GA Gauß-Allianz

Manual and Automatic Energy Tuning for HPC Codes

The Projects Score-E and READEX

European Commission | Horizon 2020 European Union funding for Research & Innovation

SPONSORED BY THE

Federal Ministry of Education and Research

Kai Diethelm

GNS Gesellschaft für numerische Simulation mbH

Michael Gerndt

TU München

Outline

- Motivation
- Score-E and READEX
- Energy Monitoring for HPC applications
- Application analysis
- Automatic energy tuning
- Applications

Exascale Energy Wall

#	Site	System	Cores (M)	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (MW)	Exascale (Factor)	Exascale (MW)
1	National Supercomputing Center, Wuxi, China	Sunway TaihuLight - Sunway 1.45GHz, Sunway, NRCPC	10.6	93.0	125.4	15.4	8	123
2	National Super Computer Center in Guangzhou, China	Tianhe-2 (MilkyWay-2), Intel Xeon E5-2.2GHz, Intel Xeon Phi 31S1P, NUDT	3.1	33.9	54.9	17.8	18	324
3	DOE/SC/Oak Ridge National Laboratory, United States	Titan - Cray XK7 , Opteron 2.2GHz, NVIDIA K20x, Cray Inc.	0.5	17.6	27.1	8,2	37	303
4	DOE/NNSA/LLNL, United States	Sequoia - BlueGene/Q, Power BQC 1.60 GHz, IBM	1.6	17.2	20.1	7,9	50	392
5	DOE/SC/LBNL/ NERSC, United States	Cori - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Cray Inc.	0.6	14.0	27.9	3,9	36	141

GA
Gauß-Allianz

Score-E

SPONSORED BY THE

Federal Ministry of Education and Research

- Partners:
  - GNS Gesellschaft für numerische Simulation mbH, Braunschweig (Coordinator)
  - RWTH Aachen
  - German Research School for Simulation Sciences, Aachen (until February 2015), TU Darmstadt (from March 2015)
  - TU Dresden
  - FZ Jülich
  - TU München
- Associated Partners:
  - Engys GmbH (Rostock), Munters Euroform GmbH (Aachen),
  - U of Oregon
- Funded by BMBF (Grant No. 01IH13001)
- October 2013 – September 2016

GA

Gauß-Allianz

Semi-automatic and Automatic Energy Tuning

READEX

Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

Funded by the European Union’s Horizon 2020 research and innovation programme under Grant Agreement No. 671657.

European Commission

Horizon 2020
European Union funding
for Research & Innovation

GA
Gauß-Allianz

READEX Project partners

- TU Dresden (Coordinator), Germany
- NTNU, Norway
- IT4I, Czech Republic
- TUM, Germany
- Intel Exascale Lab, France
- GNS Braunschweig, Germany
- ICHEC, Ireland

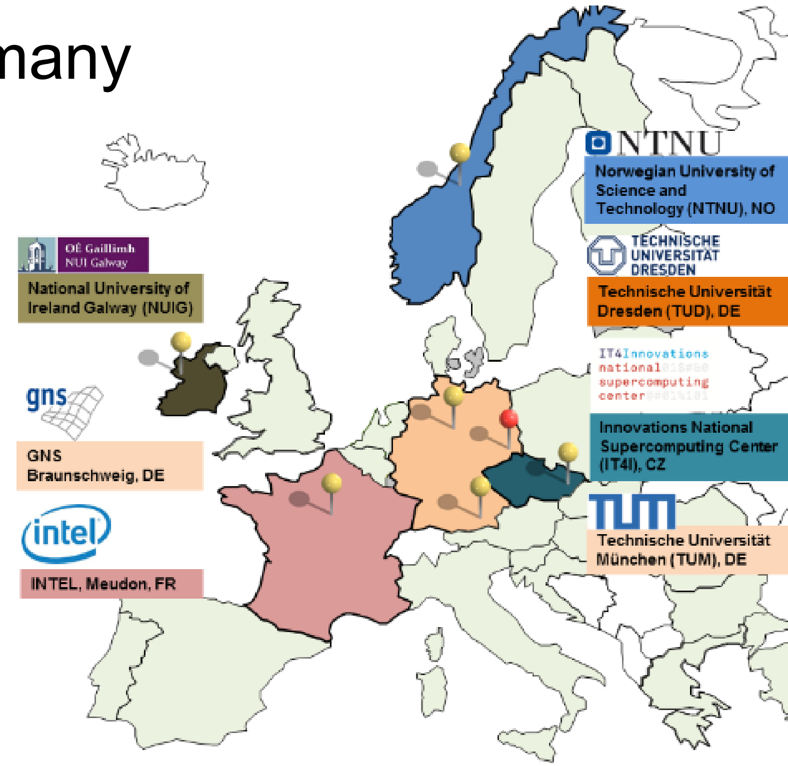

GA Gauß-Allianz

Score-E: Energy Monitoring with Score-P

Metric Plugins:

- Energy: Intel RAPL, AMD APM, HDEEM
- Temperature
- C-/P-states

Synchronization Adapter

- Shared resources

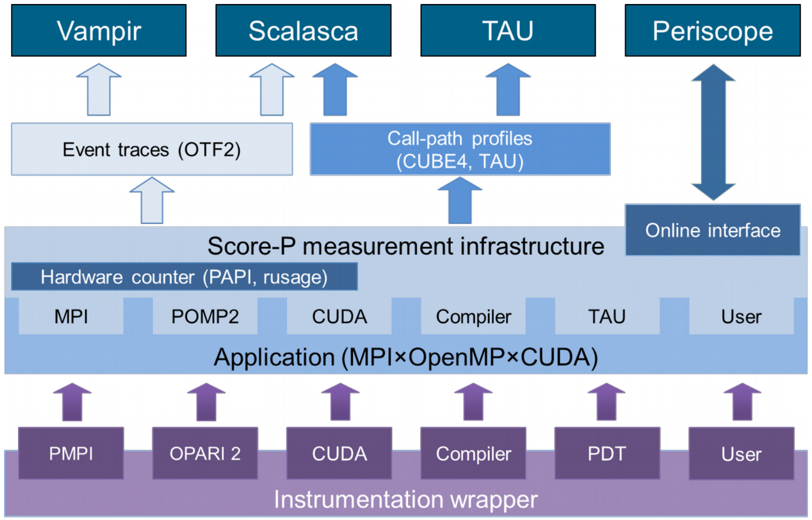

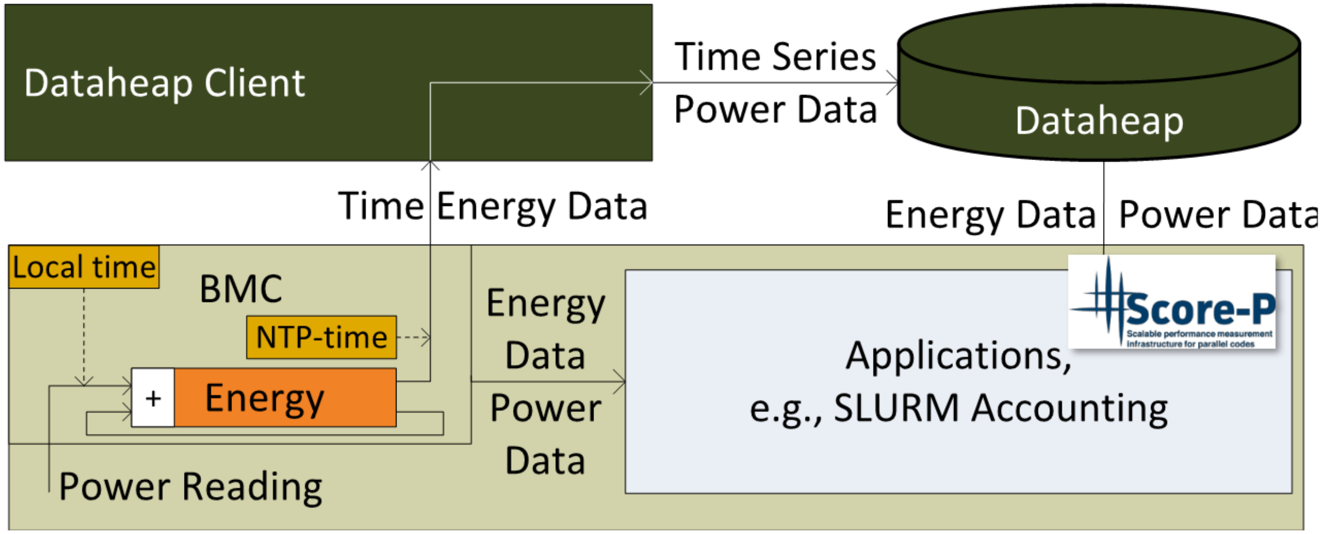

Manual Analysis of Energy Usage with Vampir

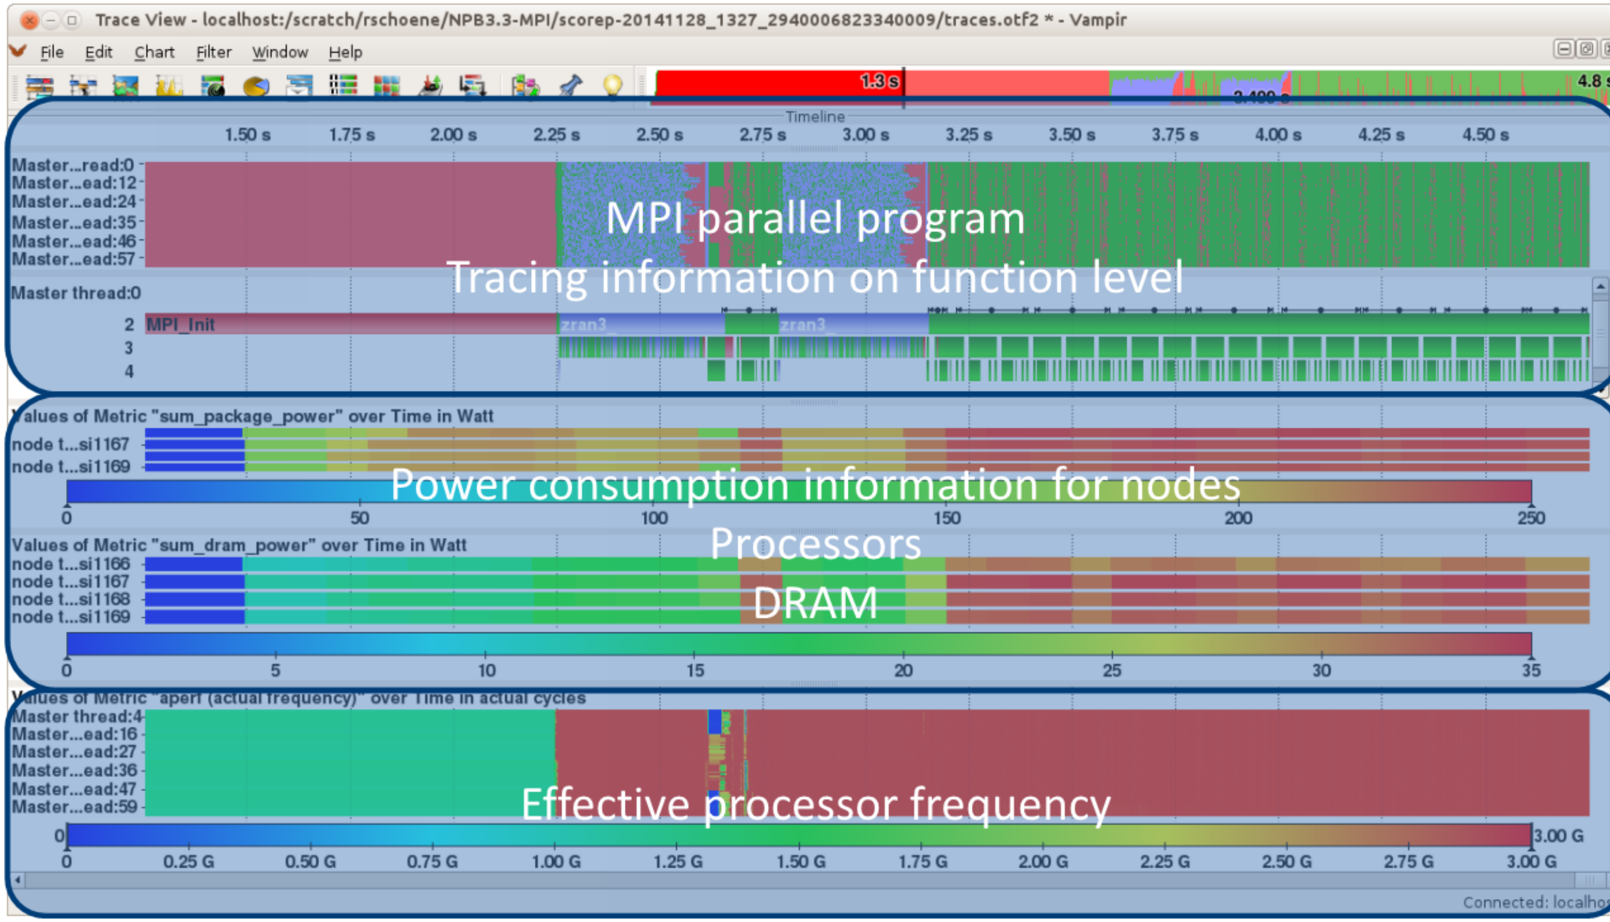

NAS Parallel Benchmarks (MG)

