


TAU PerfExplorer

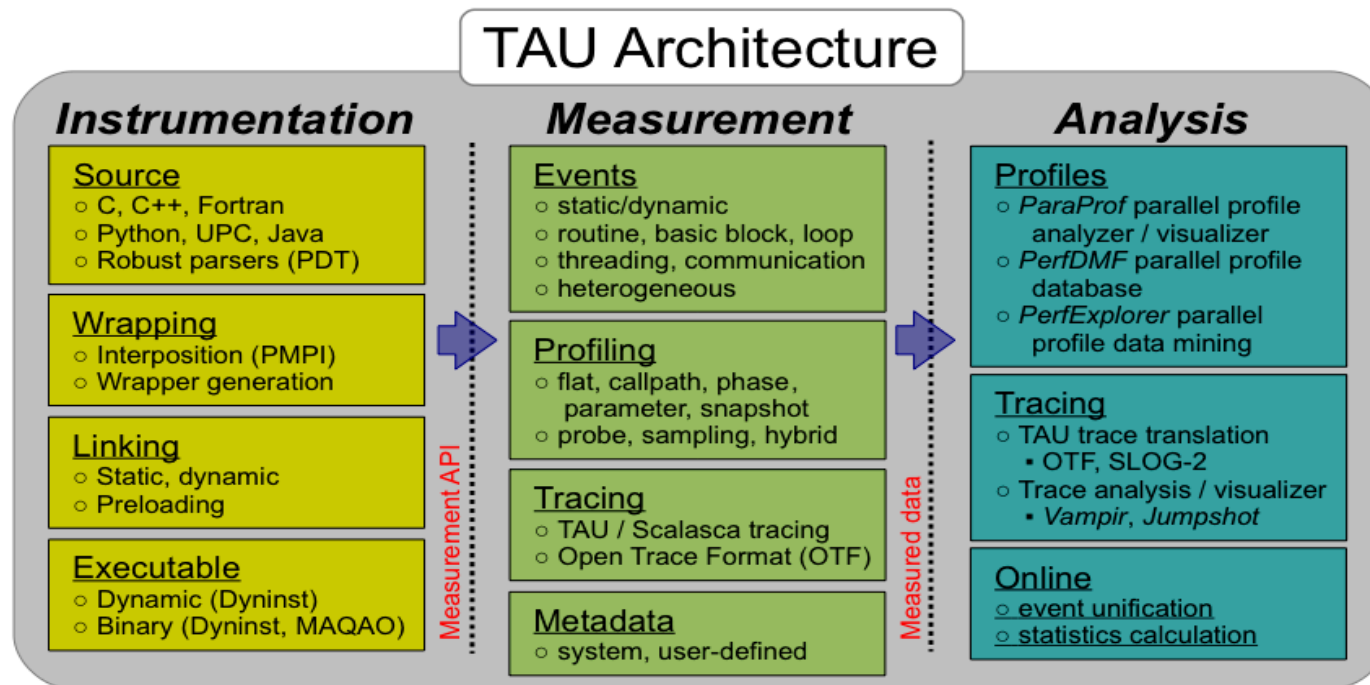


Sameer Shende
sameer@cs.uoregon.edu
 University of Oregon
<http://tau.uoregon.edu>




TAU Performance System®

- Parallel performance framework and toolkit
 - Supports all HPC platforms, compilers, runtime system
 - Provides portable instrumentation, measurement, analysis

