ERLANGEN REGIONAL COMPUTING CENTER

Components for practical performance engineering in a computing center environment: The ProPE project

Jan Eitzinger

Workshop on Performance Engineering for HPC: Implementation, Processes, and Case Studies at ISC 2017

Overview

DFG Deutsche Forschungsgemeinschaft

Call:

Performance Engineering für wissenschaftliche Software

Partners:

RWTH AACHEN UNIVERSITY

TECHNISCHE UNIVERSITÄT DRESDEN

Duration:

03/2017 – 02/2020

Coordination: Prof. G. Wellein

Current state

- HPC competence in German HPC centers distributed across country
- Gauss-Allianz is an initiative to integrate and organize TIER 2/3 HPC landscape in Germany
- Multiple local efforts and island projects:
  bwHPC, KONWIHR, HKHLR, HLRN …

Our contribution

- Similar targets as sketched in **GA Strategiepapier,** but focus on Performance-Engineering sub-topic

Integrate with and **built on** already existing efforts and further drive the final goal of an hierarchical and yet integrated German HPC infrastructure.

Major Building Blocks

- **Dissemination** – Increase publicity of project and raise general awareness for performance issues

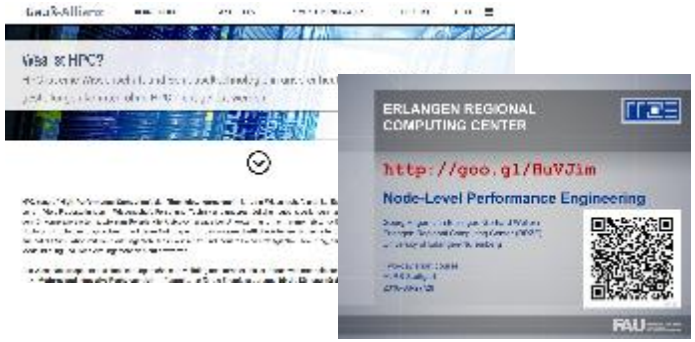

- **Documentation –** Build a central web offering, create content and provide resources to maintain it

- **Structured PE-Process** – Systematic bottleneck centric performance analysis and optimization process

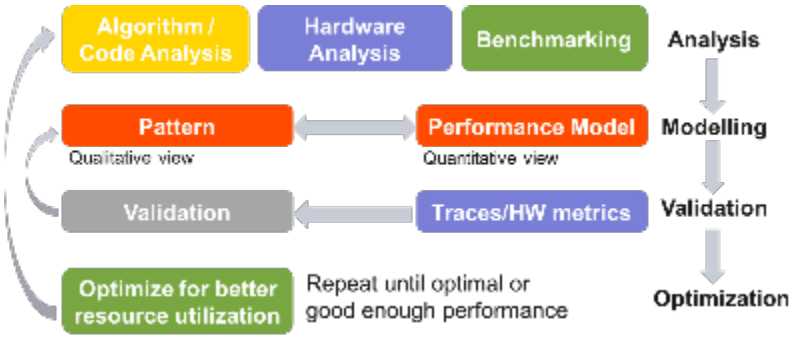

Major Building Blocks cont.

- **PE Support Infrastructure** – Process blueprint for nation-wide aligned support effort

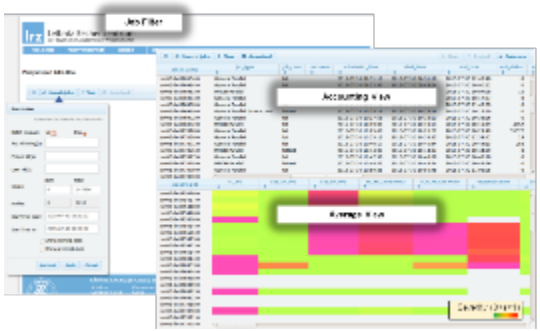

Bundesministerium für Bildung und Forschung

FEPA

- **Application Monitoring and Analysis** – Automatic profiling and bottleneck analysis for all applications running on a HPC-System

- **HPC Curriculum –** Coordinated nation-wide Workshop and Tutorial program

PE Support Infrastructure

- Multi-Tier distributed **support infrastructure** which allows to hand-over requests and allocate specialists from other centers

