Ultra-Skalierbare Multiphysiksimulationen für Erstarrungsprozesse in Metallen (SKAMPY)

HPC-Status-Konferenz der Gauß-Allianz, 29. November 2016, Hamburg

Harald Köstler, Bauer, Schornbaum, Godenschwager, Rüde, Hammer, Wellein, Hötzer, Nestler

Chair for System Simulation
Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

Karlsruhe Institute of Technology

Hochschule Karlsruhe Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

SKAMPY Project

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Karlsruhe Institute of Technology

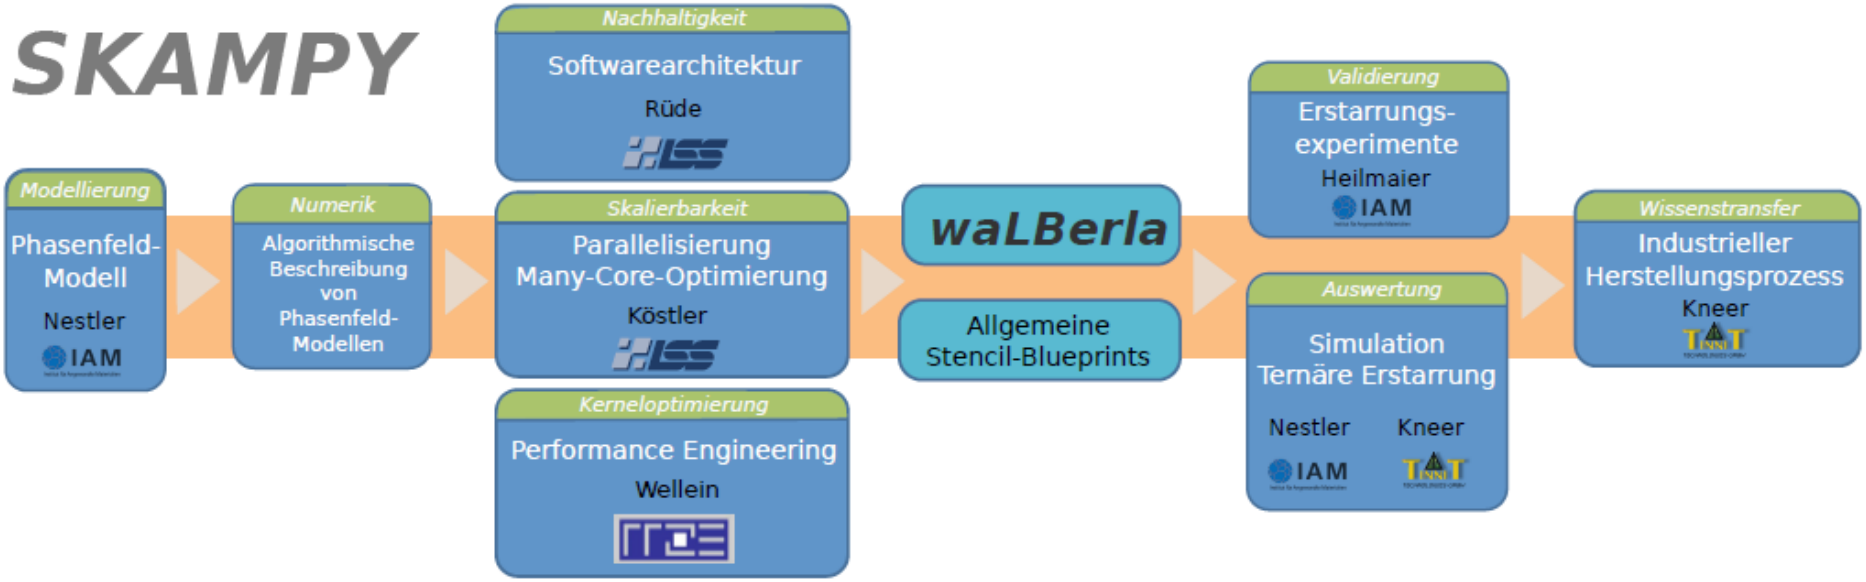

Bundesministerium für Bildung und Forschung

GA Gauß-Allianz

Outline

- waLBerla Framework
- SKAMPY Project

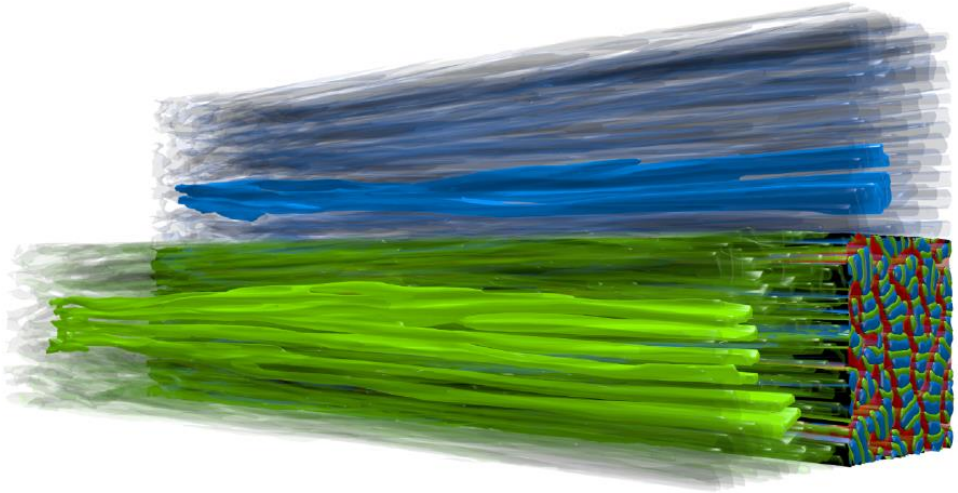

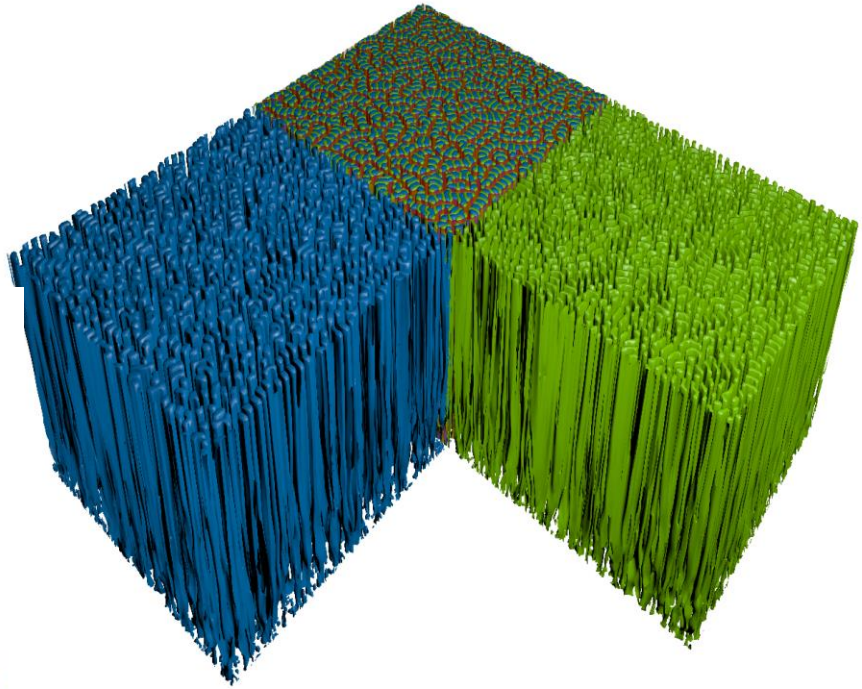

The *waLBerla* Framework

Karlsruhe Institute of Technology

Hochschule Karlsruhe
Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

waLBerla Framework

waLBerla Framework

- widely applicable Lattice-Boltzmann from Erlangen
- HPC software framework, originally developed for CFD simulations with Lattice Boltzmann Method (LBM)
- evolved into general framework for algorithms on block-structured grids
- *www.walberla.net*

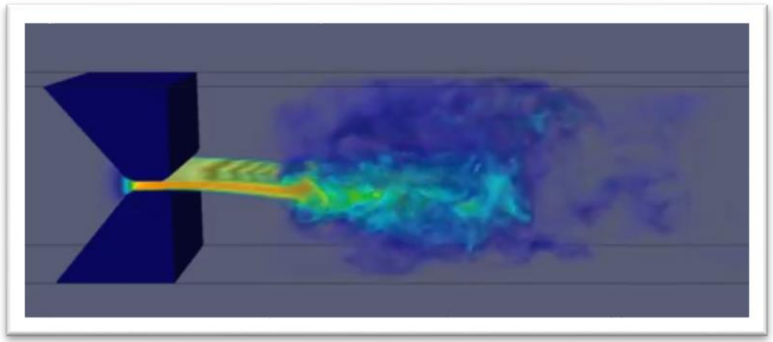

Vocal Fold Study
(Florian Schornbaum)

