EDGE: Extreme Scale Fused Seismic Simulations with the Discontinuous Galerkin Method

Alexander Breuer, Alexander Heinecke (Intel), Yifeng Cui

What is EDGE?

- **Extreme-scale Discontinuous Galerkin Environment (EDGE): Seismic wave propagation through DG-FEM**
- **Focus: Problem settings with high geometric complexity, e.g., mountain topography**
- **Written from scratch to support fused forward simulations**
- **“License”: BSD 3-Clause (software), CC0 for supporting files (e.g., user guide)**

http://dial3343.org

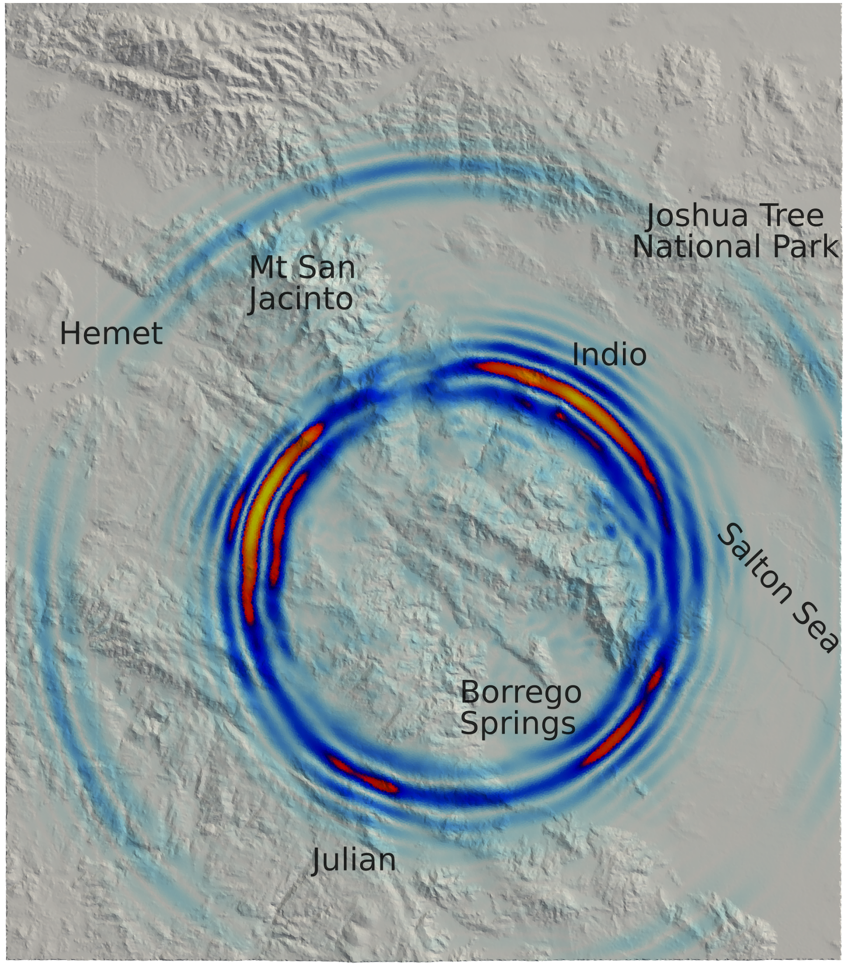

Example of hypothetical seismic wave propagation with mountain topography using EDGE. Shown is the surface of the computational domain covering the San Jacinto fault zone between Anza and Borrego Springs in California. Colors denote the amplitude of the particle velocity, where warmer colors correspond to higher amplitudes.

Getting Started: Advection Equation

- **"Simplest" hyperbolic Partial Differential Equation (PDE)**

- **Elastic wave equations similar: Linear system with variable coefficients**

$$q(x,t)_t + v \cdot q(x,t)_x = 0,\ v \in \mathbb{R}$$

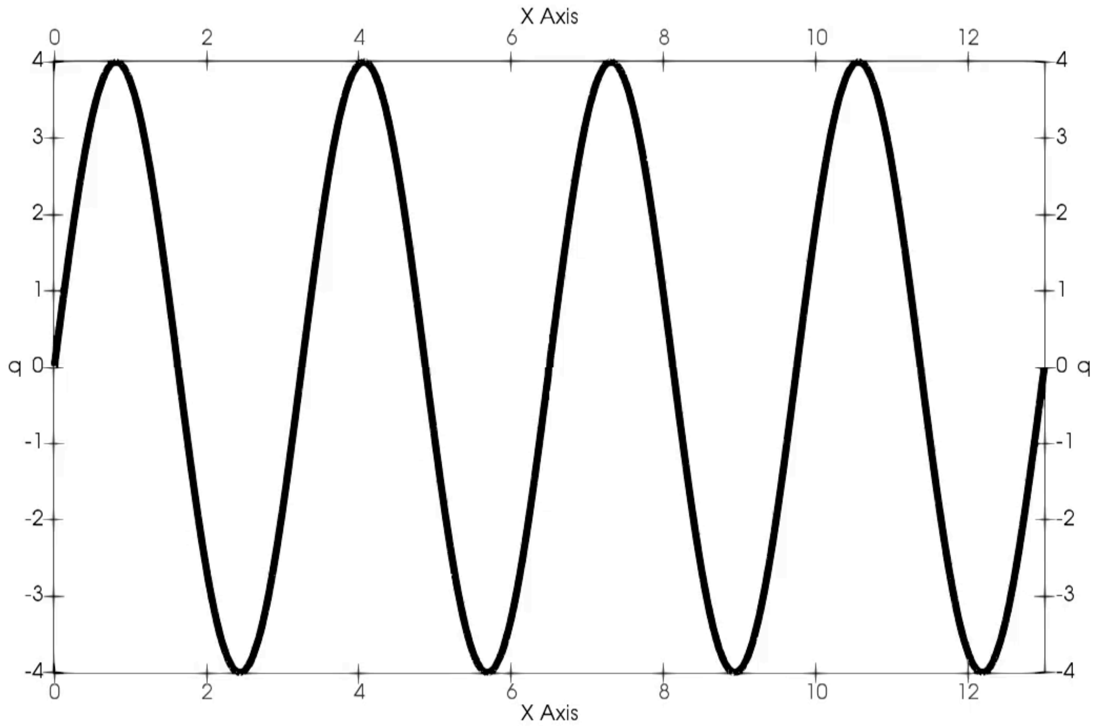

Illustration of EDGE's non-fused, third order (P2 elements) ADER-DG solver applied to the advection equation with sinusoidal initial values and periodic boundary conditions.

