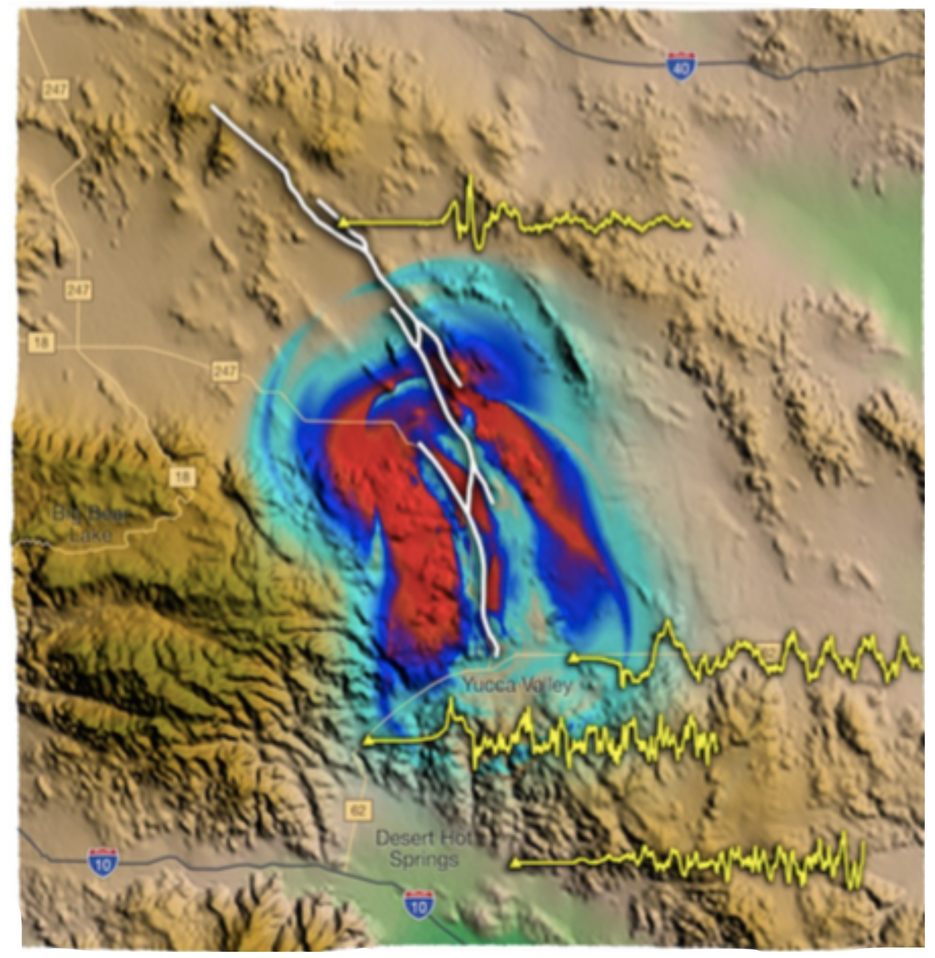

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 billion DoF, 200,000 time steps)	• Hybrid MPI+OpenMP parallelisation • Parallel I/O (HDF5, inc. mesh init.) • **Assembler-level DG kernels** • multi-physics off-load scheme for many-core architectures

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

Visco-Elastic Wave Equations:

$$q_t + Aq_x + Bq_y + Cq_z = Eq$$

with $q = (\sigma_{11}, \sigma_{22}, \sigma_{33}, \sigma_{12}, \sigma_{23}, \sigma_{13}, u, v, w, \zeta_{11}^1, \ldots, \zeta_3^3)^T \in \mathbb{R}^{9+6L}$

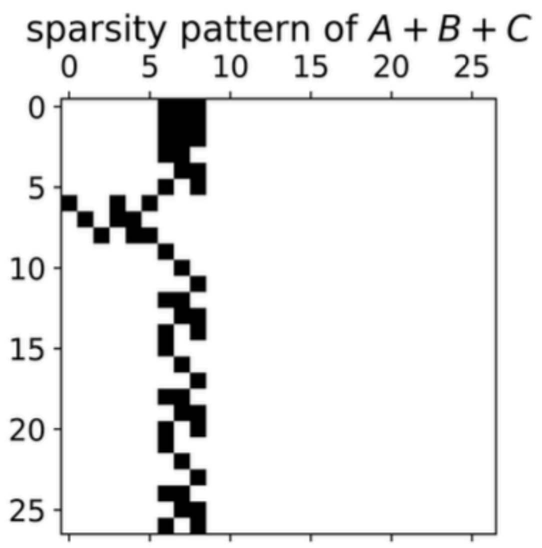

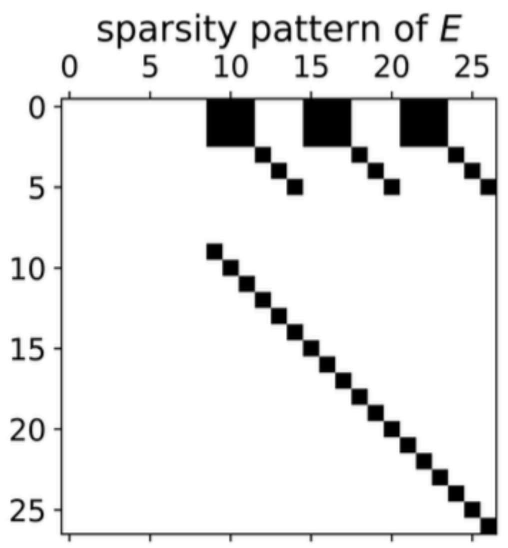

Sparse vs. Dense (Small) Matrix Operations

Determine equivalent sparsity patterns:

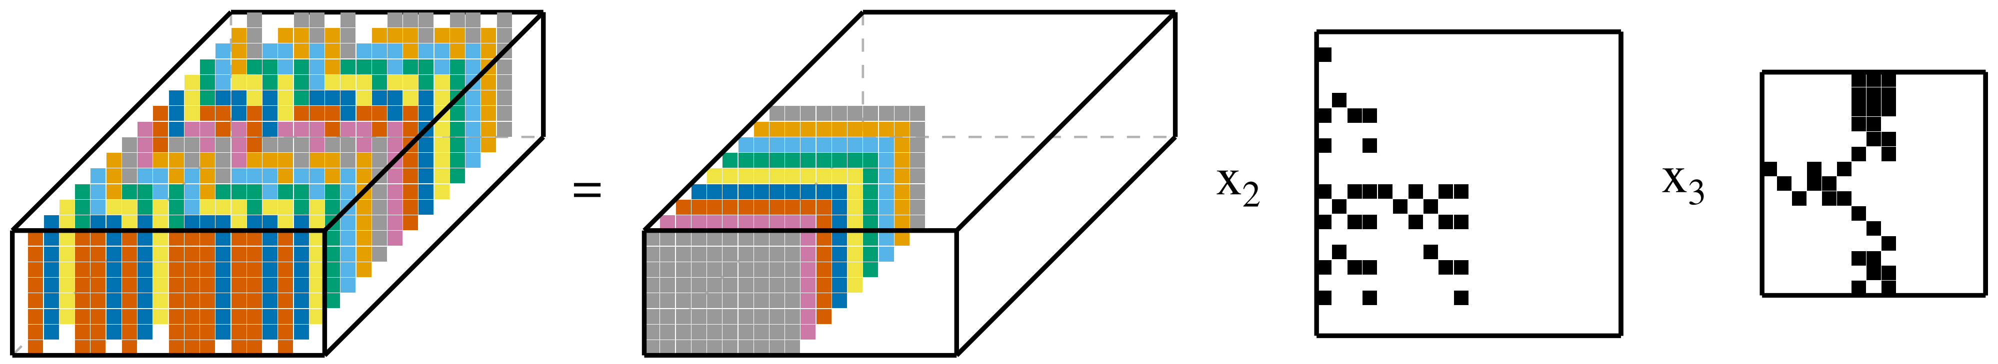

Code generation for chains of matrix/tensor multiplications:

- optimisation of chain products: find best-possible execution order (strength reduction), memory layout, index permutations, etc.

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

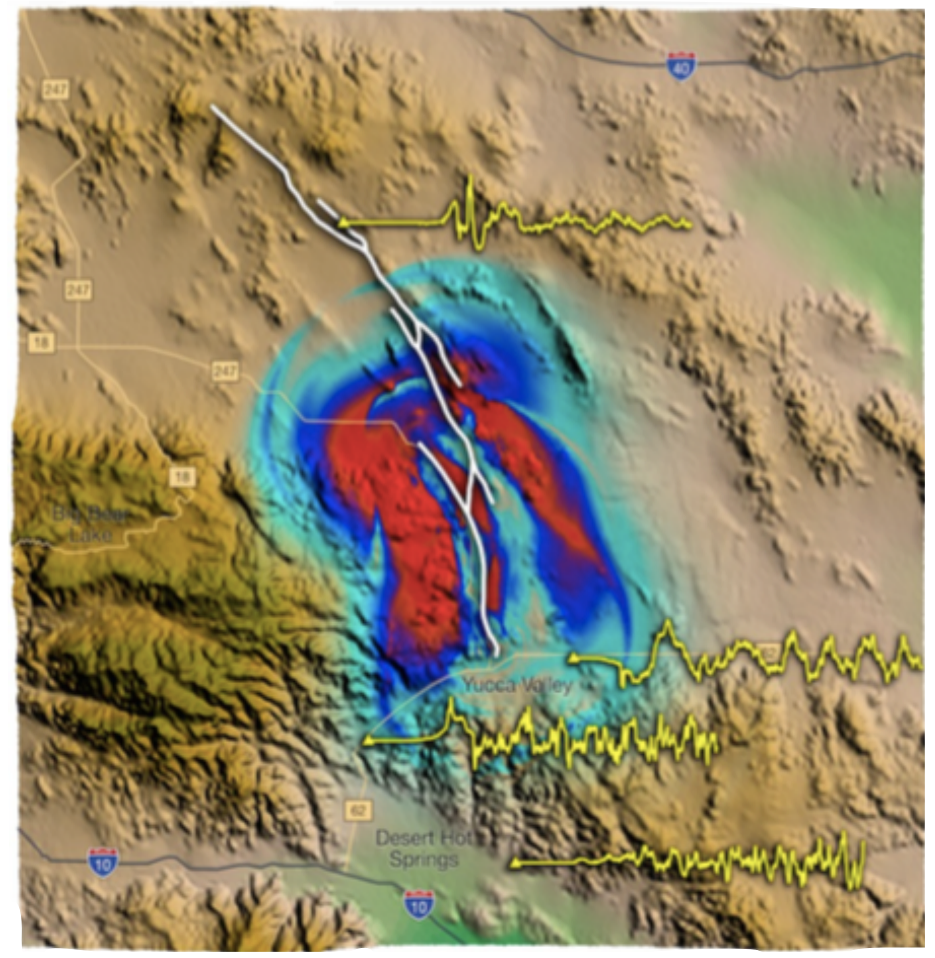

"Geophysics" Version	• Fortran 90 • MPI parallelised • Ascii based, serial I/O
Landers scenario (96 billion DoF, 200,000 time steps)	• Hybrid MPI+OpenMP parallelisation • Parallel I/O (HDF5, inc. mesh init.) • **Assembler-level DG kernels** • multi-physics off-load scheme for many-core architectures

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

Poro-Elastic Wave Equations

$$q_t + Aq_x + Bq_y + Cq_z = Eq$$

with $q = (\sigma_{11}, \sigma_{22}, \sigma_{33}, \sigma_{12}, \sigma_{23}, \sigma_{13}, u, v, w, p, u_f, v_f, w_f)^T \in \mathbb{R}^{13}$

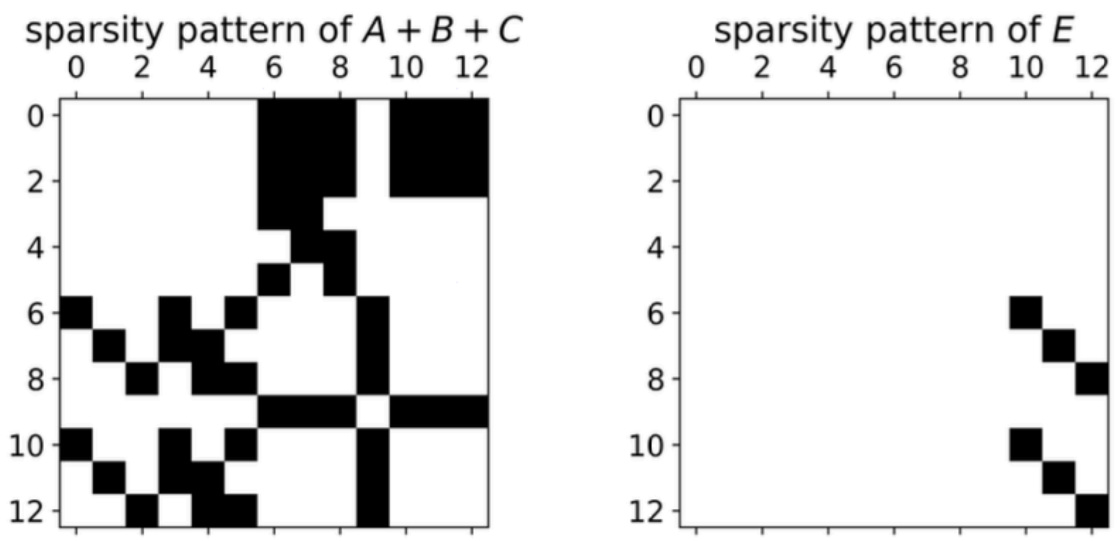

Sparse vs. Dense (Small) Matrix Operations

Determine equivalent sparsity patterns:

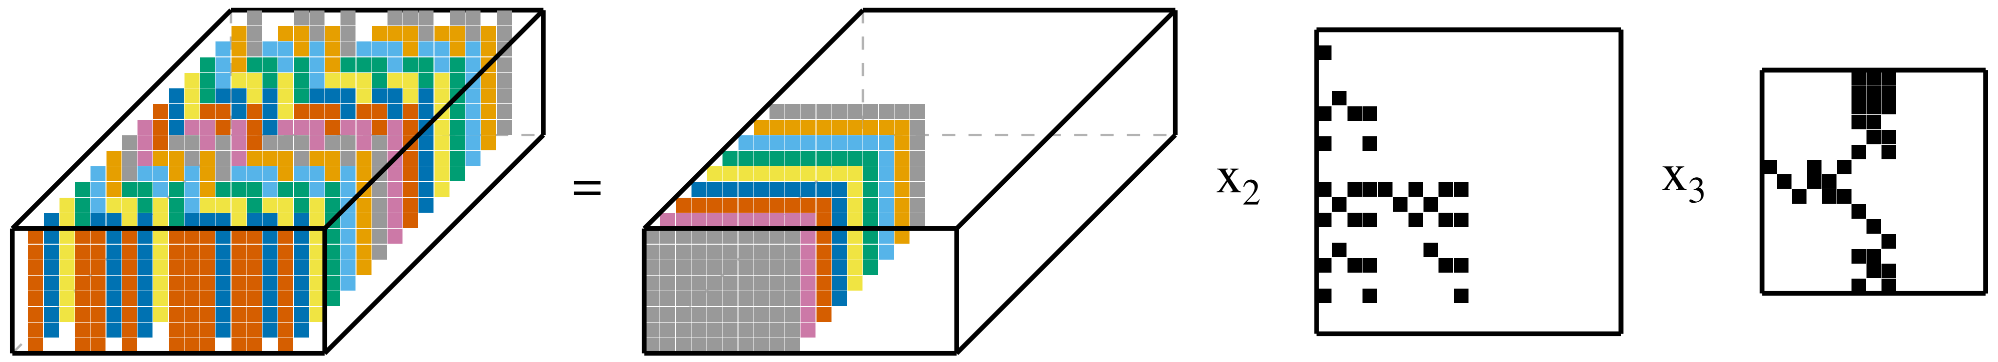

Code generation for chains of matrix/tensor multiplications:

- optimisation of chain products: find best-possible execution order (strength reduction), memory layout, index permutations, etc.

What is the cost?

- **Higher resolution models** (e.g., for broadband ground motion modeling, fault zone complexity (15 Hz require >5 h on 7000 nodes of Frontera, 20 Hz will require a mesh of 1.2 billion elements)

Computing (5 Hz Rigecrest sequence model)

- 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 CO_2 (~flying from London to Los Angeles)

Strong motivation to optimise the efficiency of seismological software!

*numbers for Xeon E5 cores
e.g. Shaheen-2 uses 30% less energy = flying from MUC -> ORD

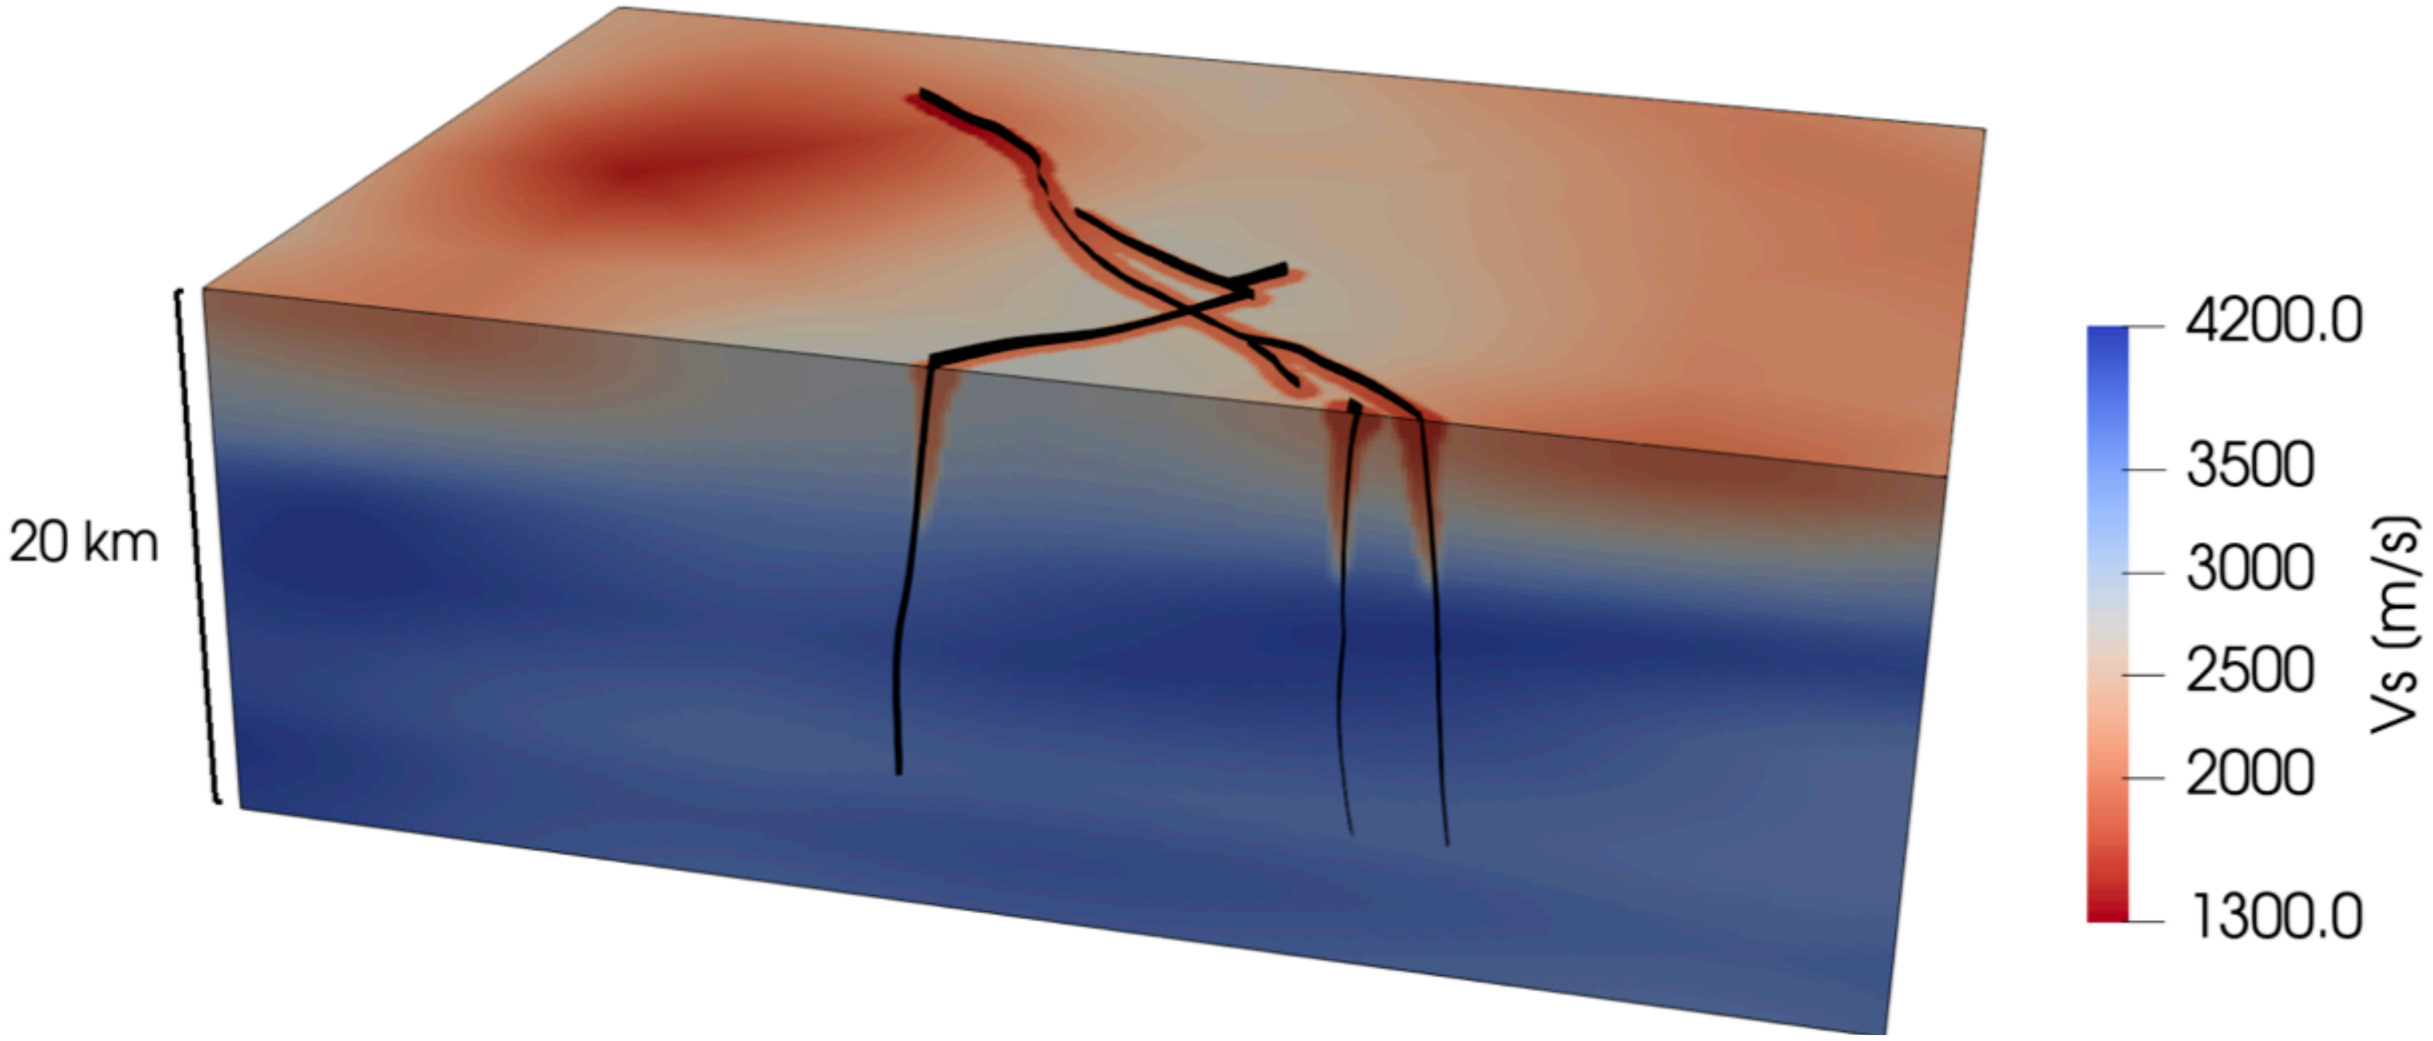

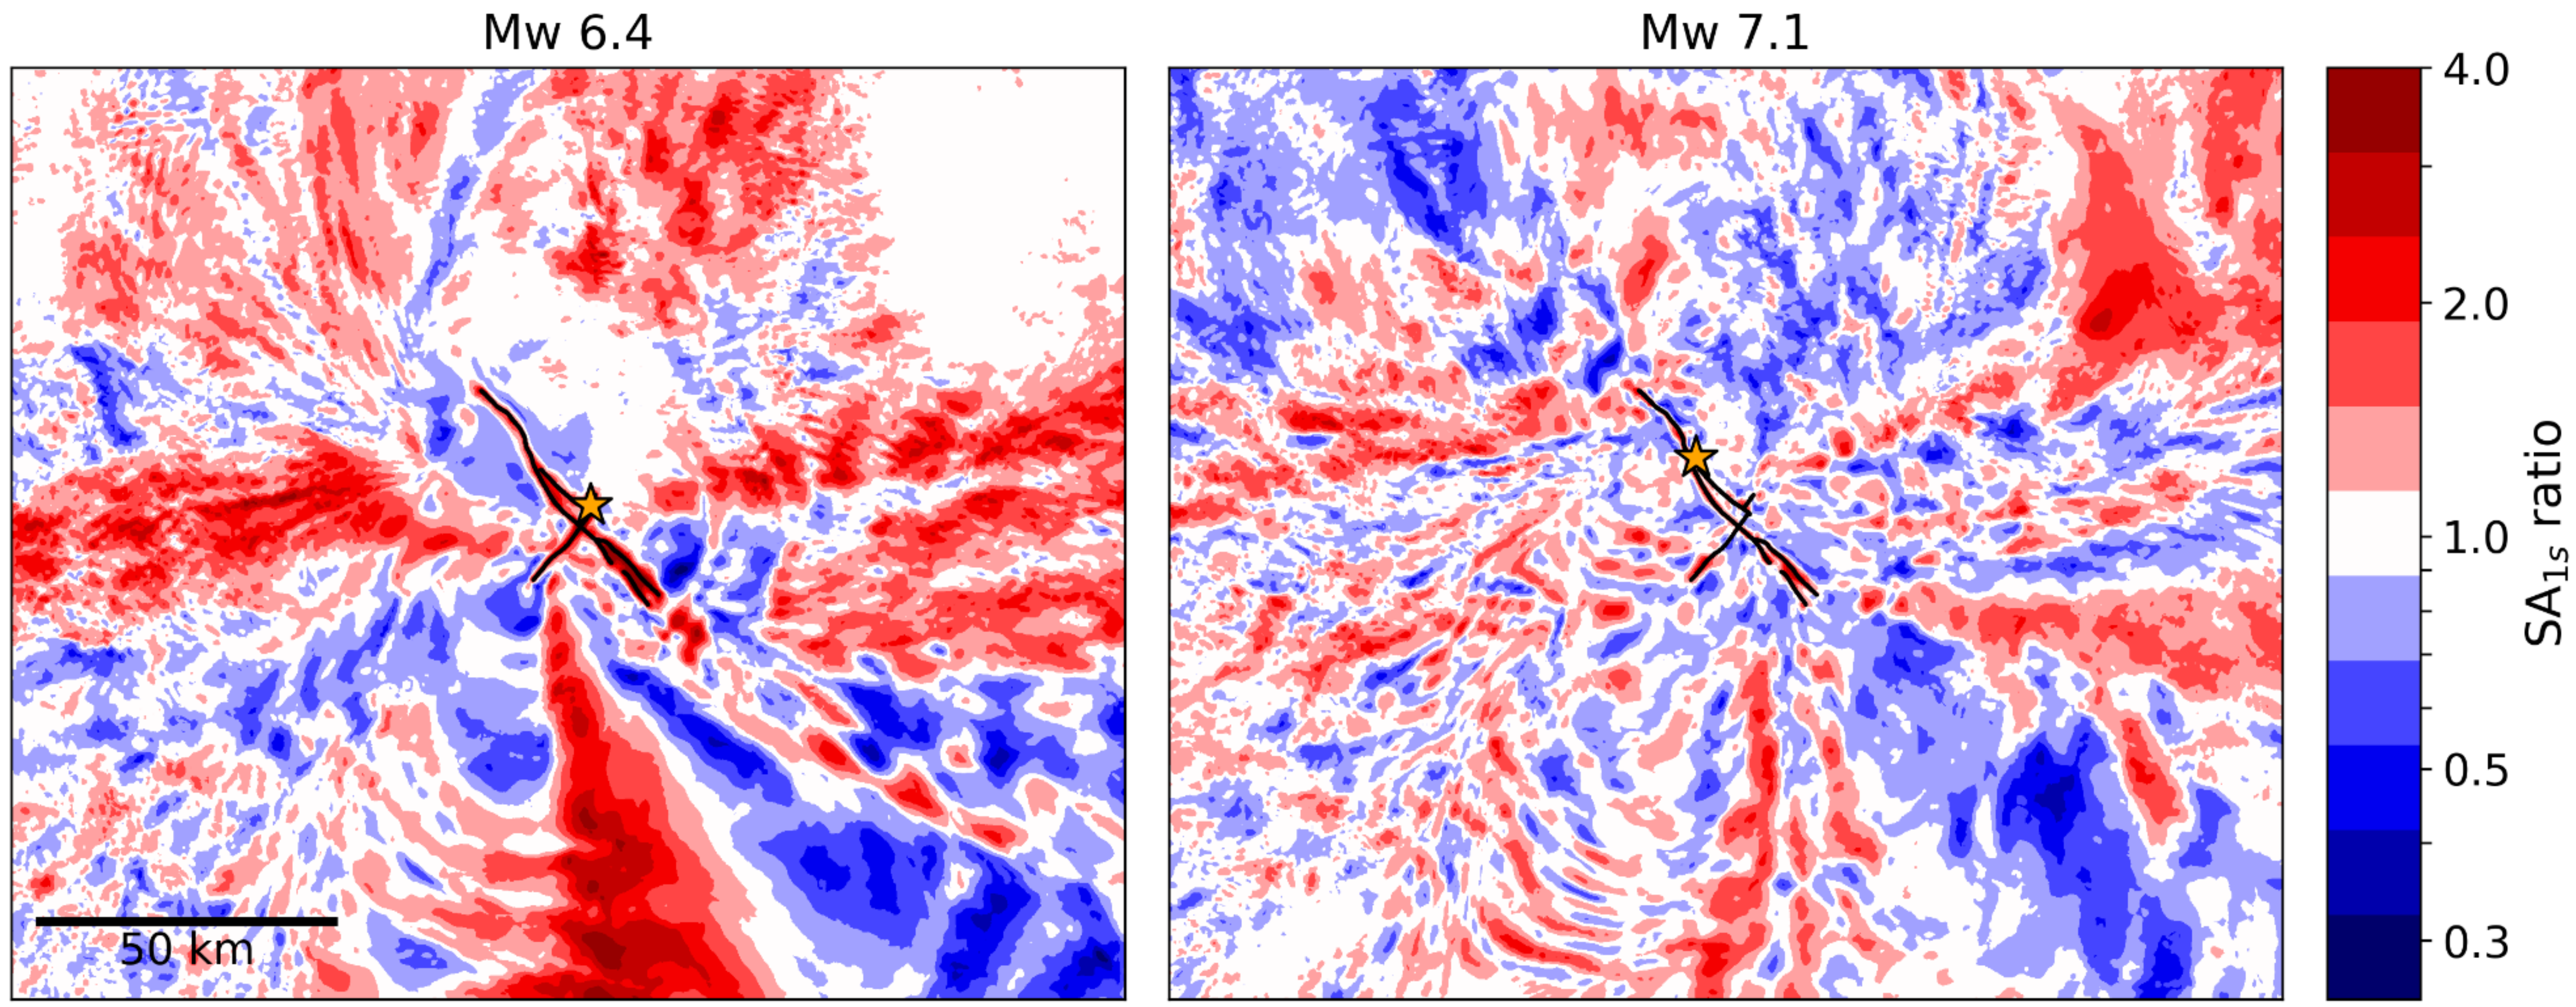

Flower-shaped fault zone affects spectral acceleration in large distance to the fault. Ratio of fault zone model spectral accelerations to model without a fault zone. Fault zone impact is not limited to the direct vicinity of the fault system but still widely affects ground motions at distances of more than 50 km to the rupturing faults. (De-) amplification patterns are highly heterogeneous and depend strongly on the regarded frequencies. (Schliwa et al., SCEC'22_

Simulation Big Data

- 3D simulation output even of small simulations is **>10s of TeraByte**
- **Reduction is possible but limited:** using modern data formats (hdf5), single precision (50%), file-system aware sequential output (10-20%), or only storying 2D output (at Earth's surface / on faults, ~ hundred GB)
- **FAIR data sharing standards** achieved by archiving simulation input & parameters (~ hundred GB, no output)

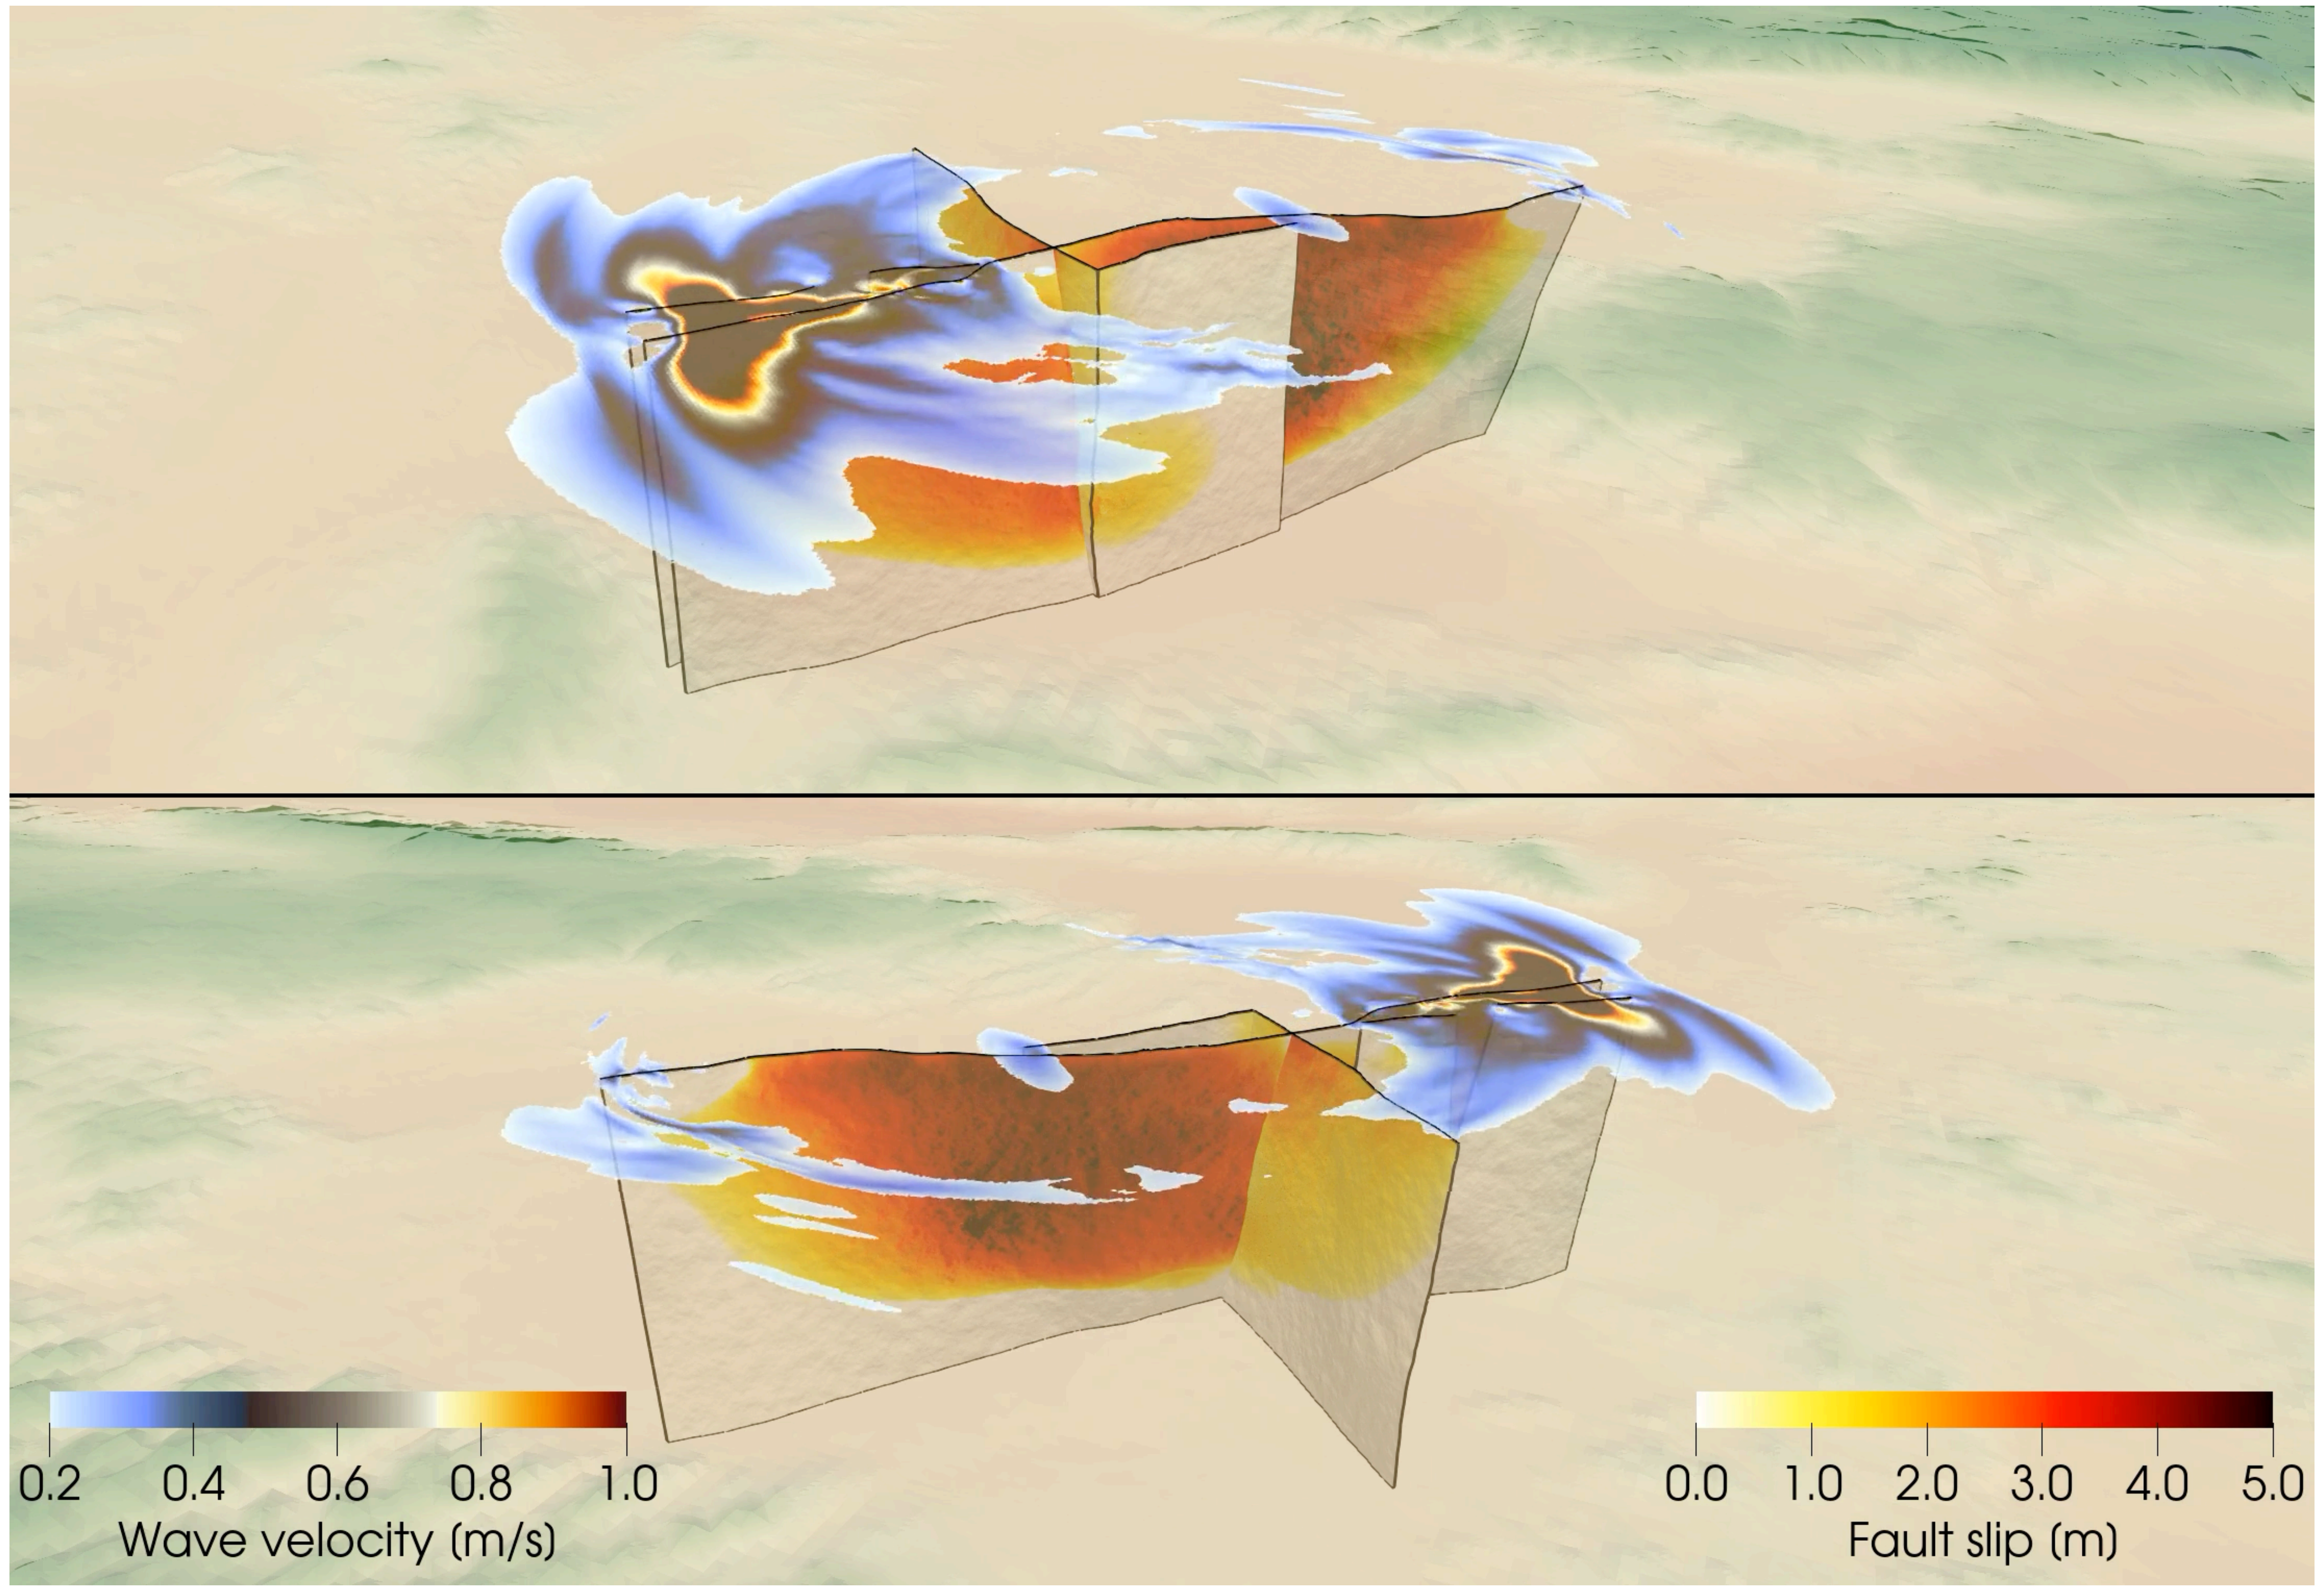

Propagation of seismic waves and "unzipping" of faults during the 2019 Ridgecrest earthquakes. Visualization of 15 TB of 3D simulation data on a supercomputer by Greg Abram and Francesca Samsel (Texas Advanced Computing Center)