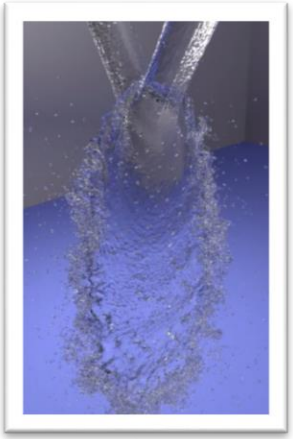

Free Surface Flow

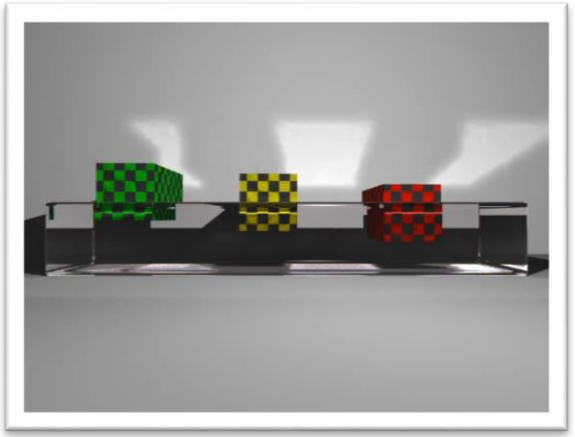

Fluid Structure Interaction
(Simon Bogner)

Block-structured Grids

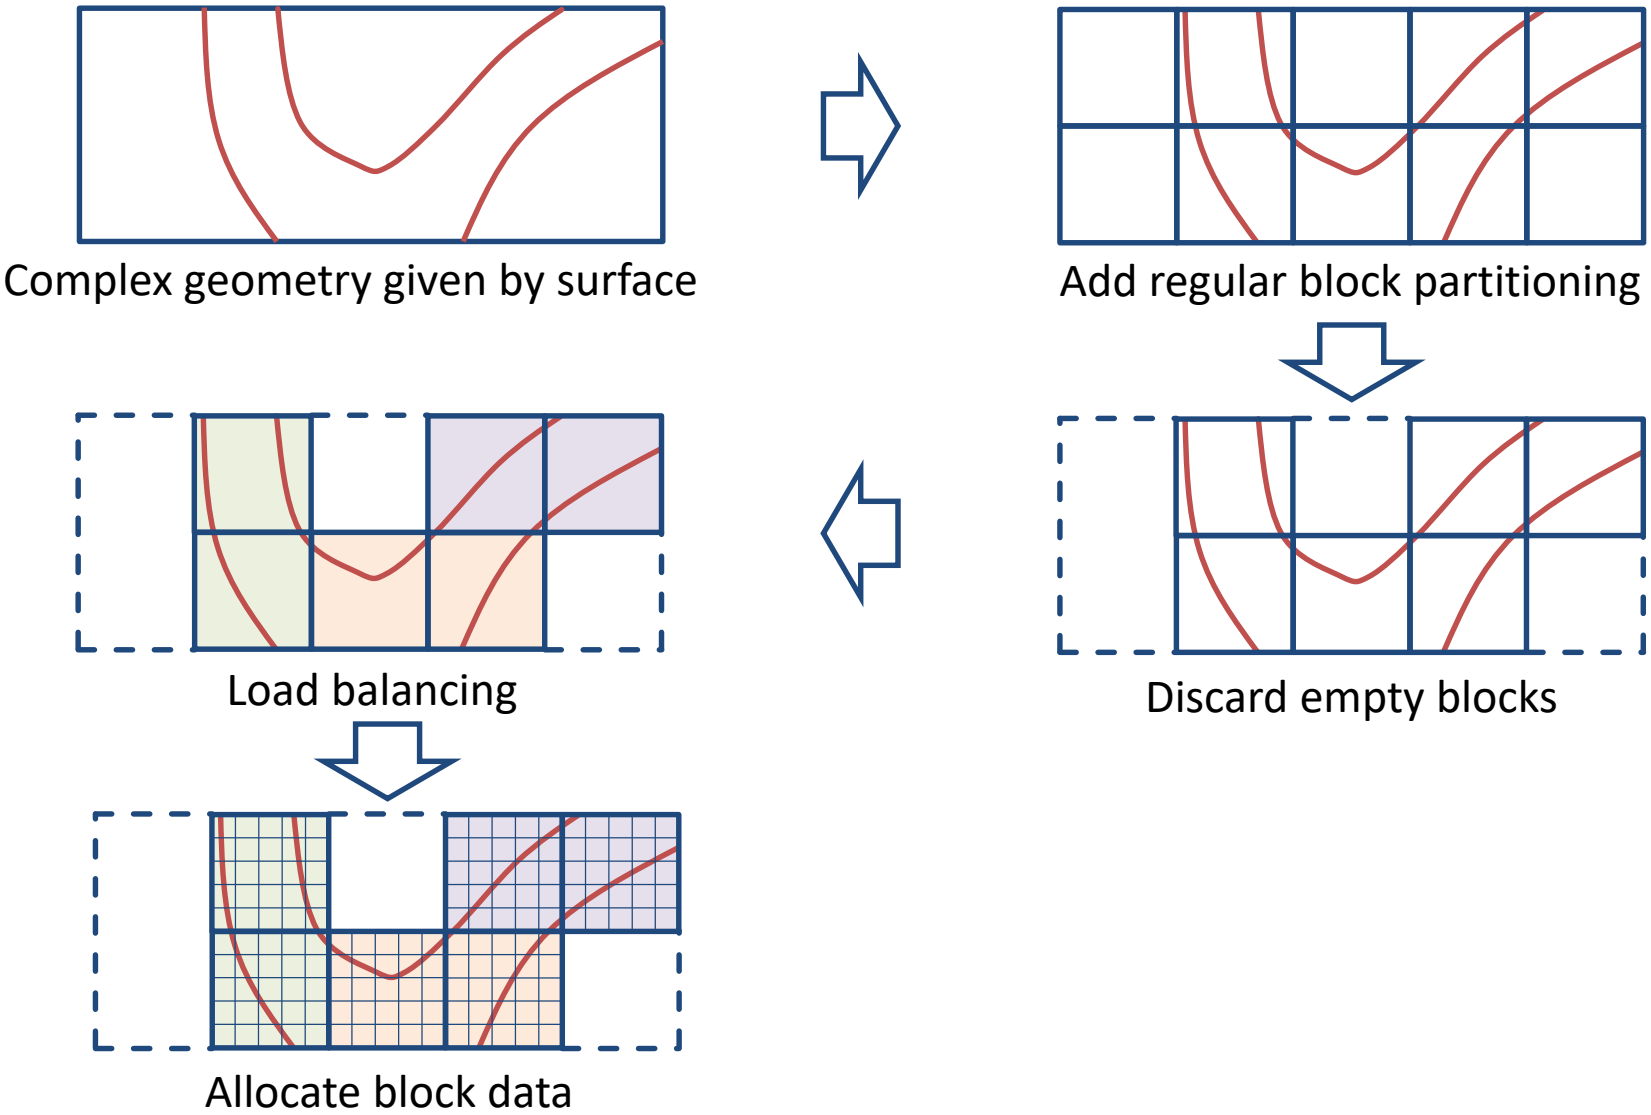

Block-structured Grids

- Domain Decomposition & Distribution to Processes:
  - regular decomposition into blocks containing uniform grids

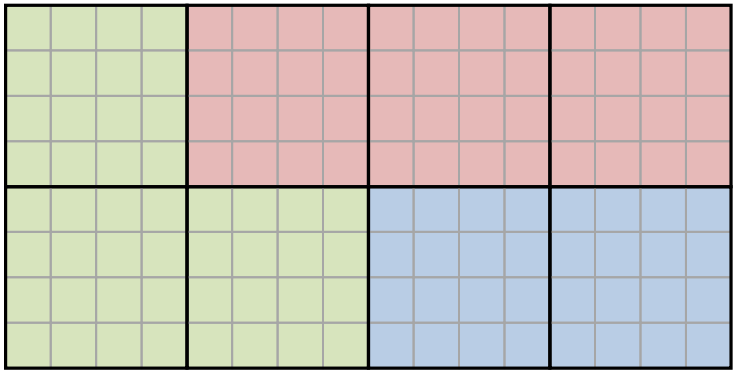

In most cases, if a regular decomposition of a uniform grid is used, exactly one block is assigned to each process.