Getting Started: Fused Solver

- **Non-Fused:** $o = s(i)$

$$o_1 = s(i_1) \quad o_4 = s(i_4)$$
$$o_3 = s(i_3) \quad o_2 = s(i_2)$$

- **Fused:** $O_m = S_m(I_m)$

$$O_4 = (o_1, o_2, o_3, o_4) = S_4(I_4)$$
$$= S_4(i_1, i_2, i_3, i_4)$$

$$q(x,t)_t + v \cdot q(x,t)_x = 0,\ v \in \mathbb{R}$$

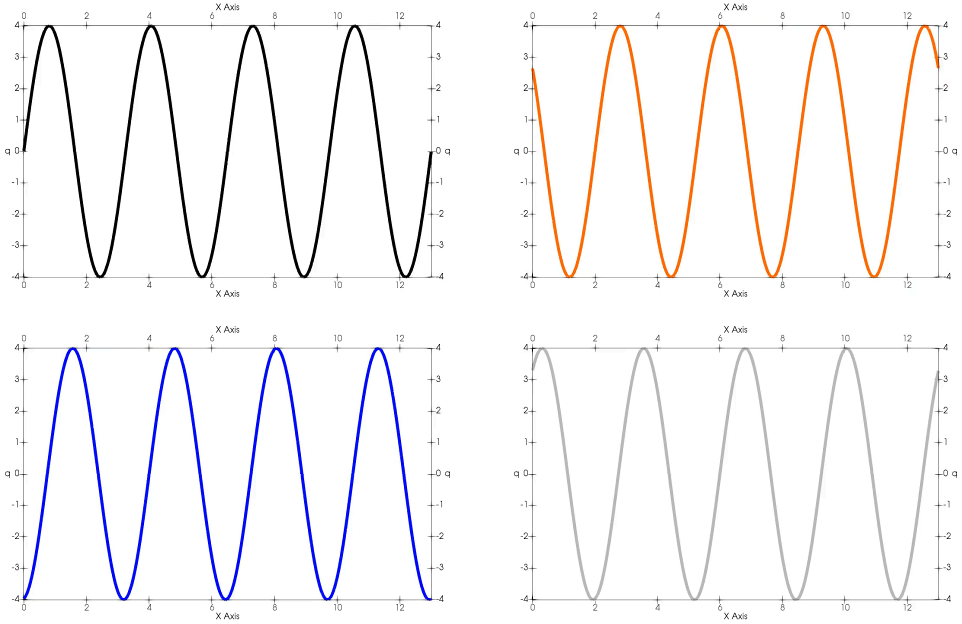

Illustration of EDGE's non-fused, third order (P2 elements) ADER-DG solver applied to the advection equation for four problem settings with sinusoidal initial values and periodic boundary conditions.

Getting Started: Fused Solver

- **Non-Fused:** $o = s(i)$

$$o_1 = s(i_1) \quad o_4 = s(i_4)$$
$$o_3 = s(i_3) \quad o_2 = s(i_2)$$

- **Fused:** $O_m = S_m(I_m)$

$$O_4 = (o_1, o_2, o_3, o_4) = S_4(I_4)$$
$$= S_4(i_1, i_2, i_3, i_4)$$

$$q(x,t)_t + v \cdot q(x,t)_x = 0,\ v \in \mathbb{R}$$

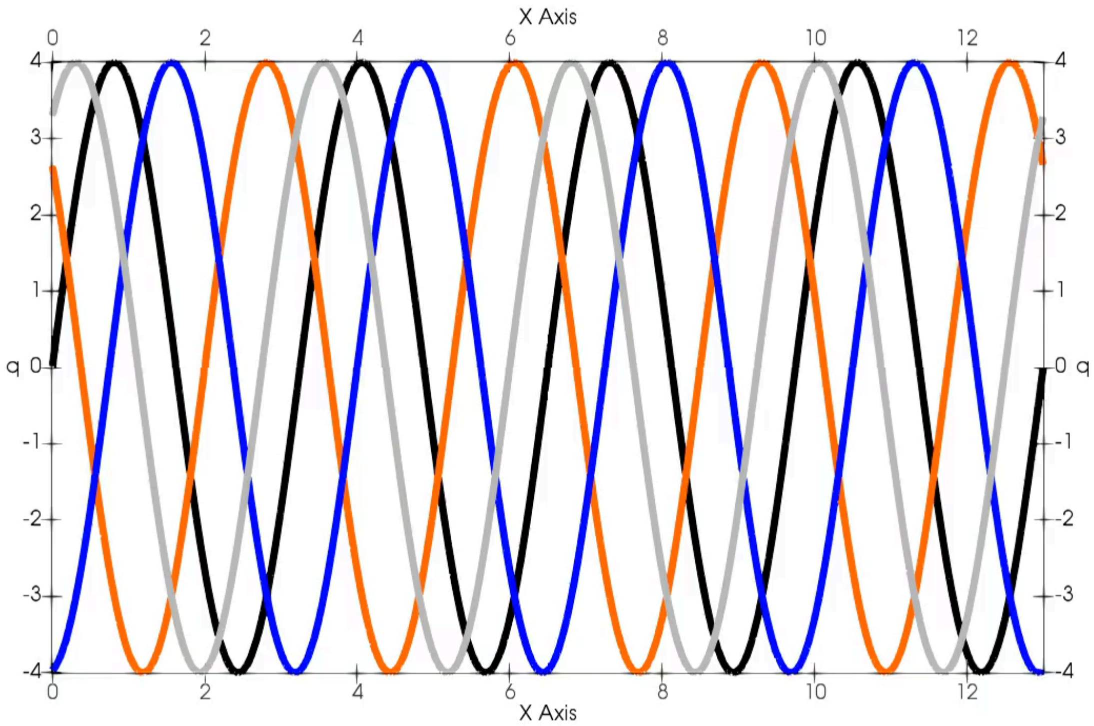

Illustration of EDGE's fused (4 simulations), third order (P2 elements) ADER-DG solver applied to the advection equation with sinusoidal initial values and periodic boundary conditions.