GA Gauß-Allianz

Manual Analysis with Cube on JURECA

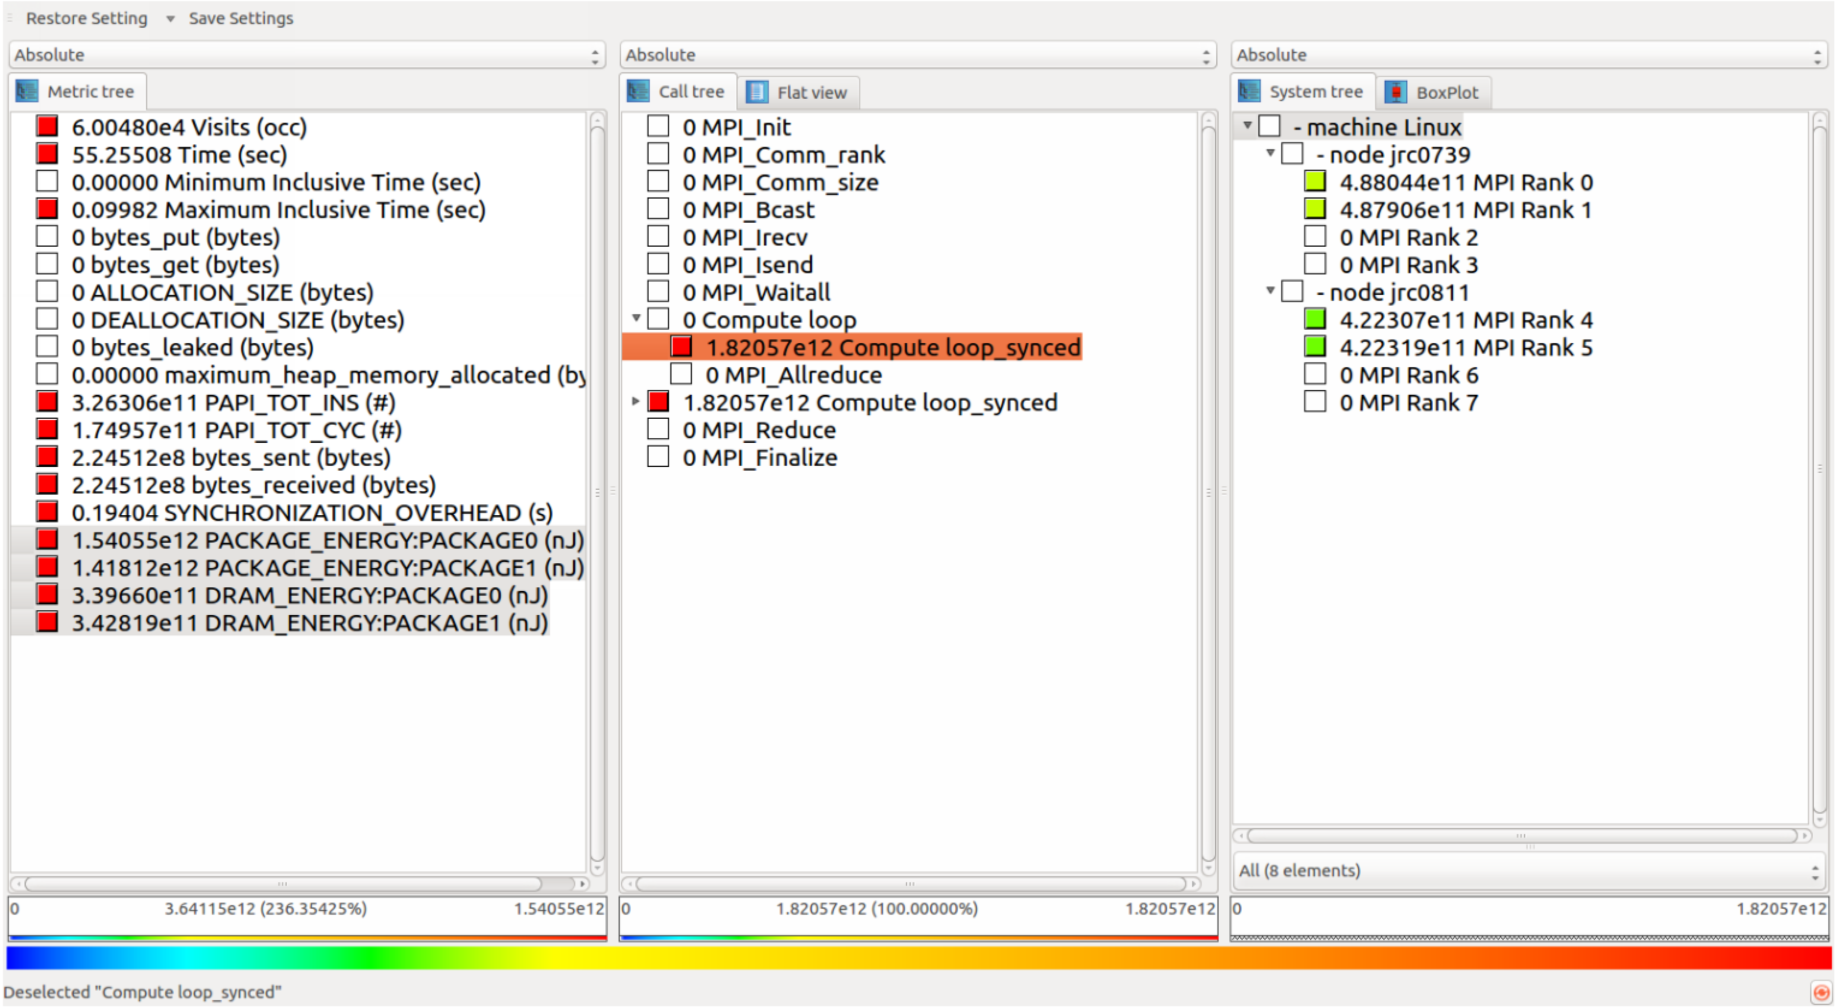

GA Gauß-Allianz

Energy Consumption Projection

Goal: Prediction of energy consumption on larger systems

- Based on Extra-P (GRS, TU Darmstadt, ETH Zürich, LLNL, FZ Jülich)
- Analytical model deduced from application profiles
- Visualization of predicted energy consumption in Cube

