GROMEX

FMM for Realistic Biomolecular Simulations on the ~~Exascale~~ Multi-Petascale

May 8, 2013 | I. Kabadshow, A. Beckmann, H. Dachsel | 1st Daresbury–Jülich Workshop

Member of the Helmholtz Association

Long Range Interactions

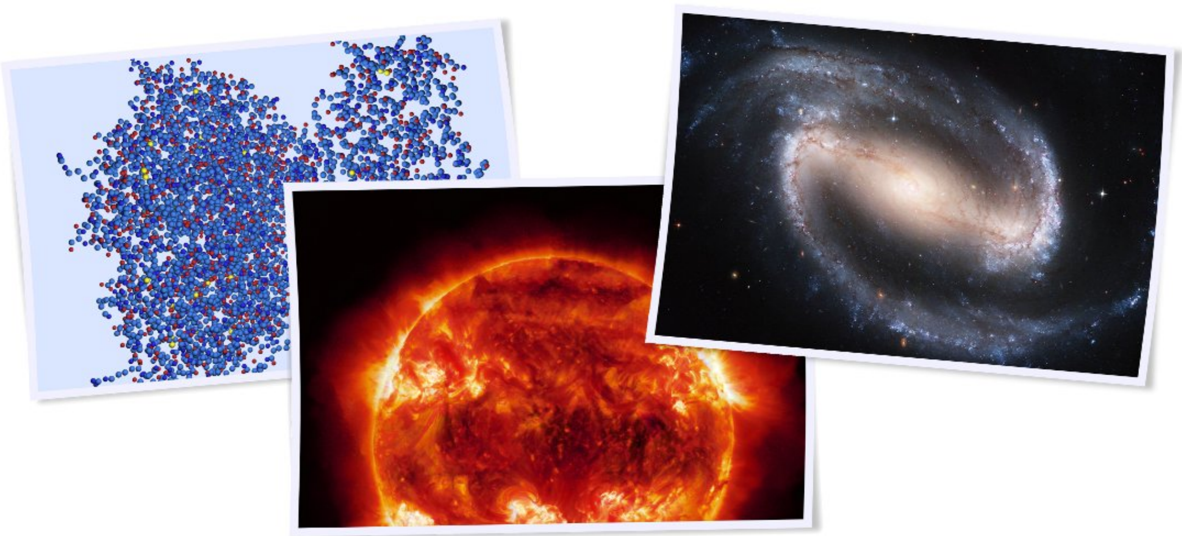

Molecular Dynamics Plasma Physics Astrophysics

Member of the Helmholtz Association

Alternatives To Direct Summation

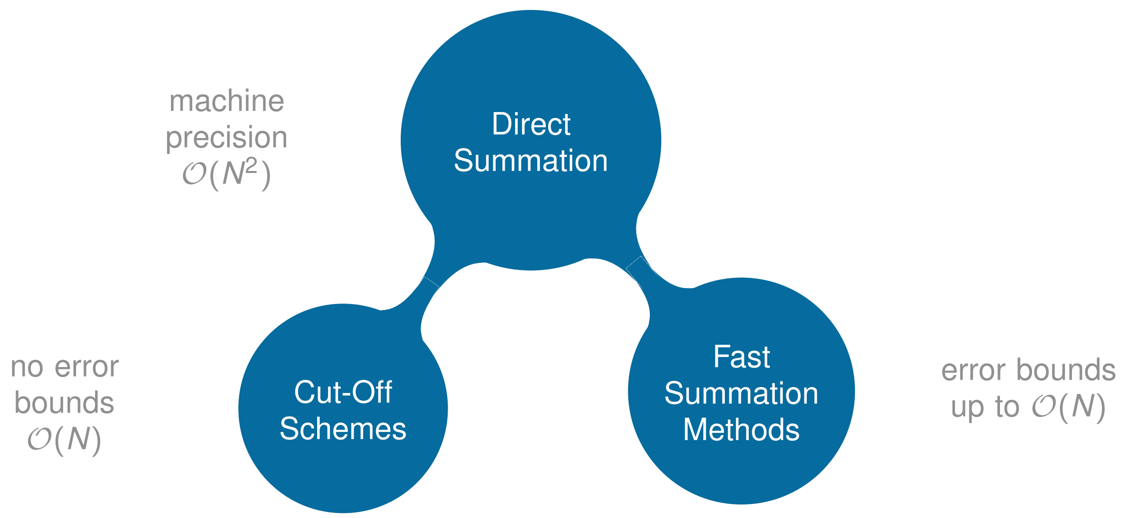

Member of the Helmholtz Association

History of Fast Summation Methods

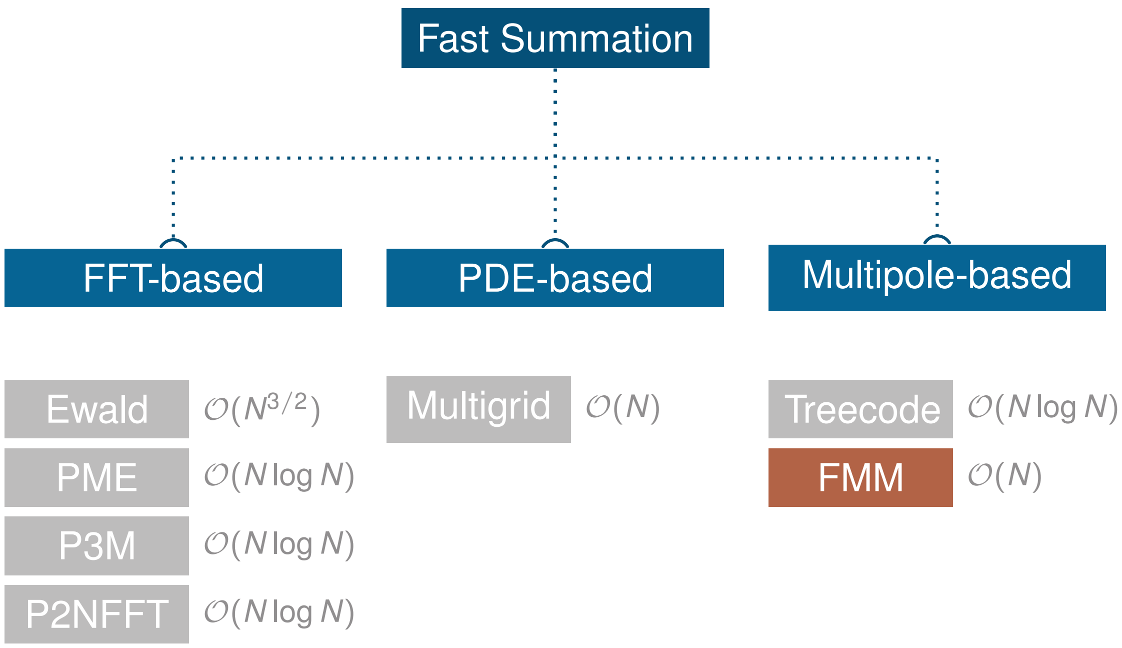

Member of the Helmholtz Association

The Fast Multipole Method

There's no such thing as a free lunch

Pro

- Complexity reduces to $\mathcal{O}(N)$

Con

- Three new parameters are introduced
  - Tree depth (d)
  - Separation Criterion (ws)
  - Number of poles (p)

Member of the Helmholtz Association

The Fast Multipole Method

Bane and Boon

Bane

- Newly introduced parameters d, ws, p need to be tuned.
- Reliable optimization scheme is essential for speedup.

Boon

- Computation time $t(\epsilon)$ can be a function of requested precision.
- Tremendous speedup possible for any requested precision.

