SDSC
SAN DIEGO SUPERCOMPUTER CENTER

UC San Diego

EDGE: Benchmarking the Seismic Wave Propagation Solver

Alexander Breuer, Alexander Heinecke (Intel), Yifeng Cui

Eleventh U.S. National Conference on Earthquake Engineering
Los Angeles, California
June 25-29, 2018

11NCEE LA2018

Extreme-scale Discontinuous Galerkin Environment (EDGE)

- Focus: Problem settings with high geometric complexity, e.g., mountain topography
- Unique support for fused simulations exploiting inter-simulation parallelism
- Rapid prototyping through support for: Line elements, quads, triangles, hexes, tets
- Parallelization: Assembly kernels for Sandy Bridge, Bulldozer, Haswell, Knights Corner, Knights Landing, EPYC, Skylake, Knights Mill; OpenMP+MPI (custom and overlapping)
- World record seismic wave propagation performance: 10.4 DP-PFLOPS on Cori Phase II
- Continuity: Continuous Integration (sanity checks), Continuous Delivery (automated convergence + benchmarks runs), code coverage, license checks, container bootstrap
- License: BSD 3-Clause (software), CC0 for supporting files, e.g., user guide

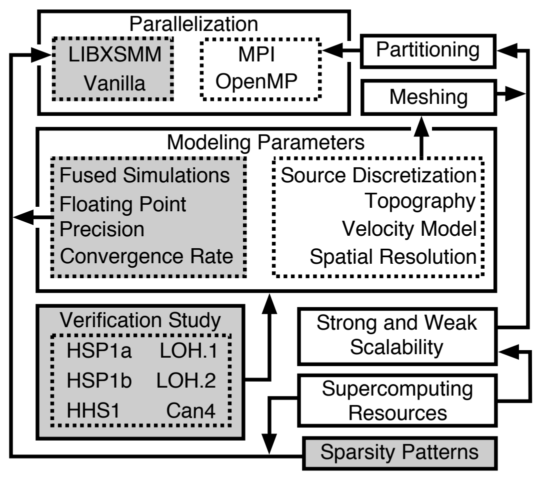

Illustration of EDGE's modeling and simulation pipeline. The parallelization is influenced by a multitude of factors. The presented verification study is key in determining modeling parameters. Together with other parts of the pipeline, highlighted in gray, the study's results drive EDGE's single-core parallelization.

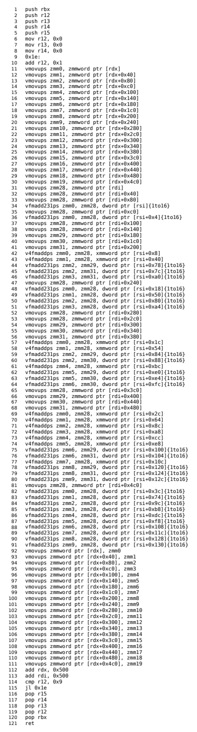

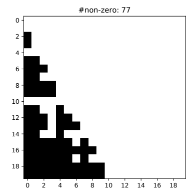

Top: One out of three stiffness matrices for tetrahedral P3 elements in EDGE. Only 77 entries of the 20x20=400 entries are non-zero.

Left: Generated assembly sparse matrix-tensor kernel for the shown stiffness matrix (top), 16 fused simulations, Knights Mill, and single precision arithmetic. The kernel reaches over 35% of non-zero peak utilization, when running out of a hot cache.

Right: EDGE's homepage http://dial3343.org and GitHub repository.

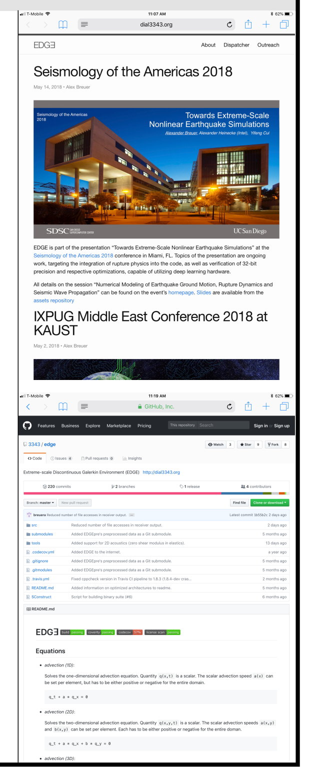

Verification study

Benchmark Problems

- Garvin's problem (2D): Acoustic and elastic, explosive point source
- Gaussian Hill and Canyon (2D): Acoustic and elastic, explosive point source
- HSP1a: Homogeneous space, near receivers
- HSP1b: Homogeneous space, distant receivers
- HHS1: Homogeneous halfspace
- LOH.1: Layer over halfspace, point source
- LOH.2: Layer over halfspace, kinematic rupture
- Can4: Simplified model of the Mygdonian basin, homogeneous layers inside the basin

