



Unified Performance Data Collection with Score-P

Bert Wesarg¹⁾

With contributions from
Andreas Knüpfer¹⁾, Christian Rössel²⁾, and Felix Wolf³⁾

¹⁾ZIH TU Dresden, ²⁾FZ Jülich, ³⁾GRS-SIM Aachen

- Several performance tools co-exist
- Separate measurement systems and output formats
- Complementary features and overlapping functionality
- Redundant effort for development and maintenance
- Limited or expensive interoperability
- Complications for user experience, support, training

Vampir

Scalasca

TAU

Periscope

VampirTrace
OTF

EPILOG /
CUBE

TAU native
formats

Online
measurement

- Start a community effort for a common infrastructure
 - Score-P instrumentation and measurement system
 - Common data formats OTF2 and CUBE4
- Developer perspective:
 - Save manpower by sharing development resources
 - Invest in new analysis functionality and scalability
 - Save efforts for maintenance, testing, porting, support, training
- User perspective:
 - Single learning curve
 - Single installation, fewer version updates
 - Interoperability and data exchange
- SILC project funded by BMBF
- Close collaboration PRIMA project funded by DOE



GEFÖRDERT VOM

Bundesministerium
für Bildung
und Forschung



- Forschungszentrum Jülich, Germany
- German Research School for Simulation Sciences, Aachen, Germany
- Gesellschaft für numerische Simulation mbH Braunschweig, Germany
- RWTH Aachen University, Germany
- Technische Universität Dresden, Germany
- Technische Universität München, Germany
- University of Oregon, Eugene, USA