  - grid refinement: octree-like decomposition

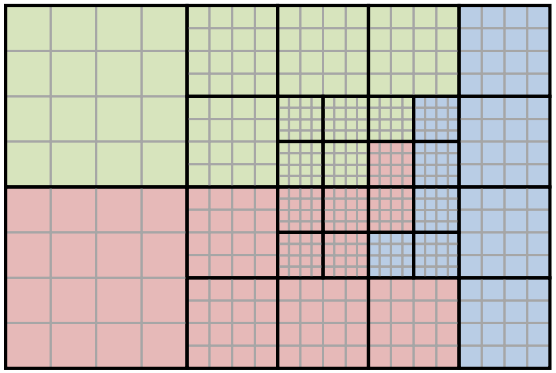

forest of octrees:
each block contains a uniform grid of the same size
→ 2:1 balance between neighboring cells on level transitions

Hybrid Parallelization

- Distributed Memory Parallelization: MPI
  - data exchange on borders between blocks via ghost layers

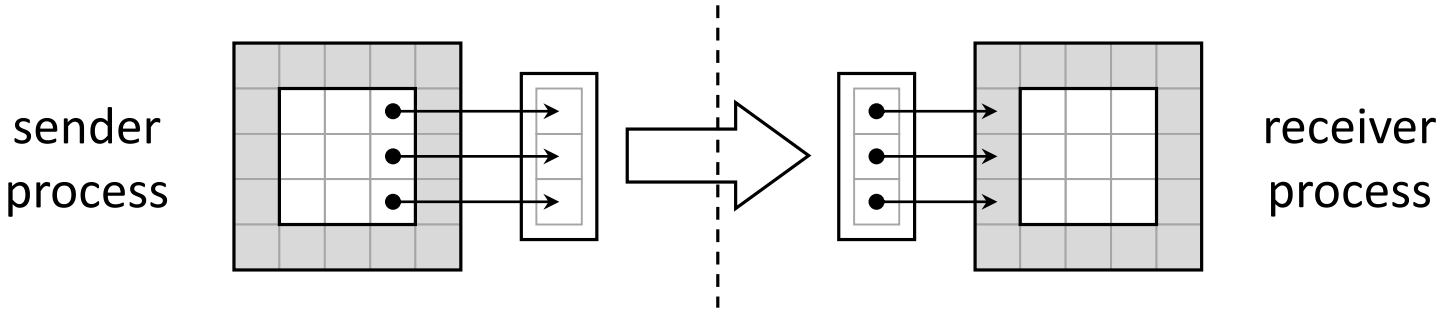

(slightly more complicated for non-uniform domain decompositions, but the same general ideas still apply)

  - support for overlapping communication and computation
  - some advanced models require more complex communication patterns (e.g. free-surface and fluid-structure interaction)

SKAMPY Project

Application

Karlsruhe Institute of Technology

Hochschule Karlsruhe
Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Overview

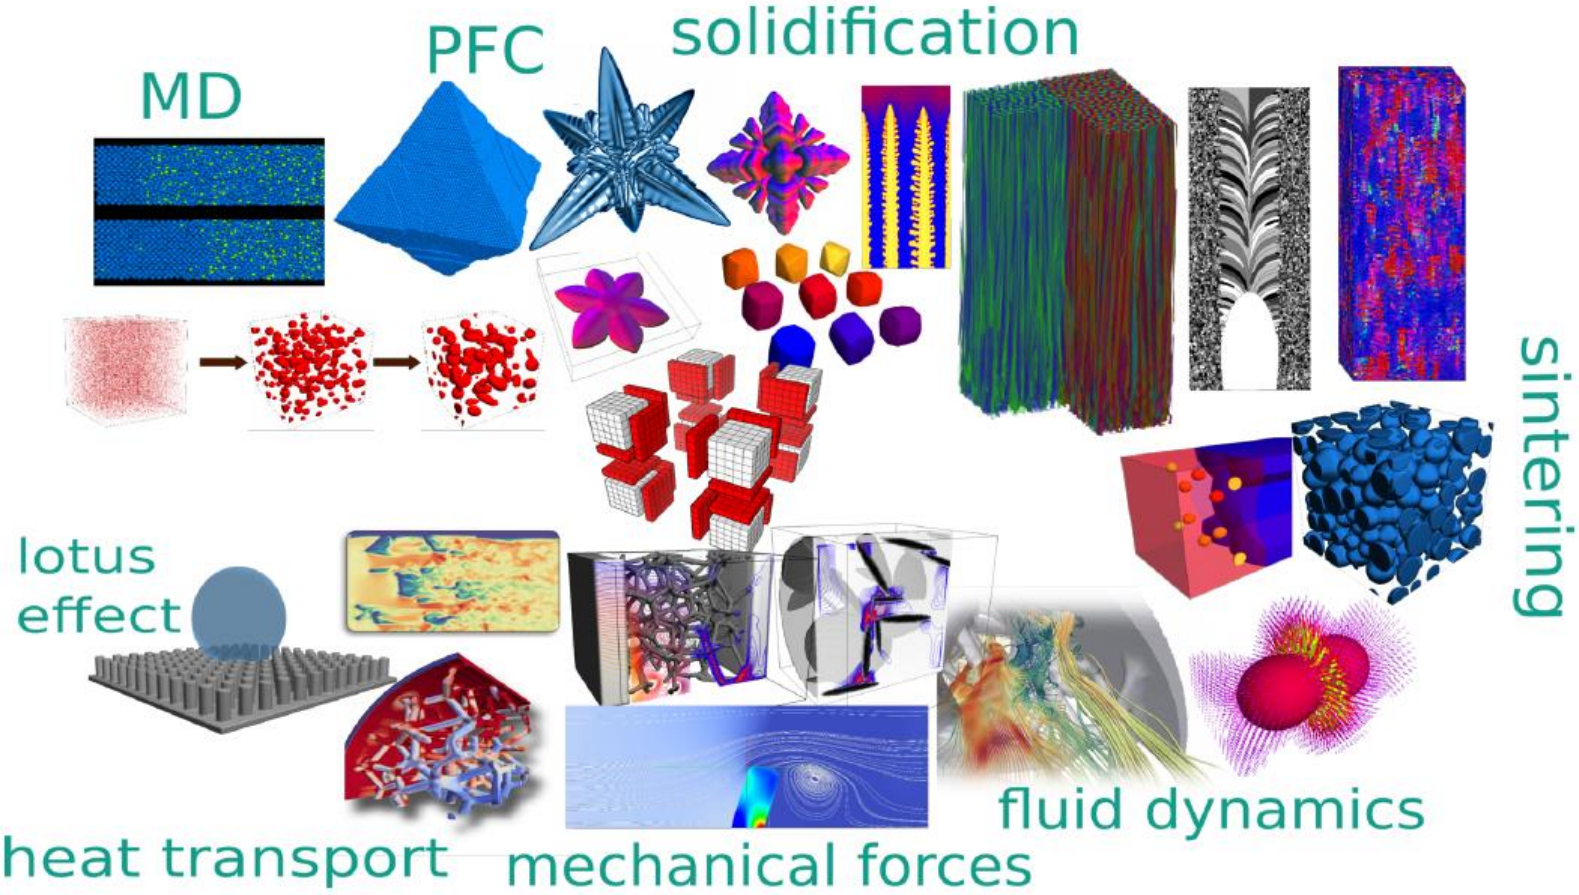
MD
PFC
solidification
sintering
lotus
effect
heat transport
mechanical forces
fluid dynamics

Application Setting

- ternary eutectic alloys
- directional solidification
- analytically moving
- temperature gradient
- massively parallel phase-field simulations
- large domain sizes

