DRIVING THE EXASCALE TRANSITION

www.flapw.de

fleur

Parallelization and optimization of the FLEUR code: new possibilities for all-electron Density Functional Theory

Uliana Alekseeva, Forschungszentrum Jülich

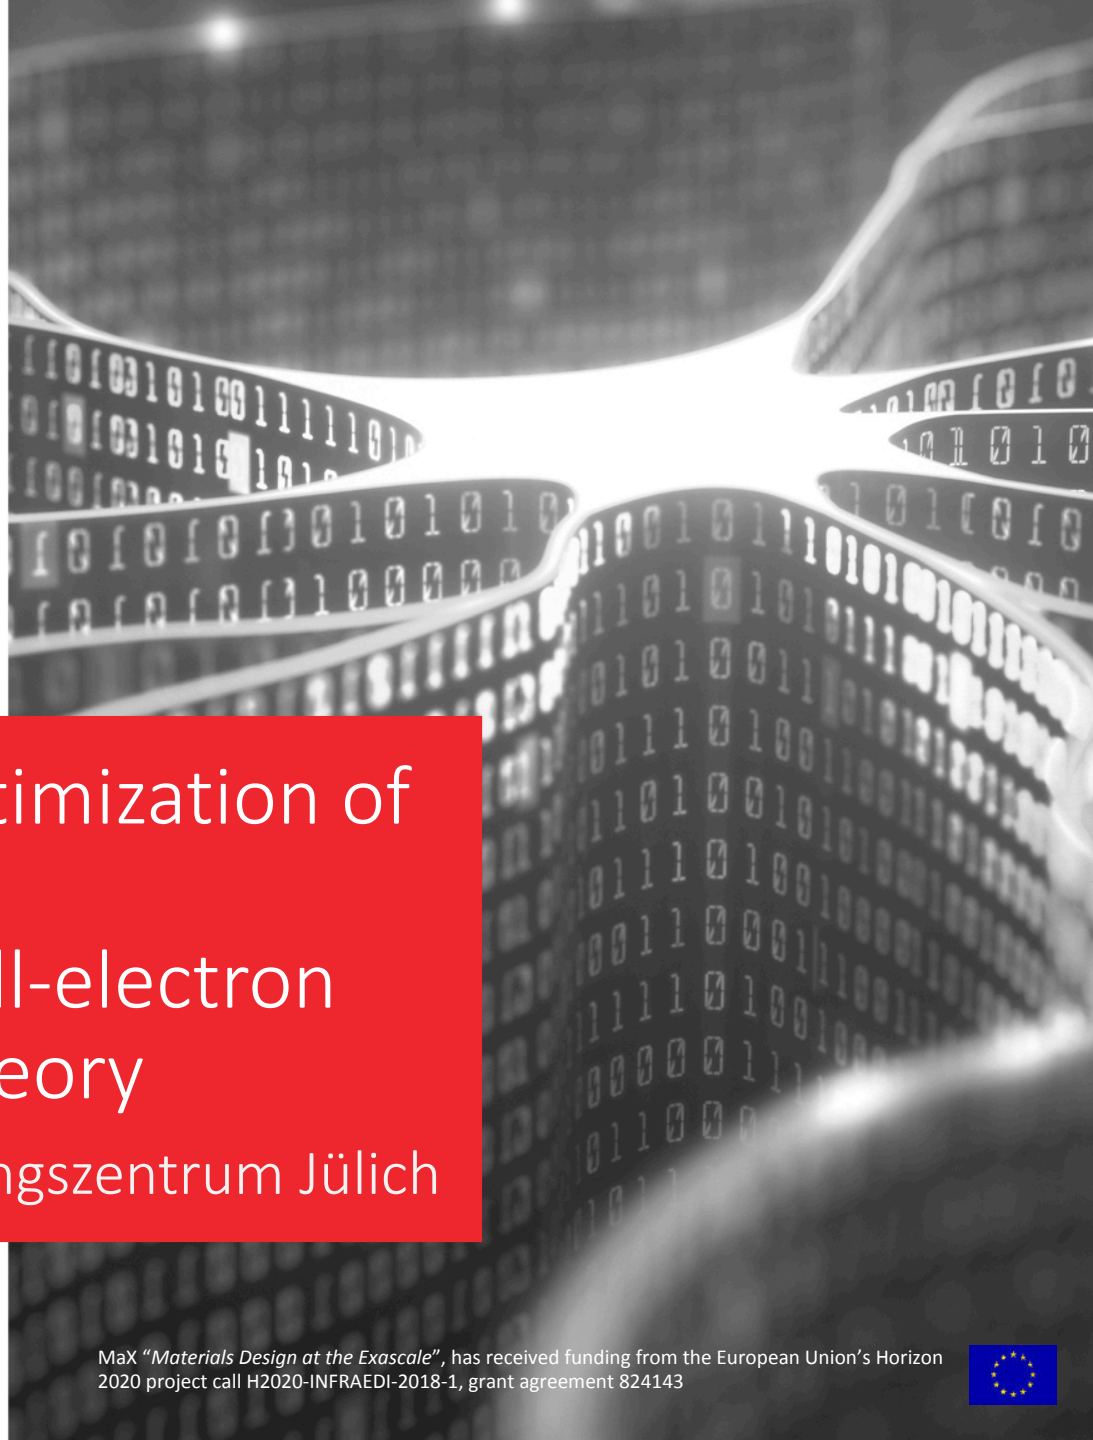

MaX "*Materials Design at the Exascale*", has received funding from the European Union's Horizon 2020 project call H2020-INFRAEDI-2018-1, grant agreement 824143

