



The Performance Evaluation Research Center (PERC)

Participating Institutions:

Argonne Natl. Lab.

Univ. of California, San Diego

Lawrence Berkeley Natl. Lab.

Univ. of Illinois

Lawrence Livermore Natl. Lab.

Univ. of Maryland

Oak Ridge Natl. Lab.

Univ. of Tennessee, Knoxville

Website: <http://perc.nersc.gov>



PERC Overview

- ✍ An “Integrated Software Infrastructure Center” (ISIC) sponsored under DoE’s SciDAC program.
- ✍ Funding: approx. \$2.4 million per year.
- ✍ Mission:
 - ✍ Develop a science of performance.
 - ✍ Engineer tools for performance analysis and optimization.
- ✍ Focus:
 - ✍ Large, grand-challenge calculations, especially SciDAC application projects.



Specific Objectives

- ✍ Understand key factors in applications that affect performance.
- ✍ Understand key factors in computer systems that affect performance.
- ✍ Develop models that accurately predict performance of applications on systems.
- ✍ Develop an enabling infrastructure of tools for performance monitoring, modeling and optimization.
- ✍ Validate these ideas and infrastructure via close collaboration with DOE Office of Science and others.
- ✍ Transfer the technology to end users.



Anticipated Benefits

Consider the economic value of improving the performance of a single high-end scientific application code by 20%.

Assume:

- ✍ \$10 million computer system lease cost per year.
- ✍ \$10 million per year in site costs, support staff, etc.
- ✍ 10-year lifetime of code.
- ✍ Code uses 5% of system cycles each year.

Savings: \$2,000,000.

Scientific benefit (additional computer runs and research) is probably much higher.



Anticipated Benefits, cont.

- ✍ We rely heavily on commercial vendors for high-performance computer systems.
- ✍ We are invited by vendors to provide guidance on the design of current and future systems.

BUT

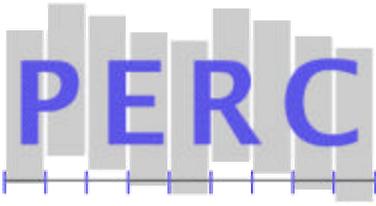
- ✍ At present we can provide only vague information – little if any quantitative data or rigorous analysis.

The performance monitoring and modeling capability to be developed in PERC will significantly improve our ability to influence future scientific computer systems.



Current State-of-the-Art

- ✍ Some tools collect performance data, but
 - ✍ They are not targeted to large parallel systems.
 - ✍ They are not able to collect performance data at individual levels of deep memory hierarchies.
- ✍ A few performance modeling techniques have been developed, but
 - ✍ They are time-consuming to generate, difficult to use, or have limited accuracy.
- ✍ Some automatic tuning techniques have been developed, but
 - ✍ They have been applied only in limited algorithm domains.
 - ✍ There is no hardened support for real-time optimization.



New Capabilities

Better Benchmarks:

-  Polished, concise versions of real user codes, representing strategic application areas.
-  Kernel benchmarks extracted from real codes reduce complexity of analyzing full-size benchmarks.
-  Low-level benchmarks measure key rates of data access at various levels of memory hierarchy.

Modern performance monitoring tools:

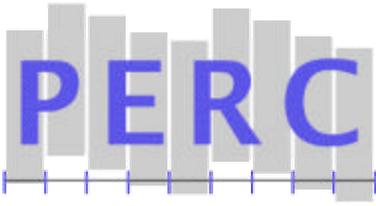
-  Flexible instrumentation systems capture hardware and software interactions, instruction execution frequencies, memory reference behavior, and execution overheads.
-  An advanced data management infrastructure tracks performance experiments and data across time and space.



New Capabilities, cont.

Performance modeling:

-  *Application signature* tools characterize applications independent of the machine where they execute.
-  *Machine signature* tools characterize computer systems, independent of the applications.
-  *Convolution* tools combine application and machine signatures to provide accurate performance models.
-  *Statistical models* find approximate performance models based on easily measured performance data.
-  *Phase models* analyze separate sections of an application, providing overall performance predictions.
-  *Performance bound* tools determine ultimate potential of an application on a given system.