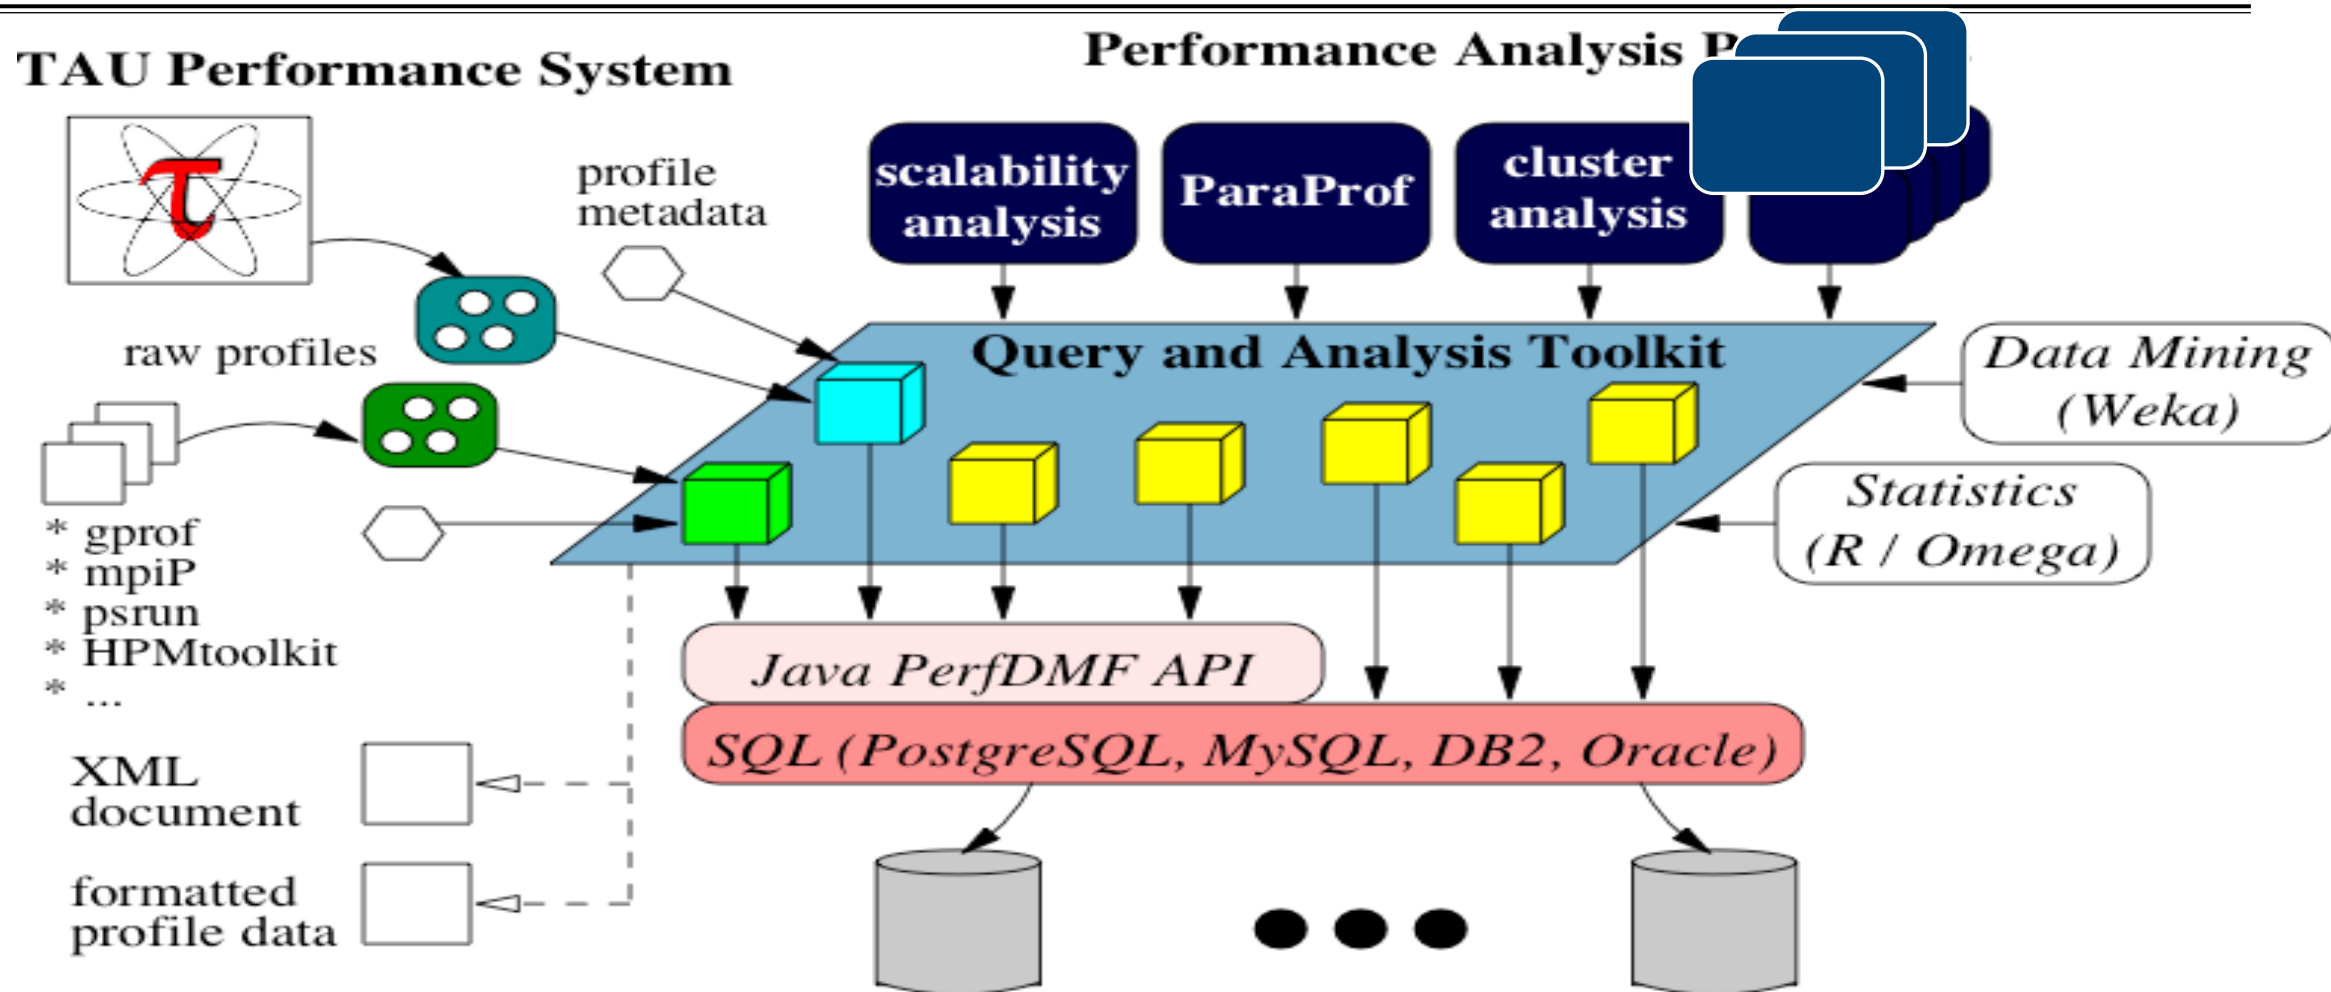


TAU's Analysis Tools: PerfExplorer

TAU Analysis



TAUdb: Performance Data Management Framework



Using TAUdb

- **Configure TAUdb (Done by each user)**

- % `taudb_configure --create-default`

- Choose derby, PostgreSQL, MySQL, Oracle or DB2
 - Hostname
 - Username
 - Password
 - Say yes to downloading required drivers (we are not allowed to distribute these)
 - Stores parameters in your `~/.ParaProf/taudb.cfg` file

- **Configure PerfExplorer (Done by each user)**

- % `perfexplorer_configure`

- **Execute PerfExplorer**

- % `perfexplorer`

Local Installation (*Archer*)

- Setup preferred program environment compilers

```
% source /home/y14/y14/shende/tau.bashrc
% cp /home/y14/y14/shende/data.tgz . ; tar zxf data.tgz
% cd data
% cat README
  and follow the steps
% cd tau
% ./upload.sh
% perfexplorer
```

Local Installation (*Isambard*)