Modeling and Simulation Pipeline

Preprocessing

1) Unstructured Meshing
2) Partitioning and mesh-reordering
3) EDGE XML-configs
4) Source discretization in Standard Rupture Format

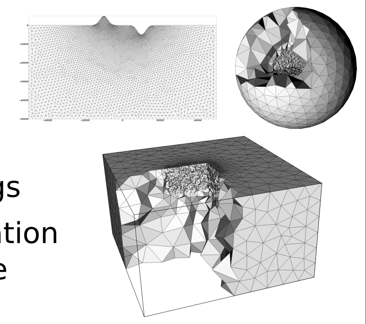

Simulation

1) Environment setup on supercomputer (here: Cori and Stampede 2)
2) Manual or automated job submission and monitoring

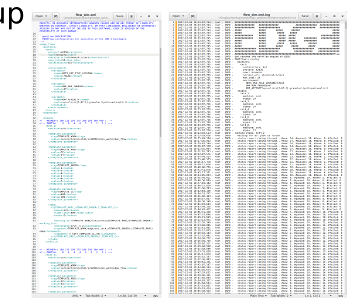

Postprocessing and Analysis

1) Manual or automated visualization and analysis of wave field and/or synthetic seismograms
2) CC0 annotation of data (“no copyright reserved”) and upload to public assets repository: http://opt.dial3343.org
3) Preparation and upload of data to SelSmic MOdeling Web INterfacE (if applicable): http://sismowine.org/

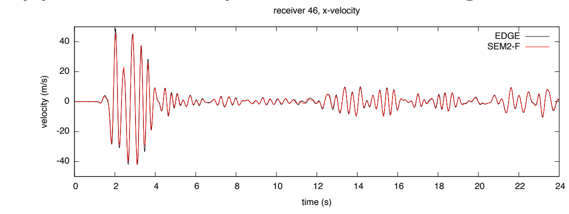

Comparison of EDGE's solution (black) to the reference solution SEM2-F [3] (red) for the first 24s of the Can4 benchmark. The receiver is above the wedge, 1km away from the northern edge of the valley.

Fused Simulations Exploit Inter-Simulation Parallelism

"Why is this a good idea?"

- Idea: Exploit input parallelism by fusing multiple, similar simulations in a single execution of the solver
- EDGE supports this idea at all levels of parallelism, starting at a single vector op
- Fusing multiples of the vector-width (KNM, SKX, KNM): 16 simulations in single precision) allows for perfect vectorization without zero ops
- Fusion of multiples of 64 bytes (16 simulations) leads to alignment to cache-lines without artificial zero-padding
- Read-only data structures are shared among all fused simulations

"Similar simulations?"

- EDGE imposes restrictions on fused seismic simulations:
  1) Identical mesh for all fused simulations
  2) Identical simulation parameters: start and end time, convergence rate, "frequency" of wave field output, "frequency" and location of seismic receivers
  3) Identical material parameters (velocity model)
  4) "Sources" mostly arbitrary: Arbitrary initial DOFs, kinematic sources: arbitrary location and moment rates, spontaneous rupture: identical friction law, other initial parameters arbitrary

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

Illustration of EDGE's fused approach (4 simulations), applied to the advection equation with sinusoidal initial values. While a traditional solver handles different initial values in multiple executions, EDGE exploits the input-parallelism and computes the 4 simulations in parallel.