Visualisation combining 2D data outputs by Nico Schliwa

Dynamic earthquake source inversion

Gallovic et al., JGR, 2019a, 2019b

- Bayesian dynamic source inversion of the **2016 Mw 6.4 Amatrice, Italy, normal faulting earthquake (300 fatalities)**
- **Dense network** of 20 three-component strong motion seismological stations within 50 km to the fault
- **Highly non-linear, very high-dimensional forward problem (and no gradients)**

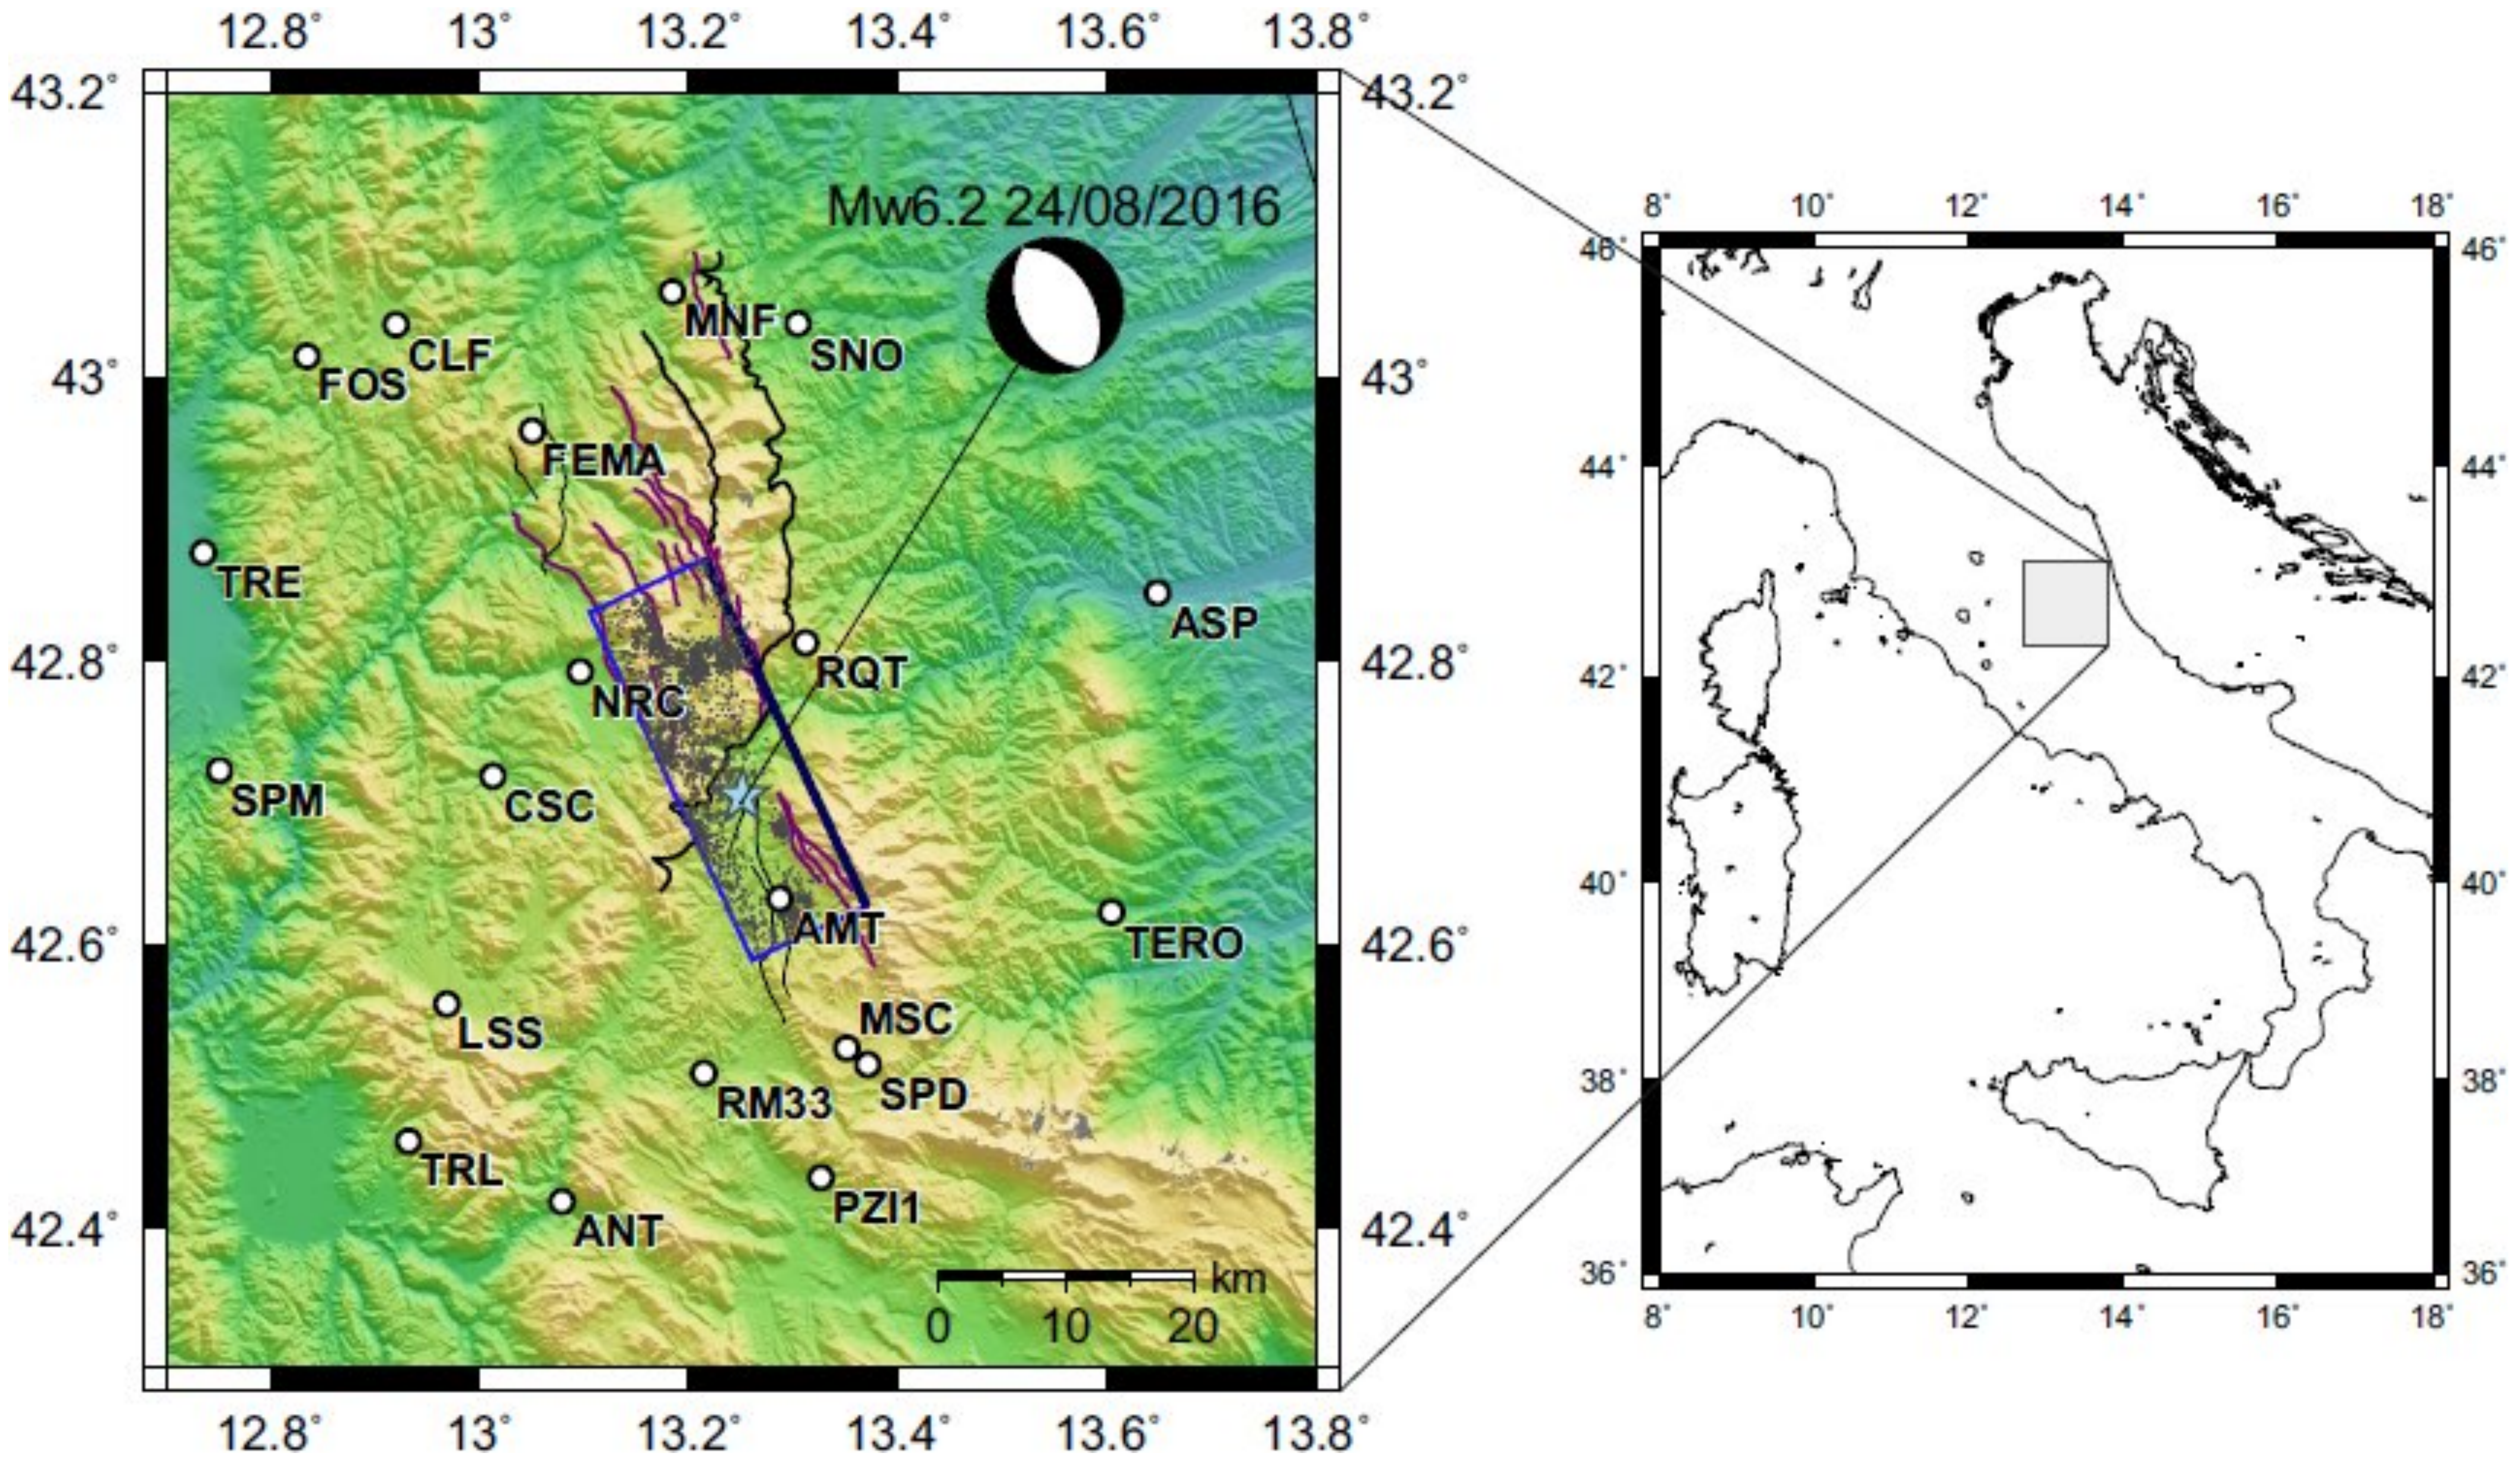

Dynamic earthquake source inversion

Gallovic et al., JGR, 2019a, 2019b

- Bayesian dynamic source inversion of the **2016 Mw 6.4 Amatrice, Italy, normal faulting earthquake (300 fatalities)**
- **Dense network** of 20 three-component strong motion seismological stations within 50 km to the fault
- **Highly non-linear, very high-dimensional forward problem (and no gradients)**

- Bayesian framework using a **Parallel Tempering Monte Carlo Markov Chain algorithm to sample the posterior PDF** (Falcioni and Deem, 1999; Sambridge, 2013, 2014)

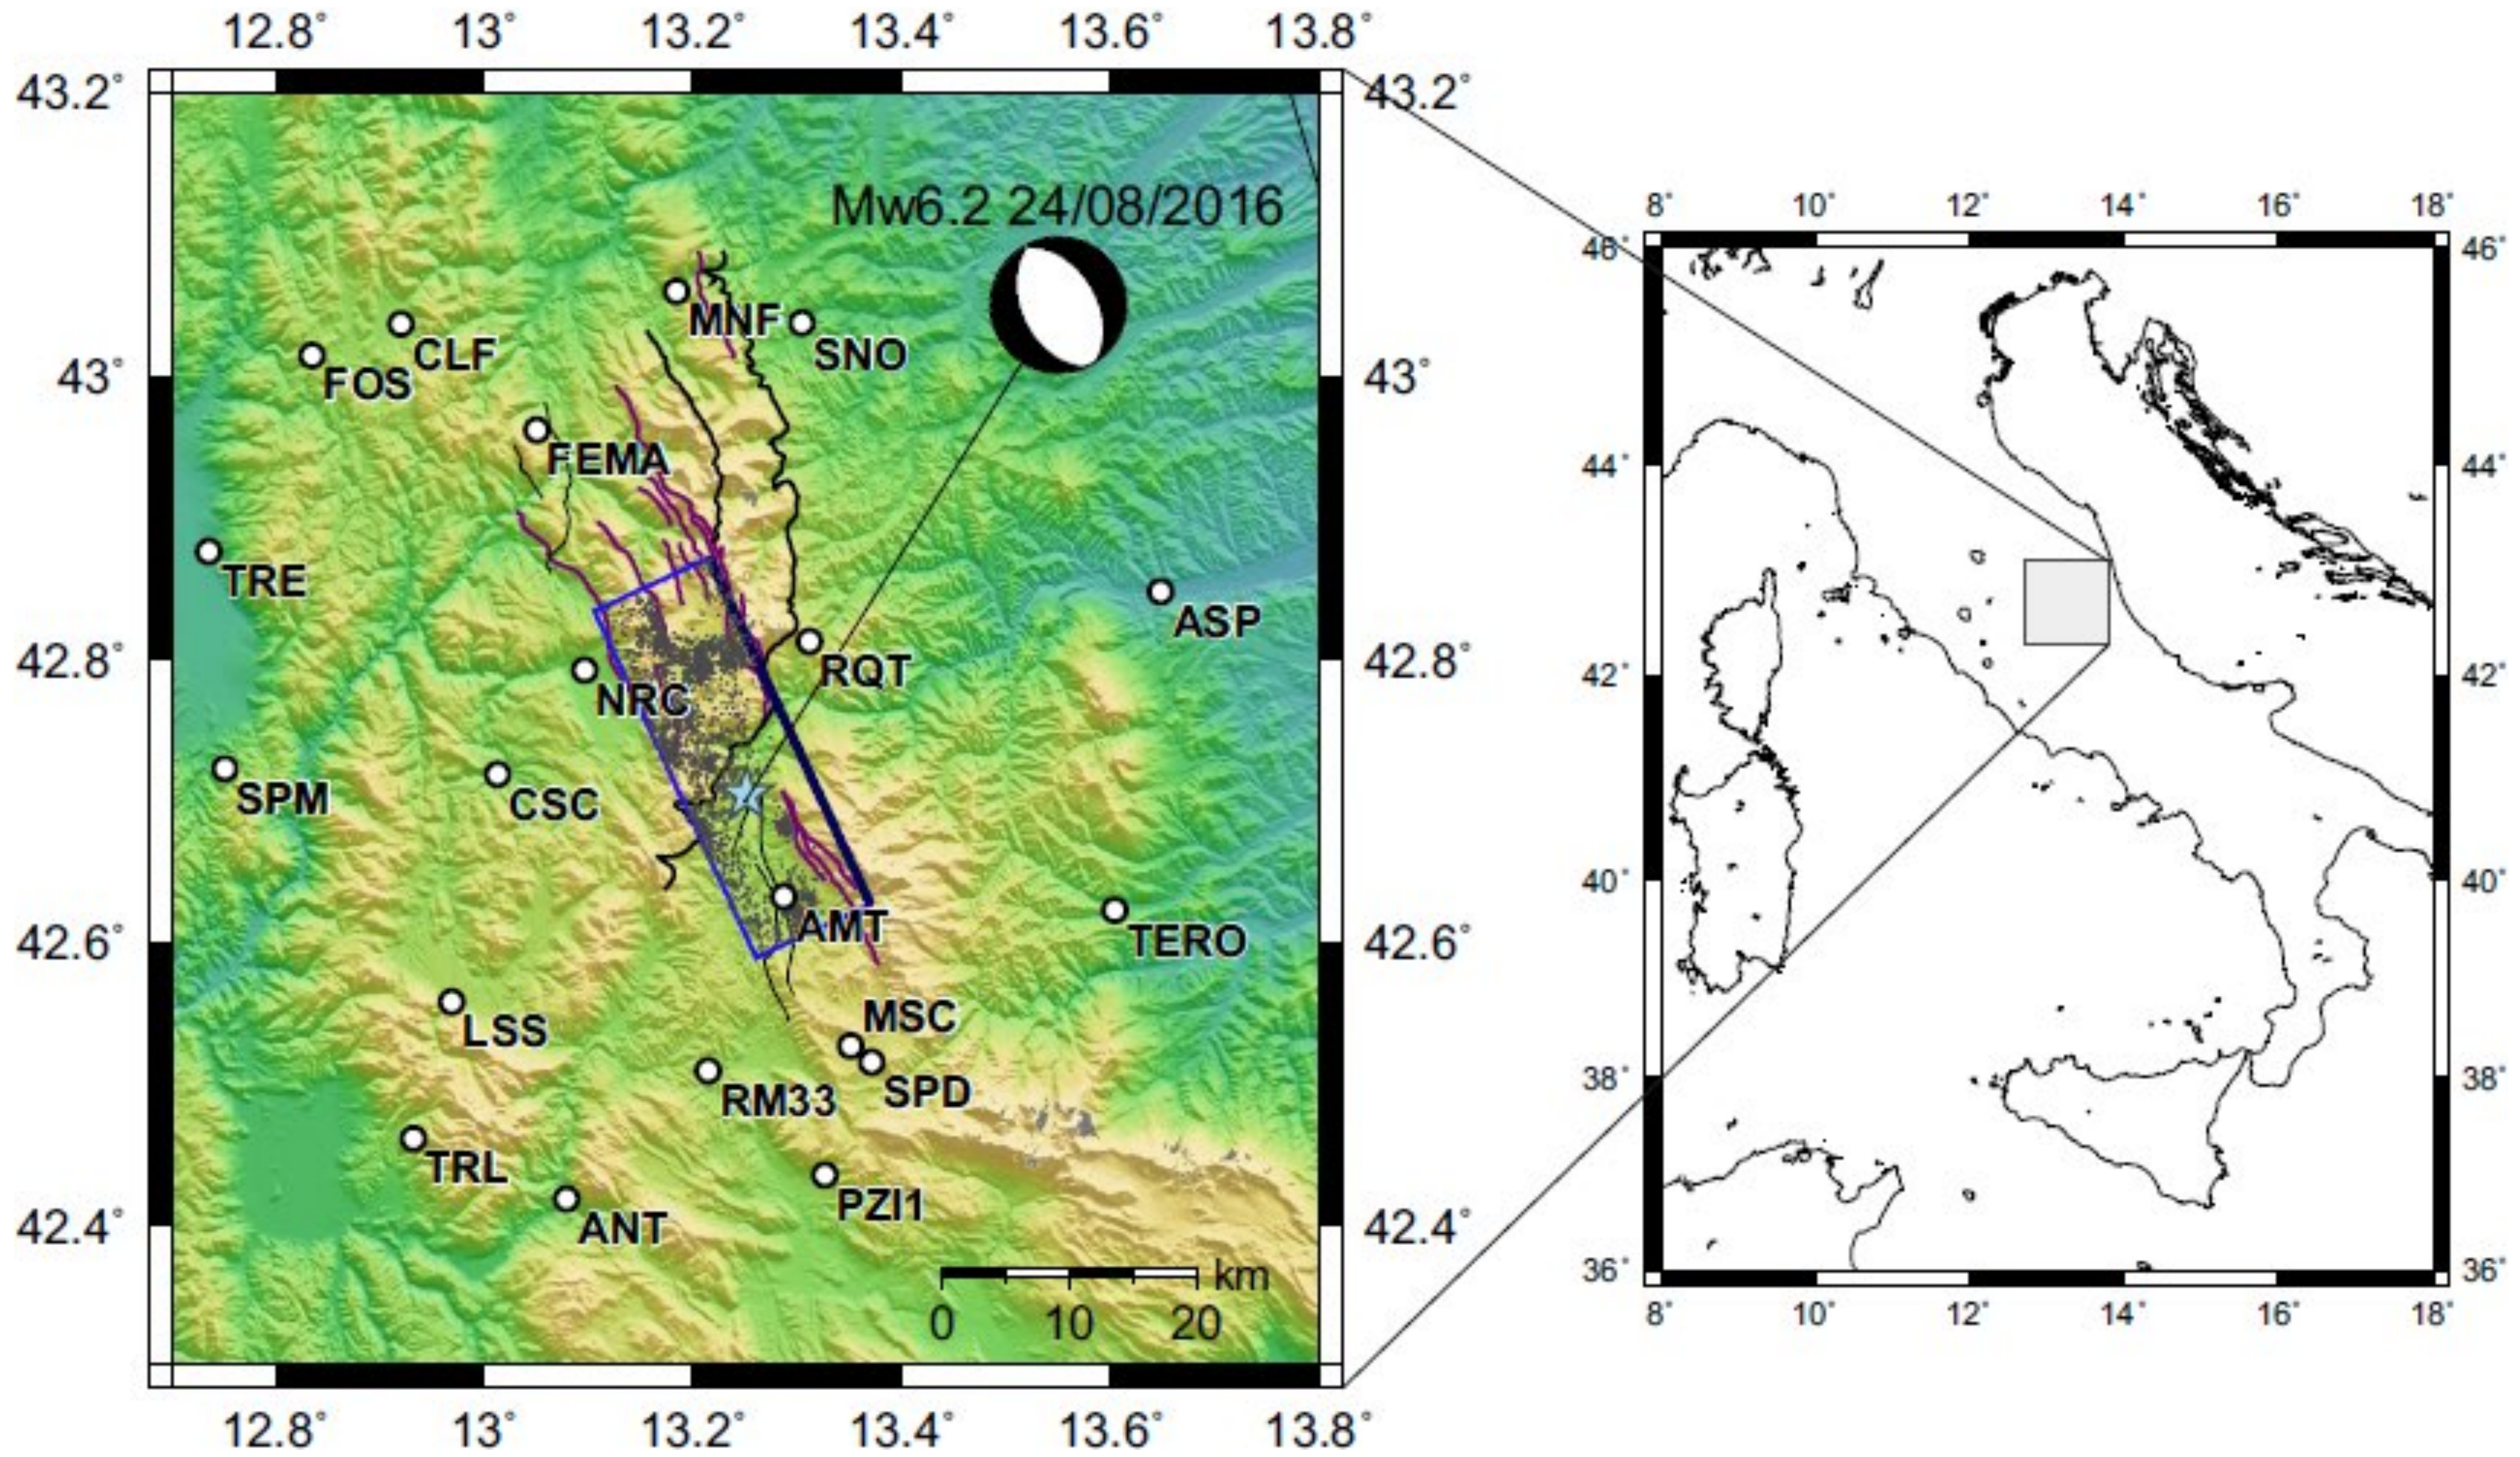

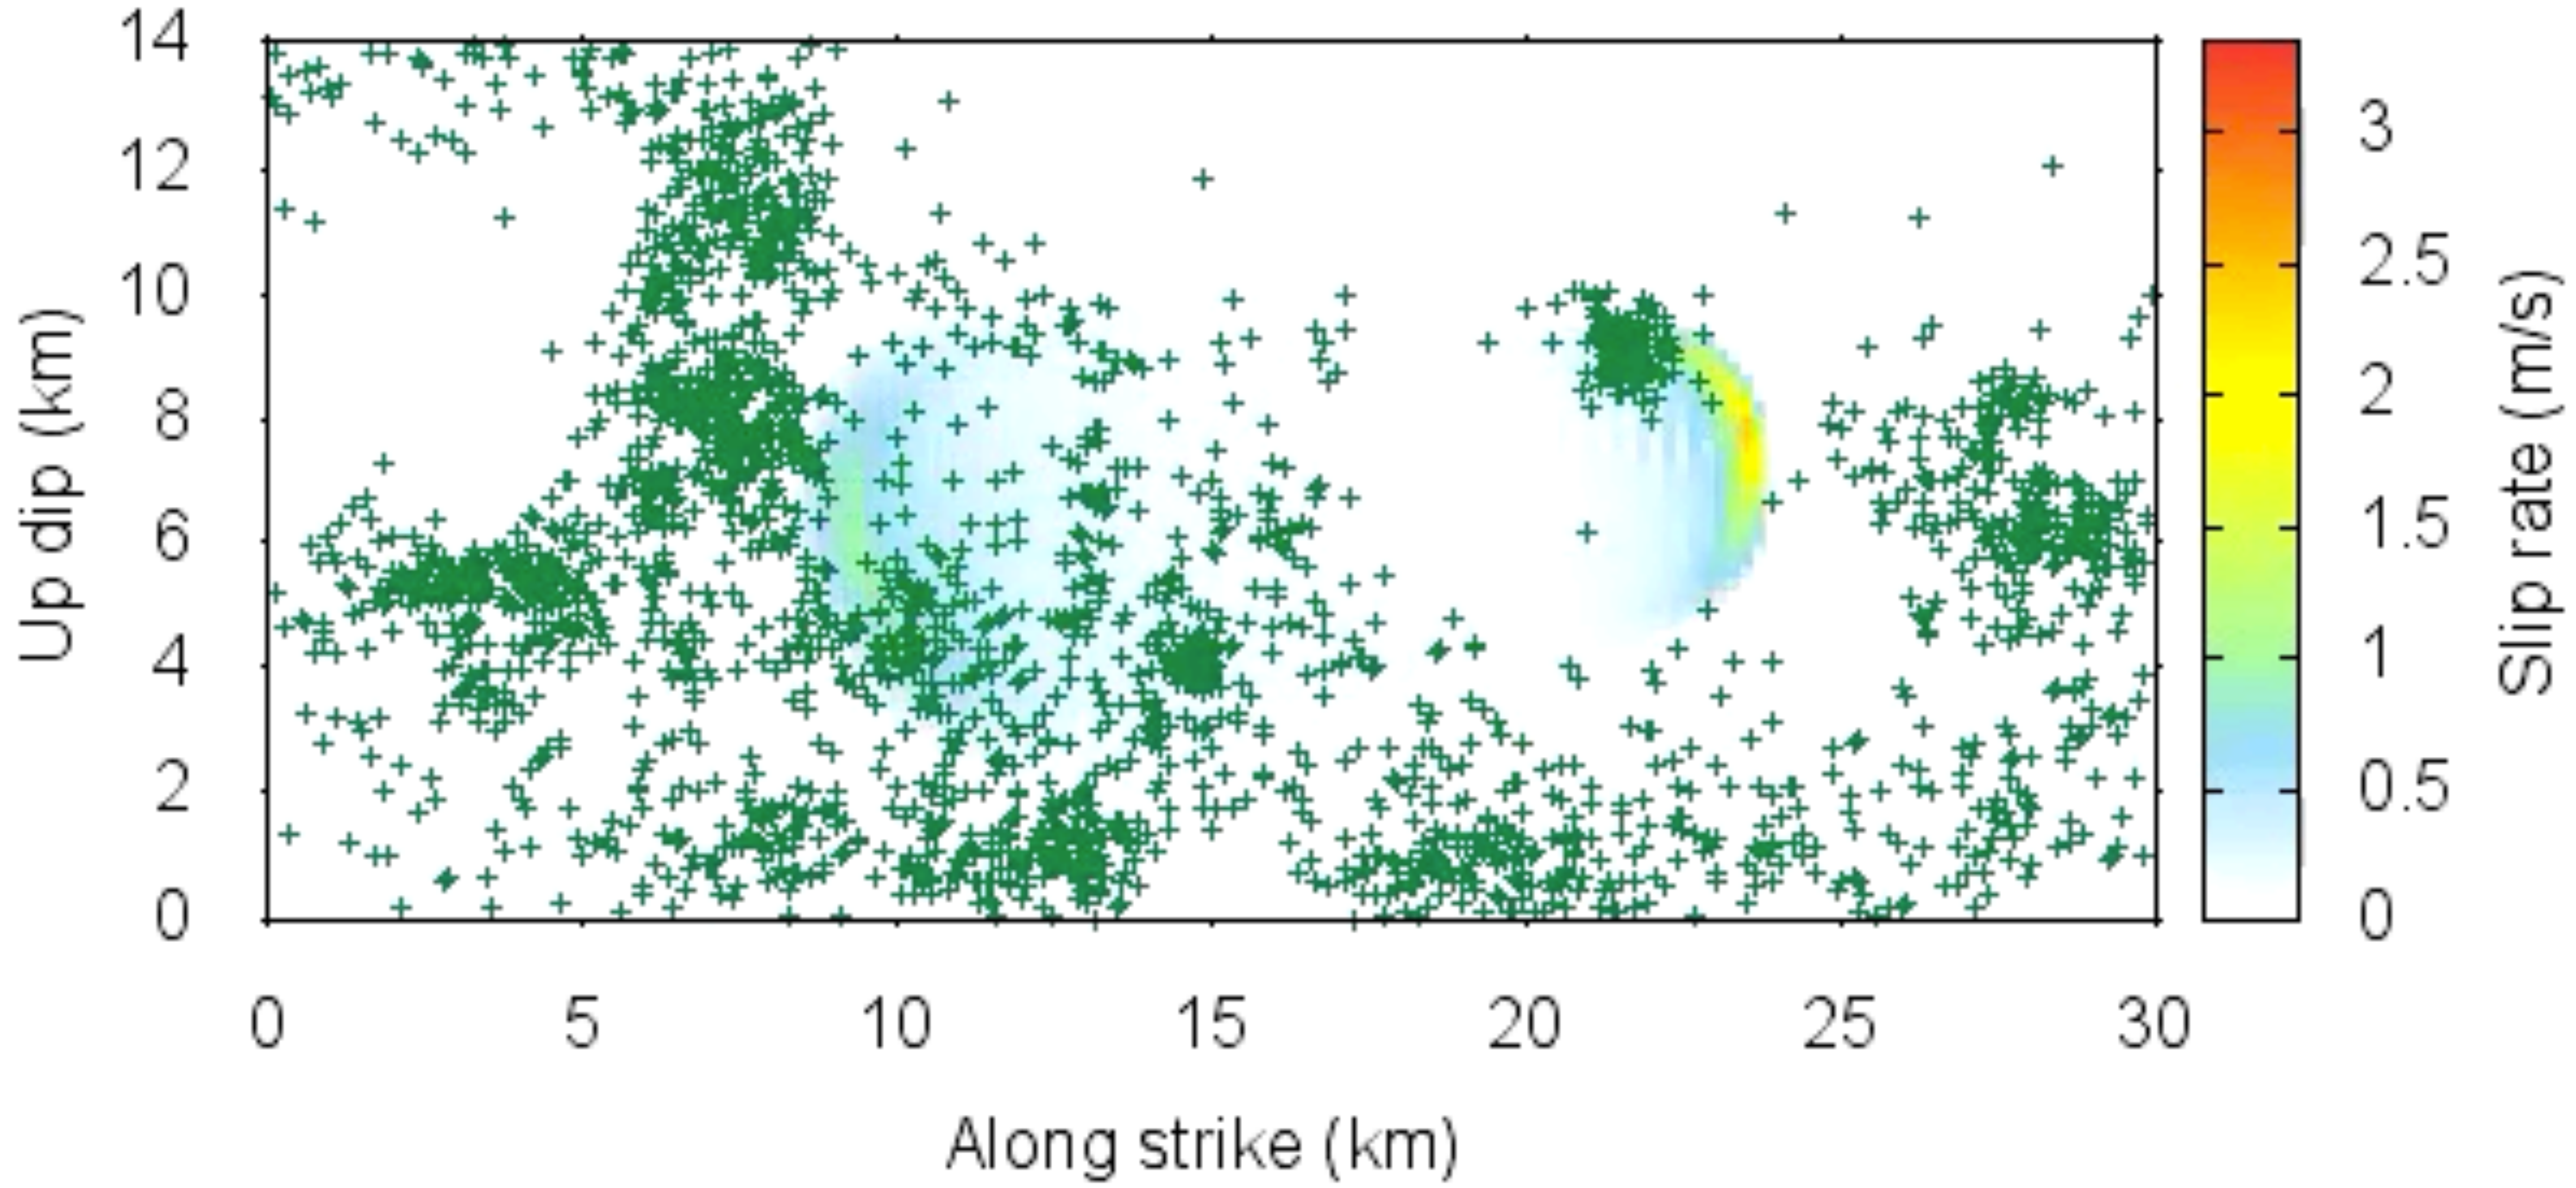

Best-fitting model of Dynamic source inversion (out of ~10⁶ visited models) in frequency range 0.05-1.0Hz (AMT and NRC stations) and 0.05-0.5Hz (others).

Dynamic earthquake rupture initiates as a weak crack with large rise times (4 s), then turns into a pulse with short rise times (0.2 s).