- Setup preferred program environment compilers

```
% source /home/ri-sshende/tau.bashrc
% cp /home/ri-sshende/data.tgz . ; tar xzf data.tgz
% cd data
% cat README
  and follow the steps
% cd tau
% ./upload.sh
% perfexplorer
```



Performance Data Mining (PerfExplorer)

- Performance knowledge discovery framework
 - Data mining analysis applied to parallel performance data
 - comparative, clustering, correlation, dimension reduction, ...
 - Use the existing TAU infrastructure
 - TAU performance profiles, taudb
 - Client-server based system architecture
- Technology integration
 - Java API and toolkit for portability
 - taudb
 - R-project/Omegahat, Octave/Matlab statistical analysis
 - WEKA data mining package
 - JFreeChart for visualization, vector output (EPS, SVG)

PerfExplorer: Using Cluster Analysis

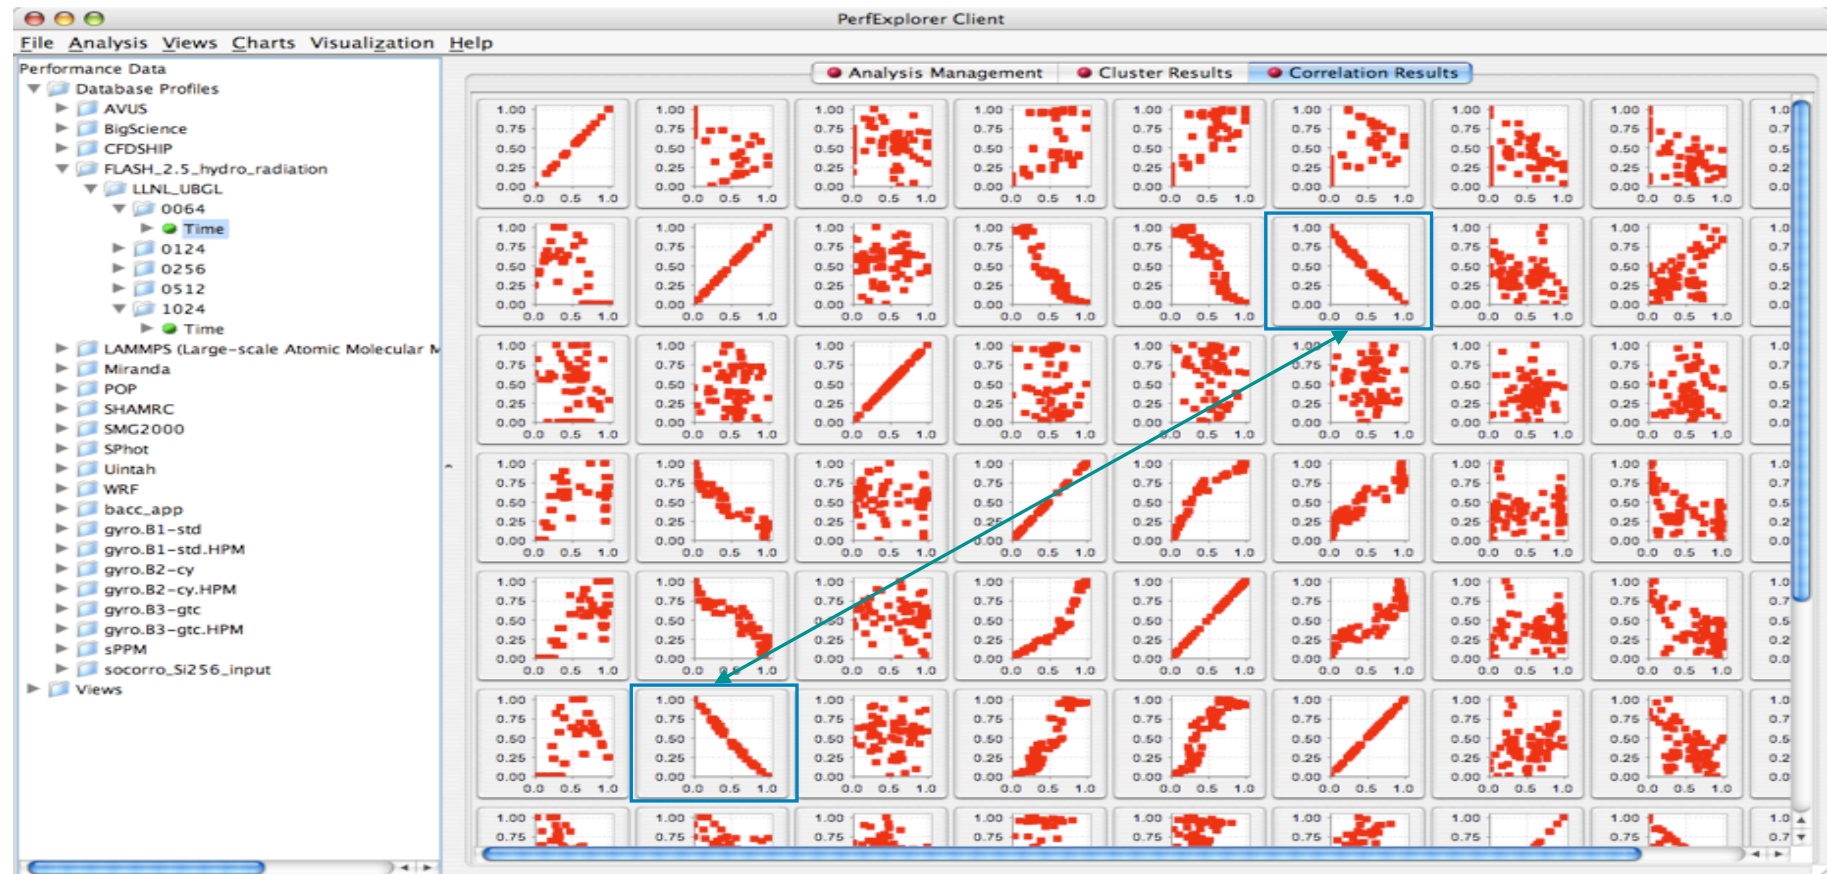
- Performance data represented as vectors - each dimension is the cumulative time for an event
- *k*-means: *k* random centers are selected and instances are grouped with the "closest" (Euclidean) center
- New centers are calculated and the process repeated until stabilization or max iterations
- Dimension reduction necessary for meaningful results
- Virtual topology, summaries constructed