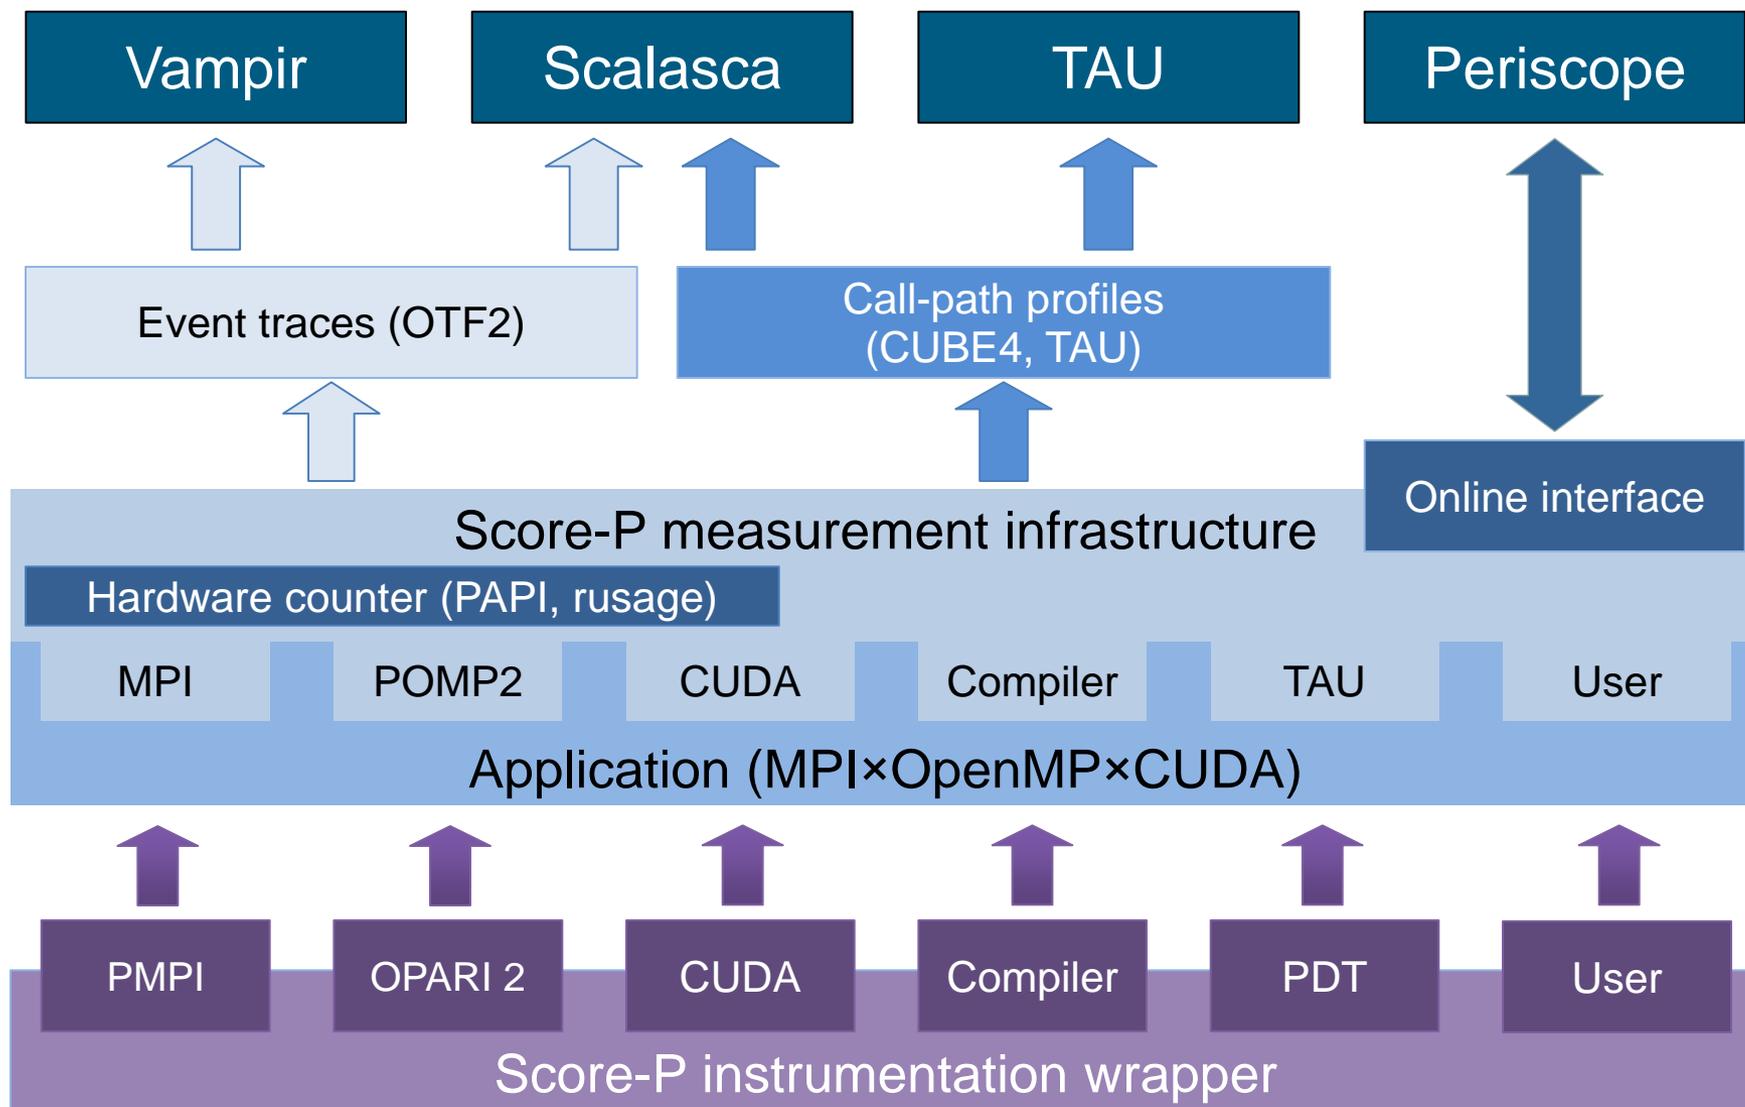
UNIVERSITY OF OREGON

- Provide typical functionality for HPC performance tools
- MPI, OpenMP, and hybrid parallelism (and serial)
- Enhanced functionality (OpenMP 3.0, CUDA, highly scalable I/O)