Member of the Helmholtz Association

FMM: Current Fortran Version

Generic Features

- Automatic tuning of FMM parameters to ensure user-requested energy accuracy ΔE_r
- Automatic runtime optimization
- Works with clustered and homogen particle distributions
- Works with open, 1D, 2D and 3D periodic boundary conditions
- Low cross-over point with direct summation (4000 particles)
- Allows different precisions $\Delta E_r = 10^{-1} \dots 10^{-30}$

Member of the Helmholtz Association

FMM Workflow

Original FMM Workflow

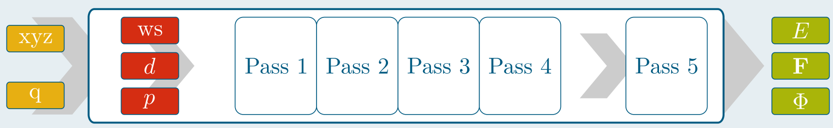

Enhanced FMM Workflow

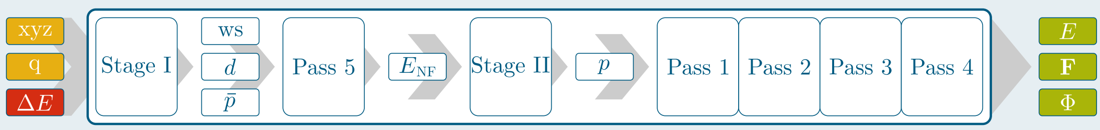

Member of the Helmholtz Association

Precision Scaling

2.097.152 Particles - Total Computation Time

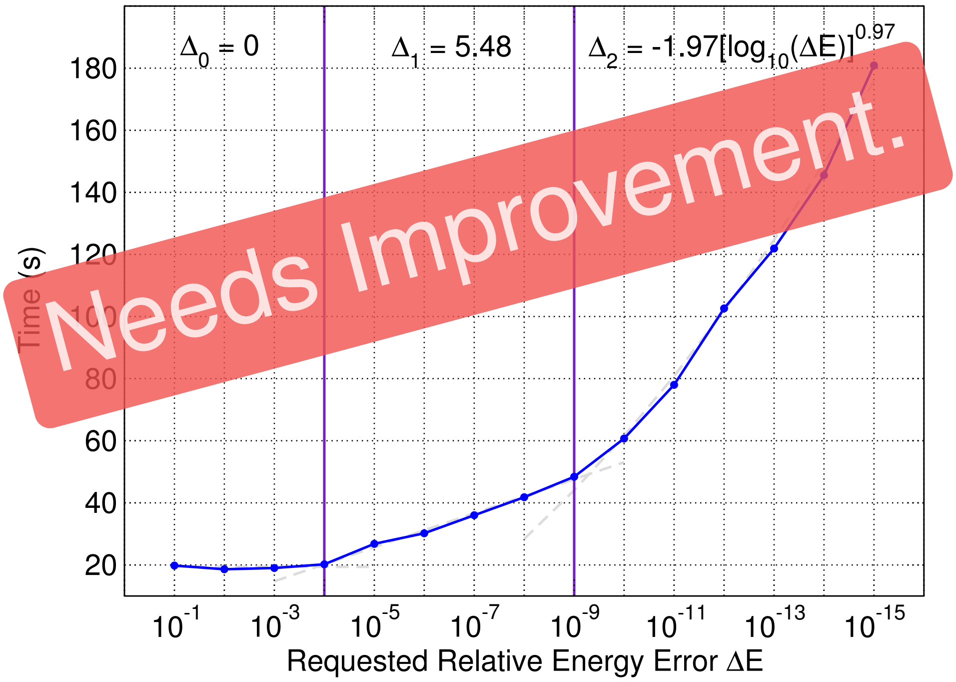

FMM Runtime Optimization

Trade `DIV` and `SQRT` for `ADD` and `MULT`

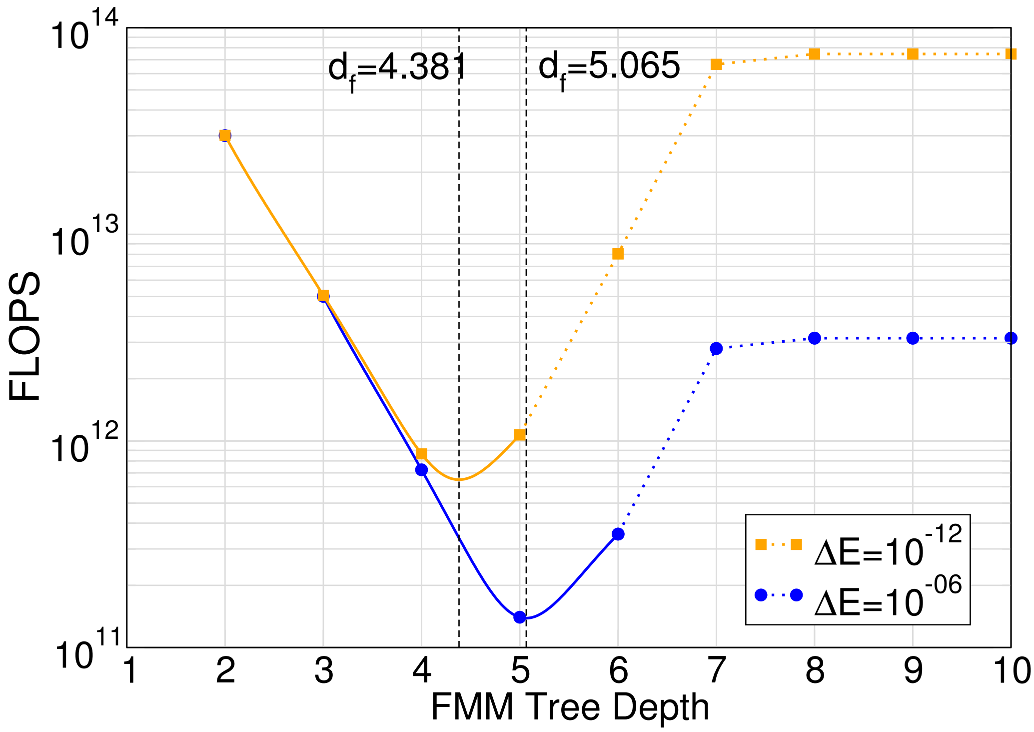

Quadruple Precision Check

1D, 2D and 3D Periodicity

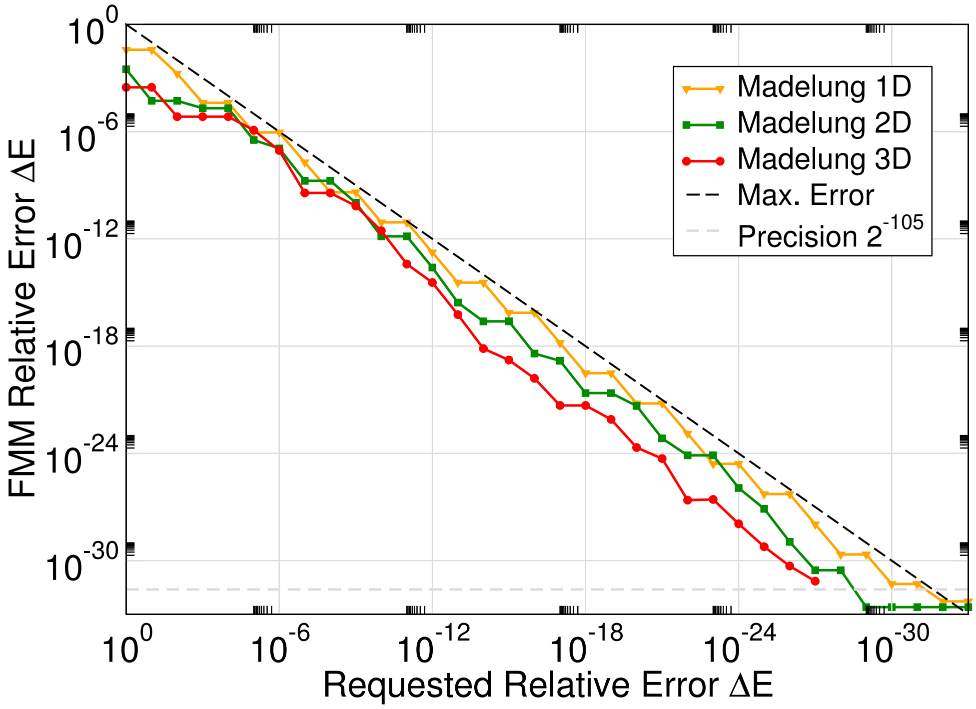

Strong Scaling On Jugene

FMM with Open Boundaries, 2.097.152 Particles, $\Delta E = 10^{-2}$

