EXA-DUNE: Flexible PDE Solvers, Numerical Methods and Applications

P. Bastian, C. Engwer, D. Göddeke, O. Iliev,
O. Ippisch, M. Ohlberger, S. Turek

Universität Heidelberg Zukunft. Seit 1386.

technische universität dortmund

Westfälische Wilhelms-Universität Münster

Fraunhofer ITWM

Universität Stuttgart

TU Clausthal
Clausthal University of Technology

Munich, March 20, 2017

SPPEXA

EXA-DUNE Project Goals

Develop open-source reusable, efficient, scalable and resilient components for the numerical solution of PDEs

Based on DUNE (Freiburg, Berlin, Heidelberg, Münster,…)

- Flexible software framework, 100+ man-years, GPL-license
- Dimension-independent, different mesh types, hierarchical local refinement, separate mesh/linear algebra, MP parallel
- Efficiency: code generation / static polymorphism in C++
- Applications: Navier-Stokes, Euler Maxwell, elasticity,…

And FEAST (Dortmund)

- Hardware oriented numeric
- Multicore/GPGPU/MPI implementation

Applications: (Multiphase) flow in porous media

SPPEXA

Main Topics covered in 2016

- Enhanced Node-level performance
  - Vectorisation for low order schemes
  - Application of sum-factorization to more complex problems
  - Combination of matrix-free and matrix-based iterative solvers
  - Hardware-aware preconditioning
- Resiliency
  - Fault tolerant multigrid
- Applications:
  - EEG inversion
  - Miscible Displacement
  - Adaptive Multiscale Methods
  - Multilevel Monte-Carlo
  - Land-Surface Model

SIMD for low-order Schemes: Vectorisation over multiple Elements

- Laplace-Equation
- Matrix assembly for conforming FEM, Q_1 Lagrange basis, patch grid
- Intel Haswell E5-2698v3, 2.3 GHz, 16 cores, AVX2, 4 lanes
- Advertised 30 GFLOPs/sec per core
- Advertised 480 GFLOPs/sec total (⟶ % peak)
- Peak without FMA: 15 GFLOPs/sec per core (⟶ % avail)
- Timing for Matrix Assembly:

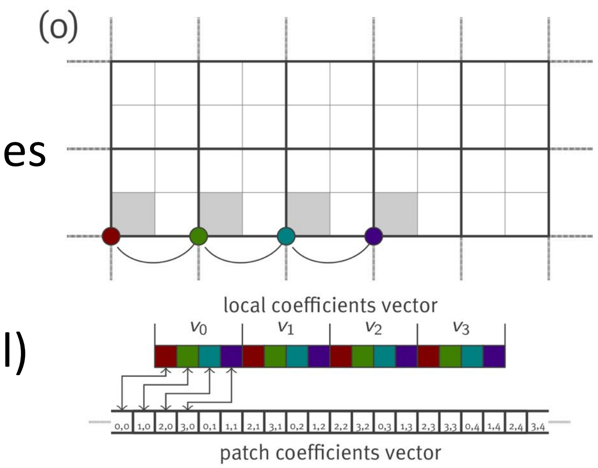

SIMD	lanes	thread	runtime	GFLOP/sec	% avail	% peak
none	1	1	38.6 s	3.0	19.2	0.6
none	1	16	2.5 s	47.3	19.4	9.7
AVX	4	1	16.6 s	5.0	32.0	1.0
AVX	4	16	1.1 s	72.9	30.0	15.0

SIMD for low-order Schemes: Vectorisation over multiple Problems

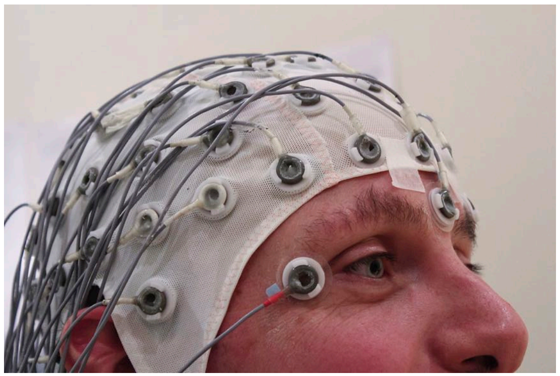

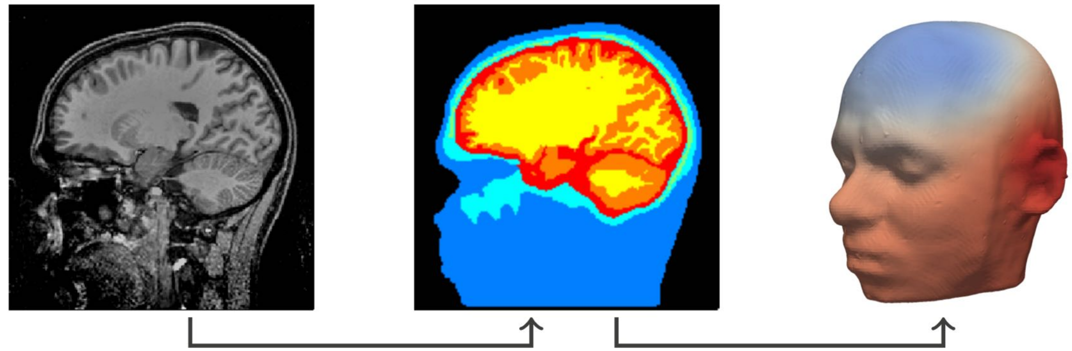

- EEG measurements, 200+ electrodes
- Goal: reconstruct brain activity from surface potential measurements
- For each electrode solve (stationary) adjoint problem to obtain sensitivity, requires solution of linear equation system with identical Matrix **A**
- Vectorise over several adjoint problems, solve multiple linear equation systems simultaneously with CG solver and AMG preconditioner