PerfExplorer - Cluster Analysis (sPPM)




PerfExplorer - Correlation Analysis (Flash)

- Describes strength and direction of a linear relationship between two variables (events) in the data



PerfExplorer - Correlation Analysis (Flash)

- -0.995 indicates strong, negative relationship
- As CALC_CUT_BLOCK_CONTRIBUTIONS() increases in execution time, MPI_Barrier() decreases



PerfExplorer - Comparative Analysis

- Relative speedup, efficiency
 - total runtime, by event, one event, by phase
- Breakdown of total runtime
- Group fraction of total runtime
- Correlating events to total runtime
- Timesteps per second

PerfExplorer - Interface

The screenshot shows the PerfExplorer Client interface. On the left is a file tree under 'gyro.B1-std'. The right pane is titled 'Analysis Management' and 'Performance Explorer'. It contains a table with 'Field' and 'Value' columns. The table lists various system and experiment metadata fields. Three red callout boxes provide context: one points to the file tree, another to the metadata table, and a third to the data structure.

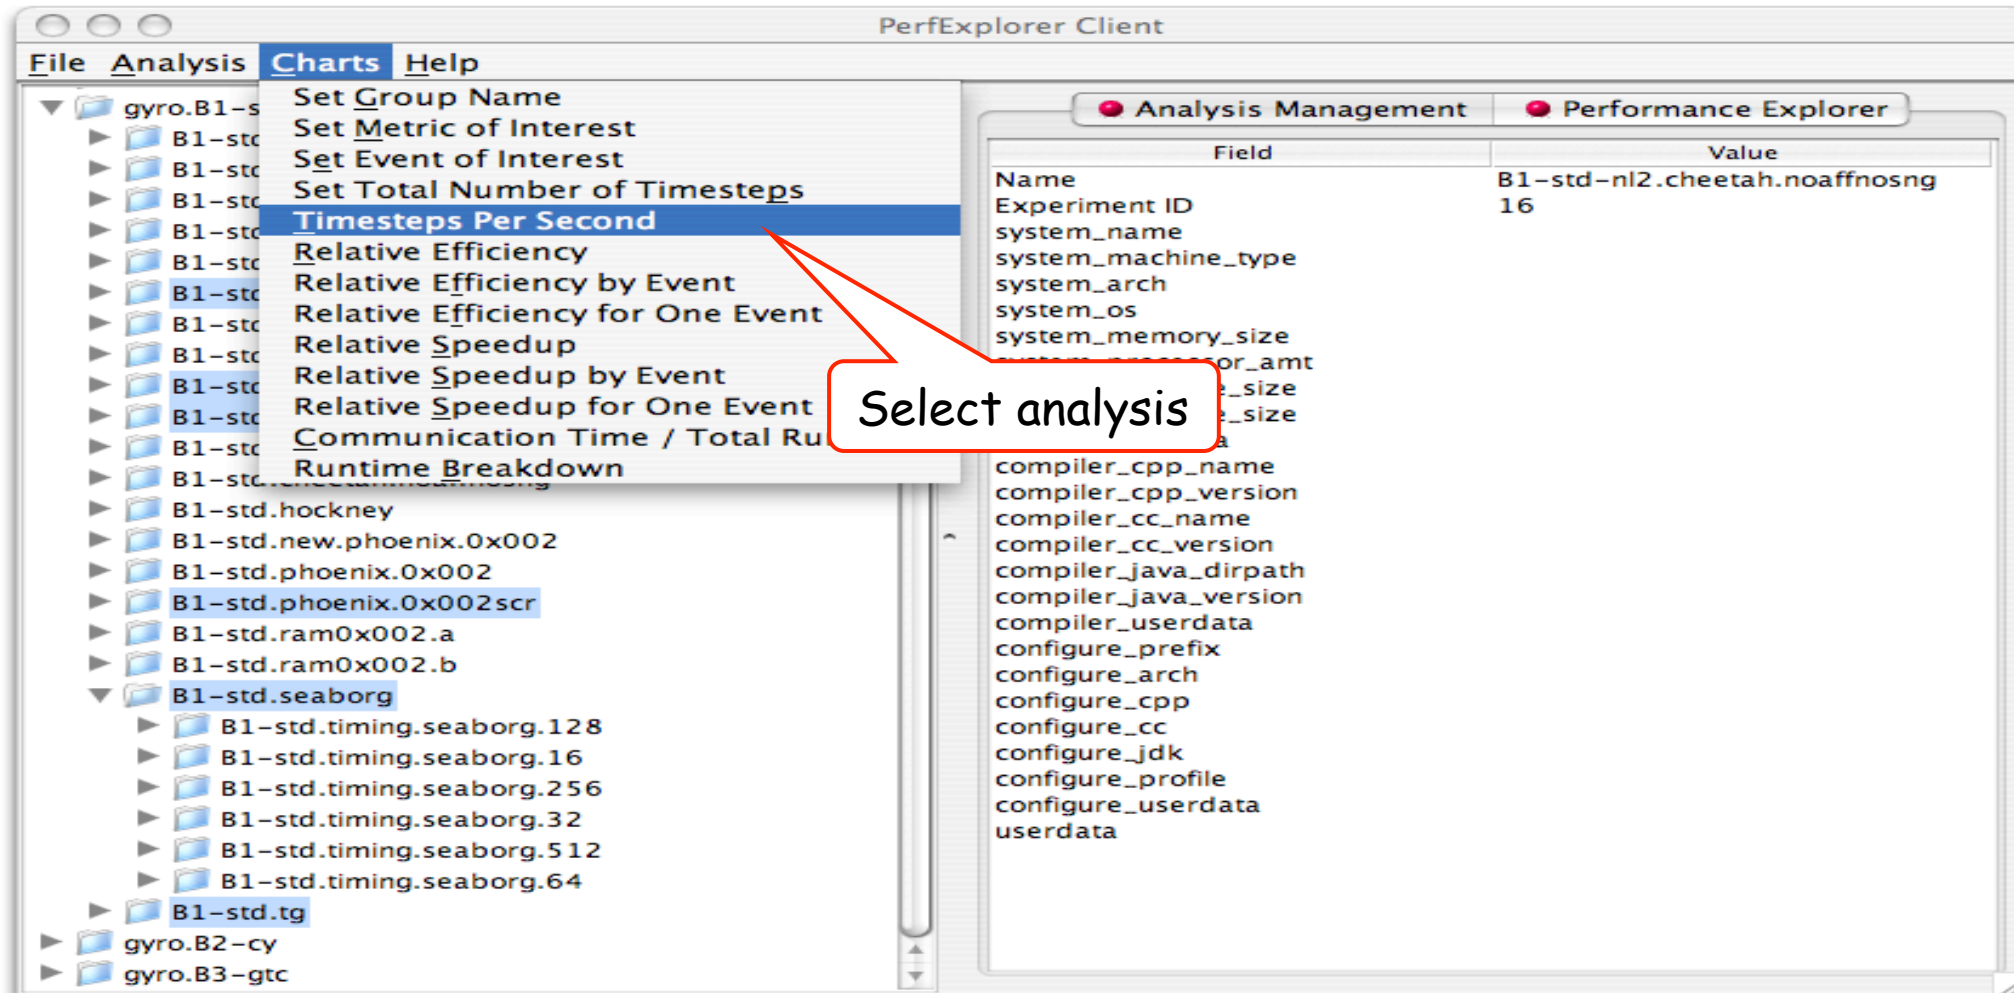
Select experiments and trials of interest