MILC: QCD code

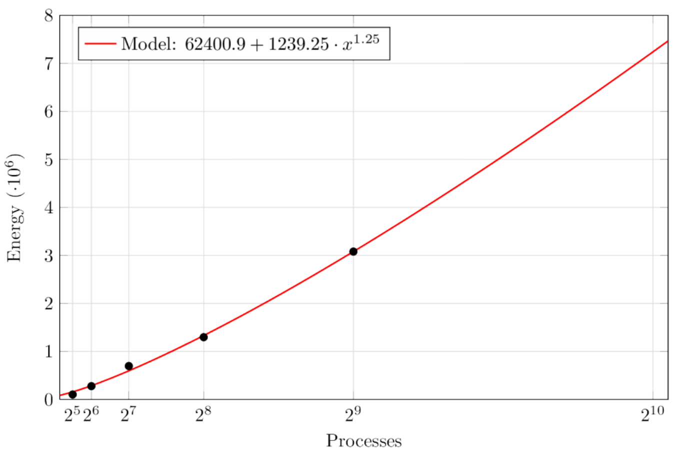

Energy consumption

Resource-Aware Visualization

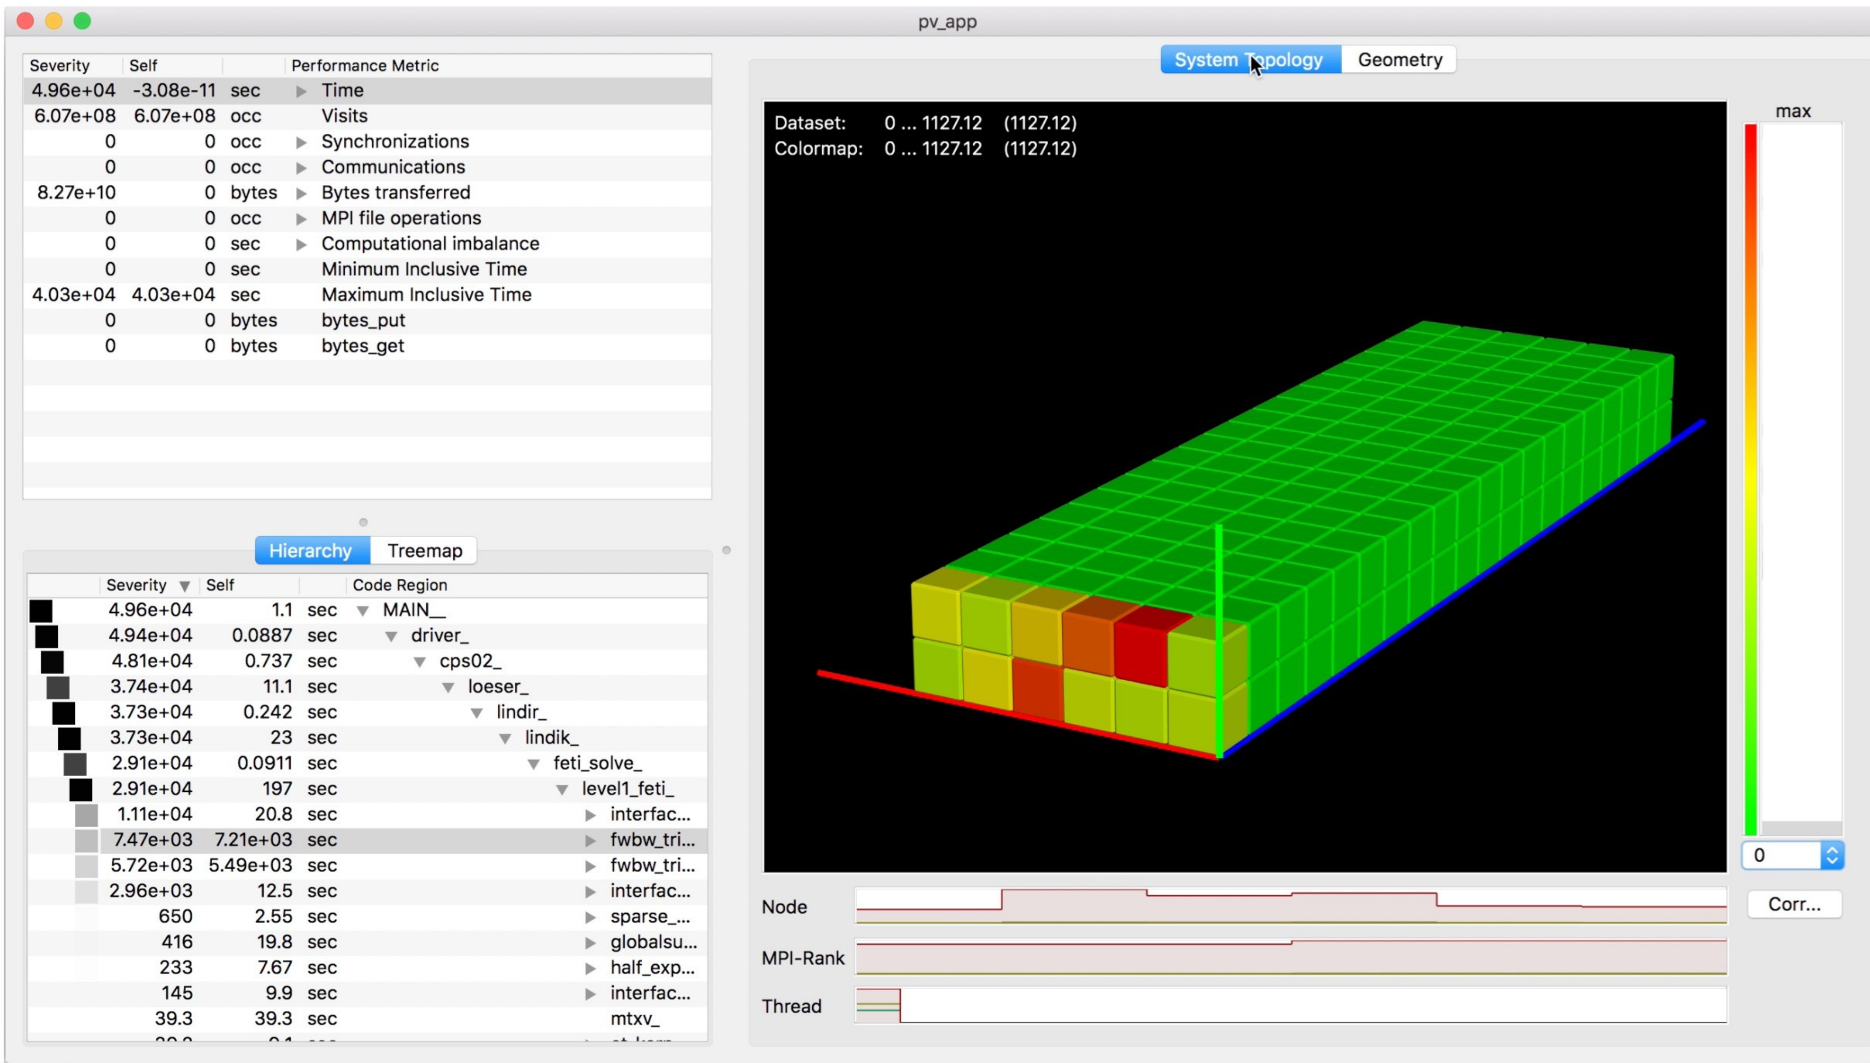

Geometry-Aware Visualization

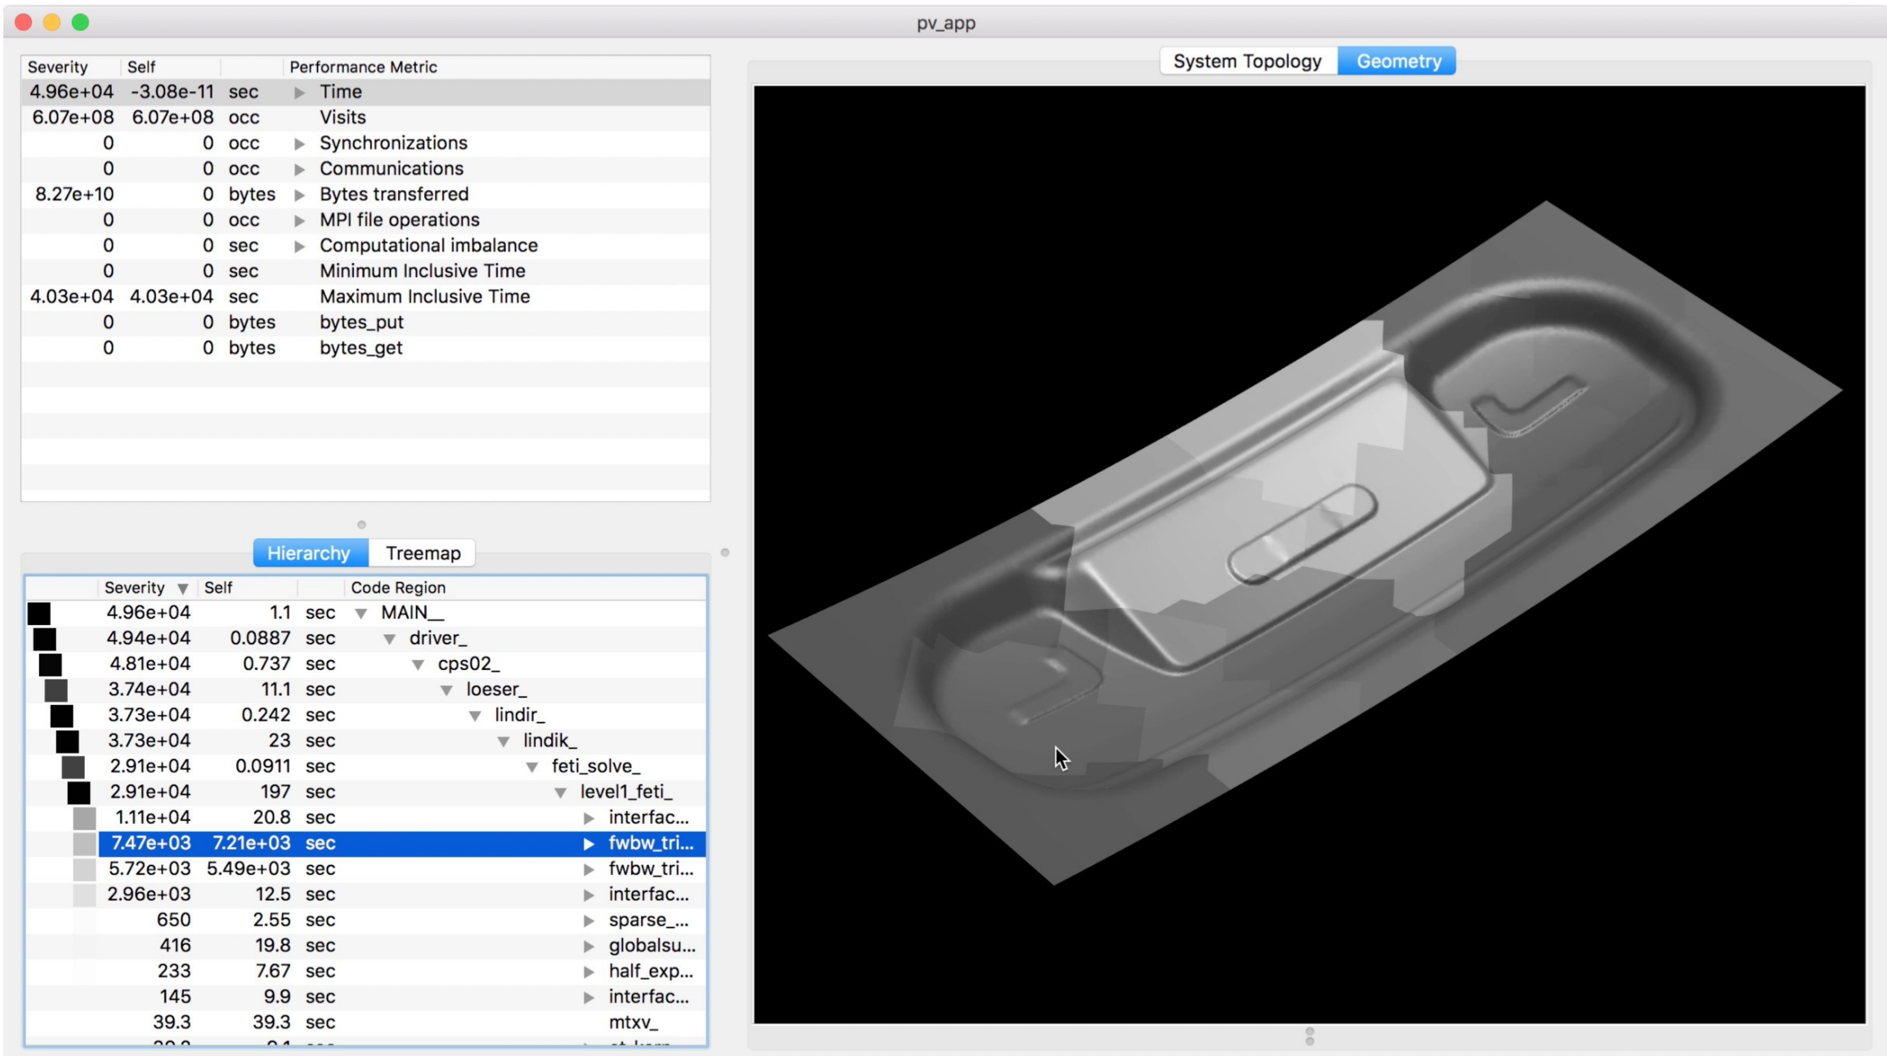

Geometry-Aware Visualization

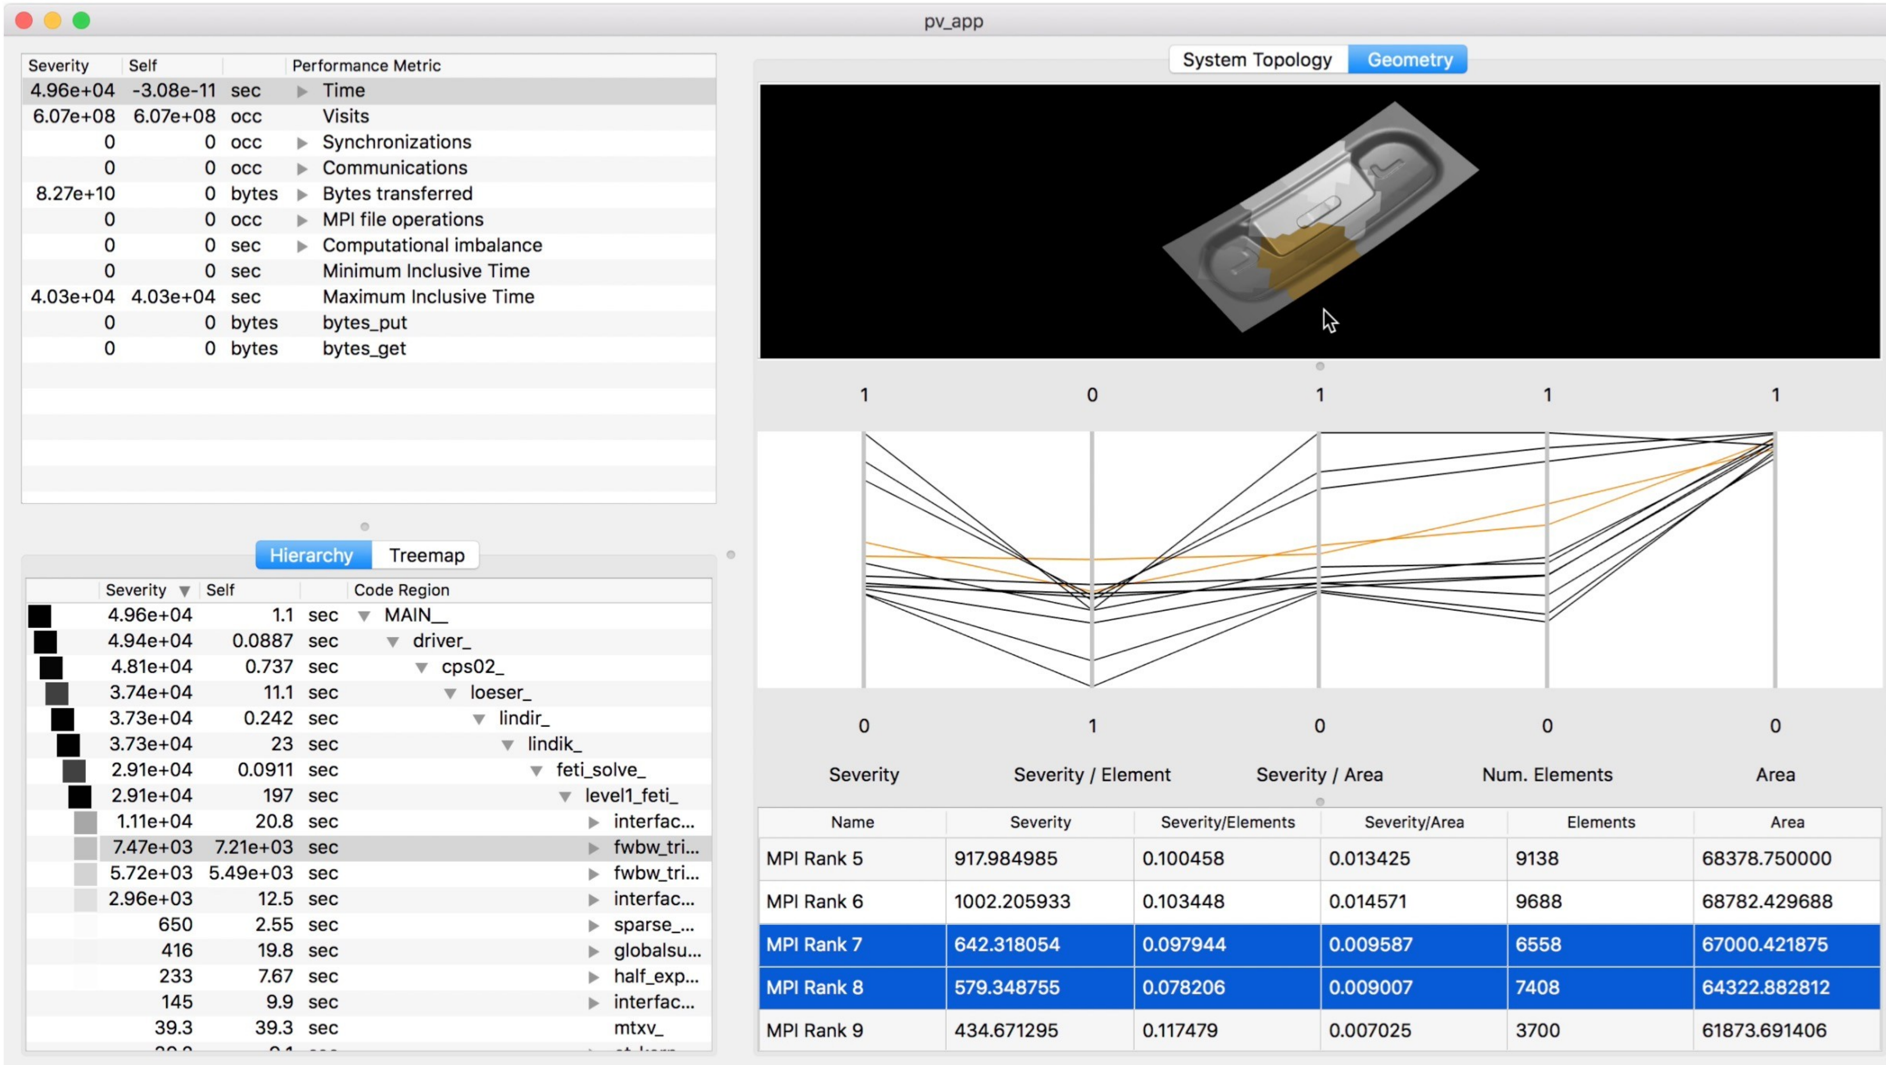

Energy Tuning with Periscope Tuning Framework

Automatic application analysis & tuning

- Tune performance and energy (statically)
- Plug-in-based architecture
- Evaluate alternatives online
- Scalable and distributed framework

Support variety of parallel paradigms

- MPI, OpenMP, OpenCL, Parallel pattern

Developed in the AutoTune EU-FP7 project

