Runtime analysis of parallel applications for industrial software development

Markus Geimer and Christian Feld, Forschungszentrum Jülich GmbH
Daniel Becker, Siemens AG

JÜLICH
FORSCHUNGSZENTRUM

Outline

Application requirements

Multicore tool support

Runtime recording with Score-P

Supportive tool stack

Summary

Siemens AG

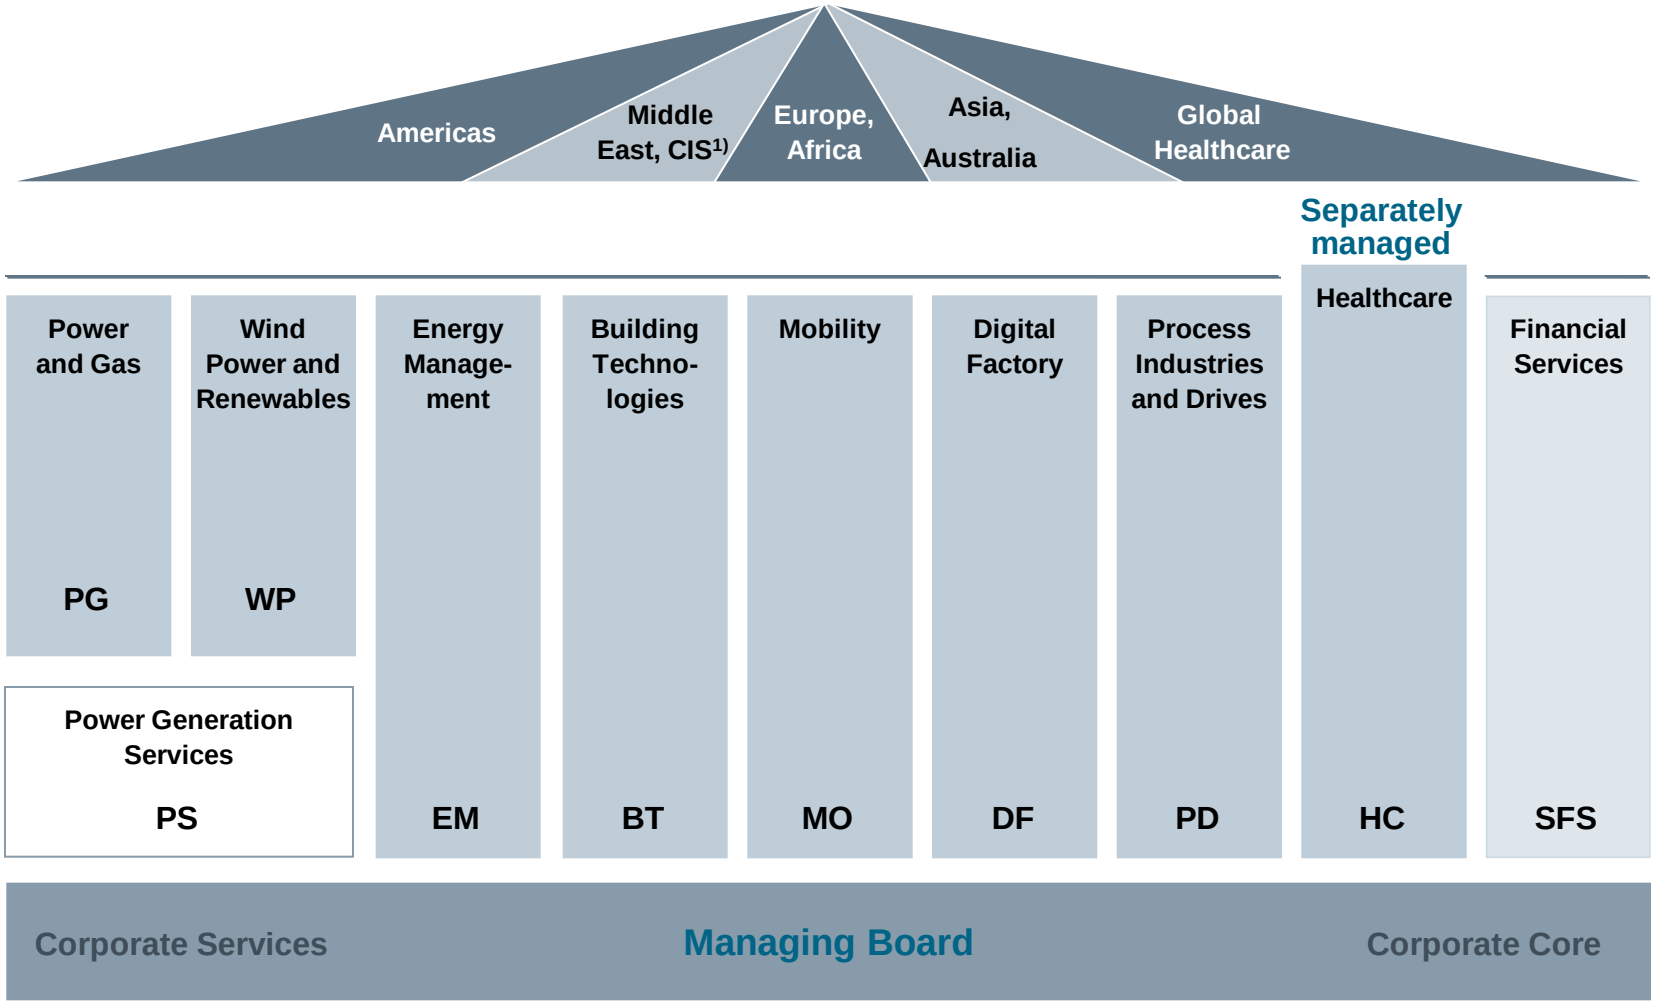

1) Commonwealth of Independent States

Typical multicore migration

Let us assume that an application has already multiple threads…

Concurrency bugs	Poor performance	High effort
• Data races • Deadlocks/livelocks • Nondeterminism • Memory model • …	• Synchronization & communication • Contention on shared resources • …	• Explicit thread management and synchronization • Limited portability • …

Multicore migration is extremely challenging

Outline

Application requirements

Multicore tool support

Runtime recording with Score-P