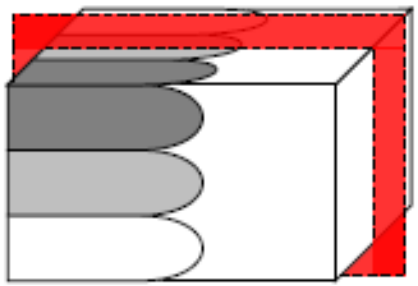

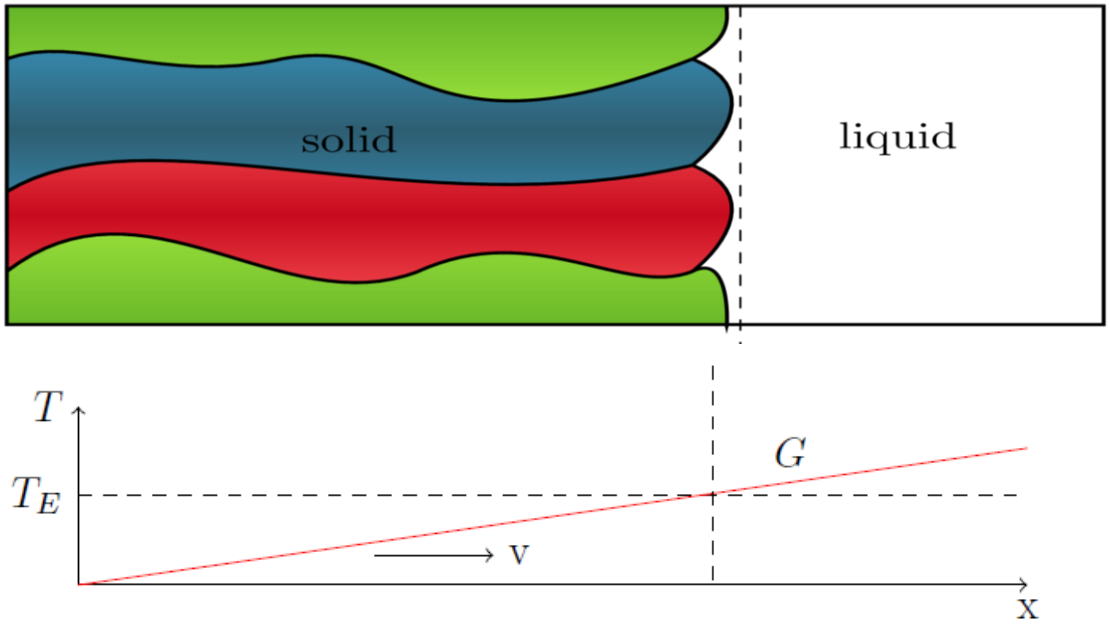

Overview

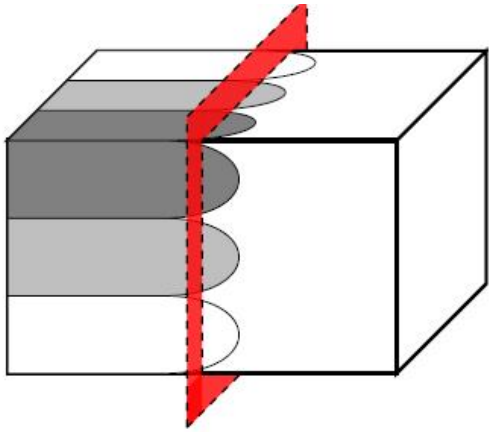

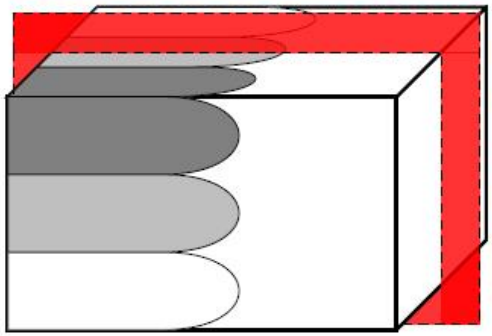

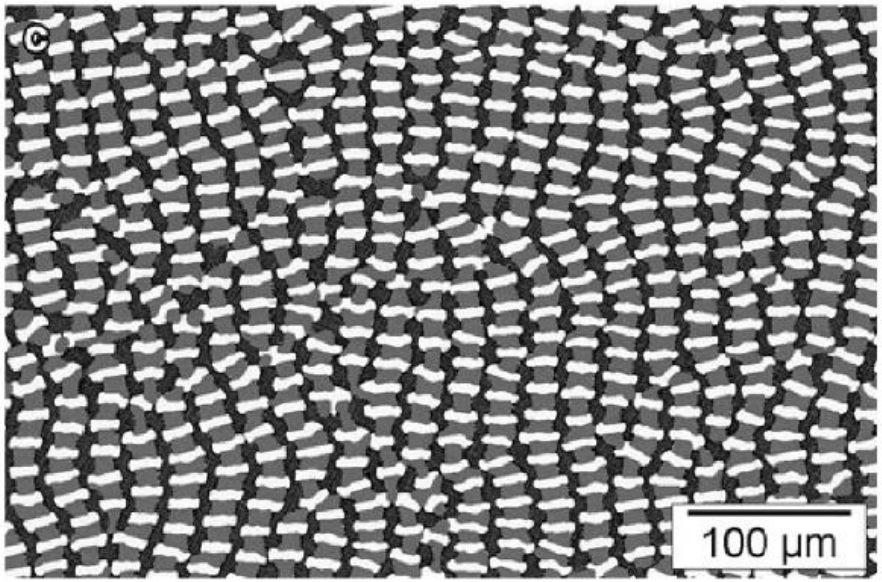

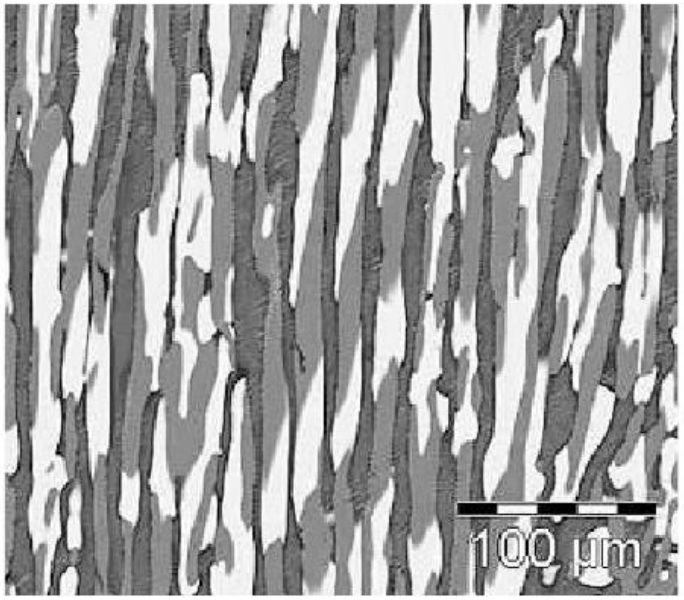

Dennstedt, A., & Ratke, L. (2012). Transactions of the Indian Institute of Metals, 65(6), 777-782.
Genau, A., & Ratke, L. (2012). International Journal of Materials Research, 103(4), 469-475.

Phase-field model

- Grand chemical potential functional:

$$\Psi(\phi, \mu, T) = \int_\Omega \underbrace{\left(\epsilon a(\phi, \nabla\phi) + \frac{1}{\epsilon}\omega(\phi) \right)}_{\text{surface energy}} + \underbrace{\psi(\phi, \mu, T)}_{\text{bulk potential}} \, d\Omega$$

- phase-field vector $\phi = (\phi_1, \phi_2, ..., \phi_N)^T$
- order parameter ϕ_α represents the volume fraction of each phase
- volumetric interface at the surface
- smooth transition between the order parameters
- Allen-Cahn type variational differentiation of the functional
- $\rightarrow$ **no interface tracking needed**

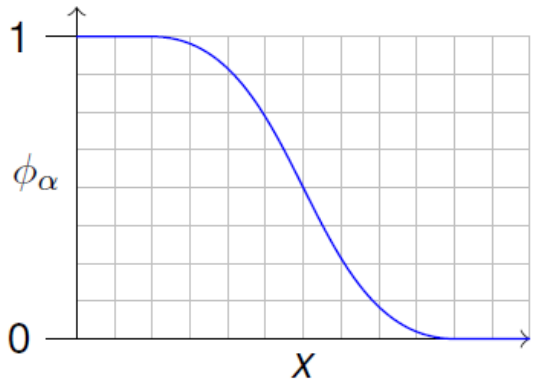

Nestler, B., Garcke, H., & Stinner, B. (2005). Physical Review E, 71(4), 041609.