DOFs: Non-Fused vs. Fused

$$o = s(i) \qquad\qquad O_4 = S_4(I_4)$$

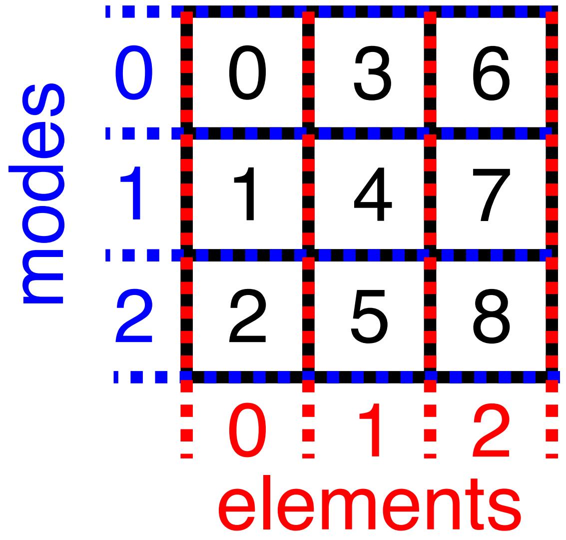

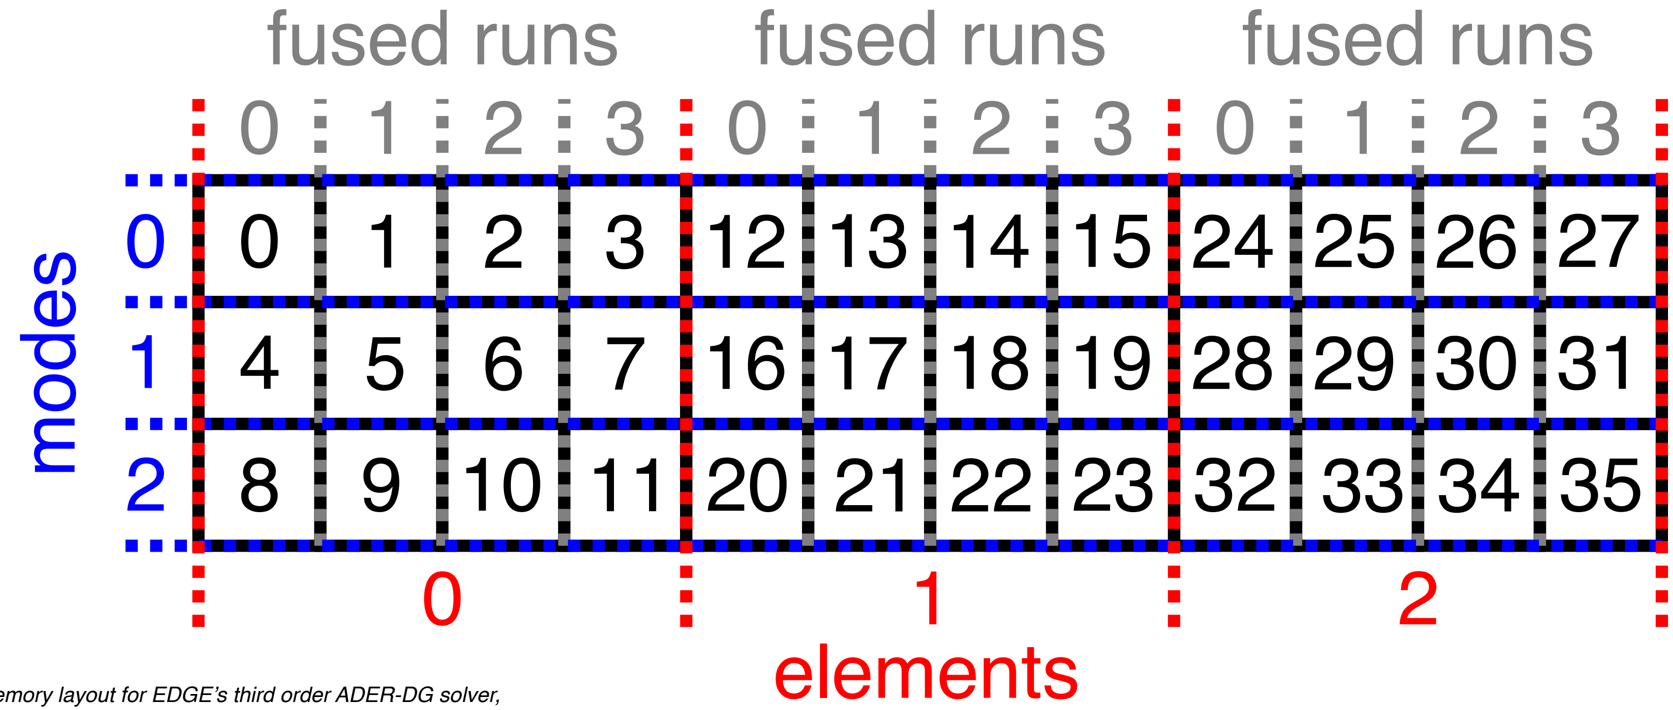

Illustration of the memory layout for EDGE's third order ADER-DG solver, line elements, and the advection equation (single quantity). Left: Non-fused memory layout, right: memory layout for 4 fused simulations.

Key Advantages

- Full vector operations, even for sparse matrix operators
- Automatic memory alignment
- Read-only data shared among all runs
- Lower sensitivity to latency (memory & network)

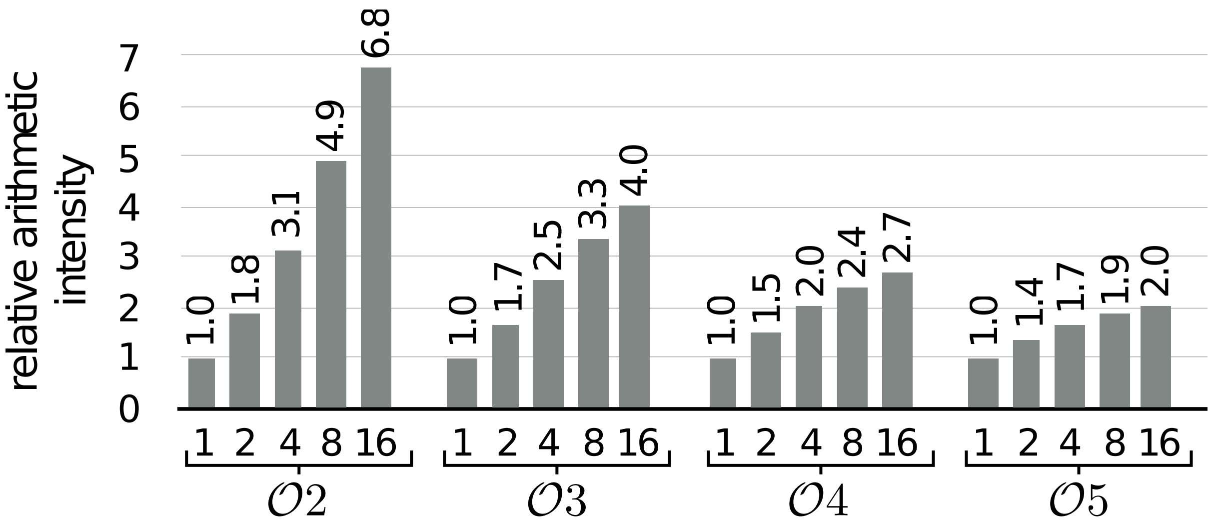

Relative arithmetic intensities. Shown are convergence rates 2-5 for the fusion of 2,4,8,16 simulations vs. a non-fused simulation for the elastic wave equations, using an ADER-DG solver. [ISC17]

“Similar Enough”: EDGE’s Approach

1. Identical mesh for all fused simulations
2. Identical simulations parameters:
    1. Start and end time
    2. Convergence rate
    3. “Frequency” of wave field output, “frequency” and location of seismic receivers
3. Identical material parameters (velocity model)
4. “Sources”:
    1. Arbitrary initial DOFs
    2. Kinematic sources: Fused or non-fused point sources
    3. Spontaneous rupture: Identical friction law, other parameters (e.g., nucleation, initial stresses, coefficients) arbitrary

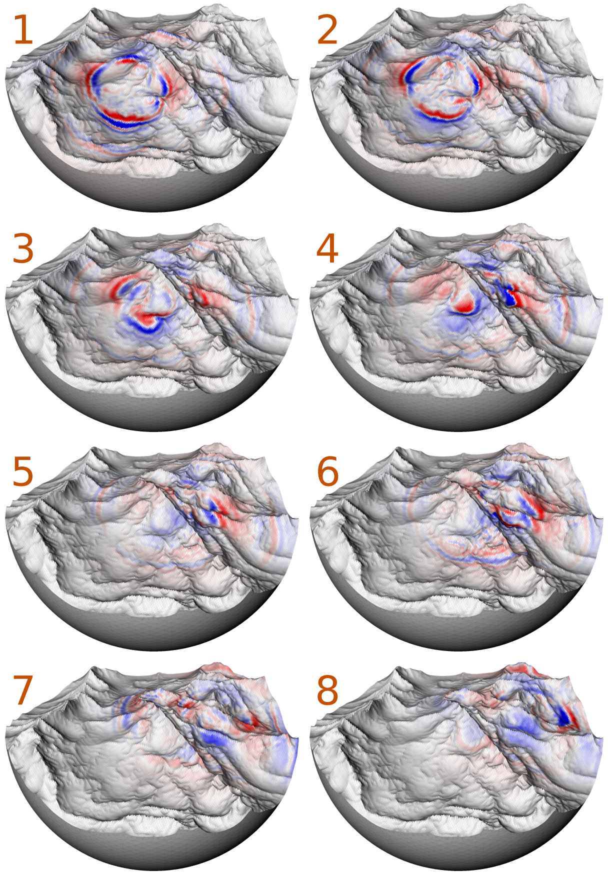

Illustration of the wave field for an exemplary fusion of eight simulations in EDGE with eight point sources at different locations.