Dynamic earthquake source inversion

Taufiqurrahman et al., GRL'22

- We augment a smooth, best-fitting **Bayesian dynamic rupture source inversion** with **fractal fault roughness, frictional heterogeneities, viscoelastic attenuation, and topography**

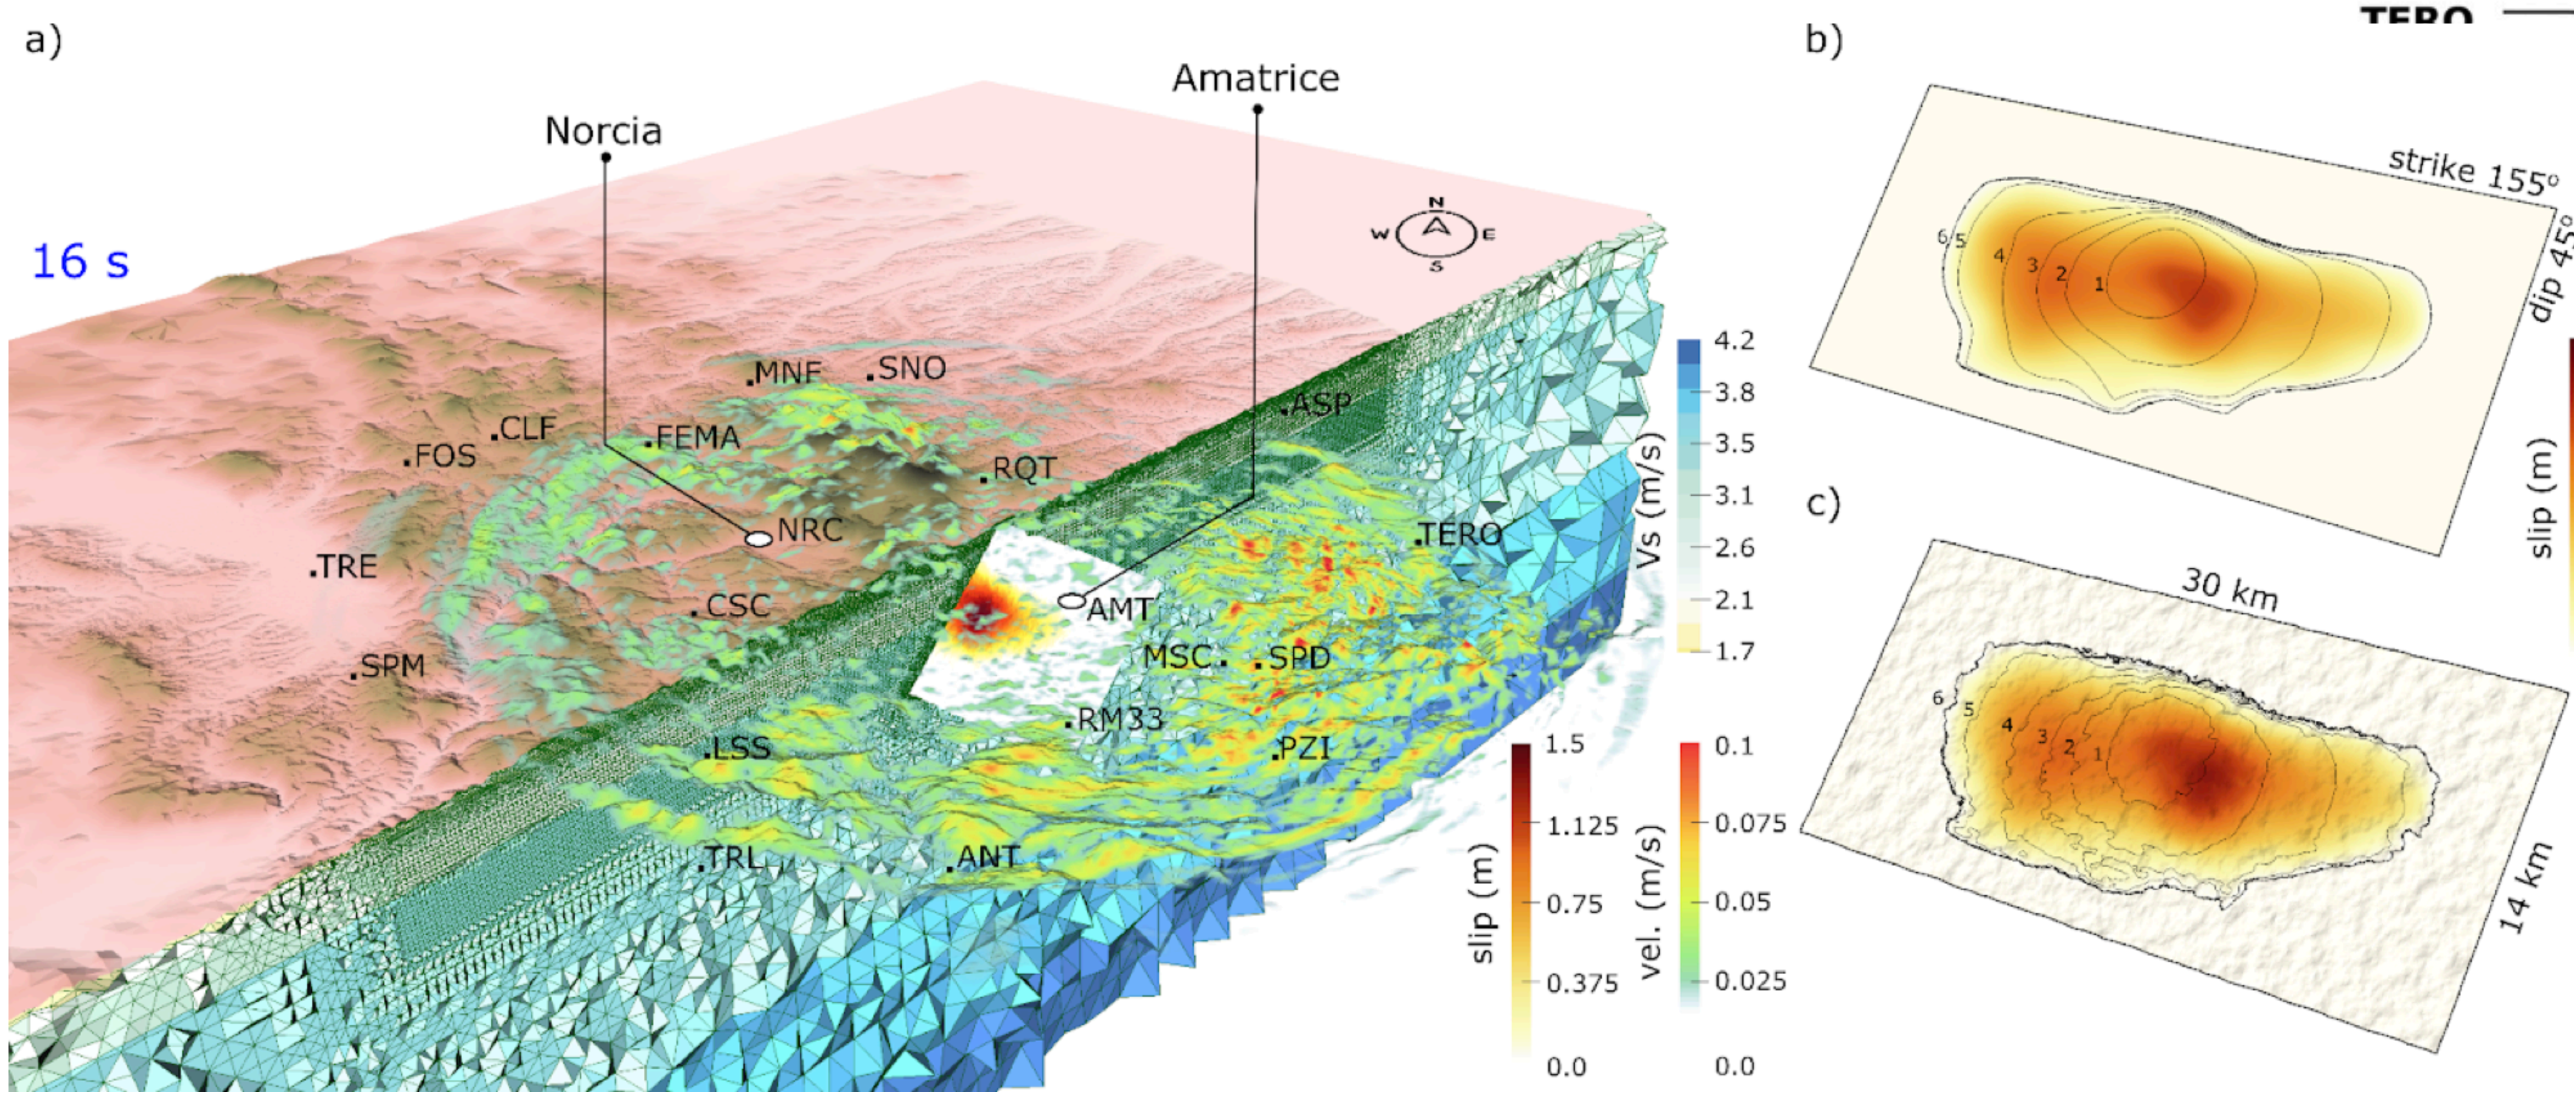

Broadband dynamic rupture modeling with fractal fault roughness, frictional heterogeneity, viscoelasticity and topography: the 2016 Mw 6.2 Amatrice, Italy earthquake

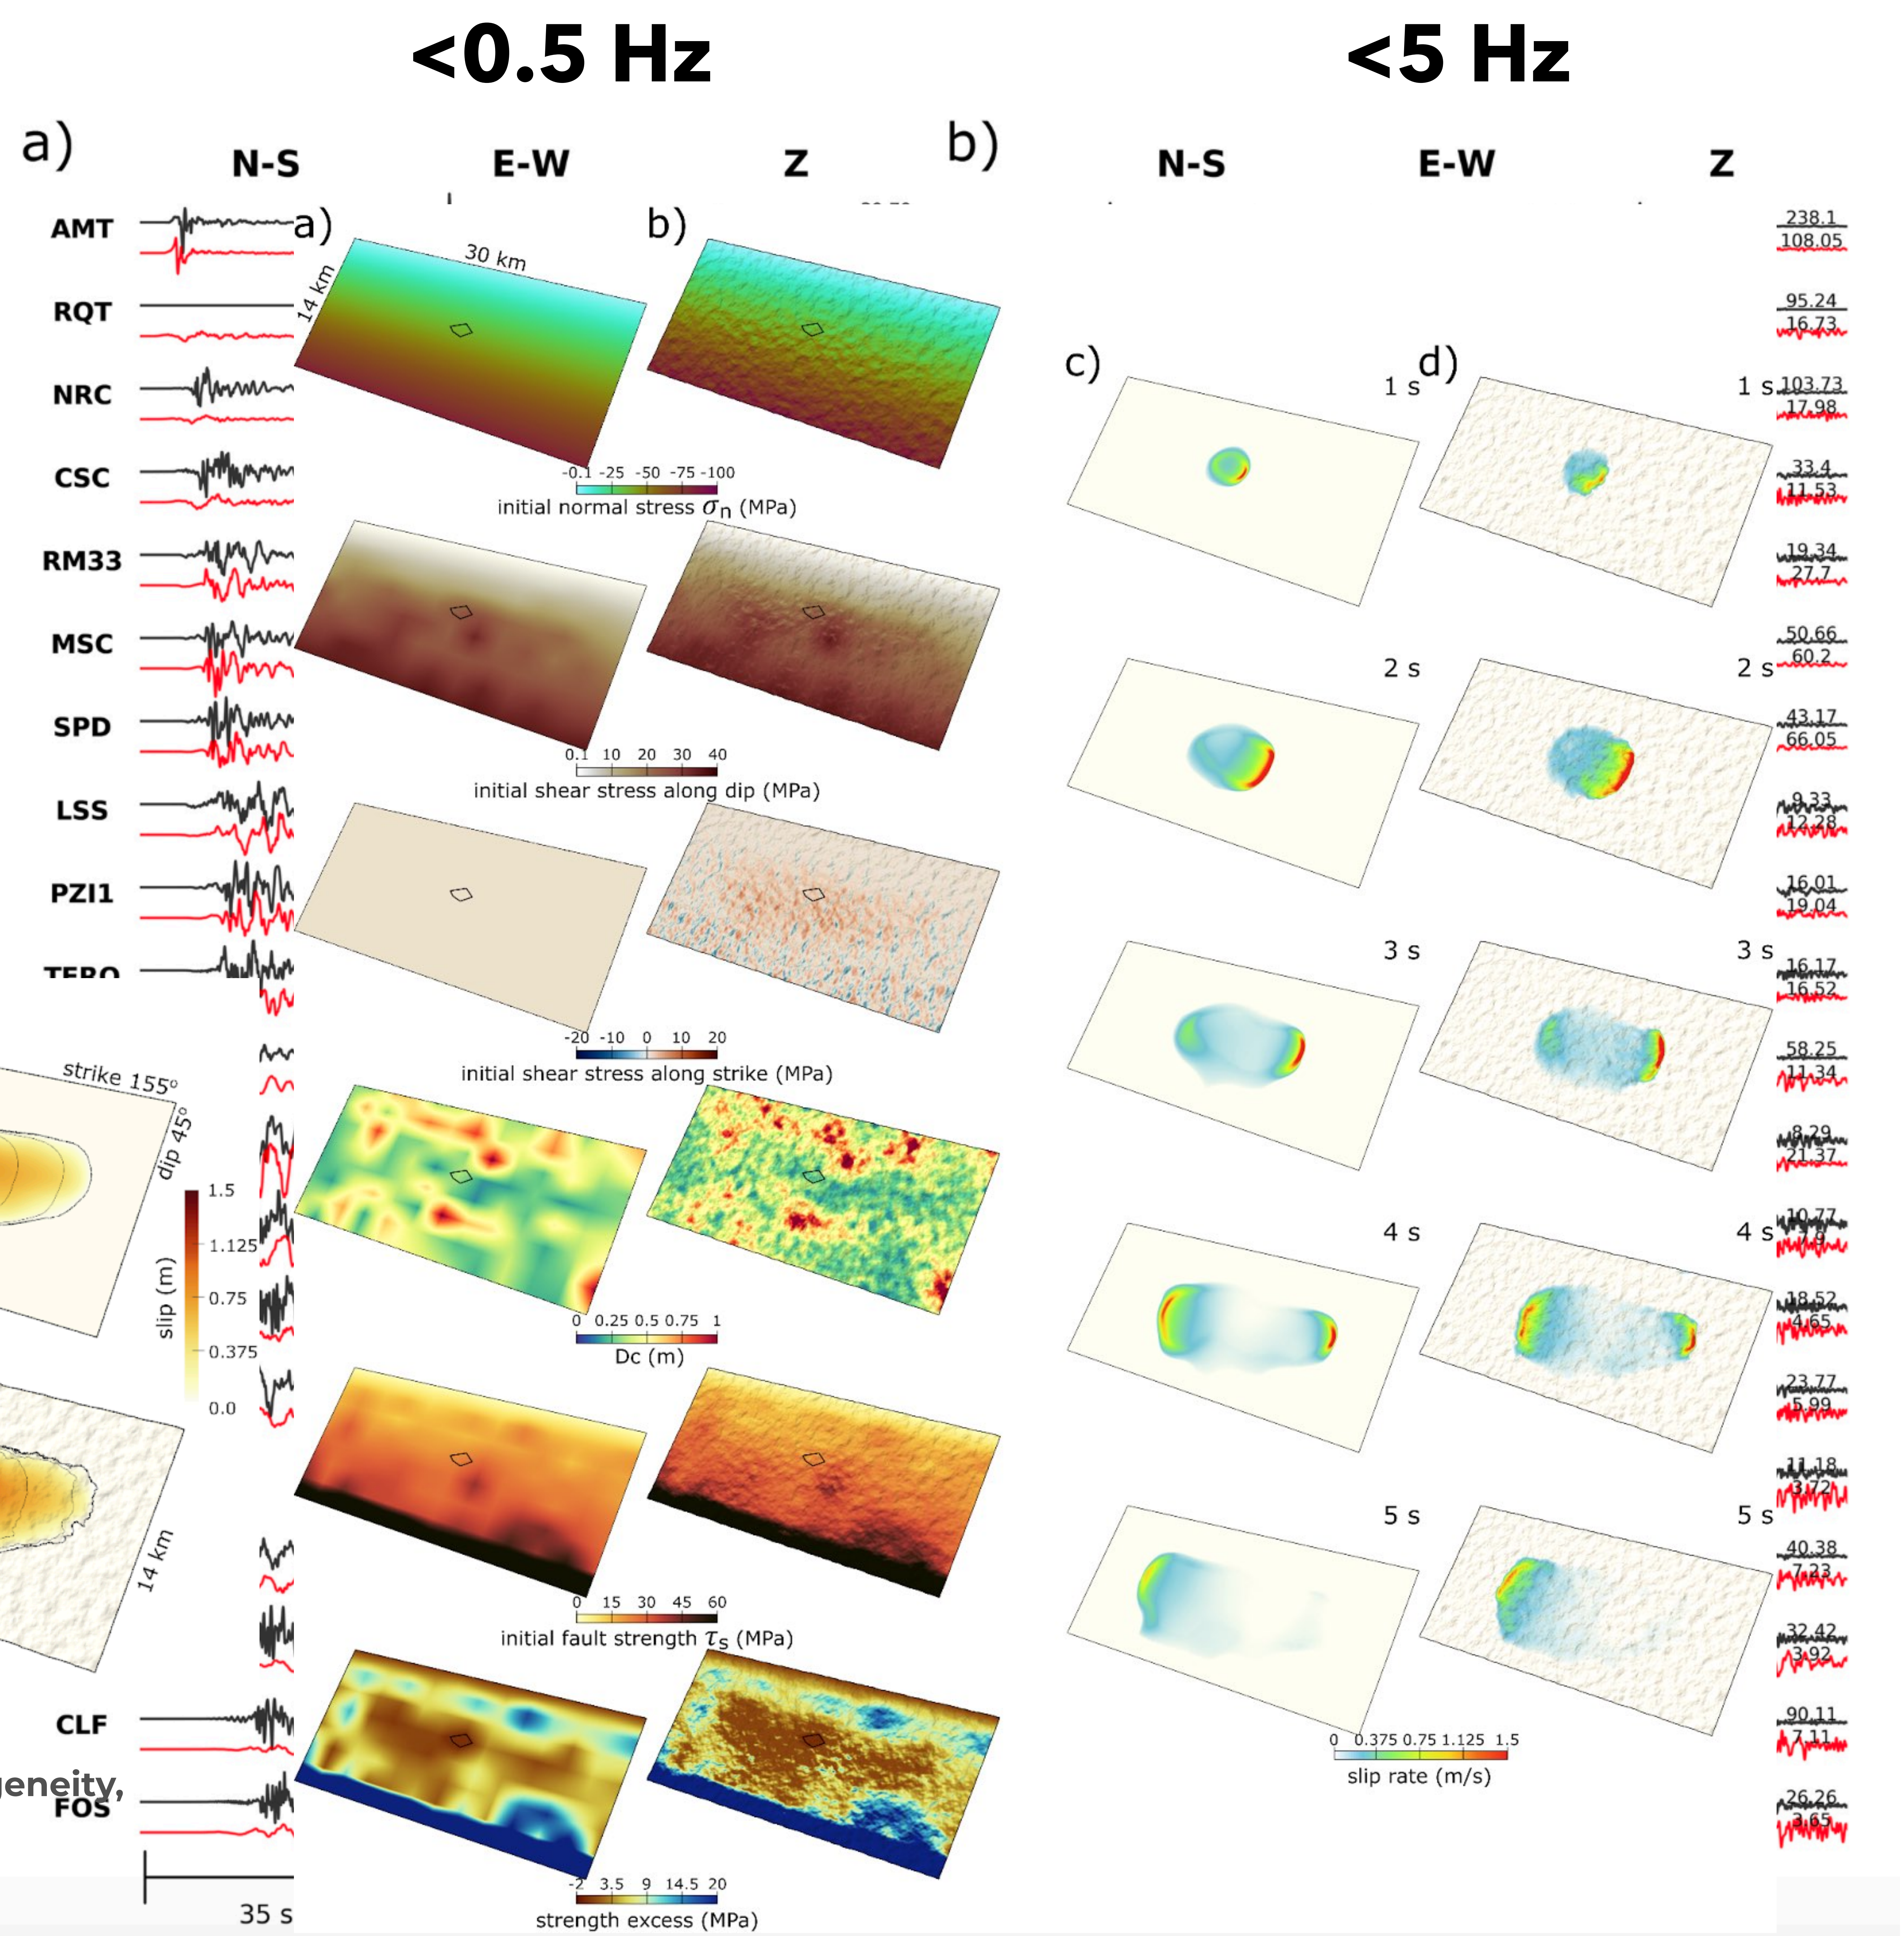

HPC enables data-driven ***and*** physics-based models

Taufiqurrahman et al., GRL’22

- We **require consistency at long periods** which allows us to **quantify the role** of earthquake dynamics, such as the 3D roughness drag, from observational broadband seismic waveforms

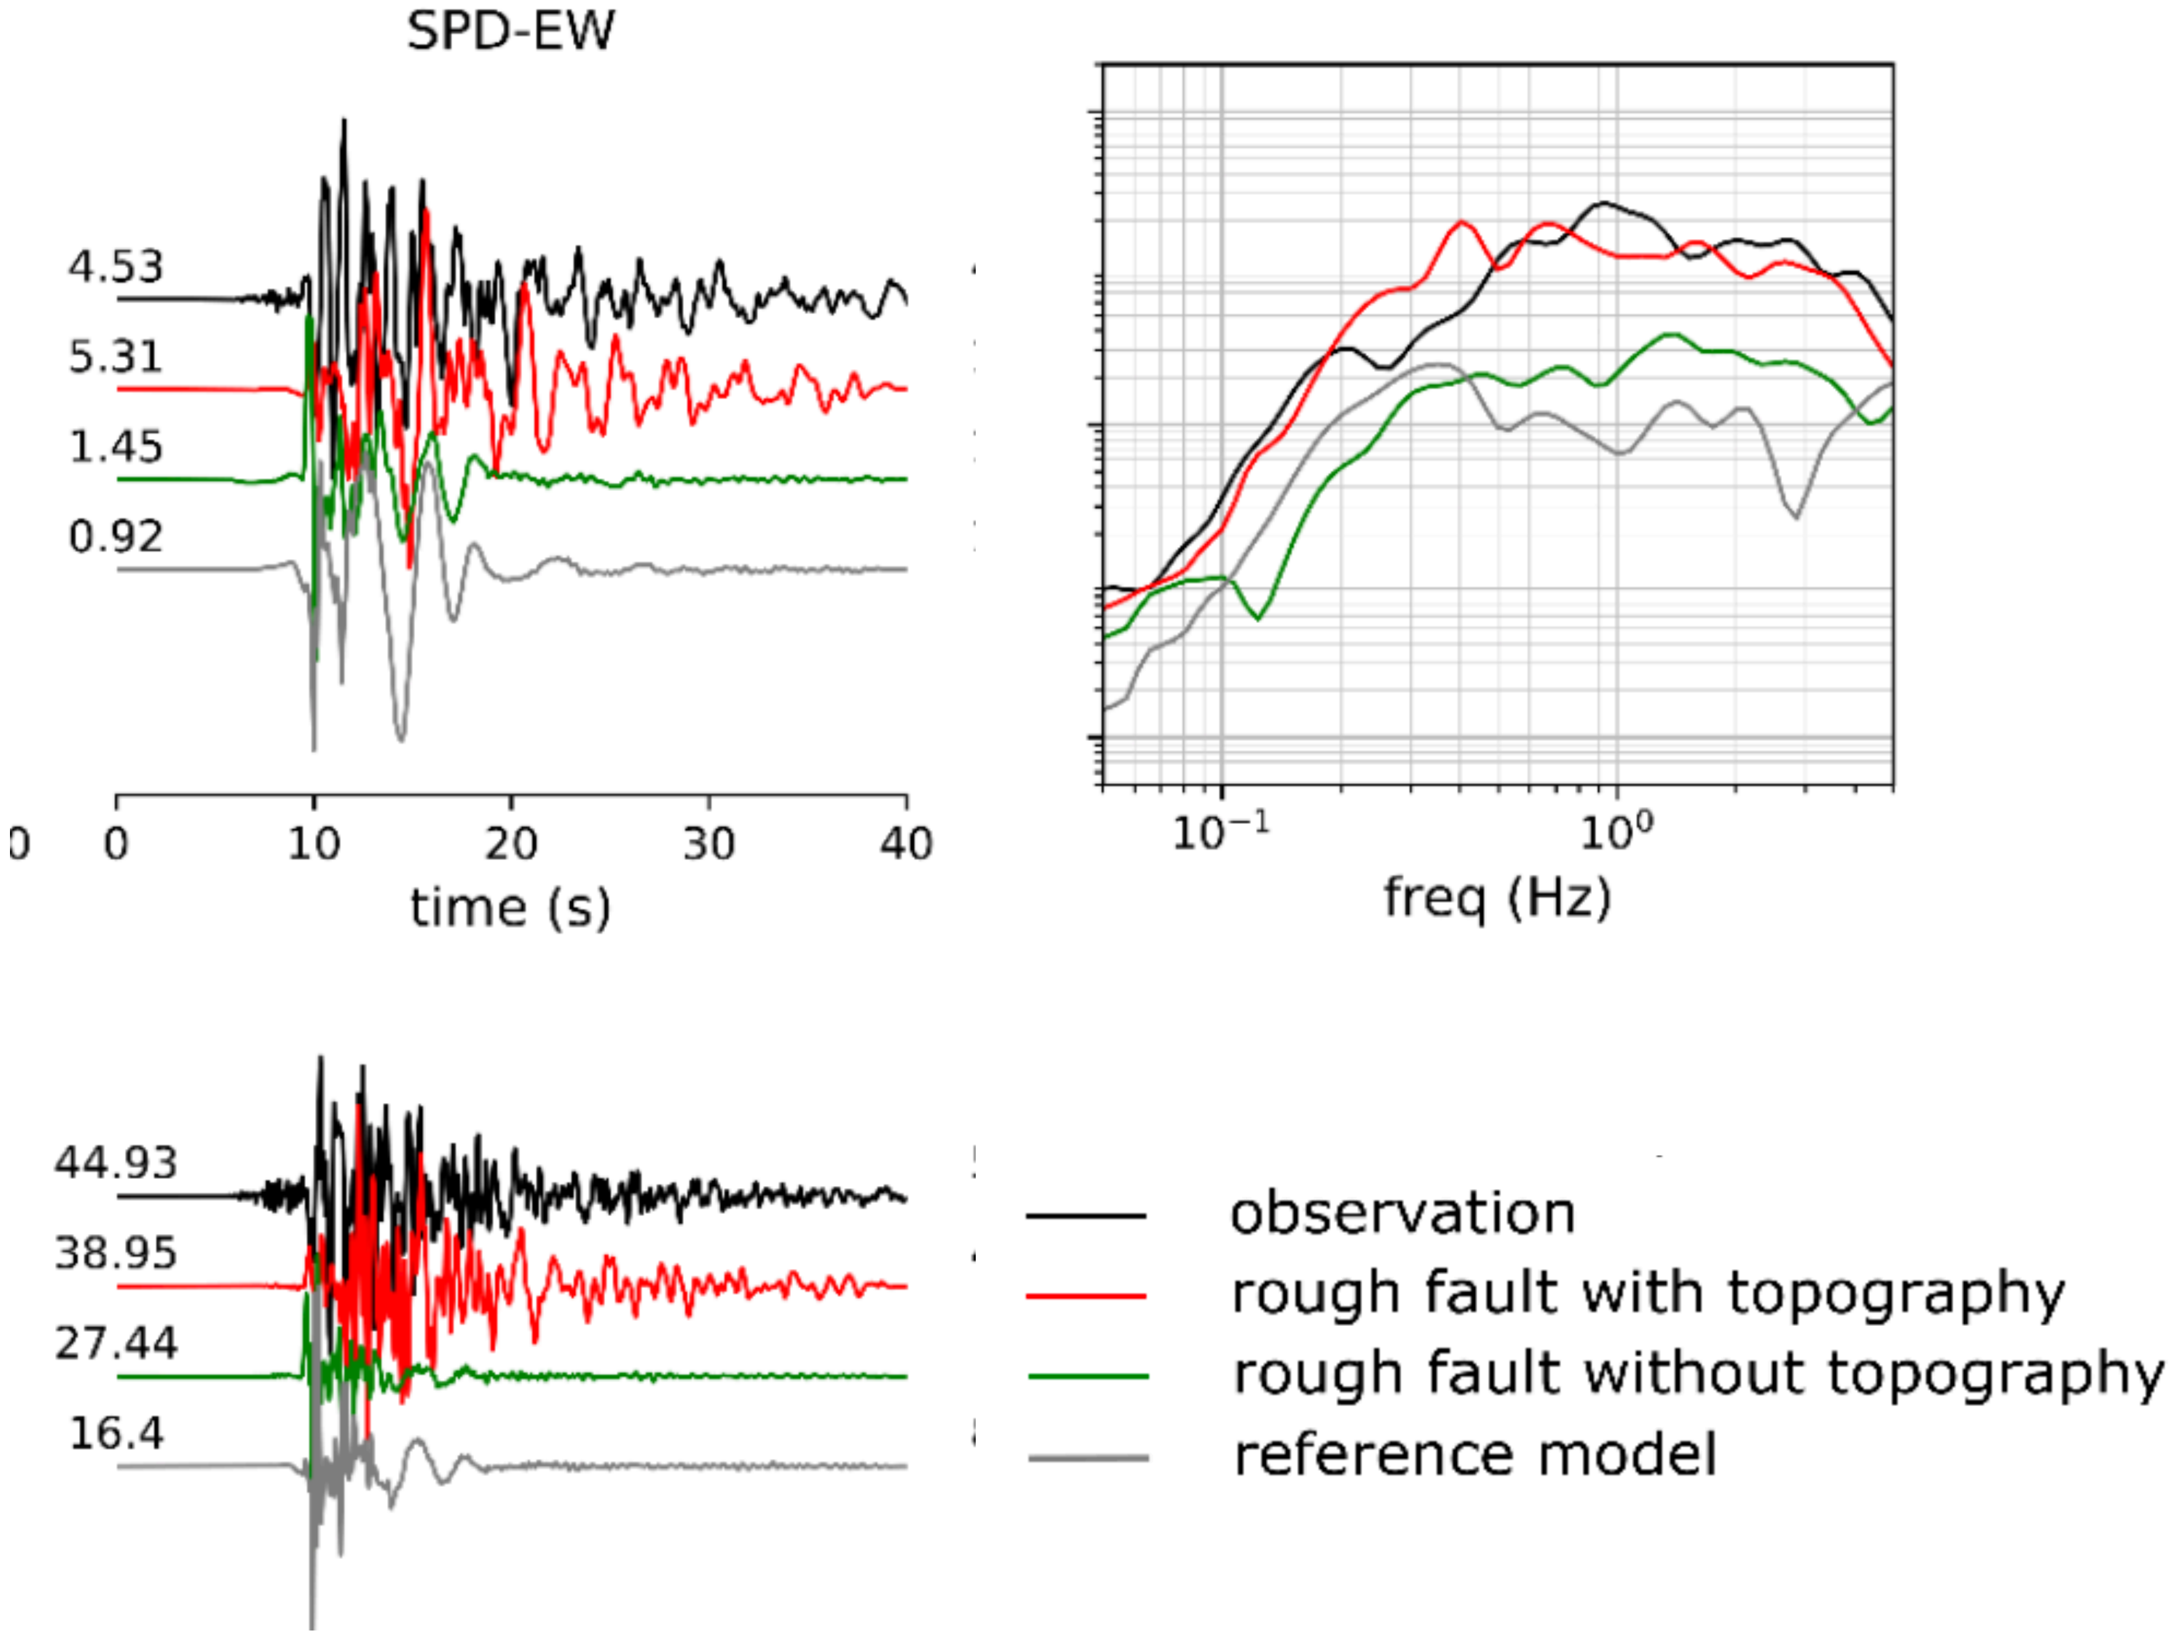

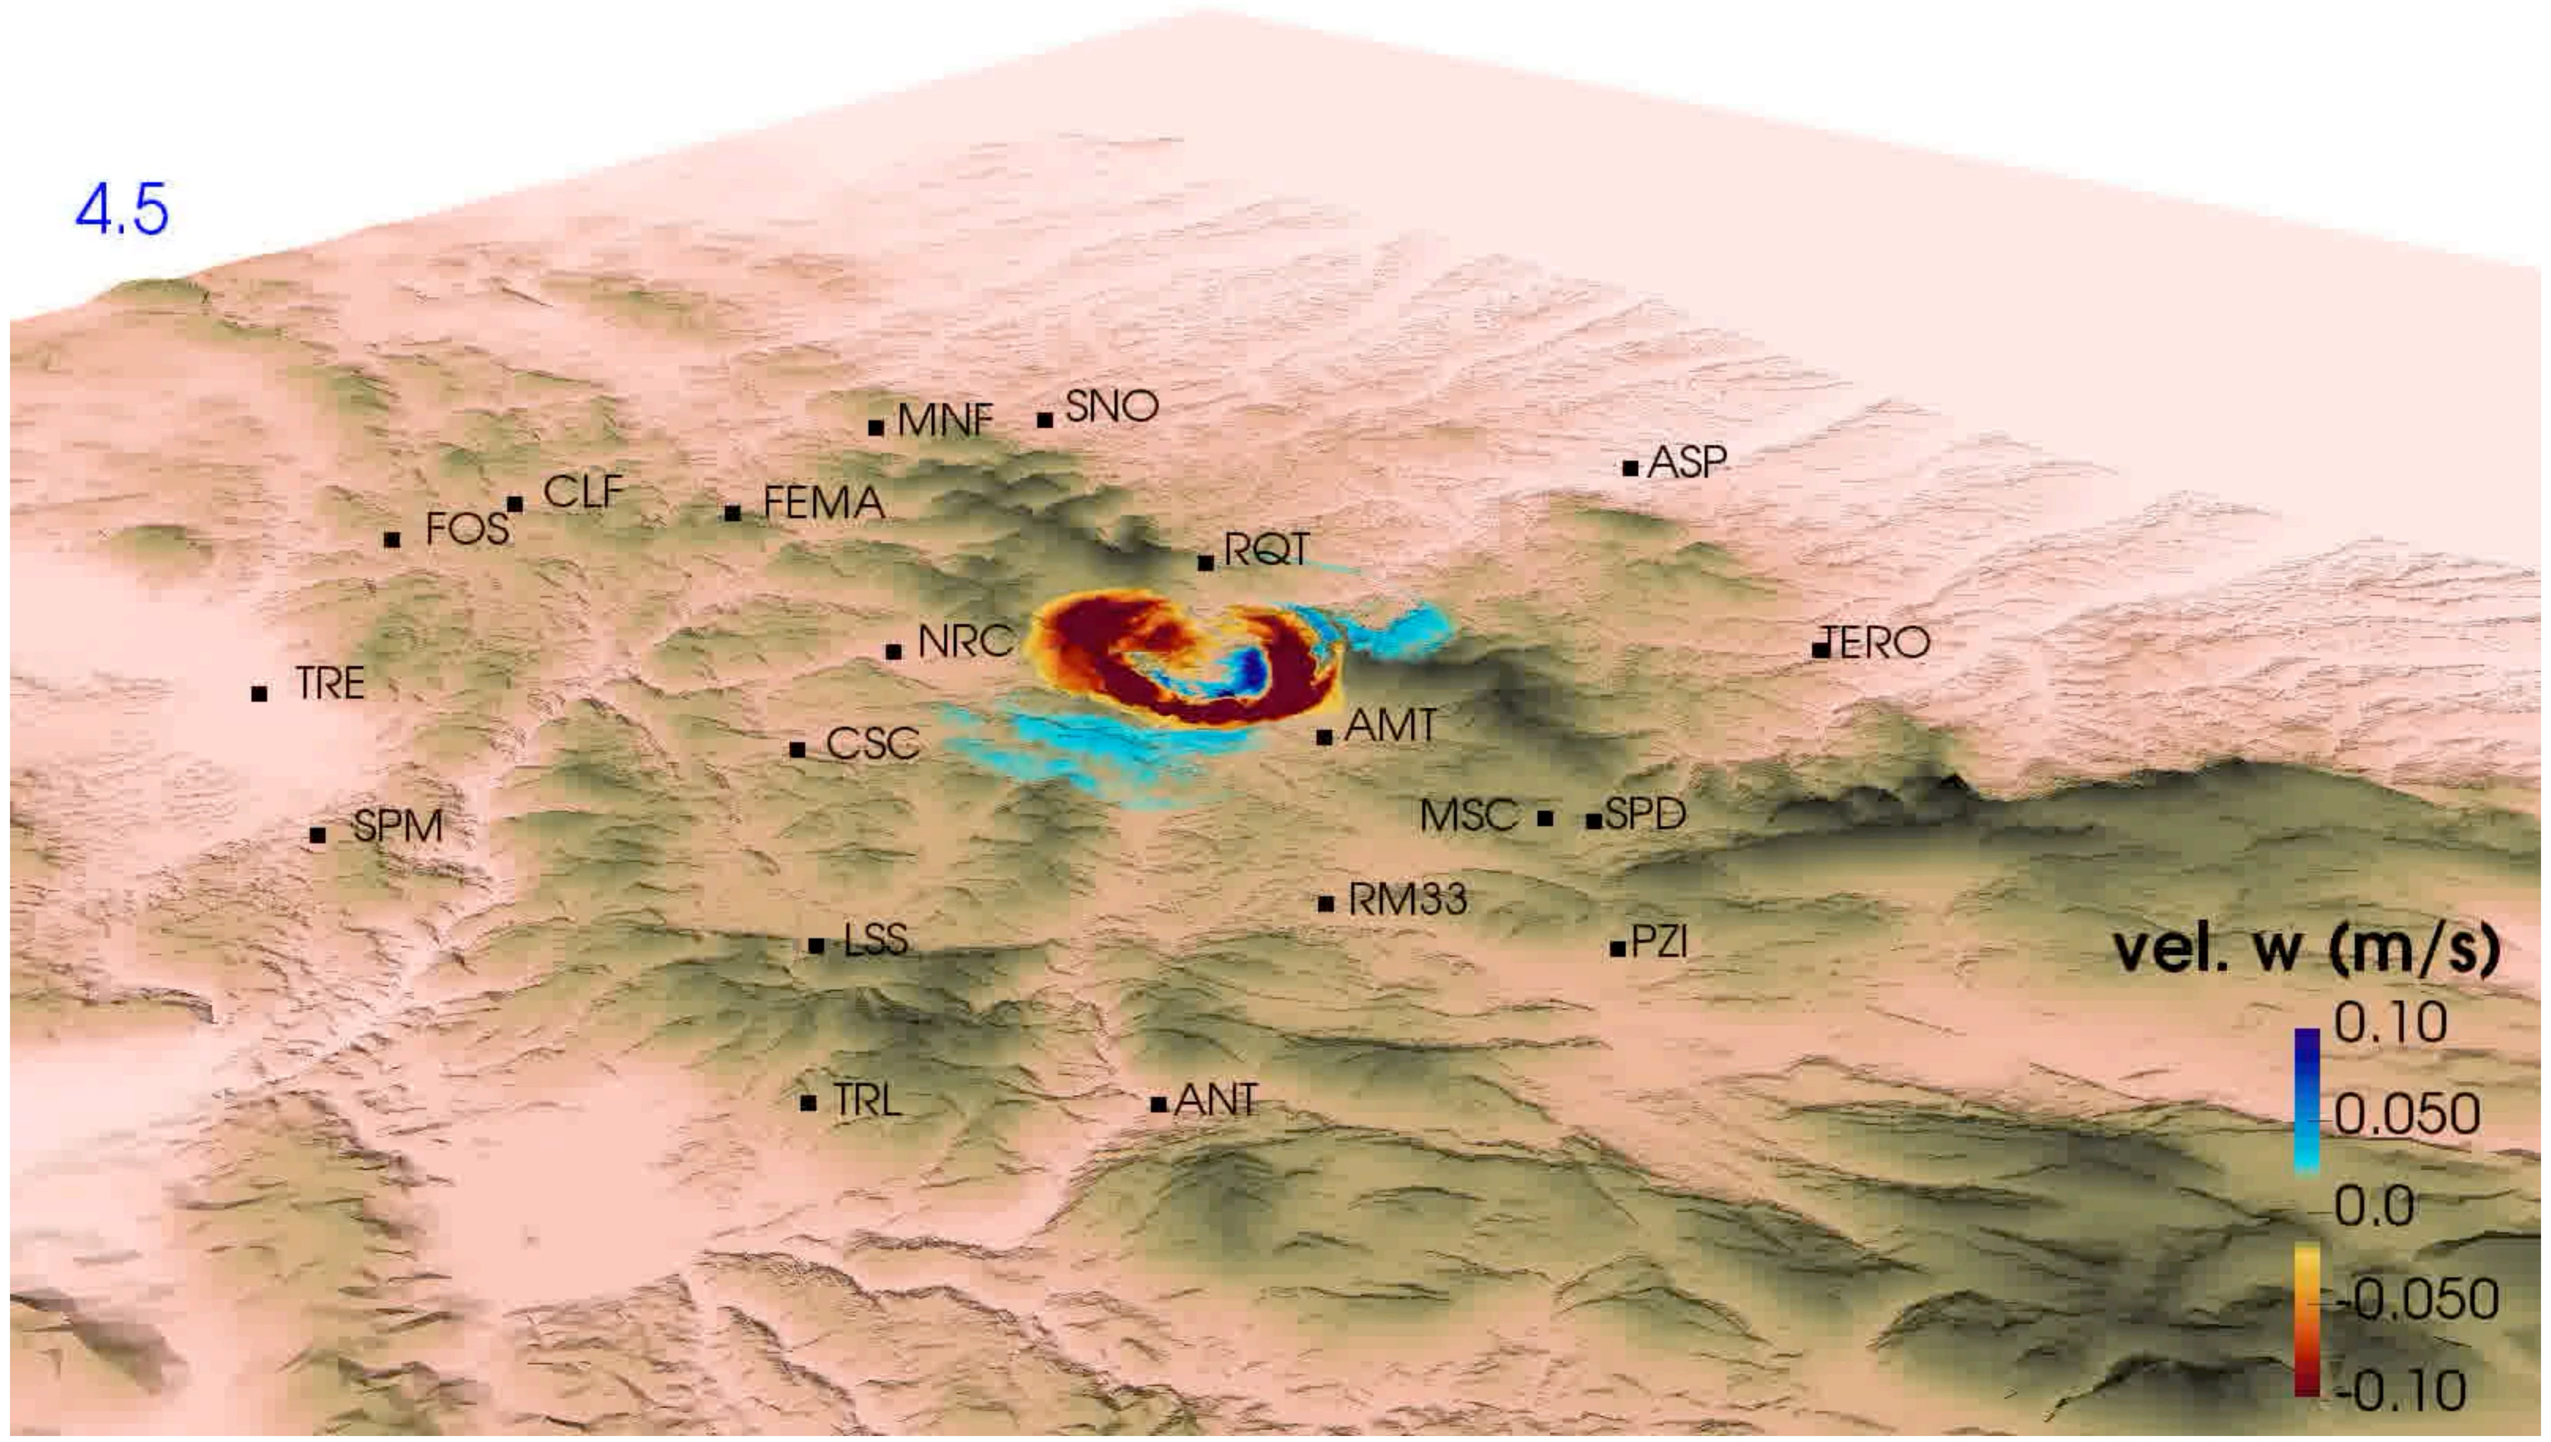

- High-frequency ground motions are **amplified early-on by fault roughness**, while **topography-induced scattered waves prolong their duration**

Bridging time scales of faulting: from coseismic to postseismic slip

Premus et al., Sci. Adv. '22; Schliwa et al., in prep.

