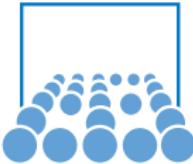


HPC and Data-Intensive Applications in Earth Sciences

Optimizing a Dynamic Rupture and Earthquake Simulation Code for SuperMUC, Tianhe-2 and Stampede

Michael Bader

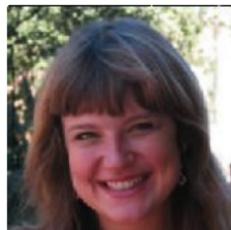
Technische Universität München



SeisSol Core Developers



Alexander
Breuer



Alice-Agnes
Gabriel



Alexander
Heinecke



Christian
Pelties



Sebastian
Rettenberger

SeisSol Core Developers and Collaborators

- TUM: A. Breuer, S. Rettenberger, M. Bader, A. Heinecke
- LMU: A.-A. Gabriel, C. Pelties
- **Intel's Parallel Computing Lab:** A. Heinecke, Karthikeyan Vaidyanathan, Mikhail Smelyanskiy and Pradeep Dubey
- **Leibniz Supercomputing Centre:** Arndt Bode
- **Texas Advanced Computing Centre:** William Barth
- **National University of Defense Technology:** Xiang-Ke Liao



Alexander
Breuer



Alice-Agnes
Gabriel



Alexander
Heinecke



Christian
Pelties



Sebastian
Rettenberger

Overview and Agenda

SeisSol:

- dynamic rupture and seismic wave propagation
- unstructured tetrahedral meshes
- high-order ADER-DG discretisation

Optimisation for Heterogeneous Petascale Platforms:

- code generation to optimize element-local matrix kernels
- hybrid MPI/OpenMP parallelisation
- offload scheme to address multiphysics

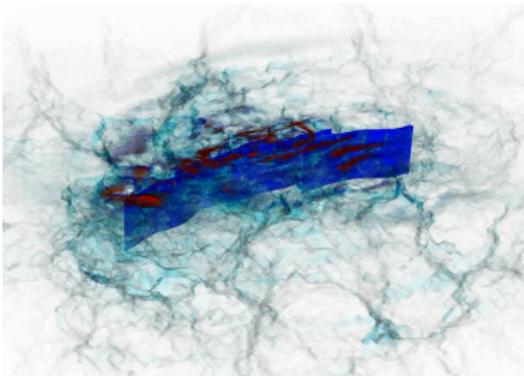
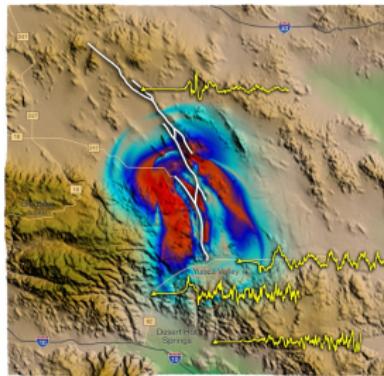
Performance on Tianhe-2, Stampede and SuperMUC:

- weak scaling of wave propagation component
- strong scaling for 1992 Landers M7.2 earthquake

Part I

Dynamic Rupture and Earthquake Simulation with SeisSol

Dynamic Rupture and Earthquake Simulation

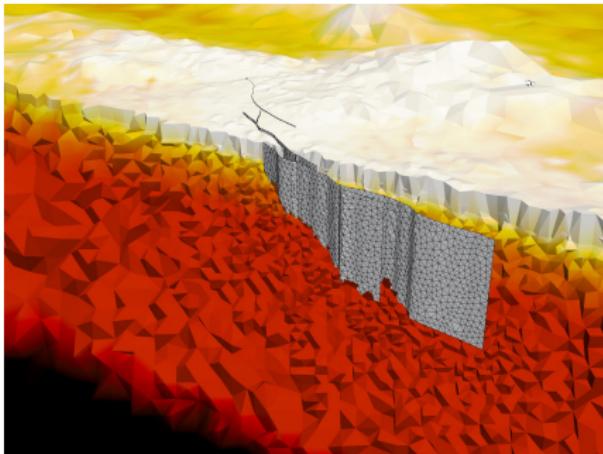
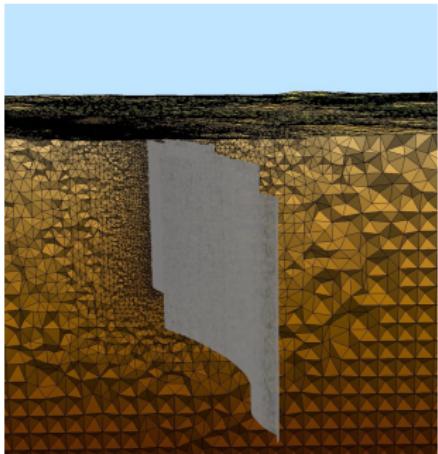


Landers fault system: simulated ground motion and seismic waves

SeisSol – ADER-DG for seismic simulations:

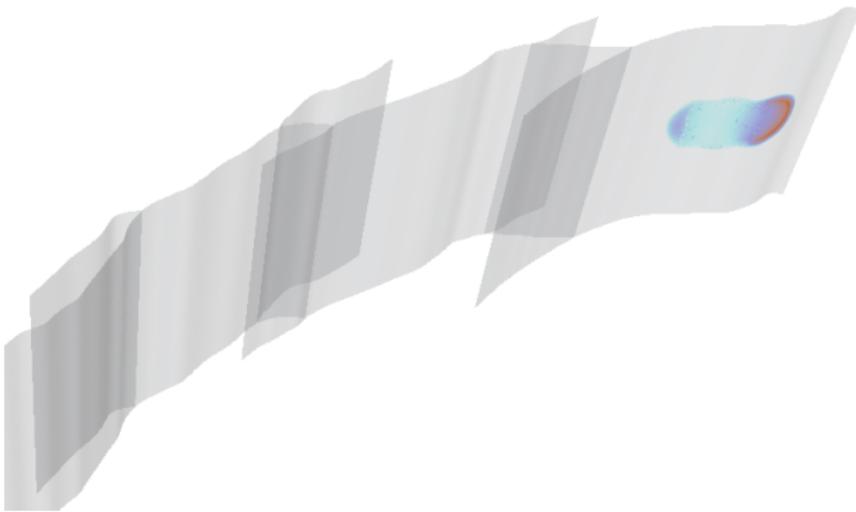
- adaptive tetrahedral meshes
→ complex geometries, heterogeneous media, multiphysics
- complicated fault systems with multiple branches
→ non-linear multiphysics dynamic rupture simulation
- ADER-DG: high-order discretization in space and time