Performance: LOH.1

- Layer Over Halfspace (LOH.1): Benchmark used for code verification
- Orders: 2-6 (non-fused), 2-4 (fused)
- Unstructured tetrahedral mesh: 350,264 elements
- Single node of Cori-II (68 core Intel Xeon Phi x200, code-named Knights Landing)
- EDGE vs. SeisSol (GTS, git-tag 201511)

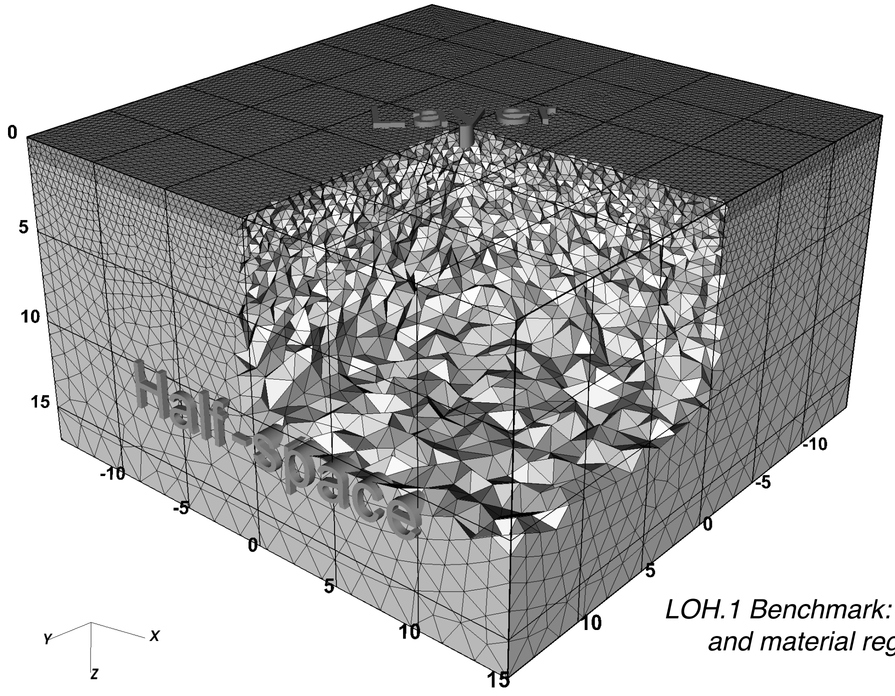

LOH.1 Benchmark: Example mesh and material regions [ISC16_1]

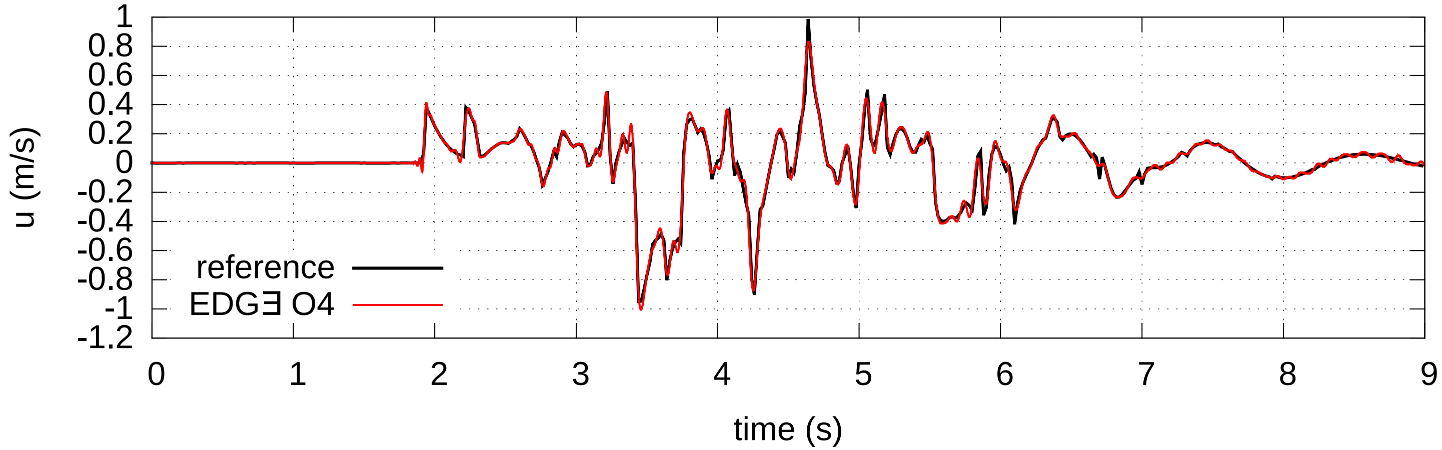

Synthetic seismogram of EDGE for quantity u at the ninth seismic receiver located at (8647 m, 5764 m, 0) in red. The reference solution is shown in black. Detailed setup: [ISC17]