Microstructure prediction Al-Ag-Cu

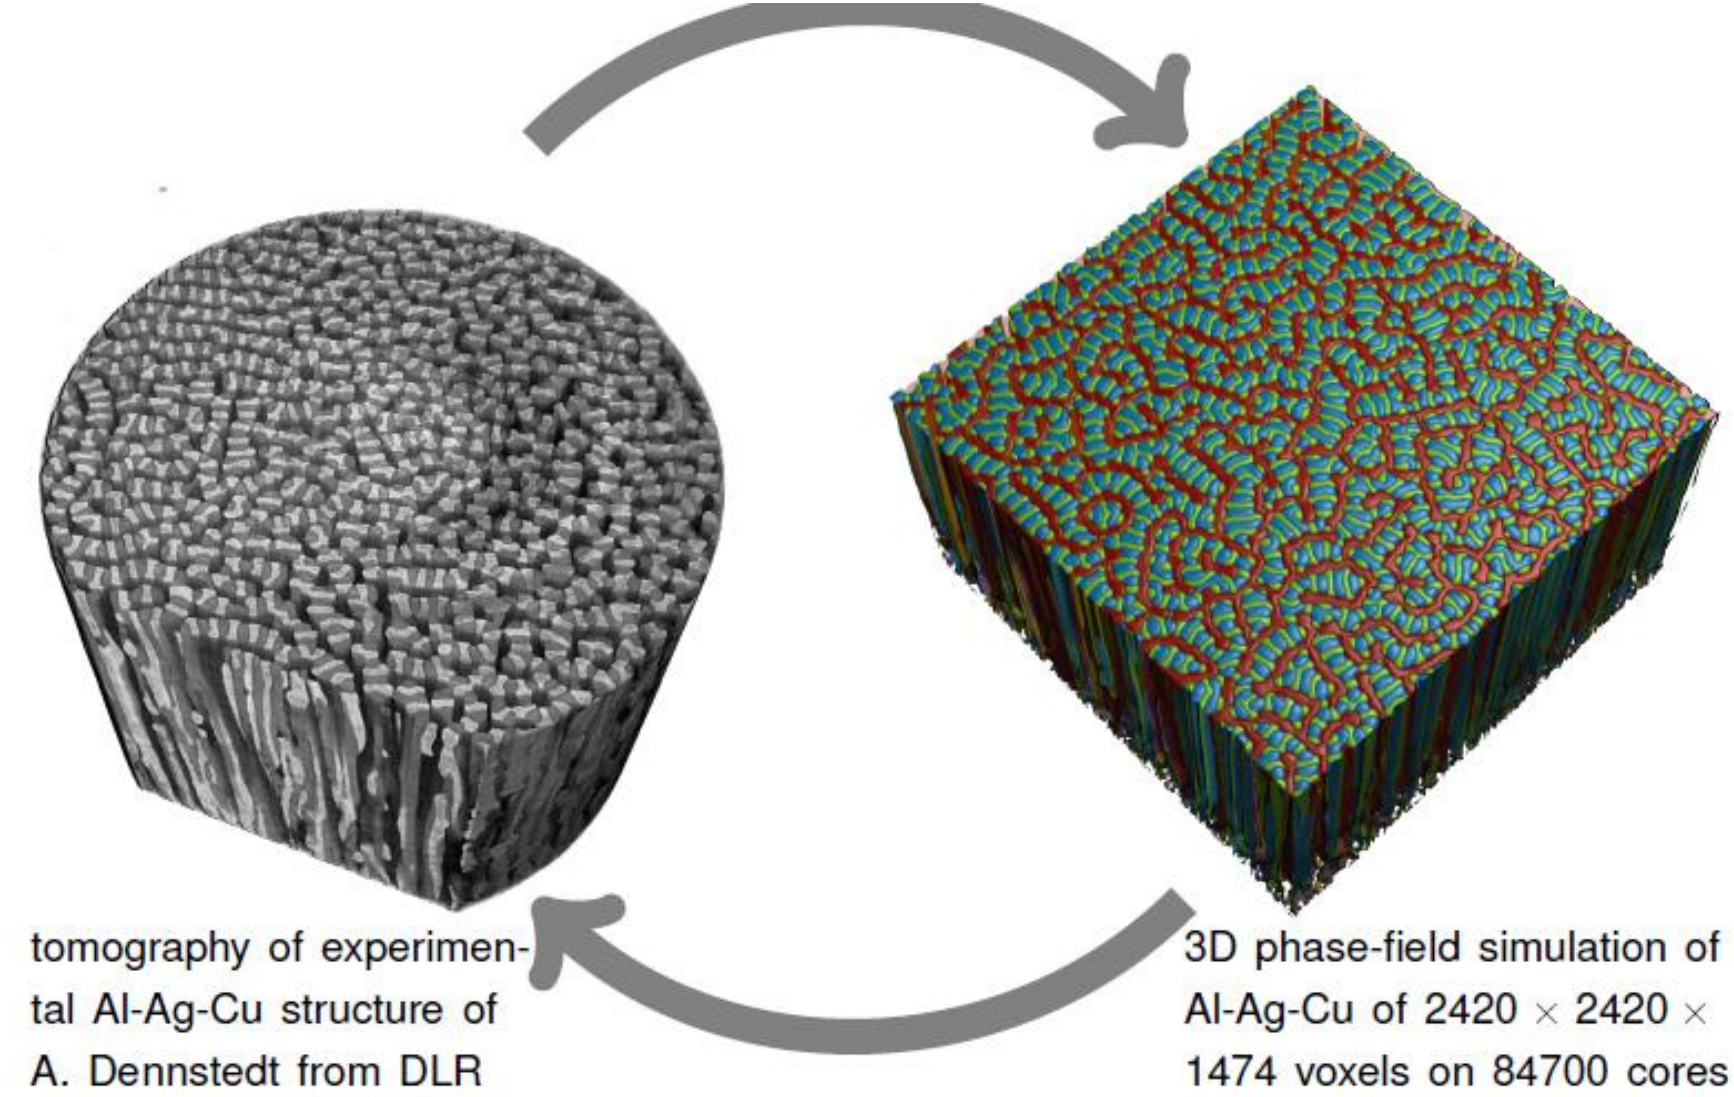

tomography of experimental Al-Ag-Cu structure of A. Dennstedt from DLR

3D phase-field simulation of Al-Ag-Cu of 2420 × 2420 × 1474 voxels on 84700 cores

Pattern features in a Al-Ag-Cu

- Experiment

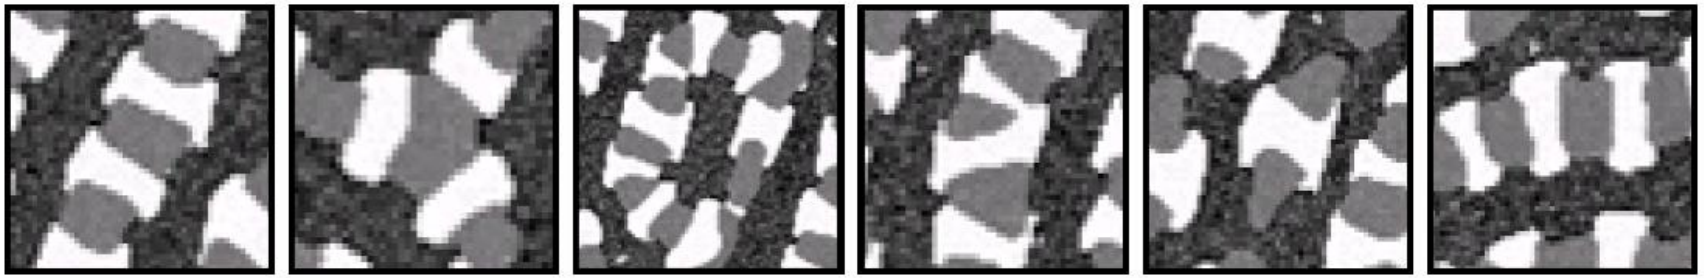

- Simulation

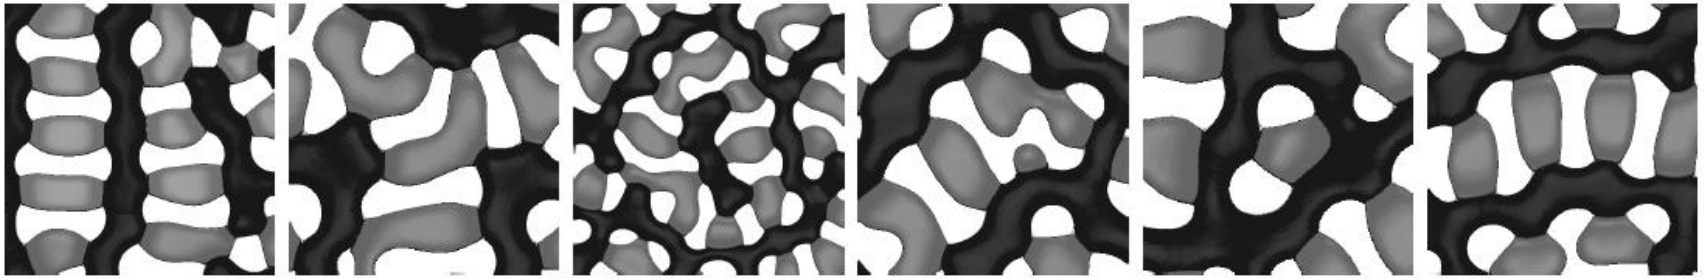

Hötzer, Jainta, Steinmetz, Nestler, Dennstedt, Genau, Bauer, Köstler, and Rüde: Acta Materialia, (2015)

Spiral growth in ternary systems

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Motivation

- Assumption, 2D tilt leads to spirals in 3D

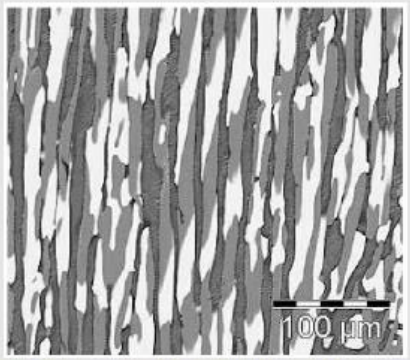

Genau, A., & Ratke, L. (2012). IJMR, 103(4), 469-475.

System parameters

- $940 \times 940 \times 2080$ cells
- system with largest 2D tilt angle

Result

- multiple spirals
- spirals part of chain
- large domains are necessary

Hötzer, Steinmetz, Jainta, Schulz, Kellner, Nestler, Genau, Dennstedt, Bauer, Köstler, and Rüde: Acta Materialia, (2015) (submitted)

SKAMPY Project

Performance Engineering

Karlsruhe Institute of Technology

Hochschule Karlsruhe
Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Work packages