1992 Landers M7.2 Earthquake



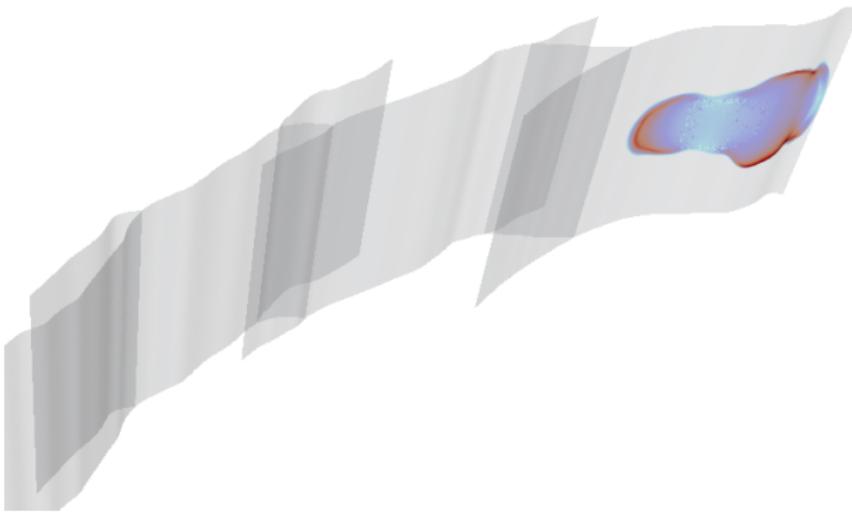
- multiphysics simulation of dynamic rupture and resulting ground motion of a M7.2 earthquake
- fault inferred from measured data, regional topography from satellite data, physically consistent stress and friction parameters
- 1D velocity structure, low velocity near surface

Multiphysics Dynamic Rupture Simulation



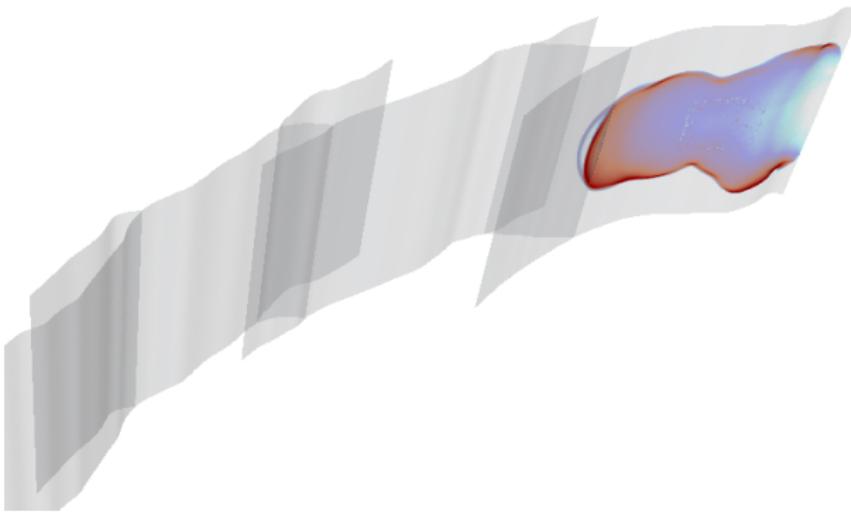
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



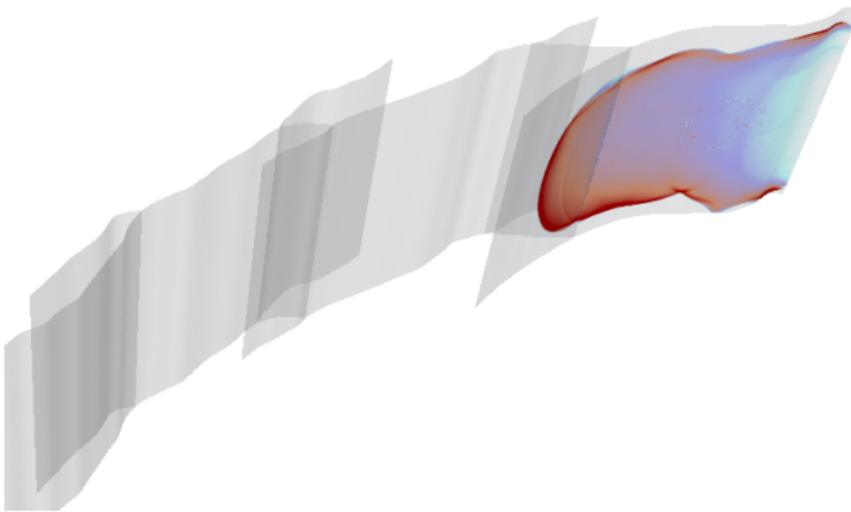
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



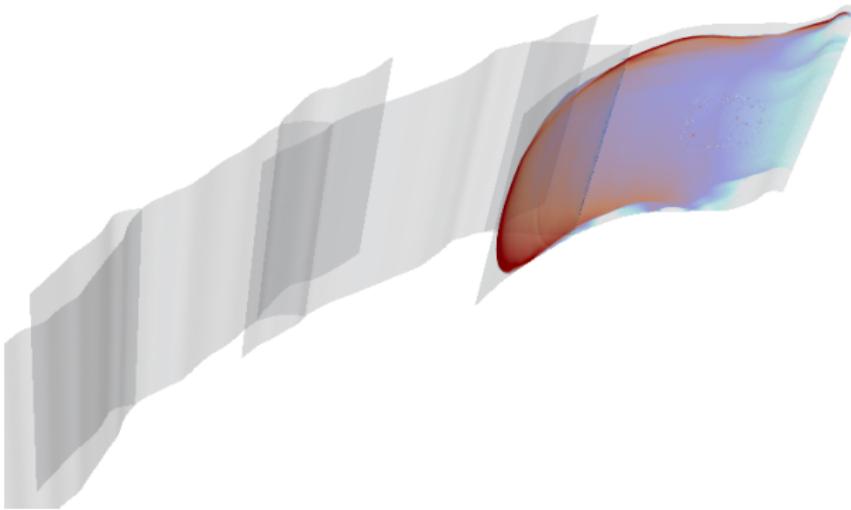
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



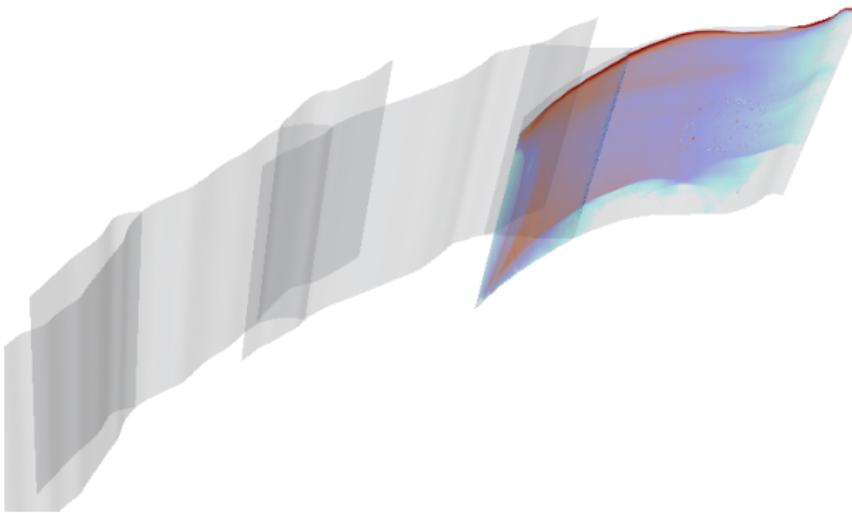
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