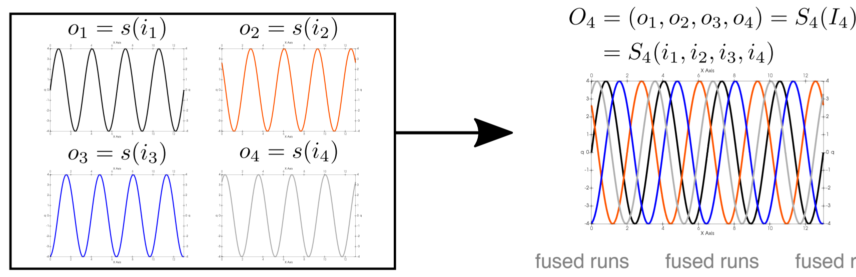

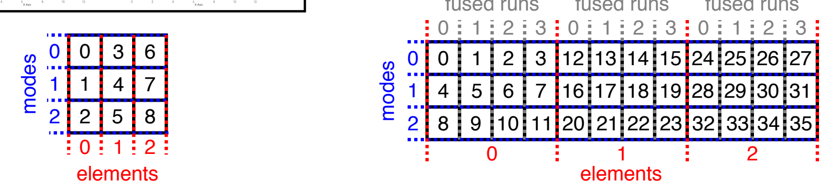

Illustration of EDGE's memory layout for a third order (P2 elements) ADER-DG solver for the advection equation. In the case of the traditional, non-fused approach, the DG-modes are the fastest dimension in memory, followed by the elements. In case of fused simulations, the fused runs are the fastest dimension, followed by the DG-modes and the elements.

Implications for High Performance Computing

- Hardware architectures are moving to an era of free computations and expensive data movement
- Deep Learning is the new number one driver for hardware developments of all major vendors
- Deep Learning hardware increases computation-intensive, reduced precision dense linear algebra capabilities, while sacrificing double precision performance
- Verification study shows: Single precision arithmetic is sufficient for seismic wave propagation using ADER-DG
- Fused simulation technology is critical to exploit single precision performance: EDGE increases throughput by 4.2x over SeisSol [3] (16 Knights Mill processors, GTS, 16 fused FP32 simulations, fifth order)

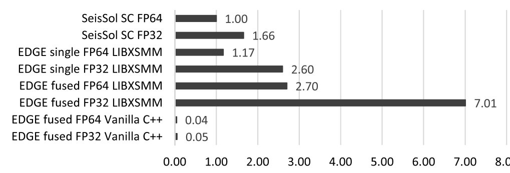

Left: Speed-up over the SC 2017 best paper award-winning work [4] on Knights Mill using fifth order.

Right: Knights Landing vs. Knights Mill VPU: a symmetric, single-pumped combo VPU is replaced by an asymmetric (single precision biased) VPU which is double-pumped for high efficiencies on the two-issue wide Xeon Phi frontend.

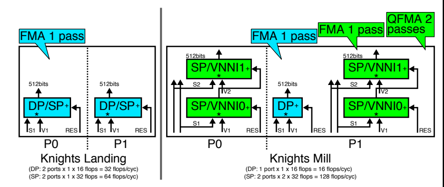

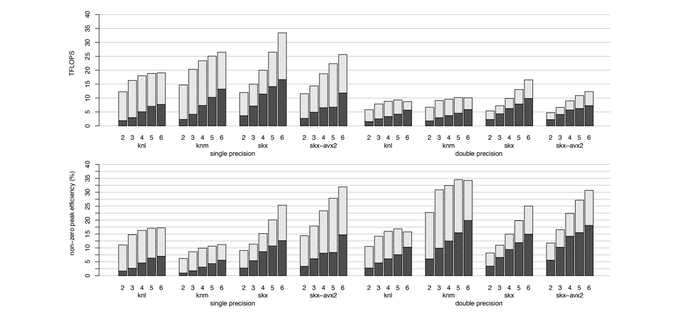

Floating point performance and peak efficiency on 16 nodes. All results are reported in terms of non-zero operations. Dark gray bars represent the performance of single seismic forward simulations, while light gray bars show the performance of fused simulations. Eight simulations are fused in double precision and sixteen in single precision arithmetic. Dimensions: a) order of convergence, b) computer architecture: Intel Xeon Phi 7250 (knl), Intel Xeon Phi 7295 (knm), and 2× Intel Scalable Xeon 8180 using AVX512 (skx) and AVX2 (skx-avx2), c) floating point precision.

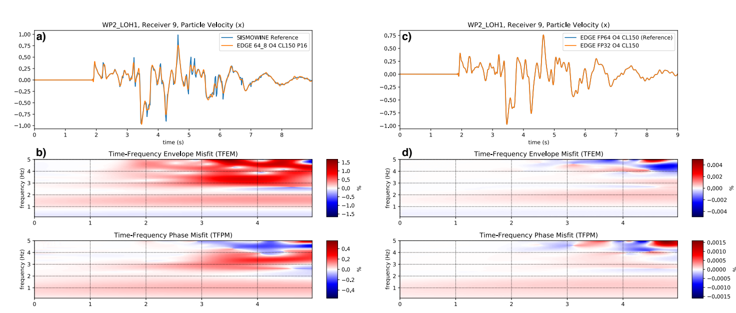

Exemplary illustration of EDGE’s fourth order solution for the ninth receiver and quantity u of the LOH.1 wave propagation benchmark. Plots a) and b) show a comparison to the reference, using double precision arithmetic. Out of the eight fused, identical solutions of the setting, only the first one is shown. Plots c) and d) show a comparison of the almost identical single and double precision results, obtained when using a single forward simulation. Dueto the low misfits, shown in d), the FP32 in and FP64 solutions are visually indistinguishable in the raw synthetics c).

