


VI-HPS



Virtual Institute – High Productivity Supercomputing



14 March 2012



Brian Wylie

Jülich Supercomputing Centre

b.wylie@fz-juelich.de

Wednesday 14th March

- 09:00 (registration)
- Welcome & Introduction to VI-HPS [Ávila, Gerndt, Wylie]
- Introduction to parallel application engineering [Gerndt]
- Introduction to parallel performance analysis [Oleyunik]
- Preparation for hands-on exercises on *levque* [Wylie]
- 11:45-14:30 (lunch)
- **Periscope** introduction & overview [Gerndt, Oleyunik]
- **Periscope** hands-on exercise on *levque*
- 15:45-16:00 (break)
- **Scalasca** introduction & overview [Wylie]
- **Scalasca** hands-on exercise on *levque*
- 17:15 (adjourn)

Thursday 15th March

- 09:00
- **Vampir** introduction & overview [Petkov]
- (break)
- **PAPI** introduction & overview [Ávila]
- 12:00-13:30 (lunch)
- Hands-on coaching with participants' codes on *levque* [all]
- 17:00 (adjourn)

Friday 16th March

- 09:00
- Sponsor presentations from IBM, Intel & SGI (OmegaSystem)
- 12:00 (adjourn)

Tools will ***not*** automatically make you, your applications or computer systems more *productive*.

However, they can help you understand ***how*** your parallel code executes and ***when / where*** it's necessary to work on *correctness* and *performance* issues.

We'd like to know a little about you, your application(s), and your expectations and desires from this tutorial

- What programming paradigms do you use in your app(s)?
 - only MPI, only OpenMP, mixed-mode/hybrid OpenMP/MPI, ...
 - Fortran, C, C++, multi-language, ...
- What platforms/systems *must* your app(s) run well on?
 - Cray XT/XE/XK, IBM BlueGene, SGI Altix, Linux cluster™, ...
- Who's already familiar with *serial* performance analysis?
 - Which tools have you used?
 - ▶ time, print/printf, prof/gprof, VTune, ...
- Who's already familiar with *parallel* performance analysis?
 - Which tools have you used?
 - ▶ time, print/printf, prof/gprof, Periscope, Scalasca, TAU, Vampir, ...

- Ensure your application codes build and run to completion with appropriate datasets
 - initial configuration should ideally run in less than 15 minutes with 1-4 compute nodes (up to 48 processes/threads)
 - ▶ to facilitate rapid turnaround and quick experimentation
 - larger/longer scalability configurations are also interesting
 - ▶ turnaround may be limited due to busyness of batch queues
- Compare your application performance on other systems
 - VI-HPS tools already installed on a number of HPC systems
 - ▶ if not, ask your system administrator to install them (or install a personal copy yourself)

Goal: Improve the quality and accelerate the development process of complex simulation codes running on highly-parallel computer systems

- Start-up funding (2006-2011) by Helmholtz Association of German Research Centres



- Activities
 - Development and integration of HPC programming tools
 - ▶ Correctness checking & performance analysis
 - Training workshops
 - Service
 - ▶ Support email lists
 - ▶ Application engagement
 - Academic workshops

www.vi-hps.org



Forschungszentrum Jülich

- Jülich Supercomputing Centre



RWTH Aachen University

- Centre for Computing & Communication



Technical University of Dresden

- Centre for Information Services & HPC



University of Tennessee (Knoxville)

- Innovative Computing Laboratory





German Research School

- Laboratory of Parallel Programming



Technical University of Munich

- Chair for Computer Architecture



University of Oregon

- Performance Research Laboratory



University of Stuttgart

- HPC Centre



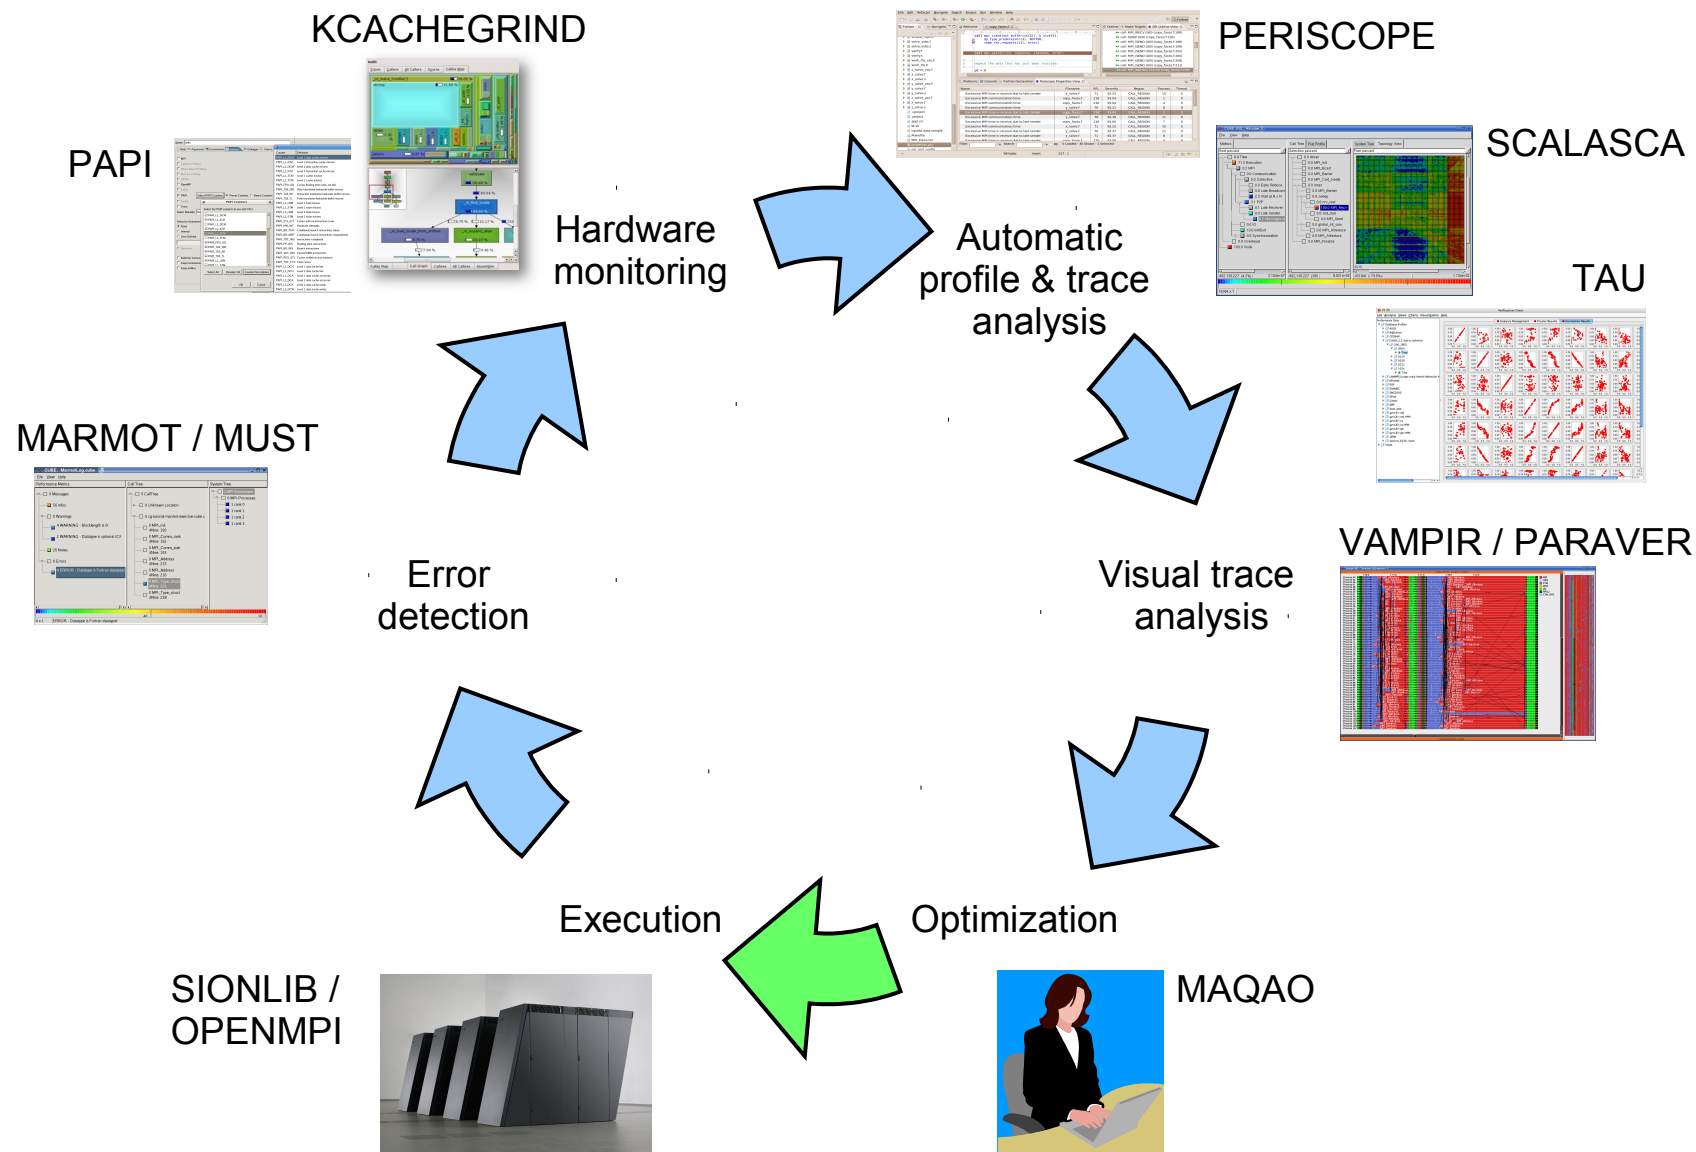
University of Versailles St-Quentin

- LRC ITACA



- **Marmot/MUST**
 - MPI correctness checking
- **PAPI**
 - Interfacing to hardware performance counters
- **Periscope**
 - Automatic analysis via an on-line distributed search
- **Scalasca**
 - Large-scale parallel performance analysis
- **TAU**
 - Integrated parallel performance system
- **Vampir/VampirTrace**
 - Event tracing and graphical trace visualization & analysis