Fused Simulations: Speedup

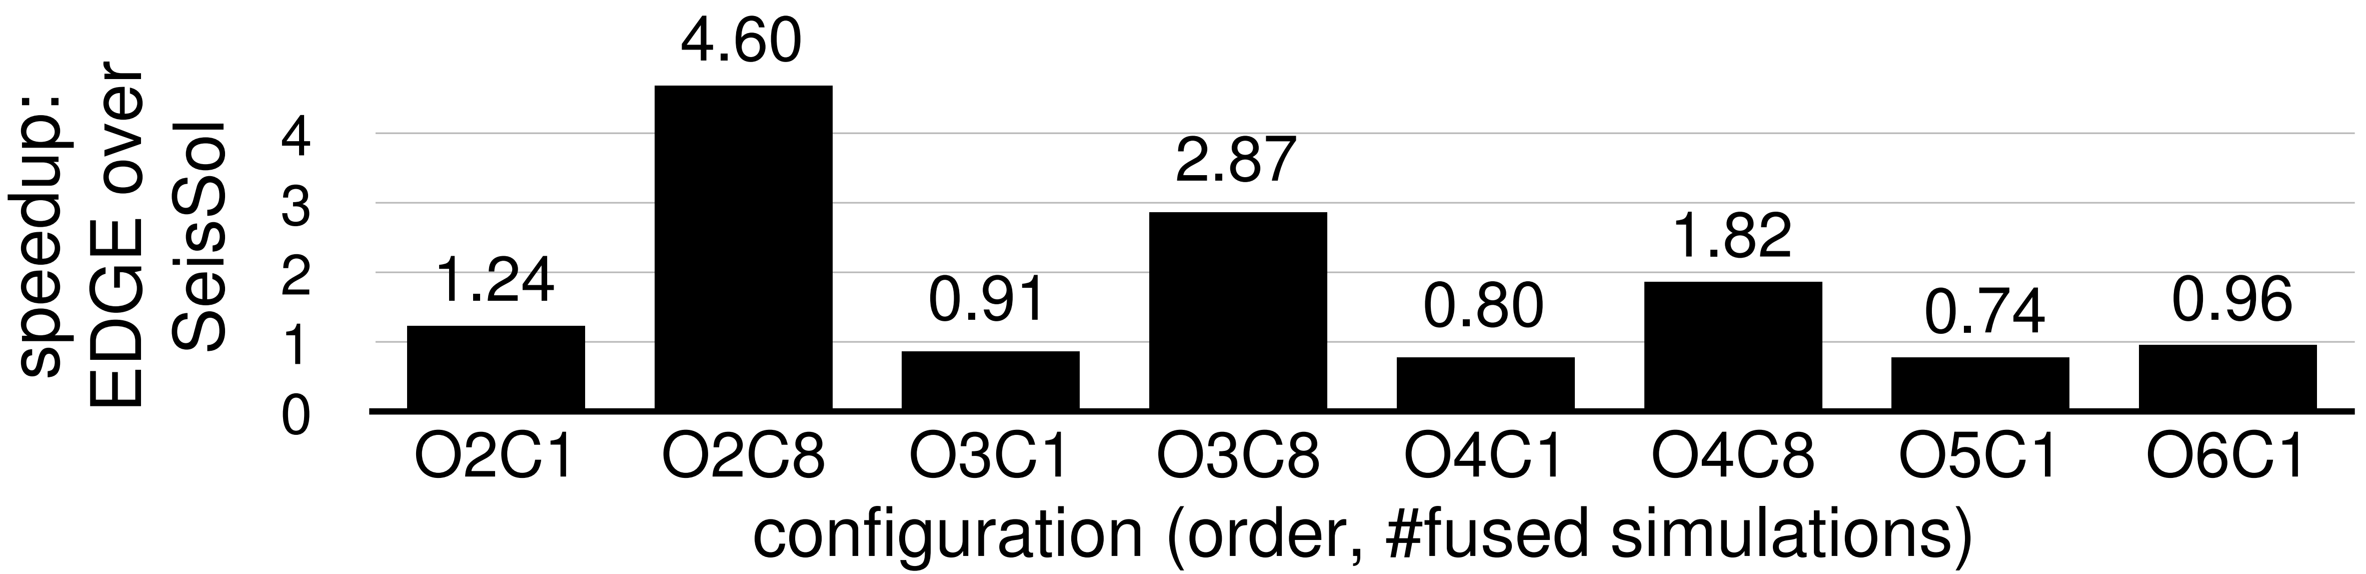

Speedup of EDGE over SeisSol (GTS, git-tag 201511). Convergence rates O2 – O6: single non-fused forward simulations (O2C1-O6C1). Additionally, per-simulation speedups for orders O2–O4 when using EDGE's full capabilities by fusing eight simulations (O2C8-O4C8). [ISC17]

Weak: Setup

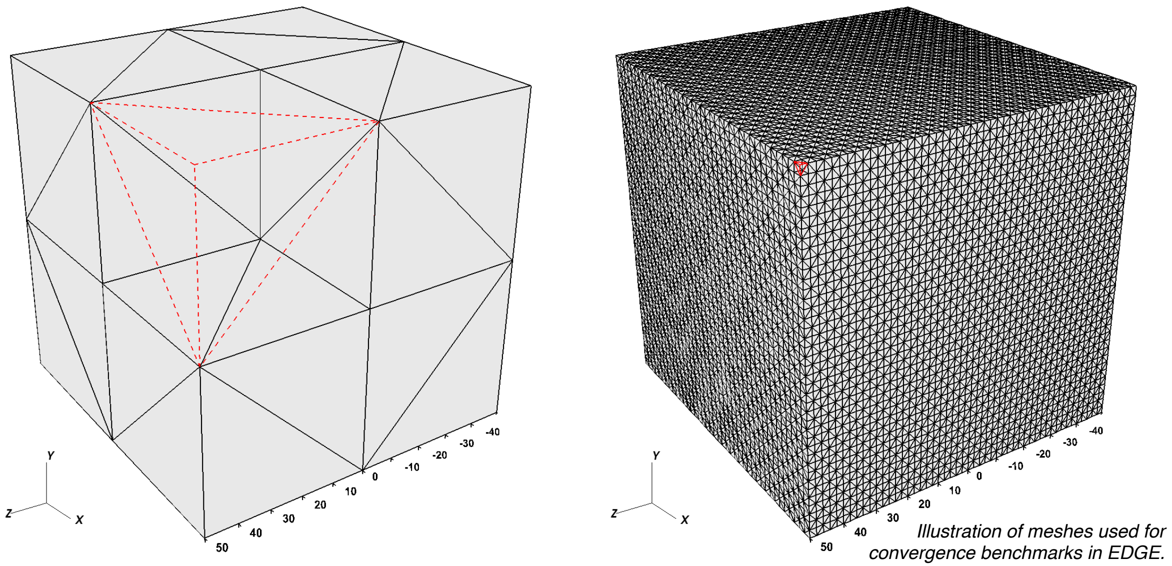

Illustration of meshes used for convergence benchmarks in EDGE.

- Regular cubic mesh, 5 Tets per Cube, 4th order (P3) and 6th order (P5)
- Imitates convergence benchmark
- 276K elements per node
- 1-9000 nodes of Cori-II (9000 nodes = 612,000 cores)

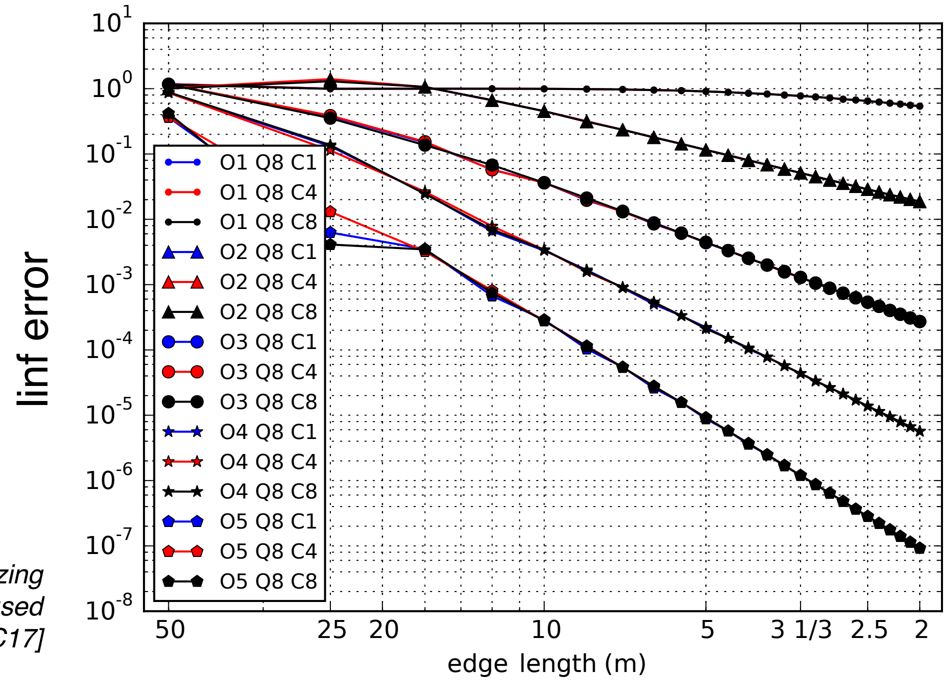

Convergence of EDGE in the L∞-norm. Shown are orders O1 – O5 for quantity v (Q8) when utilizing EDGE's fusion capabilities with shifted initial conditions. For clarity, from the total of eight fused simulations, only errors of the first (C1), fourth (C4) and last simulation (C8) are shown. [ISC17]