Supportive tool stack

Summary

Design process towards multicore software

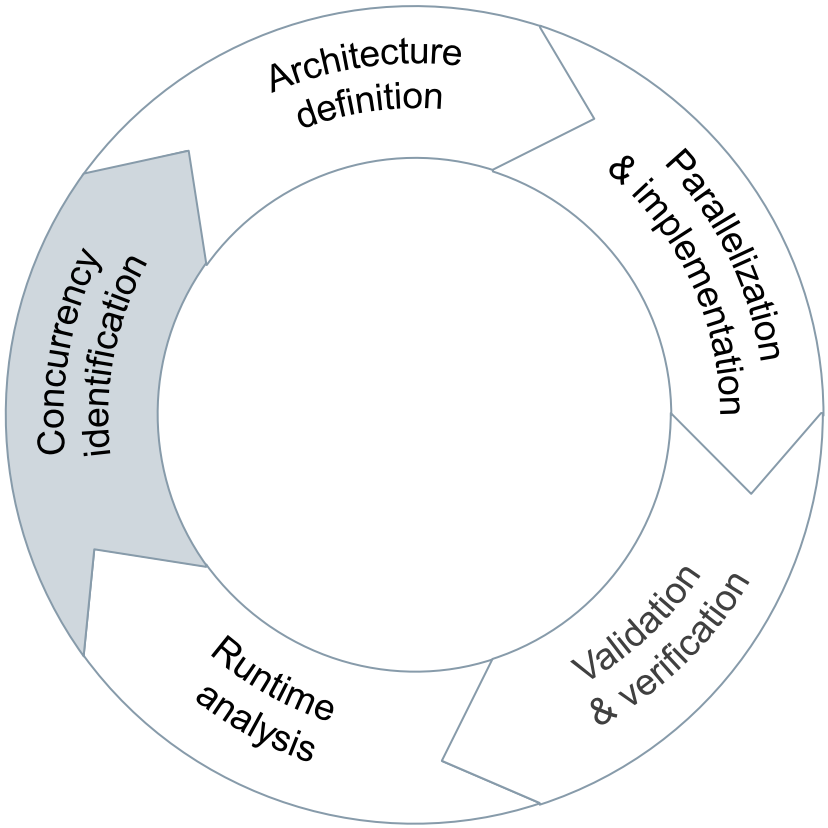

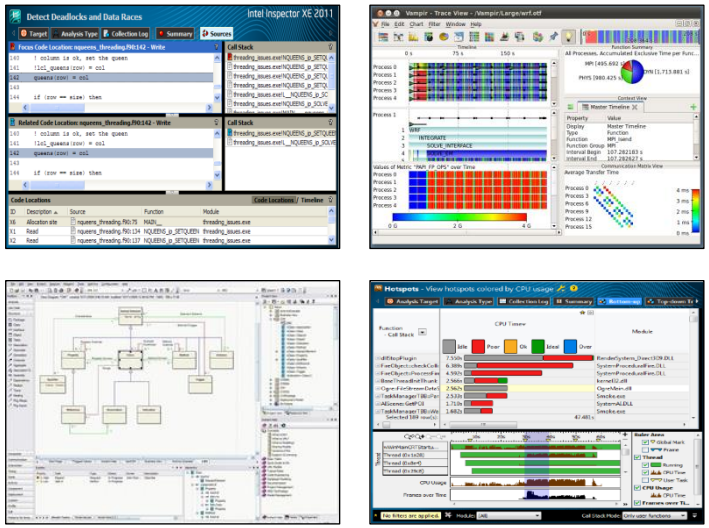

Tool support is essential for an effective and efficient parallelization

Multicore migration scenario

Multicore migration scenario

- Legacy code is code where nobody understands the details
- Runtime behavior is also complex and hard to follow
- Manual extraction of information is cumbersome
- Profiling and tracing tools are essential
  - To understand and debug
  - To engineer and optimize the runtime

Requirements for profiling & tracing systems

- Focus on understanding the application and its parallel aspects
  - Threads & processes
  - Locks & messages
- Portable to Windows, Linux
- HW independent (x86, ARM, and ppc)
- Heterogeneous system support (e.g., Intel Xeon Phi, CUDA)
- Formats enabling interoperability and custom analysis types