- Instrumentation (various methods)
- Flexible measurement without re-compilation:
 - Basic and advanced profile generation
 - Event trace recording
 - Online access to profiling data

- **Support all fundamental concepts of partner's tools**

- Portability: all major HPC platforms
- Scalability: petascale
- Low measurement overhead
- Easy and uniform installation through UNITE framework
- Robustness and QA
- Open Source: New BSD License



- 1.0 (Jan 2012)
 - Baseline with MPI and OpenMP support
 - Bugfix releases: 1.0.1 (Apr 2012) and 1.0.2 (Jun 2012)
- 1.1 (Oct 2012)
 - Basic CUDA support
 - OpenMP Task profiling
 - ARM support
- 1.2 (Jun 2013)
 - New generic trace events for RMA records (will be used for MPI 3.0 one-side, ARMCI, SHMEM, CUDA, any PGAS)
 - HMPP, OmpSs, UPC?
 - Pthreads
 - I/O?

- Scalability to maximum available CPU core count
- Support for sampling, binary instrumentation
- Support for new programming models, e.g. PGAS
- Support for new architectures

- Ensure a single official release version at all times which will always work with the tools
- Allow experimental versions for new features or research

- Commitment to joint long-term cooperation

Right now, the following performance tools support Score-P:

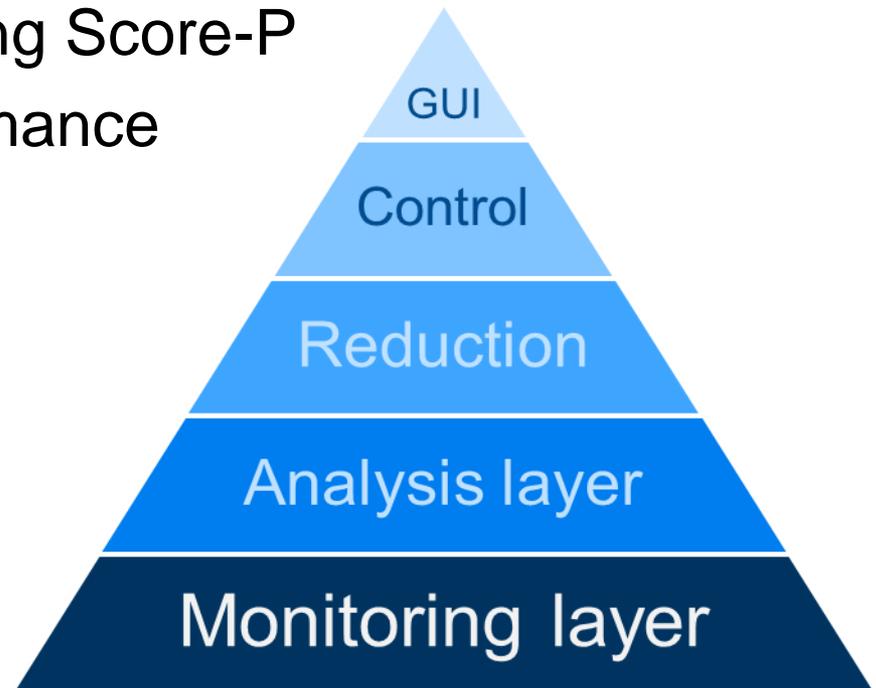
- [Periscope](#) – online bottleneck search
- [Scalasca](#) – call path profiling and wait-state analysis
- [Vampir](#) – trace visualization
- [TAU](#) – profile analysis including data mining

- Scalable distributed tree-like architecture
- Iterative on-line profiling using Score-P
- Automatic search for performance

inefficiencies:

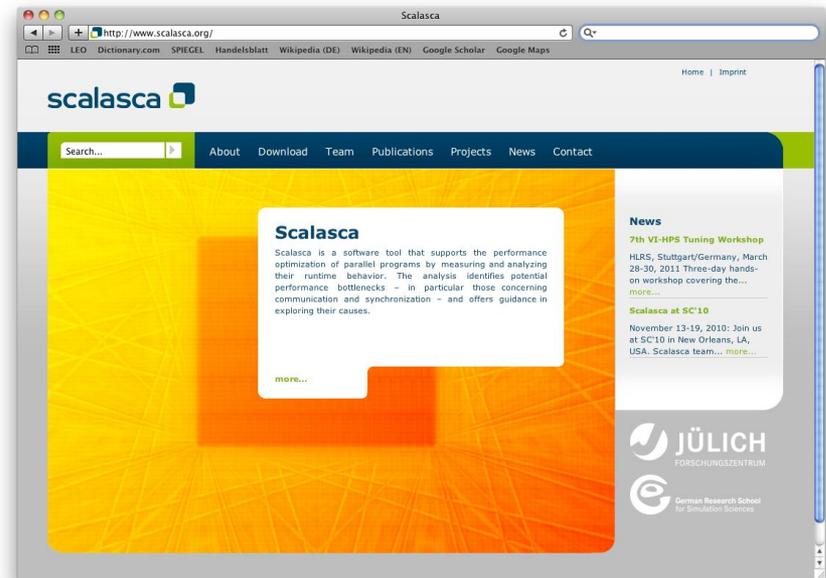
- MPI wait states
 - Single core stalls
 - OpenMP overheads
 - OpenMP scalability
- New BSD license

www.lrr.in.tum.de/periscope



- Scalable performance-analysis toolset for parallel codes
- Automatic event trace analysis to identify potential performance bottlenecks
 - Focus on communication & synchronization
- Programming models
 - MPI, OpenMP
 - Future: support for PGAS and accelerators

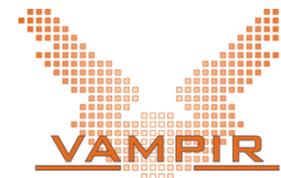
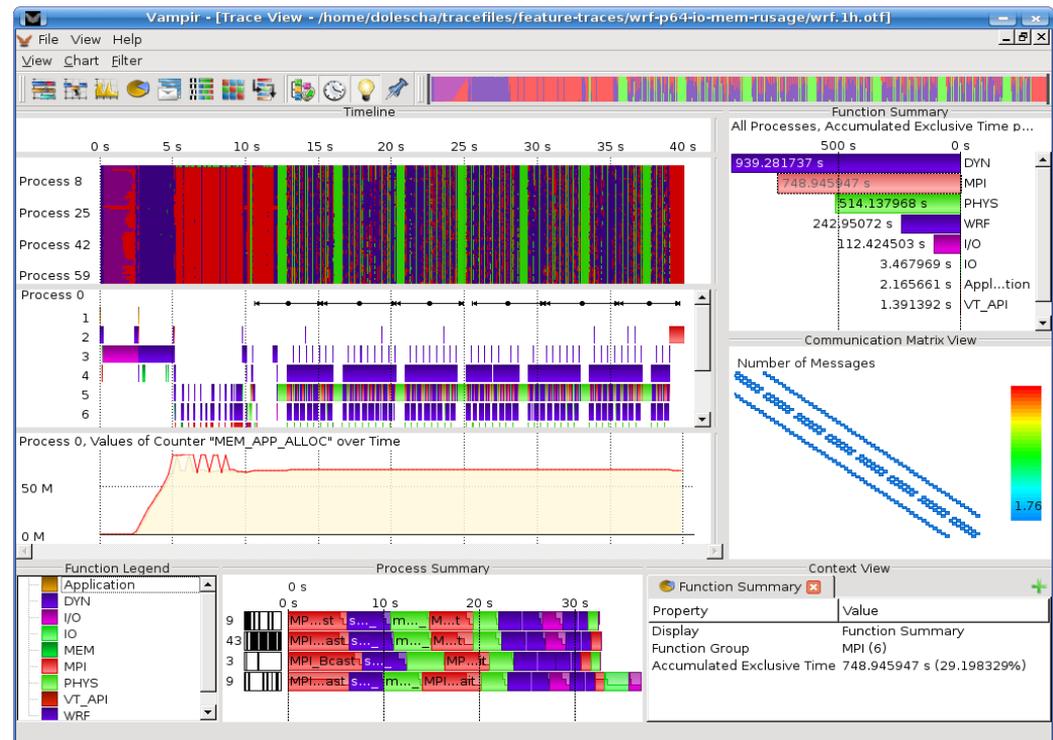
www.scalasca.org