- **Rate-and-state friction** can explain both, co-seismic (seconds, seismology & geodesy) **earthquake slip** and post-seismic (hours-weeks, geodesy) **after slip**
- **More variables to invert for but also more (and complementary!) observations**

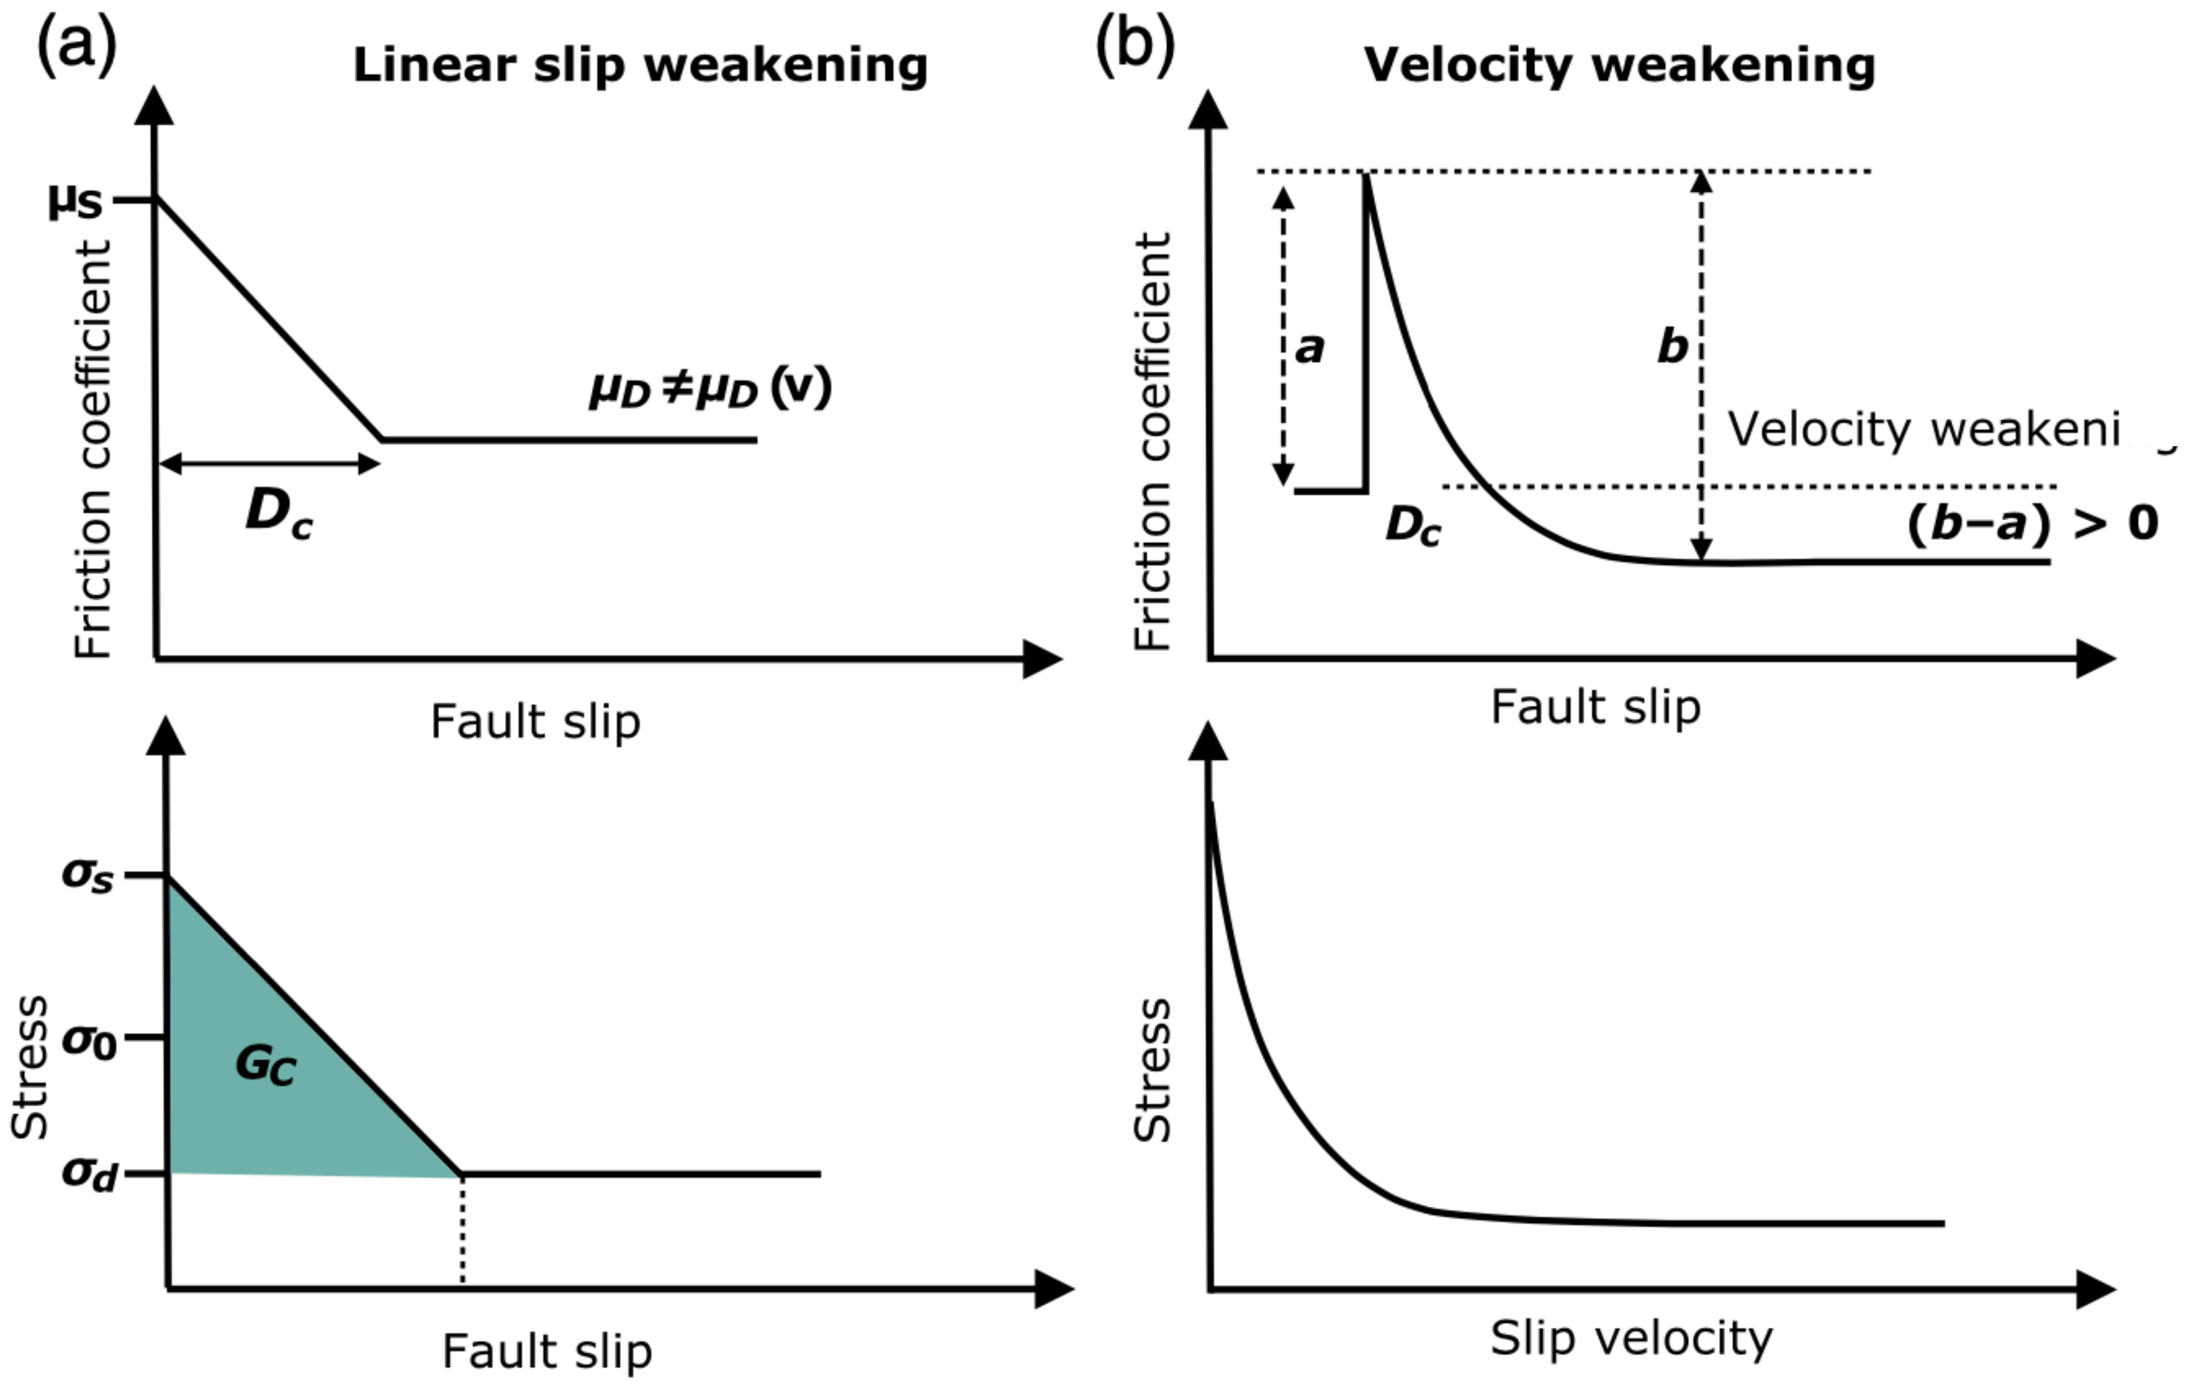

2004 Parkfield earthquake

35 seismic and 13 GPS stations recorded similarly large co- & postseismic slip

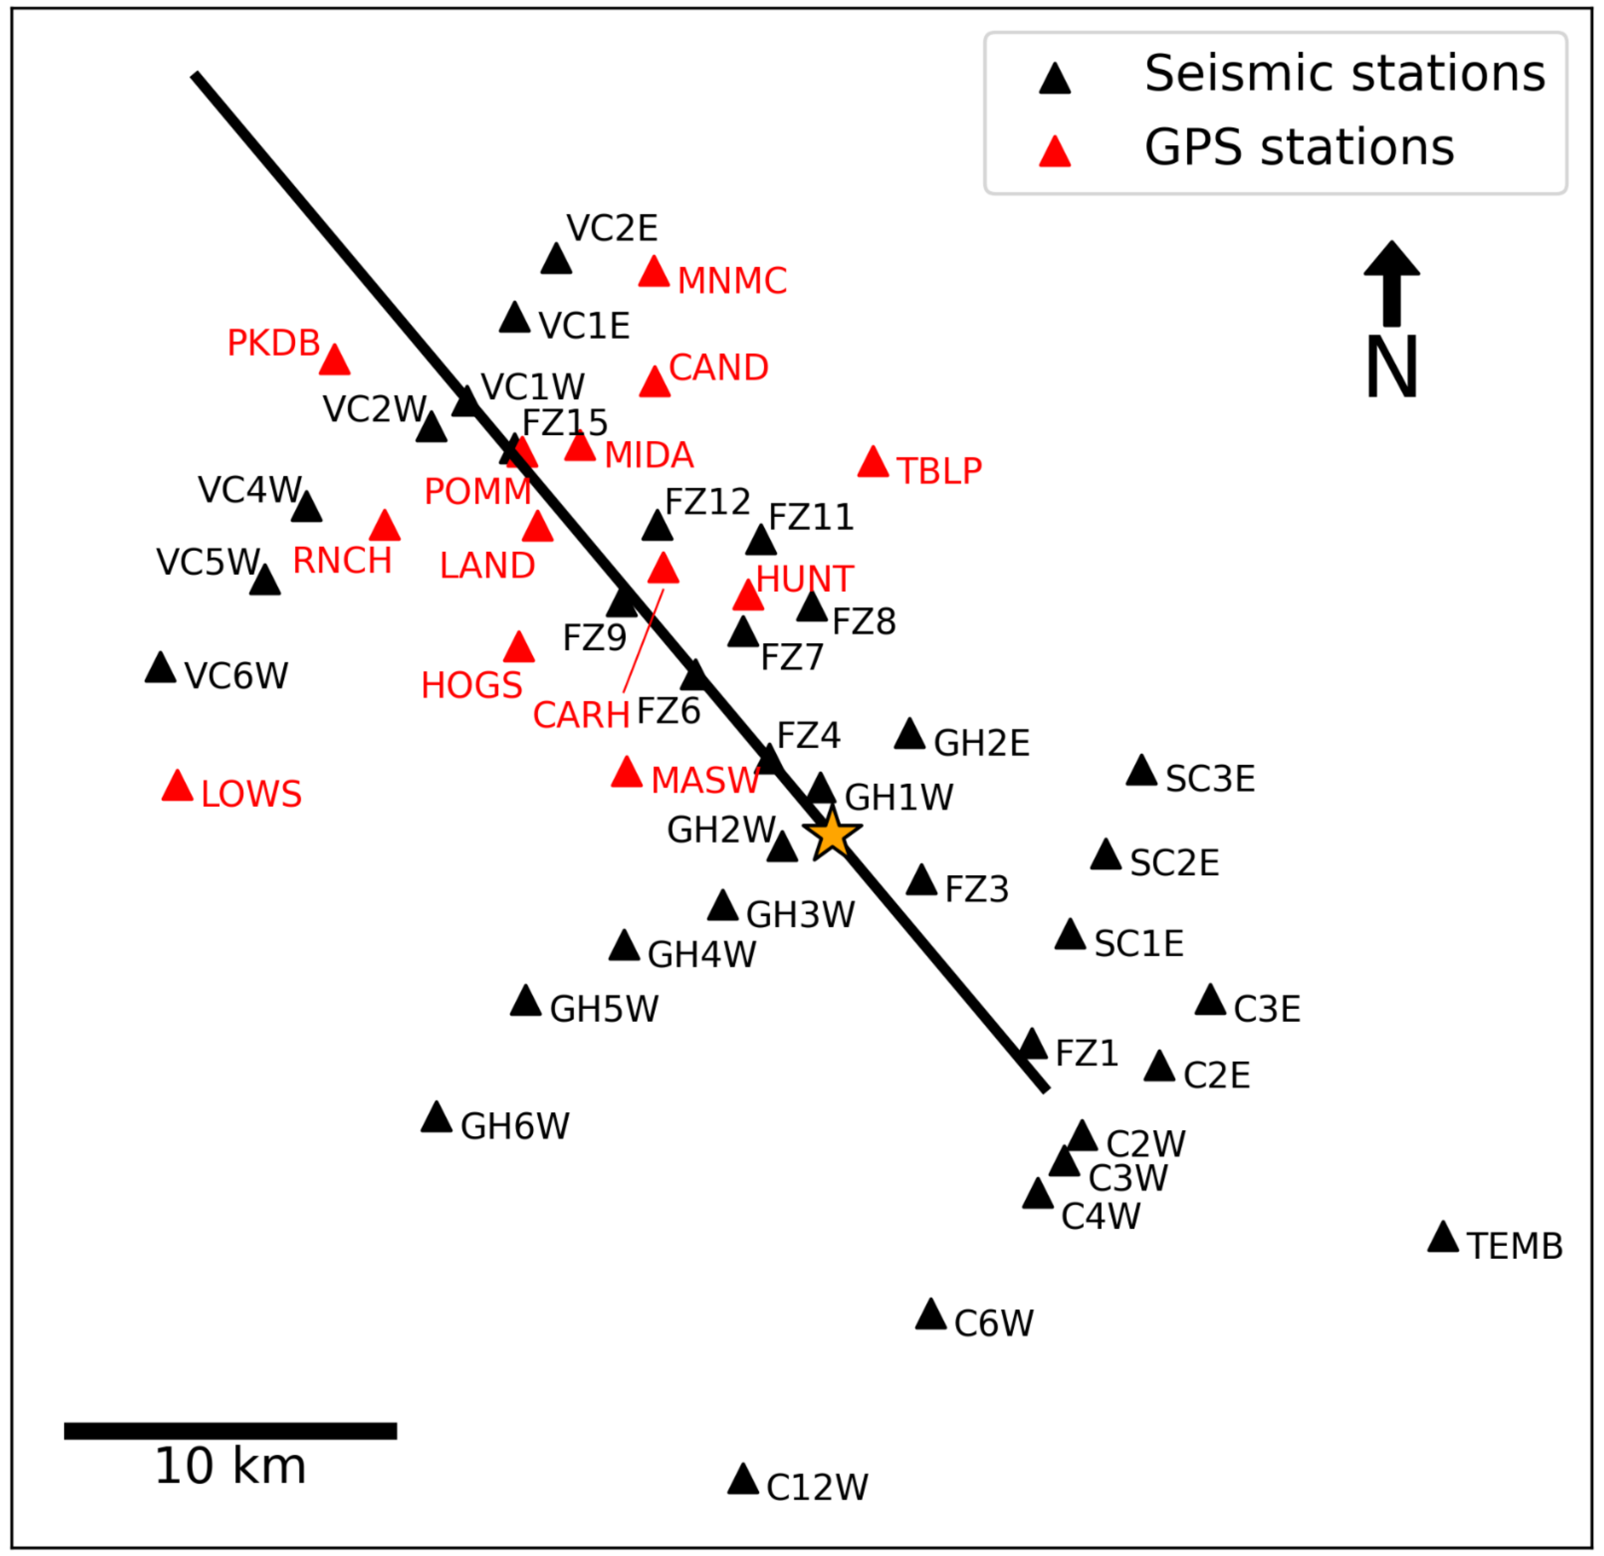

Bridging time scales of faulting: from coseismic to postseismic slip

Premus et al., Sci. Adv. ’22; Schliwa et al., in prep.

- Example: **the Mw 6.0 2014 South Napa earthquake**
- **Data:** 10 seismic broadband stations, 30 days of postseismic displacements on 7 GPS stations and surface slip measurements
- Embarrassingly parallel inversion on 3 Nvidia 2080Ti, 35 s per one forward simulation on 1 GPU ~400 000 models visited / 7500 accepted
- **Imaging complementary regions and physics of the same fault**

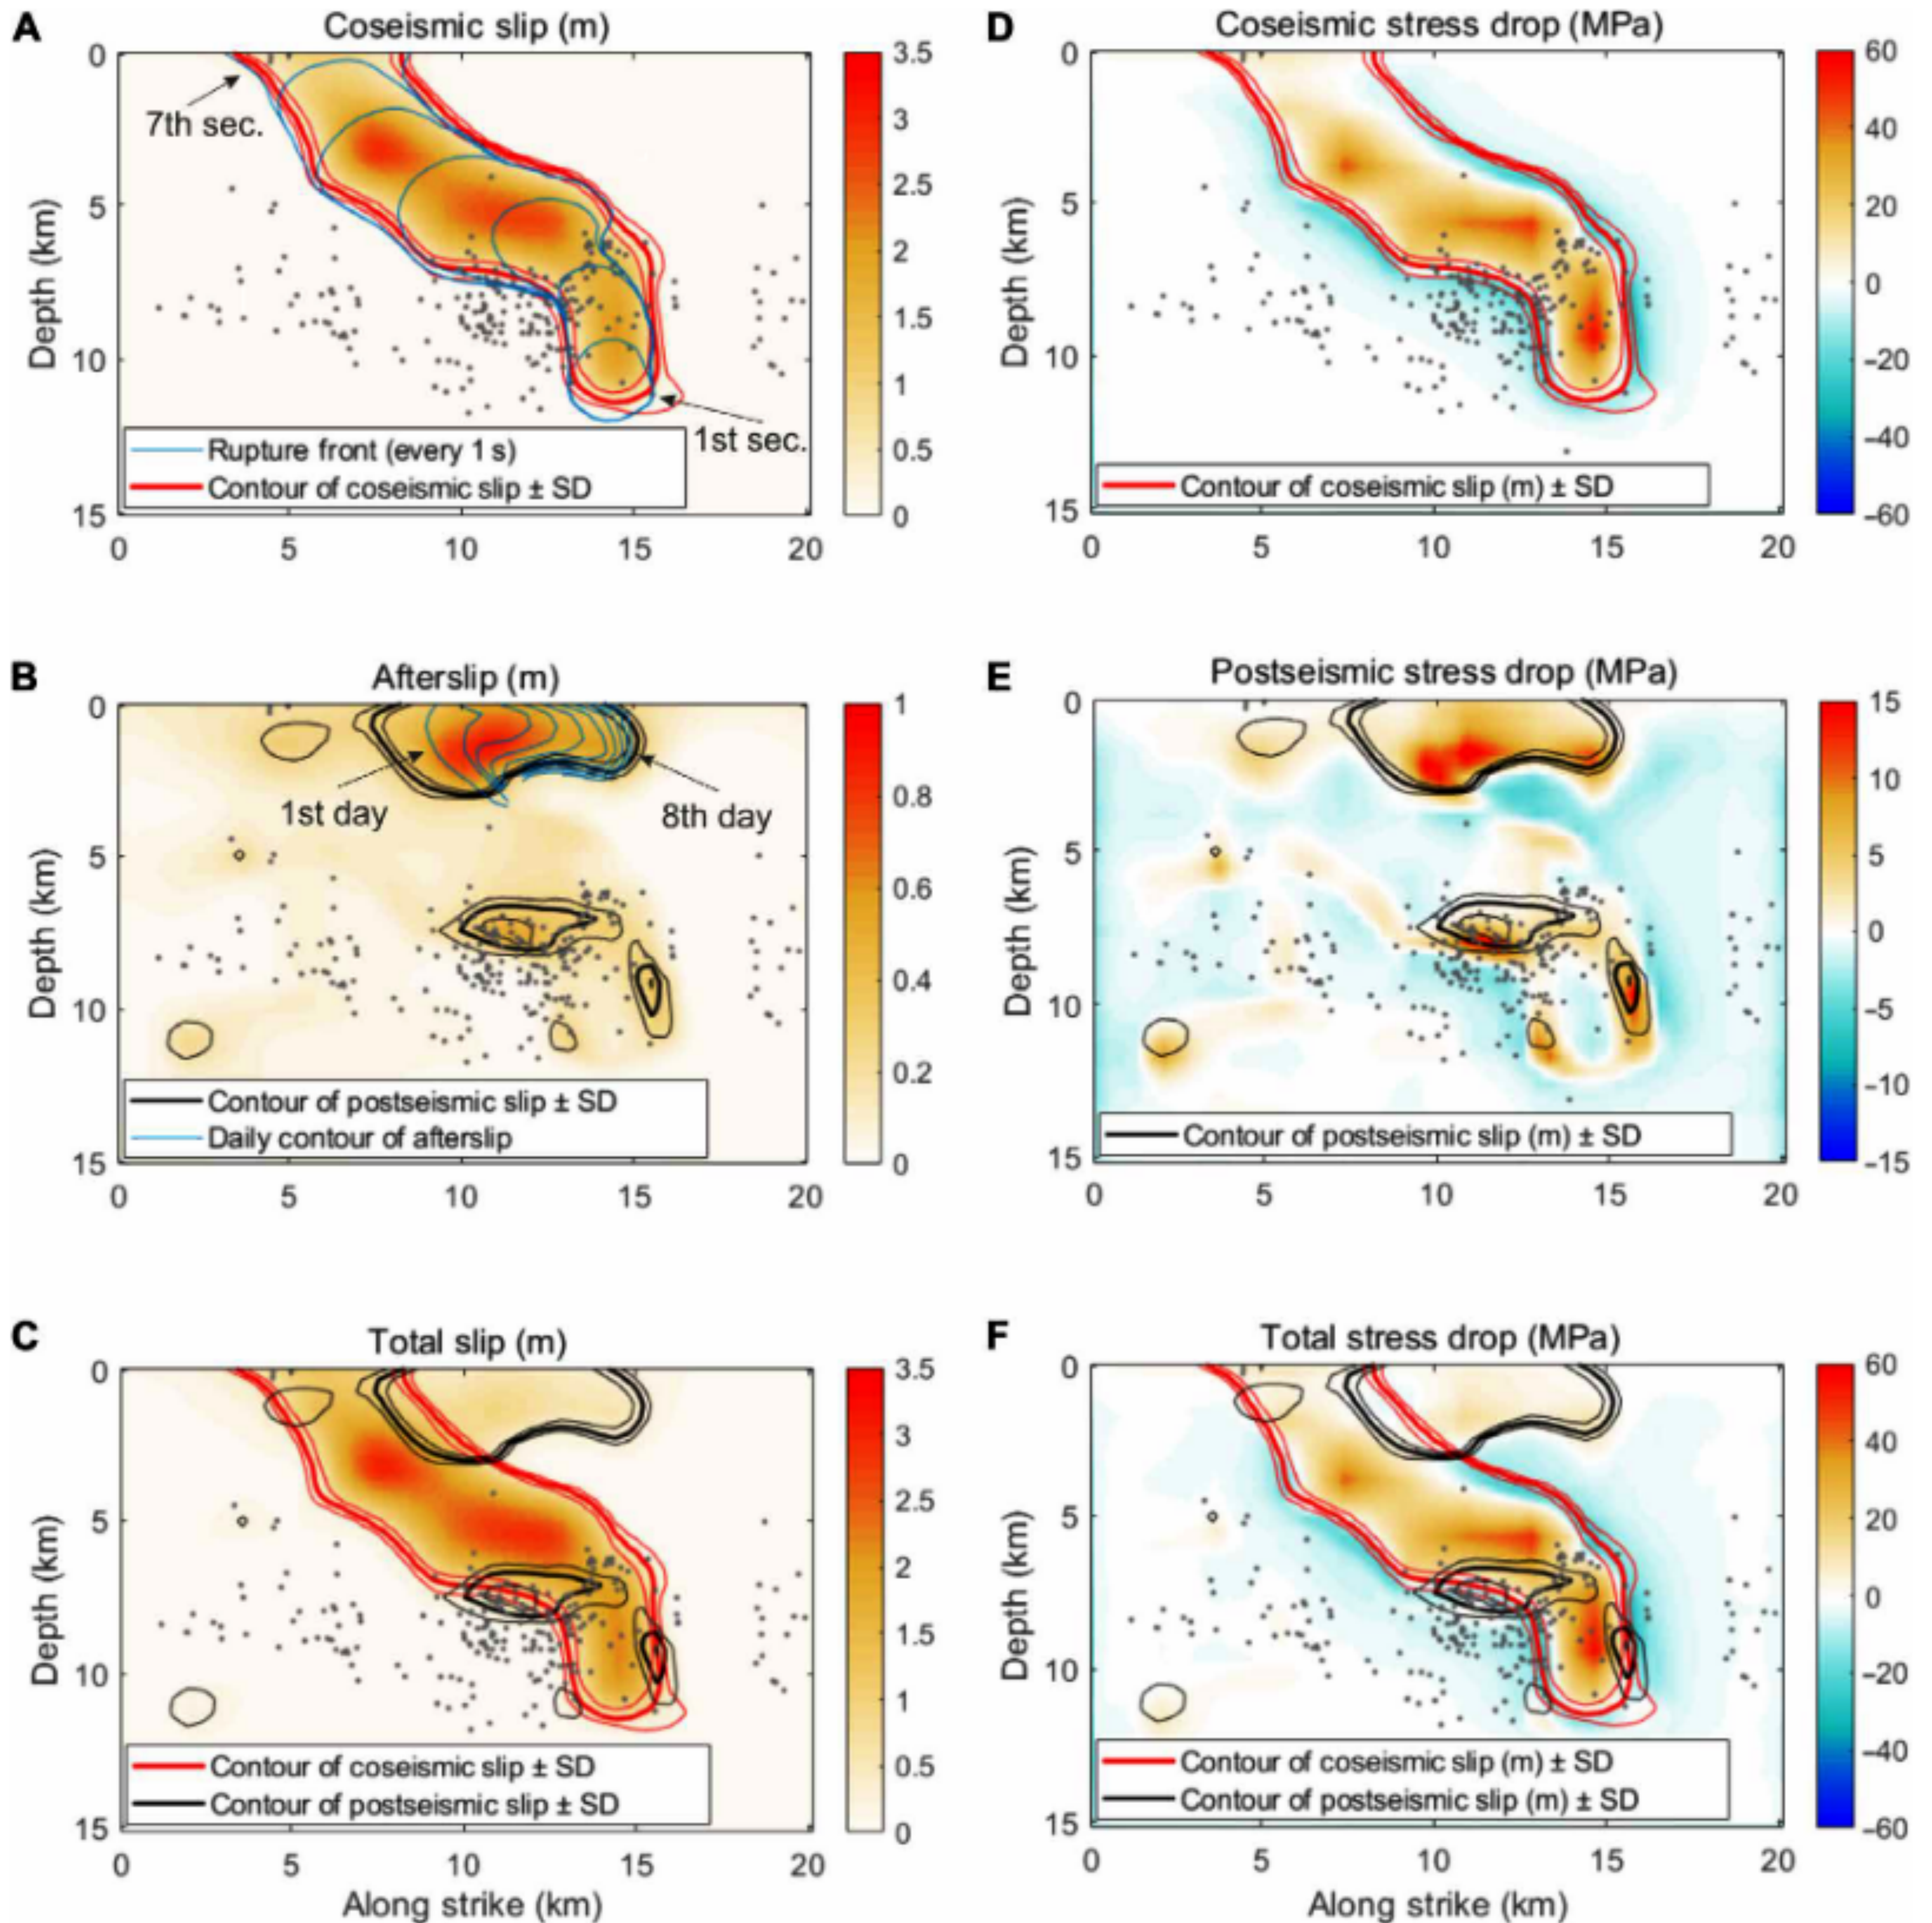

Why are geohazards more difficult than weather?

- **Global Atmospheric Models** are highly successful
- **Dynamic core:** PDEs (e.g., mass conversation) that the community agrees on
- “**Sub grid” physics:** (e.g., chemistry) where all dirty tricks are allowed
- **Digital Twin:** Combining rich, continuous and interdisciplinary data-sets in a core dynamic framework **makes hypothesis testing is easy**