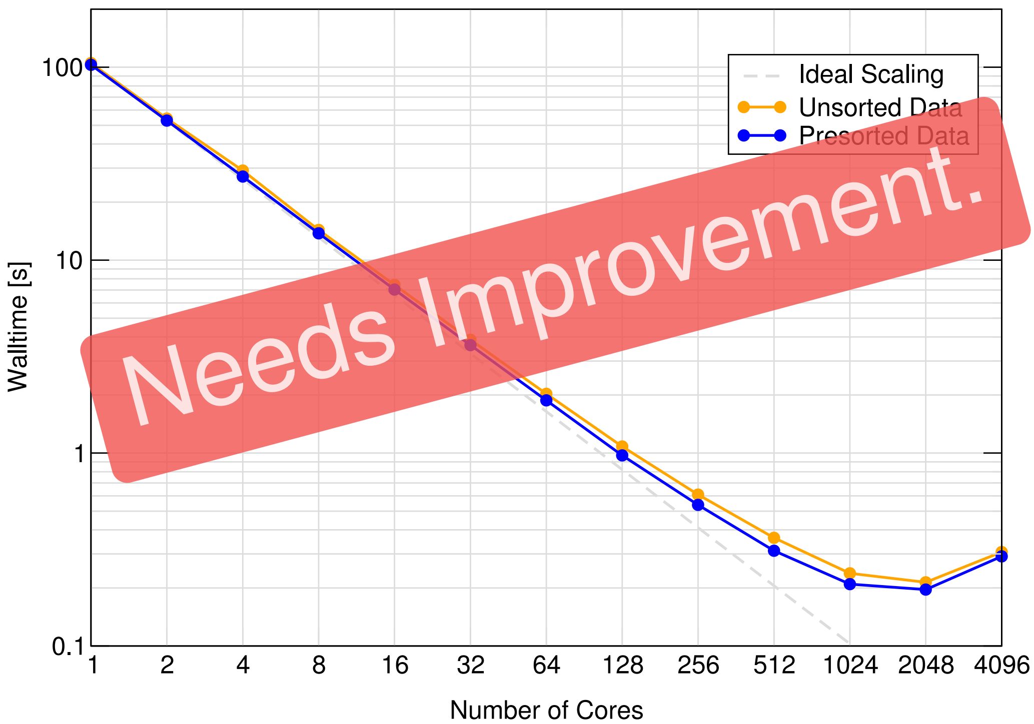

Strong Scaling

1073741824 Particles on Juropa

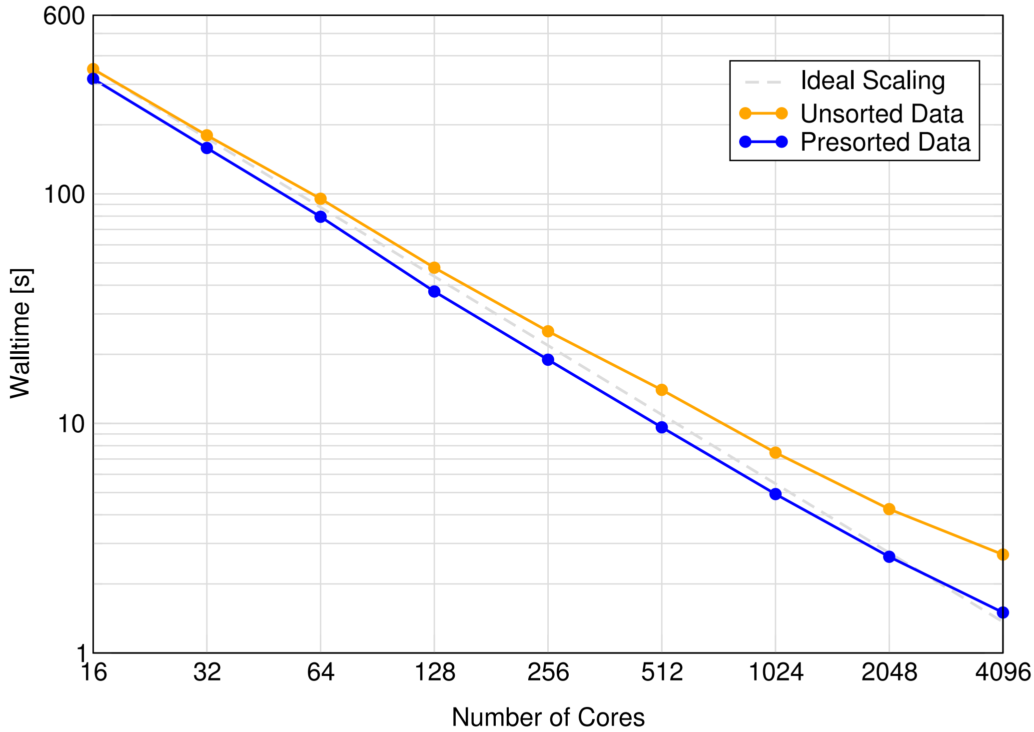

Strong Scaling

1073741824 Particles on Jugene

Three-Trillion Particles @ BGP

System Characteristics

- 3.011.561.968.121 particles
- 73728 nodes (VN mode, 294912 cores)

Results

- 3812s runtime, unsorted data (2755 particles/second/core)
- 715s runtime, presorted data (14687 particles/second/core)

Three-Trillion Particles @ BGP

Open System, 3.011.561.968.121 Particles, $\Delta E = 10^{-2}$, 44.16 Bytes/particle

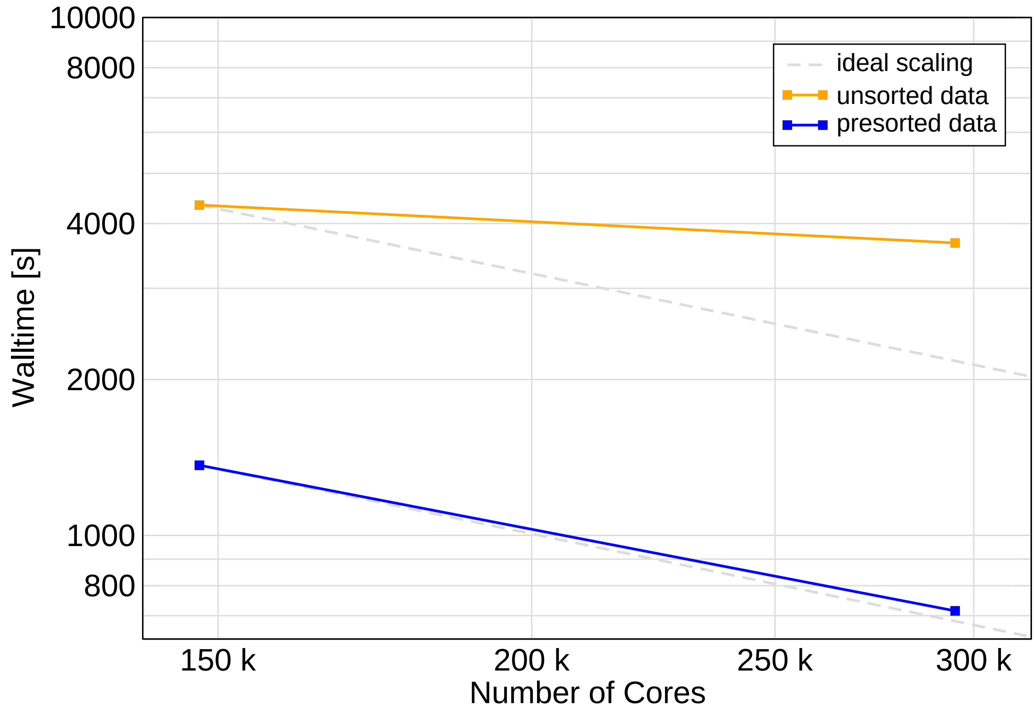

ScaFaCoS Benchmarks

FFT-based Algorithms

- P3M
- P2NFFT

PDE-based Algorithms

- PP3MG
- VMG

Multipole-based Algorithm

- FMM

Local Algorithm

- MEMD

Relative RMS Potential Error: $\epsilon_{\text{pot}} = 10^{-3}$, $\epsilon_{\text{pot}} = 10^{-5}$

$$\epsilon_{\text{pot}} := \left(\frac{\sum_{j=1}^{N} |\Phi_{\text{REF}}(\mathbf{x}_j) - \Phi_{\text{FCS}}(\mathbf{x}_j)|^2}{\sum_{j=1}^{N} |\Phi_{\text{REF}}(\mathbf{x}_j)|^2} \right)^{1/2}$$