FLEUR: all-electron code in MaX

www.flapw.de

fleur

- FLEUR is only MaX code **not** using pseudopotential approximation
- Highest possible DFT accuracy
- Reference results
- Systematic convergence

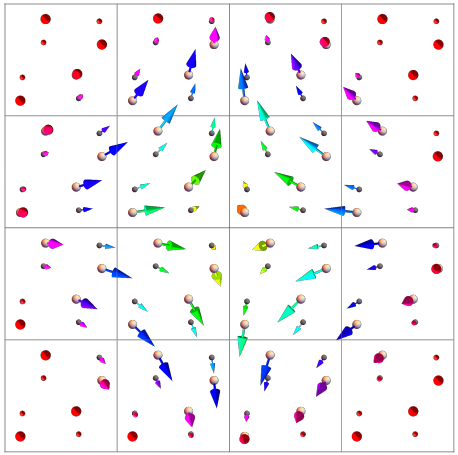

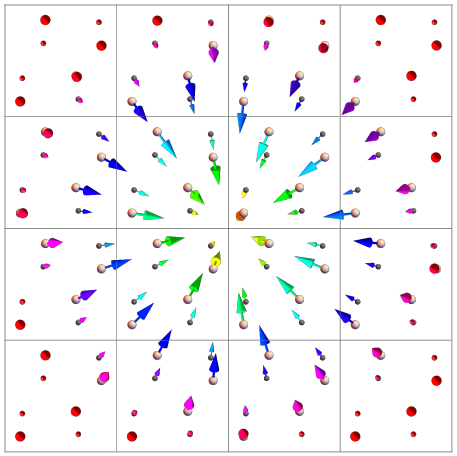

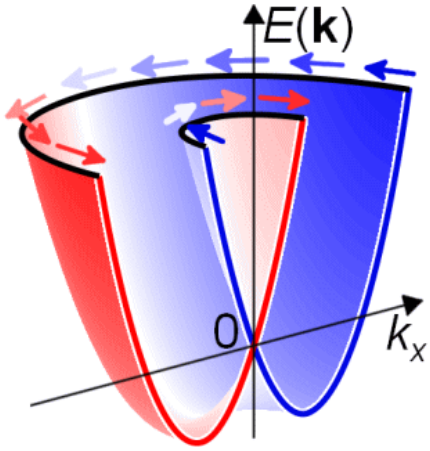

http://ishizaka.t.u-tokyo.ac.jp

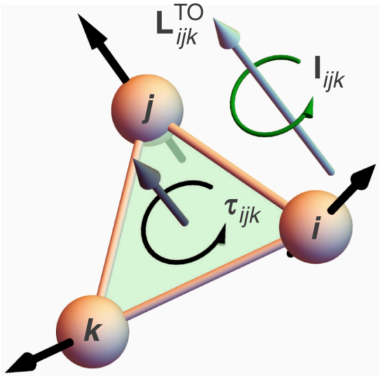

S. Grytsiuk et al. (2020)

FLEUR is the code to simulate

- **Complex non-collinear magnetic phenomena**
- Relativistic effects, Spin-orbit coupling, topological effects
- Effects involving core-electrons

MaX DRIVING THE EXASCALE TRANSITION

Our experience in tuning FLEUR

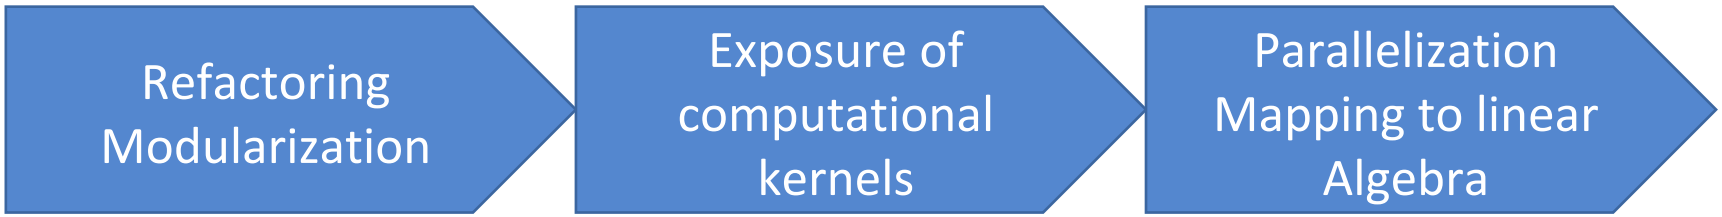

- Our codes contain a multitude of computational kernels
- Use of standard math libraries -> performance portability
- Process has to be repeated for many code sections
- Performance tuning is an iterative process
- Today: hybrid MPI+OpenMP parallelism