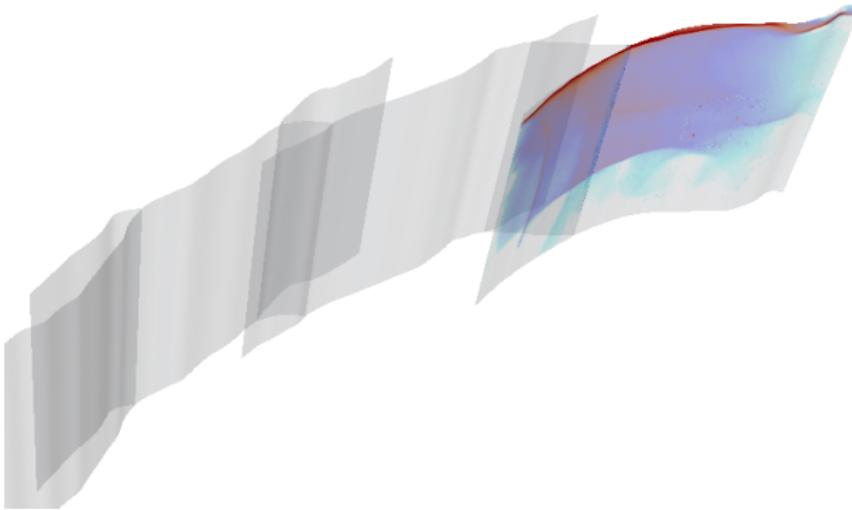
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



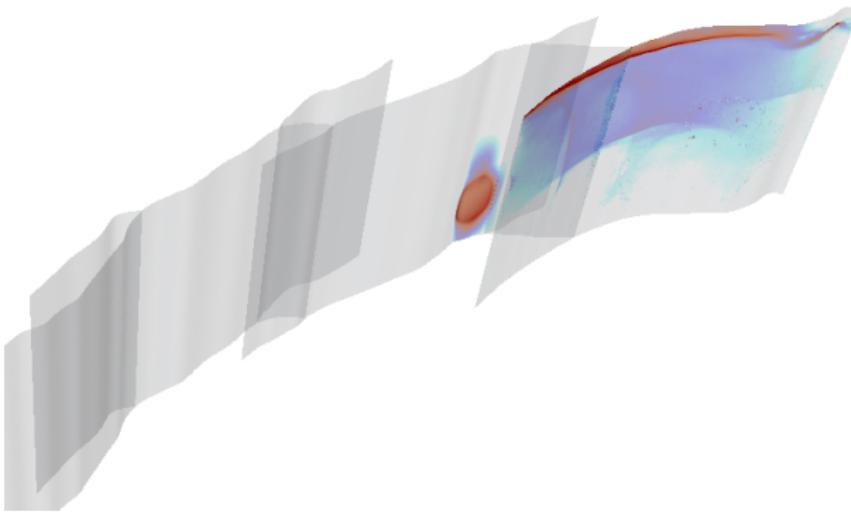
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



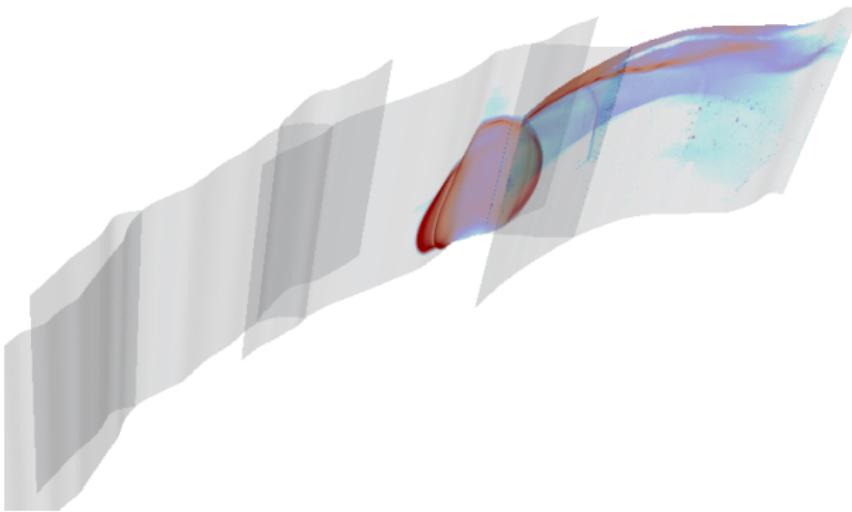
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



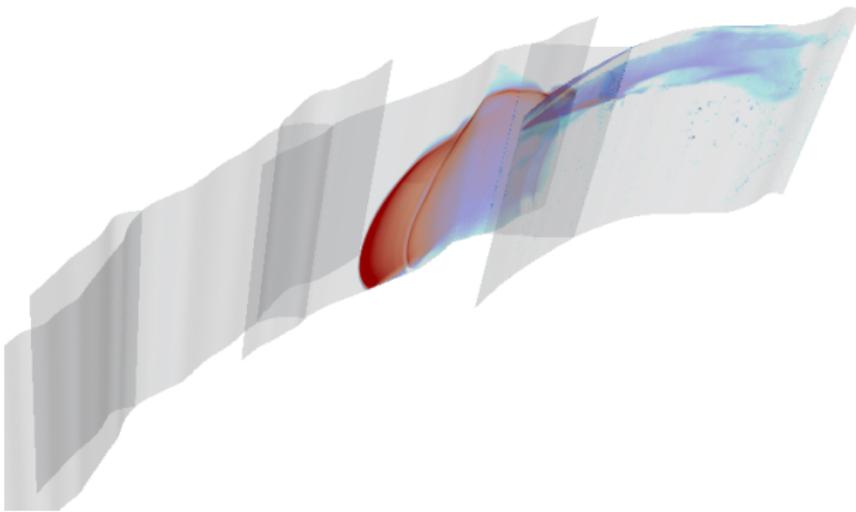
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



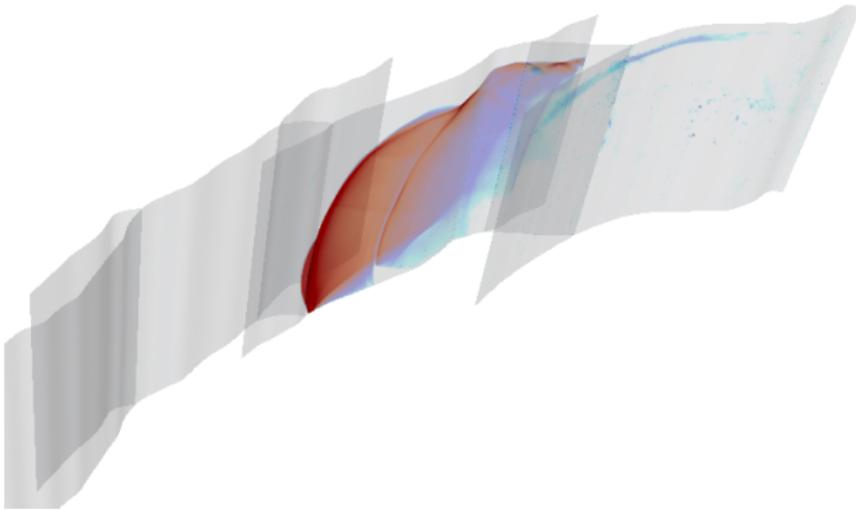
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