Highly scalable interactive performance analysis with event trace visualization

- As complement to automatic analysis
- Allows to explore and examine all aspects of dynamic parallel behavior
- Very scalable Server version

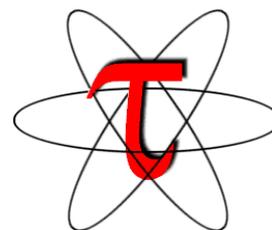
www.vampir.eu



TECHNISCHE
UNIVERSITÄT
DRESDEN

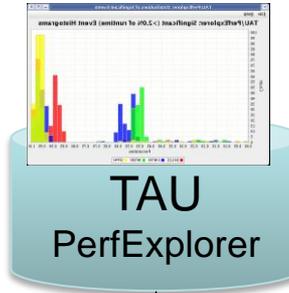
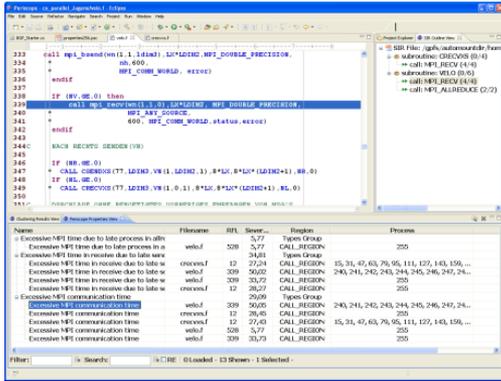
- Profiling and tracing toolkit with automatic instrumentation, measurement and analysis support
- Uses Score-P for profiling and tracing using native OTF2 generation capability
- Fortran, C++, C, UPC, Java, Python, Chapel
- MPI, Pthreads, OpenMP, CUDA, OpenCL, OpenSHMEM, and a combination of threads and MPI
- Goal: to support all HPC platforms, compilers and runtime systems. Widely ported.
- 3D profile browser, ParaProf, data mining and cross experiment analysis tool, PerfExplorer, and performance database technology (PerfDMF)
- Interfaces with Score-P, PAPI, Vampir, MAQAO, and Scalasca
- BSD style license