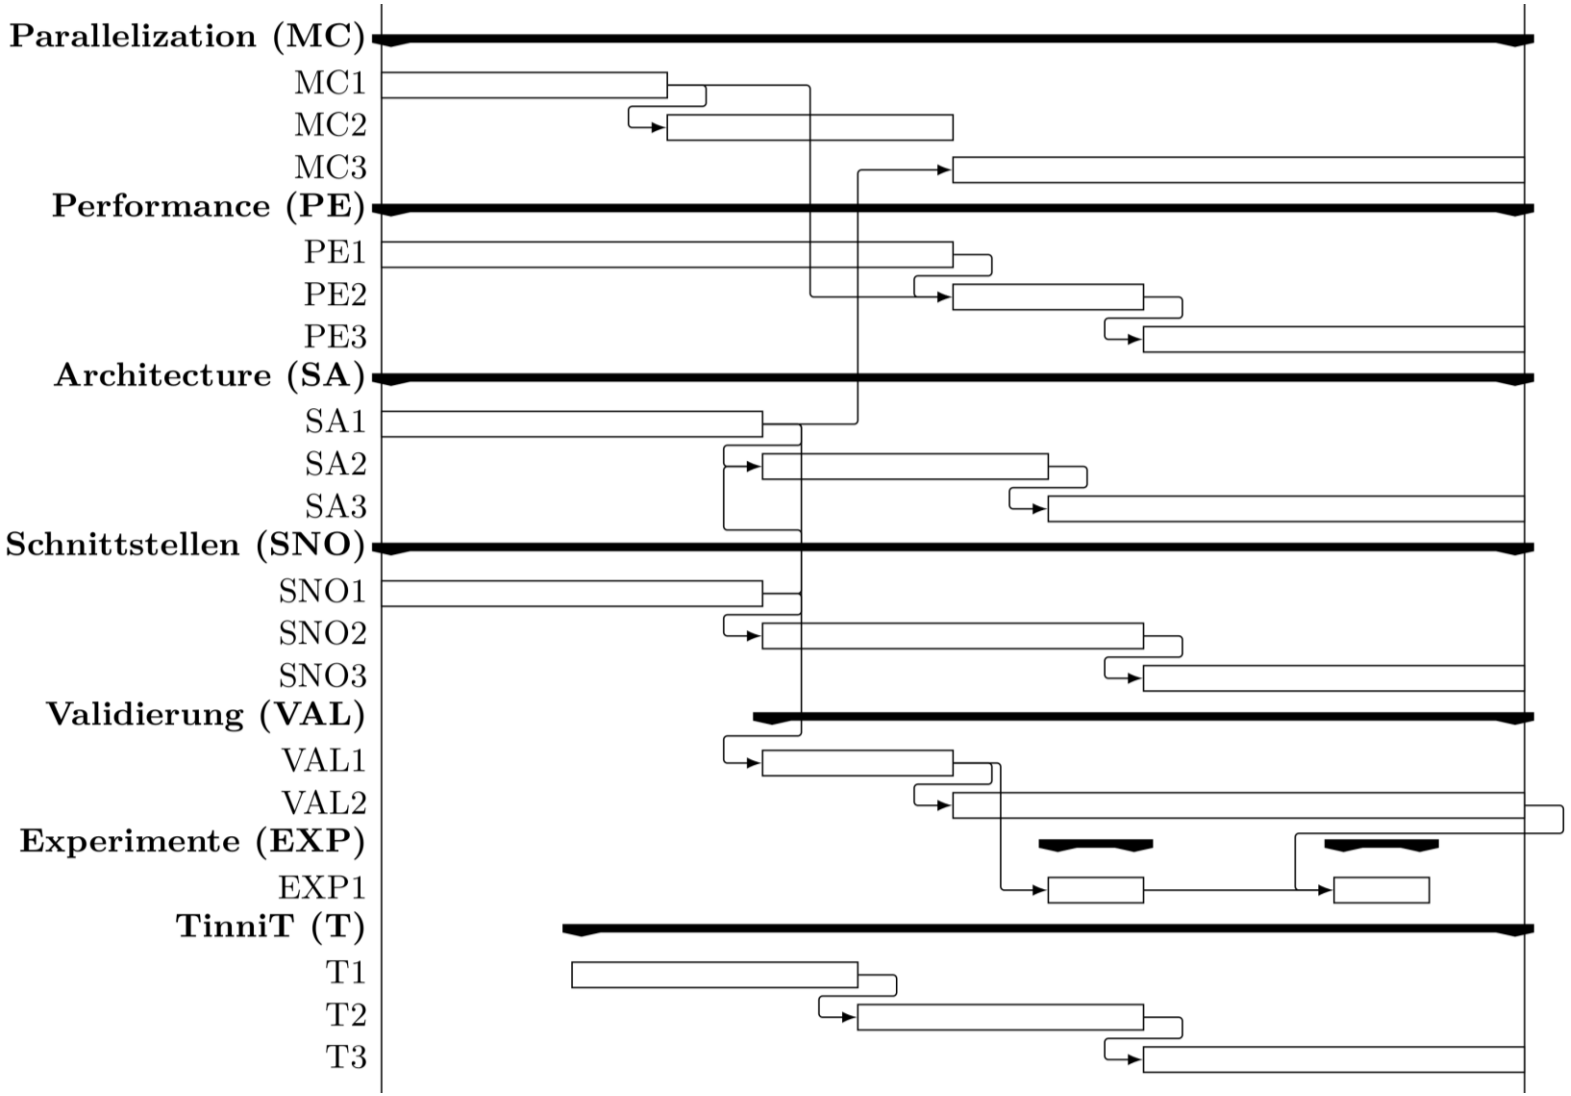

Single Node Tuning

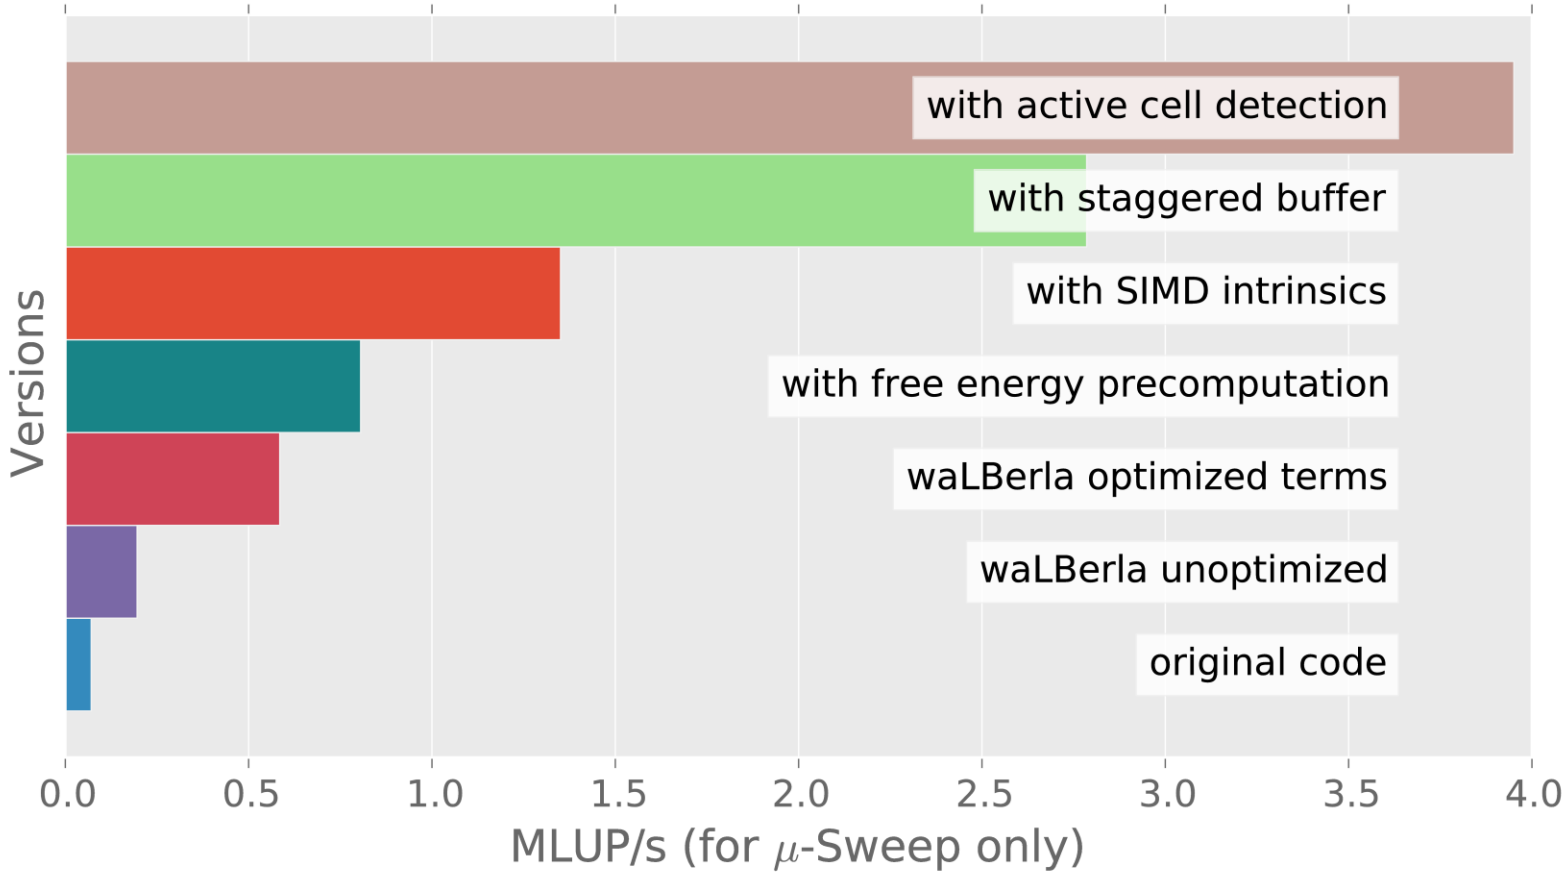

80 x faster compared to original version

Intranode Scaling

intranode weak scaling on SuperMUC

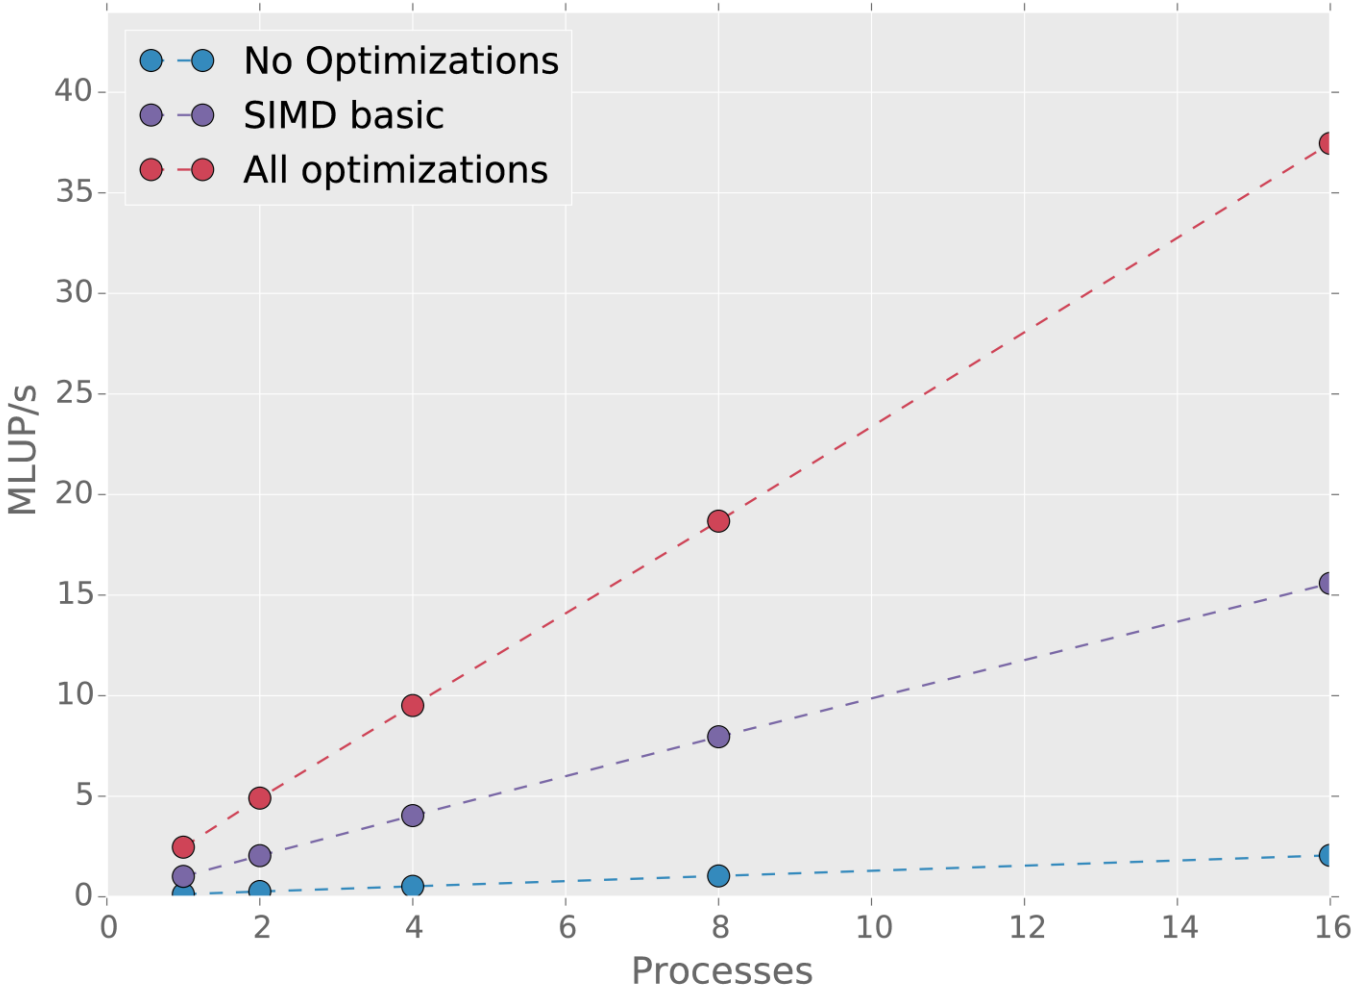

Single Node Optimization Summary

Single Node Optimizations

- replace/remove expensive operations like square roots and divisions
- pre-compute and buffer values where possible
- SIMD intrinsics

Percent Peak on SuperMUC

ϕ-Sweep	21 %
μ-Sweep	27 %
Complete Program	25%

Why not 100% Peak?

- unbalanced number of multiplications and addition
- divisions counted as 1 FLOP but they cost 43 times as much as a multiplication or addition

Scaling

- scaling on SuperMUC up to 32,768 cores