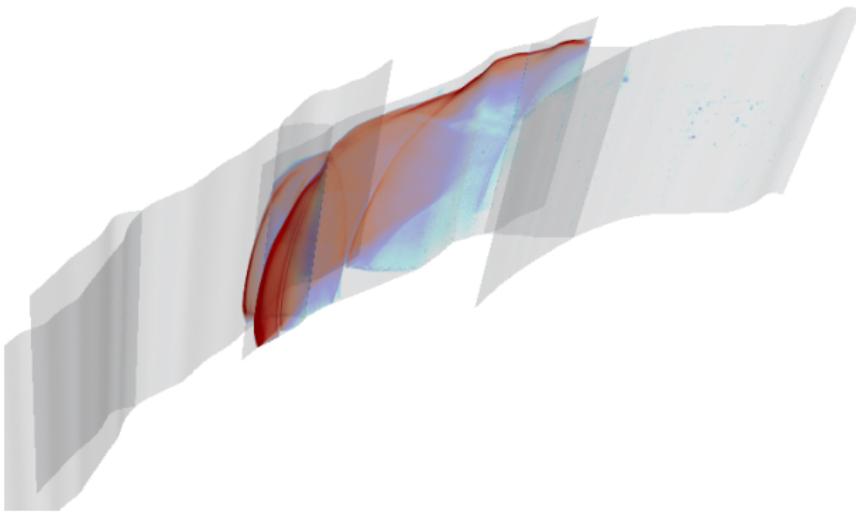
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



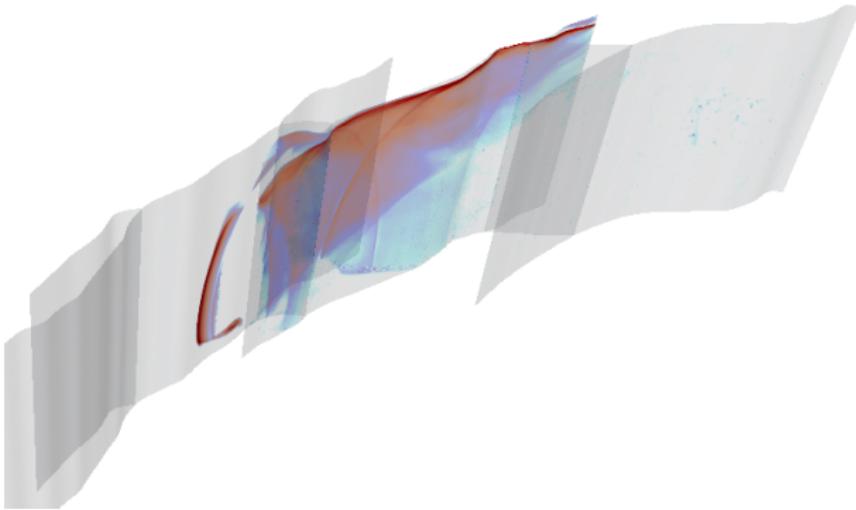
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



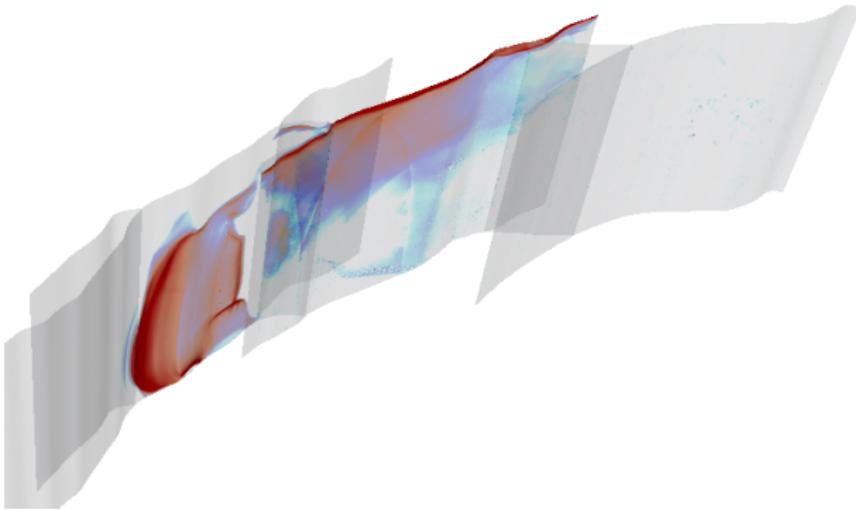
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



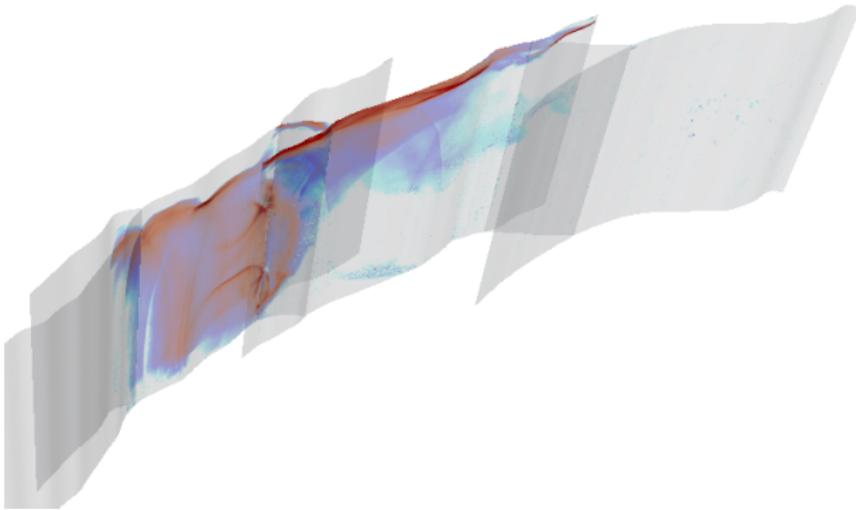
- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

Multiphysics Dynamic Rupture Simulation



- spontaneous rupture, non-linear interaction with wave-field
- featuring rupture jumps, fault branching, etc.
- tackles fundamental questions on earthquake dynamics
- realistic rupture source for seismic hazard assessment

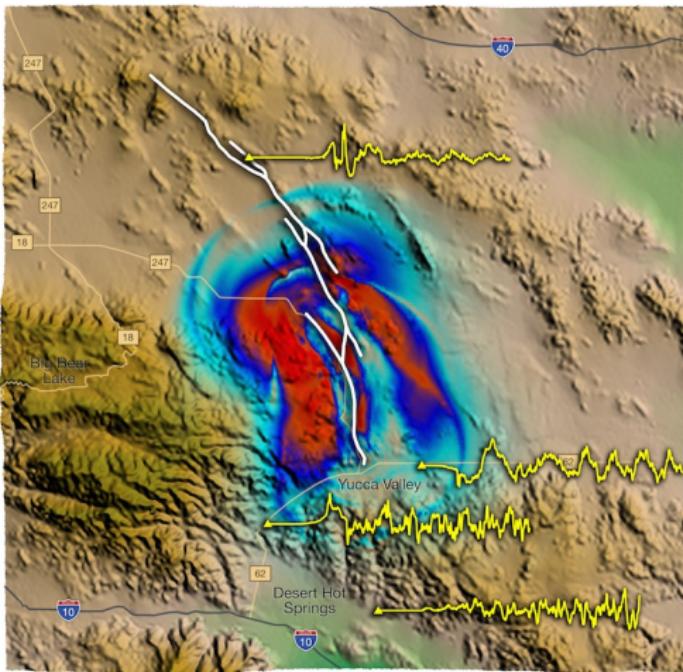
Landers Earthquake – Results

Observations:

- complex rupture dynamics (fault branching, etc.)
- high-frequency signals from rupture propagate directly into wave field
- synthetic seismograms with frequencies up to 10 Hz
- ground shaking in the engineering frequency band
- 42 s simulated time

SuperMUC Production Run:

- **1.25 PFLOPS** sustained performance
- 7 h 15 min computing time



Part II

Optimizing SeisSol for Petascale Seismic Simulations on SuperMUC, Stampede and Tianhe-2

Seismic Wave Propagation with SeisSol

Elastic Wave Equations: (velocity-stress formulation)

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

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

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & -\lambda - 2\mu & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\lambda & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\lambda & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\mu & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\mu \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\rho^{-1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\rho^{-1} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\rho^{-1} & 0 & 0 \end{pmatrix} \quad B = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\lambda & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\lambda & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\lambda & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\rho^{-1} & 0 & 0 & 0 & 0 & 0 \\ 0 & -\rho^{-1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -\rho^{-1} & 0 & 0 & 0 & 0 \end{pmatrix}$$

- high order discontinuous Galerkin discretisation
- **ADER-DG**: high approximation order in space and time:
- additional features: local time stepping, high accuracy of earthquake faulting (full frictional sliding)

→ Dumbser, Käser et al. [3,5]

SeisSol in a Nutshell – ADER-DG

Update scheme

$$\begin{aligned}
 Q_k^{n+1} = & Q_k - \frac{|S_k|}{|J_k|} M^{-1} \left(\sum_{i=1}^4 F^{-,i} I(t^n, t^{n+1}, Q_k^n) N_{k,i} A_k^+ N_{k,i}^{-1} \right. \\
 & + \sum_{i=1}^4 F^{+,i,j,h} I(t^n, t^{n+1}, Q_{k(i)}^n) N_{k,i} A_{k(i)}^- N_{k,i}^{-1} \Big) \\
 & + M^{-1} K^\xi I(t^n, t^{n+1}, Q_k^n) A_k^* \\
 & + M^{-1} K^\eta I(t^n, t^{n+1}, Q_k^n) B_k^* \\
 & + M^{-1} K^\zeta I(t^n, t^{n+1}, Q_k^n) C_k^*
 \end{aligned}$$

Cauchy Kovalewski

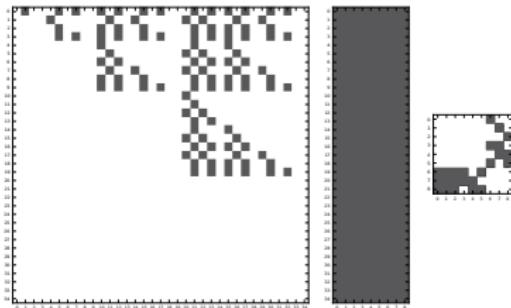
$$I(t^n, t^{n+1}, Q_k^n) = \sum_{j=0}^J \frac{(t^{n+1} - t^n)^{j+1}}{(j+1)!} \frac{\partial^j}{\partial t^j} Q_k(t^n)$$

$$(Q_k)_t = -M^{-1} ((K^\xi)^T Q_k A_k^* + (K^\eta)^T Q_k B_k^* + (K^\zeta)^T Q_k C_k^*)$$

Optimisation of Matrix Operations

Apply sparse matrices to multiple DOF-vectors Q_k

$$(K^\xi)^T \quad \frac{\partial^j}{\partial t^j} Q_k \quad A_k^*$$

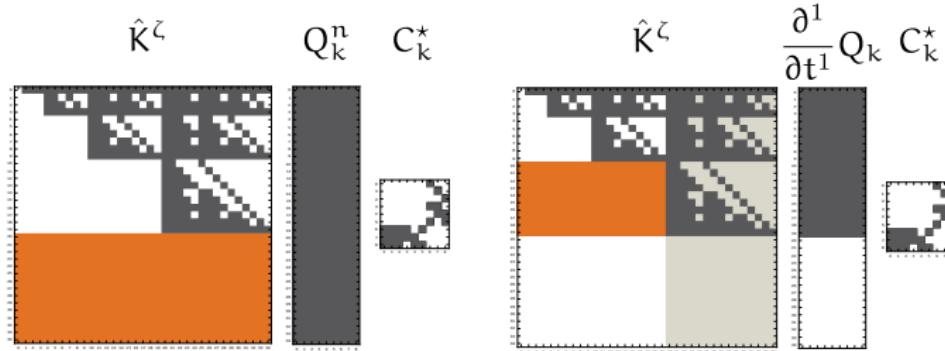


Code Generator for Sparse Kernels: (Breuer et al. [1])

- avoid overhead of CSR (or similar) data structures; store CSR elements vector, only
- full “unrolling” of all element operations using a code generator
- use intrinsics and apply blocking to improve vectorisation

Optimisation of Matrix Operations

Apply sparse matrices to multiple DOF-vectors Q_k



Dense vs. Sparse Kernels: (Breuer et al. [2])

- switch to dense kernels depending on achieved time to solution
- for sparse and dense kernels:
exploit zero-blocks generated during recursive CK computation

Performance Optimization

Switch between Sparse/Dense Kernels:

- auto-tuning approach on benchmark scenarios
- measure sparse vs. dense performance for each matrix
- select sparse vs. dense kernel based on best time to solution

order								boundary
2	sparse	13%						
3	sparse	sparse	dense	sparse	sparse	sparse	sparse	26%
4	sparse	sparse	dense	sparse	dense	dense	dense	17%
5	sparse	23%						
6	sparse	dense	dense	sparse	dense	dense	dense	9%

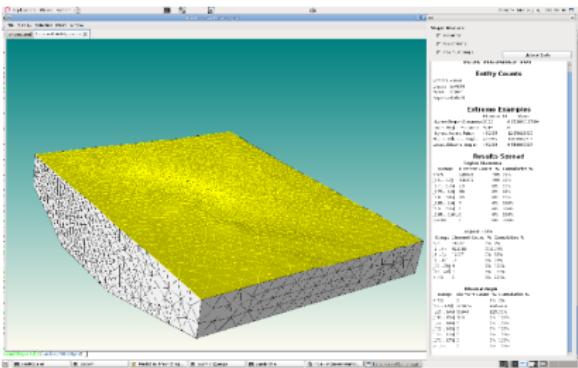
Hybrid MPI+OpenMP Parallelisation:

- careful OpenMP parallelisation of all parts (not only main kernels → communication buffers, etc.)
- targeted at manycore platforms, such as Intel Xeon Phi
- OpenMP improvements for Xeon Phi also lead to noticeable improvements for “standard” CPUs

Mesh Generation and Partitioning

Mesh Generation:

- high-quality meshes required (complicated fault structures, controllable mesh coarsening)
- with 10^8 – 10^9 grid cells
- using **SimModeler** by Simmetrix (<http://simmetrix.com/>)



Two-stage approach to provide parallel mesh partitions:

- graph-based partitioning (ParMETIS)
- create customised parallel format (based on netCDF) for mesh partitions
- highly scalable mesh input via netCDF/MPI-IO in SeisSol