SIMD for low-order Schemes: Vectorisation over multiple Problems

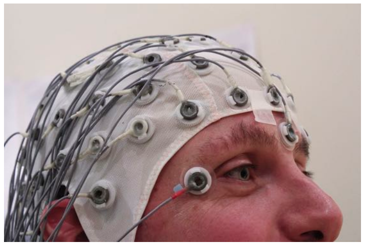

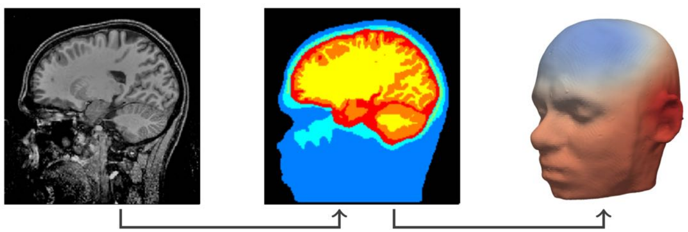

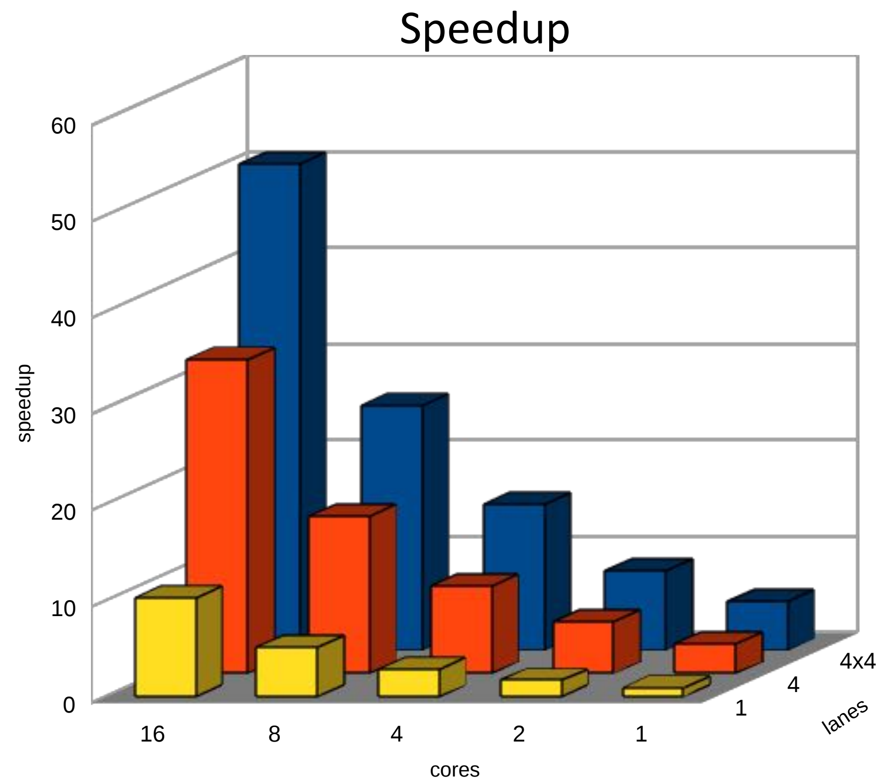

- conforming Q_1 FEM, Lagrange basis
- 256 electrodes, 300’000 grid cells, 60’000 vertices
- Speedup 50 on E5-2698v3, 2.3 GHz, 16 cores (max. theoretical speedup 64)

High-performance (high-order) DG Solver

- Extended assembly to meshes with multi-linear geometry using sum-factorized geometry evaluation
- Fully vectorised parameter evaluations and geometry calculations at quadrature points