- Integrated with Score-P in Score-E

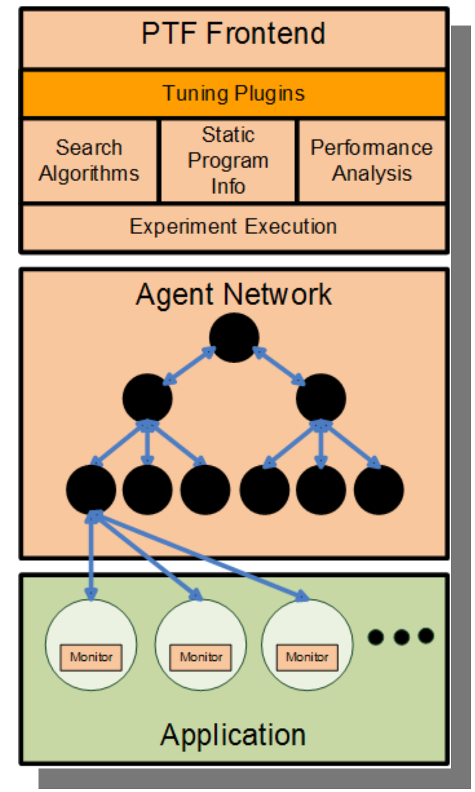

Gauß-Allianz

Tuning Plugins for Static Energy Tuning

- **PCAP**: Optimization of the thread number
  - Exploits scalability limitations
  - Reduces static and dynamic energy
  - Supports tuning of individual parallel
  - regions

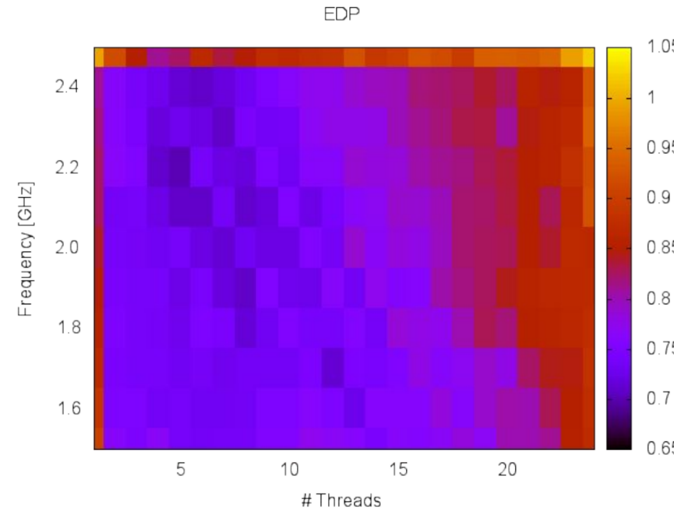

- **MPICAP**: Optimizes the number of MPI tasks
  - In addition: reduces communication energy
- **DVFS**: Computes the optimal core frequency
  - Exploits wait cycles (IO-/Memory Bound)
  - Reduces dynamic energy
- **Taurus-Energy**: Combined tuning (3D)