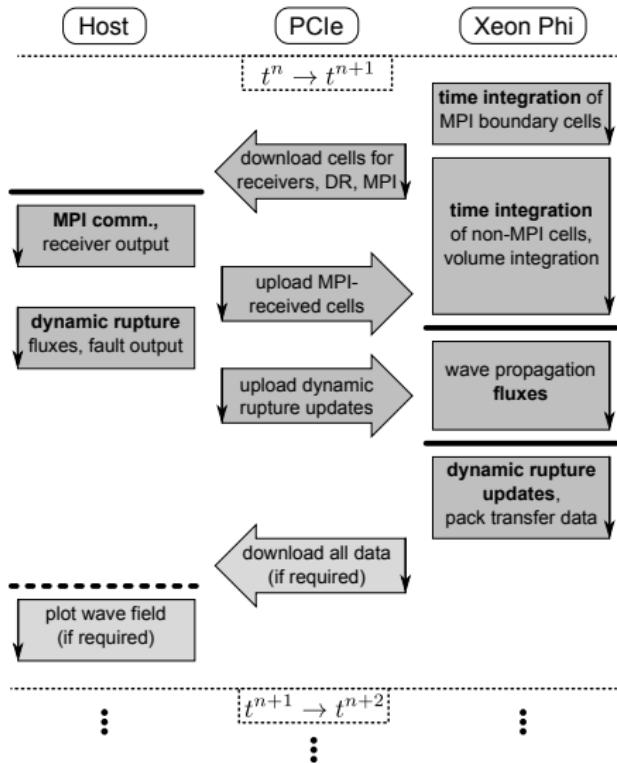
Optimization for Intel Xeon Phi Platforms

Offload Scheme:

- hides communication with Xeon Phi and between nodes
- use “heavy” CPU cores for dynamic rupture

Hybrid parallelism:

- on 1–3 Xeon Phis and host CPU(s)
- reflects multiphysics simulation
- manycore parallelism on Xeon Phi



Supercomputing Platforms

SuperMUC @ LRZ, Munich

- 9216 compute nodes (18 “thin node” islands)
- **147,456** Intel SNB-EP cores (2.7 GHz)
- Infiniband FDR10 interconnect (fat tree)
- #12 in Top 500: 2.897 PFlop/s



Stampede @ TACC, Austin

- 6400 compute nodes, **522,080 cores**
2 SNB-EP (8c) + **1 Xeon Phi SE10P** per node
- Mellanox FDR 56 interconnect (fat tree)
- #7 in Top 500: 5.168 PFlop/s

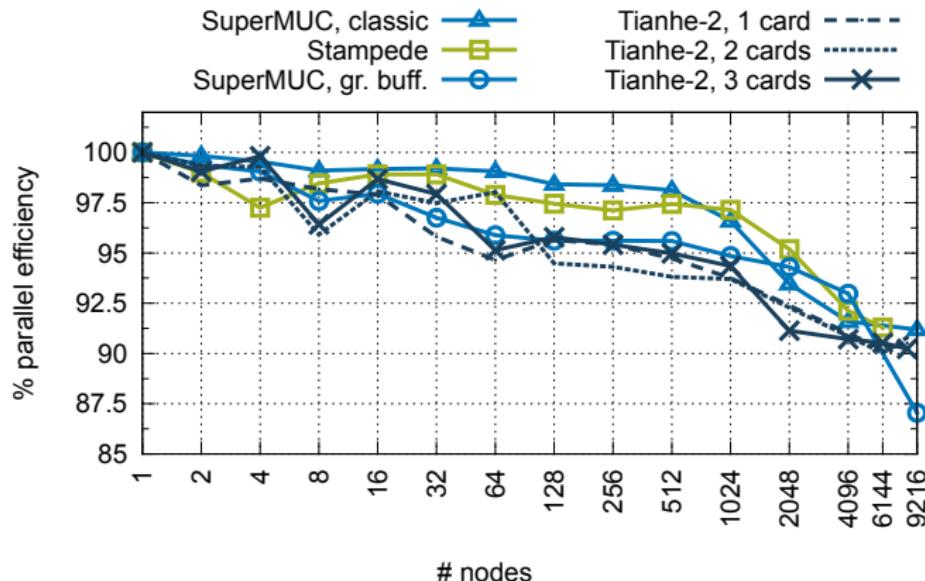


Tianhe-2 @ NSCC, Guangzhou

- 8000 compute nodes used, **1.6 Mio cores**
2 IVB-EP (12c) + **3 Xeon Phi 31S1P** per node
- TH2-Express custom interconnect
- #1 in Top 500: 33.862 PFlop/s

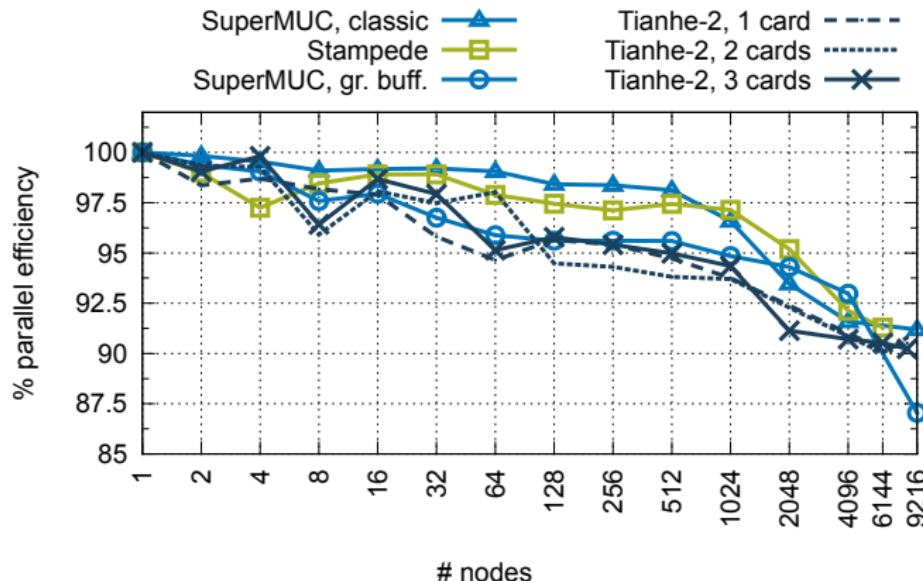


Weak Scaling of Wave Propagation



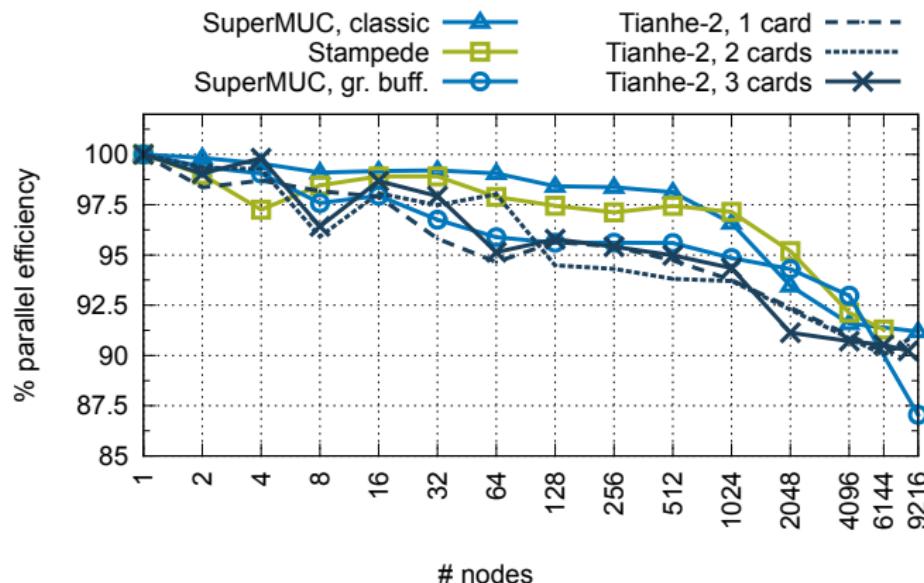
- goal: test scalability towards large problem sizes
- cubic domain, uniformly refined tetrahedral cells
- weak scaling: 400,000 elements per card/node