Member of the Helmholtz Association

Accuracy Range

Juropa, single core

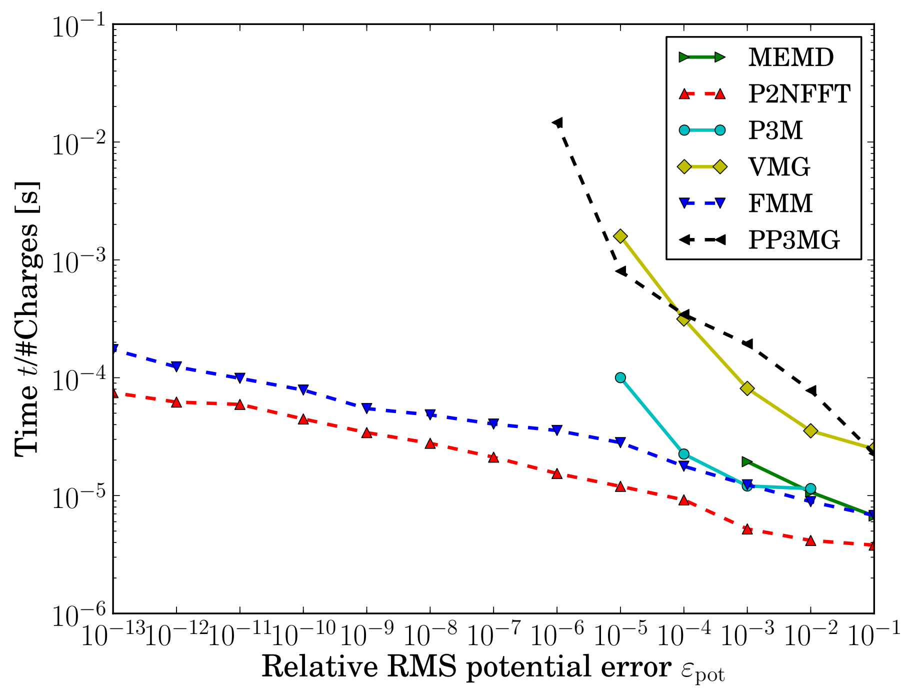

Member of the Helmholtz Association

Complexity Scaling

Juropa, single core

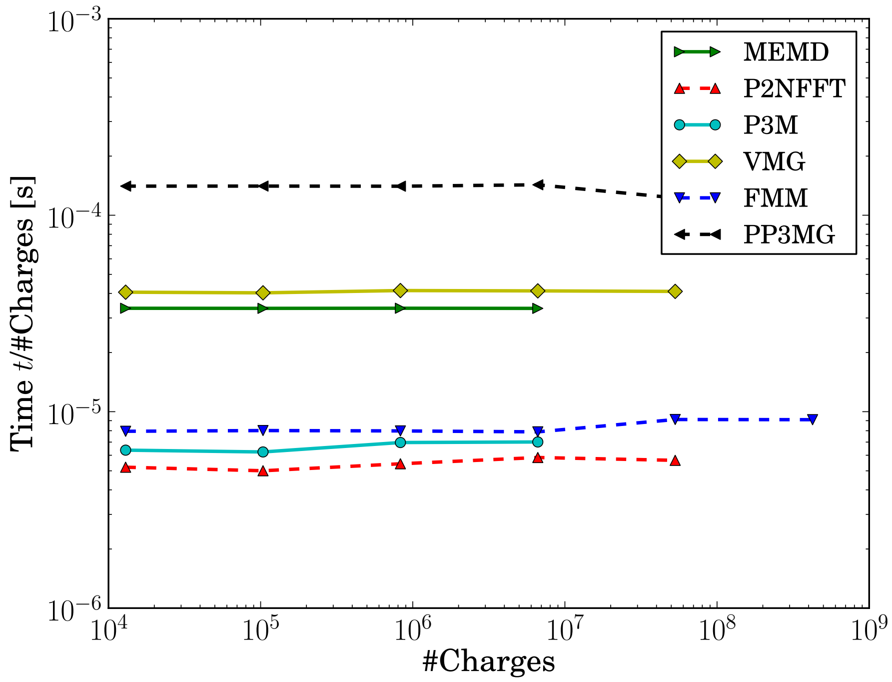

Conservation of Physical Properties

Kinetic Energy: Center of Mass Motion, $\epsilon_{\mathrm{pot}} = 10^{-3}$

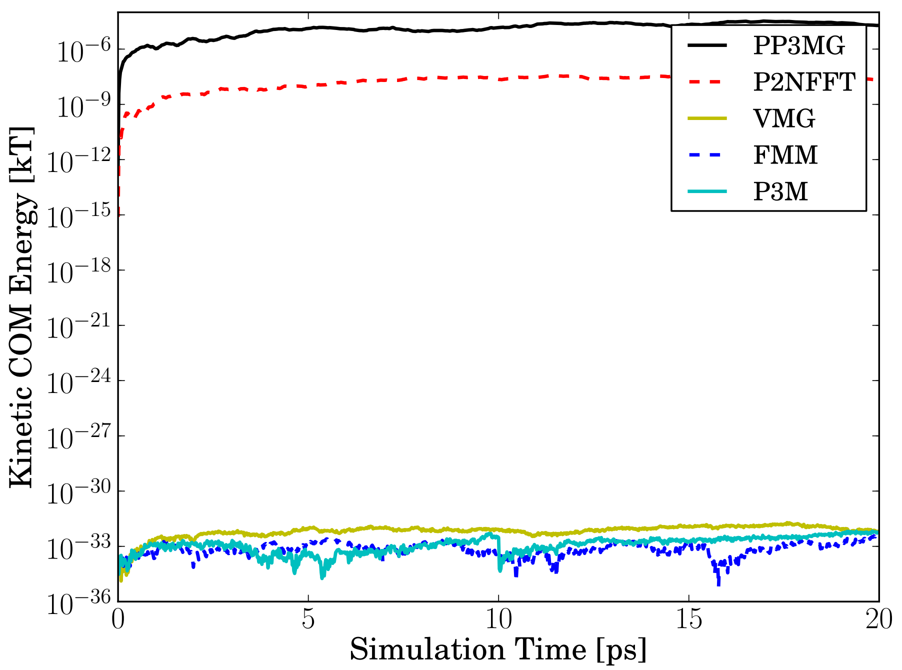

JÜLICH
FORSCHUNGSZENTRUM

Conservation of Physical Properties

Kinetic Energy: Center of Mass Motion, $\epsilon_{\text{pot}} = 10^{-5}$

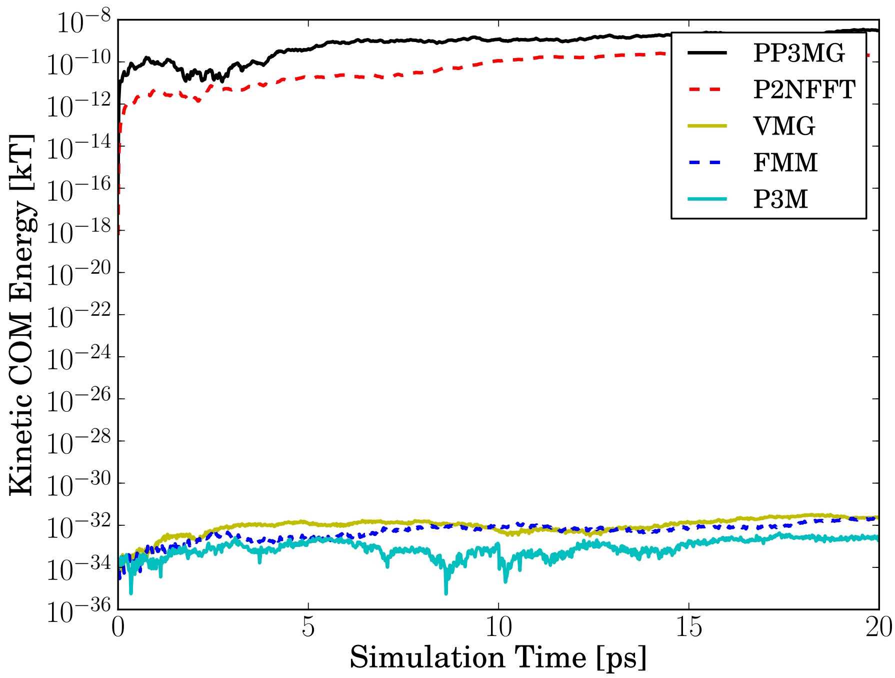

JÜLICH
FORSCHUNGSZENTRUM

Conservation of Physical Properties

Conservation of Total Energy, $\epsilon_{\text{pot}} = 10^{-3}$

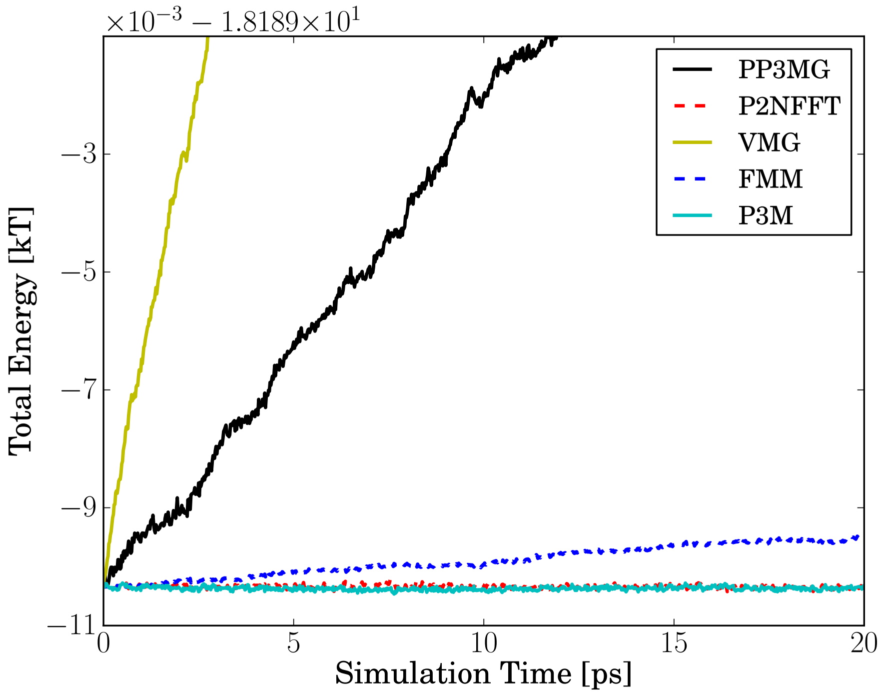

Conservation of Physical Properties

Conservation of Total Energy, $\epsilon_{\text{pot}} = 10^{-5}$

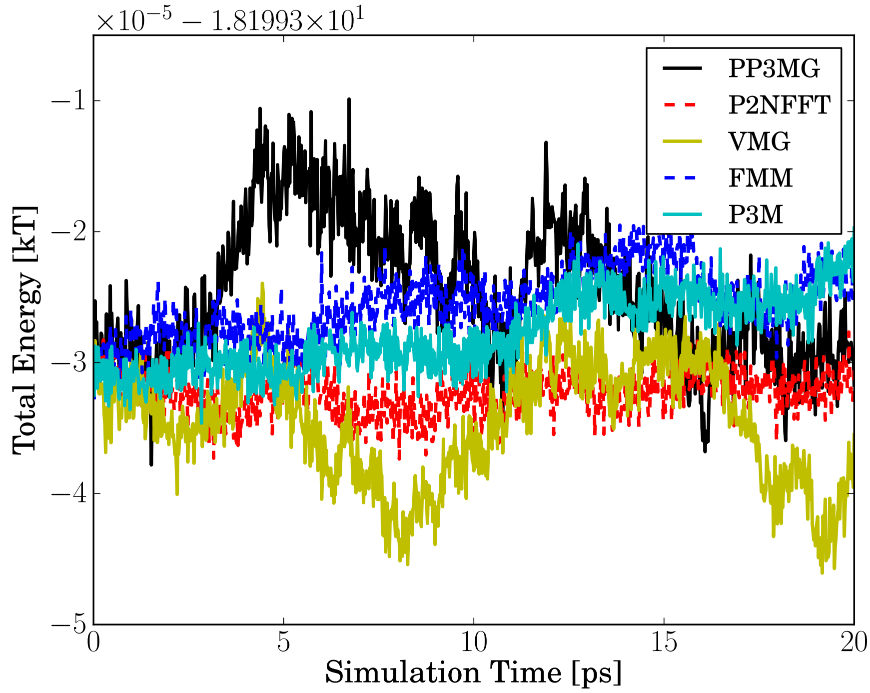

Parallel Efficiency Comparison

8100 Particles, Juropa, $\epsilon_{\text{pot}} = 10^{-3}$

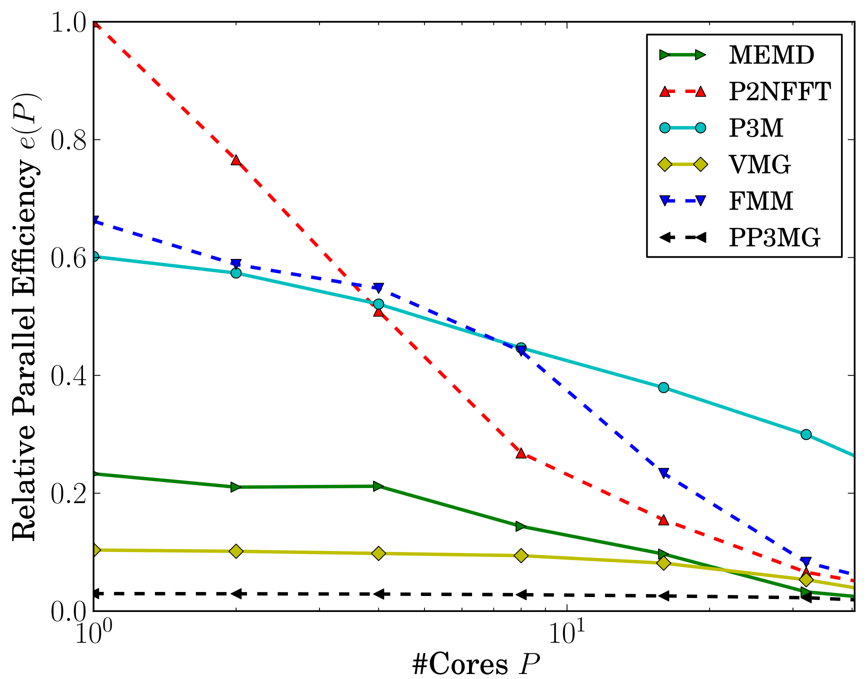

Parallel Efficiency Comparison