GA Gauß-Allianz

READEX: Beyond Static Tuning

HPC
- Automatic Tuning

Embedded
- System Scenarios

GA Gauß-Allianz

Systems Scenario based Methodology

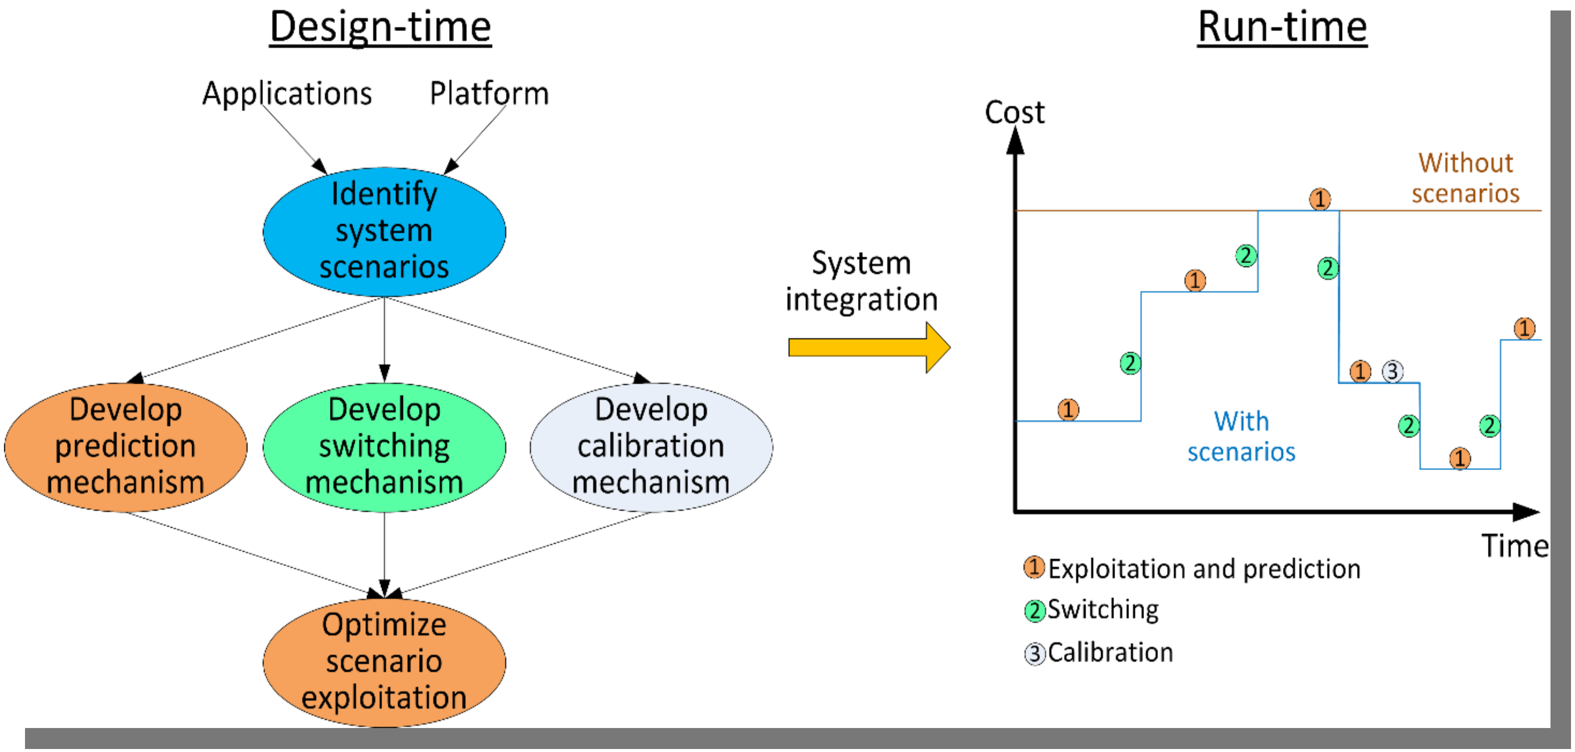

GA Gauß-Allianz

Exploit Intra- and Inter-Phase Dynamism

```
int main(void) {

  // Initialize application
  // Initialize experiment variables

  int num_iterations = 2;

  for (int iter = 1; iter <= num_ite
    // Start phase region
    // Read PhaseCharct
    laplace3D(); // significant regi
    residue = reduction(); // insign
    fftw_execute(); // significant r
    // End phase region
  }

  // Post-processing:
  // Write noise matrices to disk fo
  // Terminate application

  MPI_Finalize();
  return 0;
}
```

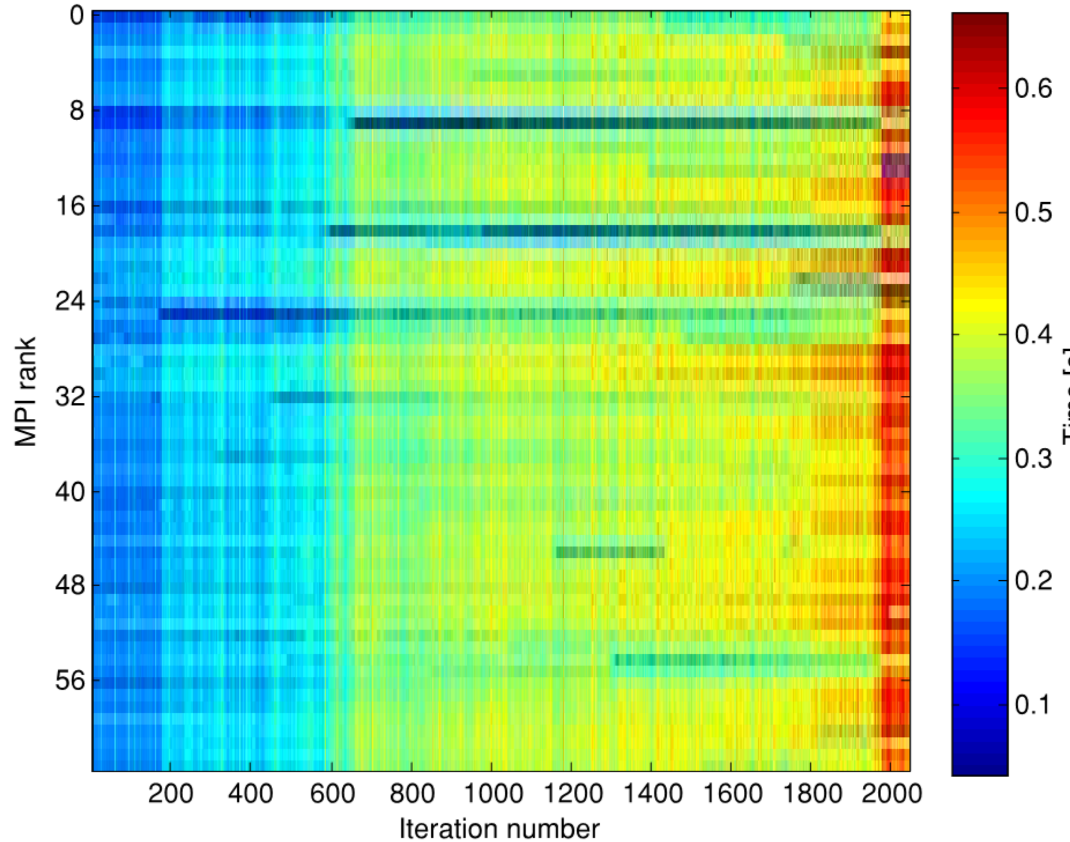

GA Gauß-Allianz

Scenario-Based Tuning

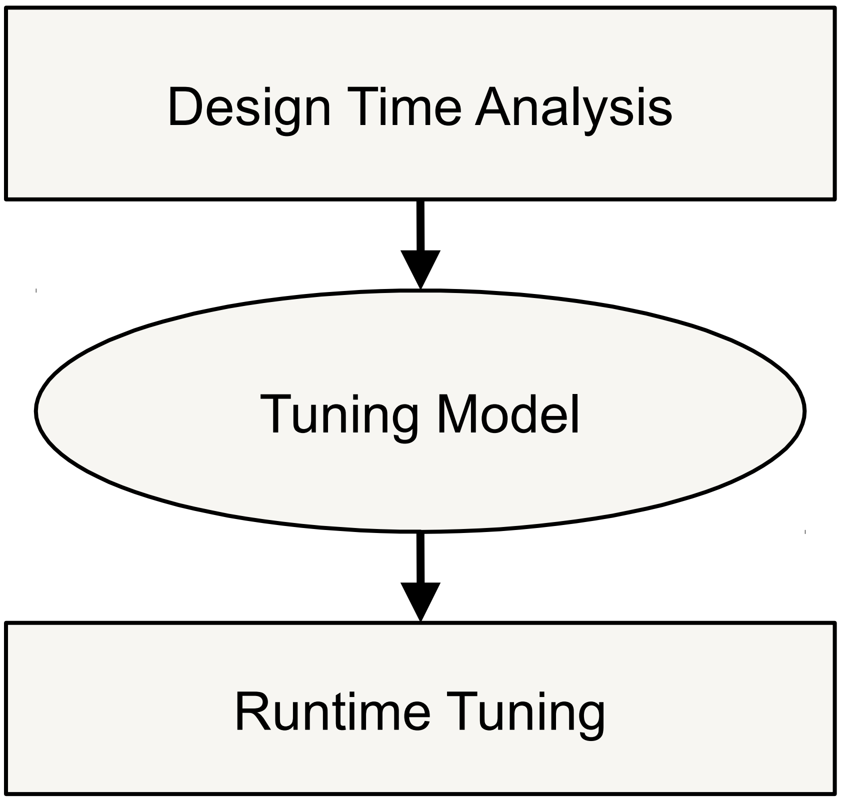

Periscope Tuning Framework (PTF)