Weak Scaling of Wave Propagation



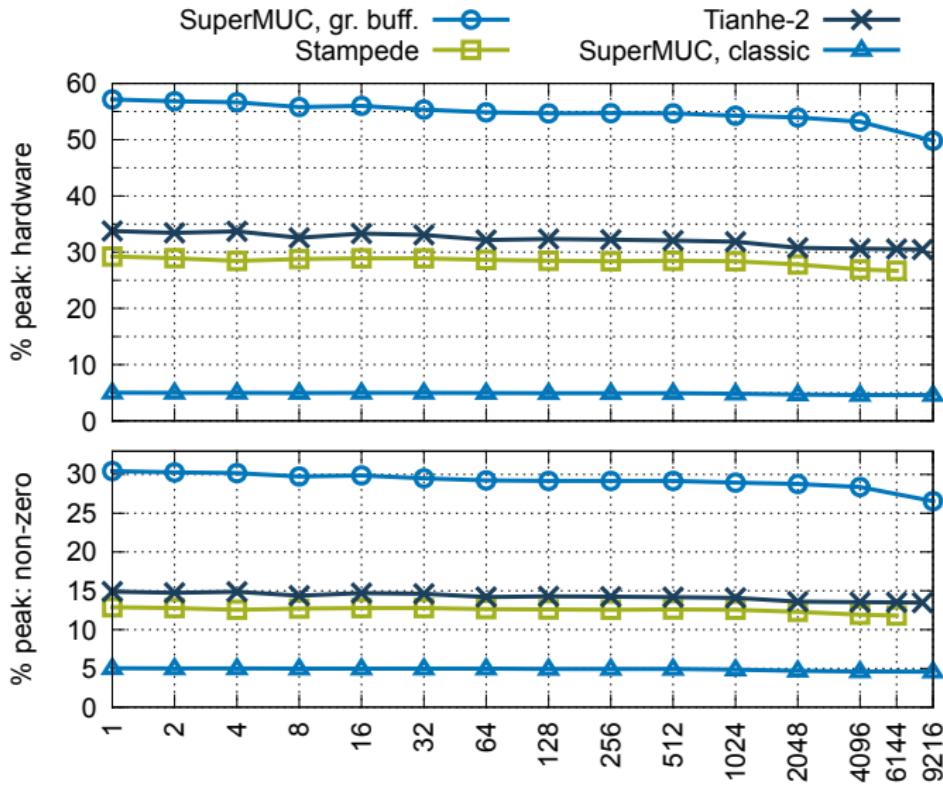
- more than 90 % parallel efficiency on Tianhe-2 and Stampede
- 87 % on full SuperMUC (no overlapping)

Weak Scaling of Wave Propagation

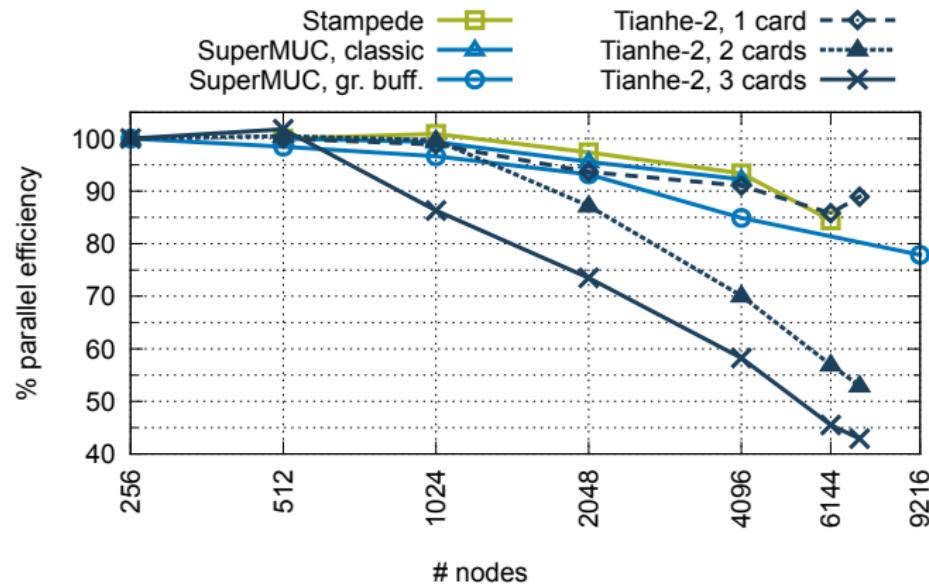


- 8.6 PFlop/s on Tianhe-2 (8000 nodes)
- 2.3 PFlop/s on Stampede (6144 nodes)
- 1.6 PFlop/s on SuperMUC (9216 nodes)