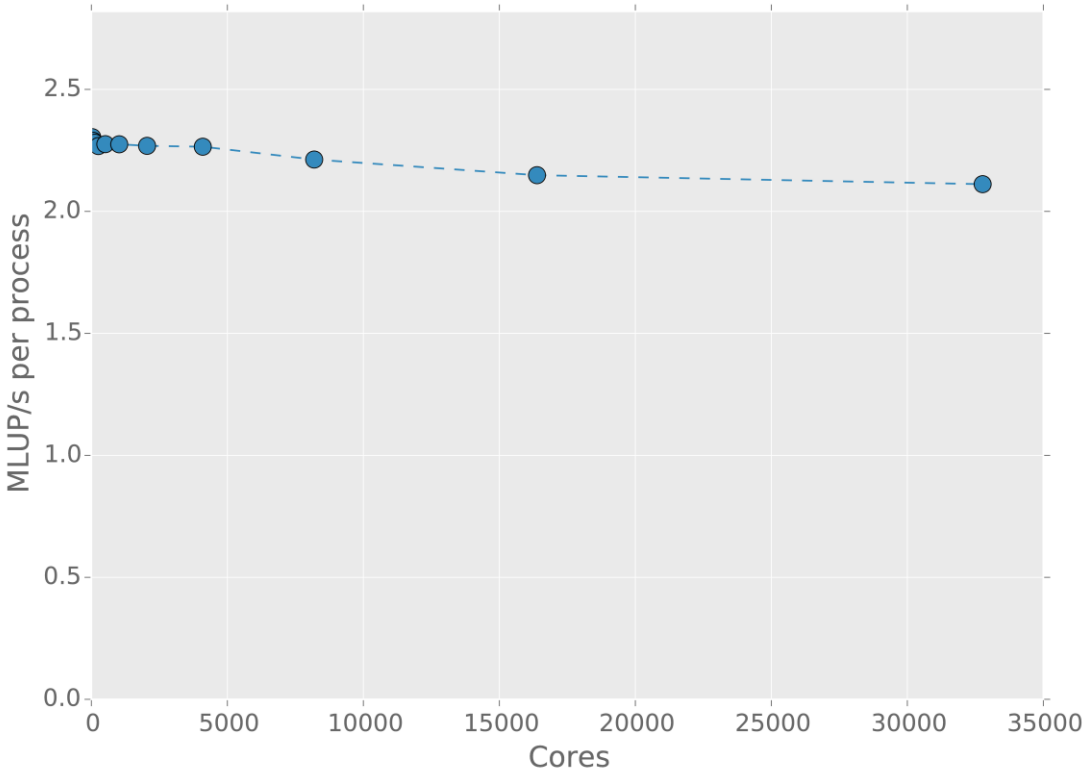

- ghost layer based communication
- communication hiding

Execution-Cache-Memory Model

Execution-Cache-Memory Model

- Automatic Layer Conditions Model

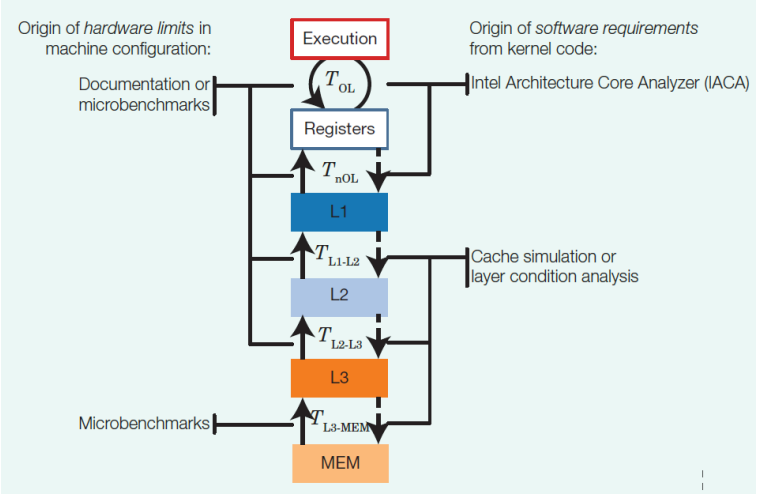

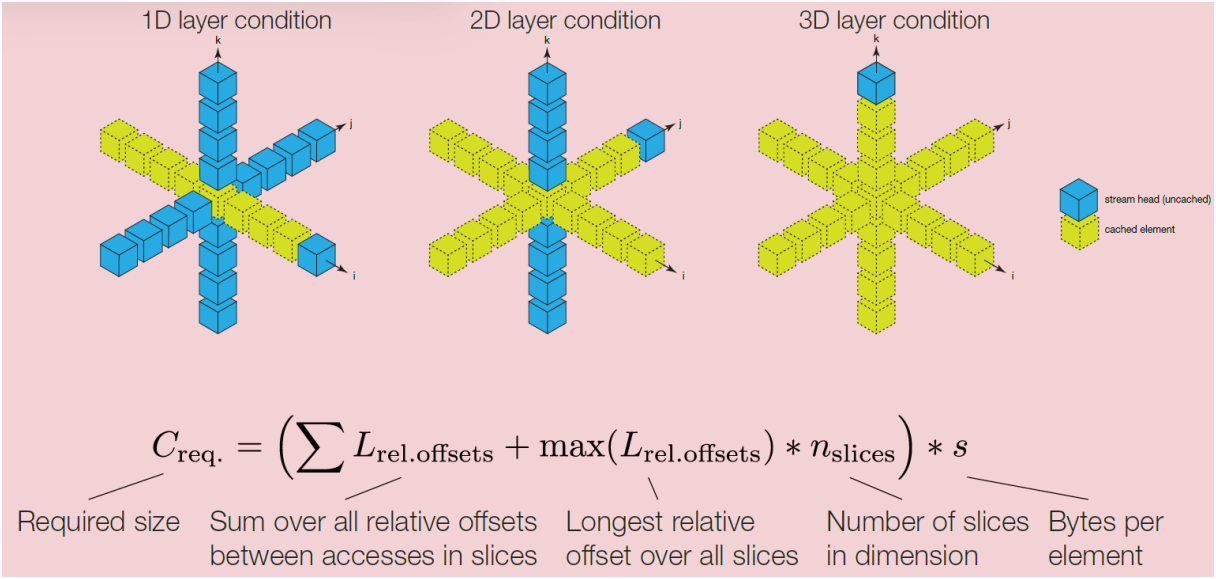

SKAMPY Project

Software Engineering and Parallelization

Karlsruhe Institute of Technology

Hochschule Karlsruhe
Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

Continuous Integration

- Written in C++ with Python extensions
- Hybridly parallelized (MPI + OpenMP)
- No data structures growing with number of processes involved
- Scales from laptop to recent petascale machines
- Parallel I/O
- Portable (Compiler/OS)
- Automated tests / CI servers
- Open Source release planned