Motivation of the Study: From Convergence to Verification

- Benchmarking is key to assess the accuracy of seismic wave propagation solvers
- EDGE has a multitude of modeling parameters:
  1) Fused vs. non-fused simulations
  2) Single vs. double precision
  3) Convergence rate in space and time
  4) Feature- and velocity-aware mesh refinement
  5) Source discretization
  6) Topography: Flat vs. DEM-derived
  7) Velocity model: Layered vs. data-input
- Choosing the right modeling parameters is crucial for best time-to-solution
- Our verification study considers the entire modeling and simulation pipeline and covers essential modeling decision (1-3 above) for best practices

Two regular tetrahedral meshes. By imposing periodic boundary conditions, the plane wave initial solution is reproduced after diagonal propagation through the domain, even in MPI-parallel settings (right).

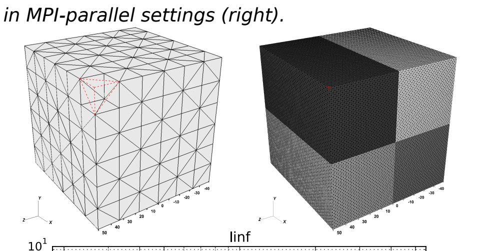

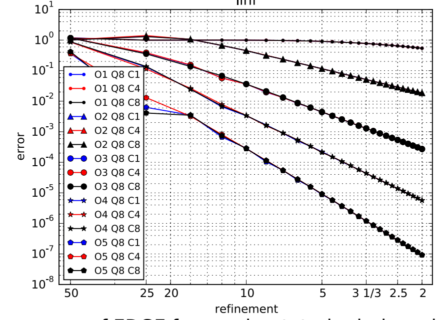

Convergence of EDGE for regular, tetrahedral meshes in the Linf-norm for the elastic wave equations in velocity-stress form. Shown are orders O1-O5 for quantity v (Q8) when utilizing EDGE’s fusion capabilities with shifted initial conditions. [1]

Convergence → Verification

LOH.1 Benchmark: Example mesh and material regions.

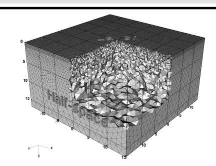

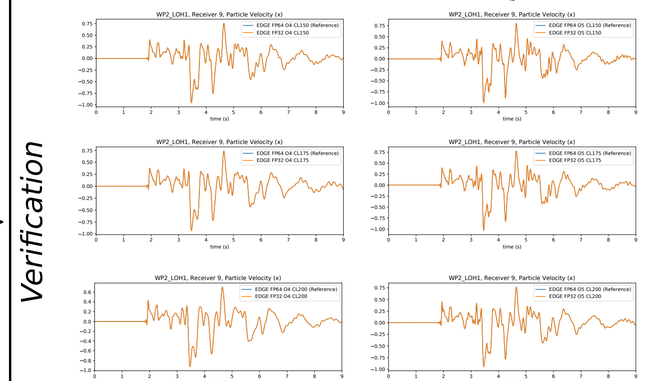

Synthetic seismograms for the ninth receiver and quantity u of the LOH.1 wave propagation benchmark. Shown is a comparison of single and double precision arithmetic in EDGE. The left synthetics show fourth order runs, the ones on the right fifth order runs. The two plots on top used a specified characteristic length of 150m, the middle plots 175m, and the bottom plots 200m. While higher resolutions and orders increase the accuracy, 32bit and 64bit precision are visually indistinguishable.

Outlook: A Comprehensive Approach

- EDGE-V: Velocity-aware mesh refinement using UCVM shear wave velocities, mesh annotation with velocity model
- EDGEcut: Feature-preserving surface meshing of mountain topography and intersecting fault geometries
- Convergence and verification study of 16bit precision for seismic ADER-DG simulations
- DG subcell-limiter to cope with large gradients in the solution, enabling EDGE for nonlinear hyperbolic PDEs
- Last but not least: If you are interested in working with us, get in touch!