Sum-Factorized DG Assembly: Results

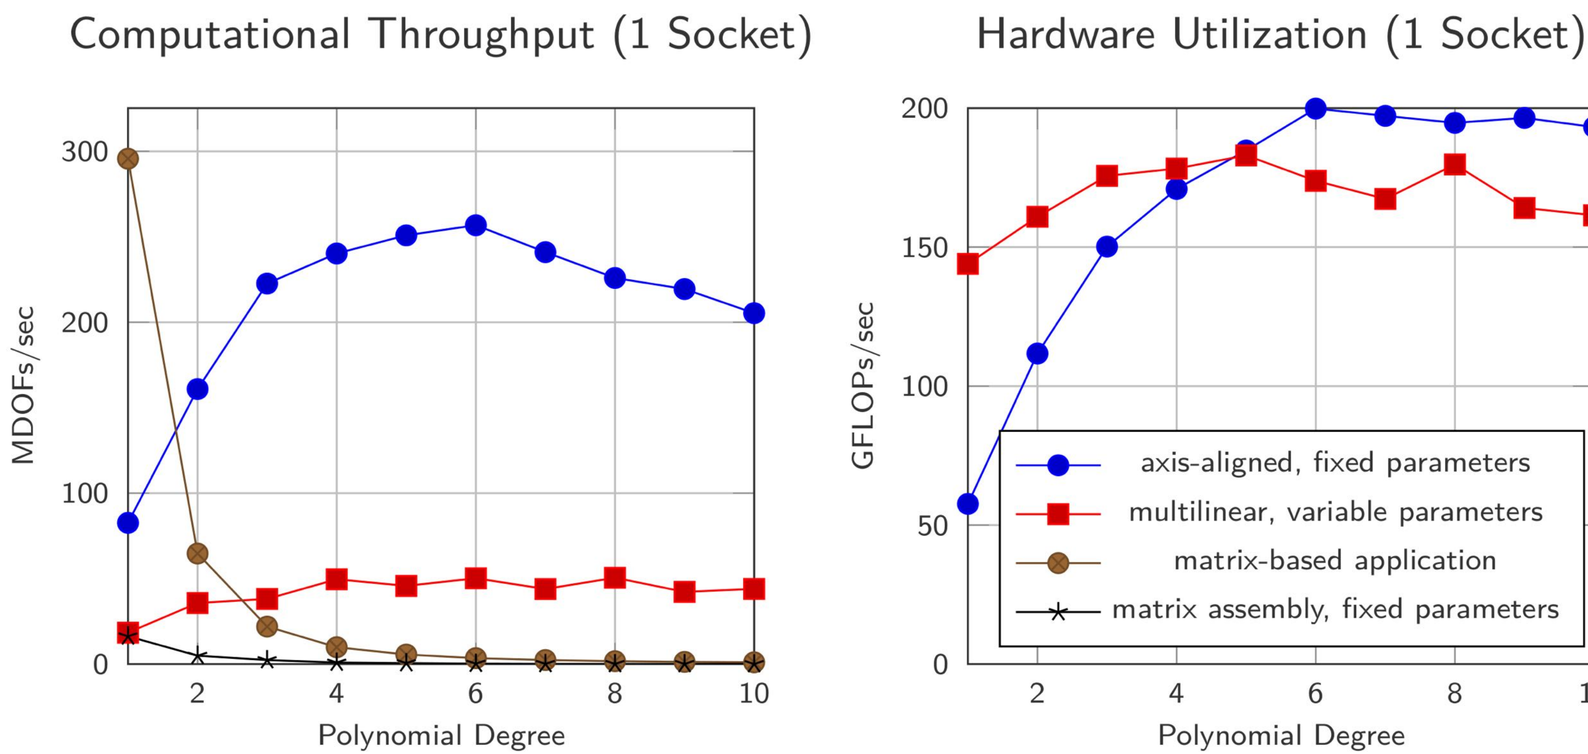

- Full operator application/assembly for two problems of different computational intensity,
  2x Intel Xeon E5-2698v3 2.3 GHz, 16 cores, AVX2, fully loaded
- Smaller problem size for matrix-based computations due to memory constraints

SPPEXA

High-performance (high-order) DG Solver

- Extended assembly to meshes with multi-linear geometry using sum-factorized geometry evaluation
- Fully vectorised parameter evaluations and geometry calculations at quadrature points
- Applying sum-factorized assembly to Navier-Stokes problems (cooperation with Marian Piatkowski, Heidelberg)
- Incorporating knowledge into code generation framework (cooperation with Dominic Kempf, René Heß, Heidelberg)

High-performance (high-order) DG Solver

- Extended assembly to meshes with multi-linear geometry using sum-factorized geometry evaluation
- Fully vectorised parameter evaluations and geometry calculations at quadrature points
- Applying sum-factorized assembly to Navier-Stokes problems (cooperation with Marian Piatkowski, Heidelberg)
- Incorporating knowledge into code generation framework (cooperation with Dominic Kempf, René Heß, Heidelberg)
- Improved inversion of diagonal blocks in partially matrix-free two-level preconditioner (matrix-free for DG + AMG on $P_{0,1}$ subspace, cooperation with Eike Müller, Bath):
  - Different types of Krylov solvers (BiCGStab, GMRES)
  - Block SSOR and tridiagonal preconditioners (for more complicated problems less amenable to block Jacobi)