Outline

Application requirements

Multicore tool support

Runtime recording with Score-P

Supportive tool stack

Summary

Runtime recording with Score-P

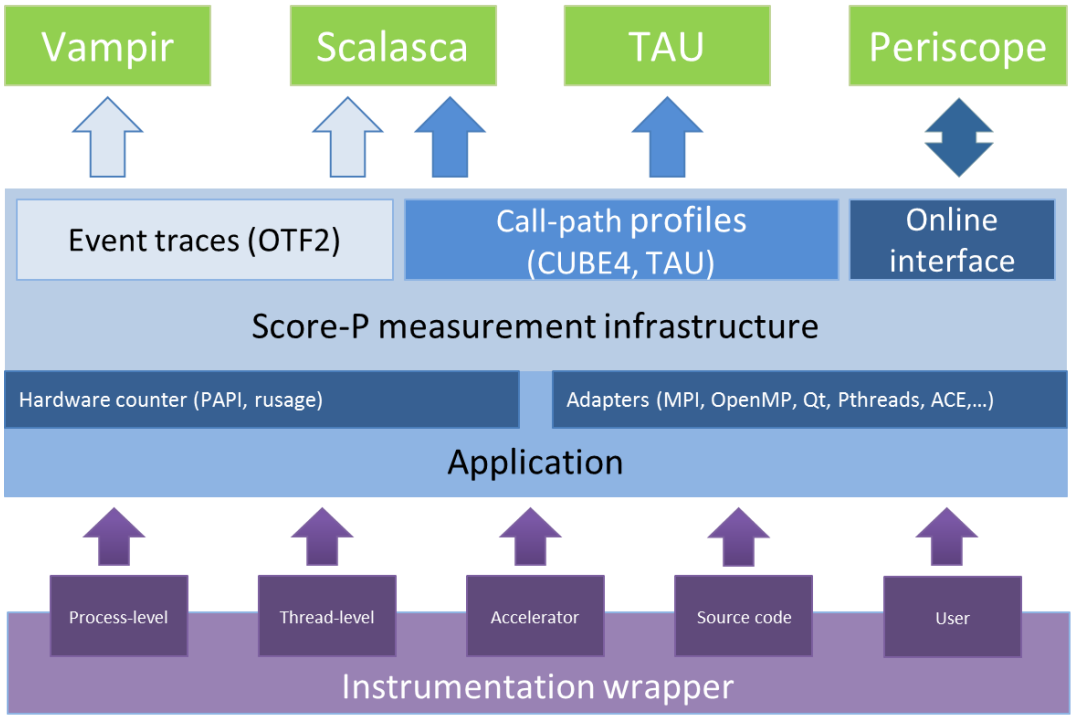

- Open source community
- Linux (& Windows)
- HW independent (x86, ARM, PPC, …)
- Heterogeneous systems (e.g., Intel Xeon Phi, CUDA)
- Open formats enabling interoperability and custom analysis types
- Extremely scalable

Bundesministerium für Bildung und Forschung

U.S. DEPARTMENT OF ENERGY | Office of Science

JÜLICH FORSCHUNGSZENTRUM | TECHNISCHE UNIVERSITÄT DRESDEN | RWTH AACHEN UNIVERSITY | TUM

German Research School for Simulation Sciences | UNIVERSITY OF OREGON

www.score-p.org

Score-P 1.3

Key features

- Provides typical functionality for HPC performance tools
  - Support for process-level parallelism using MPI/SHMEM
  - Support for thread-level parallelism using OpenMP/Pthreads
  - Support for accelerator-level parallelism using CUDA
- Based on instrumentation
  - Supports various techniques
  - Extensive runtime filtering & selective recording capabilities
- Flexible measurement with single re-compilation
  - Basic and advanced profile generation
  - Event trace recording
  - Online access to profiling data
- Scalability: Petascale
- Portability: Supports all major HPC platforms, incl.
  IBM Blue Gene, Cray XT/XE/XK/XC, Fujitsu FX10 & K computer,
  SGI Altix, Power/AIX, Linux-based clusters (x86, ARM, Power)
- Open source: 3-clause BSD license

JUQUEEN
28 rack IBM Blue Gene/Q
28,672 nodes (458,752 cores, 4-way SMT)
448 TB RAM, 5.9 Petaflop/s peak