1.012.500 Particles, Jugene, $\epsilon_{\text{pot}} = 10^{-3}$

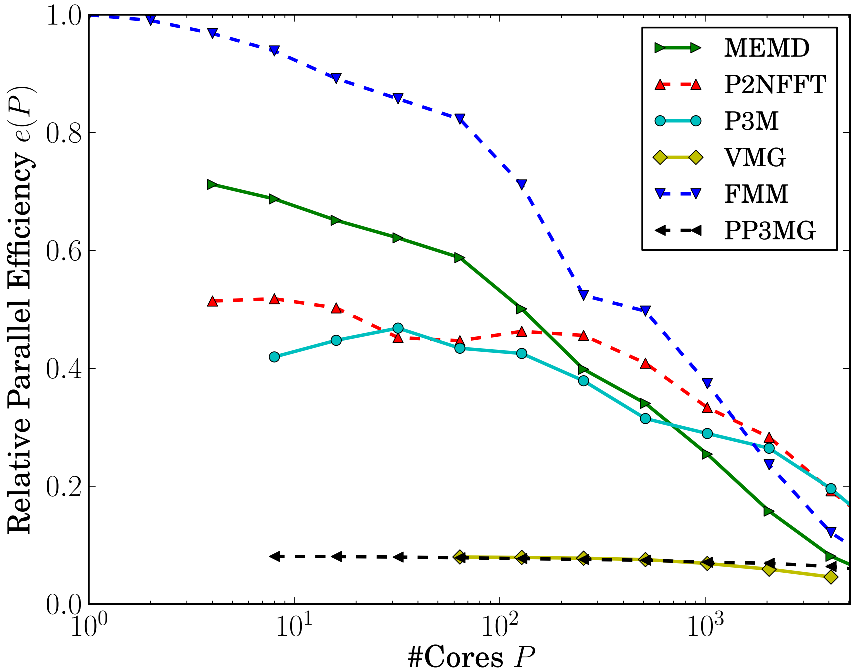

Member of the Helmholtz Association

Parallel Efficiency Comparison

9.830.400 Particles, Jugene, $\epsilon_{\text{pot}} = 10^{-3}$

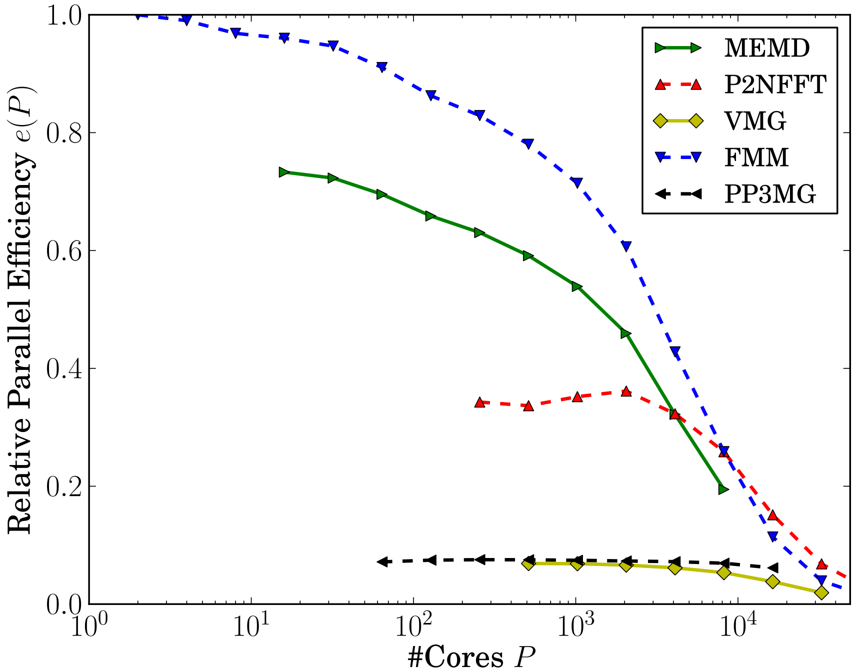

Member of the Helmholtz Association

Parallel Efficiency Comparison

1.012.500.000 Particles, Jugene, $\epsilon_{\text{pot}} = 10^{-3}$

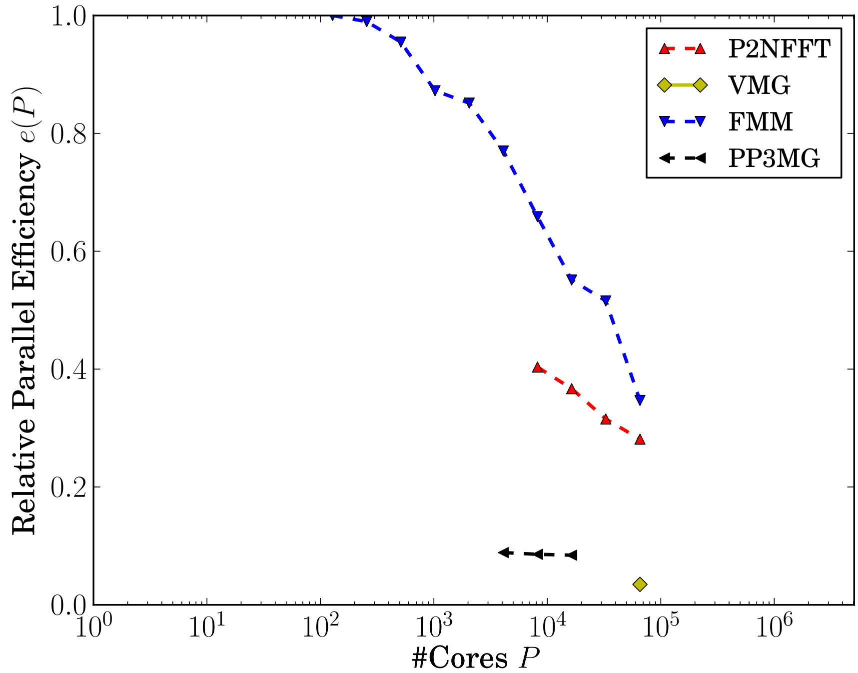

DFG Priority Programme (SPP 1648)

Software for Exascale Computing (SPPEXA)

- entering era of massive parallelism (> 10^7 processing units)
- urgent need for fundamental reseach on HPC software
- reconnect fields of computer science with CSE and HPC
- demands close cooperation and co-design
- service-driven collaborations not permitted
- (3+3) year funding period

Prediction

- exascale system expected 2018 ($\approx 10^{18}$ FLOPS)
- racks without brains strategy will not suffice

SPPEXA Research Directions

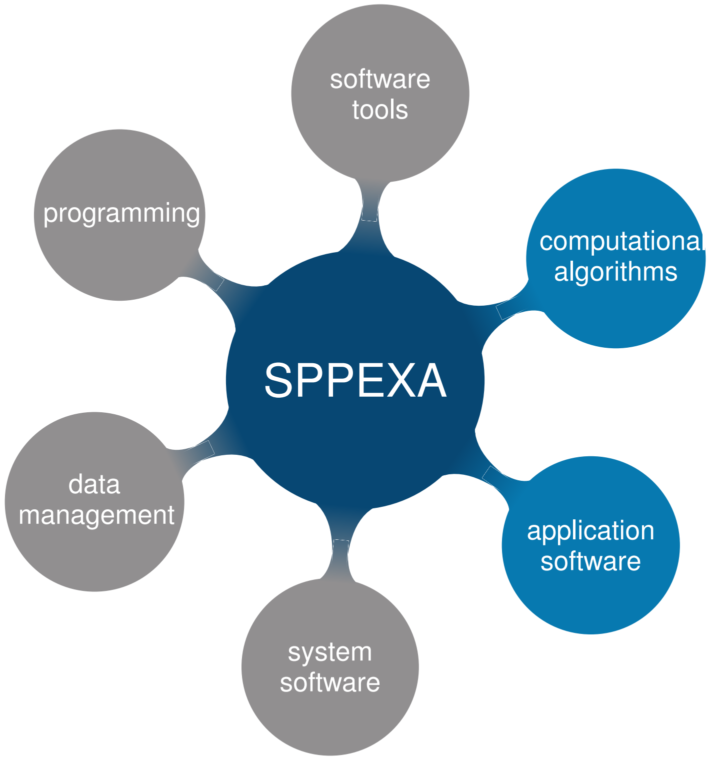

DFG Priority Programme

Software for Exascale Computing

- EXA-DUNE
  Flexible PDE Solvers, Numerical Methods, and Applications
- DASH
  Hierarchical Arrays for Efficient and Productive Data-Intensive Exascale Computing
- TERRA-NEO
  Integrated Co-Design of an Exa-Scale Earth Mantle Modeling Framework
- EXASTEEL
  Bridging Scales for Multiphase Steels
- ExaStencils
  Advanced Stencil-Code Engineering
- EXAHD
  An Exa-Scalable Two-Level Sparse Grid Approach for Higher-Dimensional Problems in Plasma Physics and Beyond
- ExaFSA
  Exascale Simulation of Fluid-Structure-Acoustics Interactions

DFG Priority Programme

Software for Exascale Computing

- EXA??*
  A fast and fault tolerant microkernel-based system for exascale computing
- ESSEX
  Equipping Sparse Solvers for Exascale
- EXASOLVERS
  Extreme scale solvers for coupled problems
- EXAMAG
  Exascale simulations of the evolution of the universe including magnetic fields
- GROMEX
  Unified Long-range Electrostatics and Dynamic Protonation for Realistic Biomolecular Simulations on the Exascale
- CATWALK
  A Quick Development Path for Performance Models