Experiment metadata


Data organized in application, experiment, trial structure (will allow arbitrary in future)

Field	Value
Name	B1-std-nl2.cheetah.noaffnosng
Experiment ID	16
system_name	
system_machine_type	
system_arch	
system_os	
system_memory_size	
system_processor_amt	
system_l1_cache_size	
system_l2_cache_size	
userdata	
cpp_name	
cpp_version	
cc_name	
cc_version	
compiler_java_dirpath	
compiler_java_version	
compiler_userdata	
configure_prefix	
configure_arch	


PerfExplorer - Interface




PerfExplorer - Relative Efficiency Plots




PerfExplorer - Relative Efficiency by Routine



PerfExplorer - Relative Speedup





PerfExplorer - Timesteps Per Second




Evaluate Scalability


- Goal: How does my application scale? What bottlenecks occur at what core counts?
- Load profiles in taudb database and examine with PerfExplorer




Evaluate Scalability




PerfExplorer



PerfExplorer



Performance Regression Testing



Performance Research Lab, University of Oregon, Eugene, USA



Support Acknowledgments


- US Department of Energy (DOE)
 - Office of Science contracts
 - SciDAC, LBL contracts
 - LLNL-LANL-SNL ASC/NNSA contract
 - Battelle, PNNL contract
 - ANL, ORNL contract
- Department of Defense (DoD)
 - PETTT, HPCMP
- National Science Foundation (NSF)
 - Glassbox, SI-2
- NASA
- CEA, France
- Partners:
 - University of Oregon
 - ParaTools, Inc., ParaTools, SAS
 - The Ohio State University
 - University of Tennessee, Knoxville
 - T.U. Dresden, GWT
 - Juelich Supercomputing Center



Acknowledgement

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation's exascale computing imperative.

Download TAU from U. Oregon



<http://tau.uoregon.edu>

<http://www.hpclinux.com> [LiveDVD, OVA]

<https://e4s.io> [Containers for Extreme-Scale Scientific Software Stack]

Free download, open source, BSD license