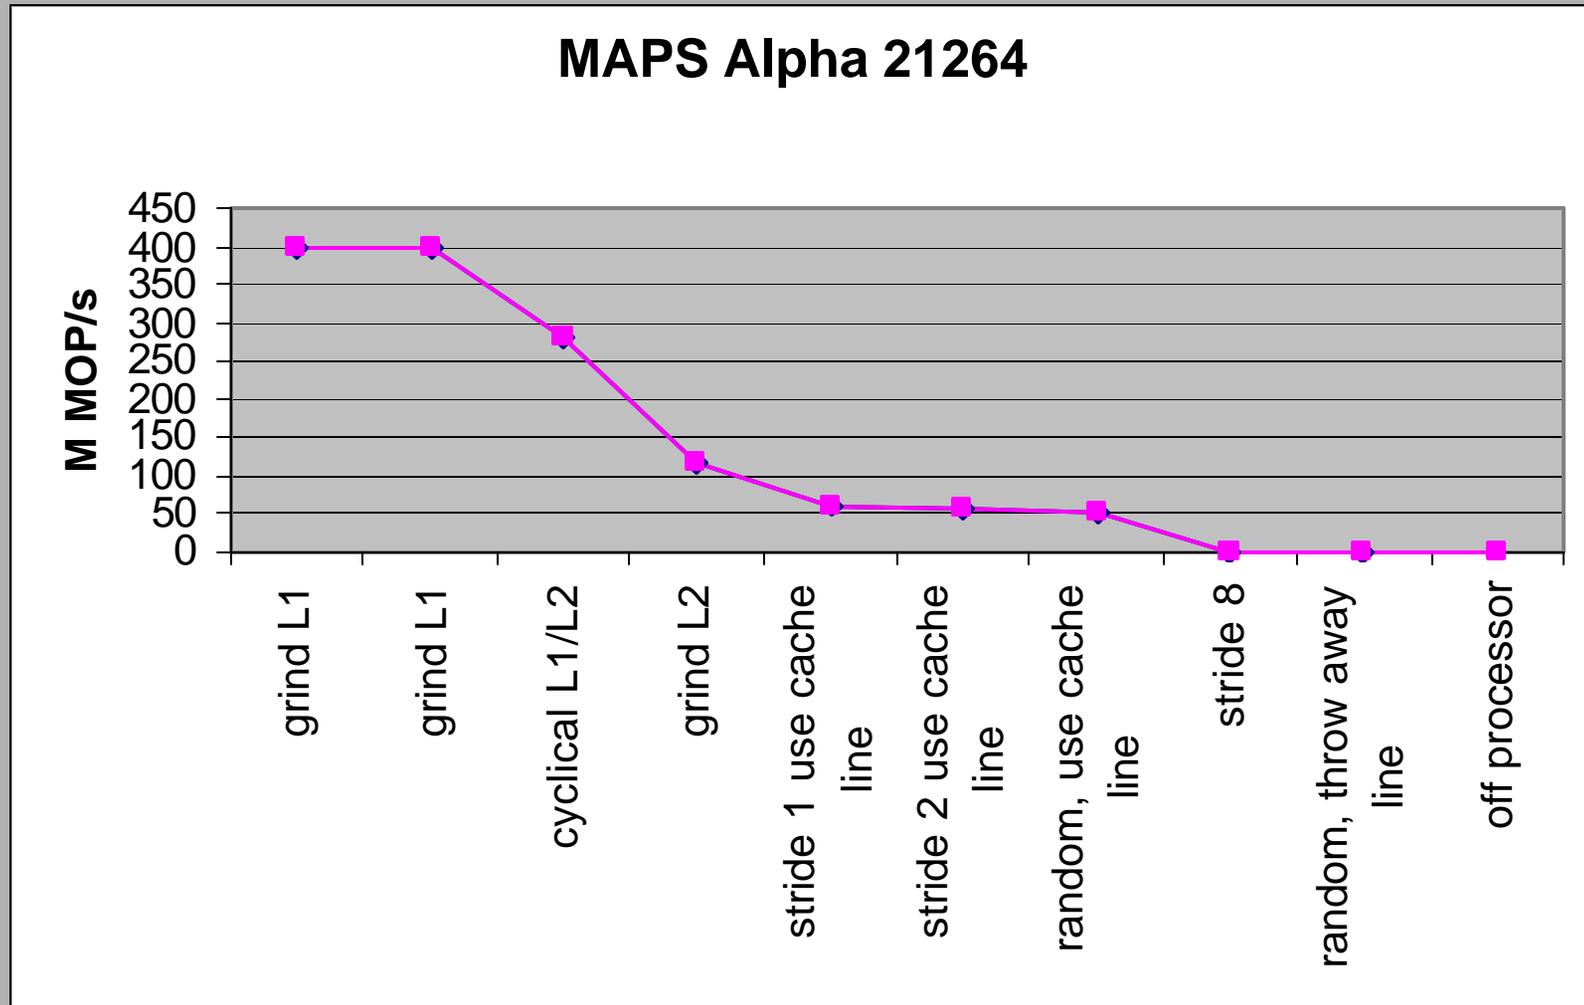


New Capabilities, cont.