Performance optimization during MaX project

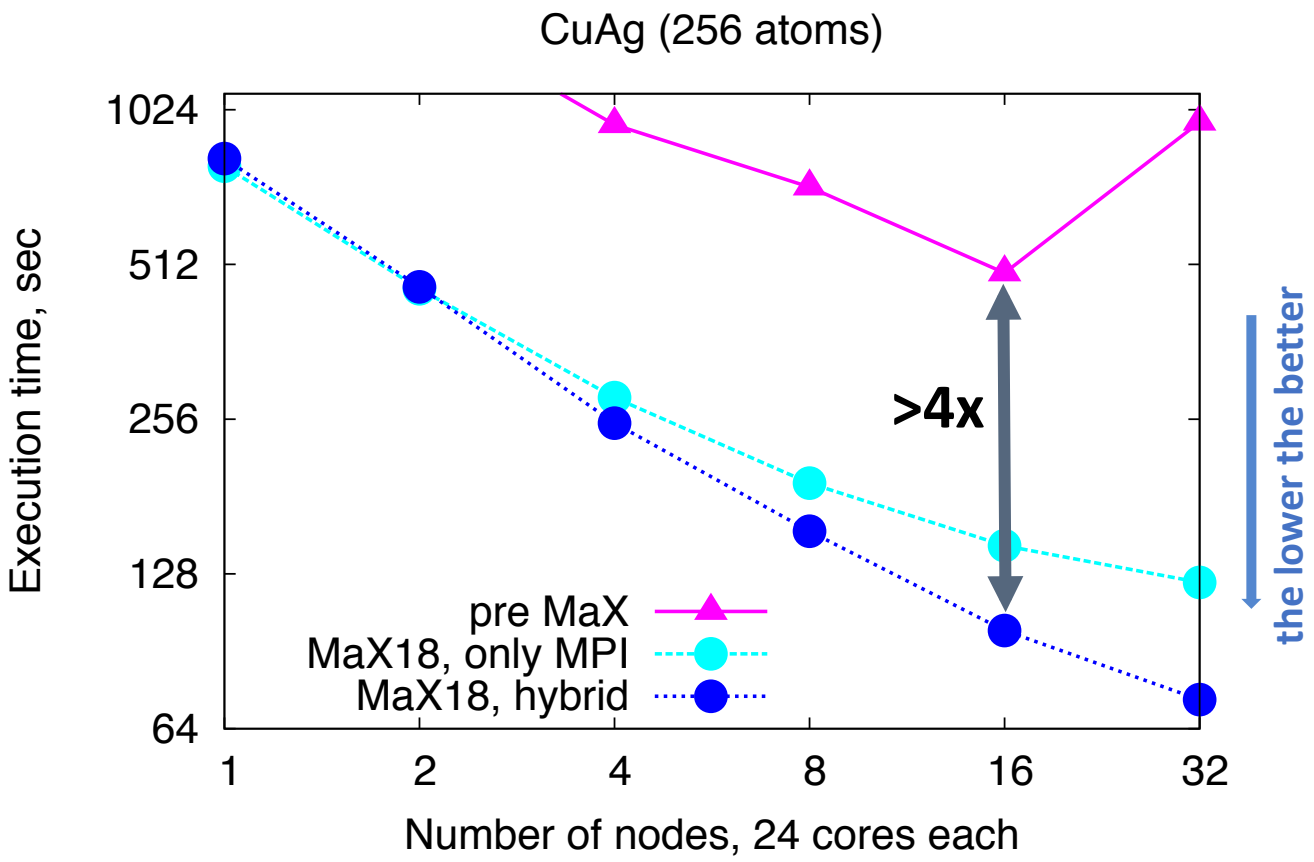

U. Alekseeva, G. Michalicek, D. Wortmann, S. Blügel (2018)
https://doi.org/10.1007/978-3-319-96983-1_52

FLEUR simulations with large unit cells

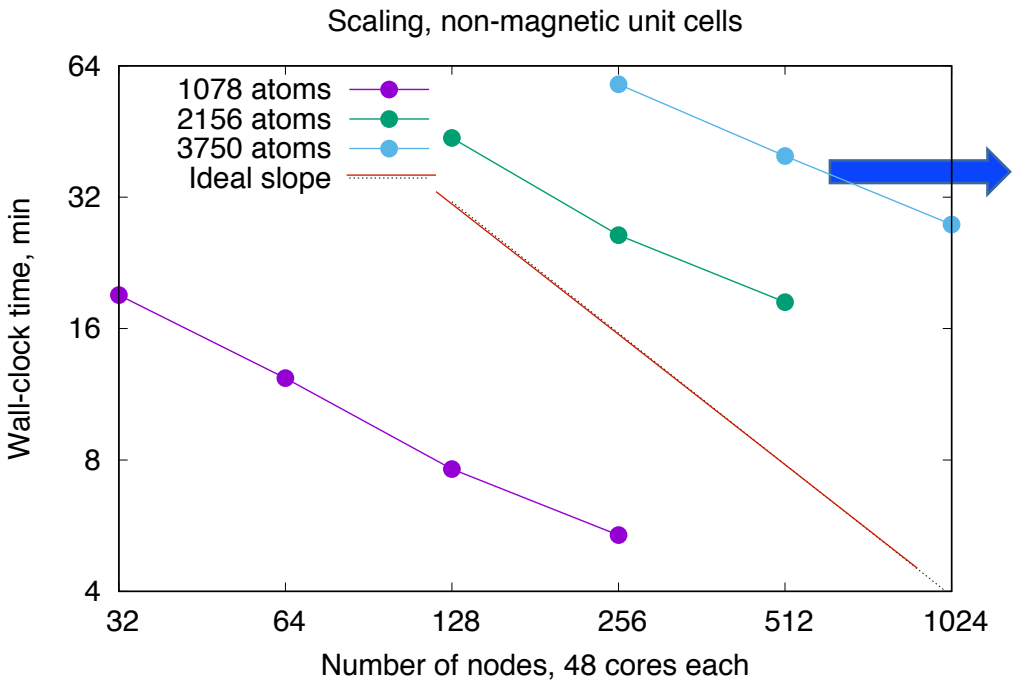

$SrTiO_3$, 3750 atoms, 4 k-points

- 1024 nodes (Intel, 48 cores)
- Size of the matrices:
  345k x 345k (dense Hermitian)
- 100 minutes/ iteration
- 45 iterations to convergence
  (800k core-hours)

Extended defect in STO