Score-P workflow

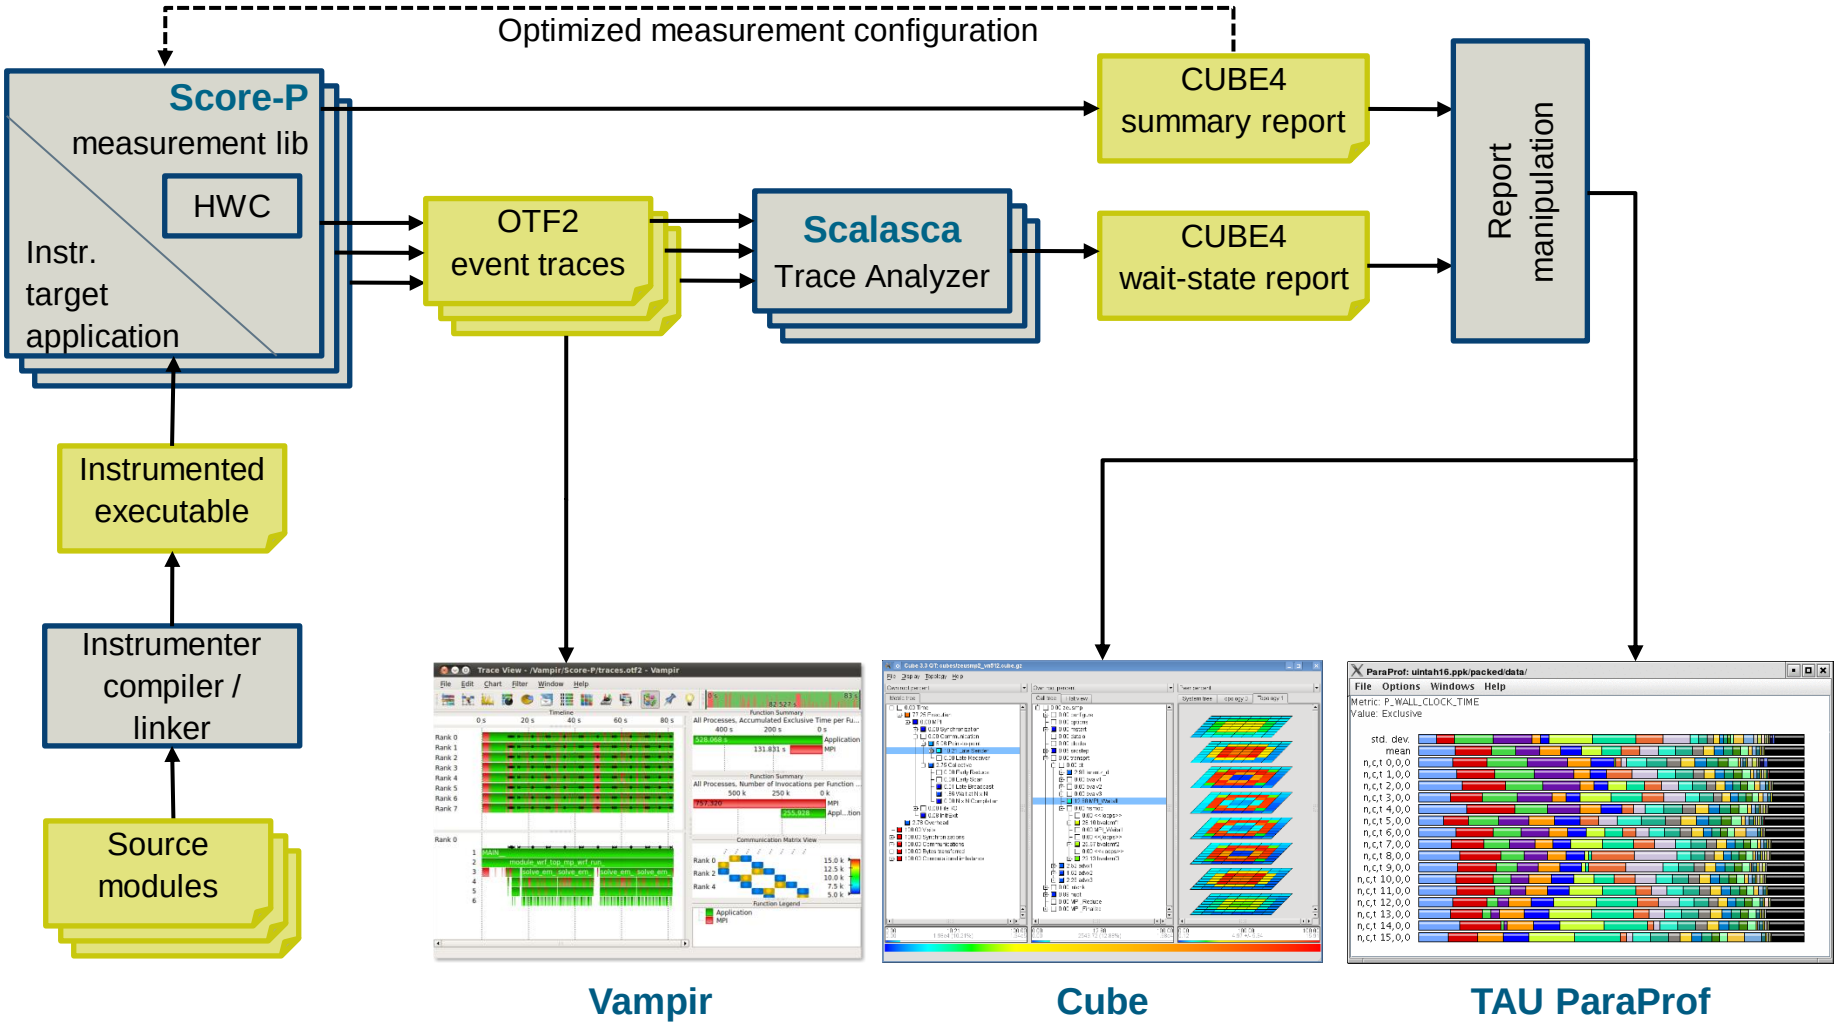

Score-P instrumentation options

Manual instrumentation

- Extensive API (C/C++/Fortran), supporting
  - Program phases
  - Functions
  - Arbitrary code regions
  - Measurement control