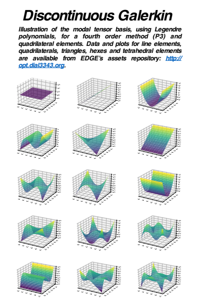

Illustration of EDGE's subcell solver for a fourth order method. Shown are the basis functions and finite volume sub-grid for quadrilateral elements.

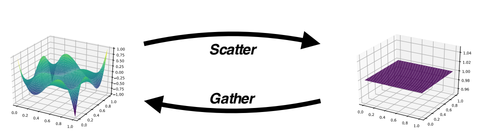

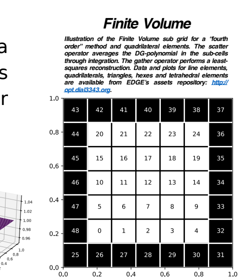

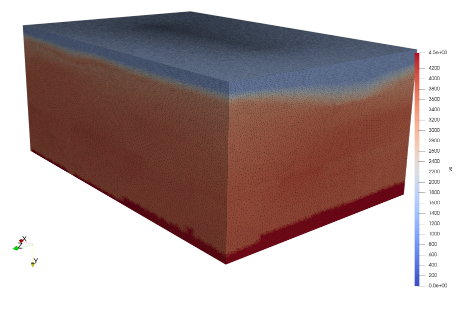

Velocity-aware mesh-refinement using EDGEcut for a modeling domain, which covers the Imperial fault. Contributors: Rajdeep Konwar, Junyi Qiu.

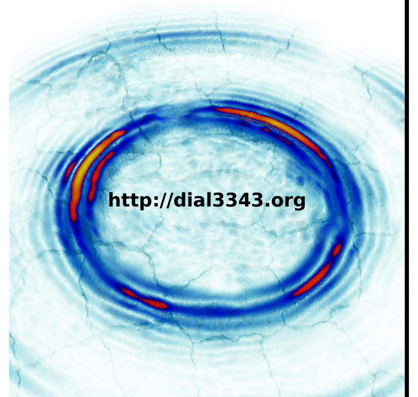

References & Support

[1] EDGE: Extreme Scale Fused Seismic Simulations with the Discontinuous Galerkin Method - A. Breuer, A. Heinecke, Y. Cui. In High Performance Computing: 32nd International Conference, ISC High Performance 2017, Frankfurt, Germany, June 18-22, 2017, Proceedings

[2] 3-D numerical simulations of earthquake ground motion in sedimentary basins: testing accuracy through stringent models - E. Chaljub, E. Maufroy, P. Moczo, J. Kristek, F. Hollender, P.-Y. Bard, E. Priolo, P. Klin, F. de Martin, Z. Zhang, W. Zhang, X. Chen. In Geophys. J. Int. (2015) 201, Issue 1, 90–111

[3] Extreme scale multi-physics simulation of the tsunamigenic 2004 Sumatra megthrust earthquake - C. Uphoff, S. Rettenberger, M. Bader, E. H. Madden, T. Ulrich, S. Wollherr, A.-A. Gabriel. In SC'17 Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis Article No. 21

This work was supported by the Southern California Earthquake Center (**SCEC**) through award #16247. This work is supported by SCEC through award #18211.

This research used resources of the National Energy Research Scientific Computing Center (**NERSC**), 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 work used the Extreme Science and Engineering Discovery Environment (**XSEDE**), which is supported by National Science Foundation grant number ACI-1053575.

EDGE heavily relies on contributions of many authors to open-source software. This software includes, but is not limited to: ASan (debugger), Catch (unit tests), CGAL (surface meshing), Clang (compiler), Cppcheck (static code analysis), Easylogging++ (logging), ExprTk (expression parsing), GCC (compiler), Git (versioning), Git LFS (versioning), gitbook (documentation), Gmsh (volume meshing), GMT (DEM pre-processing), GoCD (continuous delivery), HDF5 (I/O), jekyll (homepage), LIBXSMM (matrix kernels), METIS (partitioning), MOAB (mesh interface), NetCDF (I/O), ParaView (visualization), Proj.4 (map projections), pugixml (XML interface), SAGA-Python (automated remote job submission), SCons (build tool), TF-MISFIT GOF CRITERIA (signal analysis), UCVMC (velocity model), Valgrind (memory debugging), Visit (visualization).