READEX Runtime Library (RRL)

Gauß-Allianz

Lulesh: Individual, 3 cores per task

individual, 3 cores/task	Worst	Static Tuning	READEX	#Threads	Freq
	16; 1,2	8; 2,4			
CalcCourantConstraintForE	6075	5375	5232	10	2,4
CalcKinematicsForElems	11170	8260	8250	9	2,4
CalcMonotonicQGradients	4977	3920	3370	6	2,4
CalcMonotonicQRegionFor	11714	12470	11714	16	2,4
CalcVolumeForceForElems	56821	48111	47804	9	2,4
EvalEOSForElems	28335	18311	17082	4	2,4
SUM	**119092**	**96447**	**93452**		
		19,0%	3,1%		
Energy for phase	163891	138508			
		15,5%			

GA Gauß-Allianz

Score-E Application Example: Analysis of *Indeed*

- Finite Element code for sheet metal forming simulation
- SMP version (OpenMP) & DMP version (hybrid: OpenMP + MPI) available
- Two different application scenarios:
  - Classical use case: direct simulation of forming process
  - Recent variation: determination of optimal process parameters (in particular, tool geometry)
- Fundamental difference: mesh size for discretization of tools
- Consequences:
  - Significance of contact search algorithms changes drastically
  - Classification changes between “compute bound” and “memory bound”

GA Gauß-Allianz

Score-E Application Example: Analysis of *Indeed*

- Potential tuning parameters:
  - CPU frequency
  - Number of OpenMP threads
  - Number of MPI processes
  - Code path switching

Score-E Application Example: Analysis of *Indeed*

- Potential tuning parameters:
  - CPU frequency
  - Number of OpenMP threads
  - Number of MPI processes
  - Code path switching

Accepted by customers

Undesired by customers

Measurement Results for *Indeed*

Relative values w. r. t. default clock frequency (Sandy Bridge, classical example)

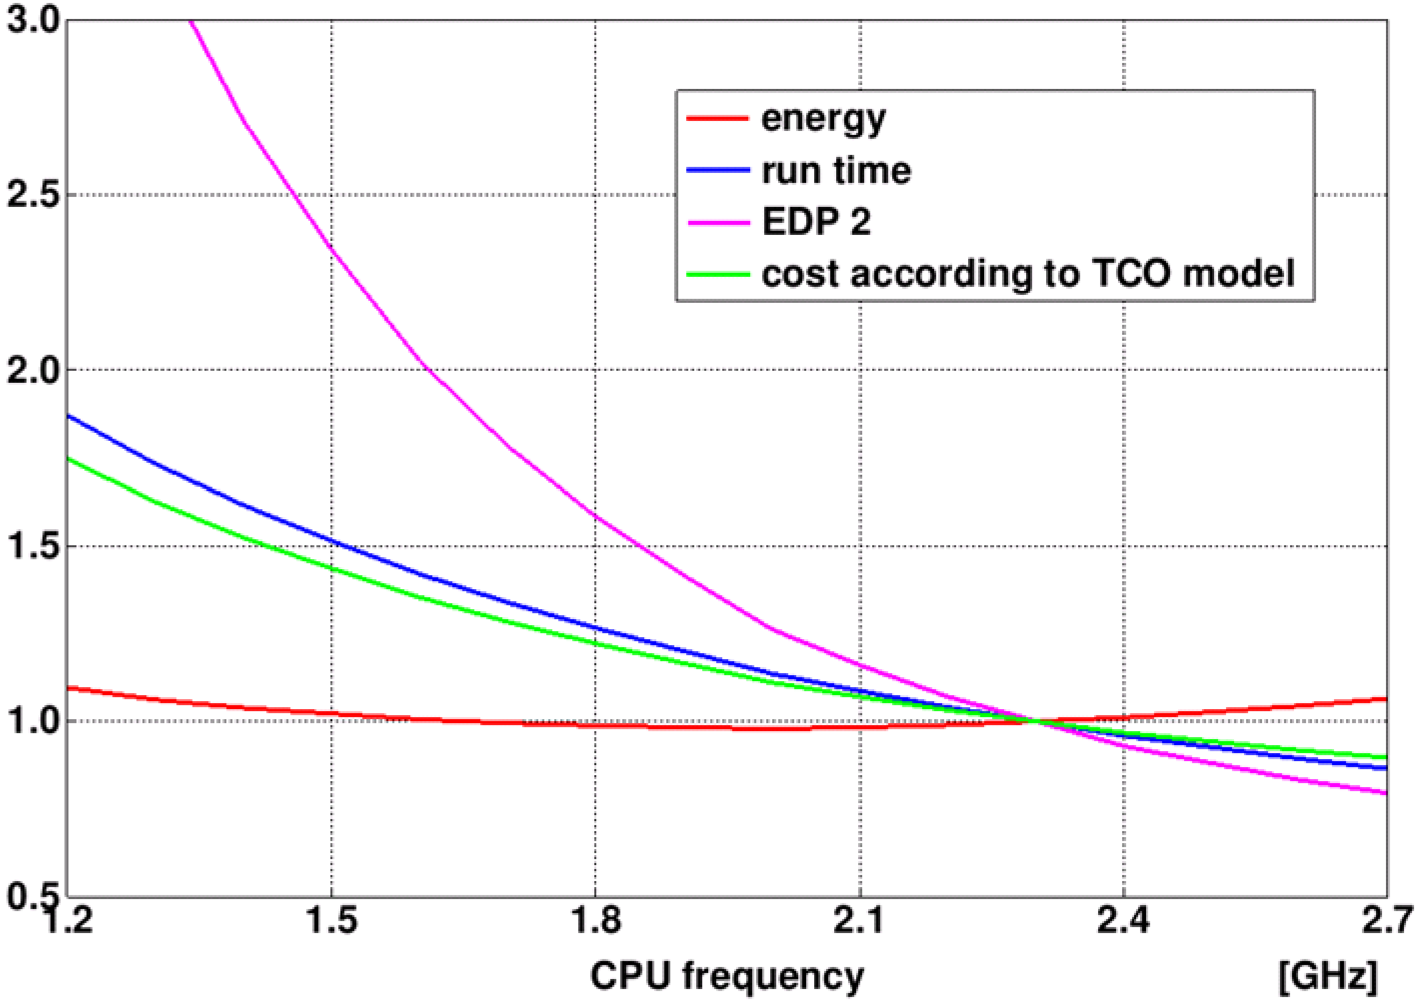

GA Gauß-Allianz

Measurement Results for *Indeed*

Relative values w. r. t. default clock frequency (Haswell, classical example)