Weak: Results

- **O6C1 @ 9K nodes: 10.4 PFLOPS (38% of peak)**
- **O4C8 vs. O4C1 @ 9K nodes: 2.0x speedup**

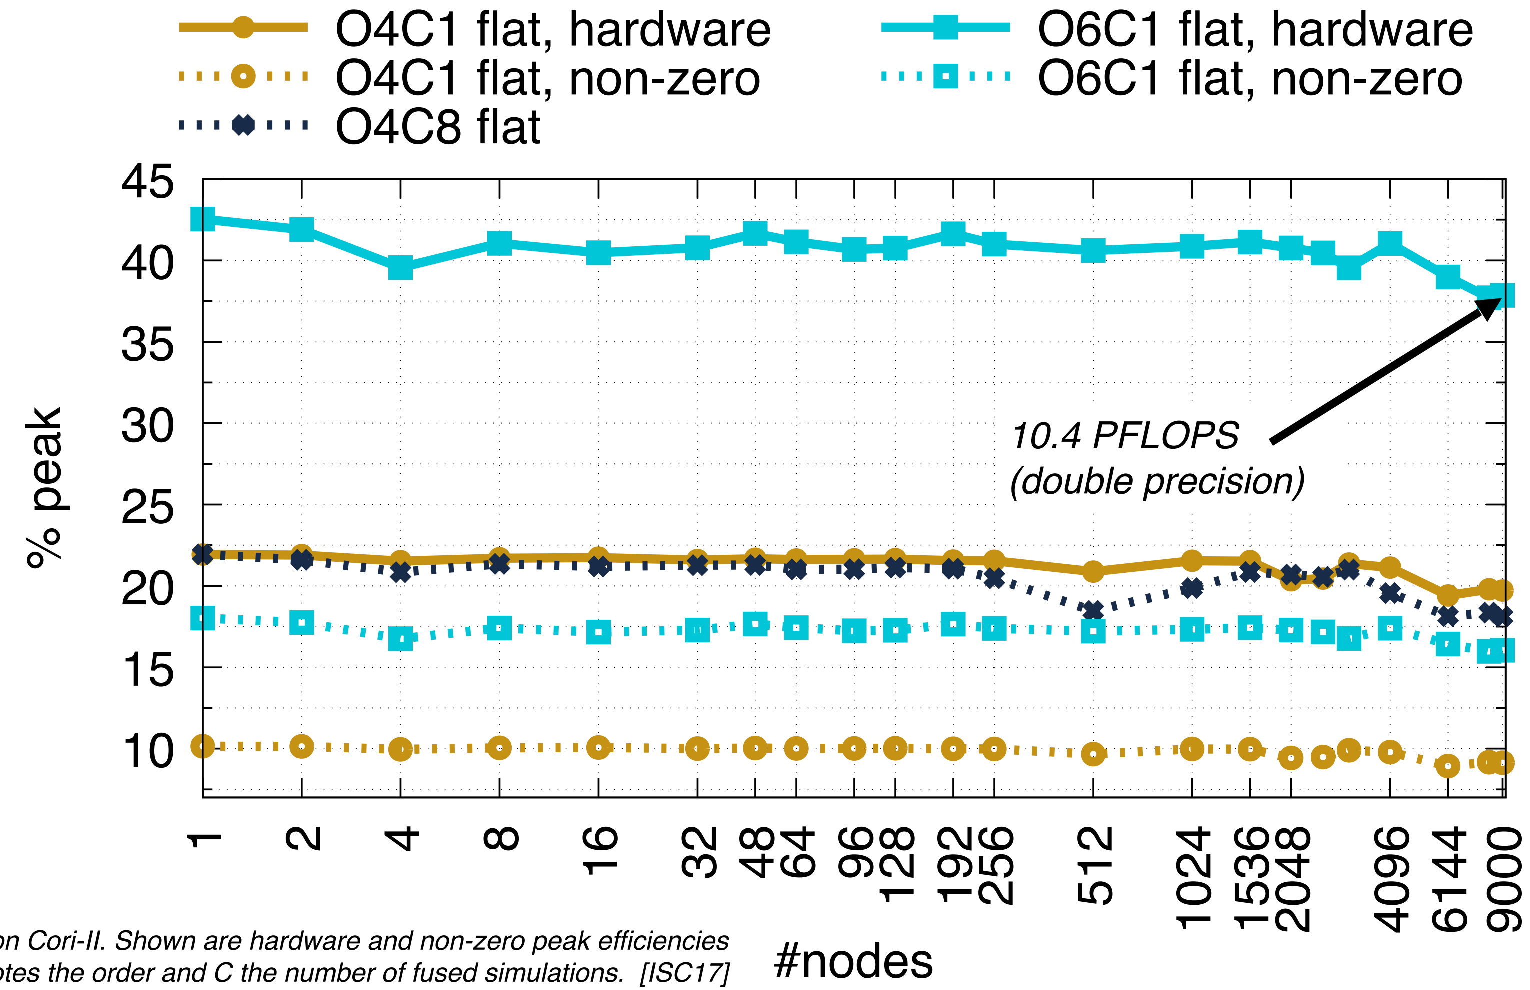

Weak scaling study on Cori-II. Shown are hardware and non-zero peak efficiencies in flat mode. O denotes the order and C the number of fused simulations. [ISC17]

Strong: LOH.1

- **Orders: 4 & 6 (non-fused), 4 (fused)**
- **Unstructured tetrahedral mesh: 172,386,915 elements**
- **32-3200 nodes of Theta (64 core Intel Xeon Phi x200, code-named Knights Landing)**
- **3200 nodes = 204,800 cores**

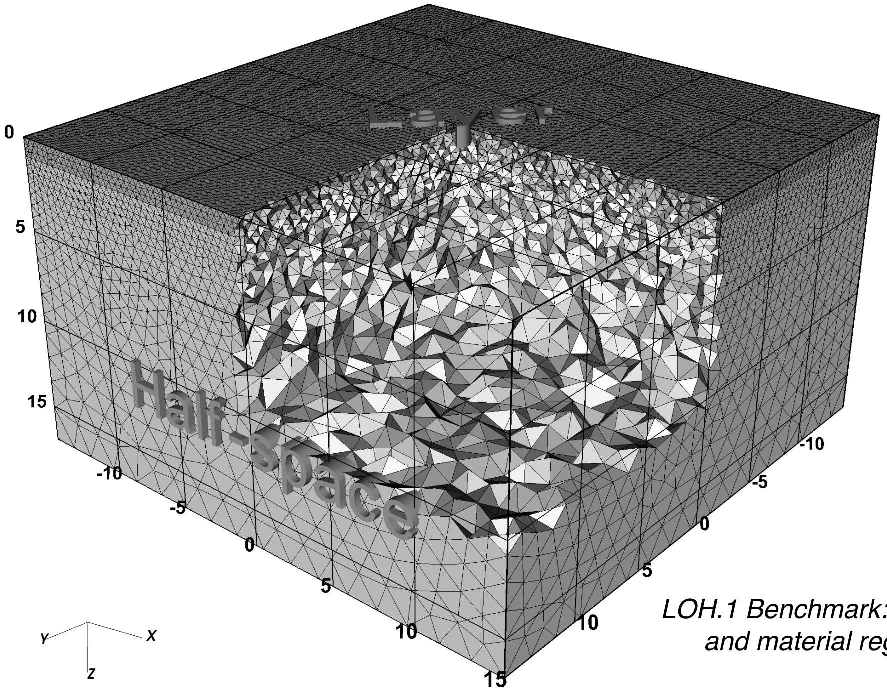

LOH.1 Benchmark: Example mesh and material regions [ISC16_1]