Sum-Factorized DG Assembly: Results

Total Solution Time

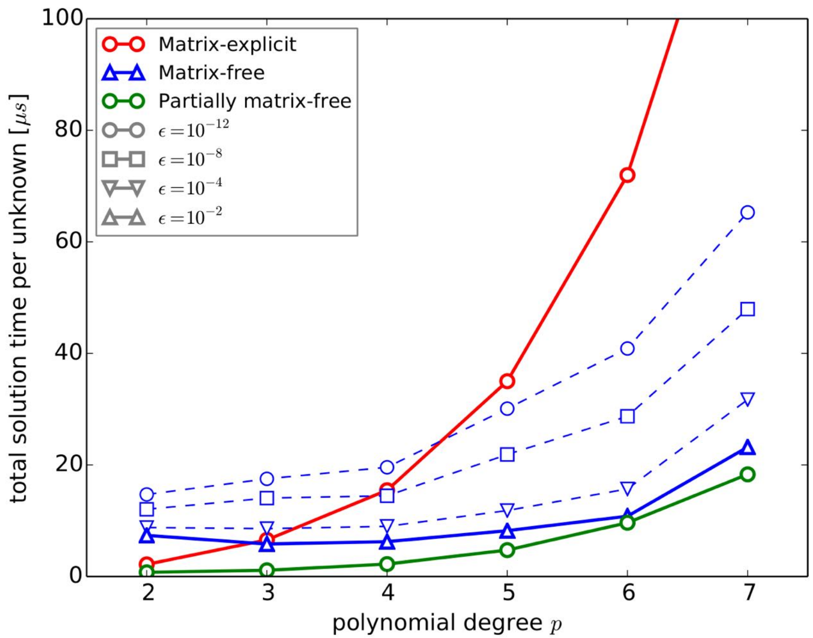

Number of block-iterations ϵ=10^{-2}

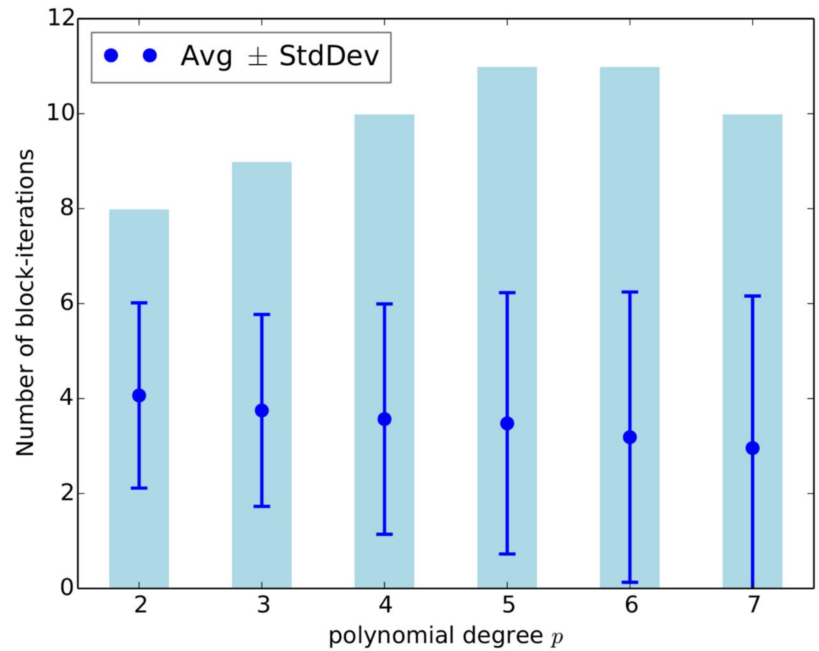

Results by Eike Müller (University of Bath)

- SIPG for 3D stationary convection diffusion equation
- GMRES + tridiagonal preconditioner (large grid Peclet number)
- Slight advantage for storing inverted blocks (partially matrix-free) but larger memory footprint