- **KCachegrind**
 - Callgraph-based cache analysis [x86 only]
- **MAQAO**
 - Assembly instrumentation & optimization [x86 only]
- **ompP**
 - OpenMP profiling tool
- **OpenMPI**
 - Memory checking
- **Paraver/Extrae**
 - Event tracing and graphical trace visualization & analysis
- **Score-P**
 - Common instrumentation & measurement infrastructure
- **SIONlib**
 - Optimized native parallel file I/O




- Goals
 - Give an overview of the programming tools suite
 - Explain the functionality of individual tools
 - Teach how to use the tools effectively
 - Offer hands-on experience and expert assistance using tools
 - Receive feedback from users to guide future development
- For best results, bring & analyse/tune your own code(s)!
- VI-HPS Tutorial series
 - SC'08, ICCS'09, SC'09, Cluster'10, SC'10, SC'11
- VI-HPS Tuning Workshop series
 - 2008 (Aachen & Dresden), 2009 (Jülich & Bremen), 2010 (Garching & Amsterdam), 2011 (Stuttgart & Aachen)
 - **2012/04/23-27 (Paris)**, 2012/10/15-19 (Garching)

- 9th VI-HPS Tuning Workshop (23-27 Apr 2012)
 - hosted by UVSQ, St.-Quentin-en-Yvelines, France
 - using PRACE Tier-0 *Curie* system at CEA / TGCC
 - Scalasca, Vampir, TAU, Periscope, KCachegrind, MAQAO, ...
- Further events to be determined
 - (one-day) tutorials
 - ▶ with guided exercises using Live DVD
 - (multi-day) training workshops
 - ▶ with your own applications on real HPC systems

Check www.vi-hps.org/training for announced events

- Contact us if you might be interested in hosting an event

- Bootable Linux installation ISO (on DVD or USB stick)
- Includes everything needed to try out our parallel tools on an x86-architecture notebook computer
 - VI-HPS tools: KCachegrind, Marmot, PAPI, Periscope, Scalasca, TAU, VT/Vampir*
 - Also: Eclipse/PTP, TotalView*, etc.
 - ▶ * time/capability-limited evaluation licences provided for commercial products
 - GCC (w/ OpenMP), OpenMPI
 - Manuals/User Guides
 - Tutorial exercises & examples
- Produced by U. Oregon PRL
 - Sameer Shende



Cachegrind: cache analysis by simple cache simulation

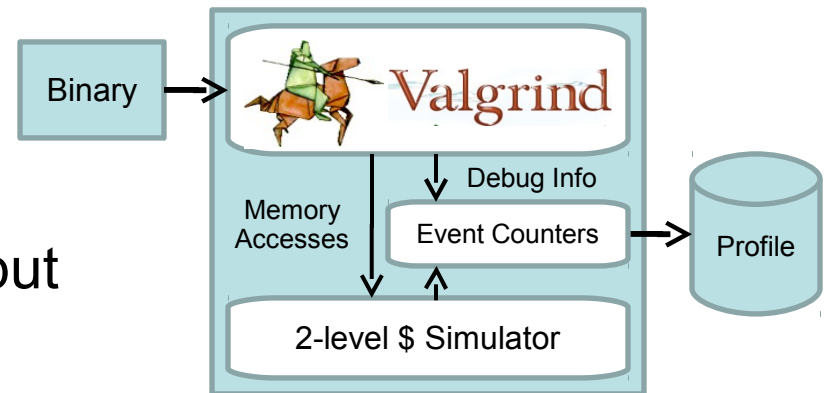
- Captures dynamic callgraph
- Based on valgrind dynamic binary instrumentation
- Runs on x86/PowerPC/ARM unmodified binaries
 - ▶ No root access required
- ASCII reports produced

[KQ]Cachegrind GUI

- Visualization of cachegrind output

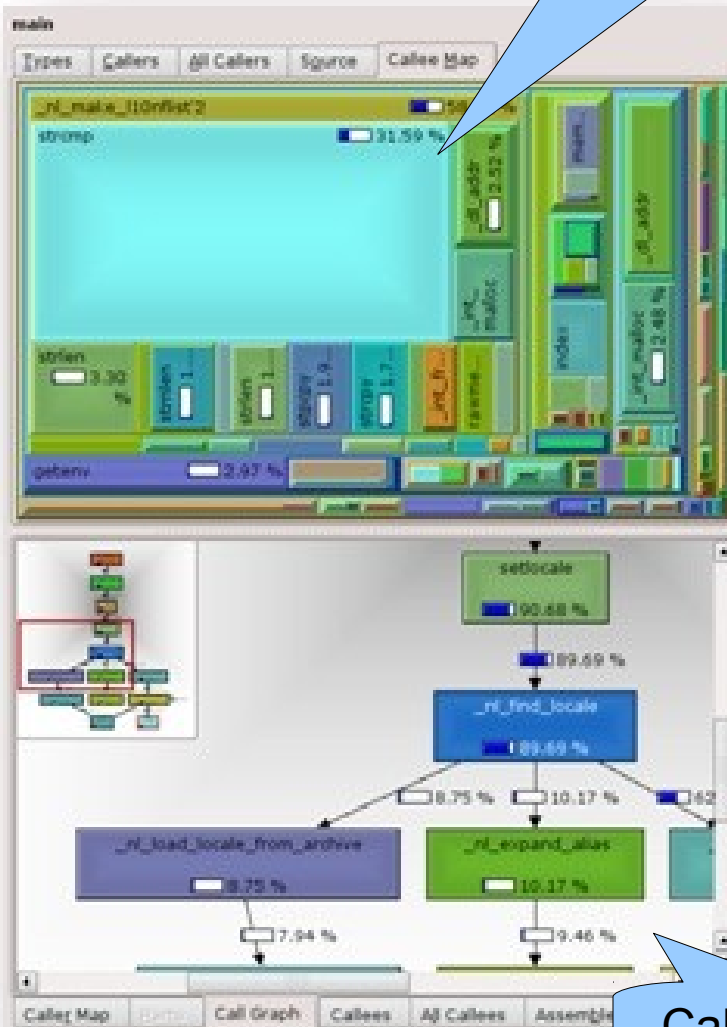
Developed by TU Munich

- Released as GPL open-source
- <http://kcachegrind.sf.net/>

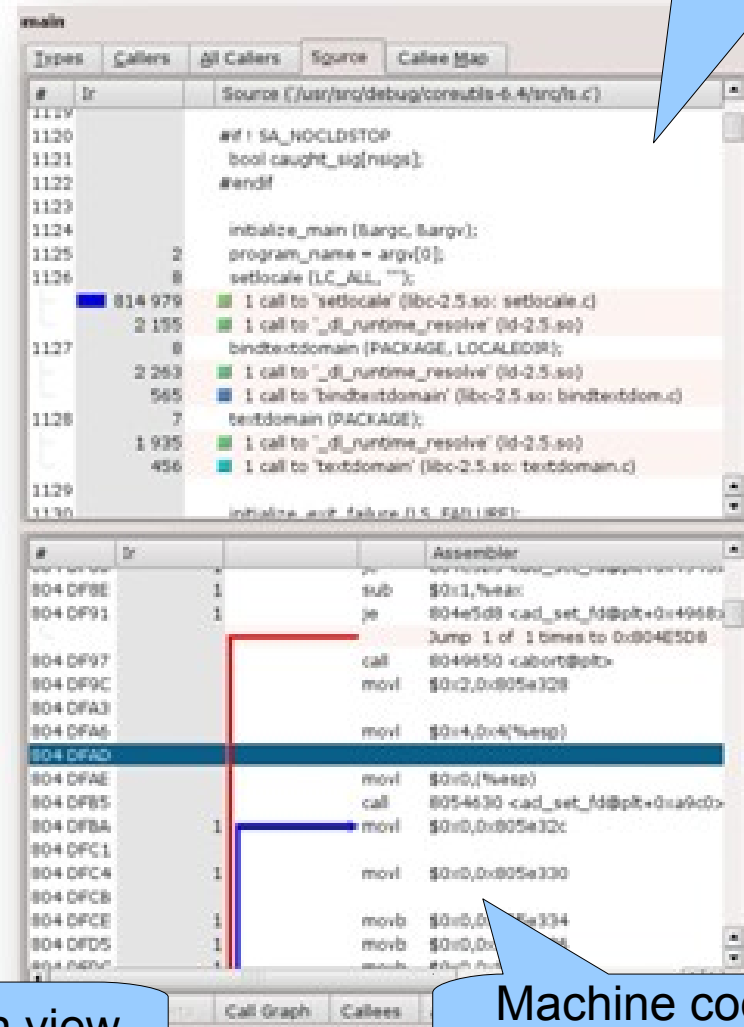


Event cost tree map

Source code view



Call graph view



Machine code annotation

Tool to check for correct MPI usage at runtime



- Checks conformance to MPI standard
 - ▶ Supports Fortran & C bindings of MPI-1.2
- Checks parameters passed to MPI
- Monitors MPI resource usage

Implementation

- C++ library gets linked to the application
- Does not require source code modifications
- Additional process used as DebugServer
- Results written in a log file (ASCII/HTML/CUBE)

Developed by HLRS & TU Dresden

- Released as open-source
- <http://www.hlrs.de/organization/av/amt/projects/marmot>

Marmot logfiles

```

livetau@localhost:Exercise
1 (localhost.localdomain)
for MPI-Standard information see:/usr/local/packages/marmot-2.3.0/share/doc/marmot-2.3.0/MPI-STANDARD/marmot_err/node164.html

3: Warning global message with Text: Processes 0 and 1 both run on localhost.localdomain
for MPI-Standard information see:/usr/local/packages/marmot-2.3.0/share/doc/marmot-2.3.0/MPI-STANDARD/marmot_err/node165.html

```

```

10: Error from rank 0(Thread: 0) with Text: ERROR: MPI_Send: datatype is not valid!