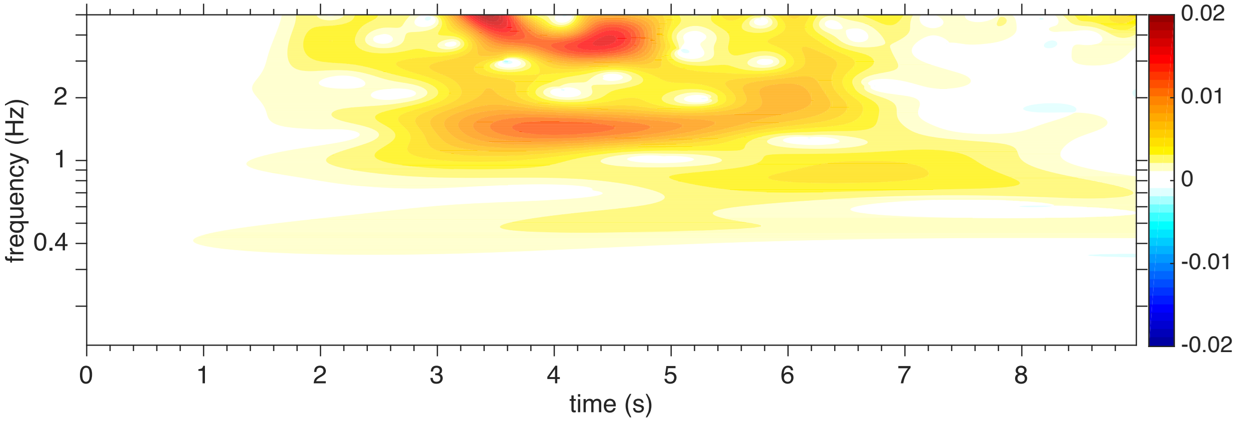

Time-frequency misfit for quantity u at the ninth seismic receiver located at (8647 m, 5764 m, 0) and in a frequency range between 0.13Hz and 5Hz. Detailed setup: [ISC17], Visualization: TF-MISFIT_GOF_CRITERIA, http://nuquake.eu

Strong: Results

- **O6C1 @ 3.2K nodes: 3.4 PFLOPS (40% of peak)**
- **O4C8 vs. O4C1 @ 3.2K nodes: 2.0x speedup**

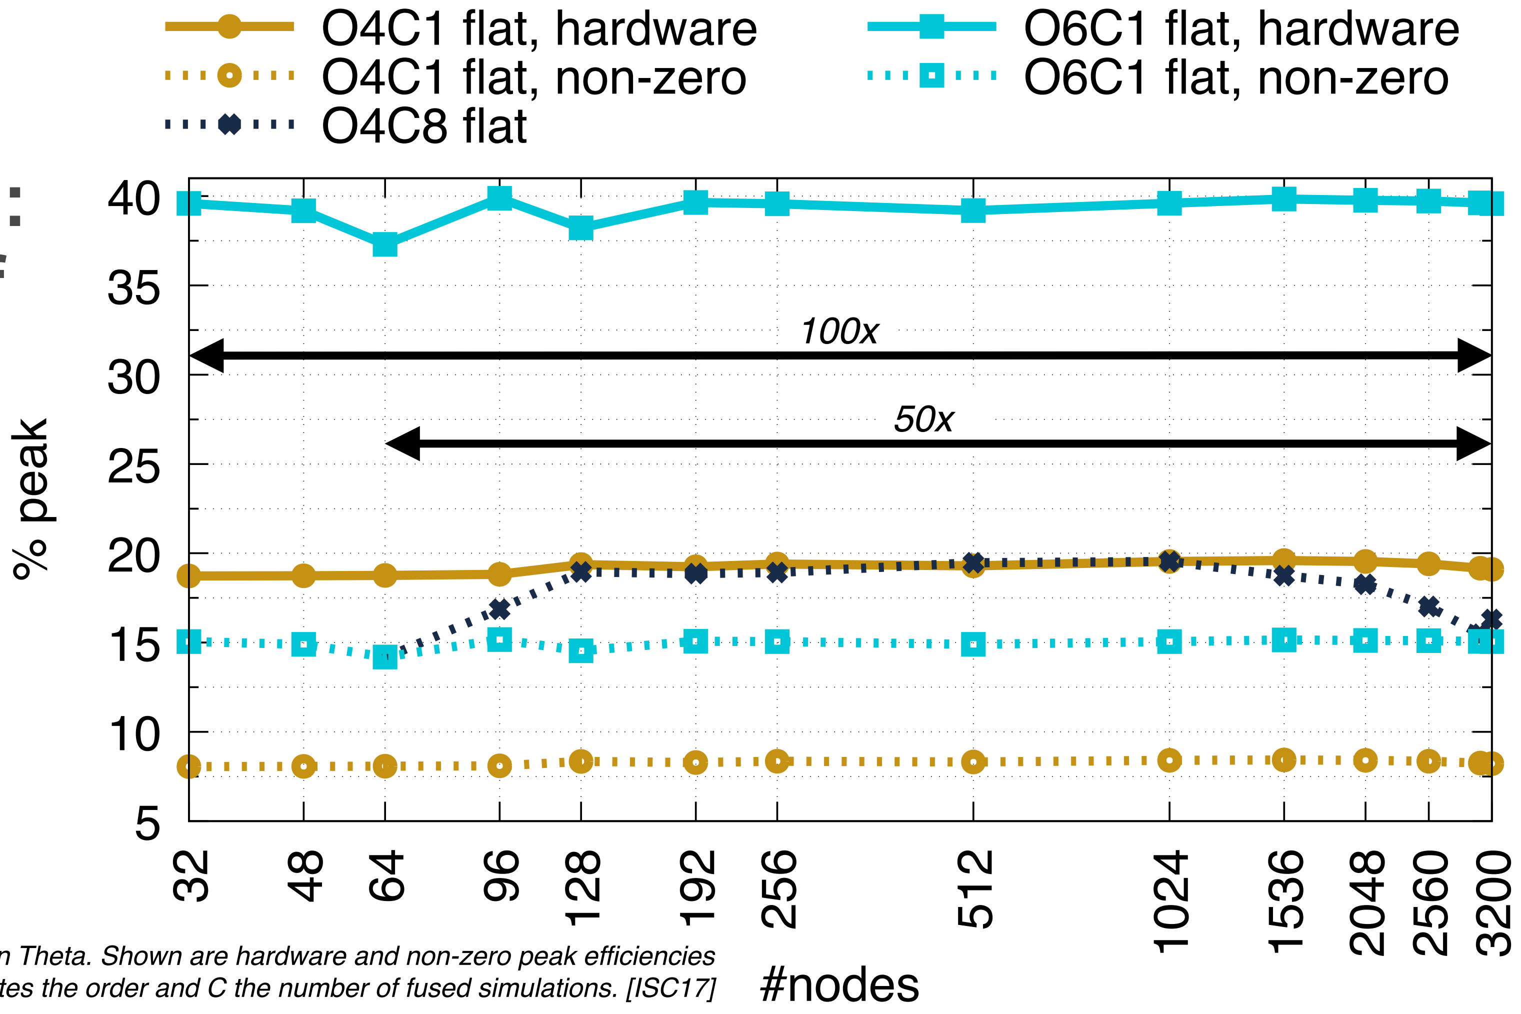

Strong scaling study on Theta. Shown are hardware and non-zero peak efficiencies in flat mode. O denotes the order and C the number of fused simulations. [ISC17]