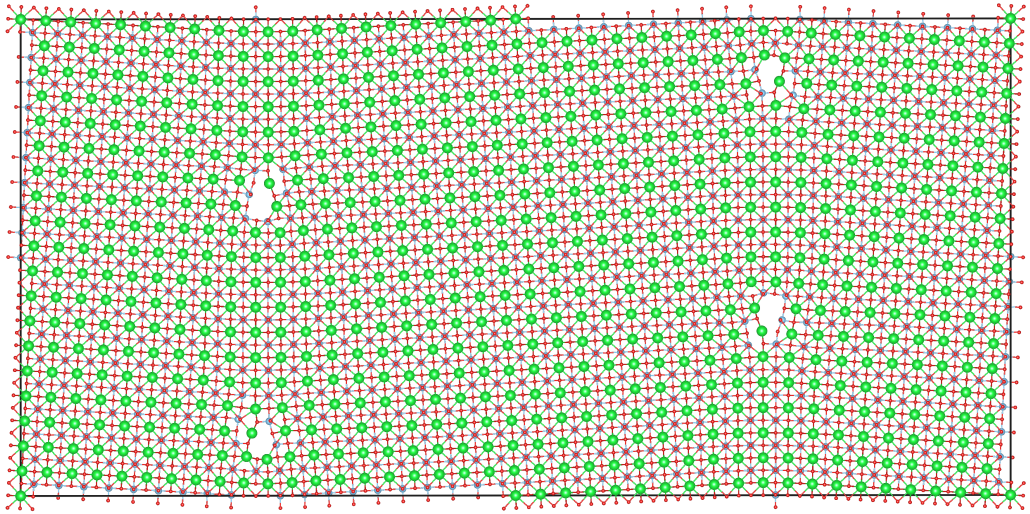

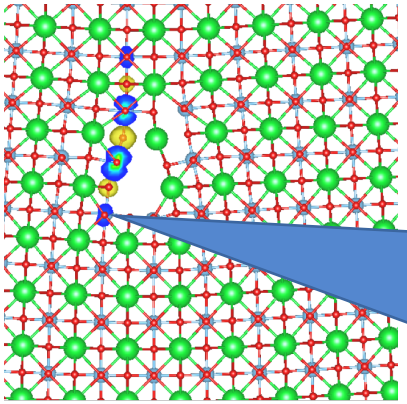

Defect state:
Resistive switching
e.g. ReRAM

Simulations of nano-sized magnetic objects

Nanomagnetism:

- Complex spin-structures on the atomic scale
- Topological protection leads to stabilization
- Possible candidates for e.g. data-processing, neuromorphic or reservoir computing

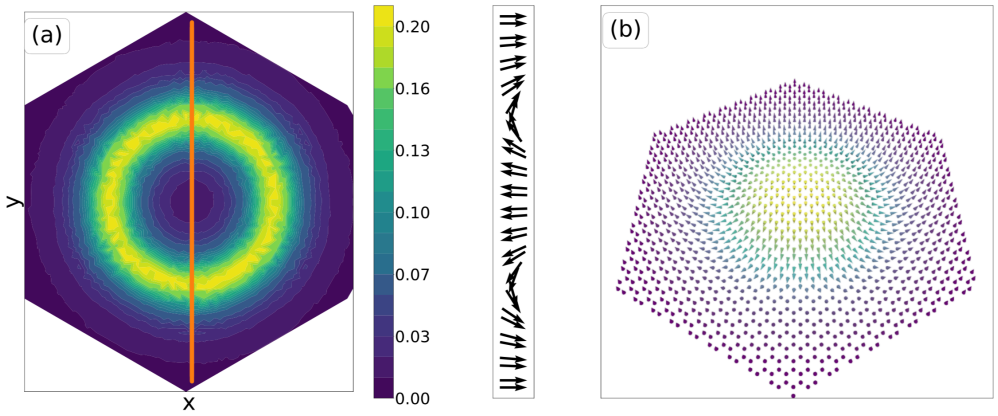

M.Redies, Masterthesis 2019 RWTH

Basic concepts can be studied by multiscale approach:

1. DFT to extract material properties
2. Spin-dynamic simulations

Large scale FLEUR simulations:

- Verify multiscale model
- Magnetism influences with electronic structure

Large magnetic setups: a globule

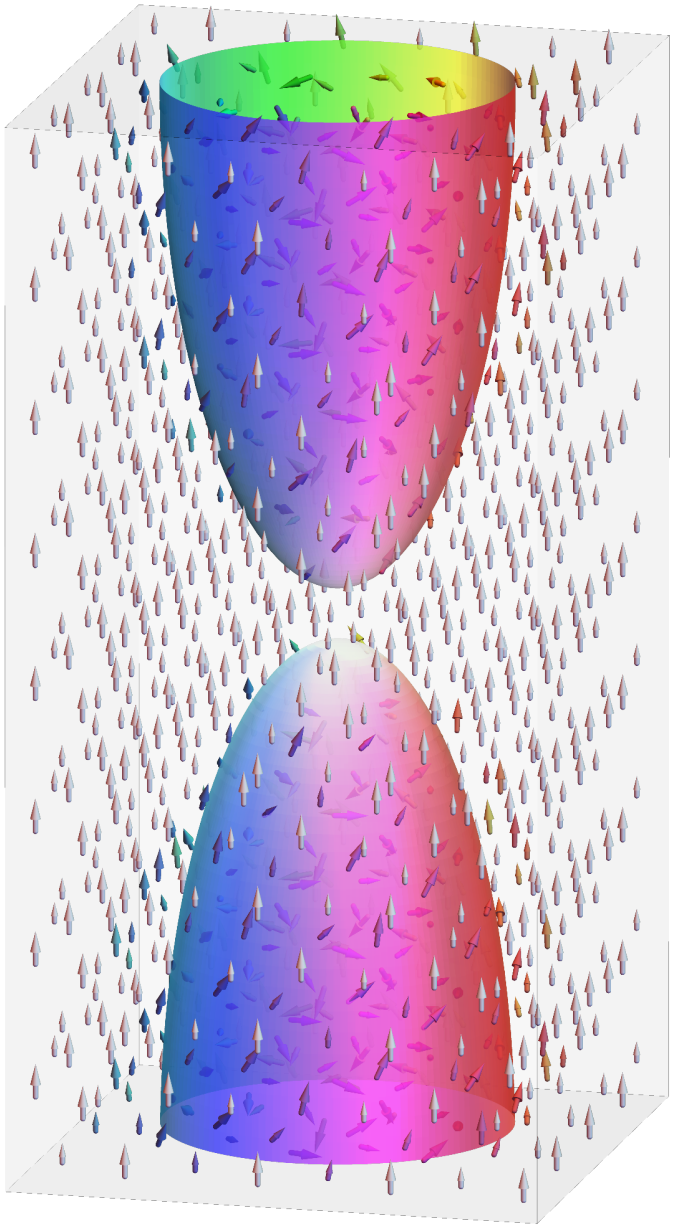

MnGe Supercell 4x4x8, 1024 atoms, 1 k-point

- 256 nodes (Intel, 48 cores)
- Size of the matrices:
  156k x 156k (dense Hermitian)
- 25 minutes / iteration
- 100 iterations to convergence
  (2 Mio core-hours)

SuperMUC-NG

A Path to Exascale Computing

Today:

Single magnetic structure

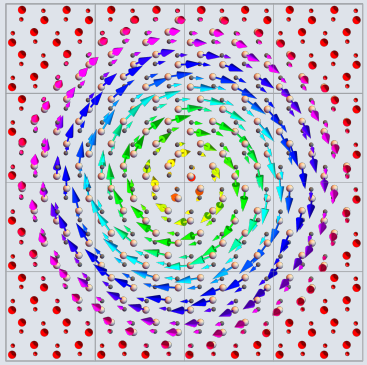

- several hundred nodes
- Mio core-hour

Magnetic structures are manifold

- Easy upscaling with HTC
- Intrinsically parallel

AiiDA

Magnetic structures evolve

- Time dependence easily requires orders of magnitude more computational effort

In a nutshell

FLEUR code

- all-electron, full-potential DFT

www.flapw.de

fleur

Optimization for HPC computers

- hybrid (MPI+OpenMP) parallelization
- unit cells with >1000 atoms

HPC examples

- simulations of large setups

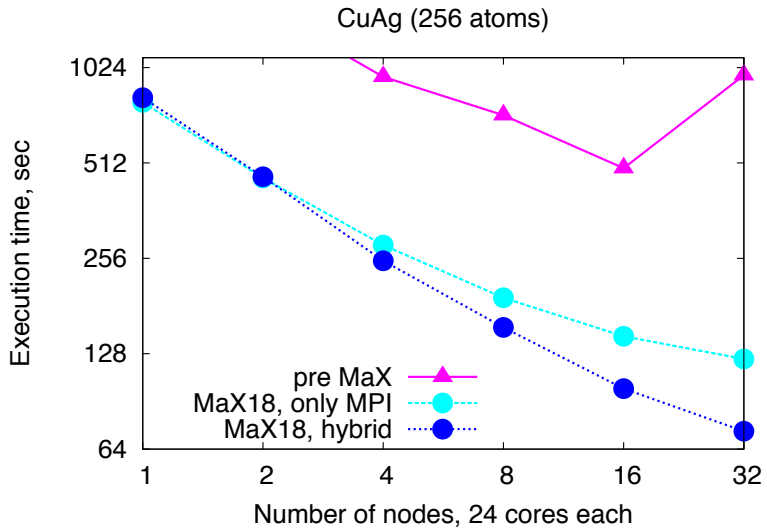

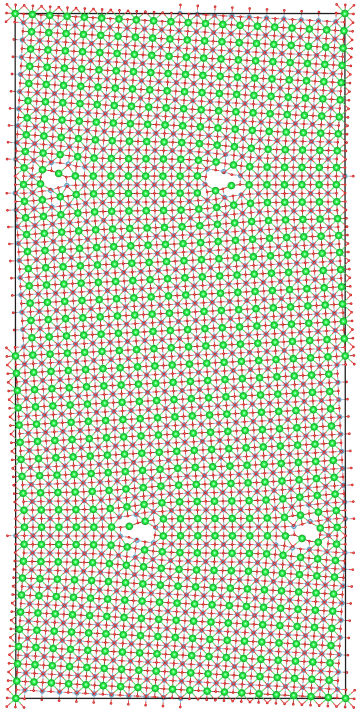

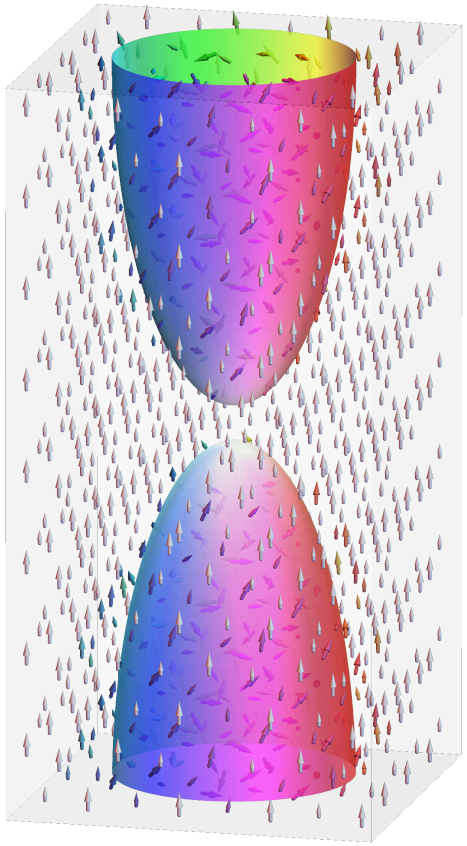

Computationally intensive parts of FLEUR

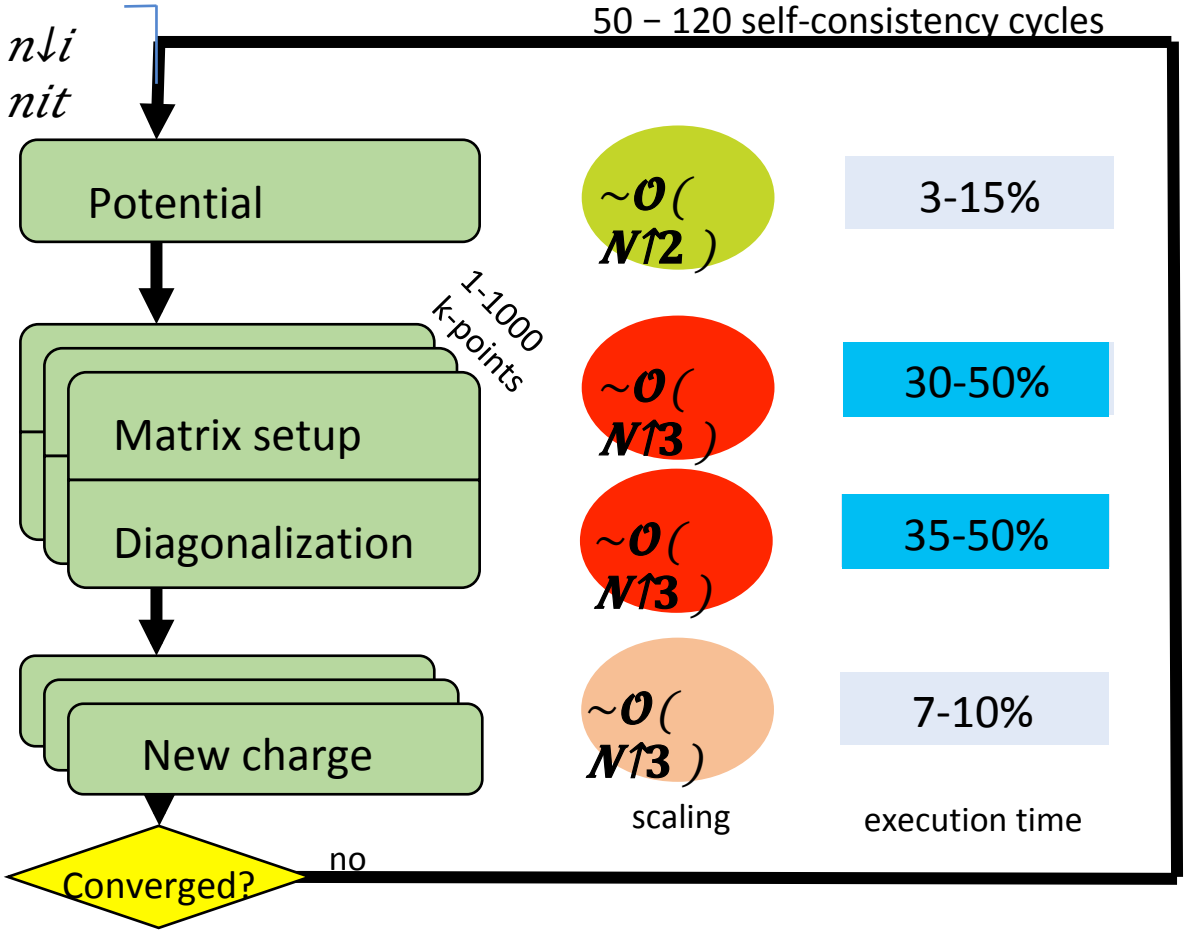

Parallelization

Levels of parallelization:

- MPI over k-points
- MPI eigenvector parallelization
- OpenMP parallelization
- SIMD Vectorization

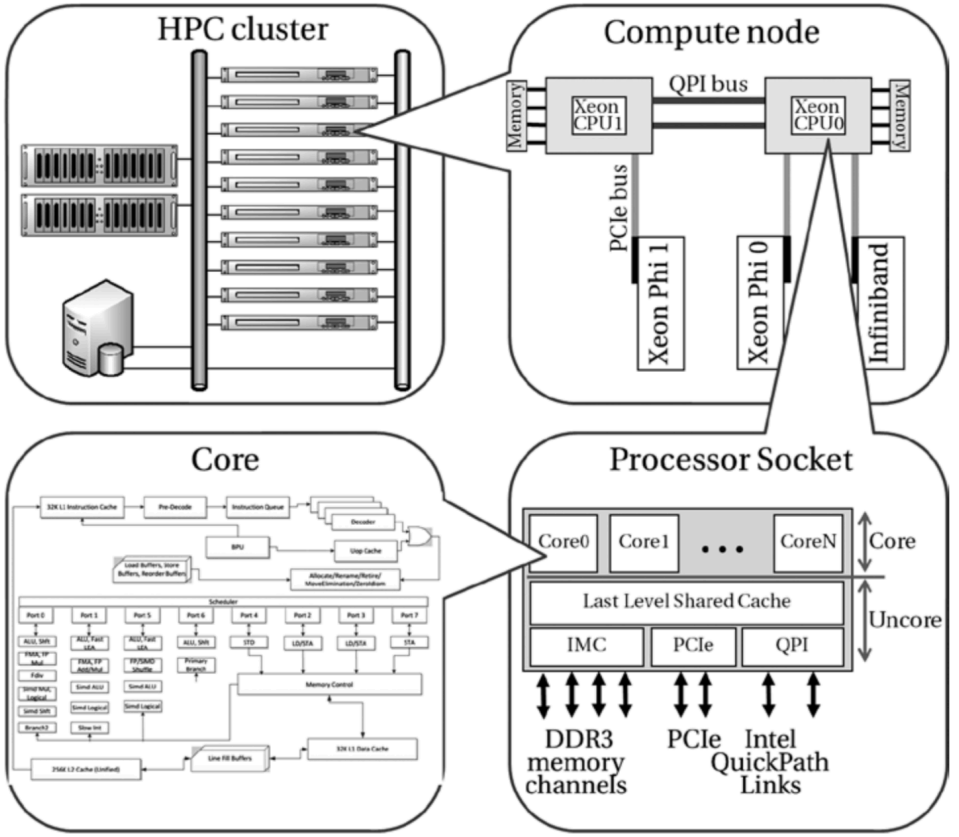

Supalov, Semin, Klemm, Dahnken
"Optimizing HPC applications with Intel Cluster Tools"