```

```

On Call: MPI Send From: datatype.c line: 53 for MPI-Standard information see:/usr/local/packages/marmot-2.3.0/MPI-STANDARD/marmot_err/node28.html

```

```

10: Error from rank 1(Thread: 1) with Text: ERROR: MPI_Recv: datatype is not valid!

```

```

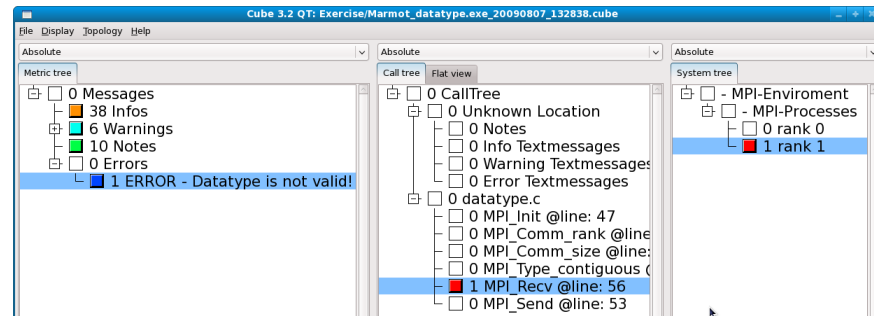
On Call: MPI Recv From: datatype.c line: 56 for MPI-Standard information see:/usr/local/packages/marmot-2.3.0/MPI-STANDARD/marmot_err/node28.html

```

```

[livetau@localhost:Exercise]

```



MARMOT HTML Logfile - Konqueror						
/home/livetau/workshop-marmot/Exercise/Marmot_datatype.exe_20090807_130509.html						
				default: 1000 microseconds)		
0	Global	0	Information	Text: MARMOT_MAX_TIMEOUT_ONE = 0 (maximum message time, default: 0 microseconds)	Unknown	
0	Global	0	Information	Text: MARMOT_MAX_TIMEOUT_TWO = 0 (maximum message time, default: 0 microseconds)	Unknown	
0	Global	0	Information	Text: MARMOT_LOGFILE_PATH = (path of Marmot log file output, default:)	Unknown	
0	Global	0	Information	Text: MARMOT_ERRCODES_SET = (not set) (not functional yet)	Unknown	
0	Global	0	Information	Text: End of the environmental variables info.	Unknown	
0	Global	0	Information	Text: Thread Synchronisation is disabled.If you are using multiple threads errors might occur	Unknown	
3	Global	0	Warning	Text: Debugserver runs on same node as process 0 (localhost.localdomain)	Unknown	Infos see MPI-Standard
3	Global	0	Warning	Text: Debugserver runs on same node as process 1 (localhost.localdomain)	Unknown	Infos see MPI-Standard
3	Global	0	Warning	Text: Processes 0 and 1 both run on localhost.localdomain	Unknown	Infos see MPI-Standard
10	0	0	Error	Text: ERROR: MPI_Send: datatype is not valid! Call: MPI_Send	datatype.c line: 53	Infos see MPI-Standard
10	1	0	Error	Text: ERROR: MPI_Recv: datatype is not valid! Call: MPI_Recv	datatype.c line: 56	Infos see MPI-Standard

Next generation MPI runtime error detection tool

- Successor of the Marmot and Umpire tools
- Initial merge of Marmot's many local checks with Umpire's non-local checks
- Improved scalability expected in future

Developed by TU Dresden, LLNL & LANL

- to be released as open-source (BSD license)
- currently in beta-testing for first release in November 2011
- <http://tu-dresden.de/.../must>

Portable performance counter library & utilities

- Configures and accesses hardware/system counters
- Predefined events derived from available native counters
- Core component for CPU/processor counters
 - ▶ instructions, floating point operations, branches predicted/taken, cache accesses/misses, TLB misses, cycles, stall cycles, ...
 - ▶ performs transparent multiplexing when required
- Extensible components for off-processor counters
 - ▶ InfiniBand network, Lustre filesystem, system hardware health, ...
- Used by multi-platform performance measurement tools
 - ▶ Periscope, Scalasca, TAU, VampirTrace, ...

Developed by UTK-ICL

- Available as open-source for most modern processors
<http://icl.cs.utk.edu/papi/>



PAPI preset counters (and their definitions)



```
juropa$ papi_avail
```

Available events and hardware information.

```
-----
PAPI Version           : 4.1.0.0
Vendor string and code : GenuineIntel (1)
Model string and code  : Intel(R) Xeon(R) CPU
                        X5570 @ 2.93GHz (26)
CPU Revision           : 5.000000
CPUID Info             : Family: 6  Model: 26
                        Stepping: 5
CPU Megahertz          : 1600.000000
CPU Clock Megahertz    : 1600
Hdw Threads per core   : 2
Cores per Socket       : 4
NUMA Nodes             : 2
CPU's per Node         : 8
Total CPU's            : 16
Number Hardware Counters : 16
Max Multiplex Counters : 512
-----
```

Name	Code	Avail	Deriv	Description
PAPI_L1_DCM	0x80000000	Yes	No	Level 1 data cache misses
PAPI_L1_ICM	0x80000001	Yes	No	Level 1 instruction cache misses

...

Of 107 possible events, 35 are available, of which 9 are derived.

```
juropa$ papi_avail -d
```

```
...
Symbol          Event Code  Count  |Short Descr.|
|Long Description|
|Developer's Notes|
|Derived|
|PostFix|
Native Code[n]: <hex> |name|
PAPI_L1_DCM      0x80000000  1 |L1D cache misses|
|Level 1 data cache misses|
||
|NOT_DERIVED|
||
Native Code[0]: 0x40002028 |L1D:REPL|
PAPI_L1_ICM      0x80000001  1 |L1I cache misses|
|Level 1 instruction cache misses|
||
|NOT_DERIVED|
||
Native Code[0]: 0x40001031 |L1I:MISSES|
PAPI_L2_DCM      0x80000002  2 |L2D cache misses|
|Level 2 data cache misses|
||
|DERIVED_SUB|
||
Native Code[0]: 0x40000437 |L2_RQSTS:MISS|
Native Code[1]: 0x40002037 |
L2_RQSTS:IFETCH_MISS|
...
```

PAPI native counters (and qualifiers)



```
juropa$ papi_native_avail
```

Available native events and hardware information.

...

Event Code	Symbol	Long Description
------------	--------	------------------

0x40000000	UNHALTED_CORE_CYCLES	count core clock cycles whenever the clock signal on the specific core is running (not halted). Alias to event CPU_CLK_UNHALTED:THREAD
------------	-----------------------------	--

0x40000001	INSTRUCTION_RETIRED	count the number of instructions at retirement. Alias to event INST_RETIRED:ANY_P
------------	----------------------------	---

...

0x40000086	UNC_SNP_RESP_TO_REMOTE_HOME	Remote home snoop response - LLC does not have cache line
40000486	:I_STATE	Remote home snoop response - LLC does not have cache line
40000886	:S_STATE	Remote home snoop response - LLC has cache line in S state
40001086	:FWD_S_STATE	Remote home snoop response - LLC forwarding cache line in S state.
40002086	:FWD_I_STATE	Remote home snoop response - LLC has forwarded a modified cache line
40004086	:CONFLICT	Remote home conflict snoop response
40008086	:WB	Remote home snoop response - LLC has cache line in the M state
40010086	:HITM	Remote home snoop response - LLC HITM

Total events reported: 135

Automated profile-based performance analysis

- Iterative on-line performance analysis
 - ▶ Multiple distributed hierarchical agents
- Automatic search for bottlenecks based on properties formalizing expert knowledge
 - ▶ MPI wait states
 - ▶ Processor utilization hardware counters
- Clustering of processes/threads with similar properties
- Eclipse-based integrated environment

Supports

- SGI Altix Itanium2, IBM Power and x86-based architectures

Developed by TU Munich

- Released as open-source
- <http://www.lrr.in.tum.de/periscope>



MPI

- Excessive MPI communication time
- Excessive MPI time due to many small messages
- Excessive MPI time in receive due to late sender
- ...

Hardware performance counters (platform-specific)

- Cycles lost due to cache misses
 - ▶ High L1/L2/L3 demand load miss rate
- Cycles lost due to store instructions
- Cycles lost due to address translation misses
- Cycles lost due to no instruction to dispatch
- ...

Periscope plug-in to Eclipse environment

VI-HPS

The screenshot shows the Eclipse IDE with the Periscope plug-in. The interface is divided into several panes:

- Project view:** Located on the left, it shows a hierarchical view of the project files. The file `g_sca_128_install.psc` is selected.
- Source code view:** The central pane displays the source code of the selected file. It shows a Fortran subroutine `field_solve_kxky` with various arguments and local variables.
- SIR outline view:** Located on the right, it provides a high-level overview of the code structure, showing subroutines and their calls. The subroutine `FIELD_SOLVE_KXKY` is highlighted.
- Properties view:** Located at the bottom, it displays a table of performance metrics. The table has columns for Name, Process, Severity, Filename, Confidence, and Extra.