Miscible Displacement Problem

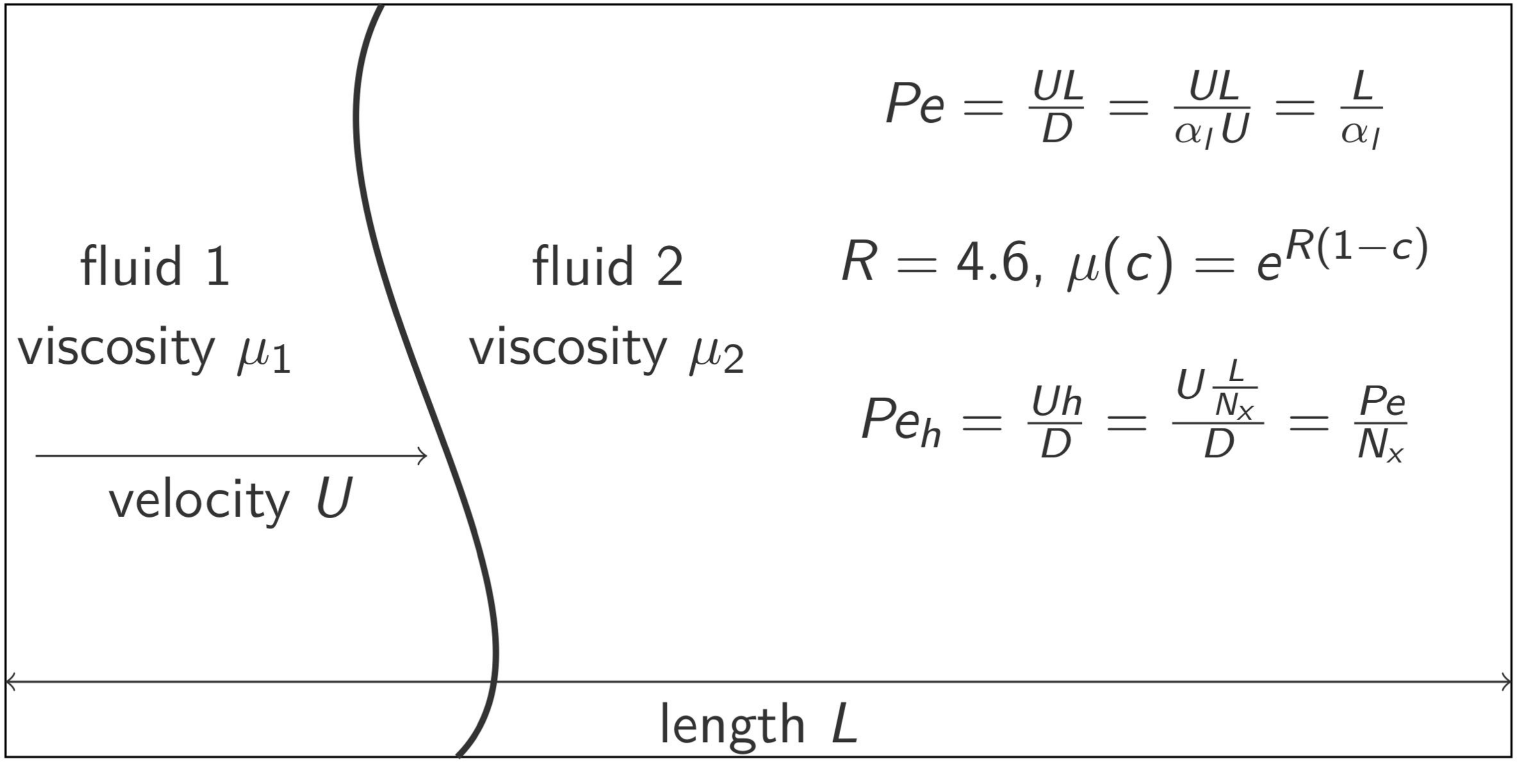

- Application in CO_2 sequestration or polymer flooding
- unstable flow for $R > 0$
- Needs $Pe_h \approx 1$ to avoid unphysical oscillations
- Aspect ratio 3 : 2, periodic boundary conditions in *y*-direction