Performance optimization:

-  *Compile-time optimization* mechanisms analyze source code to improve performance.
-  *Self-tuning software* automatically tunes code based on real-time measurements of hardware environment.
-  *Performance assertions* permit user-specified run-time tests to possibly change the course of the computation depending on results.
-  *Performance portability* programming techniques to insure that code runs at near-optimal performance across a variety of modern systems.

Measured Memory Access Patterns





SvPablo Graphical Interface

The screenshot displays the SvPablo graphical interface. The main window is titled "svPablo" and contains several panels:

- Project Description:** PCTM on IBM-SP (seaborg)
- Source Files:** fcd.F, fcd_setup.F
- Performance Contexts:** IBM-SP, 64 procs, other Metrics; IBM-SP, 64 procs, other Metrics, 10 days; IBM-SP, 32 procs, other Metrics, 10 days; IBM-SP, 16 procs, other Metrics, 10 days; IBM-SP, 8 procs, other Metrics, 10 days; IBM-SP, 4 procs, other Metrics, 10 days
- Routines in Source File:** fcd, fcd_setup, fcd_timer_clear, fcd_timer_start, fcd_init
- Routines in Performance Data:** ccm3, oceanstep, icestep, mpi_c, mpi_c
- Source File:** /u0/cmendes/PCTM/pcm2/src/sources/fcd.F

The source code is displayed in the center, with a grid of performance data on the left. Several "Specific Metric" dialog boxes are overlaid on the code, showing hardware statistics by line for various instructions:

- Call Statistics count: 240.0000 -- ccm3
- Call Statistics Duration: 211.2399 -- ccm3
- HW Statistics by Line Floating Point Instructions: 12295122047.0000 -- ccm3
- HW Statistics by Line Load Misses in D1: 300348320.0000 -- ccm3
- HW Statistics by Line Branch Instructions: 6944481284.0000 -- ccm3
- HW Statistics by Line TLB misses: 151118598.0000 -- ccm3