tau.uoregon.edu

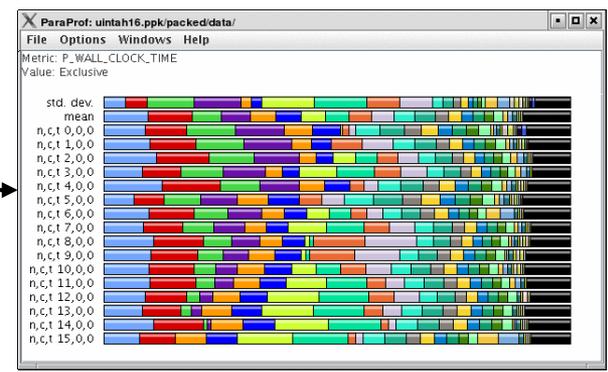


UNIVERSITY OF OREGON

Periscope



TAU ParaProf

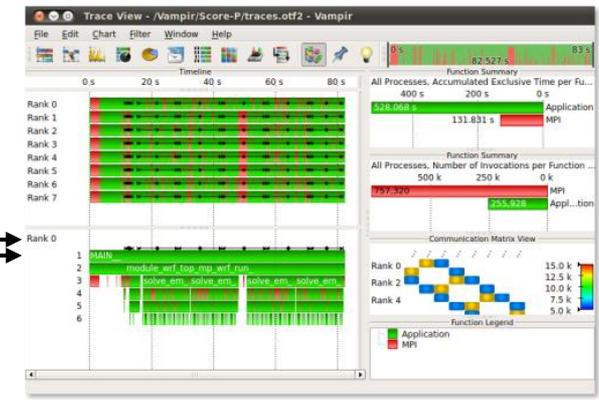
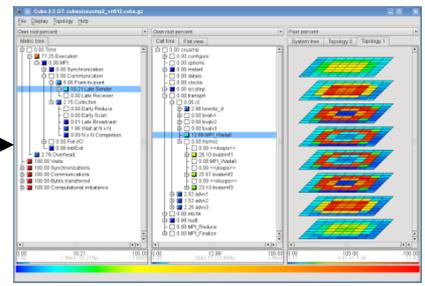


CUBE4 report

CUBE

Scalasca wait-state analysis

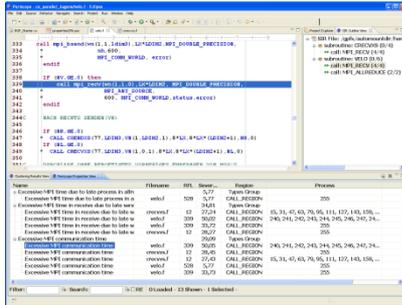
Vampir



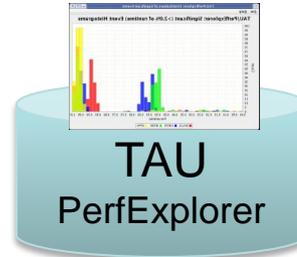
Online interface
Score-P
PAPI
Instr. target application

OTF2 traces

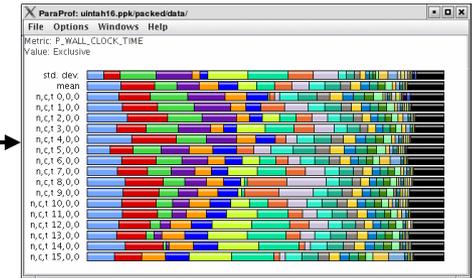
Periscope



TAU PerfExplorer



TAU ParaProf



Online interface

Score-P

PAPI

Instr. target application

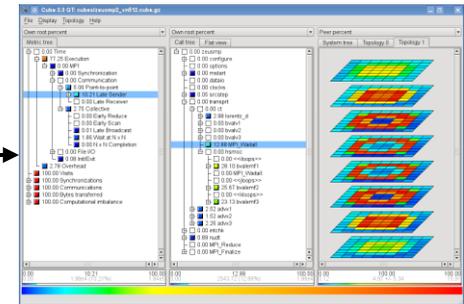
CUBE4 report

CUBE4 report

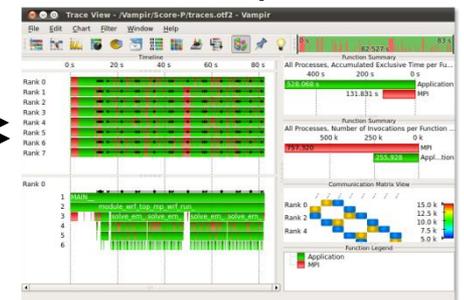
Scalasca wait-state analysis

OTF2 traces

CUBE



Vampir



- We need an extreme-scale performance monitor for extreme-scale performance tools
- This can't be done alone
- So ...

Come Join Us!

Visit www.score-p.org