Miscible Displacement Pe =7200, Pe_h=2

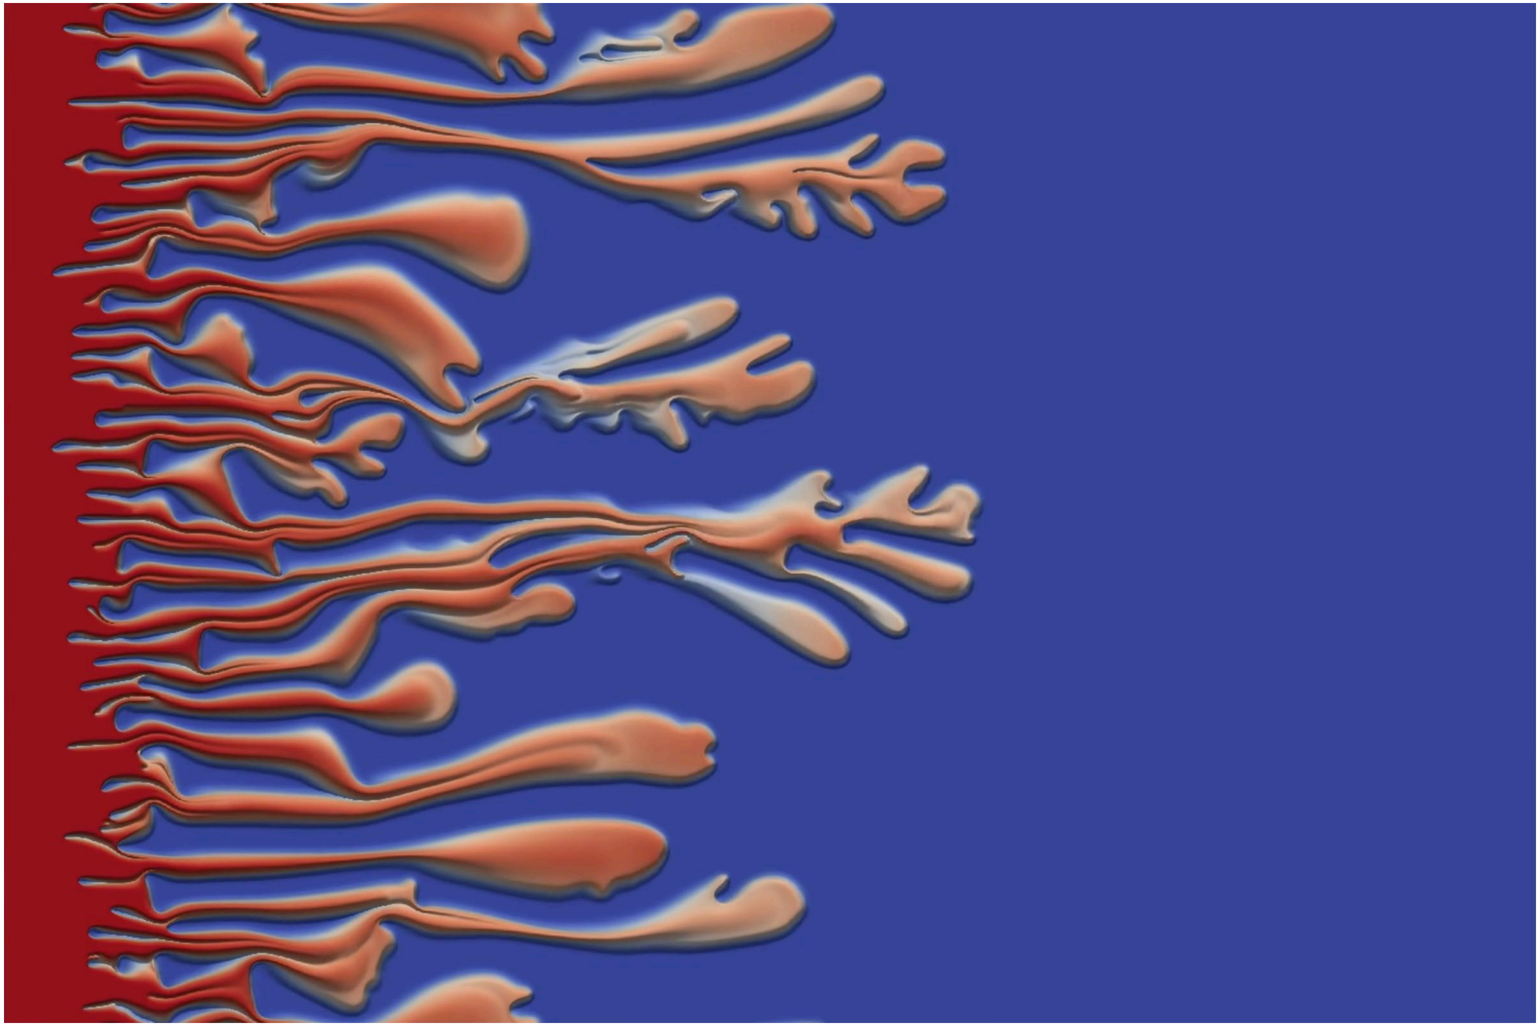

$155 \cdot 10^6$ DOF in 2D, p=1, 2nd order Alexander-scheme, 17’000 time steps
128 Xeon E5-2630 v3, 8 core, 32h

Hardware-aware preconditioning

A versatile, accelerator-exploiting preconditioner for DUNE based on approximate inverses

- First phase: $\mathbf{u} \leftarrow \mathbf{M}\mathbf{v},\ \mathbf{M} = \mathrm{SPAI}(\mathbf{A}) \approx \mathbf{A}^{-1}$
  Proved to be numerically effective
- 2016: assemble $\mathbf{M}$ efficiently on CPU and GPU:
  Different approaches featured:

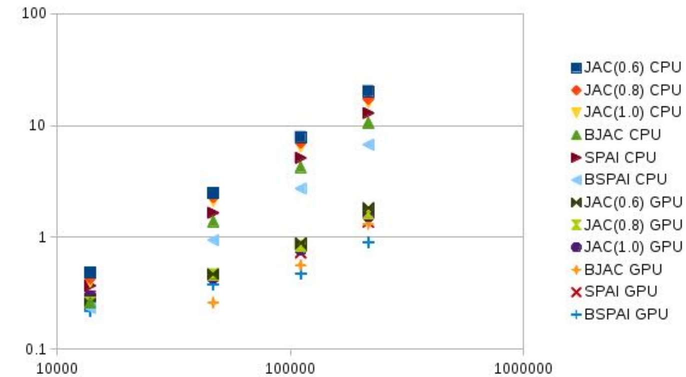

Standard

Many Householder transform-based QR decompositions, batched on GPU

cuBLAS, cuSPARSE and MKL
~4x GPU v. CPU

Monte Carlo

Partial randomization via Markov Chains

$$m_{ij} = \lim_{N\to\infty} \frac{1}{N} \sum_{s=1}^{N} \left(\sum_{q=0}^{\infty} \frac{W_q}{a_{ii}} \delta_{s_q j} \right)$$

CUDA and MKL

~5x GPU v. CPU

Machine Learning

Treat matrices as discrete functions in a function regression train a Deep Learning Net with pairs $(\mathbf{A}, \mathbf{A}^{-1})$
Make inference for $\mathbf{M}$

First results with fixed sparsity, Implementation with TensorFlow CPU and GPU

Fault Tolerance

- New fault-tolerant multigrid (geometric and algebraic)
- Basic idea: use properties from Full Approximation Scheme-Multigrid to protect the linear case
- At the moment targets bit flips, node loss still work in progress
- Local failure local recovery, no communication for faulty inner DOF
- Fault detection and correction within one Multigrid cycle
- Enables using Multigrid (V, F, W) as a preconditioner
- Very few false positives, < 10% overhead
- Parallel implementation in FEAT and DUNE::ISTL almost completed

V-cycle	**poisson**	**dico**	**andi**	**andicore**
fault-free	4	6	14	7
classic	4.225(272)	6.268(335)	15.111(850)	7.466(439)
ftmg	4.038	6.007	14.007	7.017
false-positive	13	21	27	25
worse	15	1	0	1

Averaged iterations from 4000 test runs, error probability 10^{-7} per DOF per smoothing step

Q_1-Finite Elements, 10^6 DOF

Application: Adaptive Multiscale Methods

- Strong scaling of Multiscale-Finite Element Code: $96^3 = 884'736$ cubes on coarse scale with 12^3 fine cells per coarse cell, $1.5 \cdot 10^9$ cells in total
- Phase 2 target: Adaptive multilevel hybrid implementation of the localised reduced basis method
  - Prototype for parallel SWIPDG discretization on arbitrarily partitioned grids
  - Realized Python bindings as a precursor to implementing model reduction in pyMOR