llvm/clang

Microsoft
Visual Studio 2012

Continuous Integration

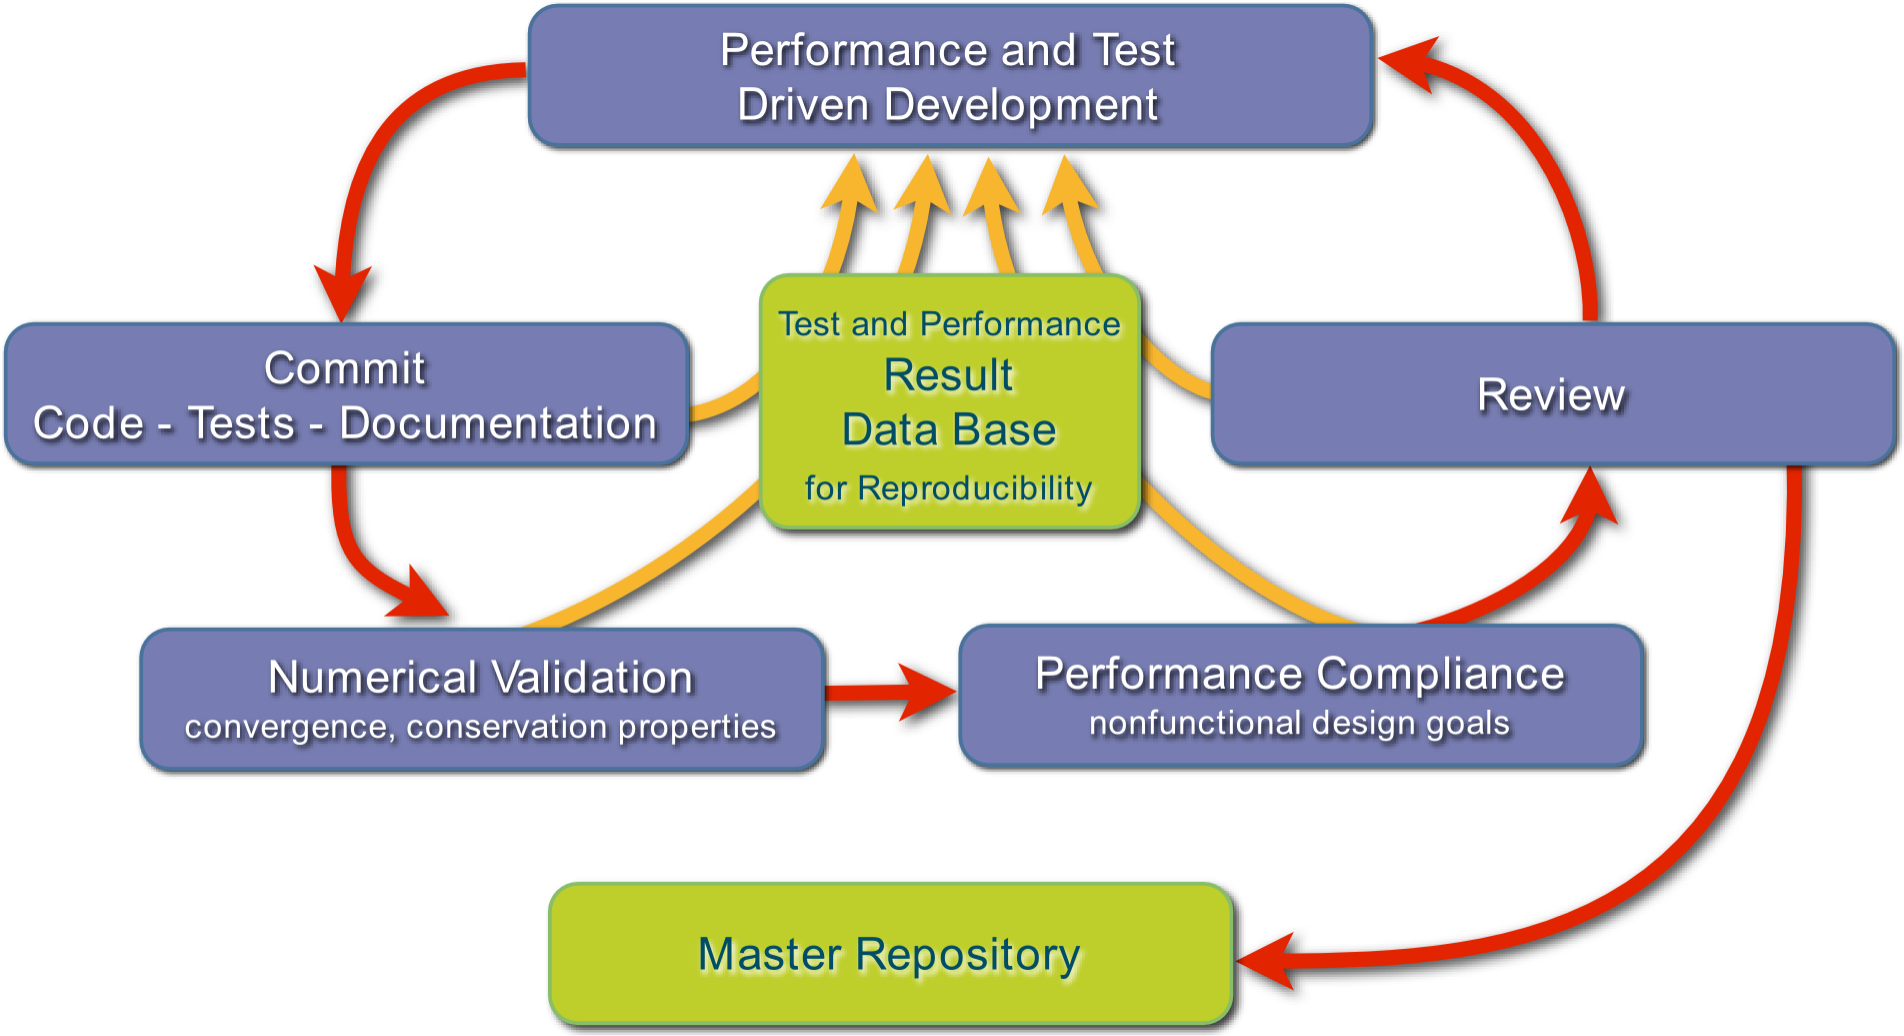

Outlook: Selected Work packages I

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

SA2: Steering & Prototyping Interface

- based on Python

SA3: Adaptivity and Dynamic Load Balancing

Outlook: Selected Work packages II

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT

MC1: Many-Core Data Structures

MC2: Communication Optimization and Multi-scale Data Reduction

- Communication strategies for heterogeous systems
- Problem specific data compression

MC3: Asynchronous Execution and Resilience

Acknowledgements

- Funded by
  - Bundesministerium für Bildung und Forschung

Bundesministerium für Bildung und Forschung

des Bundesministeriums für Bildung und Forschung von Richtlinien zur Förderung von Forschungsvorhaben auf dem Gebiet "HPC-Software für skalierbare Parallelrechner" im Rahmen des Förderprogramms "IKT 2020 - Forschung für Innovationen"

  - KONWIHR. Bavarian project

KONWIHR II

  - DFG SPP 1648/1 – Software for Exascale computing

DFG Deutsche Forschungsgemeinschaft

ExaStencils

http://www.exastencils.org/

SPPEXA

  - Industry

P&G

NVIDIA CUDA RESEARCH CENTER

abberior INSTRUMENTS

SIEMENS

  - Supercomputing centers

HLRS

lrz

JÜLICH FORSCHUNGSZENTRUM

Thank you!

Questions?

Karlsruhe Institute of Technology

Hochschule Karlsruhe
Technik und Wirtschaft
UNIVERSITY OF APPLIED SCIENCES

FRIEDRICH-ALEXANDER
UNIVERSITÄT
ERLANGEN-NÜRNBERG
TECHNISCHE FAKULTÄT