- Create a process for ***Performance Projects*** allowing to
  - **Keep track** of and **transfer** projects between sites and find the right expert for a specific problem
  - **Carry out** and **document** efforts and results in a **standardized coeherent way**
  - **Pack** an already started **project** between sites so that experts can pick it up right away

Application Performance Monitoring

Global automatic **application performance monitoring** is essential to improve **efficient** usage of HPC systems

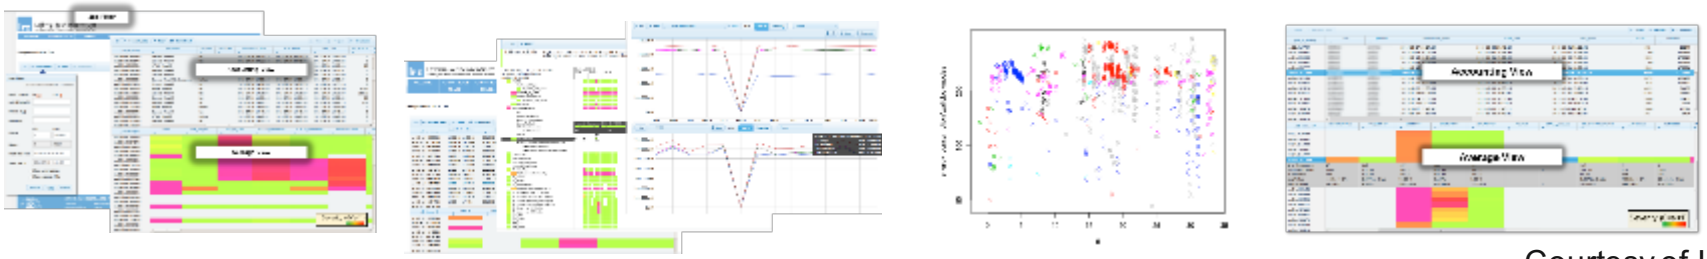

Courtesy of LRZ

Targets:

- Give user immediate feedback on job runs
- Identify applications with high optimization potential or pathological performance behavior
- Create databases with ***performance footprints*** and ***performance maps*** to characterize applications and track HPC usage statistics

Performance Engineering Tasks: Software side

Optimizing software for a specific hardware requires to align several orthogonal targets.