EDGE: Current and Upcoming

- Elements: Line, rectangular quads, 3-node triangles, rectangular hexes, <u>4-node tets</u>
- Equations: Advection (FV+ADER-DG: 1D, 2D, 3D), Shallow Water (FV: 1D), <u>Elastic Wave Equations</u> (FV+ADER-DG: 2D, <u>3D</u>)
- Parallelization: Assembly kernels for WSM, SNB, HSW, KNC (non-fused), <u>KNL (fused & non-fused)</u>, OpenMP (custom), MPI (overlapping)
- Continuity: Continuous Integration (sanity checks), Continuous Delivery incl. automated convergence + benchmarks runs, automated code coverage, automated license checks, container bootstrap
- “License”: BSD 3-Clause (software), CC0 for supporting files (e.g., user guide)

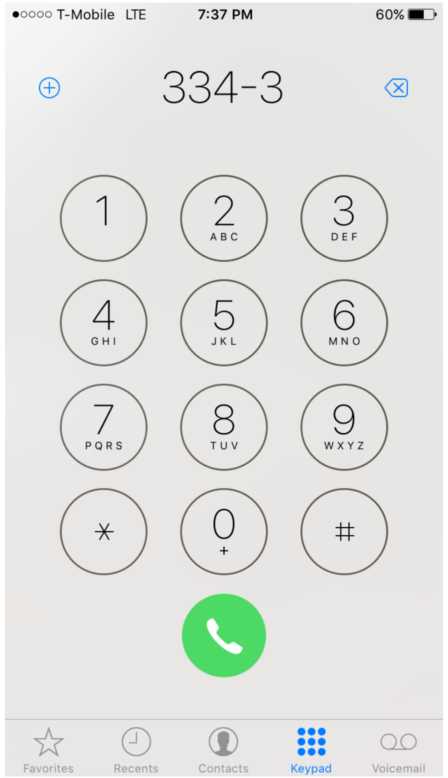

http://dial3343.org

- Sparse, fused assembly kernels for orders 5+ (✓)
- Kinematic Sources (Standard Rupture Format): Support for fused and non-fused source (✓) descriptions
- Spontaneous Rupture Simulations (✓)
- Grouped Local Time Stepping
- EDGEcut: Automated surface and volume meshing

References

- [ISC17] A. Breuer, A. Heinecke, Y. Cui: EDGE: Extreme Scale Fused Seismic Simulations with the Discontinuous Galerkin Method. To appear in proceedings of International Super Computing ISC High Performance, available online during the conference
- [ISC16_1] A. Heinecke, A. Breuer, M. Bader: High Order Seismic Simulations on the Intel Xeon Phi Processor (Knights Landing). High Performance Computing: 31st International Conference, ISC High Performance 2016, Frankfurt, Germany, June 19-23, 2016, Proceedings. http://dx.doi.org/10.1007/978-3-319-41321-1_18
- [ISC16_2] A. Heinecke, A. Breuer, M. Bader: Chapter 21 - High Performance Earthquake Simulations. In Intel Xeon Phi Processor High Performance Programming Knights Landing Edition.
- [IPDPS16] A. Breuer, A. Heinecke, M. Bader: Petascale Local Time Stepping for the ADER-DG Finite Element Method. In Parallel and Distributed Processing Symposium, 2016 IEEE International. http://dx.doi.org/10.1109/IPDPS.2016.109
- [ISC15] A. Breuer, A. Heinecke, L. Rannabauer, M. Bader: High-Order ADER-DG Minimizes Energy- and Time-to-Solution of SeisSol. In 30th International Conference, ISC High Performance 2015, Frankfurt, Germany, July 12-16, 2015.
- [SC14] A. Heinecke, A. Breuer, S. Rettenberger, M. Bader, A.-A. Gabriel, C. Pelties, A. Bode, W. Barth, X.-K. Liao, K. Vaidyanathan, M. Smelyanskiy and P. Dubey: Petascale High Order Dynamic Rupture Earthquake Simulations on Heterogeneous Supercomputers. In Supercomputing 2014, The International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE, New Orleans, LA, USA, November 2014. Gordon Bell Finalist.
- [ISC14] A. Breuer, A. Heinecke, S. Rettenberger, M. Bader, A.-A. Gabriel and C. Pelties: Sustained Petascale Performance of Seismic Simulations with SeisSol on SuperMUC. In J.M. Kunkel, T. T. Ludwig and H.W. Meuer (ed.), Supercomputing — 29th International Conference, ISC 2014, Volume 8488 of Lecture Notes in Computer Science. Springer, Heidelberg, June 2014. 2014 PRACE ISC Award.
- [PARCO13] A. Breuer, A. Heinecke, M. Bader and C. Pelties: Accelerating SeisSol by Generating Vectorized Code for Sparse Matrix Operators. In Parallel Computing — Accelerating Computational Science and Engineering (CSE), Volume 25 of Advances in Parallel Computing. IOS Press, April 2014.

Acknowledgements

*Only the great support of experts at **NERSC** and **ALCF** made our extreme-scale results possible. In particular, we thank **J. Deslippe, S. Dosanjh, R. Gerber,** and **K. Kumaran**. This work was supported by the Southern California Earthquake Center (**SCEC**) through contribution #16247. This work was supported by the **Intel Parallel Computing Center** program. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. This work used the Extreme Science and Engineering Discovery Environment (**XSEDE**), which is supported by National Science Foundation grant number ACI-1053575. This research is part of the **Blue Waters** sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.*

EDGE heavily relies on contributions of many authors to open-source software.
*This software includes, but is not limited to: **ASan** (https://clang.llvm.org/docs/AddressSanitizer.html, debugging), **Catch** (https://github.com/philsquared/Catch, unit tests), **CGAL** (http://www.cgal.org, surface meshes), **Clang** (https://clang.llvm.org/, compilation), **Cppcheck** (http://cppcheck.sourceforge.net/, static code analysis), **Easylogging++** (https://github.com/easylogging/, logging), **GCC** (https://gcc.gnu.org/, compilation), **gitbook** (https://github.com/GitbookIO/gitbook, documentation), **Gmsh** (http://gmsh.info/, volume meshing), **GoCD** (https://www.gocd.io/, continuous delivery), **jekyll** (https://jekyllrb.com, homepage), **libxsmm** (https://github.com/hfp/libxsmm, matrix kernels), **MOAB** (http://sigma.mcs.anl.gov/moab-library/, mesh interface), **ParaView** (http://www.paraview.org/, visualization), **pugixml** (http://pugixml.org/, XML interface), **SCons** (http://scons.org/, build scripts), **Valgrind** (http://valgrind.org/, memory debugging), **Visit** (https://wci.llnl.gov/simulation/computer-codes/visit, visualization).*