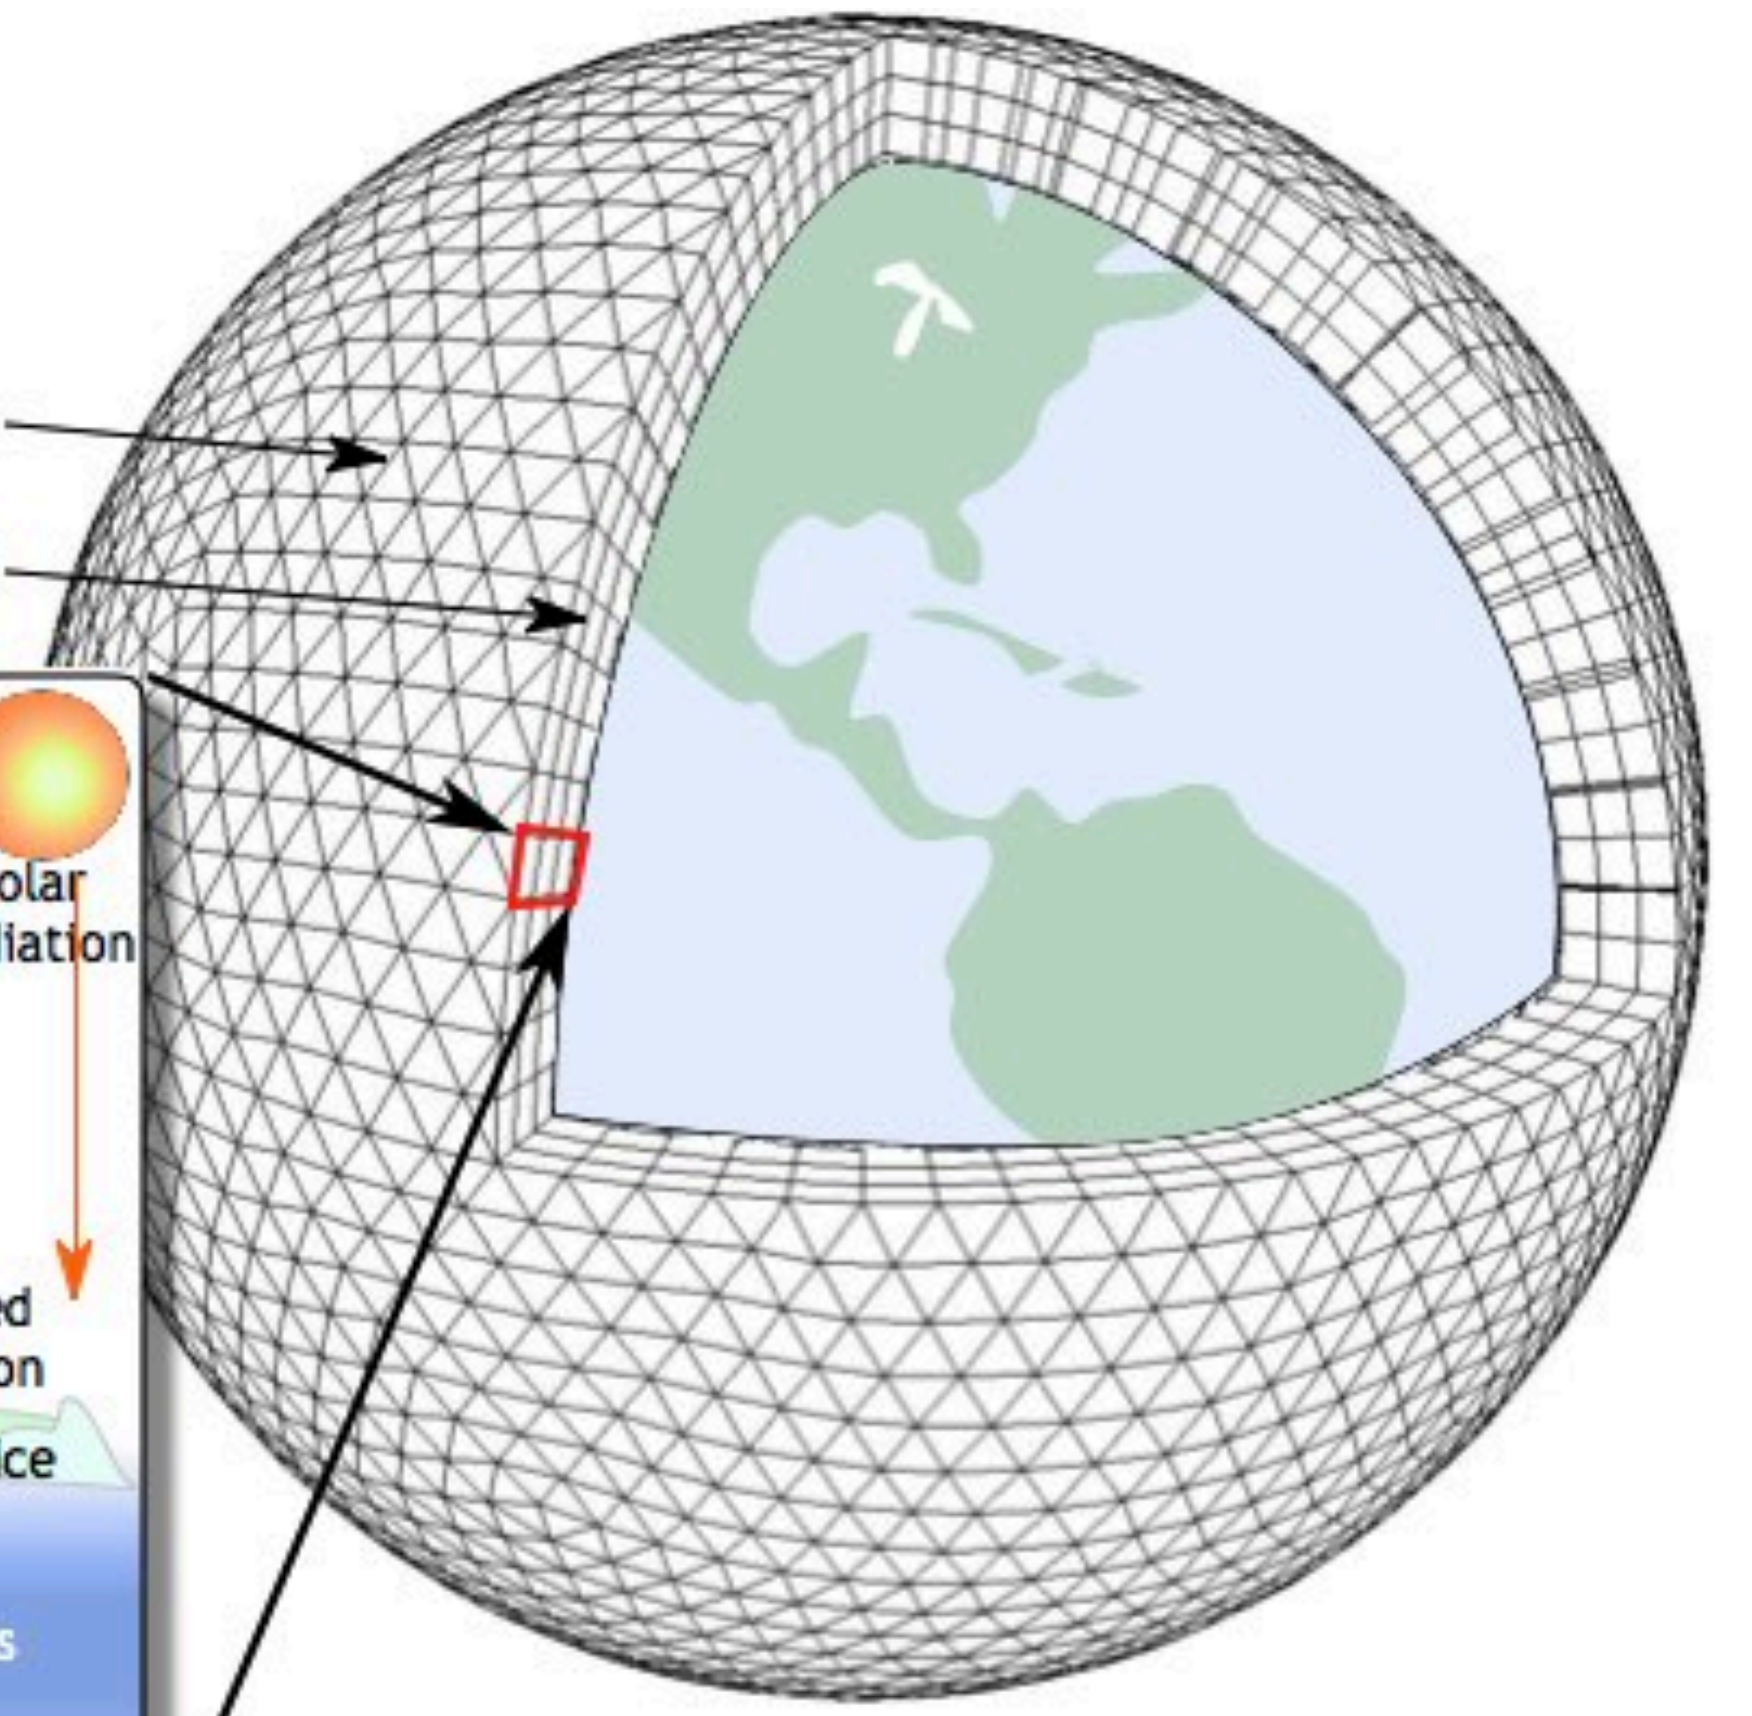

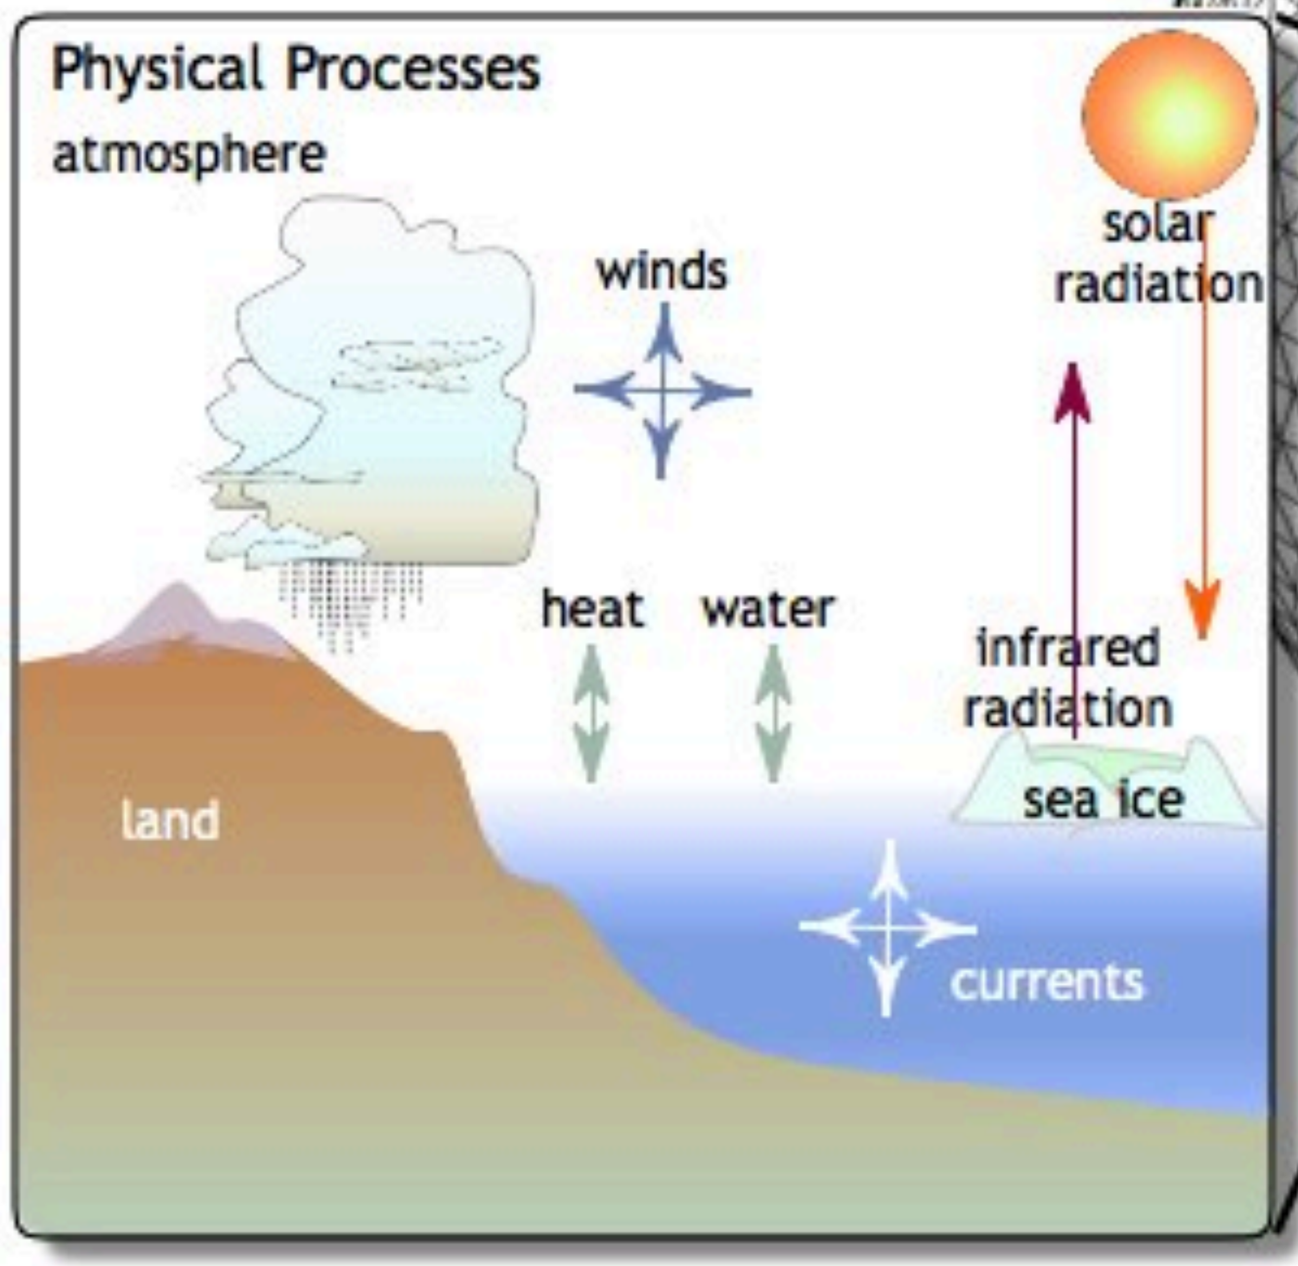

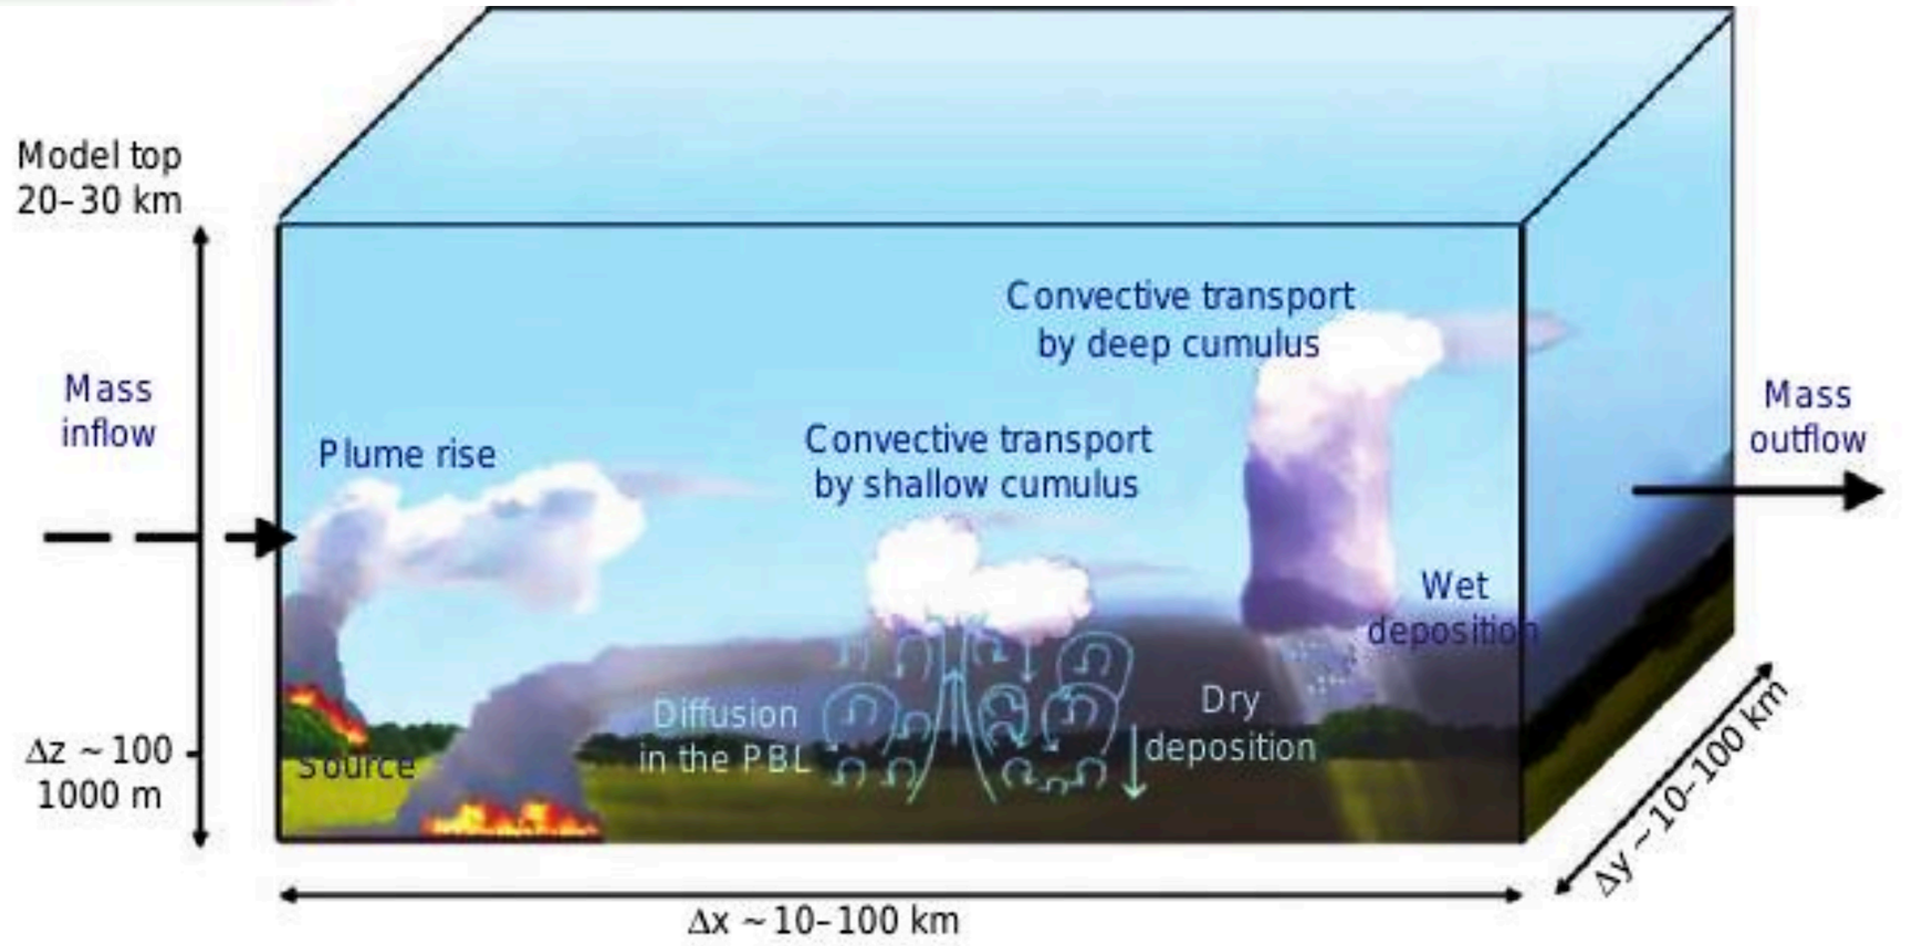

Why are geohazards more difficult than fire?

- **Short-term and long-term predictability of wildfires** is based on **data assimilation**, i.e. is a function of data availability

- **Seismological data streams** are sparse in space, but more importantly **sparse in time, challenging data assimilation**

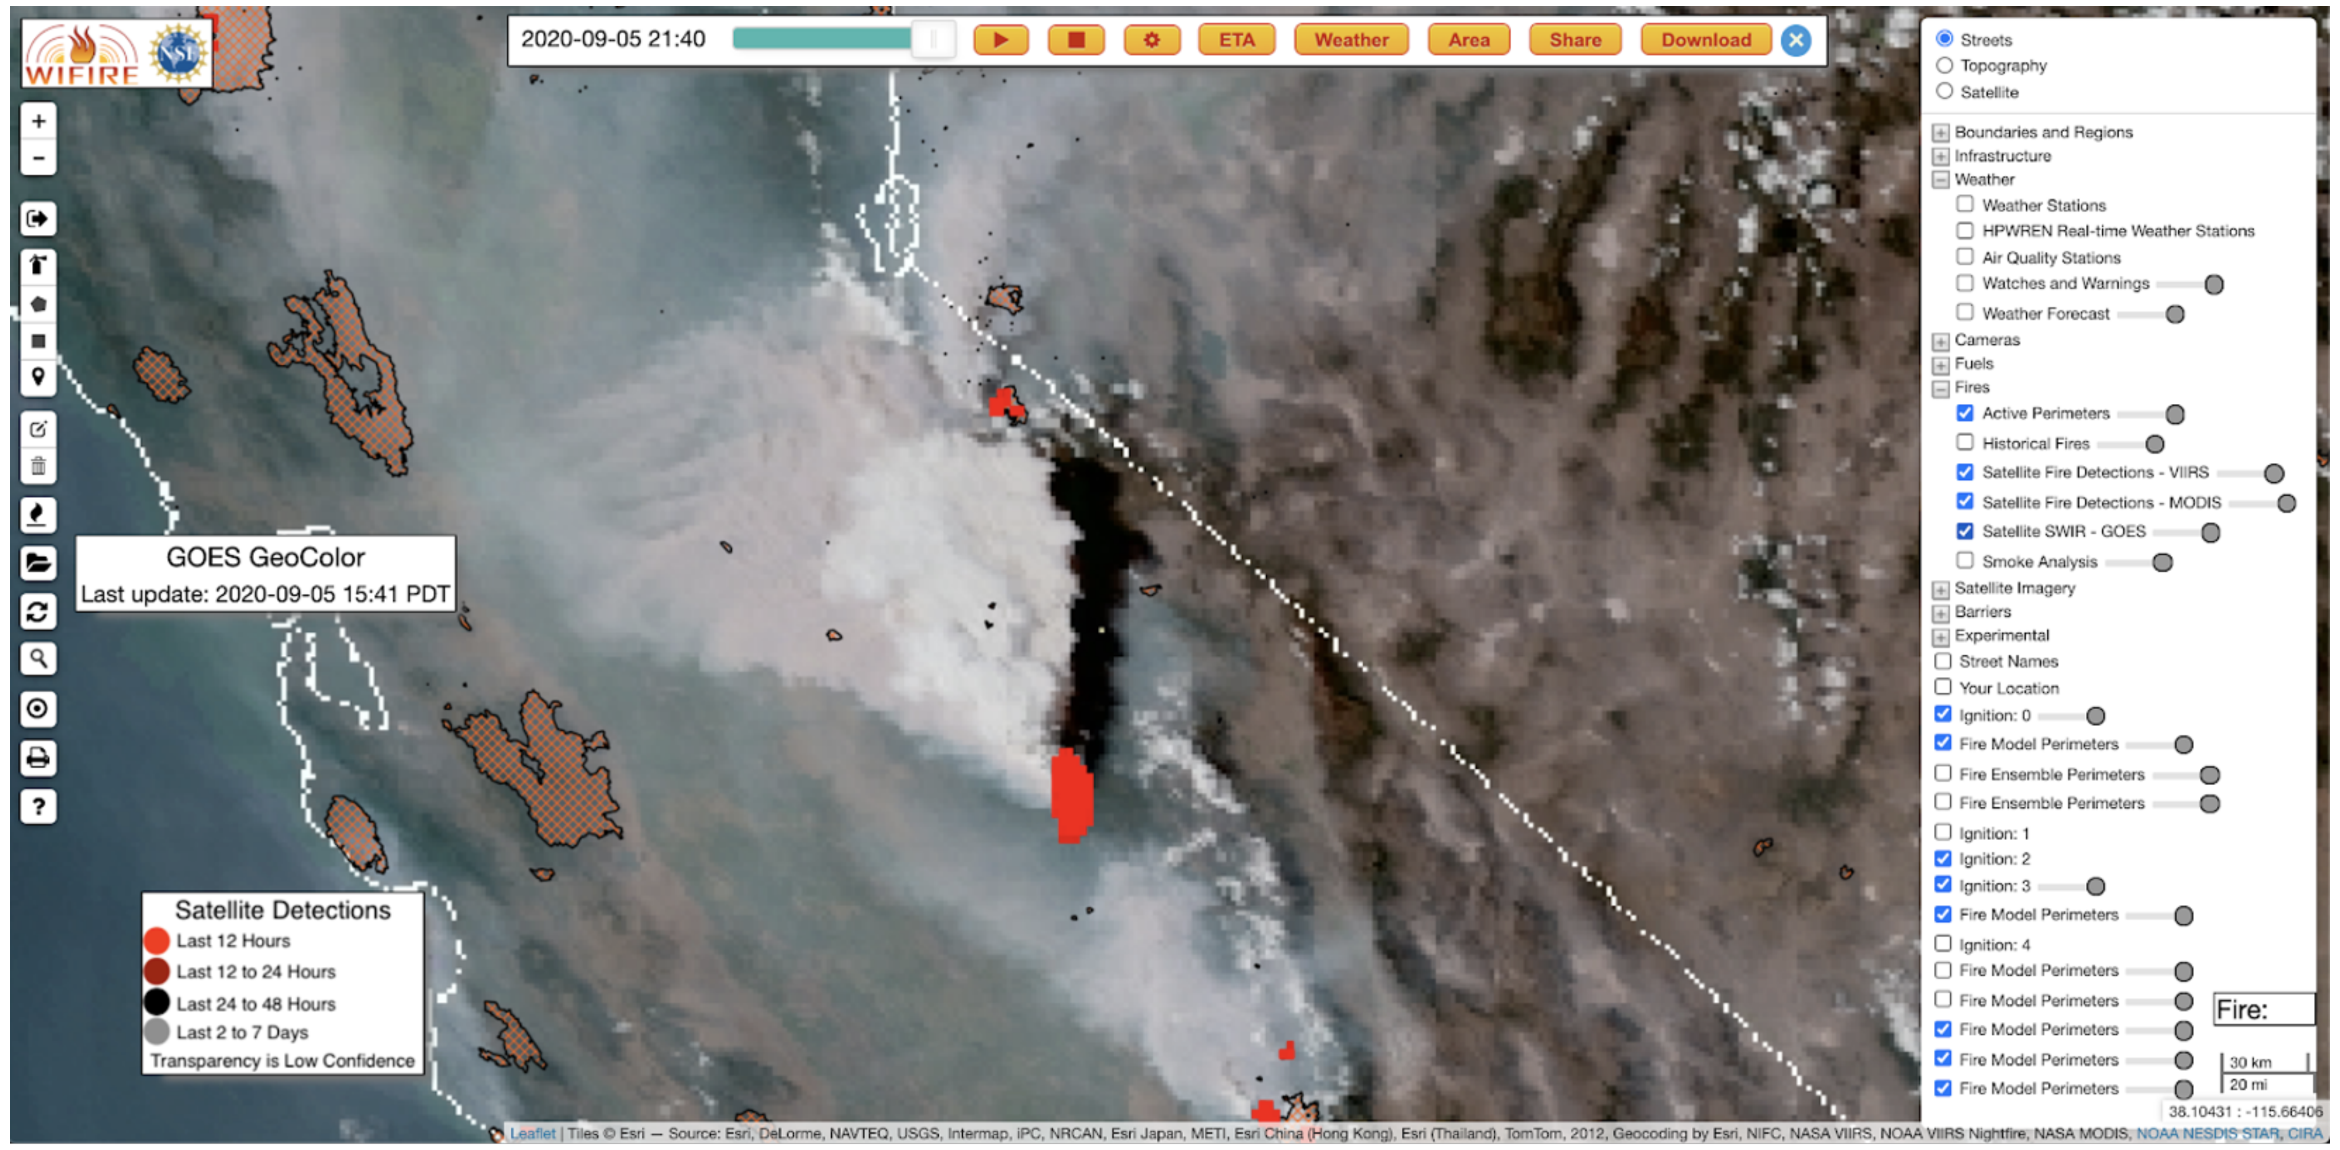

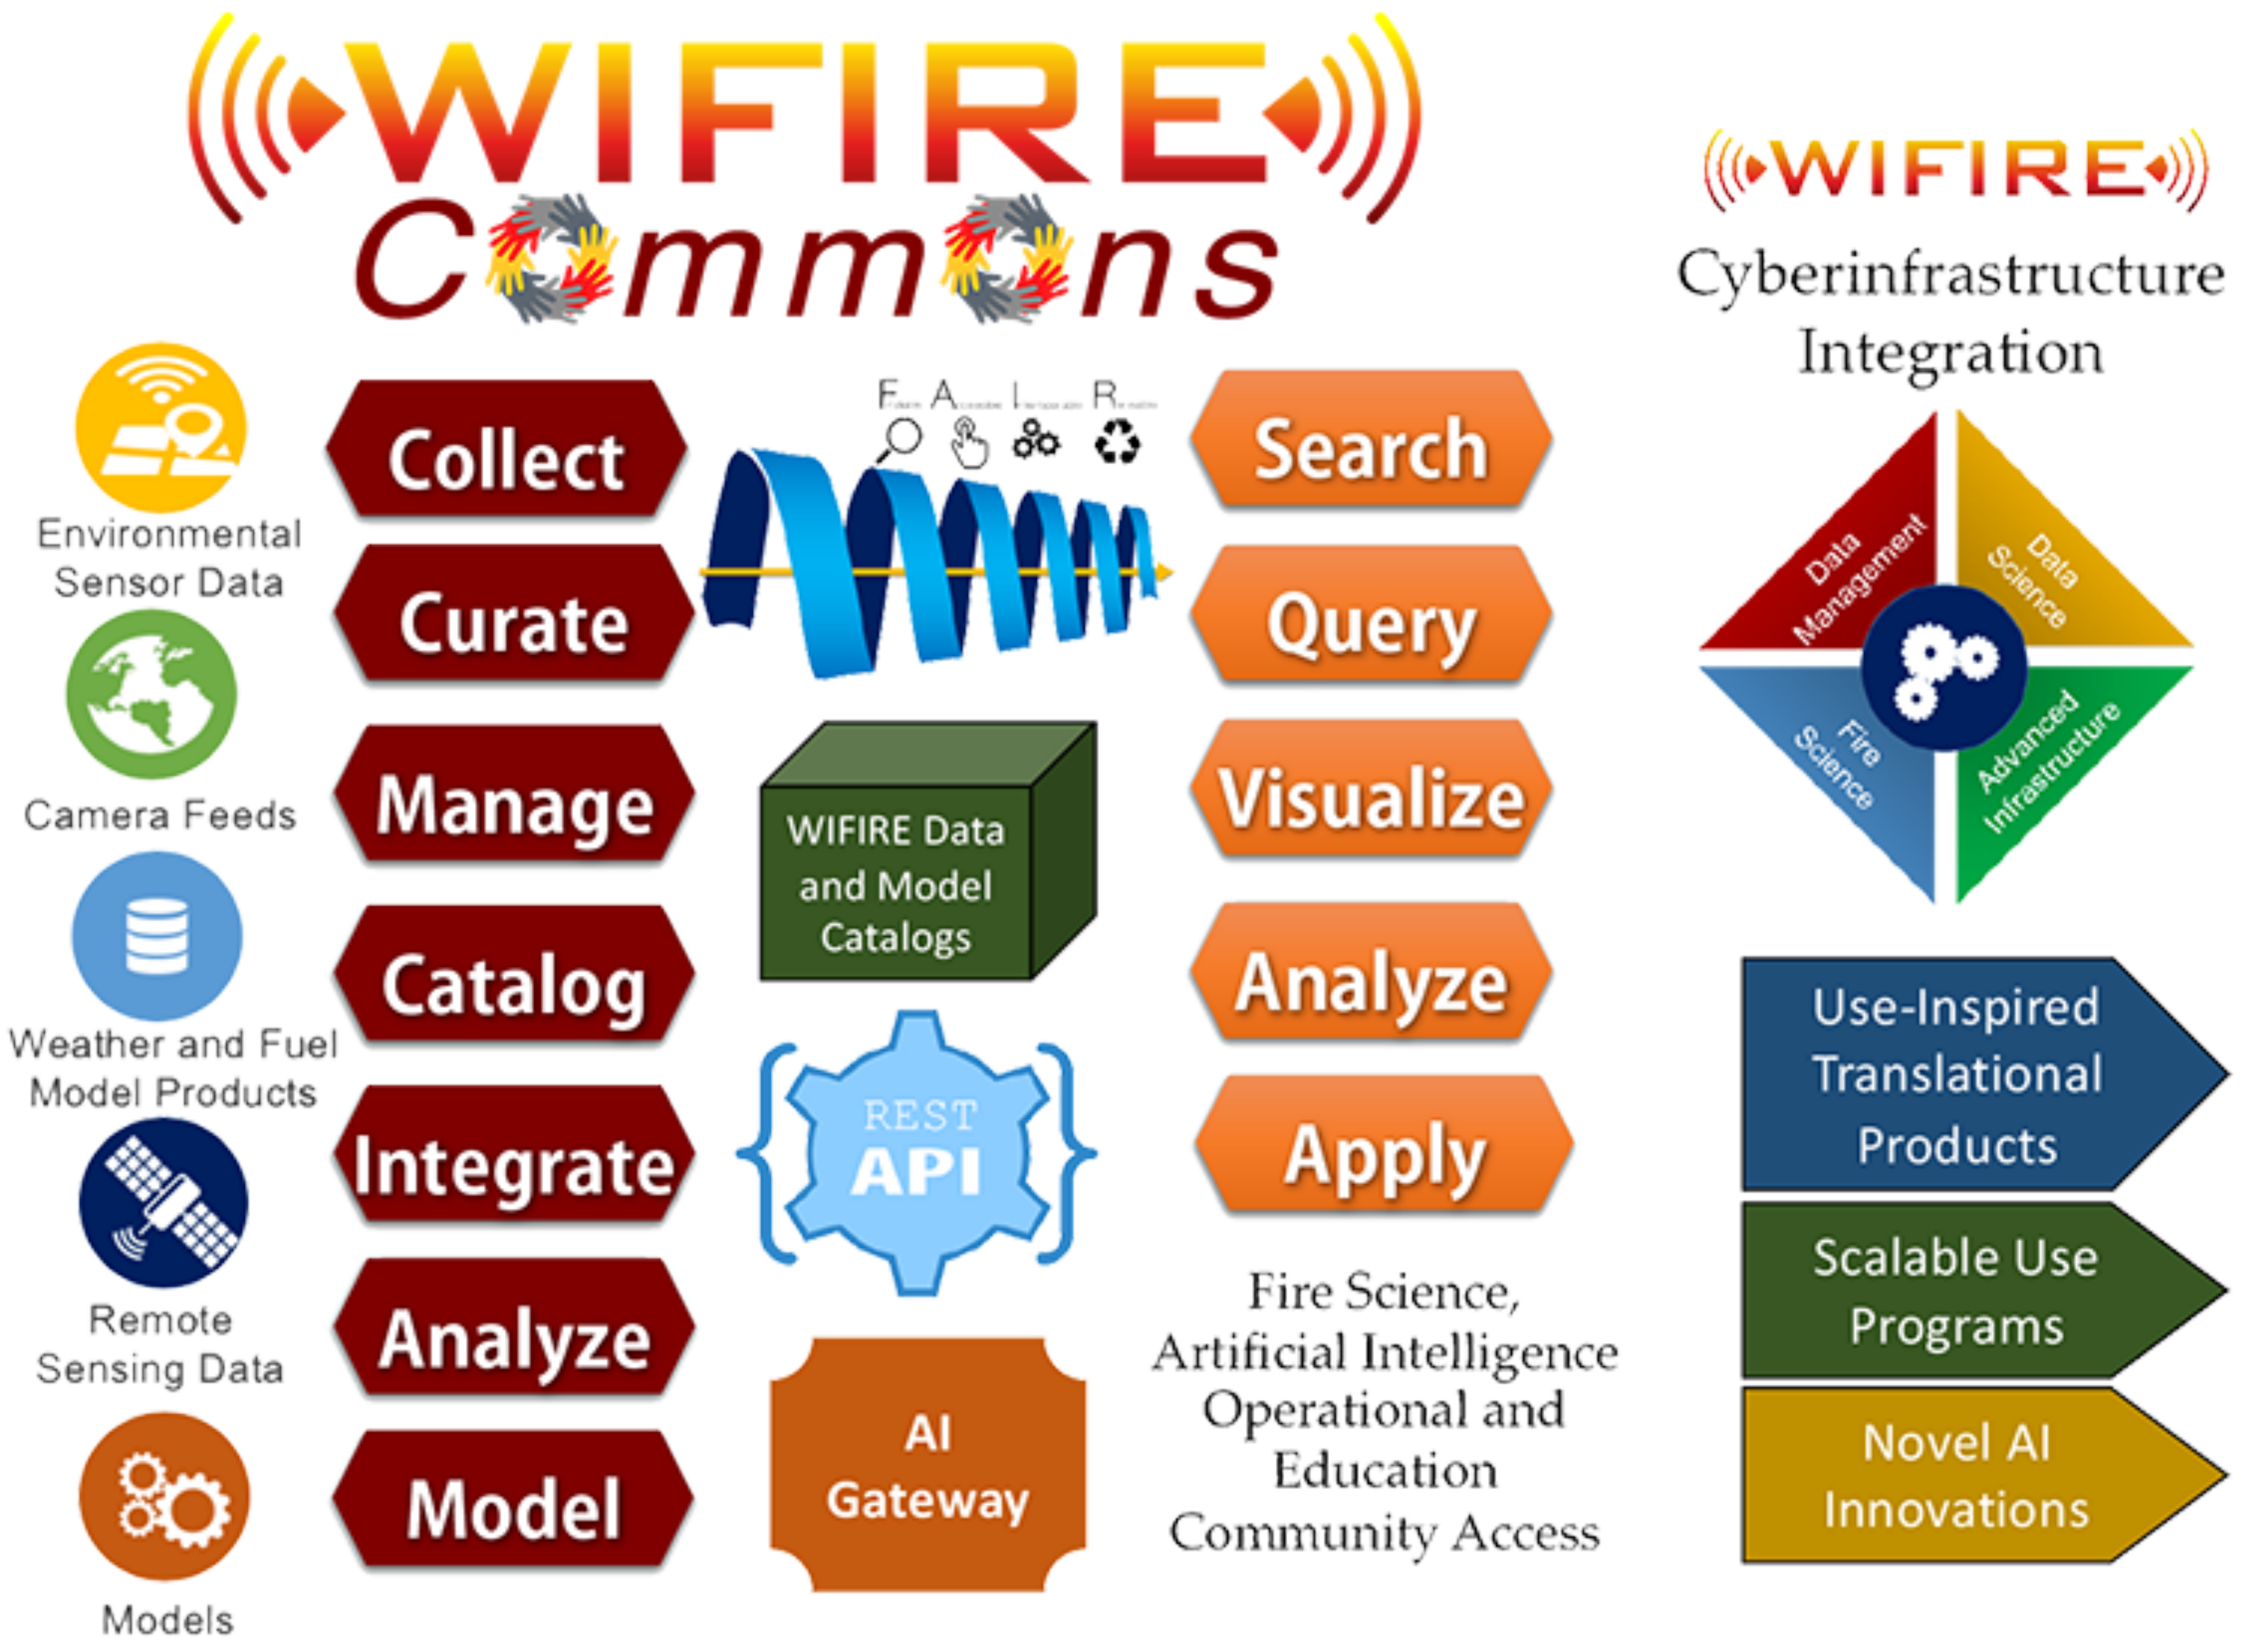

Wifire Commons, real time wildfire forecasting at SDSC,

Altintas et al., 2015

On the path to *exascale** supercomputing ..