Weak Scaling – Peak Efficiency

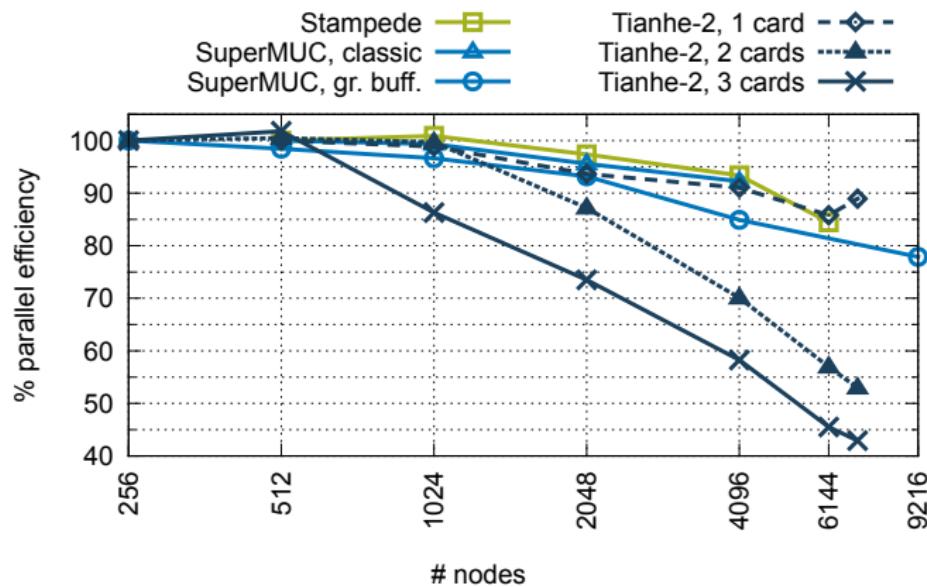


Strong Scaling of Landers Scenario



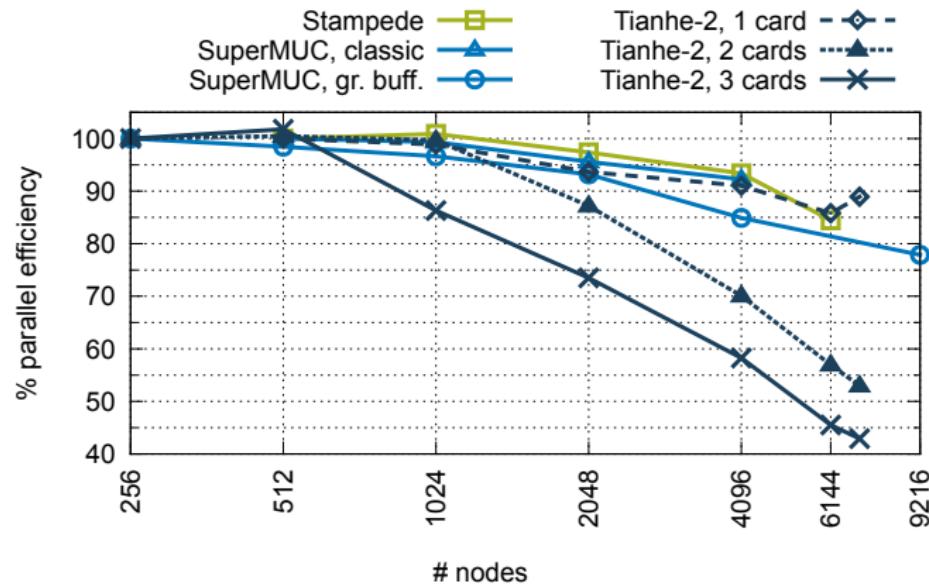
- 191 million tetrahedrons; 220,982 element faces on fault
- 6th order, 96 billion degrees of freedom

Strong Scaling of Landers Scenario



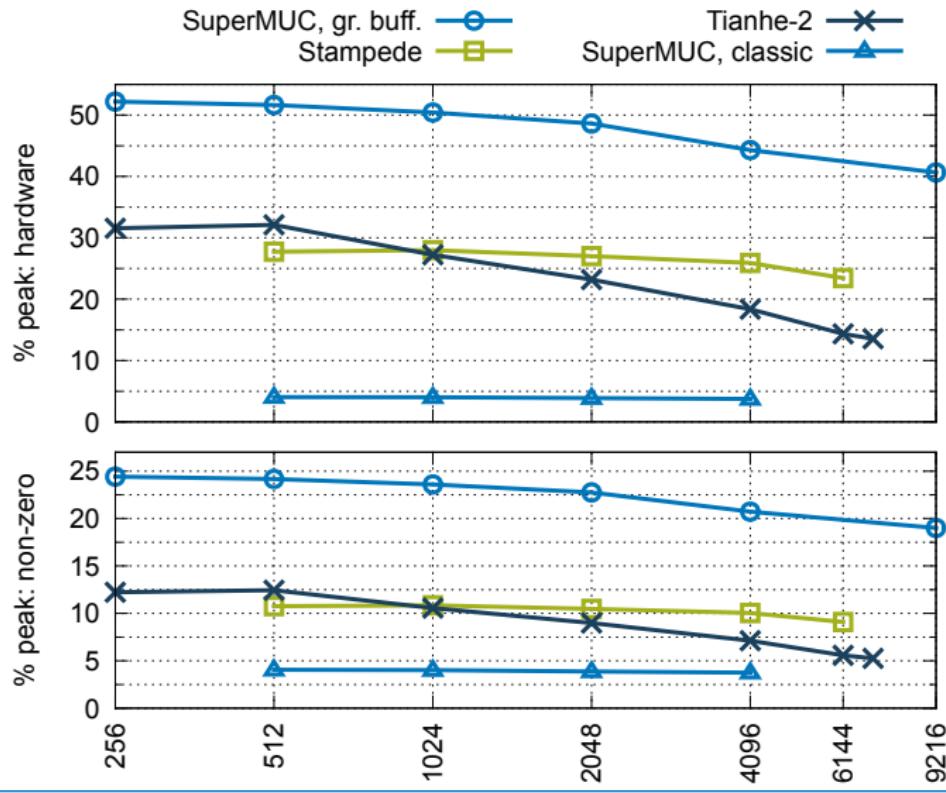
- more than 85 % parallel efficiency on Stampede and Tianhe-2 (when using only one Xeon Phi per node)
- multiple-Xeon-Phi performance suffers from MPI communication

Strong Scaling of Landers Scenario



- 3.3 PFlop/s on Tianhe-2 (7000 nodes)
- 2.0 PFlop/s on Stampede (6144 nodes)
- 1.3 PFlop/s on SuperMUC (9216 nodes)

Landers Strong Scaling – Peak Efficiency



SeisSol: Earthquake Simulation @ Petascale

Multiphysics Dynamic Rupture Simulations with SeisSol:

- high-order ADER-DG on complicated geometries
- non-linear interaction of rupture process and seismic waves
- physics-based seismic hazard analysis

Petascale Performance on Heterogeneous Platforms:

- scalable mesh input of more than 9 billion cells
- exploits high computational intensity of ADER-DG
- requires careful tuning of the entire simulation pipeline
- code generation to accelerate element kernels
- offload-scheme for multiphysics with Xeon Phi
- extreme (hybrid) parallelism with approx. 1.6 million cores

<http://seissol.geophysik.uni-muenchen.de/>

References

- [1] A. Breuer, A. Heinecke, M. Bader, C. Pelties: *Accelerating SeisSol by generating vectorized code for sparse matrix operators*. In: Advances in Parallel Computing 25, IOS Press, 2014. Proceedings of ParCo 2013
- [2] A. Breuer, A. Heinecke, S. Rettenberger, M. Bader, A.-A. Gabriel, C. Pelties: *Sustained Petascale Performance of Seismic Simulations with SeisSol on SuperMUC*. In: Supercomputing, LNCS 8488, p. 1–18. PRACE ISC Award 2014.
- [3] M. Dumbser, M. Käser: *An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes – II. The three-dimensional isotropic case*. Geophys. J. Int. 167(1), 2006.
- [4] A. Heinecke, A. Breuer, S. Rettenberger, M. Bader, A.-A. Gabriel, C. Pelties, A. Bode, W. Barth, X.-K. Liao, K. Vaidyanathan, M. Smelyanskiy, P. Dubey: *Petascale High Order Dynamic Rupture Earthquake Simulations on Heterogeneous Supercomputers*. Gordon Bell Prize Finalist 2014.
- [5] M. Käser, M. Dumbser, J. de la Puente, H. Igel: *An arbitrary high-order Discontinuous Galerkin method for elastic waves on unstructured meshes – III. Viscoelastic attenuation*. Geophys. J. Int. 168(1), 2007.
- [6] C. Pelties, J. de la Puente, J.-P. Ampuero, G. B. Brietzke, M. Käser: *Three-dimensional dynamic rupture simulation with a high-order discontinuous Galerkin method on unstructured tetrahedral meshes*. J. Geophys. Res.: Solid Earth, 117(B2), 2012.