Software side: Reduce algorithmic and processor **work**

Performance Engineering Tasks: Hardware

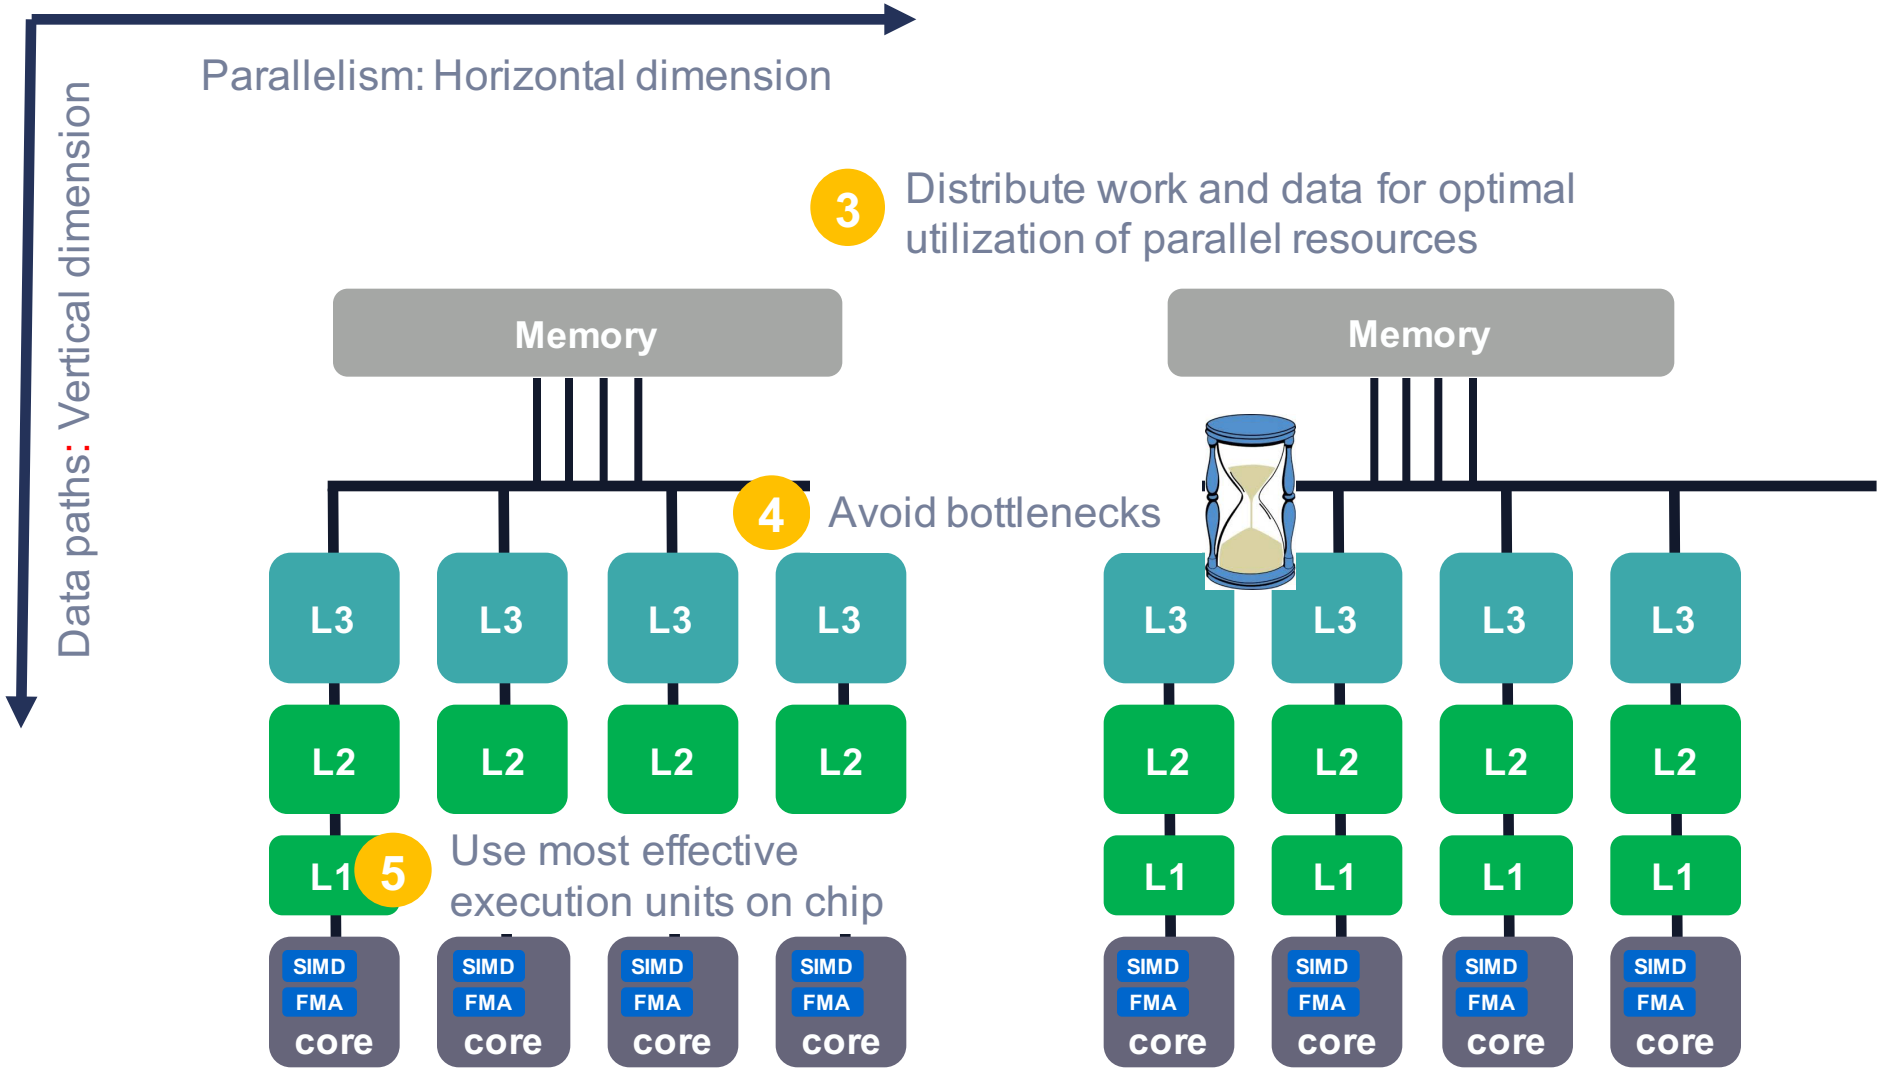

History of Intel hardware developments

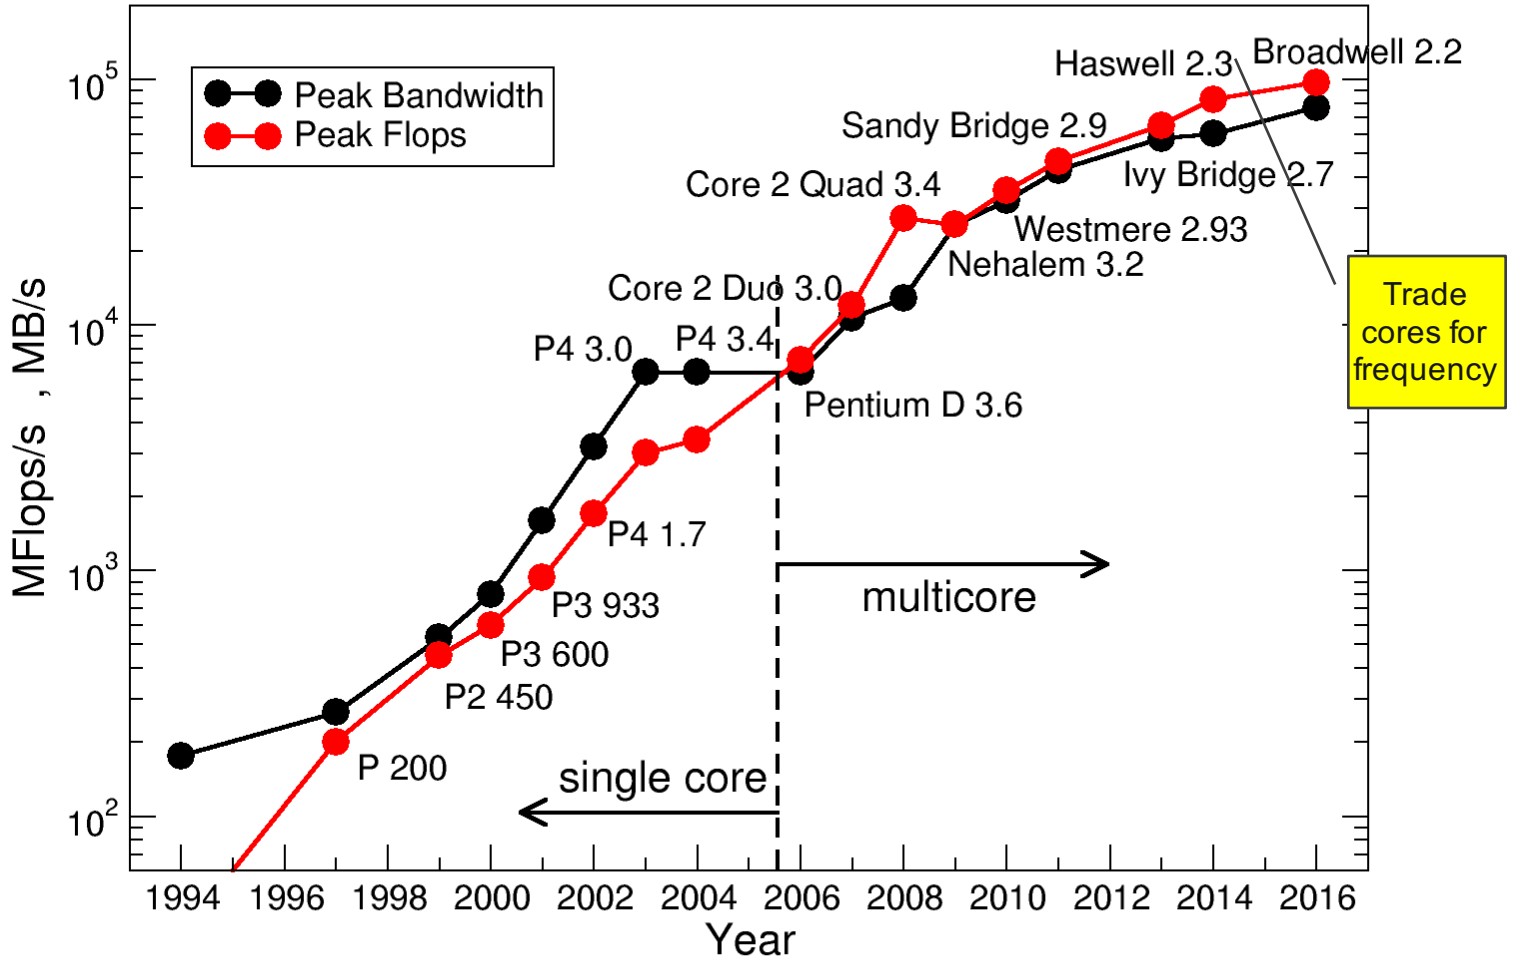

The real picture

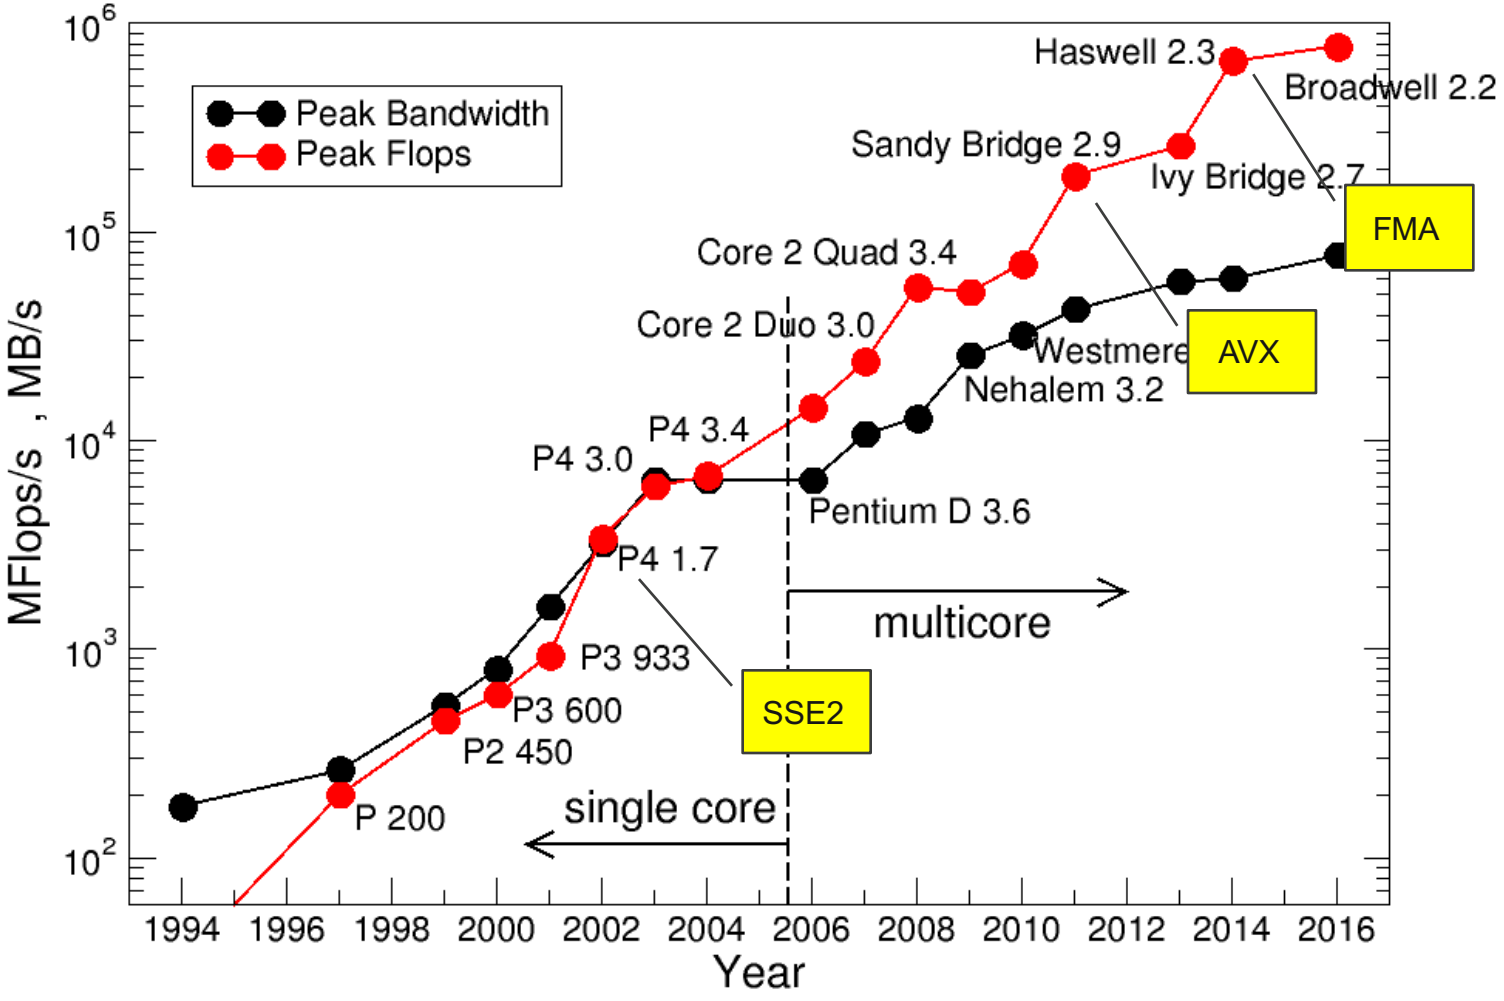