Name	Process	Severity	Filename	Confidence	Extra
Stalls due to waiting for data delivery to register	46	30.22	field_solve_kxky.psc.f90	1.00	
Stalls due to waiting for data delivery to register	5	30.32	field_solve_kxky.psc.f90	1.00	
Stalls due to waiting for data delivery to register	45	30.41	field_solve_kxky.psc.f90	1.00	
L2 misses	102	30.53	field_solve_kxky.psc.f90	1.00	es=221330 L2Misses=164831 L3Misses=
Stalls due to waiting for data delivery to register	17	31.11	field_solve_kxky.psc.f90	1.00	
IA64 Pipeline Stall Cycles	4	31.14	field_solve_kxky.psc.f90	1.00	
IA64 Pipeline Stall Cycles	56	31.38	field_solve_kxky.psc.f90	1.00	
IA64 Pipeline Stall Cycles	50	31.65	field_solve_kxky.psc.f90	1.00	
IA64 Pipeline Stall Cycles	49	31.68	field_solve_kxky.psc.f90	1.00	

Filter: Search: RE 131 Shown - 1 Selected - Sort: [Severity (FWD)]

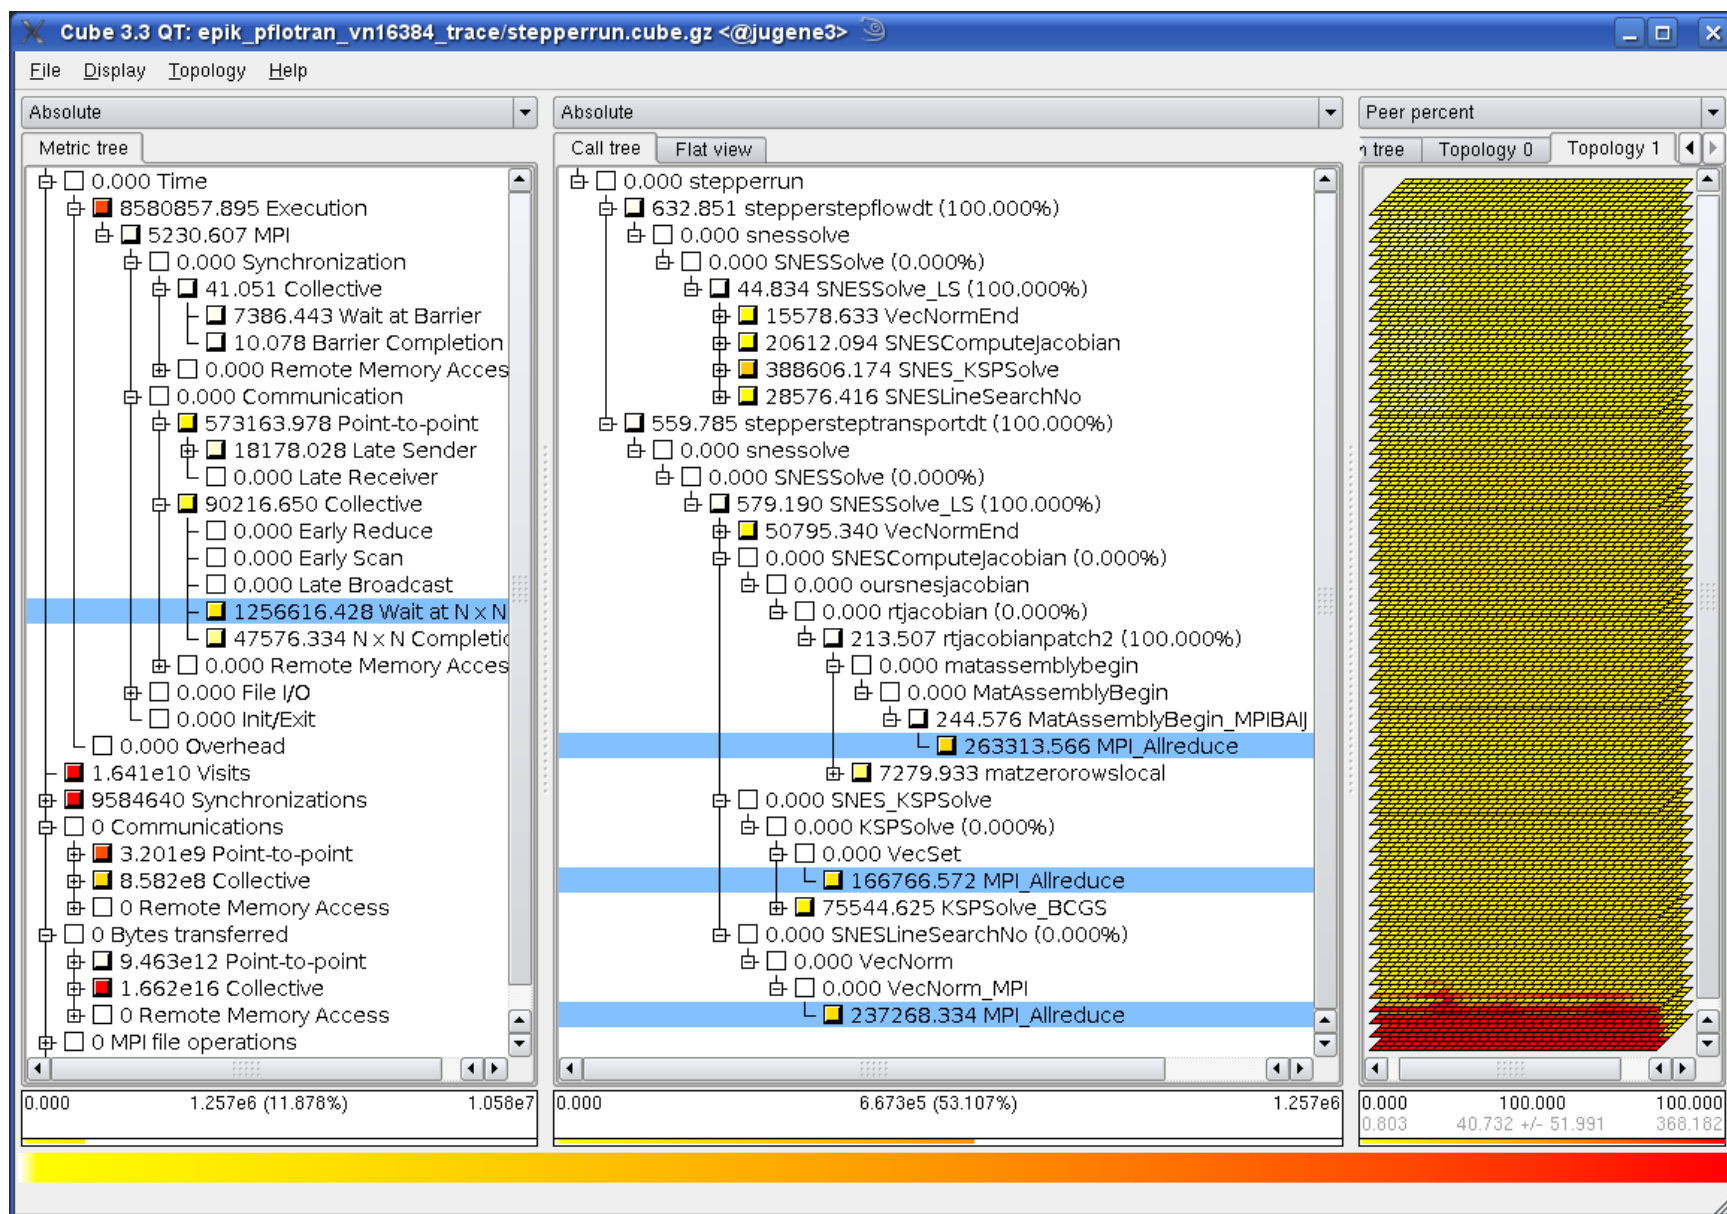
Automatic performance analysis toolset

- Scalable performance analysis of large-scale applications
 - ▶ particularly focused on MPI & OpenMP paradigms
 - ▶ analysis of communication & synchronization overheads
- Automatic and manual instrumentation capabilities
- Runtime summarization and/or event trace analyses
- Automatic search of event traces for patterns of inefficiency
 - ▶ Scalable trace analysis based on parallel replay
- Interactive exploration GUI and algebra utilities for XML callpath profile analysis reports