Member of the Helmholtz Association

GROMEX within SPPEXA

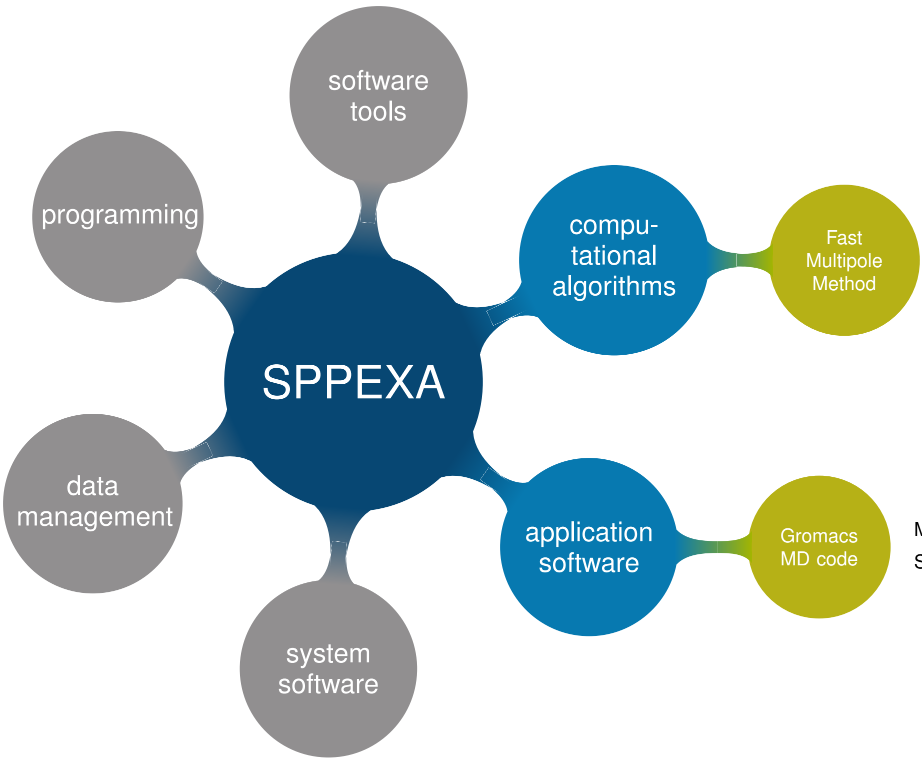

Jülich Supercomputing Centre

MPI Göttingen
Stockholm University

Member of the Helmholtz Association

GROMEX Toolbox

Flexible exascale solver for long-range interactions

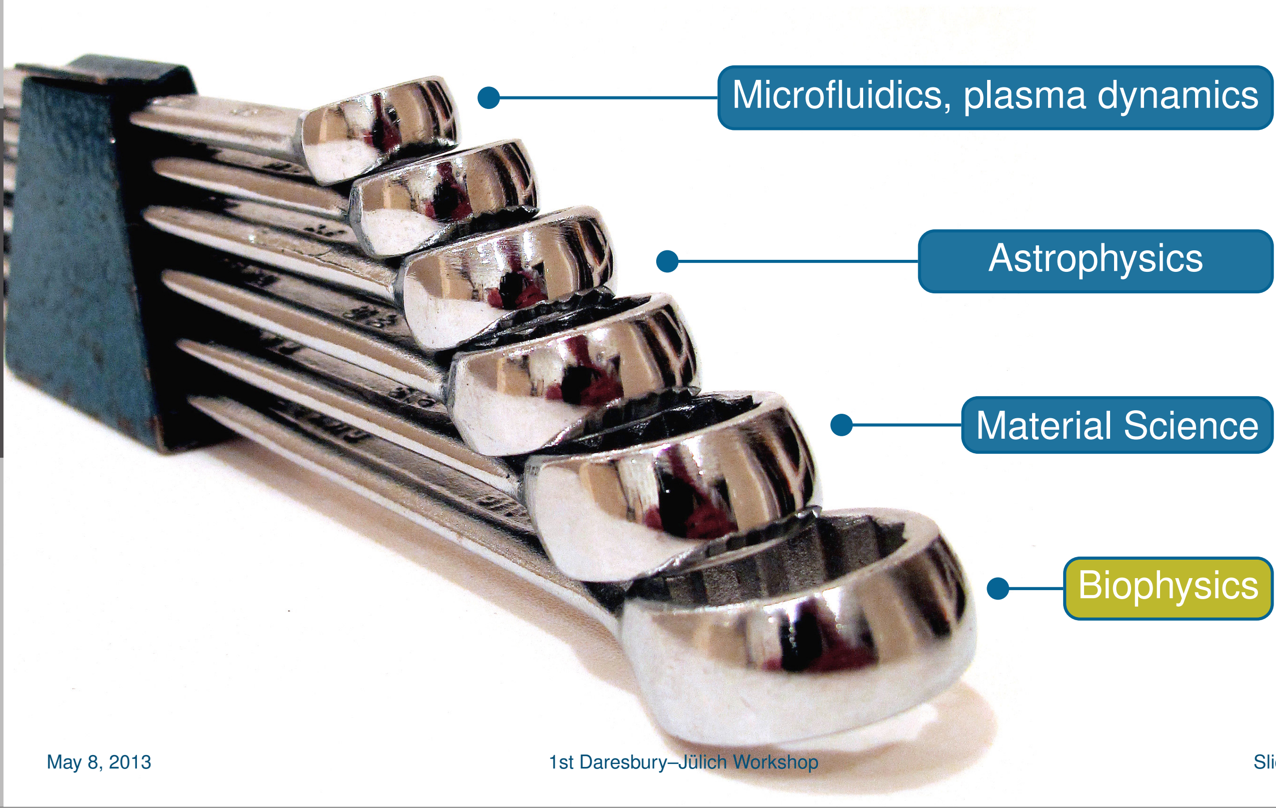

FMM Algorithmic Challenges

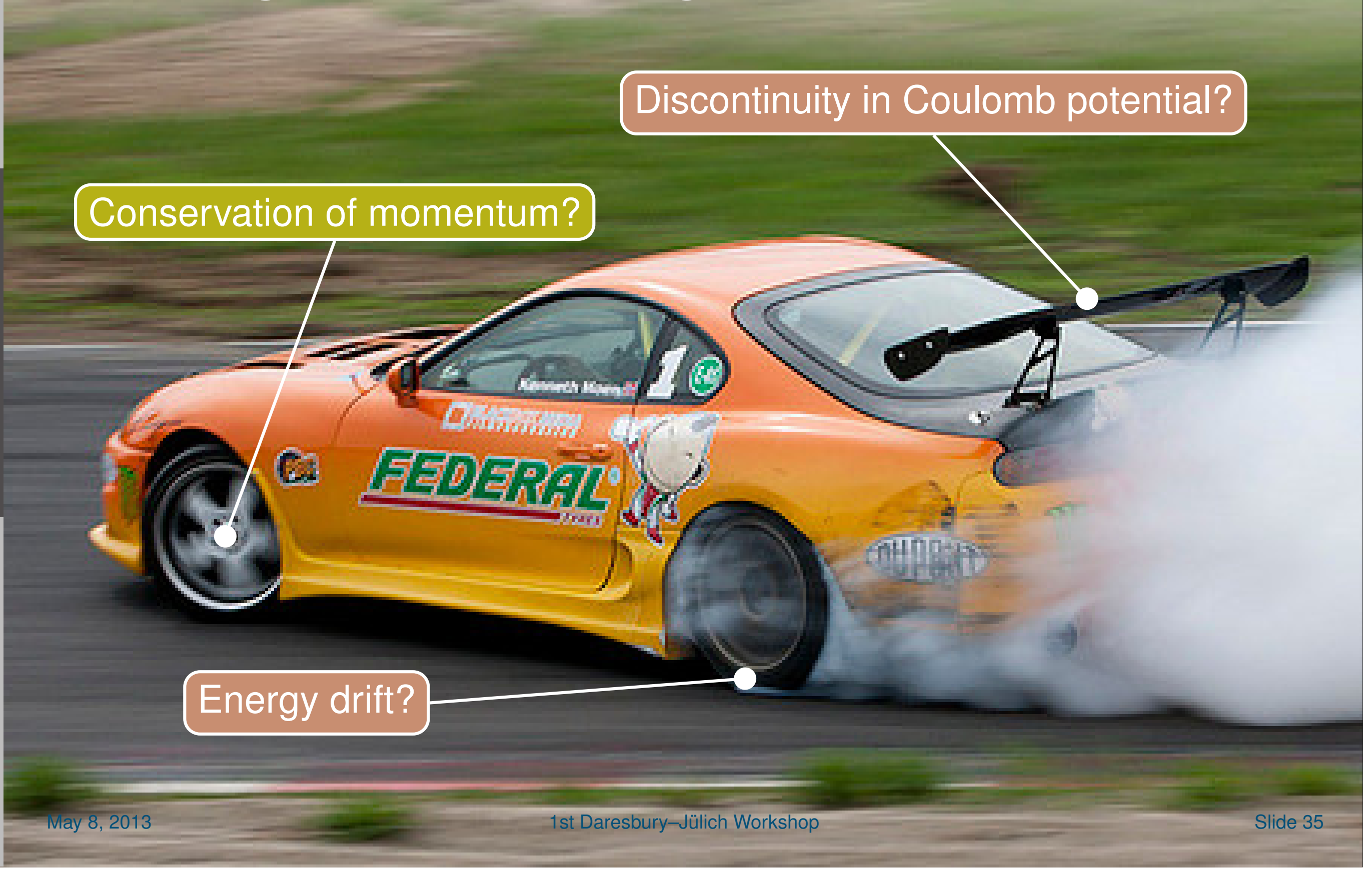

New Algorithmic Features

Dynamic Protonation, (Material provided by C. Kutzner, MPI Göttingen)