** 10^{19} floating point operations per second*

given we properly address vectorization, energy limits, specialised/ heterogenous cores, optimal accuracy per degree of freedom

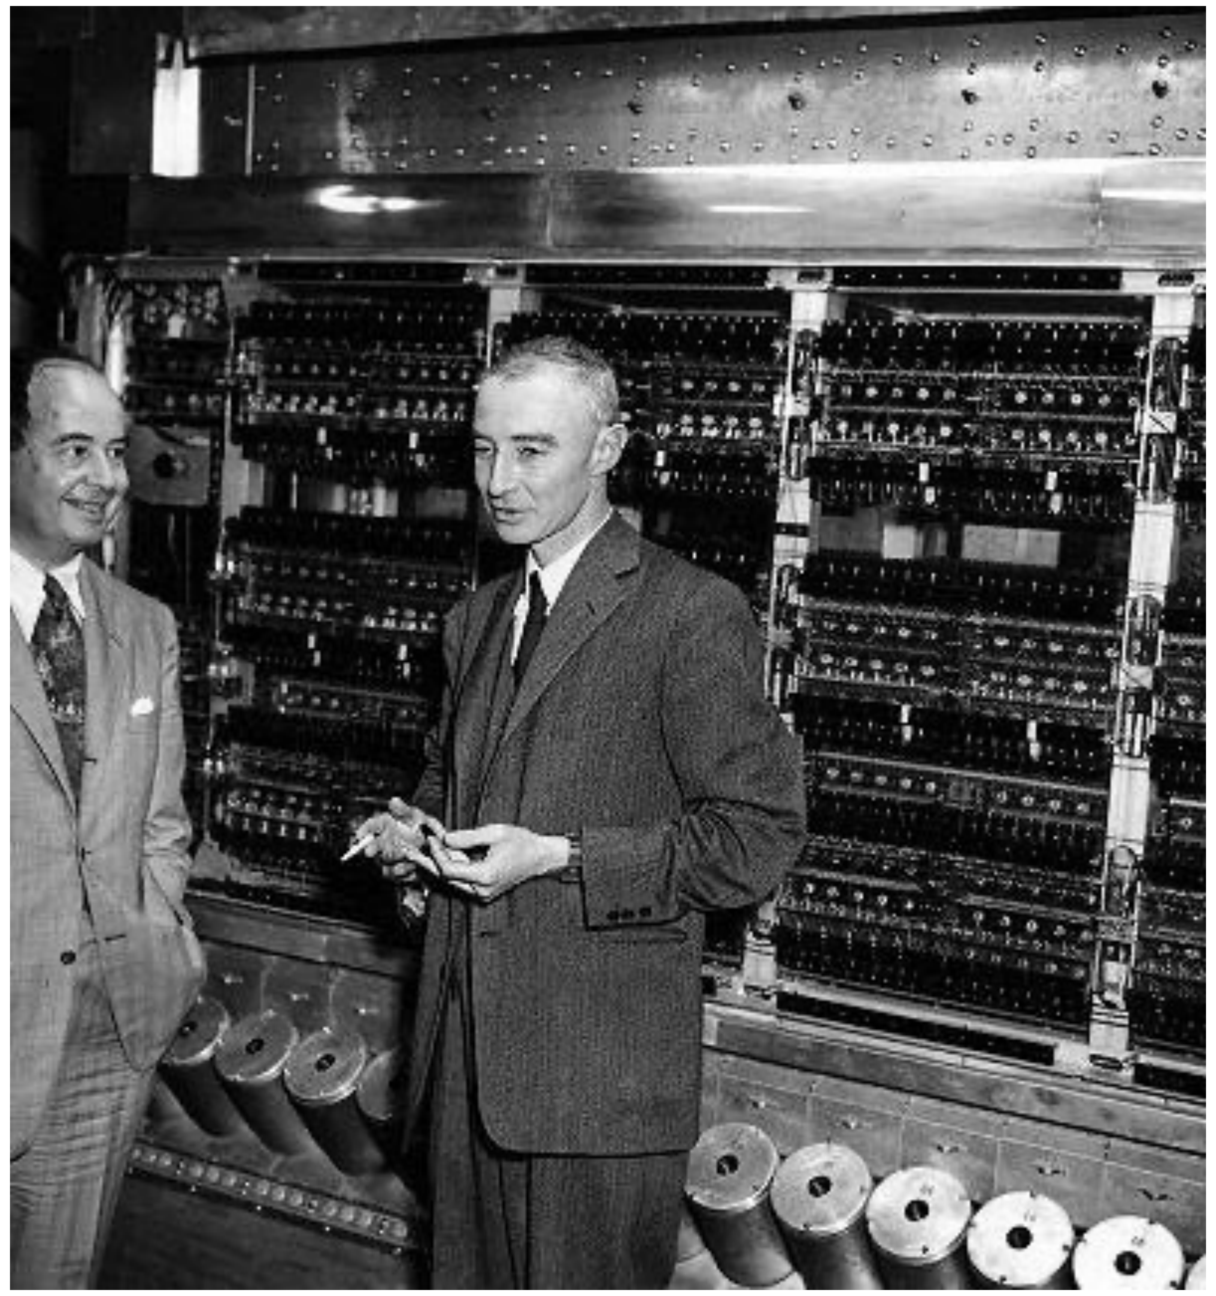

John von Neumann and Robert Oppenheimer, Princeton, IAS, 1957

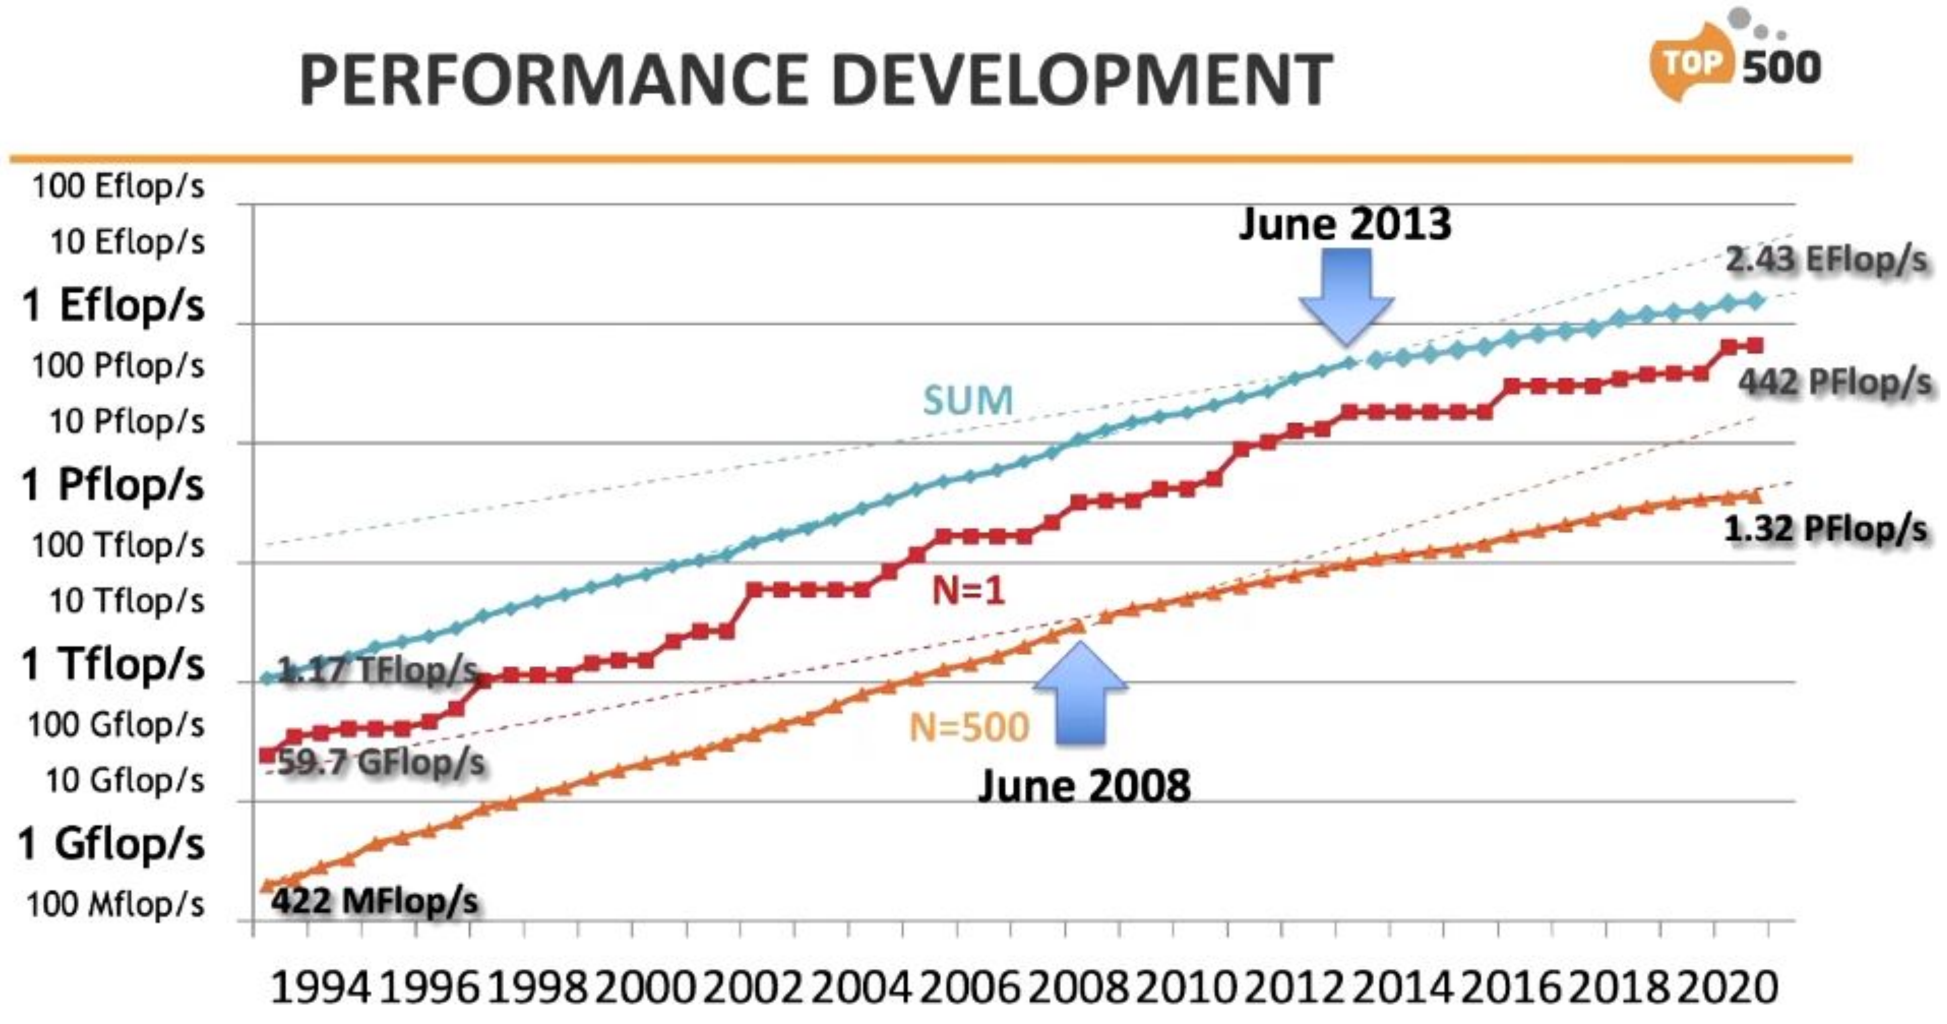

(Top 500, SC'20)

Field work, TACC, Austin, 2021

EUROPE CLEARS PATH TO 2023 EXASCALE SUPERCOMPUTER

January 4, 2022 Nicole Hemsoth

Two systems in China achieved 1.3 ExaFLOPS peak performance and around 1.05 ExaFLOPS (or higher) sustained performance in Linpack benchmark in March 2021

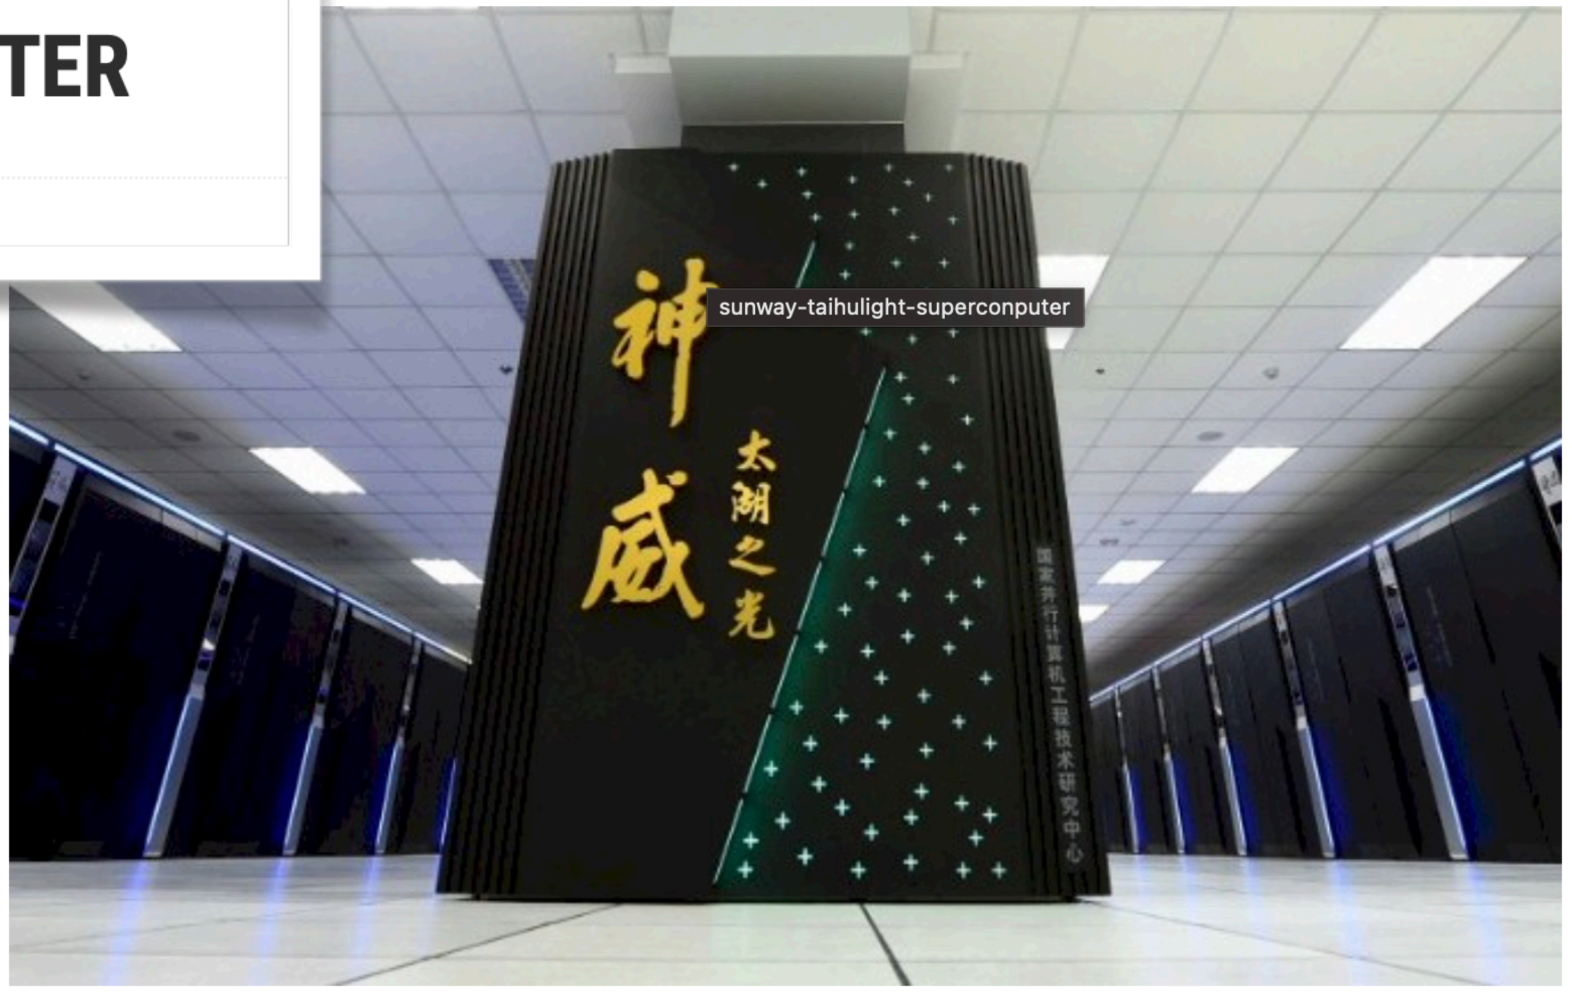

On the path to *exascale** supercomputing ..

*** 10^{19} floating point operations per second**
given we properly address vectorization, energy limits, specialised/ heterogenous cores, optimal accuracy per degree of freedom

.. and to new, dense earthquake *observatories**

if funded

RUFZO, SZ4D, CRESCENT, IPOC, DONET, … :
Dense **arrays with a mix of instrumentation** ranging from:
- Ocean bottom pressure sensors
- Strong ground motion
- cGNSS
- Creep, strainmeters
- Cameras

e.g., DONET and S-net
- 5800 km of seafloor cable
- 150 nodes (30 km spacing).
- Each node:
  - 2x Pressure sensor
  - 4x accelerometers

Entering a golden Age of Geodetic Imaging Data
- Sentinel-1 and NISAR

Slide adapted from Chris Milliner

RUFZO – near fault dense observatory (proposed)

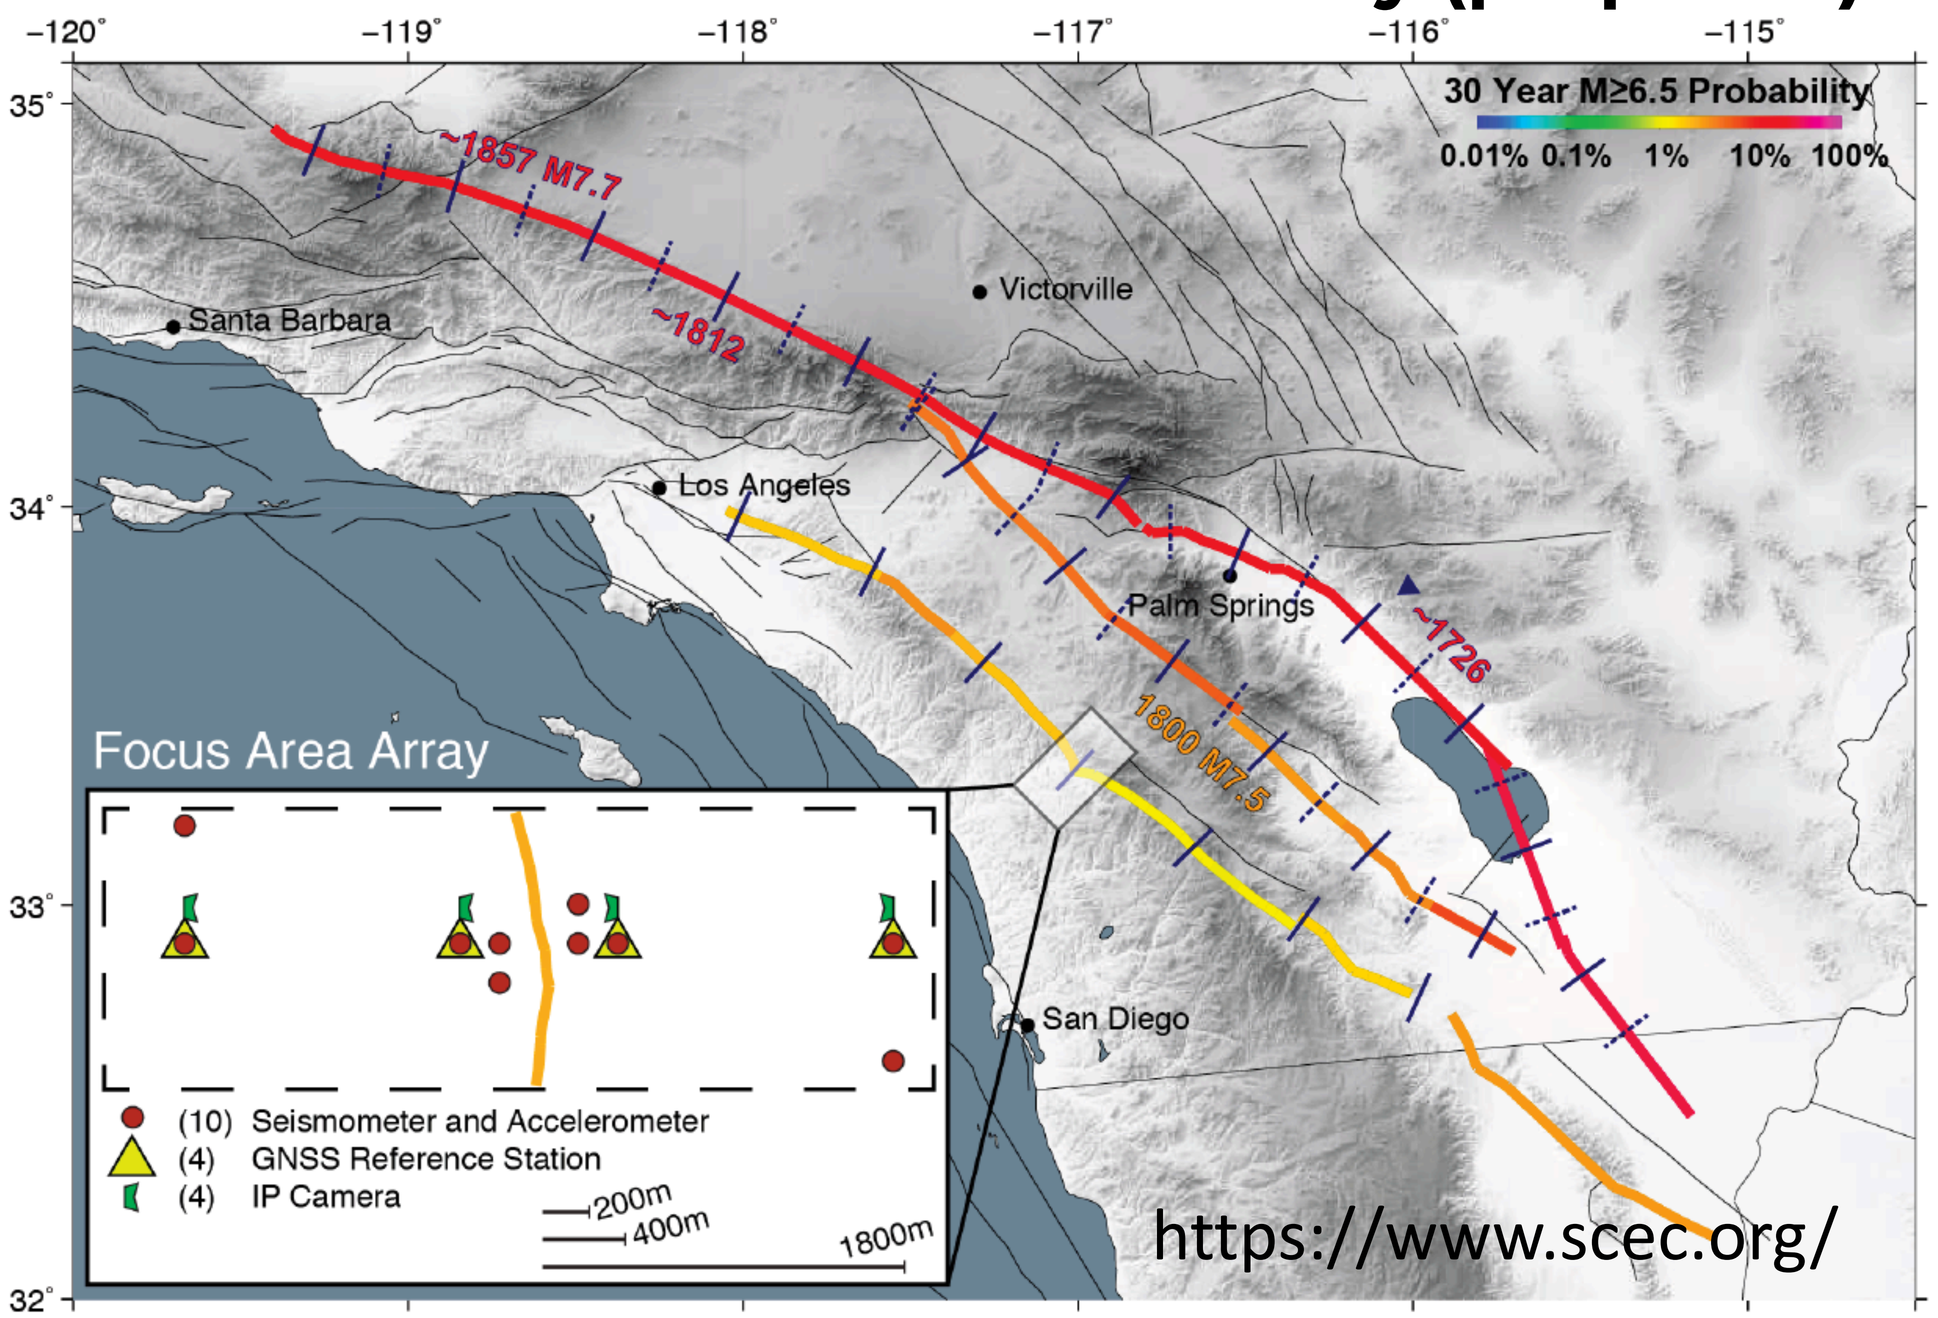

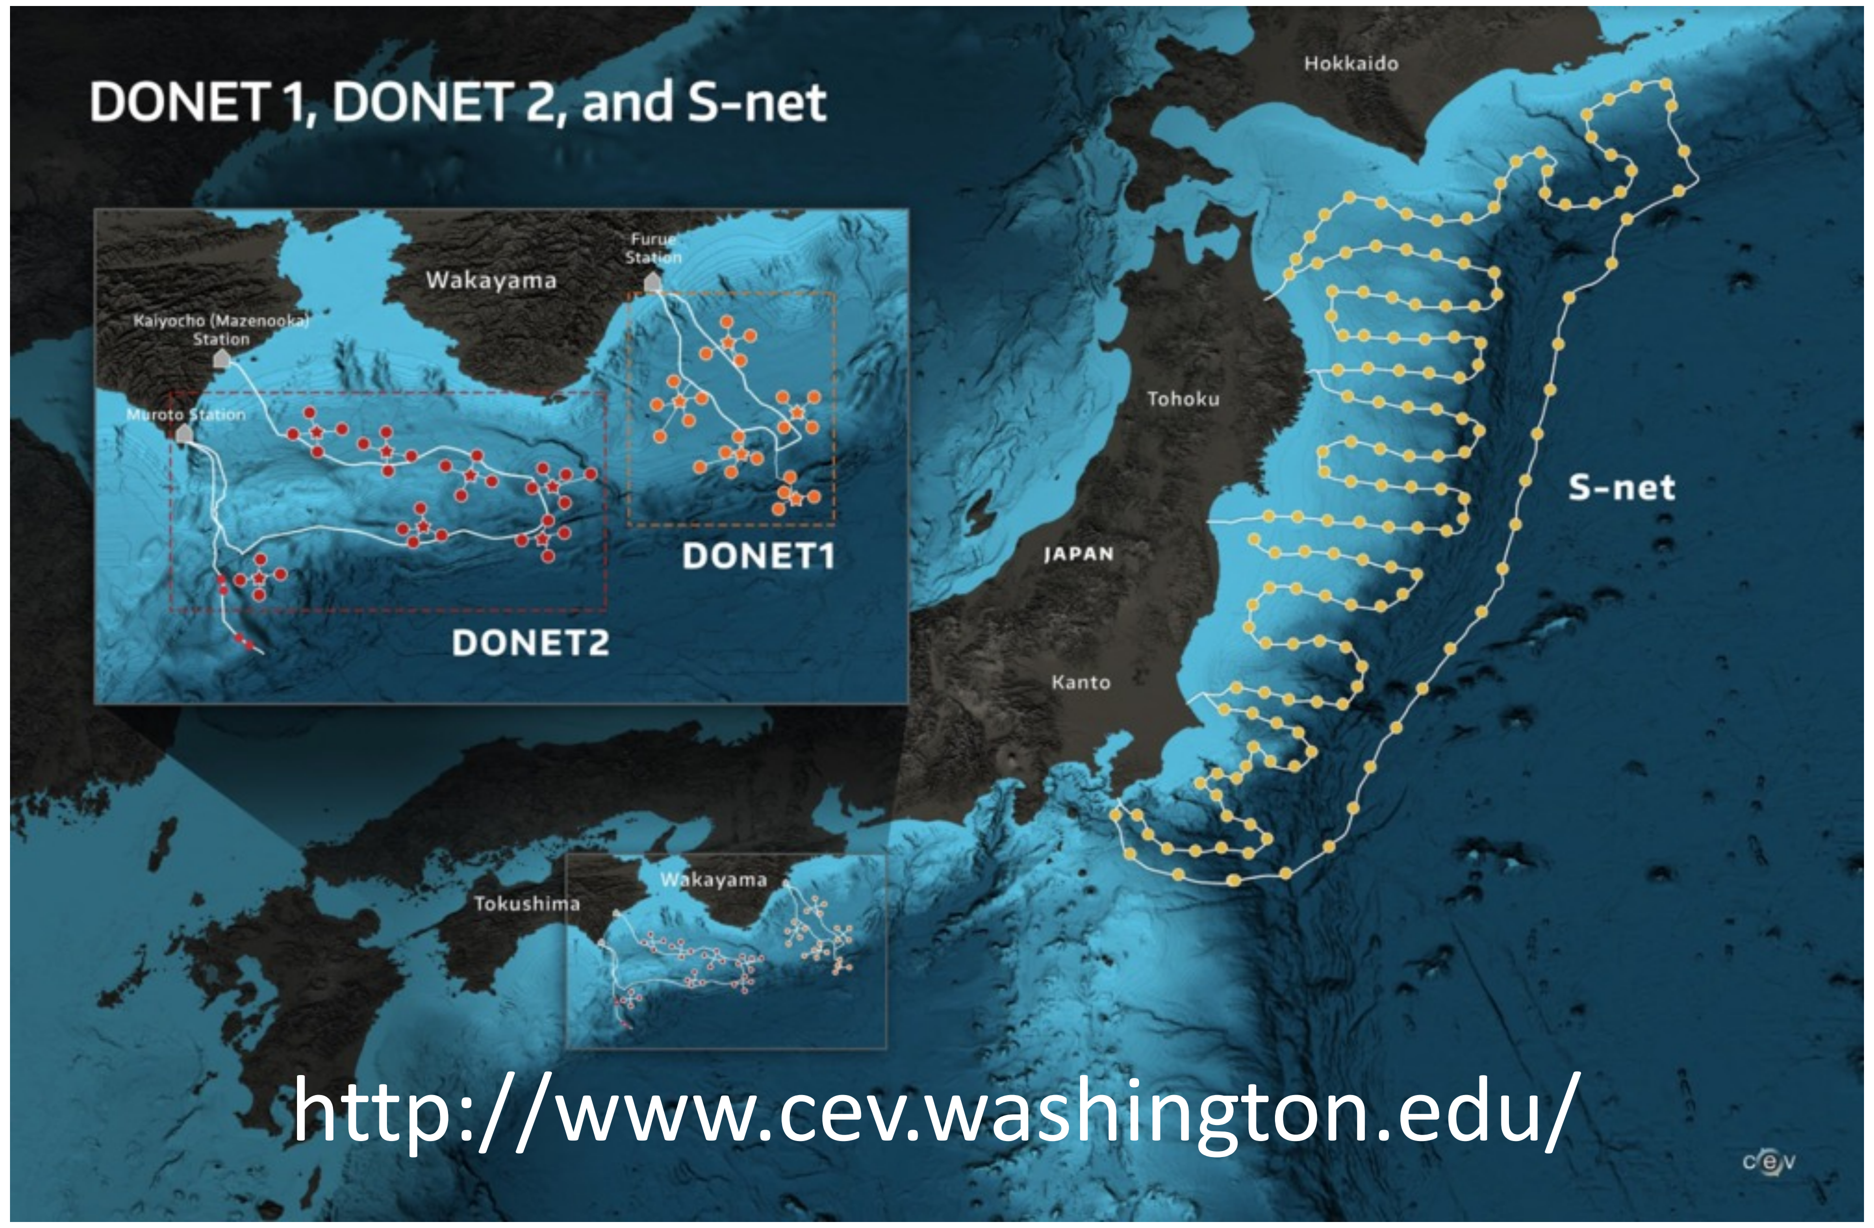

SZ4D – (proposed)

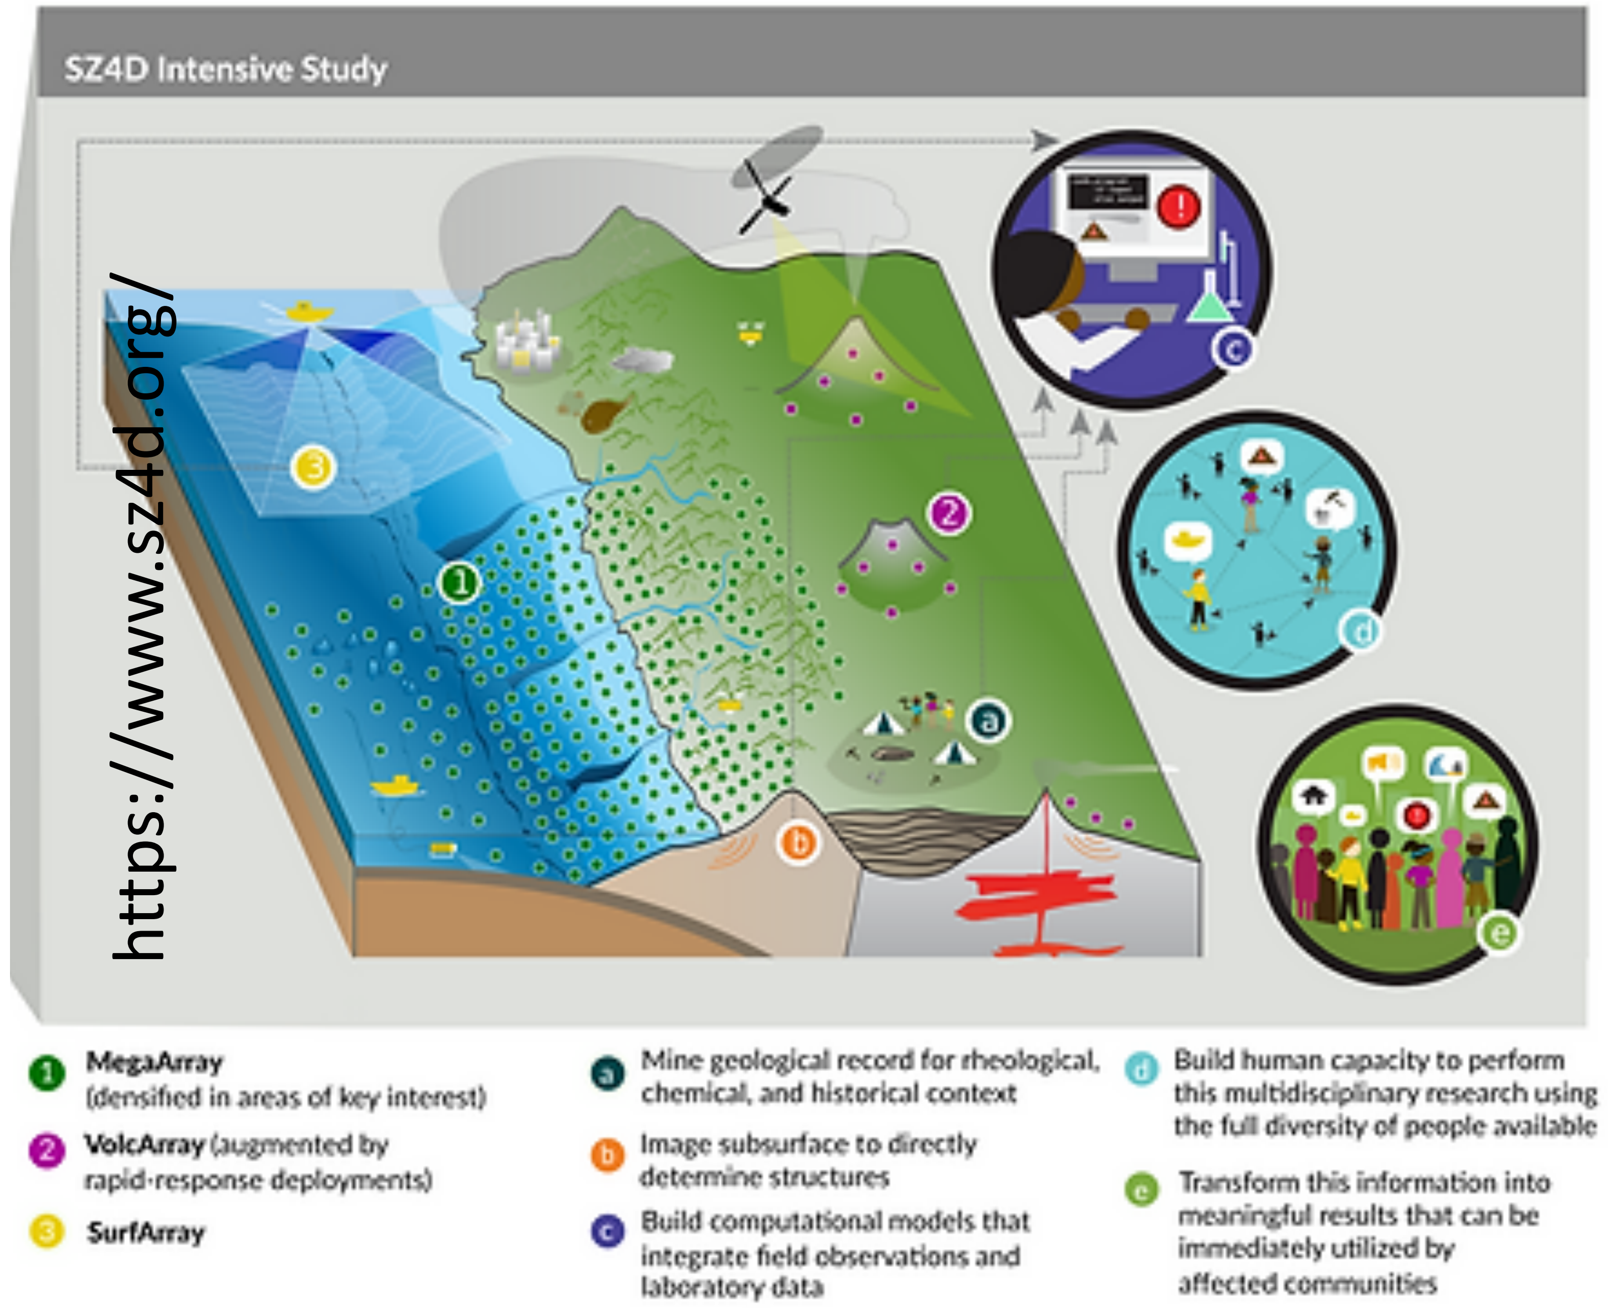

IPOC

https://www.ipoc-network.org/

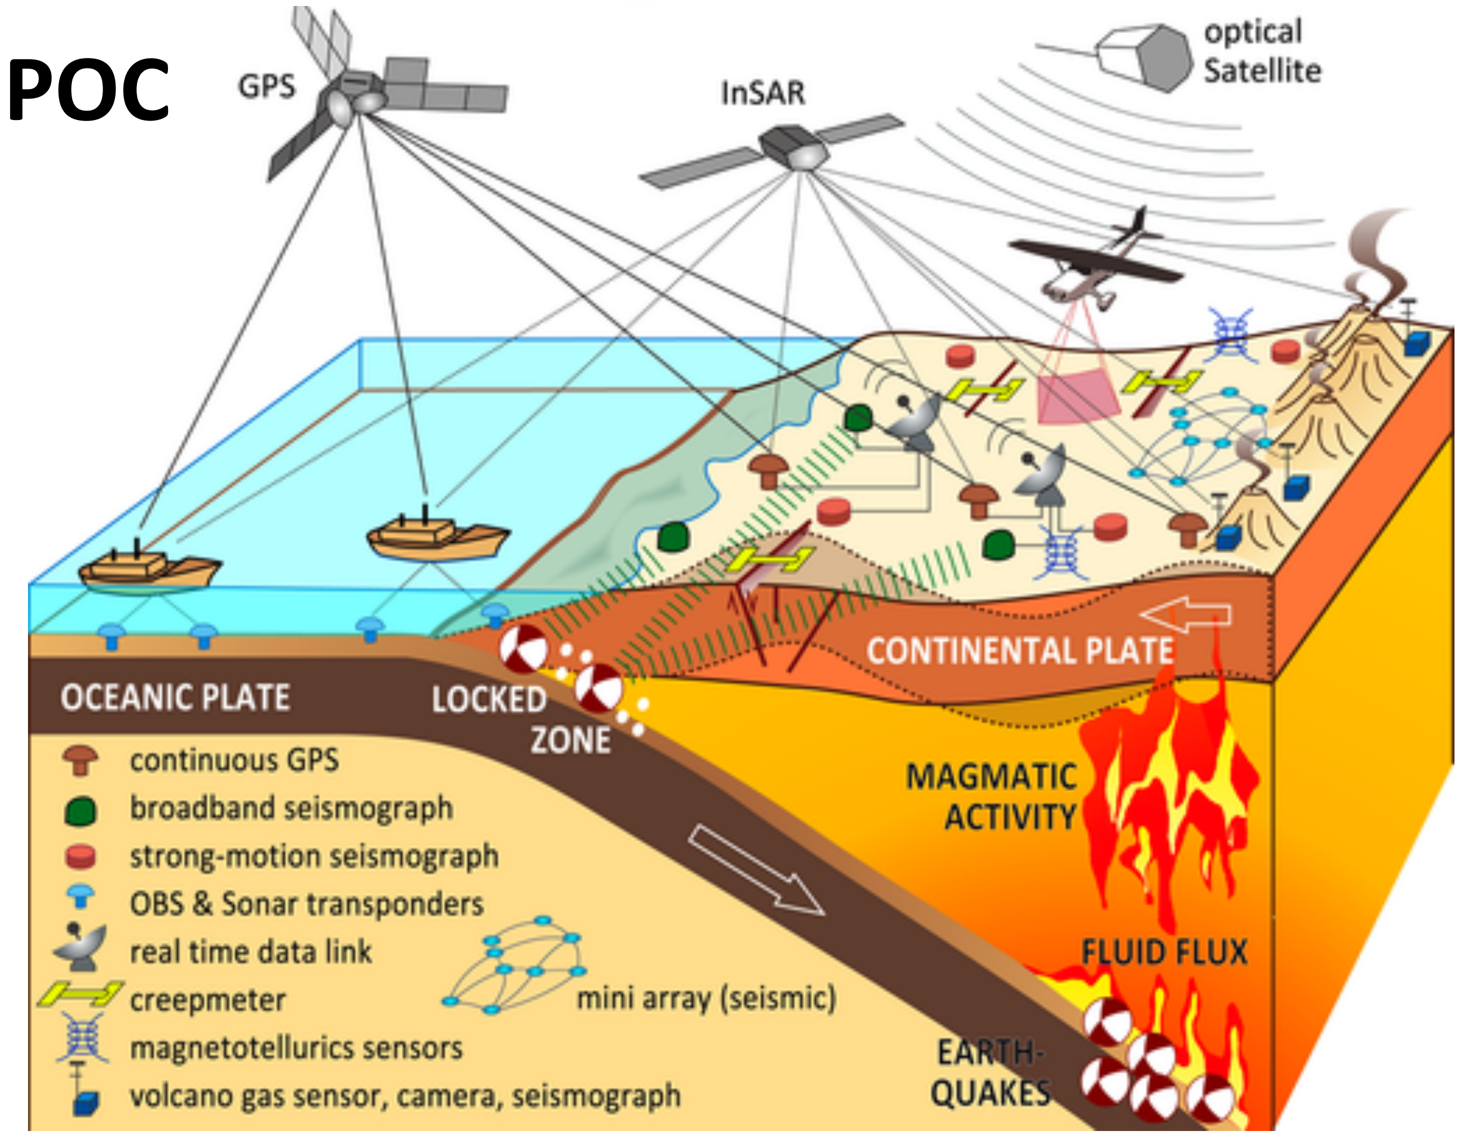

Integrated Plate Boundary Observatory Chile

What is the future?