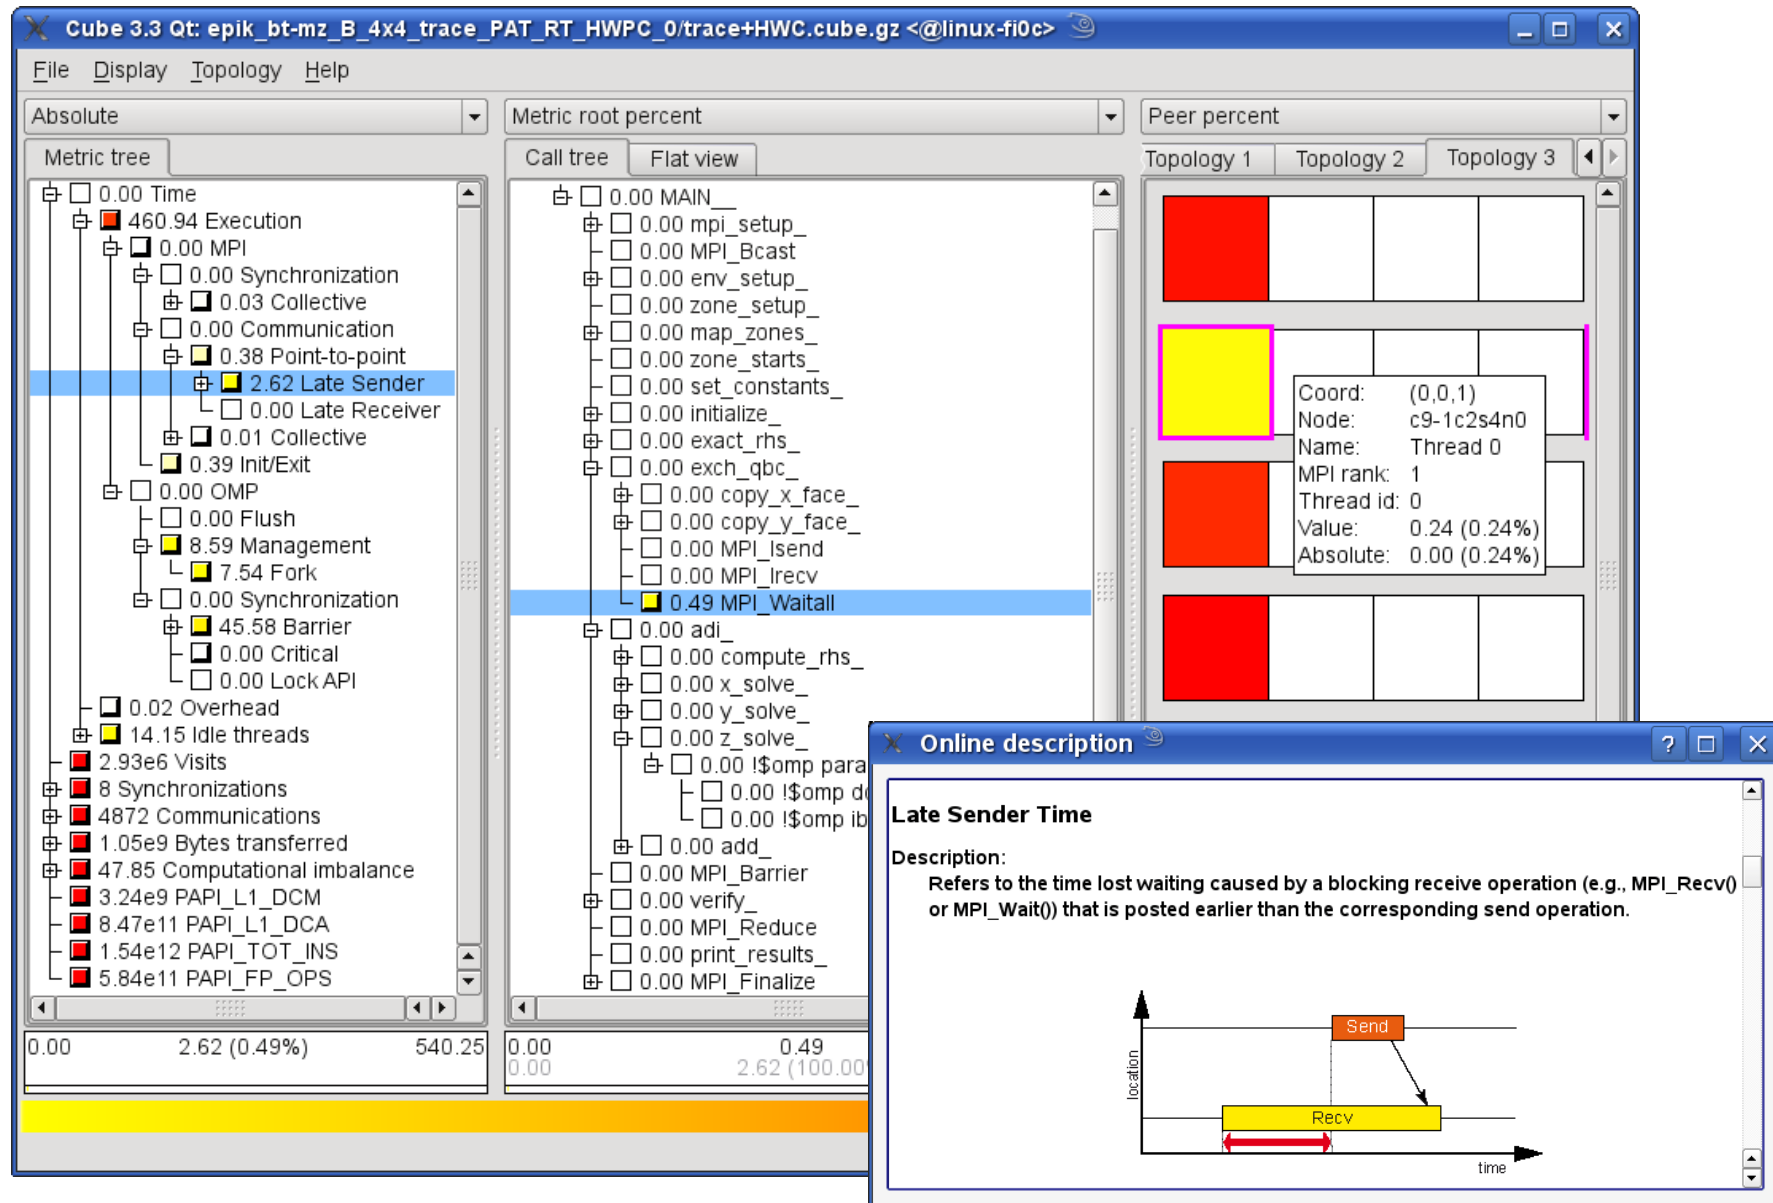
Developed by JSC & GRS

- Released as open-source
- <http://www.scalasca.org/>

Scalasca automatic trace analysis report



VÍ-HPS




Integrated performance toolkit

- Instrumentation, measurement, analysis & visualization
 - ▶ Highly customizable installation, API, envvars & GUI
 - ▶ Supports multiple profiling & tracing capabilities
- Performance data management & data mining
- Targets all parallel programming/execution paradigms
 - ▶ Ported to a wide range of computer systems
- Performance problem solving framework for HPC
- Extensive bridges to/from other performance tools
 - ▶ PerfSuite, Scalasca, Vampir, ...

Developed by U. Oregon/PRL

- Broadly deployed open-source software
- <http://tau.uoregon.edu/>



V-HPS

The diagram illustrates the TAU Architecture, which is organized into three main horizontal layers: **Instrumentation**, **Measurement**, and **Analysis**.

- Instrumentation Layer (Blue):** This layer handles the initial setup. It includes components for **source code**, **object code**, **library wrapper**, **binary code**, and **virtual machine**. An **event selection** process feeds into the source code, and **event information** is output from the virtual machine. All these components interface with the **MEASUREMENT API**.
- Measurement Layer (Green):** This layer is responsible for data collection and management. It is divided into:
 - Event creation and management:** Includes event identifier, entry/exit events, atomic events, event mapping, and event control.
 - Profiling:** Includes statistics, atomic profiles, entry/exit profiles, phase profiles, I/O profiles, and profile sampling.
 - Tracing:** Includes trace buffering, record creation, trace I/O, timestamp generation, trace filtering, and trace merging.
 - Performance data sources:** Includes timing, hardware counters, system counters, and kernel.
 - OS and runtime system modules:** Includes threading, interrupts, runtime system, and I/O.
- Analysis Layer (Yellow/Gold):** This layer processes the collected data. It is divided into:
 - Profile Data Management (PerfDMF):** Includes profile translators, Metadata (XML), and a profile database.
 - Trace Data Management:** Includes trace translators and trace storage.
 - Profile Analysis (ParaProf):** Displays various visualization charts for profile analysis.
 - Profile Data Mining (PerfExplorer):** Displays charts for mining profile data.
 - Trace Visualizers:** Includes Vampir, JumpShot, and Paraver.
 - Trace Analyzers:** Includes Expert, ProfileGen, and Vampir Server.

Data Flow: The flow starts with **event selection** entering the **Instrumentation** layer. Data moves through the **Measurement** layer to the **Analysis** layer. Within the Analysis layer, **Instrumentation** and **Measurement** components interact. **event information** is passed from the Measurement layer to the Analysis layer. **profiles** and **traces** are generated and then processed by the respective management and analysis modules. A **symbol table** is also utilized in the analysis phase.

Program Analysis

Application / Library

C / C++ parser

Fortran parser F77/90/95

IL

IL

C / C++ IL analyzer

Fortran IL analyzer

Program Database Files

DUCTAPE

PDBhtml

Program documentation

SILOON

Application component glue

CHASM

C++ / F90/95 interoperability

Data Instrumentation

Automatic source instrumentation

Parallel Profile Analysis

ParaDMF

Performance Analysis Programs

ParaProf

TAU Performance System

Query and Analysis Toolkit

Java PerfDMF API

SQL (PostgreSQL, MySQL, DB2, Oracle)

Performance Data

PerfDMF

ParaProf

Call Graphs

Histograms

Call Trees

Bar Charts

Comparative Displays

Text Displays

Vis Package

3D Displays

Scripting Interface

Jython

TAU, mpiP, ompP, HPMToolKit, Cube, HPCIToolkit, gprof, Dynaprof, PSRun

Runtime Data Collection

Supermon, MRNet

DBMS

PostgreSQL, MySQL, Oracle, DB2, Derby

scalability analysis

ParaProf

cluster analysis

Data Mining (Weka)