Automatic instrumentation

- By using the compiler
  (GCC, Intel, PGI, IBM xl, Cray, Fujitsu)
- By using source-to-source translation
  (Opari2, PDToolkit)
- By linking against a pre-instrumented library
  (MPI, SHMEM)
- By using GNU linker symbol renaming
  (POSIX threads, SHMEM)

Automatic binary instrumentation using Pin

- Pin: dynamic binary instrumentation tool from Intel
- Flexible instrumentation at image-load time
  - No recompilation necessary
  - But only supports x86 architectures

- **Current prototype supports**
  - Function wrapping
  - Incl. pre-runtime filtering (i.e., file level, region level, shared-object level)
  - Replacement of threading API routines (by calls to Score-P measurement system)
    - POSIX threads
    - Windows threads
    - Qt thread API
    - ACE threads

Future work: Support for MTAPI (Multicore Association), Intel TBB

Outline

Application requirements

Multicore tool support

Runtime recording with Score-P

Supportive tool stack

Summary

Cube profile viewer

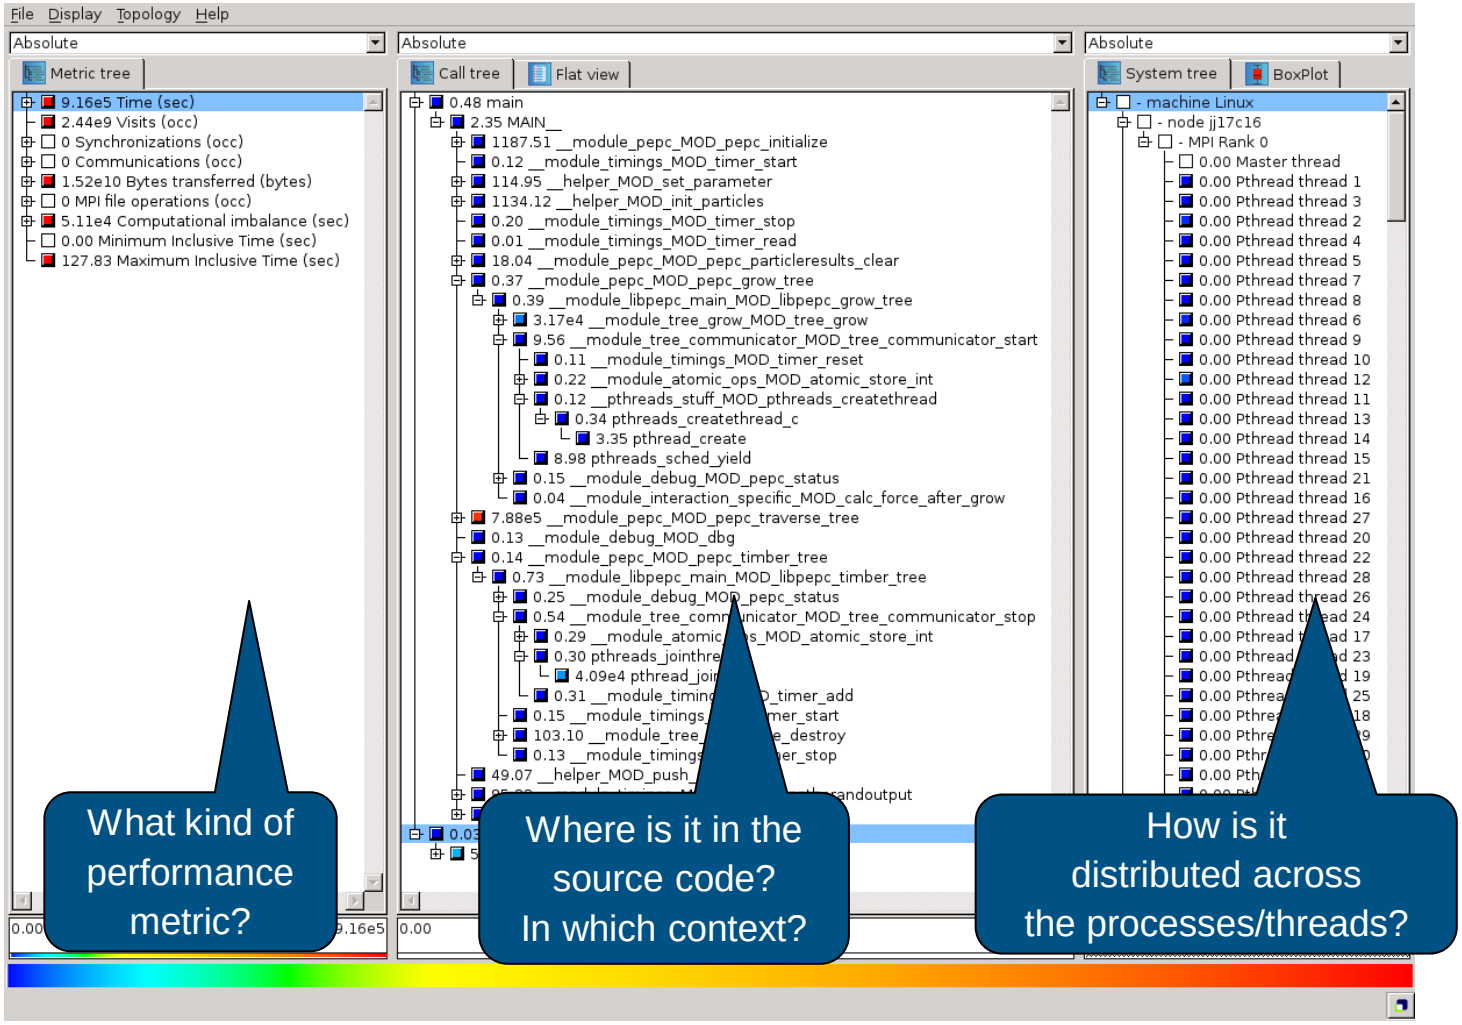

Open Trace Format 2 (OTF2)

Key facts

- Successor to OTF (Vampir) and EPILOG (Scalasca)
- Very flexible, scalable, and space-efficient multi-file format
- Supports many event types
  - Enter/Exit regions/phases/…
  - Point-to-point, collective, one-sided communication, and synchronization
  - Fork/join and create/wait threading
  - HW + SW counters
  - Extensible event attributes
- Well-defined read/write C API

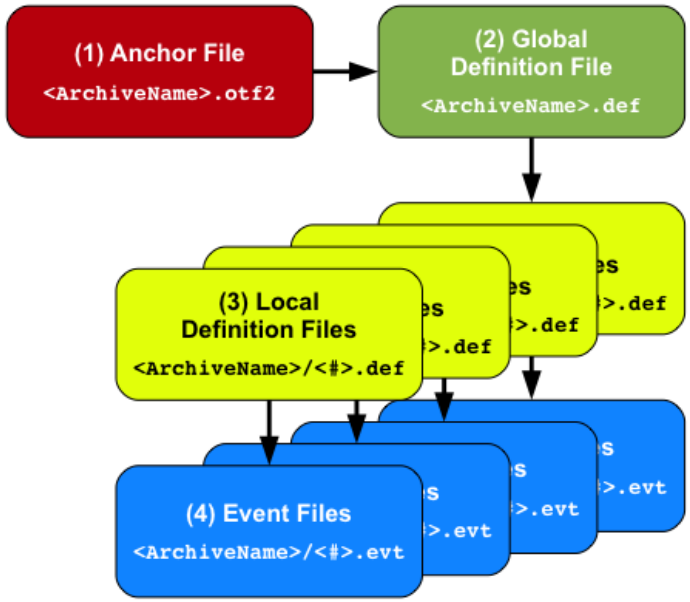

Time-line visualization in Vampir (TU Dresden)