1. **Digital Twins (no more hero runs?).** **Hybrid ML/HPC. Physics-based PSHA. Dynamic source inversion. Adjoint sensitivity analysis. Uncertainty quantification. Optimal observational network design.**

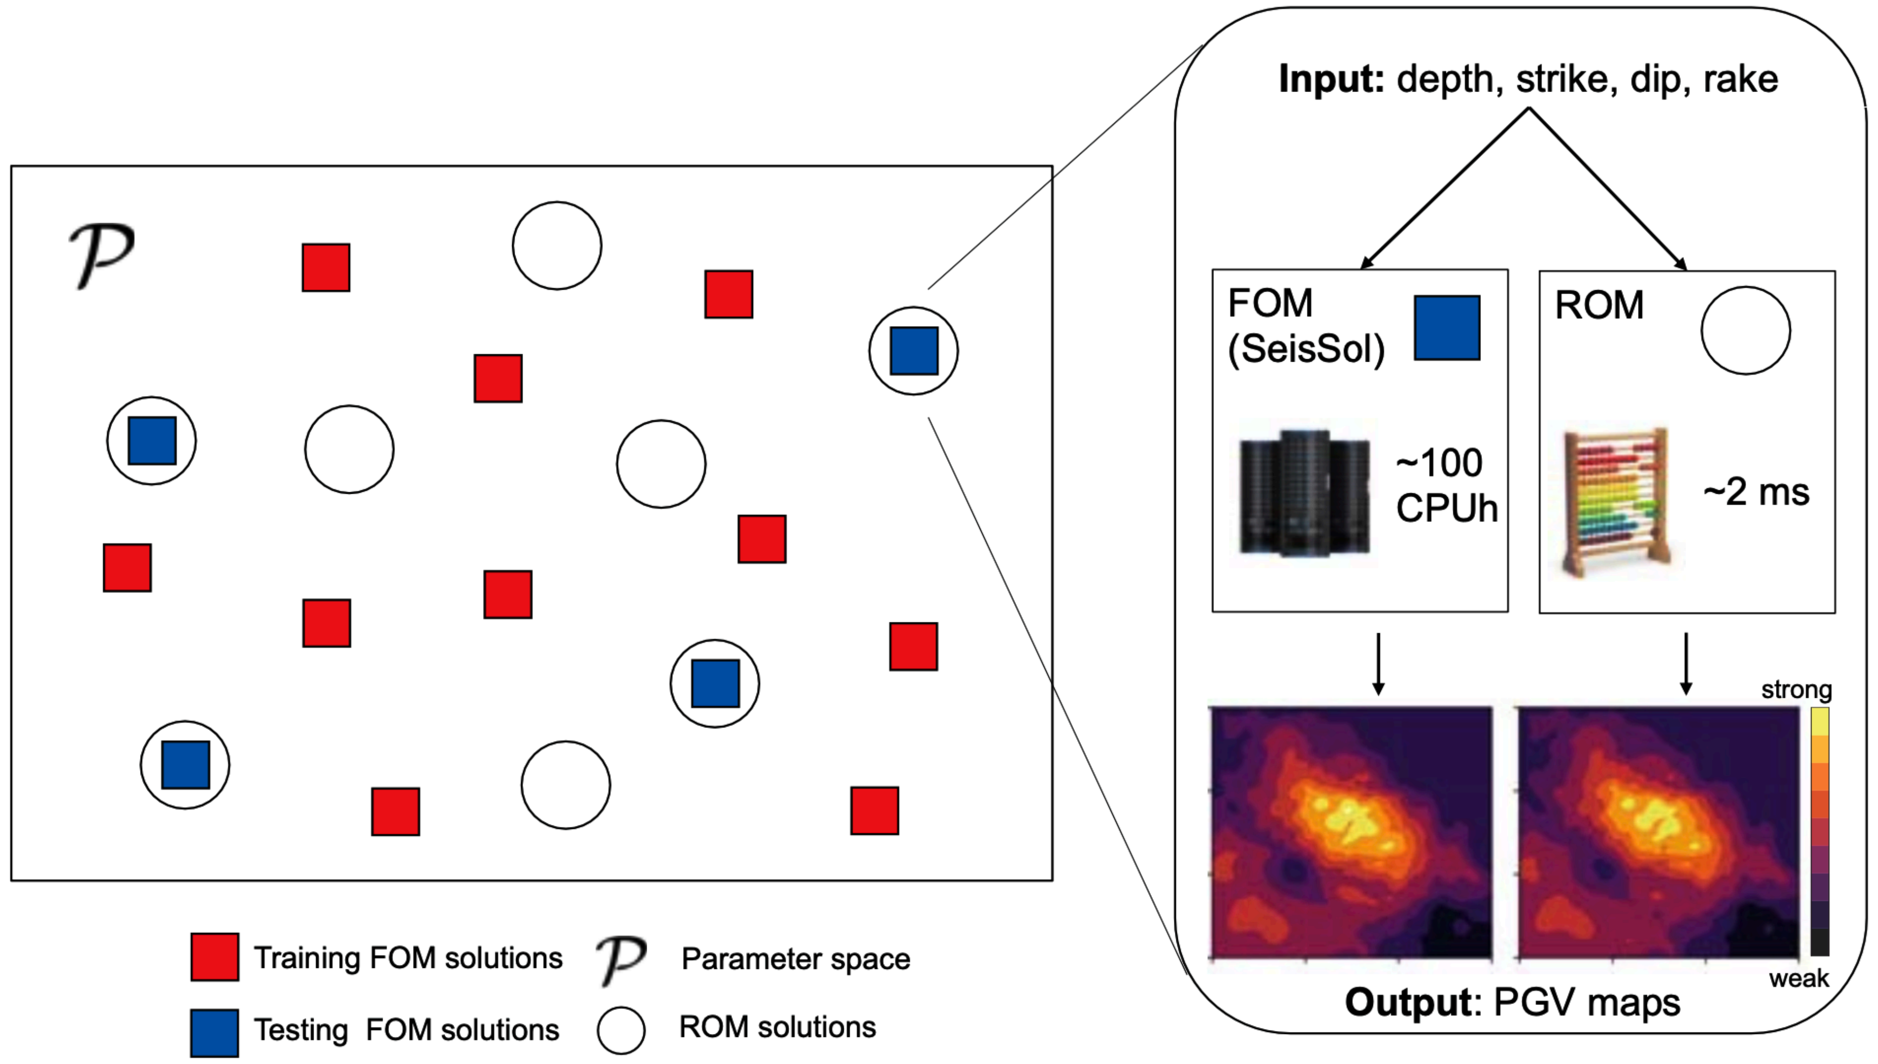

Constructing reduced order surrogate models combined with machine learning for ground motion modeling (John Rekoske et al., Arxiv)

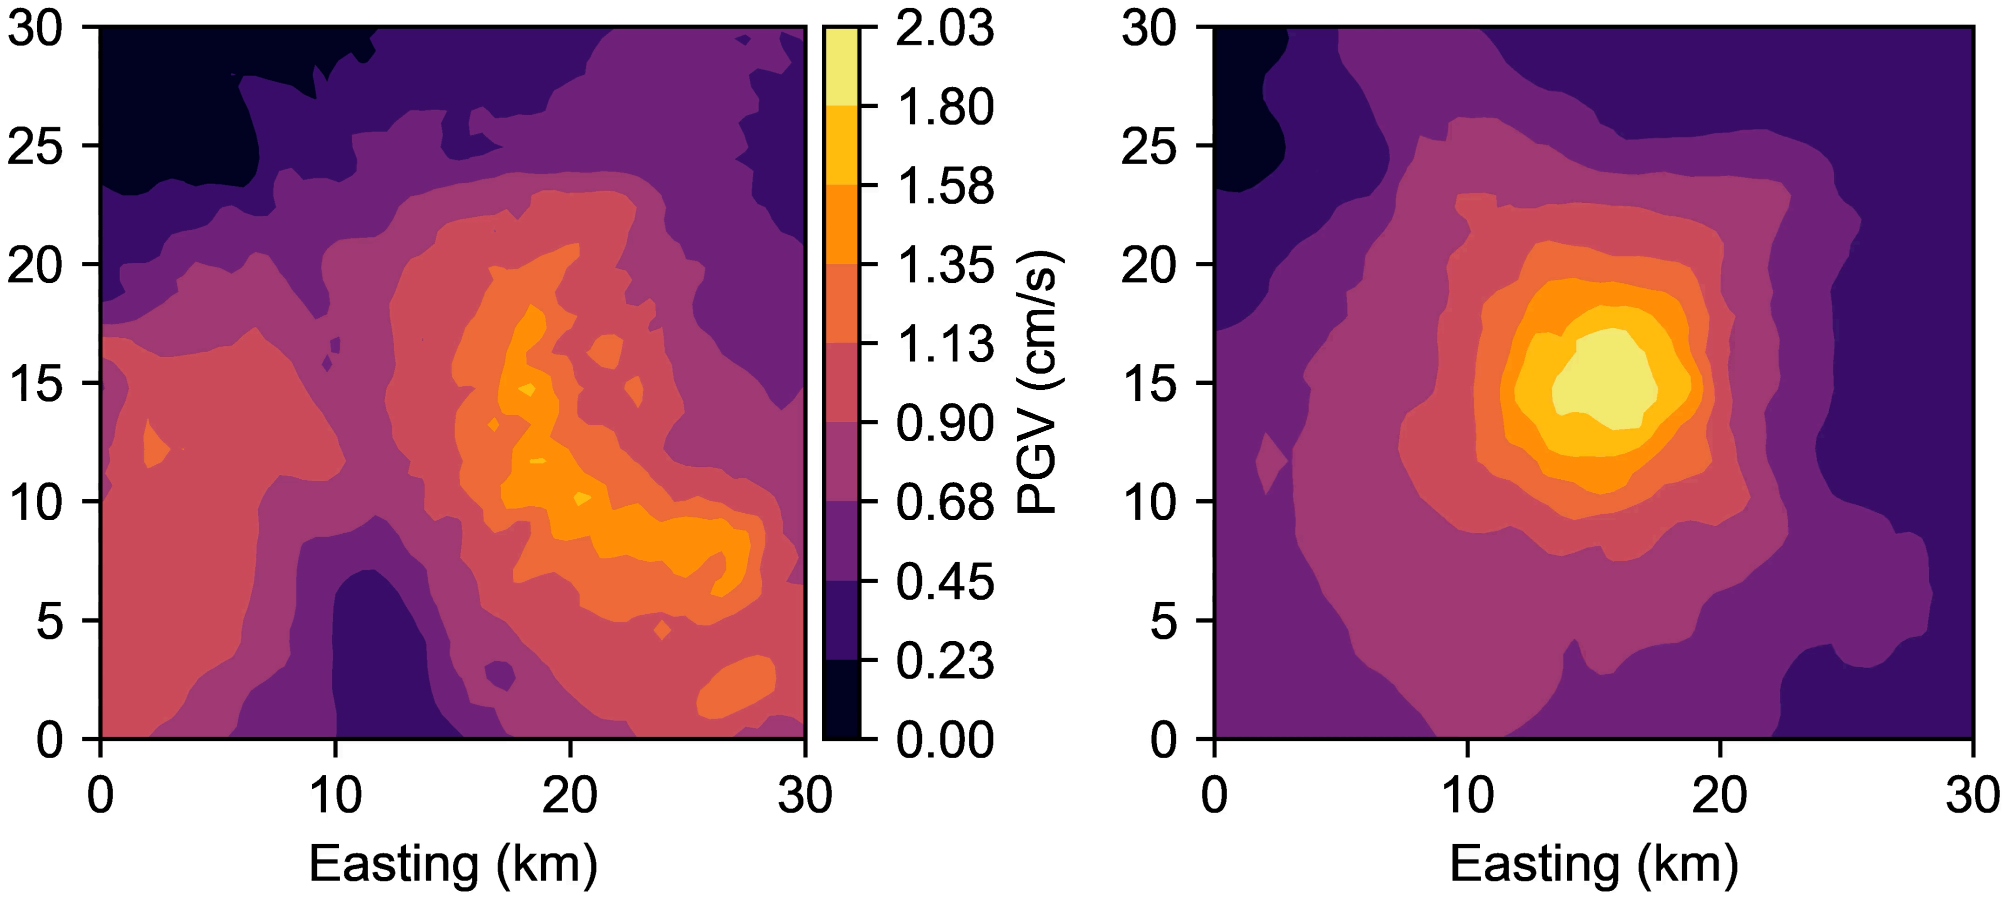

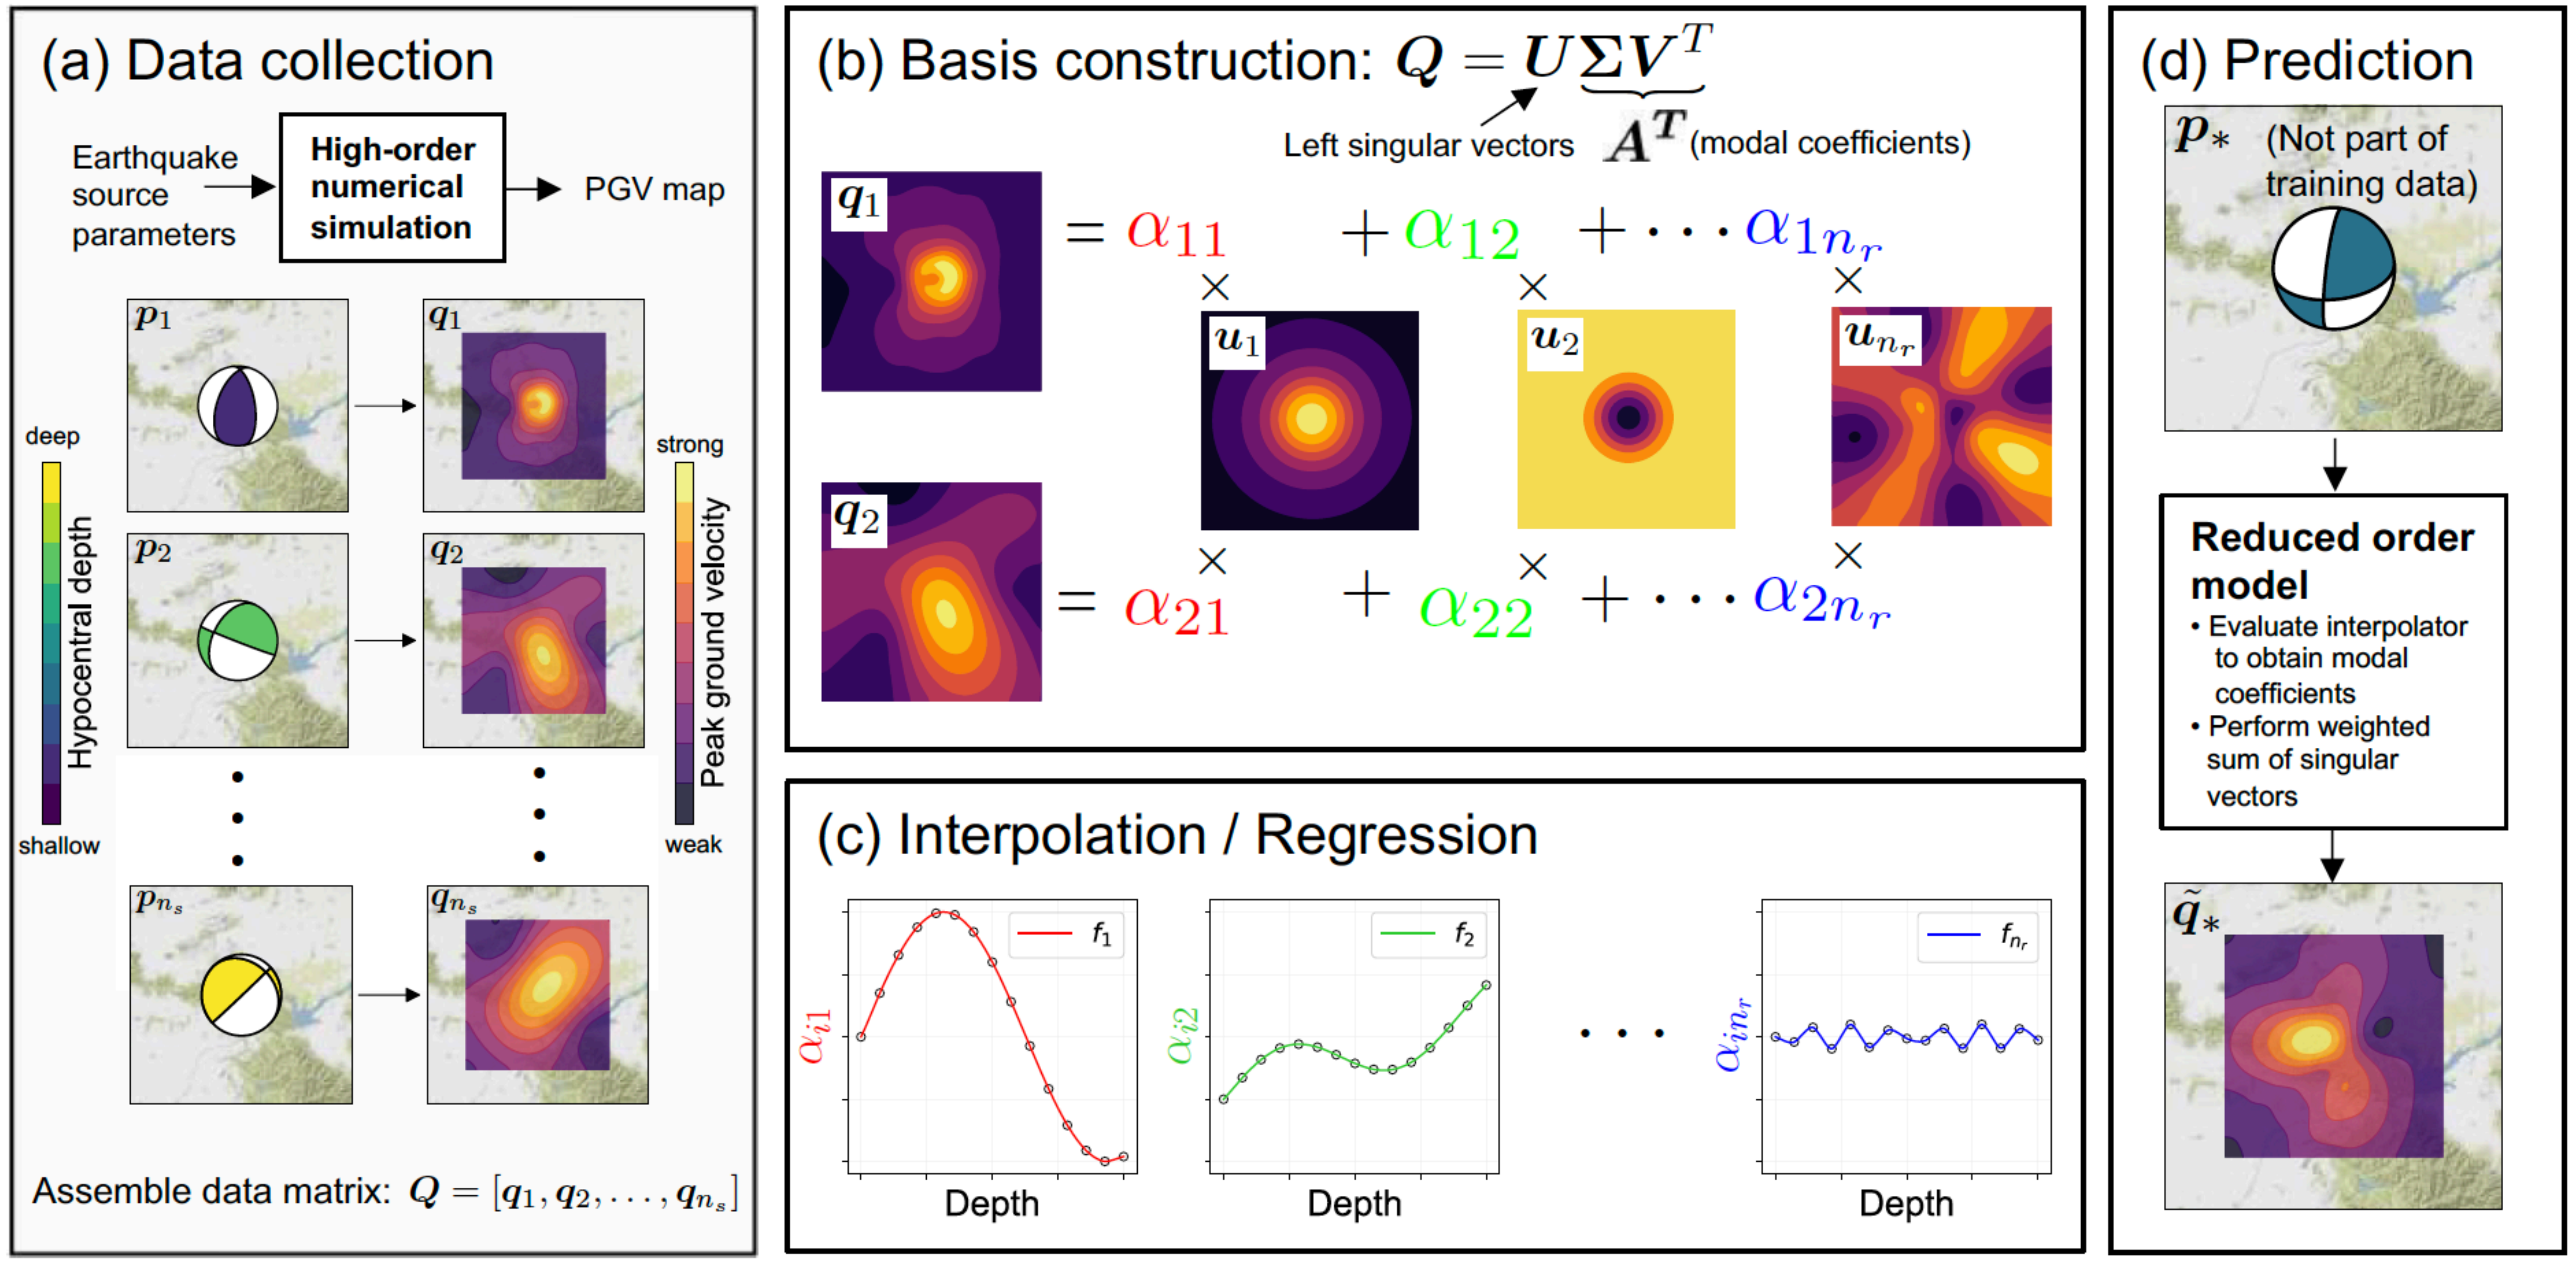

What is the future?

1. **Digital Twins (no more hero runs?).** Hybrid ML/HPC. Physics-based PSHA. Dynamic source inversion. Adjoint sensitivity analysis. Uncertainty quantification. Optimal observational network design.
2. **Rapid data-driven and physics-based simulations.** Integrate physics-based simulations into rapid geohazard response.

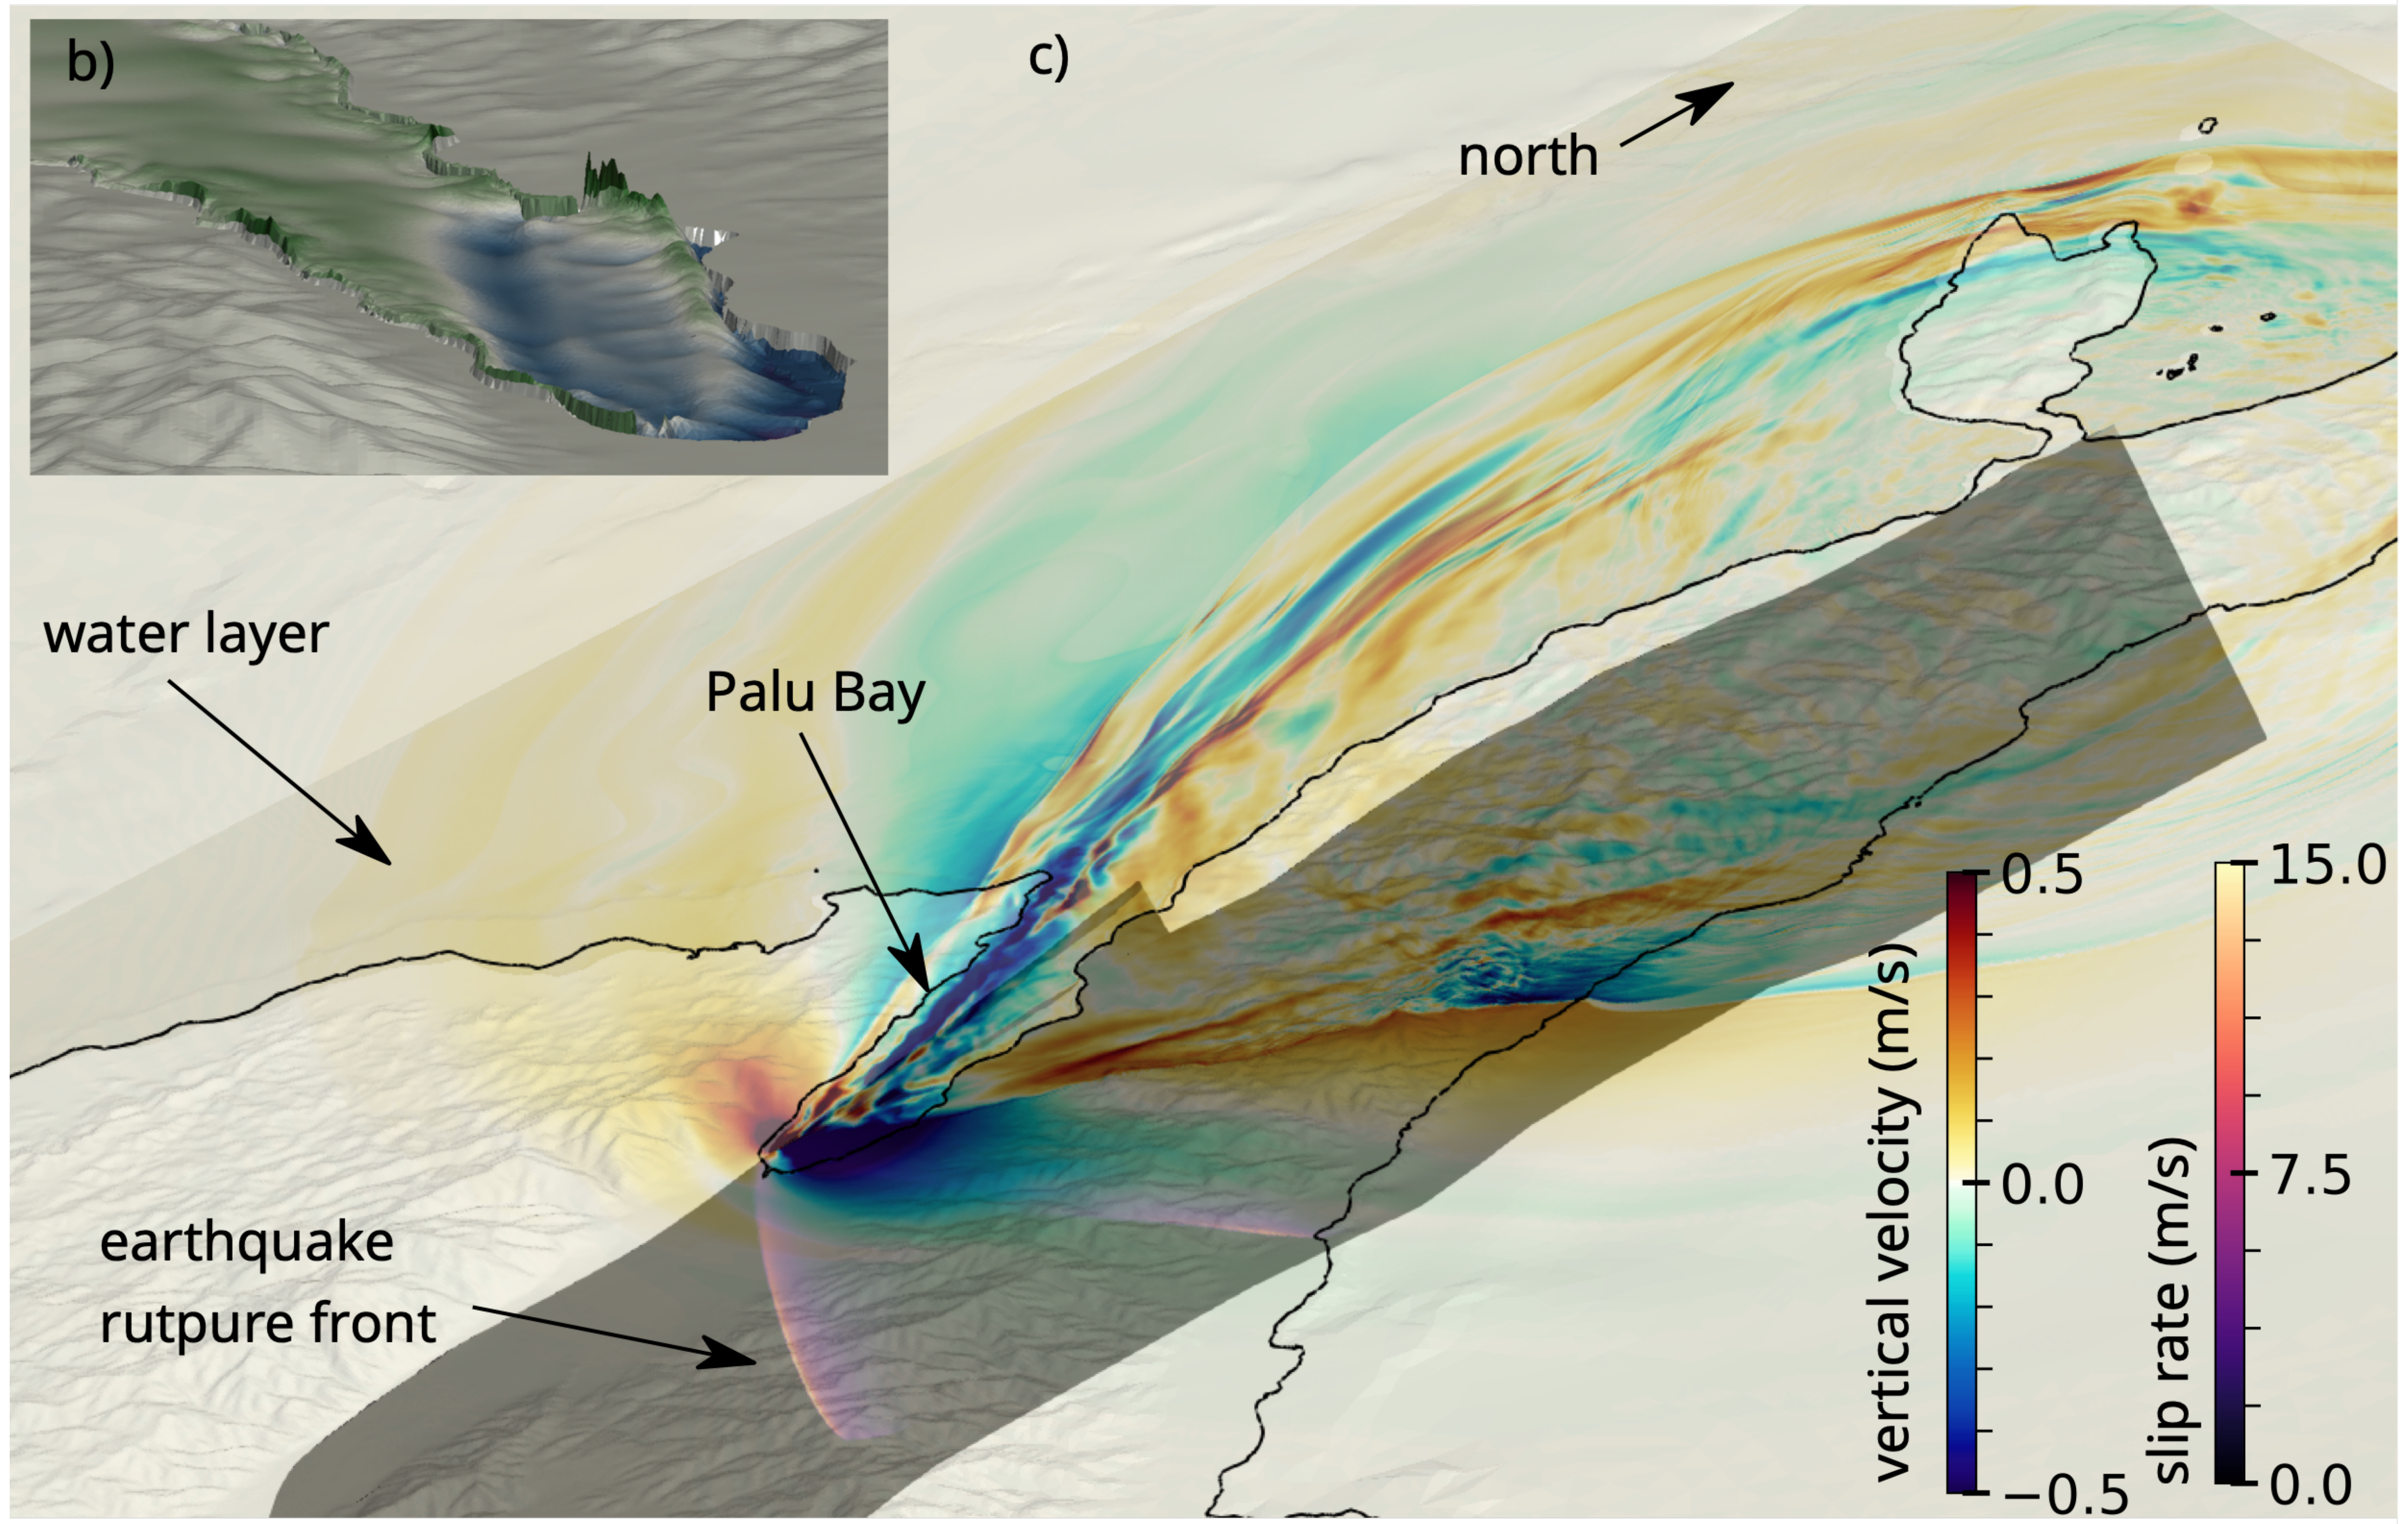

Krenz et al. SC'2021. Fully-coupled earthquake, acoustic and tsunami model for the 2018 Palu events.

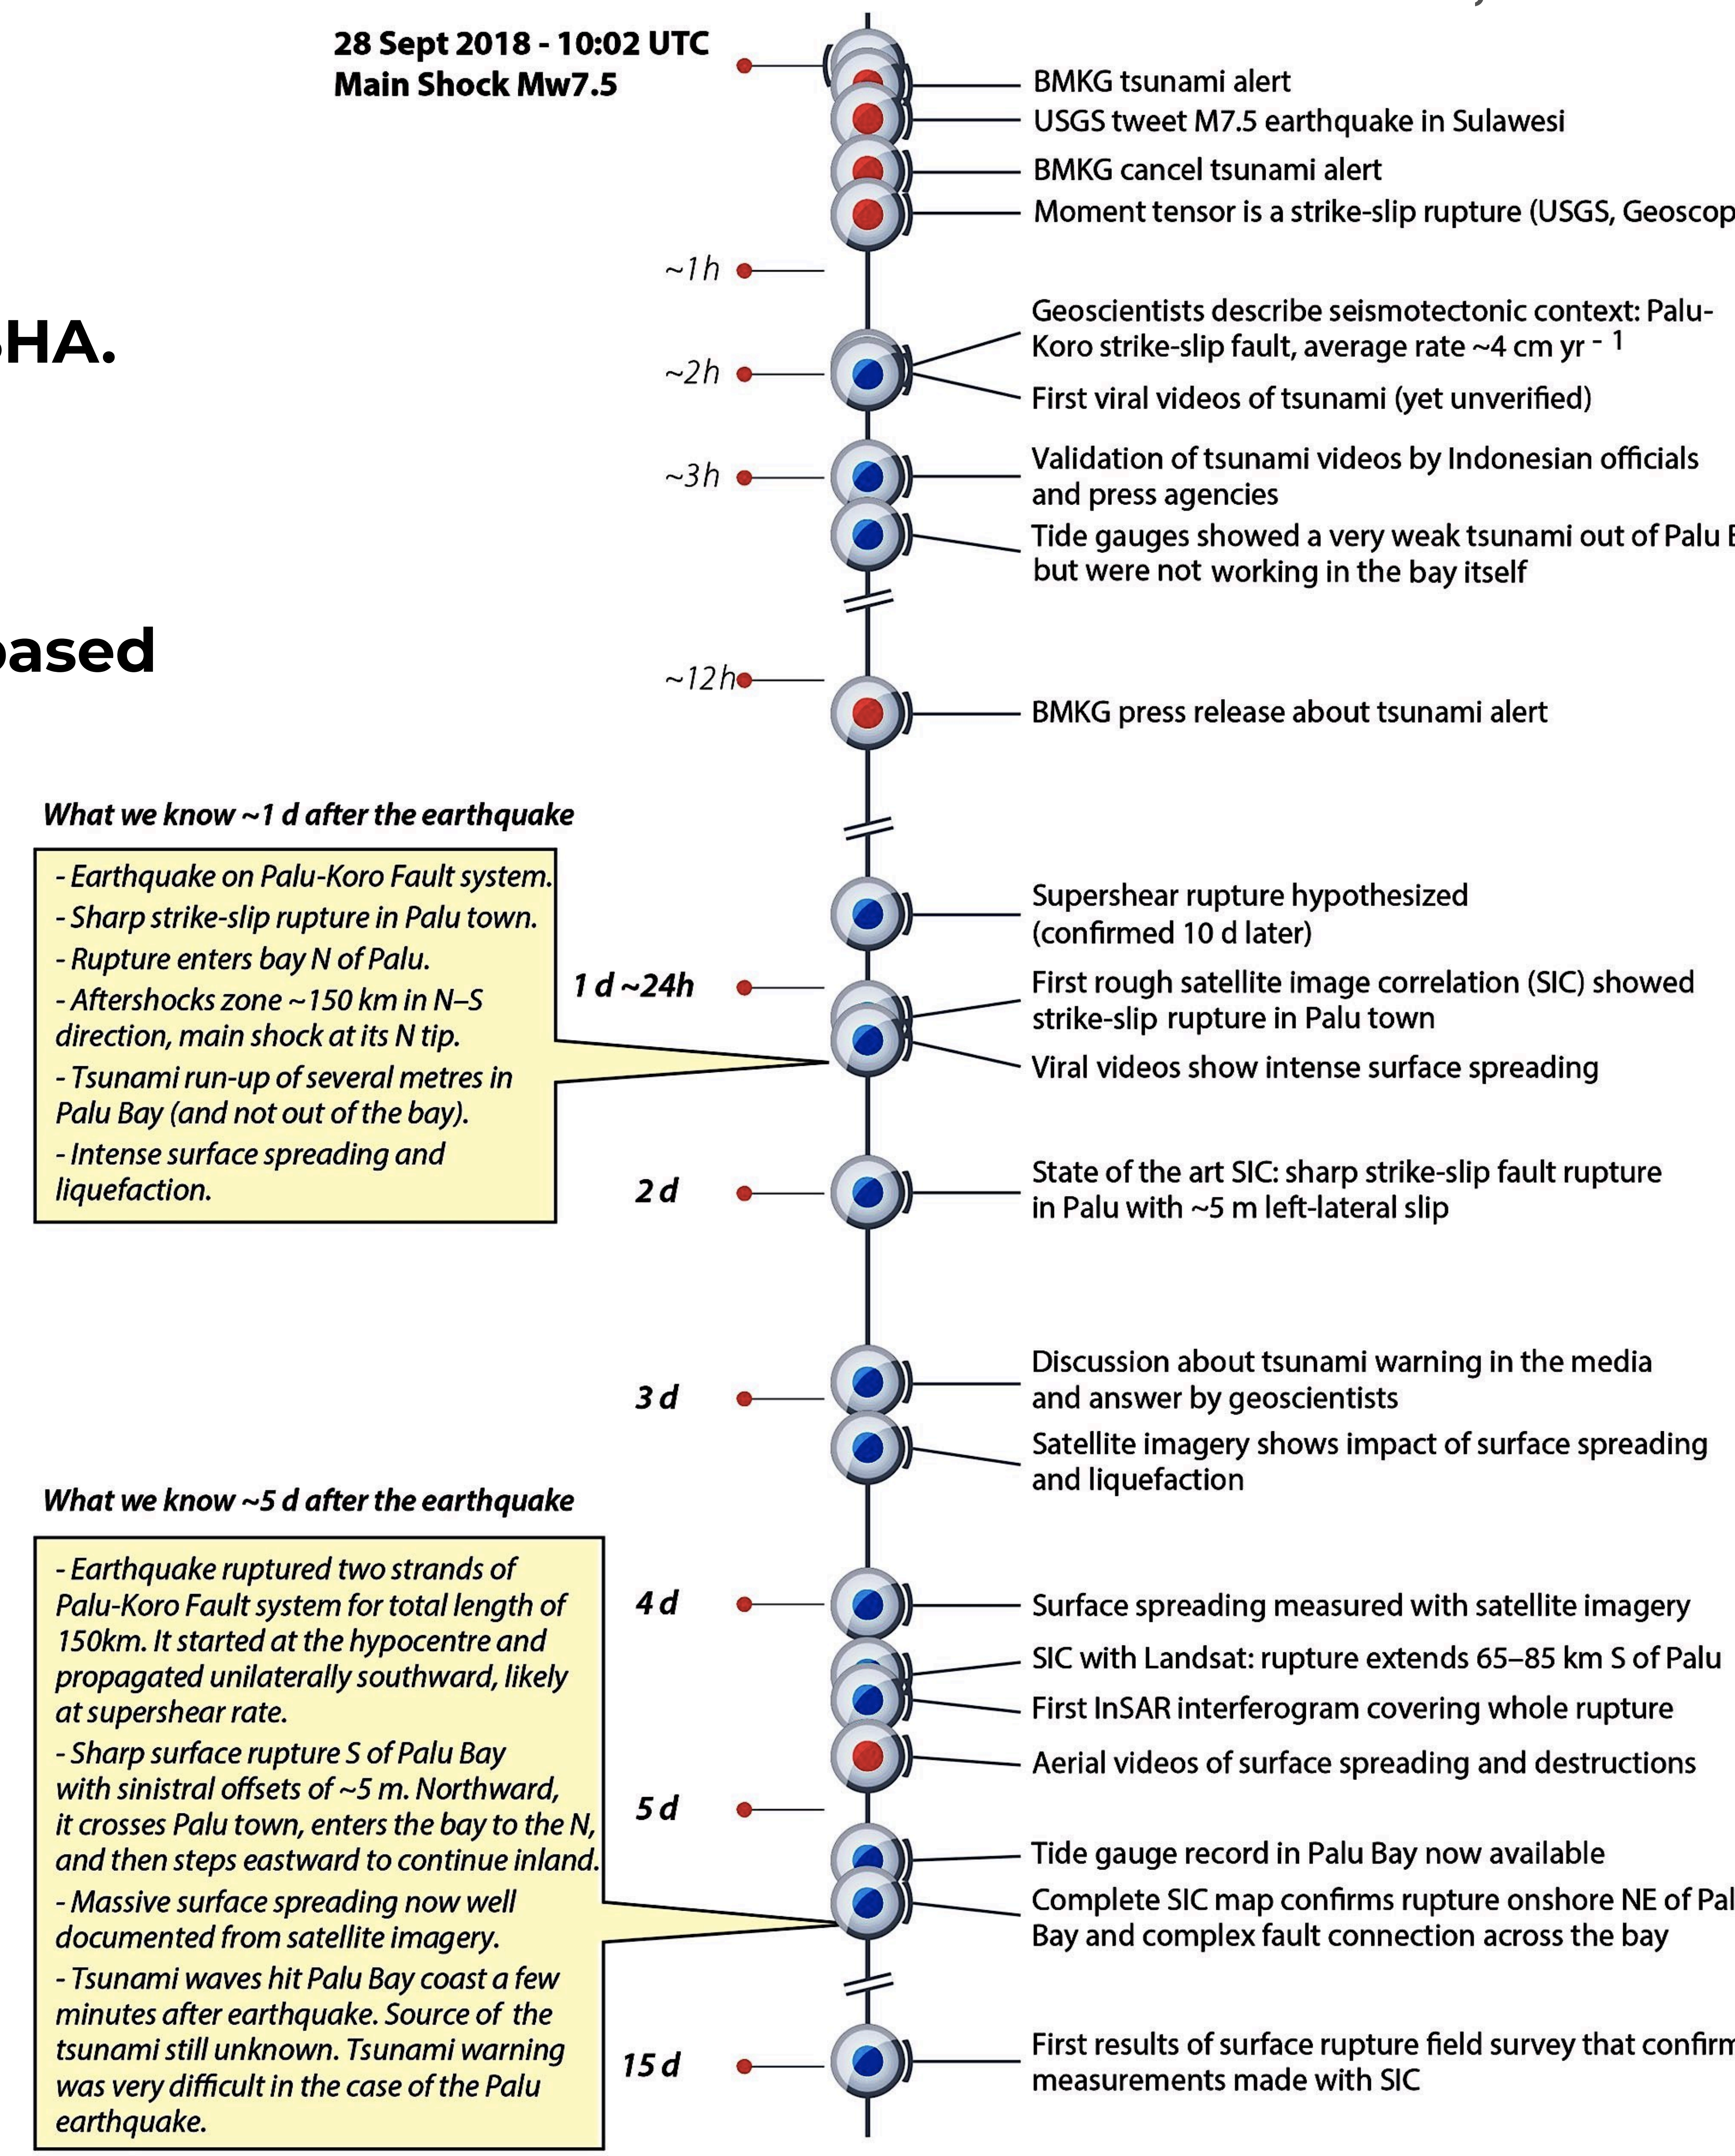

What is the future?