Statistics (R / Omega)

raw profiles

profile metadata

*** gprof**

*** mpiP**

*** psrun**

*** HPMtoolkit**

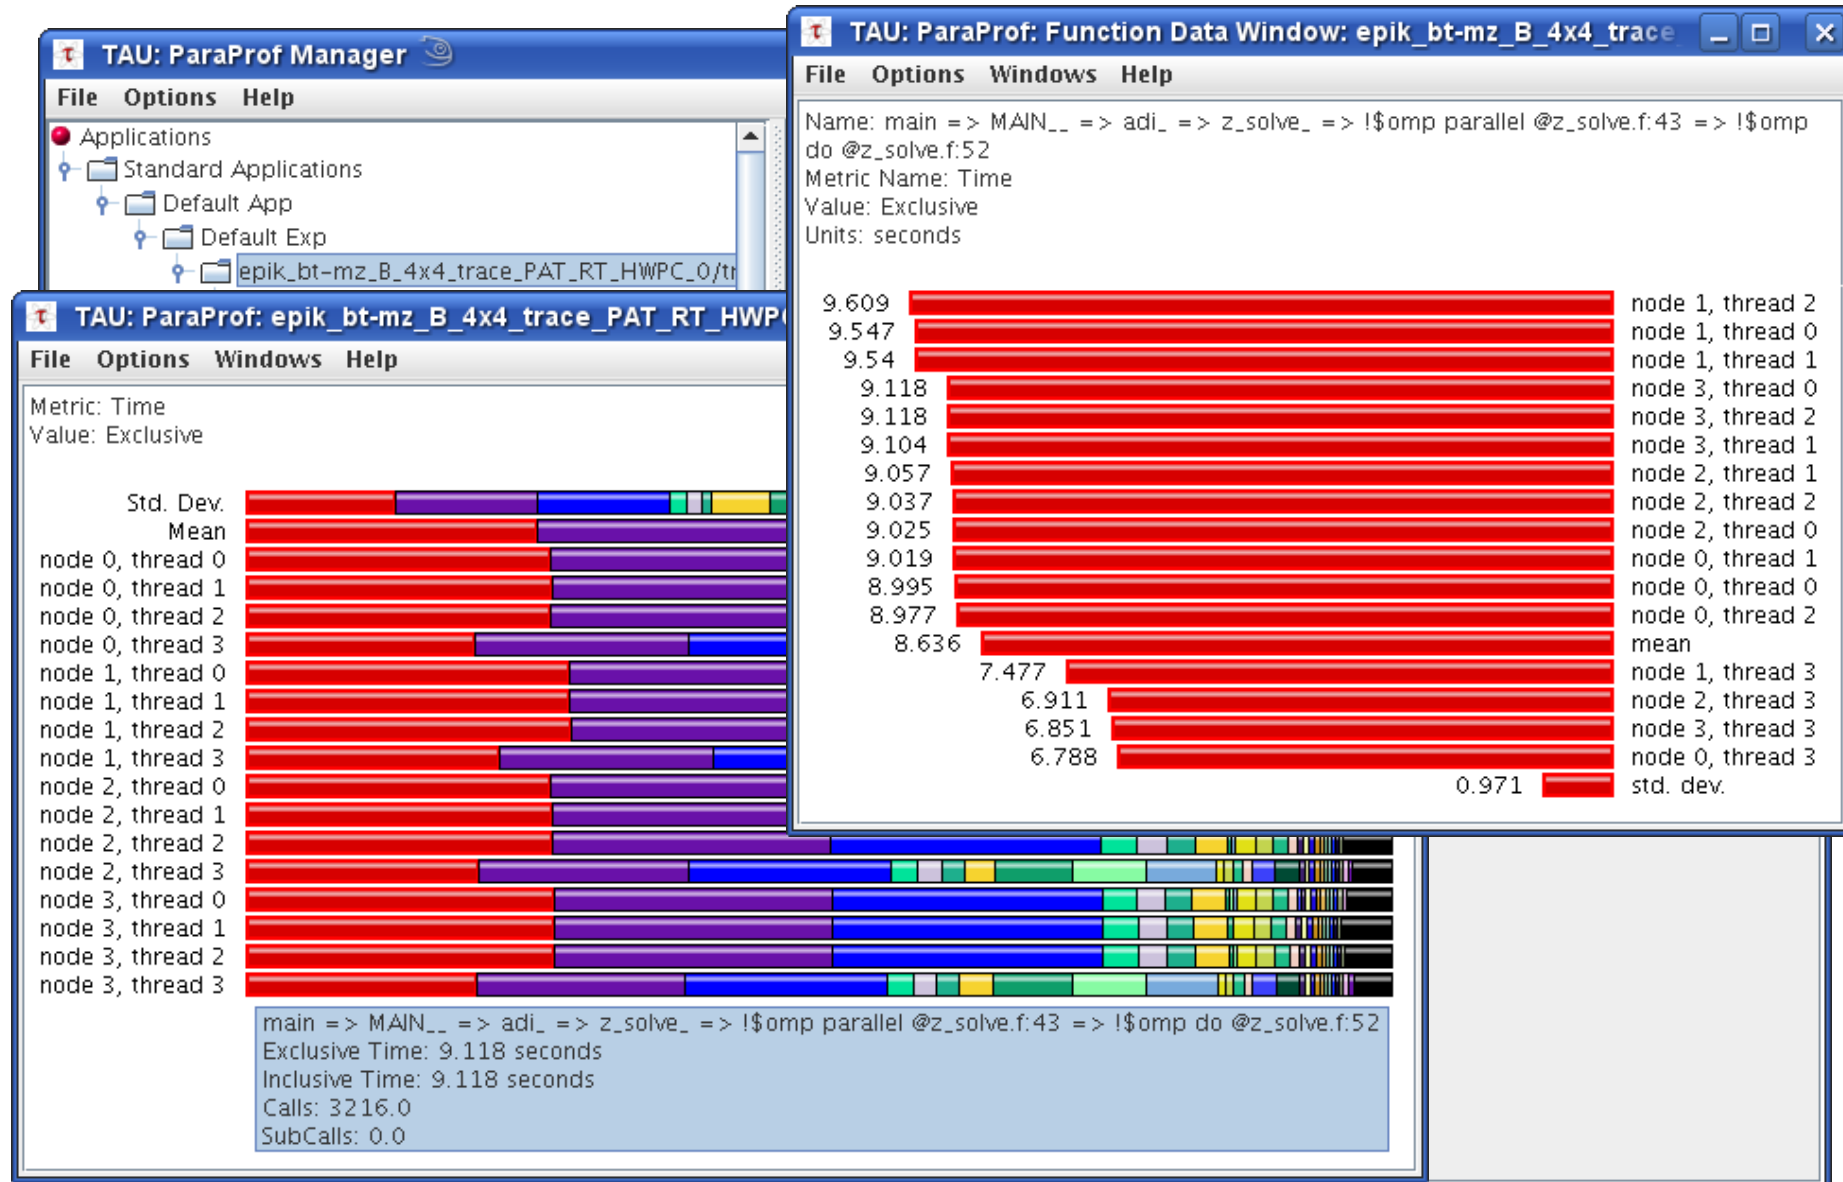
...