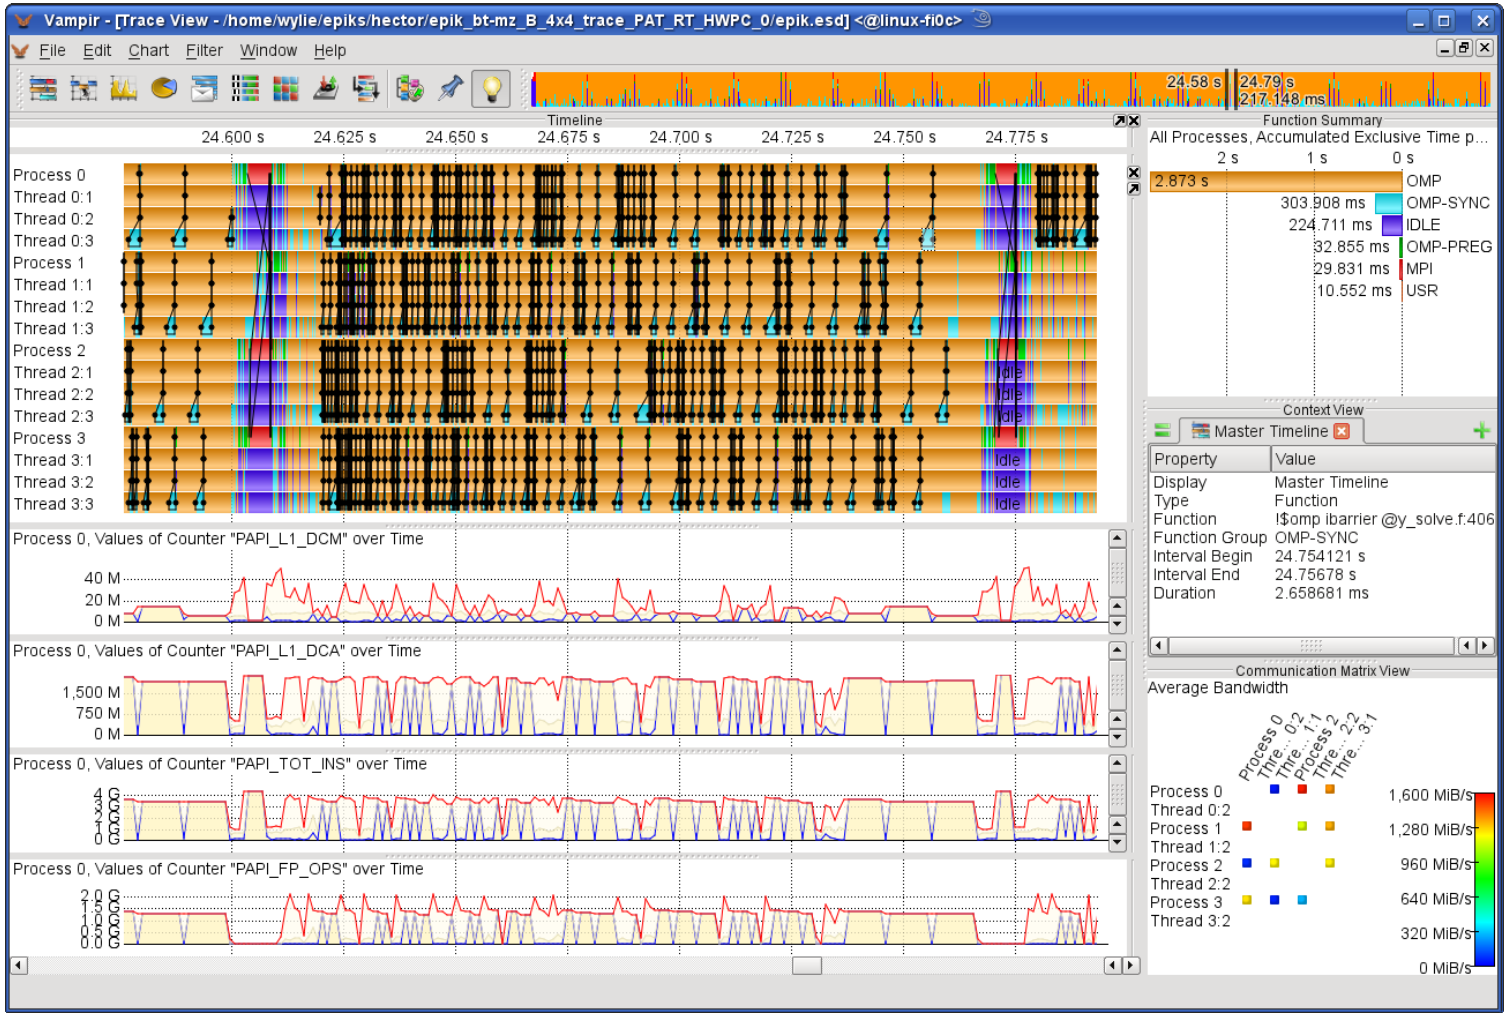

Automatic trace analysis with Scalasca

- **Idea**
  - Automatic search for patterns of inefficient behavior
  - Classification of behavior & quantification of significance

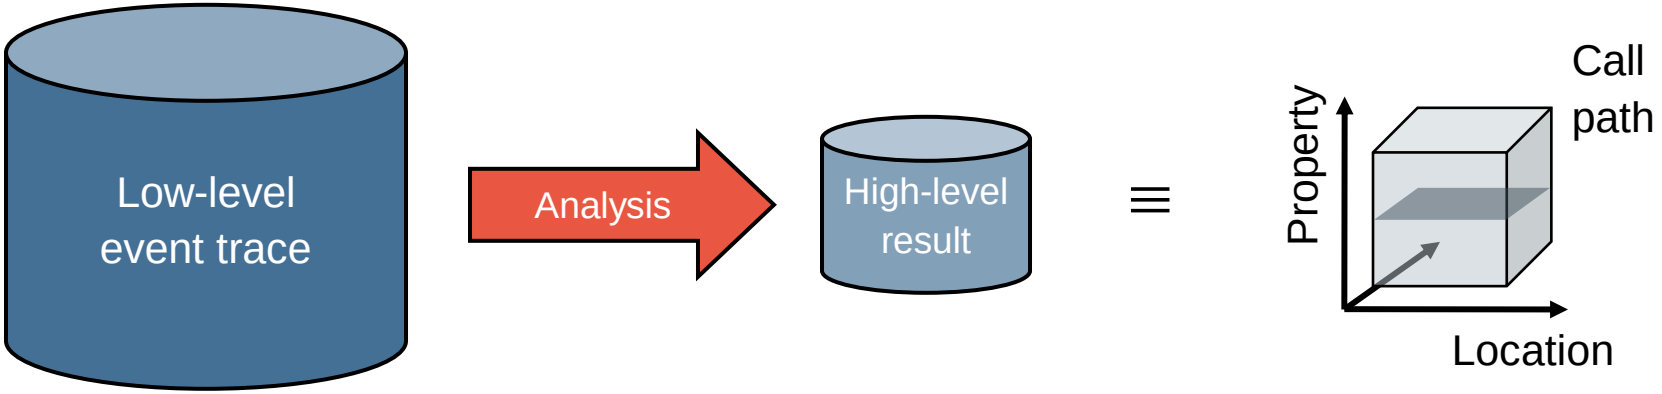

- **Advantages**
  - Guaranteed to cover the entire event trace
  - Quicker than manual/visual trace analysis
  - Parallel replay analysis exploits available memory & processors to deliver scalability