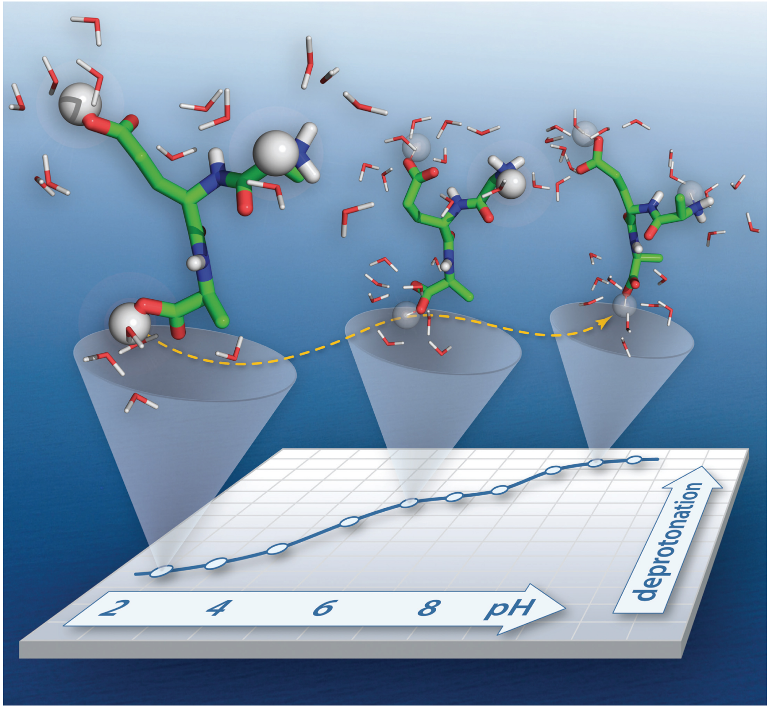

Histidine Protonation States

Material provided by C. Kutzner, MPI Göttingen

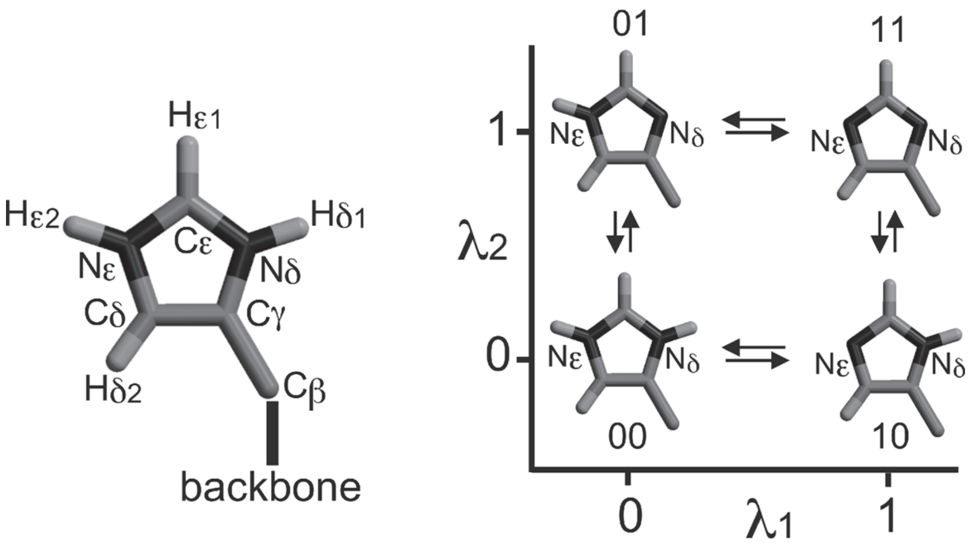

Dynamic Protonation States

Material provided by C. Kutzner, MPI Göttingen

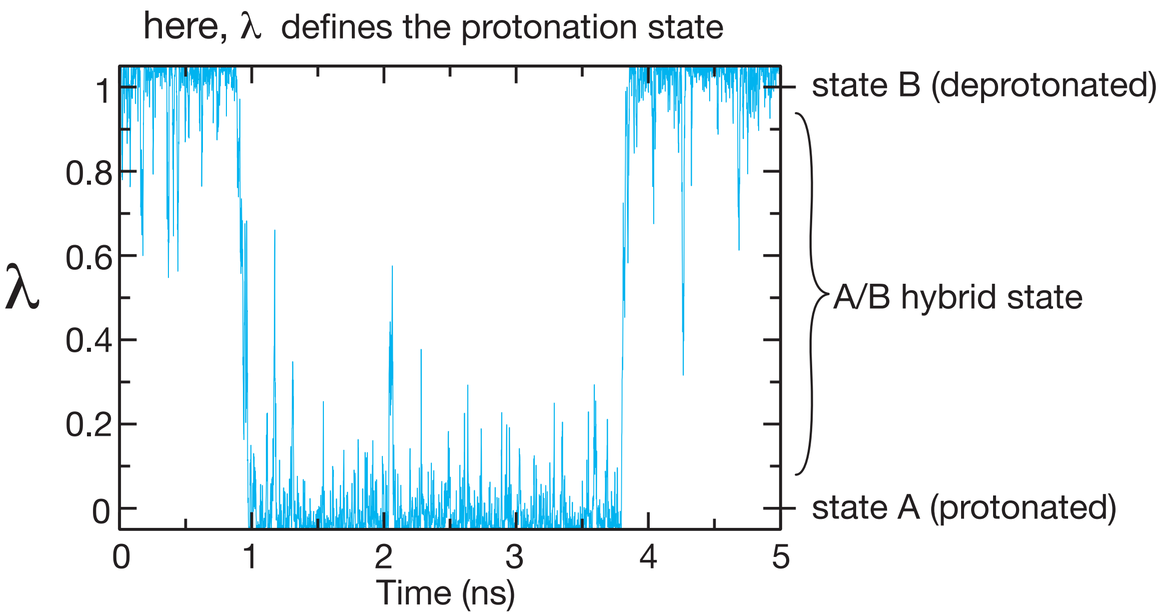

Member of the Helmholtz Association

GROMEX – Next Steps

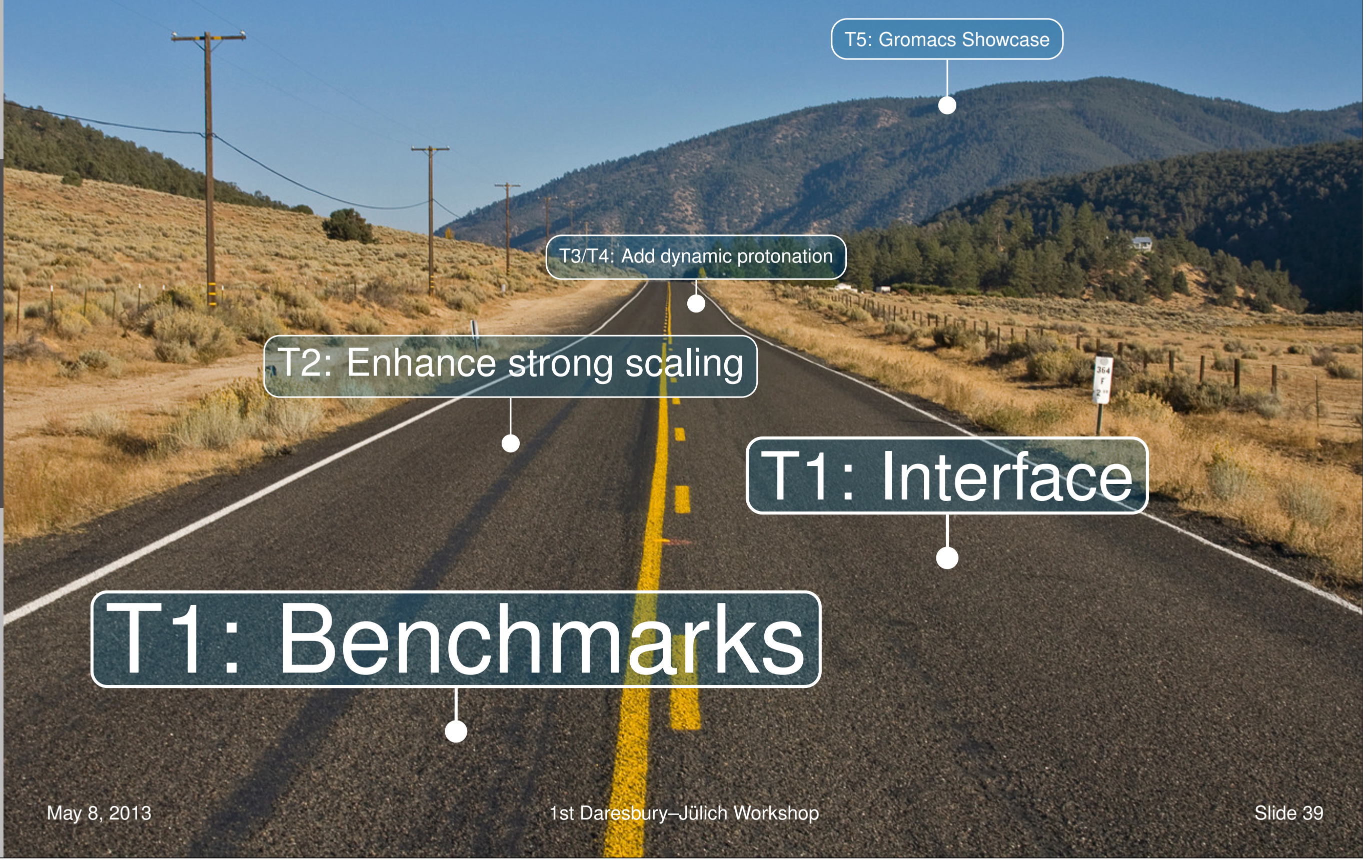

Member of the Helmholtz Association