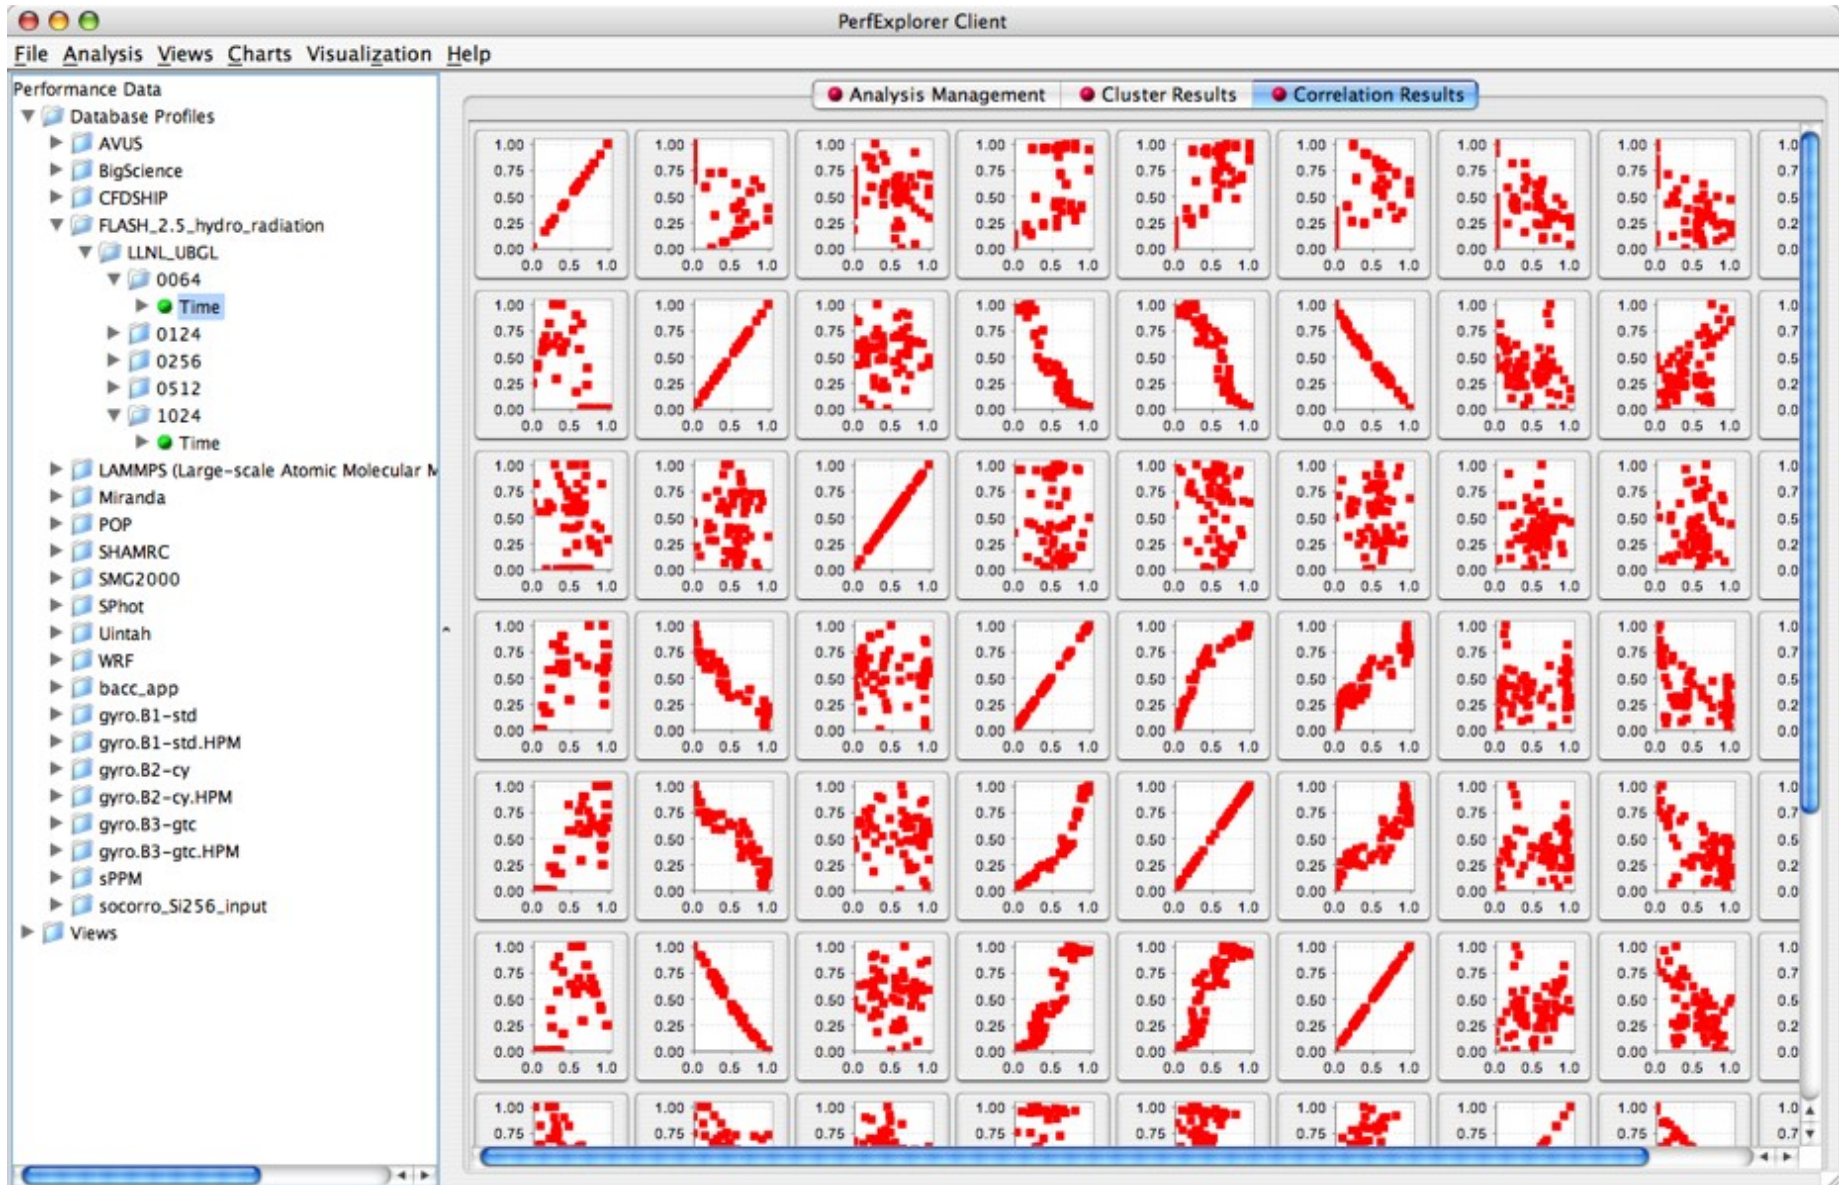
XML document

formatted profile data

[illegible]

TAU ParaProf GUI displays (selected)





Interactive event trace analysis

- Alternative & supplement to automatic trace analysis
- Visual presentation of dynamic runtime behaviour
 - ▶ event timeline chart for states & interactions of processes/threads
 - ▶ communication statistics, summaries & more
- Interactive browsing, zooming, selecting
 - ▶ linked displays & statistics adapt to selected time interval (zoom)
 - ▶ scalable server runs in parallel to handle larger traces


Developed by TU Dresden ZIH

- Open-source VampirTrace library bundled with OpenMPI 1.3
- <http://www.tu-dresden.de/zih/vampirtrace/>
- Vampir Server & GUI have a commercial license
- <http://www.vampir.eu/>



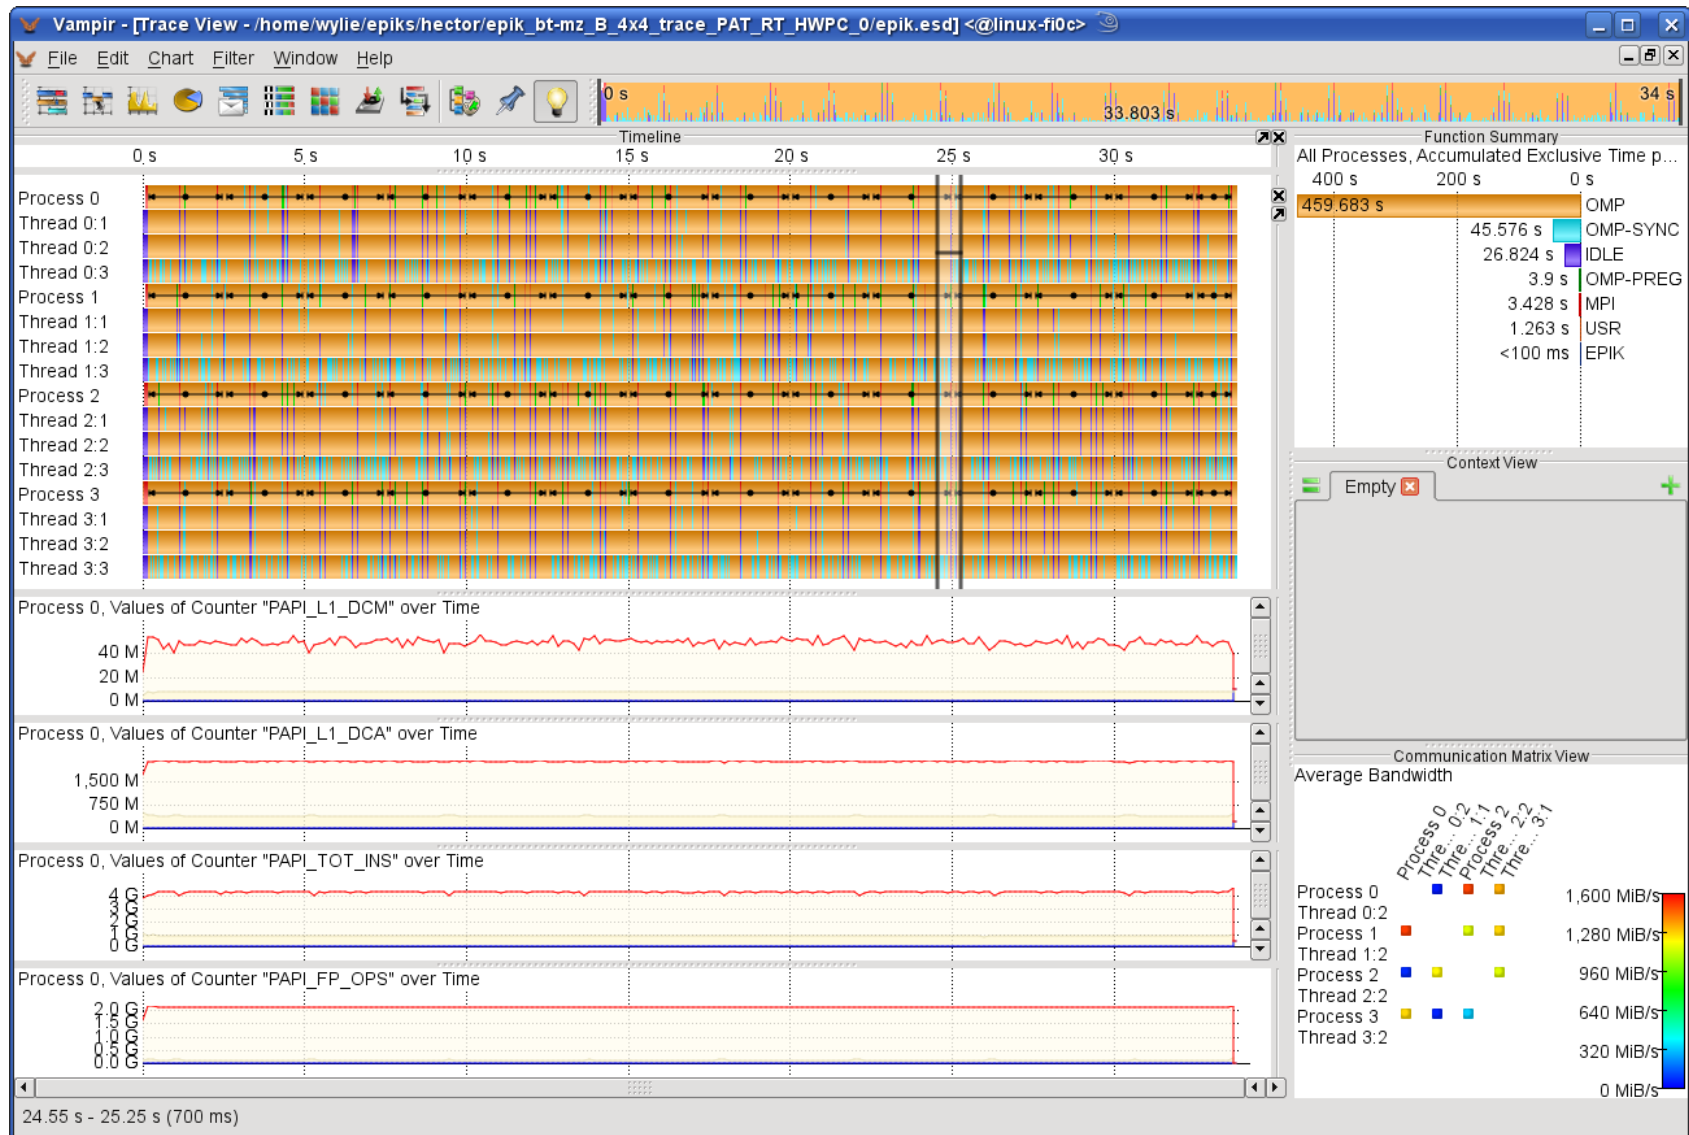
Vampir interactive trace analysis GUI

VI-HPS




Vampir interactive trace analysis GUI

VI-HPS



Vampir interactive trace analysis GUI (zoom)