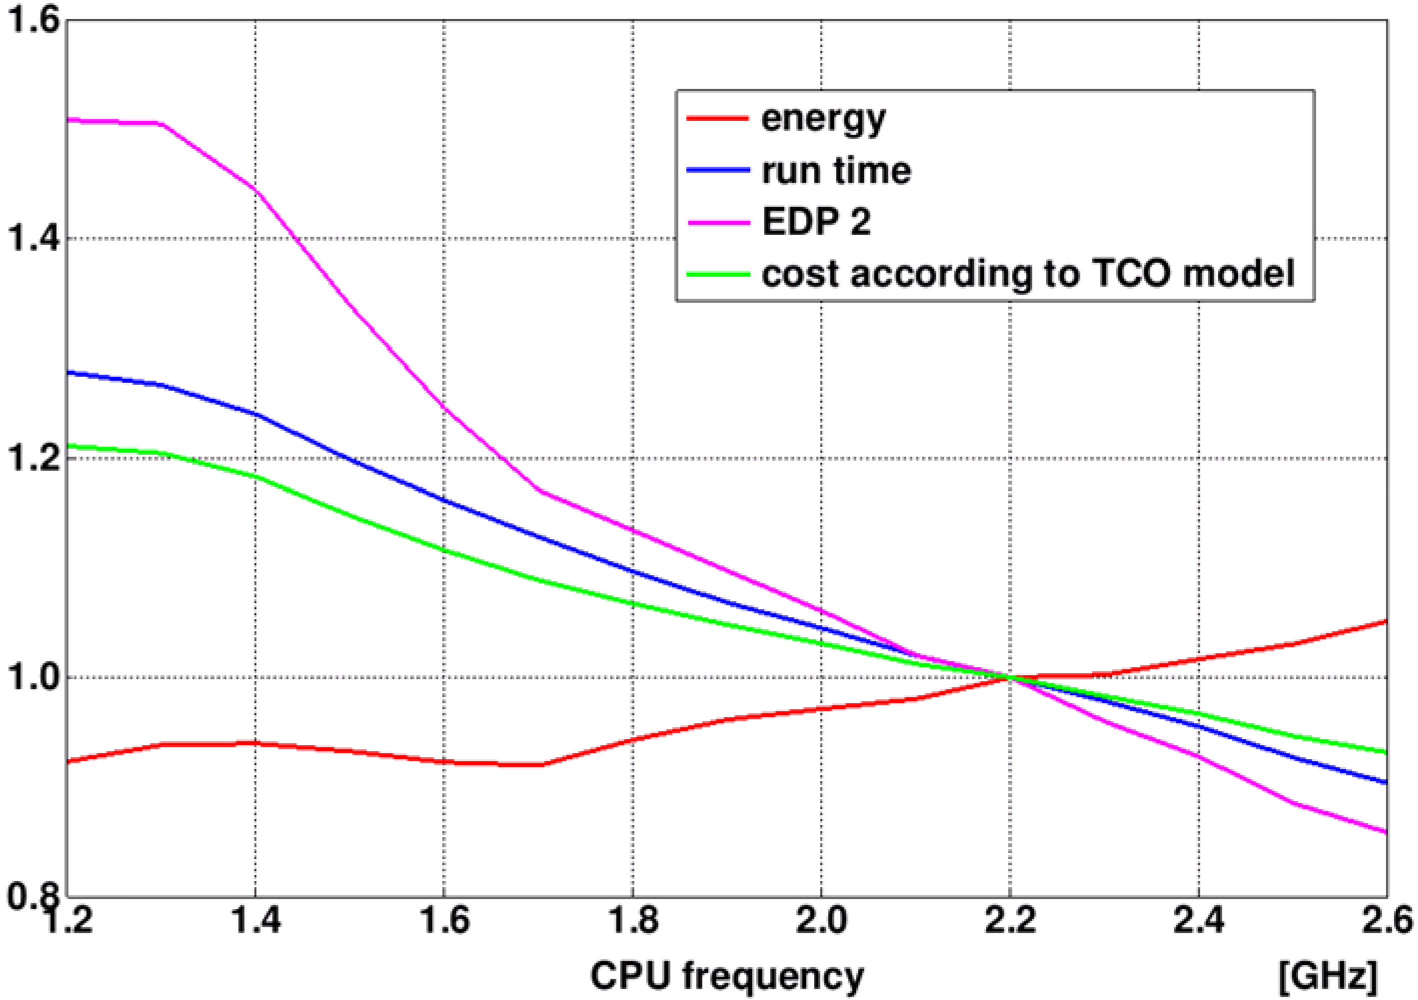

Measurement Results for *Indeed*

Relative values w. r. t. default clock frequency (Haswell, nonclassical example)

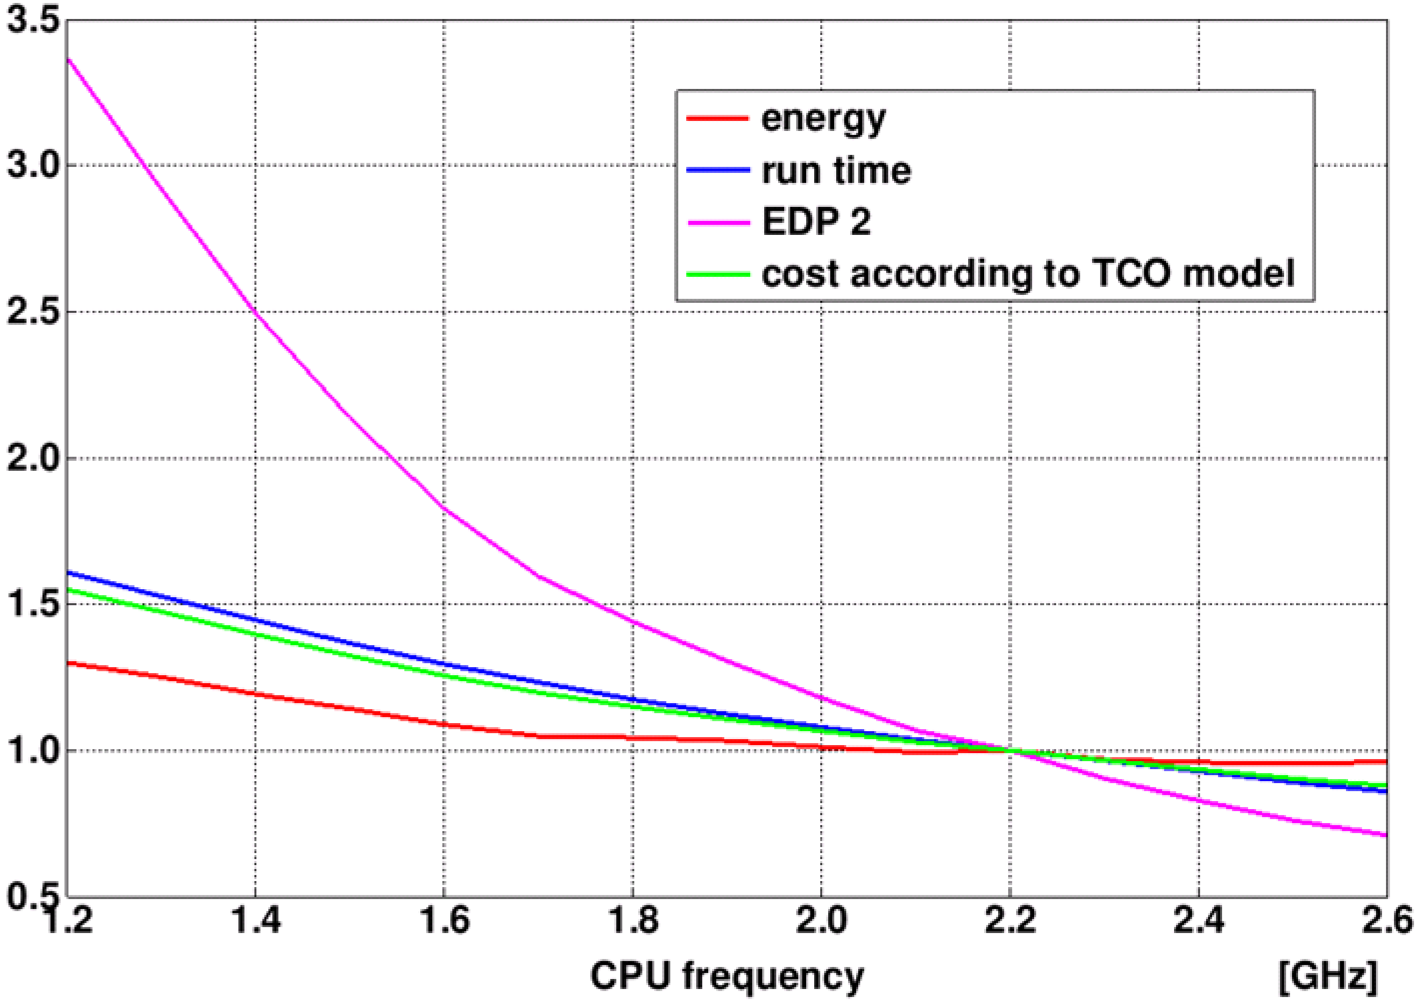

Conclusions for *Indeed*

- Choose high clock frequency
- Exact optimal value depends on
  - Processor architecture
  - Characteristics of input deck
  - Concrete choice of objective function
  - Static tuning very attractive on Sandy Bridge;
- further improvements possible with dynamic tuning
- Optimization on Haswell requires further methods and tools
  (automatic dynamic tuning at runtime → READEX)
- Viability of platform-dependent approach
  from commercial point of view?

GA Gauß-Allianz

Future Work in READEX and Expected Outcome

- Use tools based on PTF for design-time analysis (search for significant regions, dynamism detection, ...)
- Runtime application tuning based on READEX Runtime Library (RRL)
- Alpha version for full READEX tool suite scheduled for spring 2017
- Later step:
  Programming paradigm for expressing application dynamism
  → further improve automatic dynamic energy tuning
- Significant increase in energy efficiency without prohibitively high effort for software developers

GA Gauß-Allianz

Acknowledgements

We thank

- BMBF for providing the funding for Score-E,
- DLR PT-SW for the support and cooperation with respect to Score-E,
- the European Union for the funding for READEX.

Further information

- www.vi-hps.org/projects/score-e
- www.readex.eu

GA

Gauß-Allianz