Large magnetic setups: skyrmion tubes

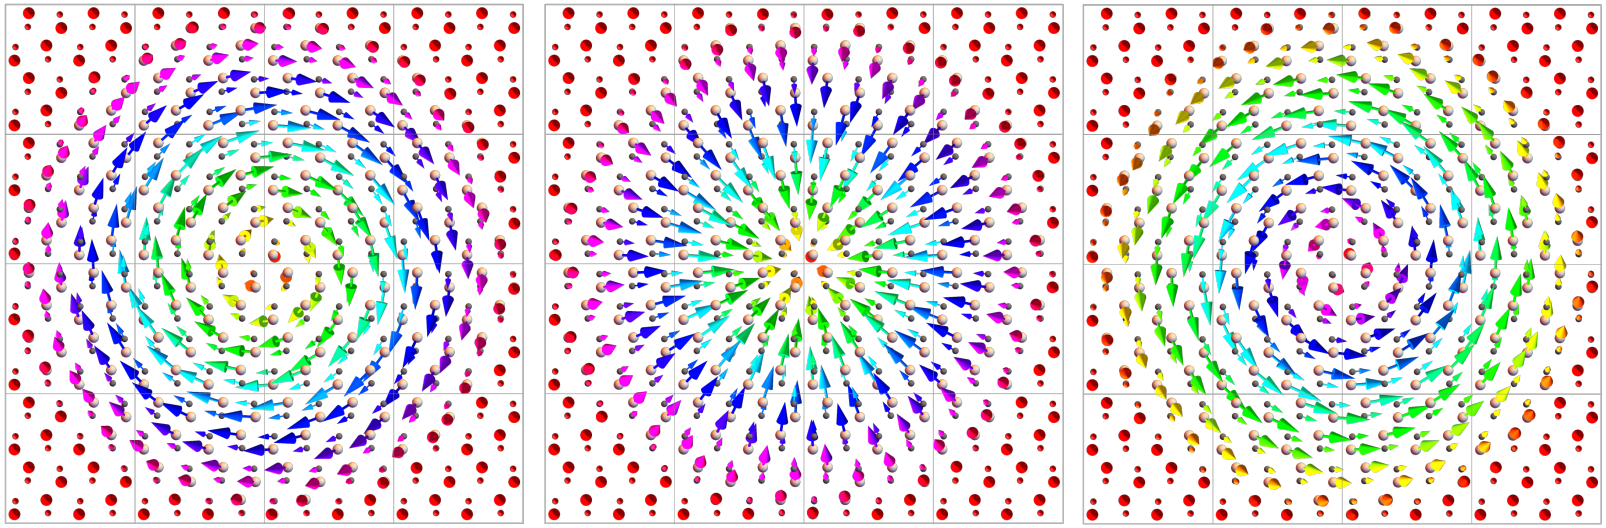

MnGe 8x8x1 (512 atoms), SuperMUC-NG

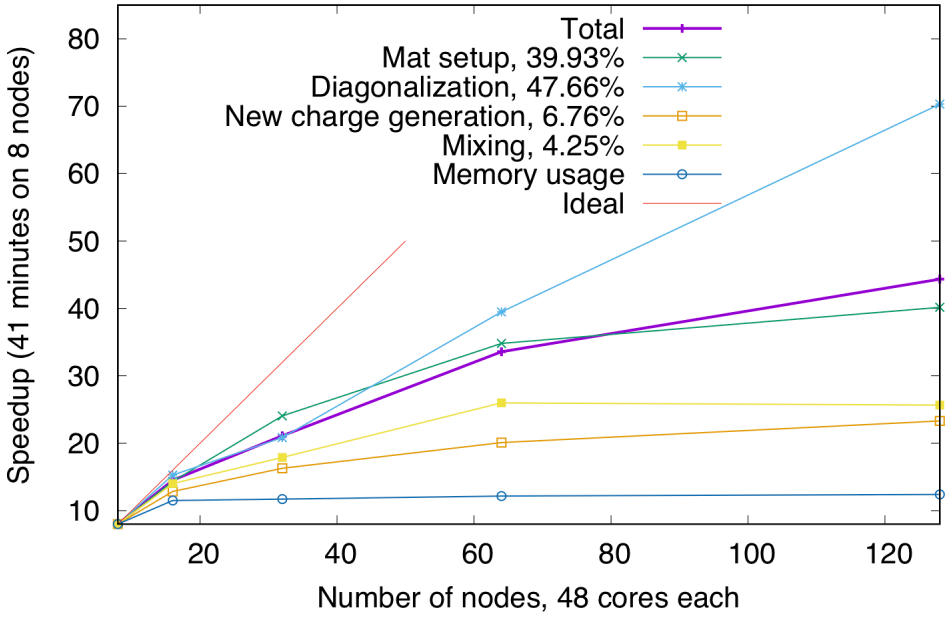

MnGe Supercell 8x8x1, 512 atoms, 4 k

- 256 nodes (64 nodes/ k-point)
- Size of the matrices:
  78k x 78k (dense Hermitian)
- 10 min/ iteration
- 66 iterations to convergence
  (135k core-hours)