1. Digital Twins (no more hero runs?). Hybrid ML/HPC. Physics-based PSHA. Dynamic source inversion. Adjoint sensitivity analysis. Uncertainty quantification. Optimal observational network design.
2. Rapid data-driven and physics-based simulations. Integrate physics-based simulations into rapid geohazard response.
3. More realistic (dynamic core) models. Solving more/different equations. Bridging time scales and multi-physics.

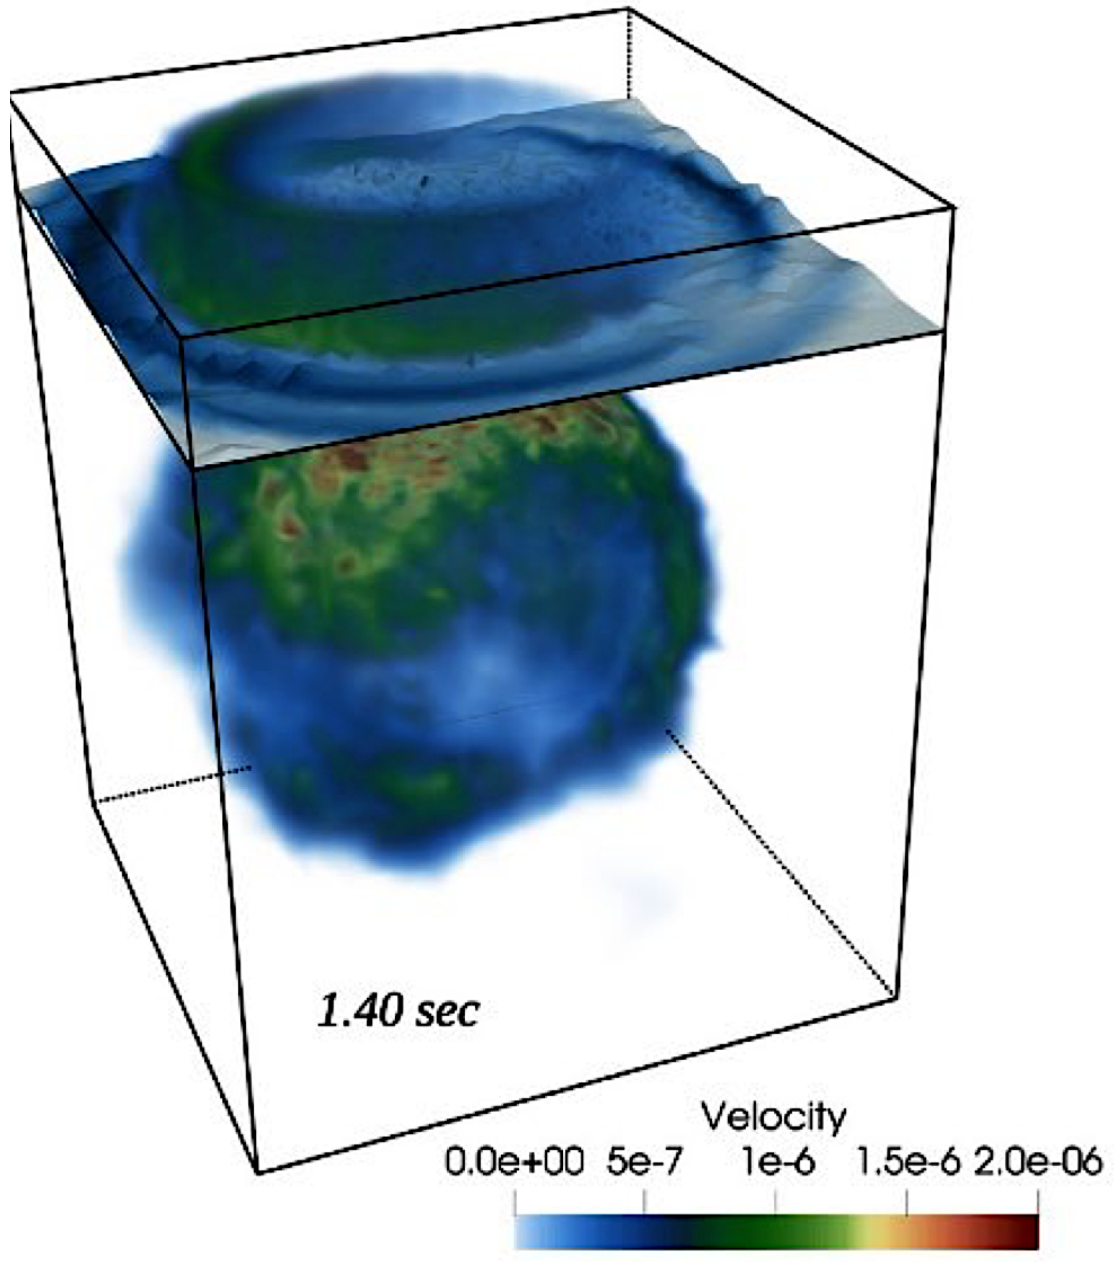

Seismic-acoustic simulation of the 'bang' that a M1.8 induced earthquake produces in Helsinki (Krenz et al., under review, ArXiv)

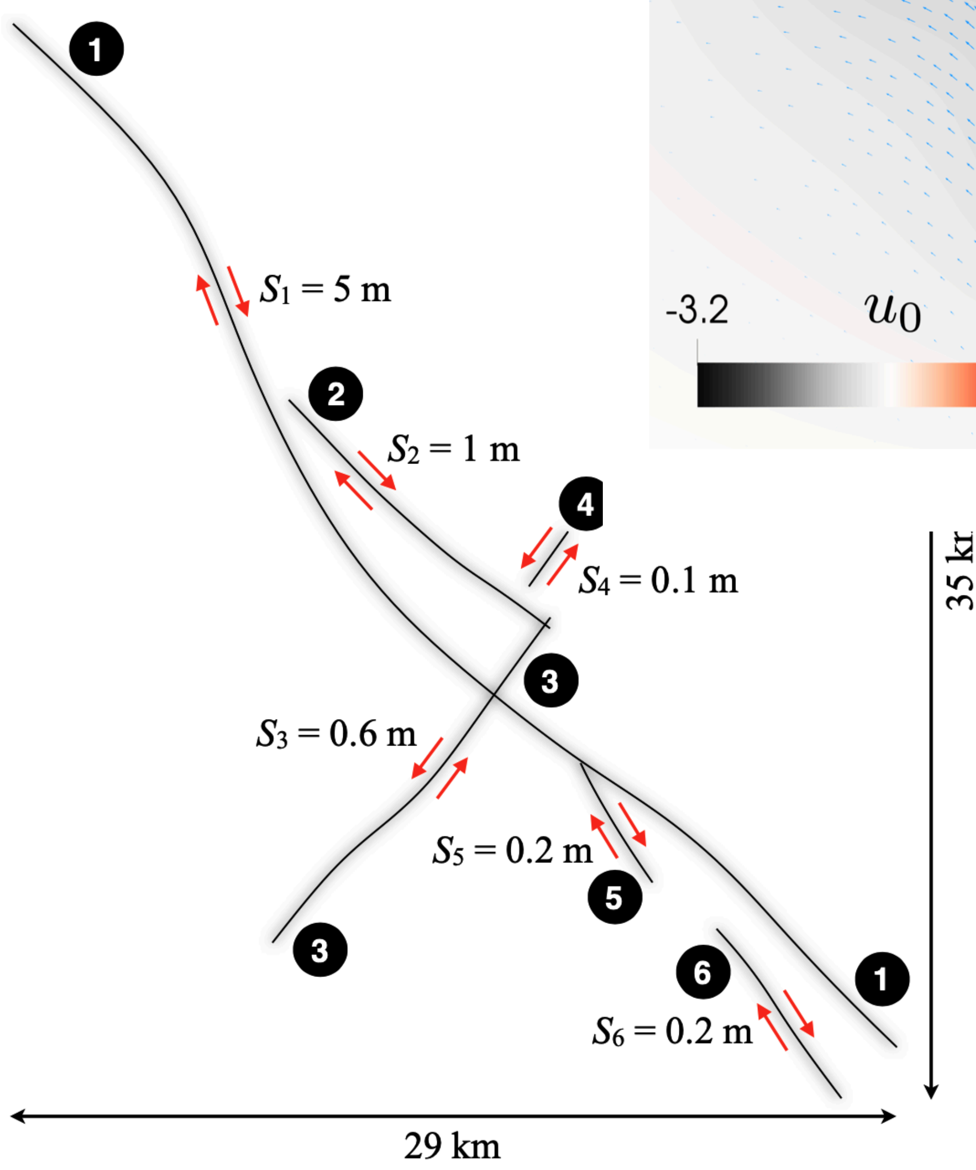

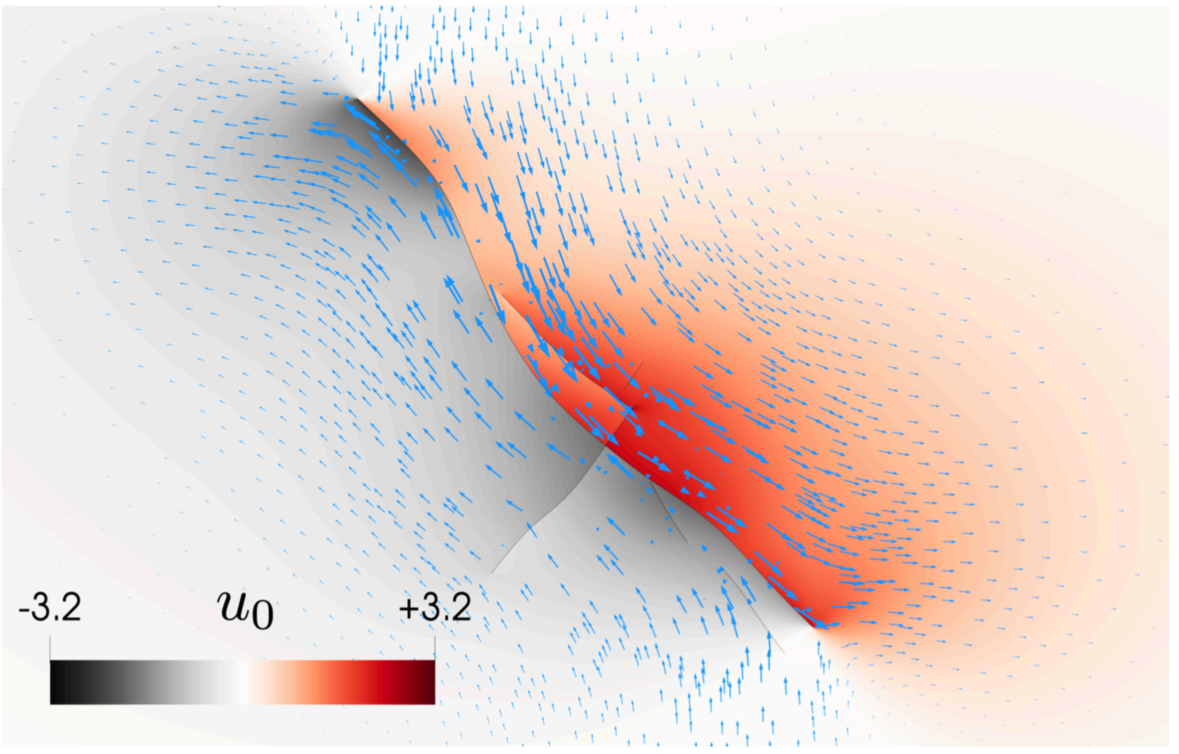

(a) x displacement (u_0)

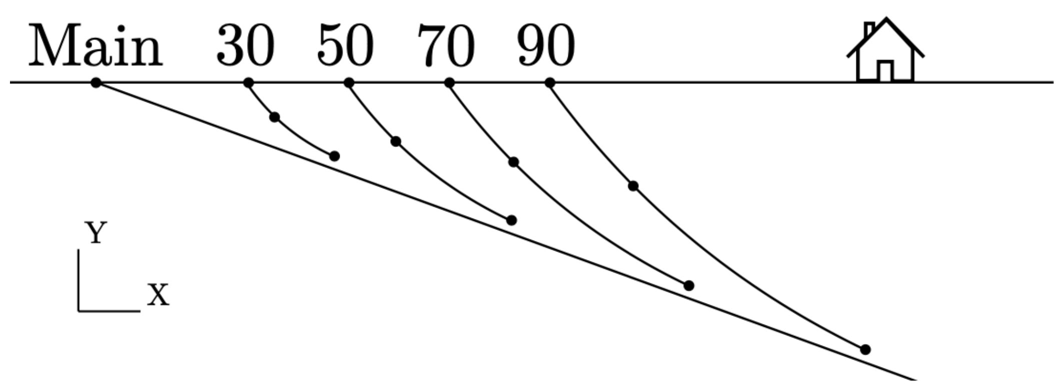

Multi-fault volumetric seismic cycle model on a shallowly dipping normal fault with four curved splay faults and for the Ridgecrest fault system using the elliptic PDE HPC solver tandem (Uphoff et al., GJI 2022)

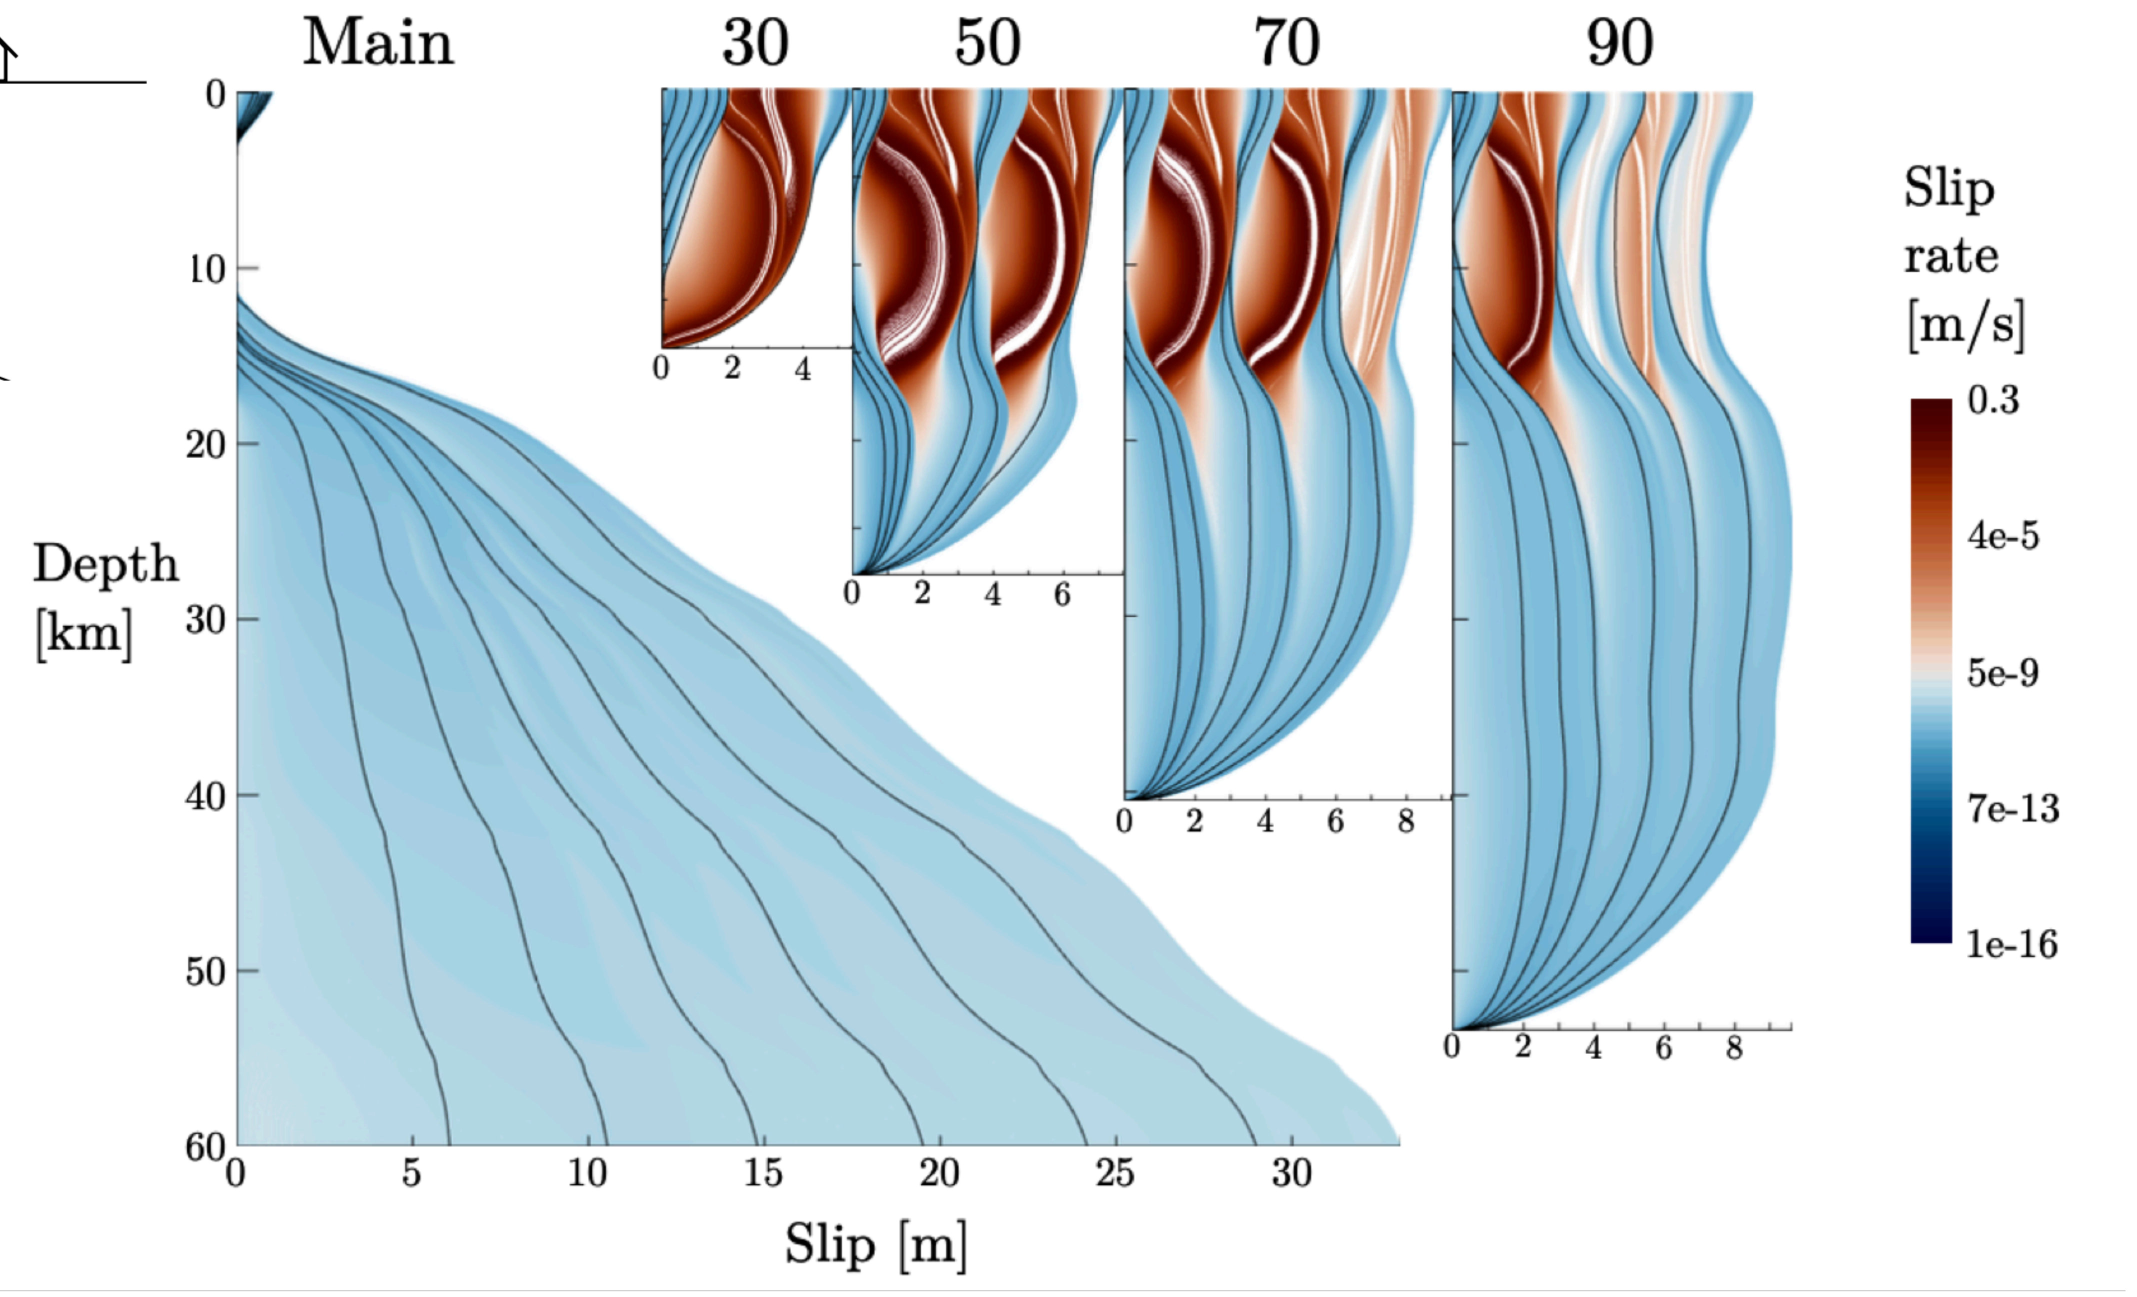

What is the future?

1. Digital Twins (no more hero runs?). Hybrid ML/HPC. Physics-based PSHA. Dynamic source inversion. Adjoint sensitivity analysis. Uncertainty quantification. Optimal observational network design.
2. Rapid data-driven and physics-based simulations. Integrate physics-based simulations into rapid geohazard response.
3. More realistic (dynamic core) models. Solving more/different equations. Bridging time scales and multi-physics.
4. Embracing Interdisciplinarity. HPC/Cloud optimization. Services. Workflows.

} Need for Simulation Data Lakes

DT GEO

Geo-INQUIRE
Geosphere **IN**frastructures for **QU**estions into **I**ntegrated **R**esearch

Geo-INQUIRE

SCOPED

ChEESE
Center of Excellence for Exascale in Solid Earth

MCS RCN

CIG COMPUTATIONAL INFRASTRUCTURE for GEODYNAMICS

Summary

- 3D physics-based dynamic rupture modeling provides mechanically viable insight into the (regional) physical conditions that govern the dynamics of earthquakes
- **Observational constraints can be routinely included, unified and probed for dynamic plausibility reducing non-uniqueness and helping to constrain competing views and to bridge scales**
- HPC allows for physics-based and data-driven models
- **Advances in high-performance computing and dense observations will allow us to go beyond scenario-based analysis, aiming for, e.g., multi-physics, fully non-linear source-path-site effects, urgent response, data-driven dynamic source inversion, (Bayesian) uncertainty quantification, hybrid HPC & ML modeling...**
- **... motivating Simulation Data Lakes**

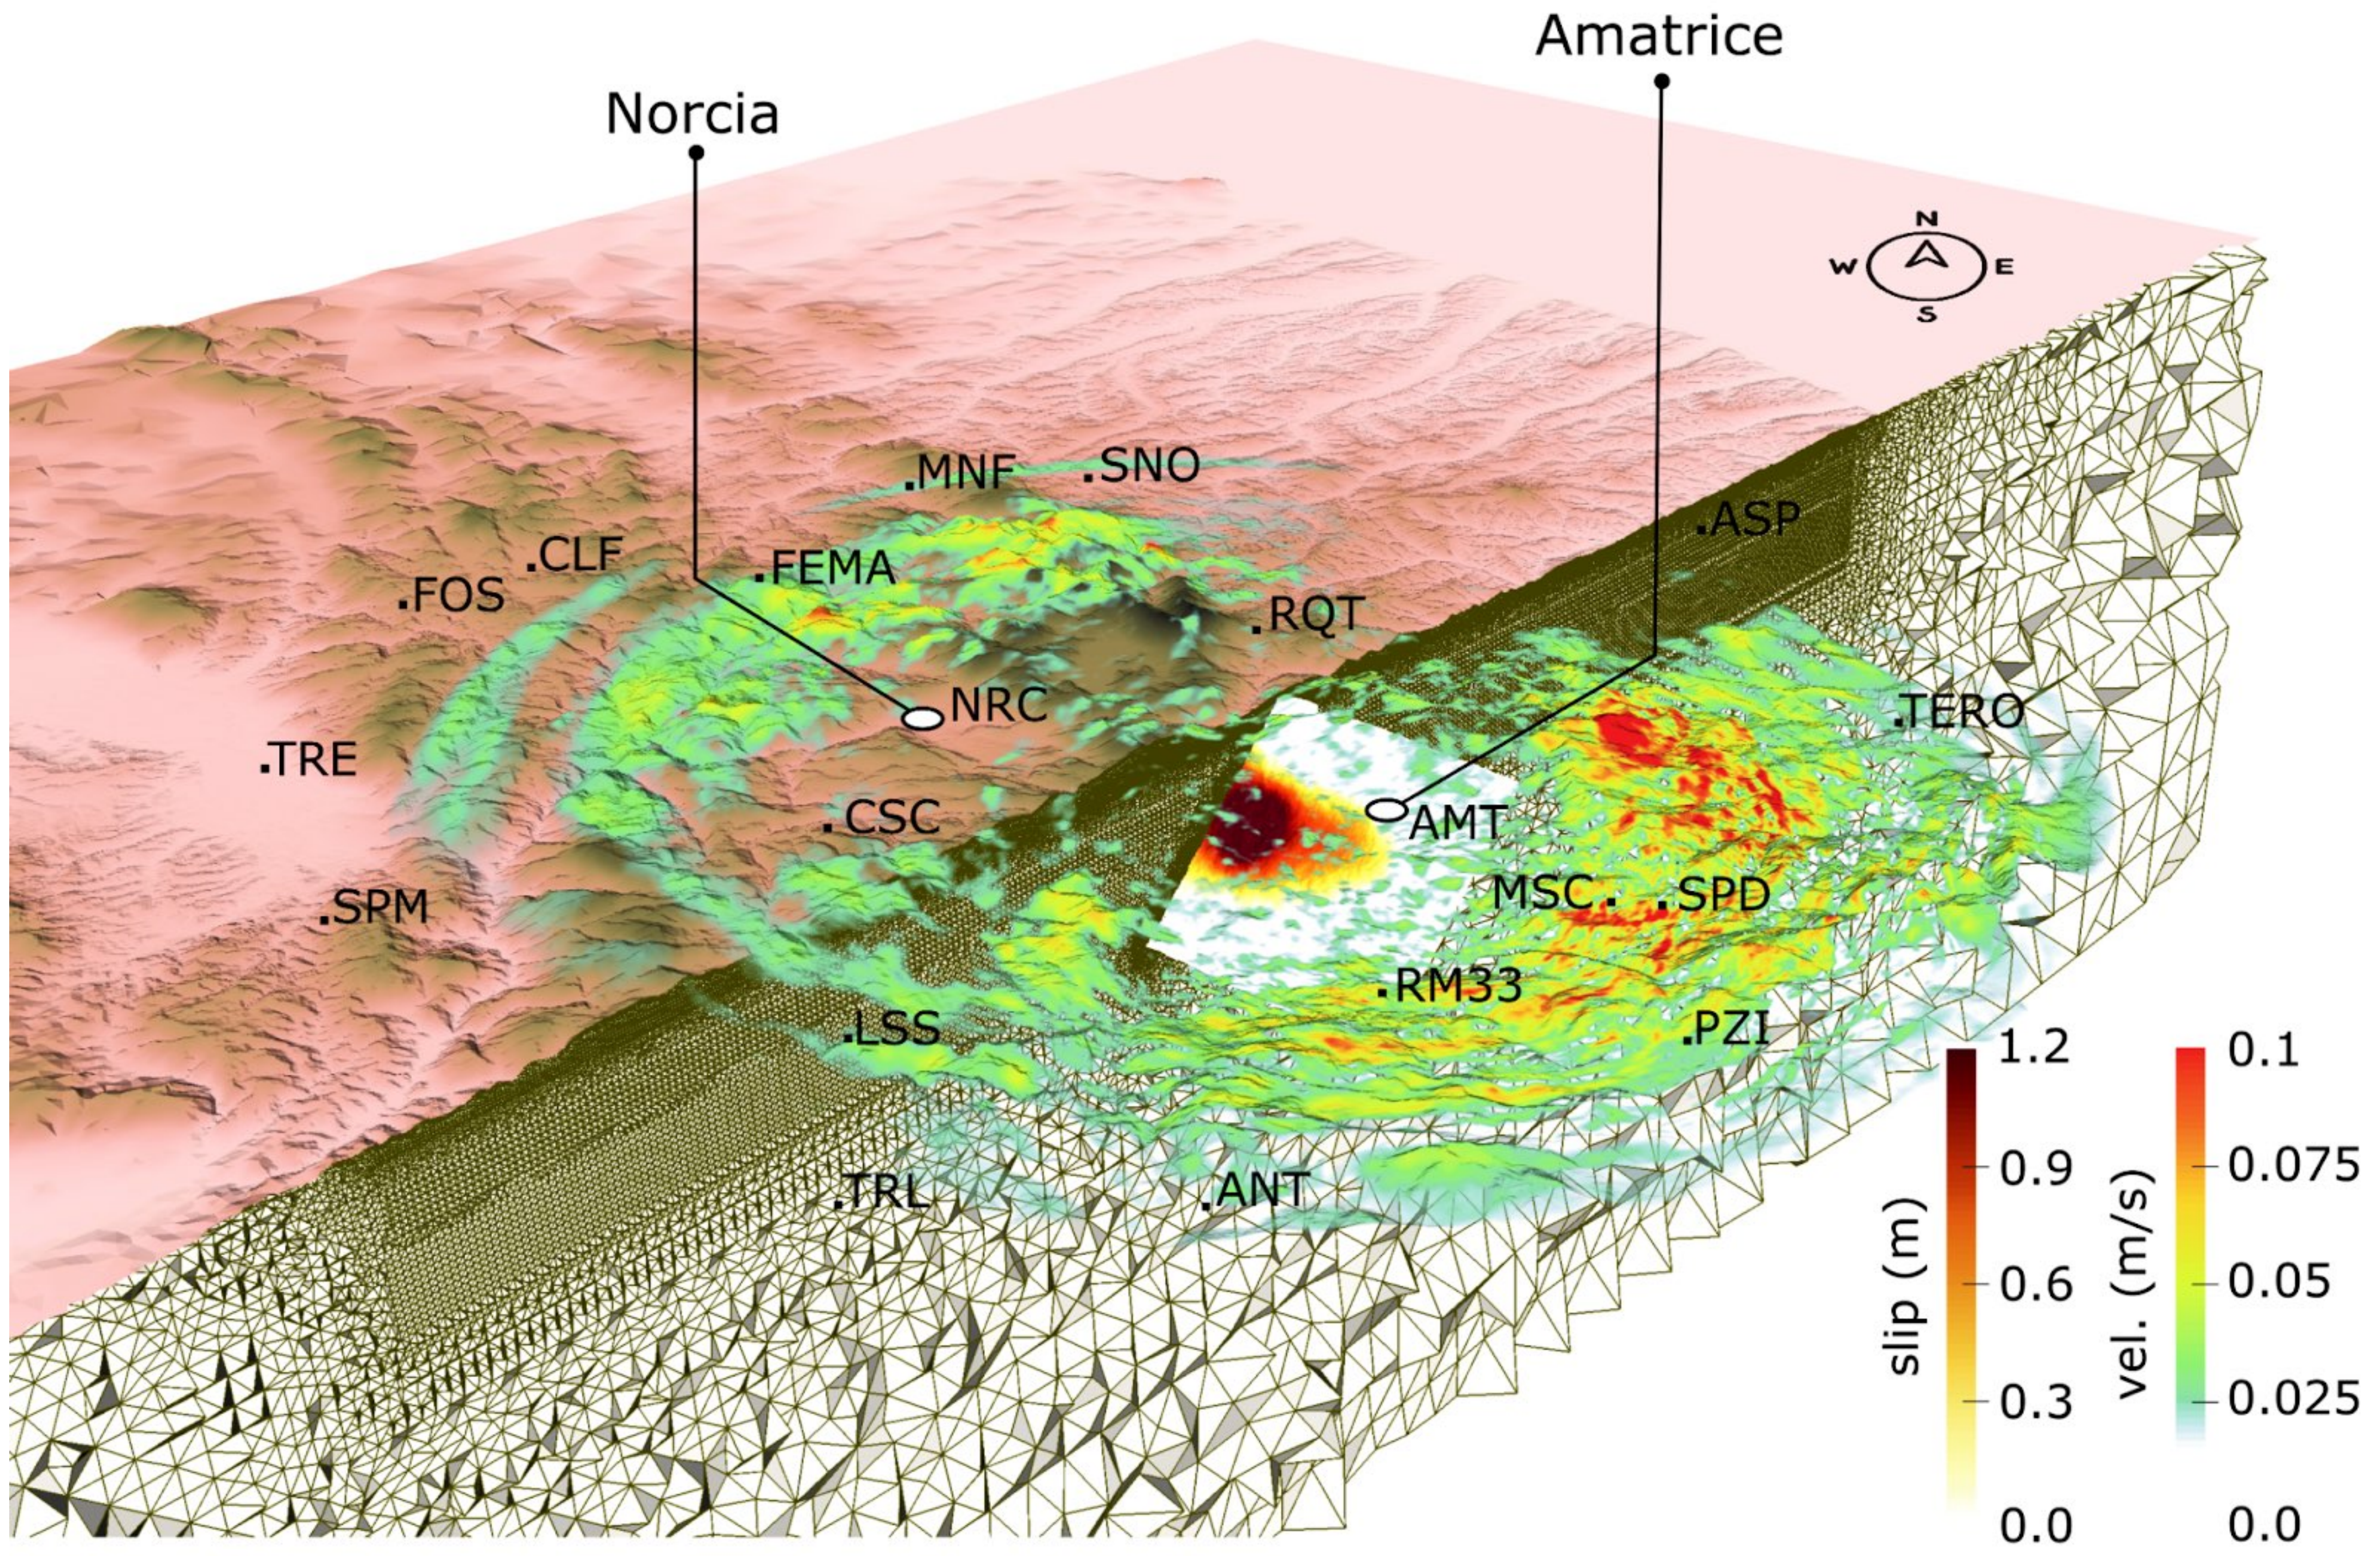

algabriel@ucsd.edu
gabriel@geophysik.uni-muenchen.de

twitter @InSeismoland