Technologies Driving Performance

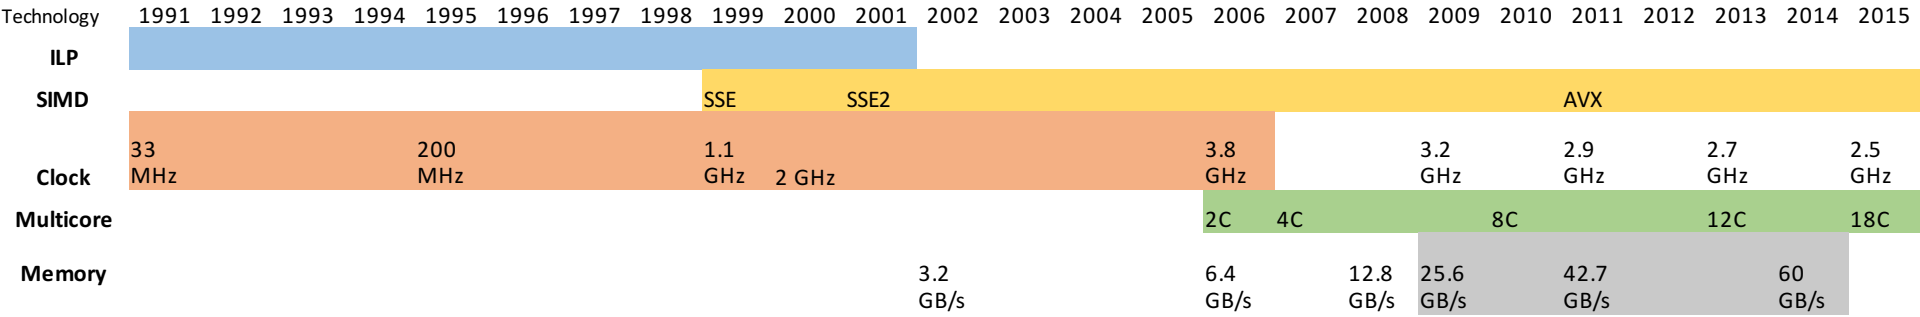

- ILP **Obstacle**: Not more parallelism available
- Clock **Obstacle**: Heat dissipation
- Multi- Manycore **Obstacle**: Getting data to/from cores

Flavors of improvements

- Pure speed increase: Clock
- **Transparent** solutions
- **Explicit** solutions

Common strategies

- Parallelism
- Specialisation

What will the future bring?

- **System configuration** (frequency, placement strategies, COD, prefetching settings, cache coherence settings, memory configuration)

- **Specialization**
  - Special purpose instructions
  - Special purpose memories
  - Special purpose cores
  - Configurable execution units

Main questions

- Will this work transparently?
- Who will be responsible?
  - Tool chain
  - Developer
  - User
  - System provider

ERLANGEN REGIONAL COMPUTING CENTER

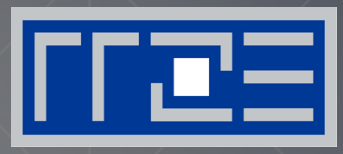

Thank You.