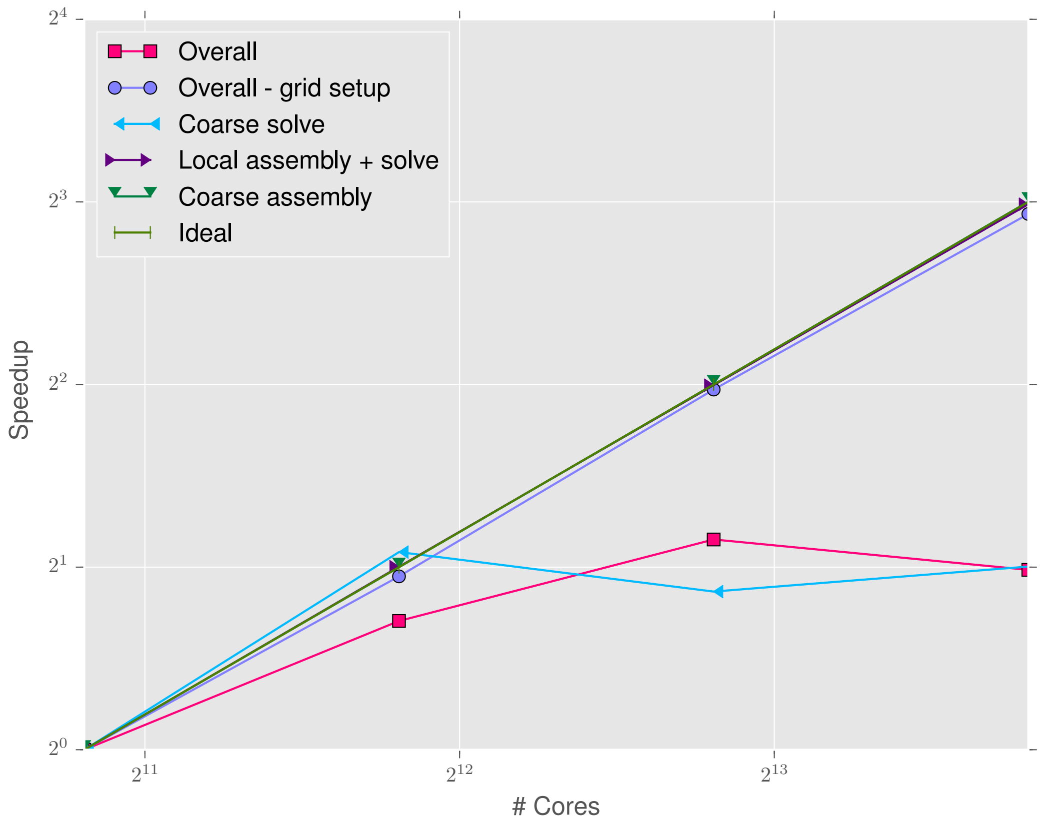

Application: Multilevel Monte-Carlo

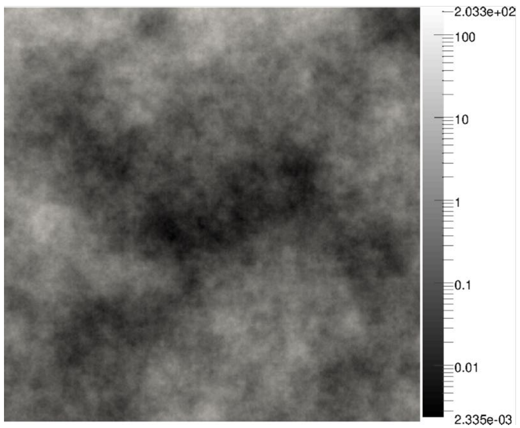

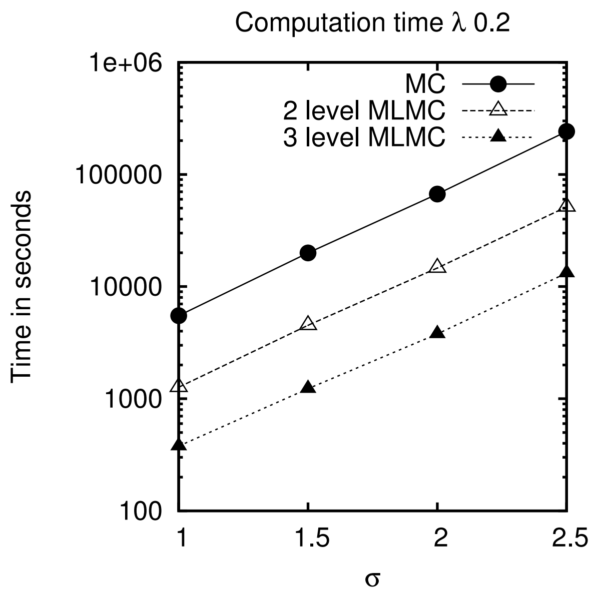

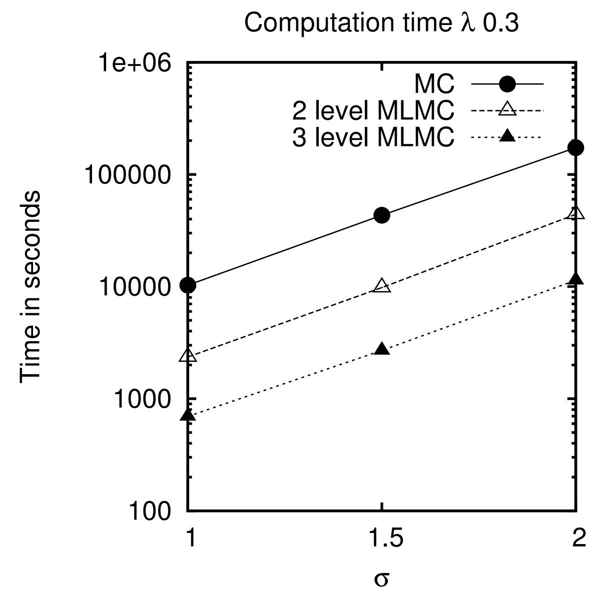

$$-\nabla \cdot [k(x,\omega)\nabla p(x,\omega)] = 0$$

- Good acceleration with Multilevel Monte-Carlo compared to standard Monte-Carlos scheme
- Ongoing work on improvement of parallelization and permeability generator