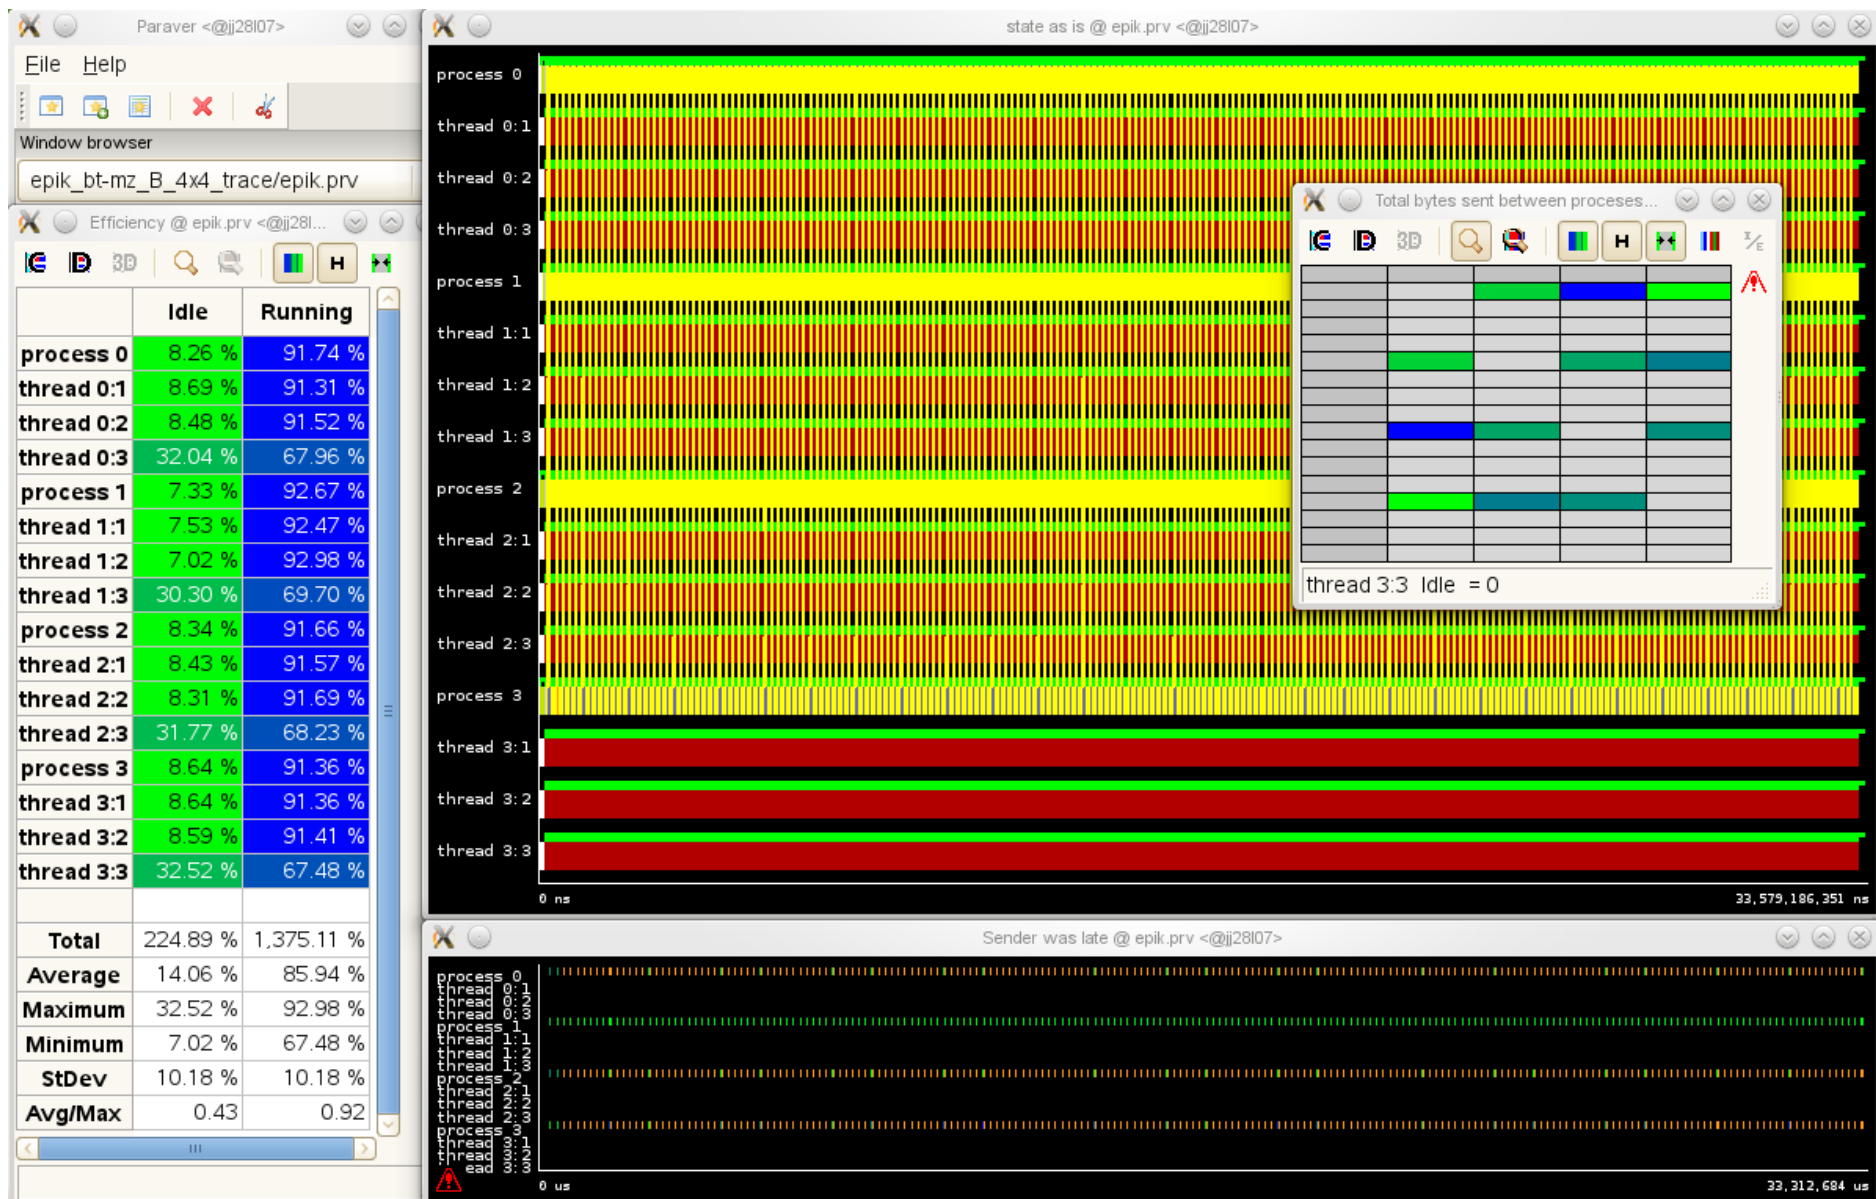
VI-HPS



- Interactive event trace analysis
 - Visual presentation of dynamic runtime behaviour
 - ▶ event timeline chart for states & interactions of processes
 - ▶ Interactive browsing, zooming, selecting
 - Large variety of highly configurable analyses & displays
- Developed by Barcelona Supercomputing Center
 - Paraver trace analyser and Extrae measurement library
 - Open source available from <http://www.bsc.es/paraver/>

Paraver interactive trace analysis GUI

VI-HPS



Key tool components also provided as open-source

- Program development environment
 - ▶ Eclipse PTP ETFw, [UNITE](#)
- Program/library instrumentation
 - ▶ COBI, OPARI, PDTToolkit
- Runtime measurement systems
 - ▶ [Score-P](#), UniMCI
- Scalable I/O
 - ▶ [SIONlib](#)
- Libraries & tools for handling (and converting) traces
 - ▶ EPILOG, PEARL, OTF
- Analysis algebra & hierarchical/topological presentation
 - ▶ CUBE

Scalable performance measurement infrastructure

- Supports instrumentation, profiling & trace collection, as well as online analysis of HPC parallel applications
- Works with Periscope, Scalasca, TAU & Vampir prototypes
- Based on updated tool components
 - ▶ CUBE4 profile data utilities & GUI
 - ▶ OA online access interface to performance measurements
 - ▶ OPARI2 OpenMP & pragma instrumenter
 - ▶ OTF2 open trace format

Created by German BMBF SILC & US DOE PRIMA projects

- JSC, RWTH, TUD, TUM, GNS, GRS, GWT & UO PRL
- Available as BSD open-source from <http://www.score-p.org/>

Portable native parallel I/O library & utilities

- Scalable massively-parallel I/O to task-local files
- Manages single or multiple physical files on disk
 - ▶ optimizes bandwidth available from I/O servers by matching blocksizes/alignment, reduces metadata-server contention
- POSIX-I/O-compatible sequential & parallel API
 - ▶ adoption requires minimal source-code changes
- Tuned for common parallel filesystems
 - ▶ GPFS (BlueGene), Lustre (Cray), ...
- Convenient for application I/O, checkpointing,
 - ▶ Used by Scalasca tracing (when configured)

Developed by JSC

- Available as open-source from
<http://www.fz-juelich.de/jsc/sionlib/>

Uniform integrated tool environment

- Manages installation & access to program development tools
 - ▶ based on software environment management “modules”
 - ▶ commonly used on most cluster and HPC systems
 - ▶ configurable for multiple MPI libraries & compiler suites
- Specifies how & where tools packages get installed
 - ▶ including integrating tools where possible
- Defines standard module names and different versions
- Supplies pre-defined module files
- Configurable to co-exist with local installations & policies

Developed by JSC, RWTH & TUD

- Available as open-source from
<http://www.vi-hps.org/projects/unite/>