Example: Lock contention analysis

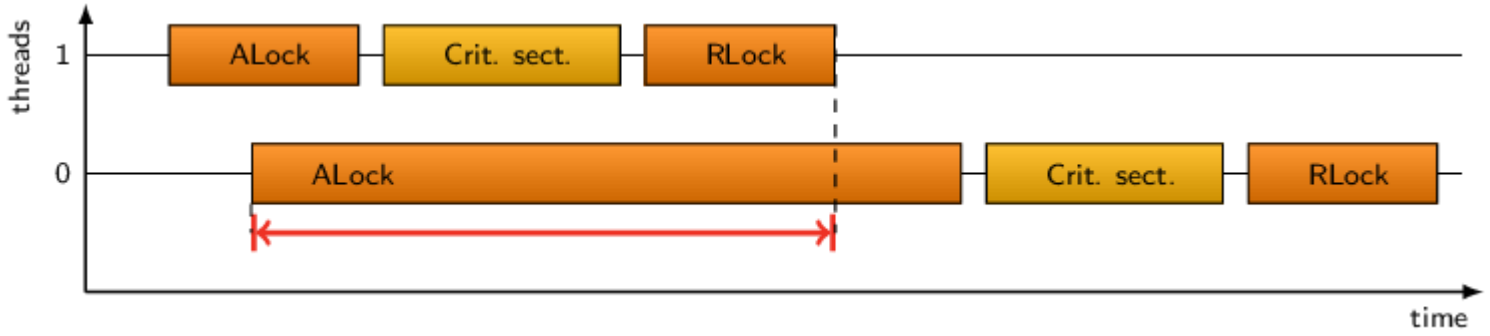

- Frequent and severe problem during multicore migration
- Automatic determination of waiting times in acquire lock operations
- Easy identification of blocking threads in different call paths

- Currently supports
  - OpenMP critical sections & lock API
  - Pthread mutex & condition variable APIs

- Future work
  - Determine root-cause of waiting time

Outline

Application requirements

Multicore tool support

Runtime recording with Score-P

Supportive tool stack

Summary

Summary

- Parallelizing legacy software is tedious
  - Developers have to understand all parts of the software
  - Getting synchronization right *and* efficient is challenging
- Profiling and tracing tools
  - Enable developers to understand, debug, engineer, and optimize their application
  - Have to be portable and extendable
  - Should allow to focus only on relevant parts of the application
- Score-P fulfills these requirements and comes with an supportive tool stack enabling effective and efficient multicore migrations

More information and contacts

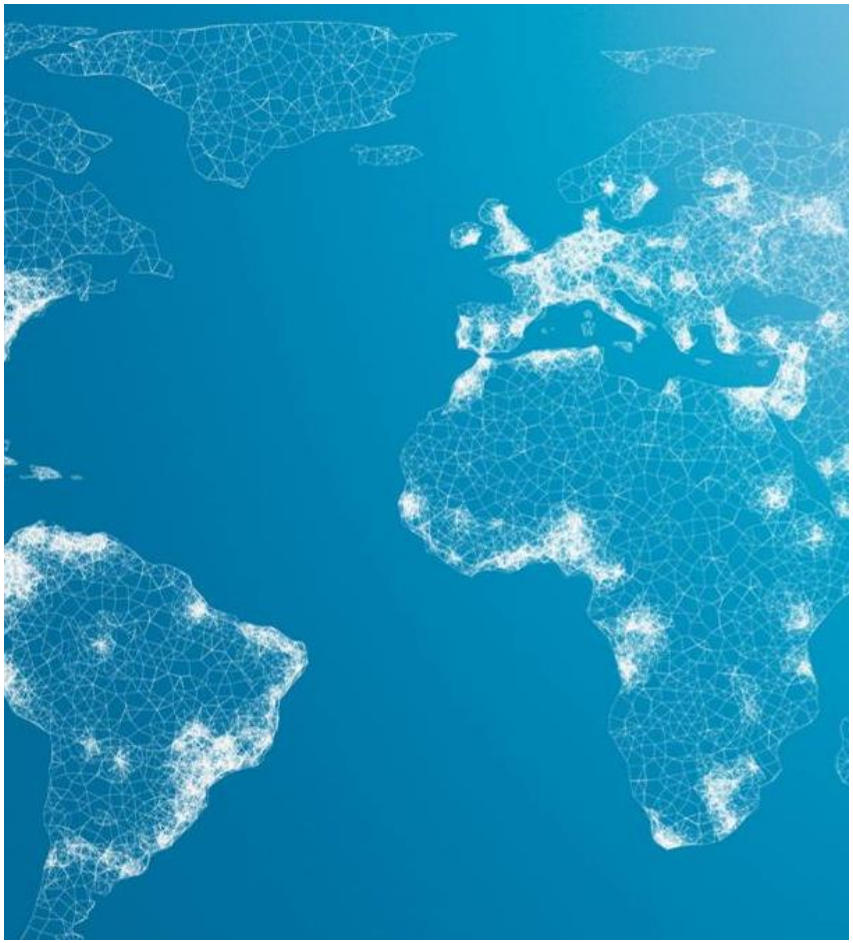

Multicore Expert Center, Siemens AG

- Dr. Daniel Becker
- becker.daniel@siemens.com

Score-P & OTF2

- http://www.score-p.org
- support@scorep.org

Scalasca & Cube

- http://www.scalasca.org
- scalasca@fz-juelich.de