Application: Land-Surface-Model

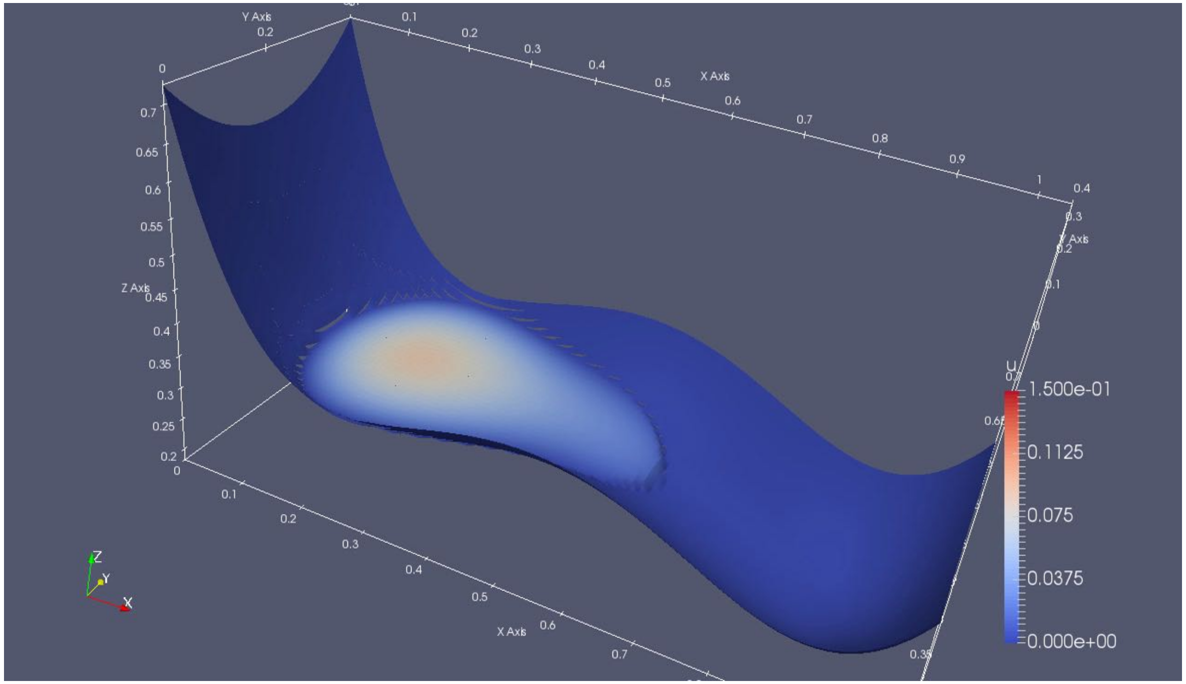

- Diffusive Wave Approximation for surface runoff
  - Parallel arbitrary order DG Implementation
  - Factor 10'000 speedup with semi-implicit Time-Discretization compared to explicit scheme
- Richards-Equation for subsurface flow
  - Parallel arbitrary order DG Implementation
  - Diagonally-implicit Runge-Kutta in time, Newton for linearization
- Operator-Splitting, surface/subsurface coupling with special form of Dirichlet-Neumann Coupling (ongoing work)

Activities and Outreach

Conferences and Workshops

- HPC-FEA: High-Performance Computing in Finite Element Applications, Minisymposium, Pilsen (Czech Republic), together with **SPPEXA** and **TERRA-NEO**
- Exa-scale ready PDE solvers, Minisymposium, Algorithmy 2016, Podbanske (Slovakia)

Summerschools and Tutorials

- Numerical GPGPU (2 times), Dortmund
- Accelerator Computing (with **SPPEXA** and **DASH)**, Obernai (France)
- DUNE/PDELab Course, Heidelberg
- WUCSSS 2016, Bangkok (Thailand)
- ESSAM School on Mathematical Modelling, Numerical Analysis and Scientific Computing, Kacov (Czech Republic)
- Summer School on Applied Mathematics and Scientific Computing, Manila (Philipines)

Invited Guests

Michael Ortiz (Caltech) , Svetozar Margenov (Bulgarian Academy of Science)

Cooperation

with **EXASTENCILS, TERRA-NEO, AIMES, ADA-FS, EXA-DG** and **EXAMAG**

10 Papers, 25 Keynotes, Talks and Posters related to SPPEXA in 2016

Activities and Outreach

Conferences and Workshops

- HPC-FEA: High-Performance Computing in Finite Element Applications, Minisymposium, Pilsen (Czech Republic), together with **SPPEXA** and **TERRA-NEO**
- Exa-scale ready PDE solvers, Minisymposium, Algorithmy 2016, Podbanske (Slovakia)

Summerschools and Tutorials

- Numerical GPGPU (2 times), Dortmund
- Accelerator Computing (with **SPPEXA** and **DASH)**, Obernai (France)
- DUNE/PDELab Course, Heidelberg
- WUCSSS 2016, Bangkok (Thailand)
- ESSAM School on Mathematical Modelling, Numerical Analysis and Scientific Computing, Kacov (Czech Republic)
- Summer School on Applied Mathematics and Scientific Computing, Manila (Philipines)

Invited Guests

Michael Ortiz (Caltech) , Svetozar Margenov (Bulgarian Academy of Science)

Cooperation

with **EXASTENCILS, TERRA-NEO, AIMES, ADA-FS, EXA-DG** and **EXAMAG**

10 Papers, 25 Keynotes, Talks and Posters related to SPPEXA in 2016

Christian Engwer successfully completed Tenure-track at University of Münster