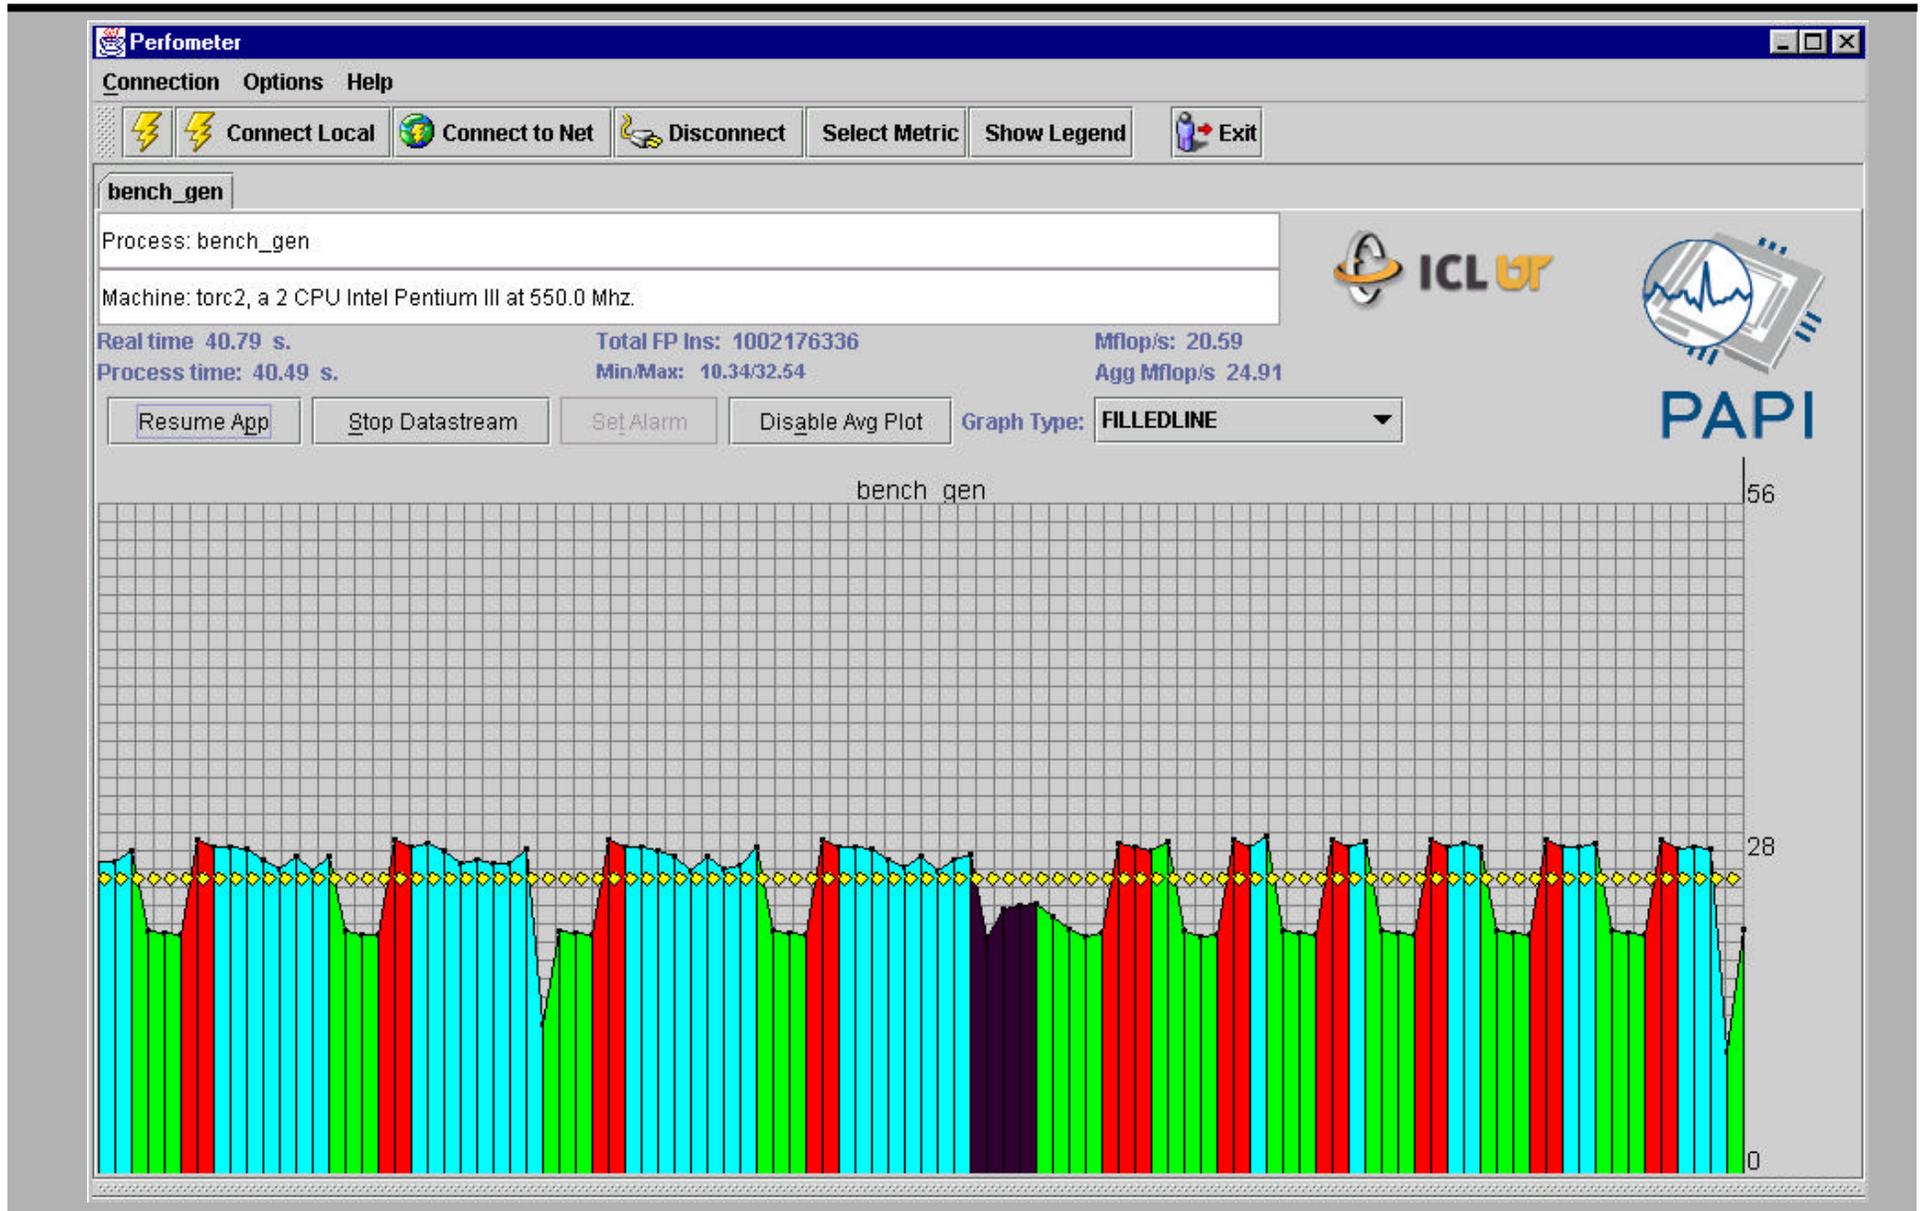
On the right, a "Legend: Source Code Metrics" panel lists ten columns of metrics with corresponding color-coded bars:

- Column 1: Call Statistics count (240, 10)
- Column 2: Call Statistics Duration (211.24, 119.525)
- Column 3: Loop Statistics count (0, 0)
- Column 4: Loop Statistics Duration (0, 0)
- Column 5: HW Statistics by Line Floating Point Instr (1.65722e+10, 0)
- Column 6: HW Statistics by Line Load Misses in D1 (8.62196e+08, 2.46132e+08)
- Column 7: HW Statistics by Line Branch Instructions (6.94448e+09, 0)
- Column 8: HW Statistics by Line Load Instructions (3.10653e+10, 0)
- Column 9: HW Statistics by Line Instruction Cache M (9.55654e+07, 6.80955e+07)
- Column 10: HW Statistics by Line TLB misses (1.51119e+08, 4.0955e+07)

At the bottom, there are buttons for "Dismiss" and "Help" for the legend, and radio buttons for "Instrument/Clear Line" and "View Line Data".

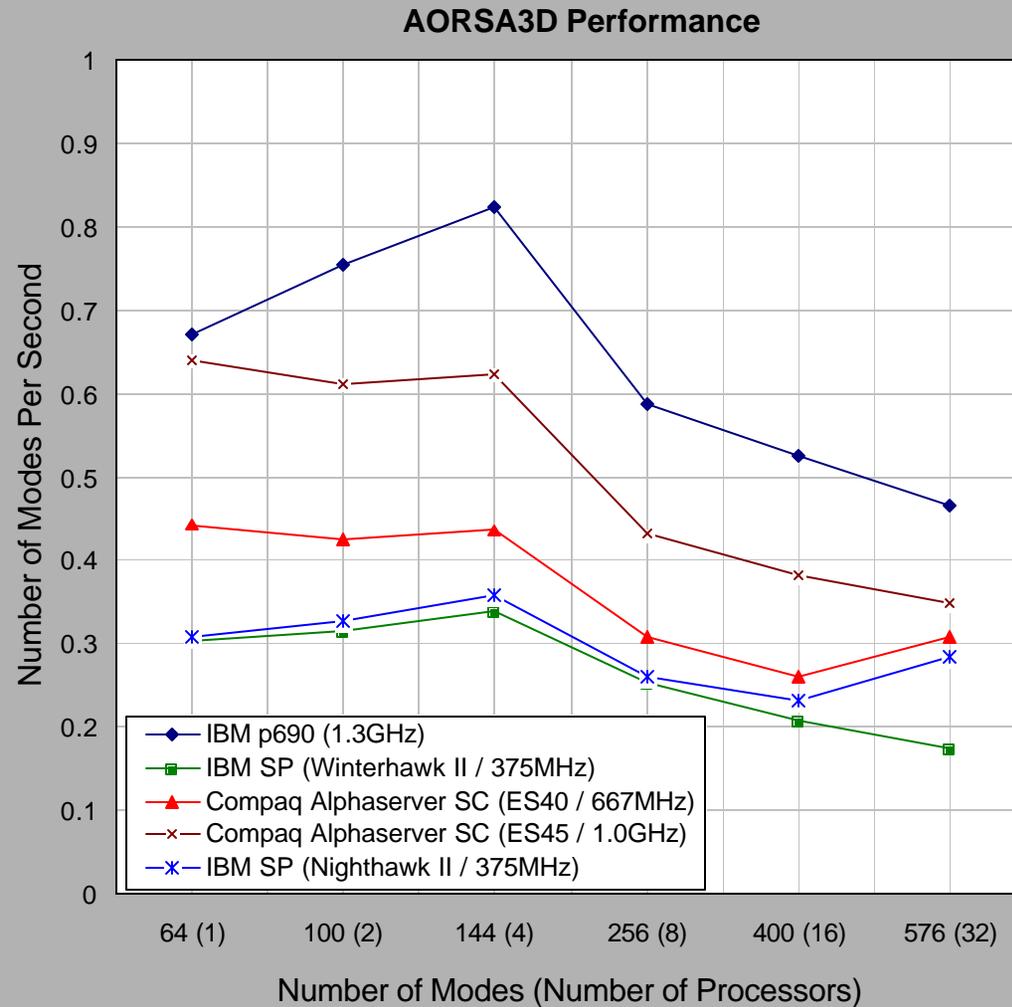


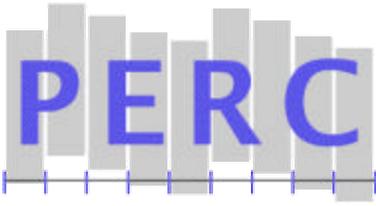
PAPI Perfometer Interface





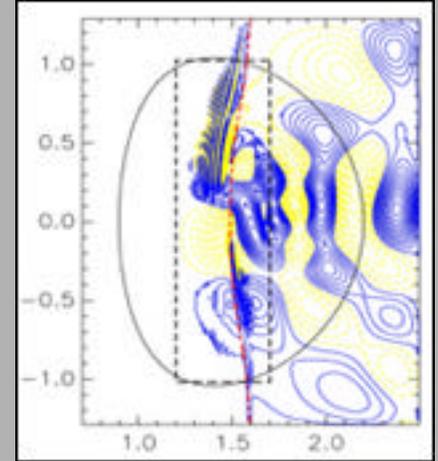
Performance Study of SciDAC Application





Performance Measurements for AORSA3D – a fusion code

- ✍ AORSA-3D solves for the wave electric field and heating in a 3D stellerator plasma heated by radio frequency waves, using an all orders spectral algorithm. It represents an important kernel in the "Numerical Computation of Wave-Plasma Interactions in Multi-Dimensional Systems" SciDAC project.
- ✍ A Fortran code that uses ScaLAPACK to solve a dense set of linear equations. ScaLAPACK routines handle the MPI communication and account for most of the execution time.

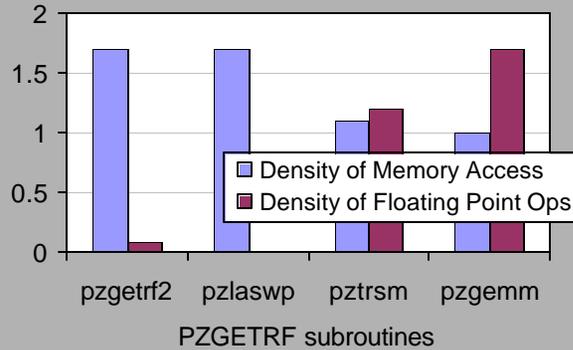


Performance data on next viewgraph:

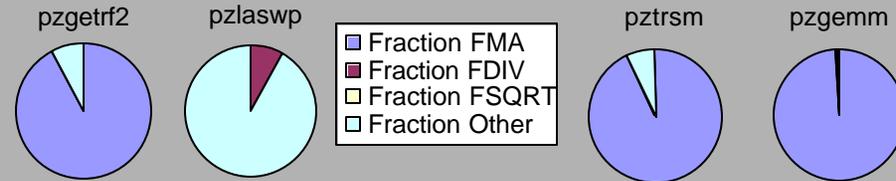
- ✍ Performance data for AORSA-3D on an IBM SP (Nighthawk II/375MHz).
- ✍ Hardware event counts for a 16 processor run collected with PAPI.
- ✍ Trace information collected for a 4 processor run with Vampirtrace.
 - ✍ Data for the ScaLAPACK LU factorization routine PZGETRF.
 - ✍ PZGETRF is composed of an initialization step, and a loop that calls 4 ScaLAPACK subroutines.



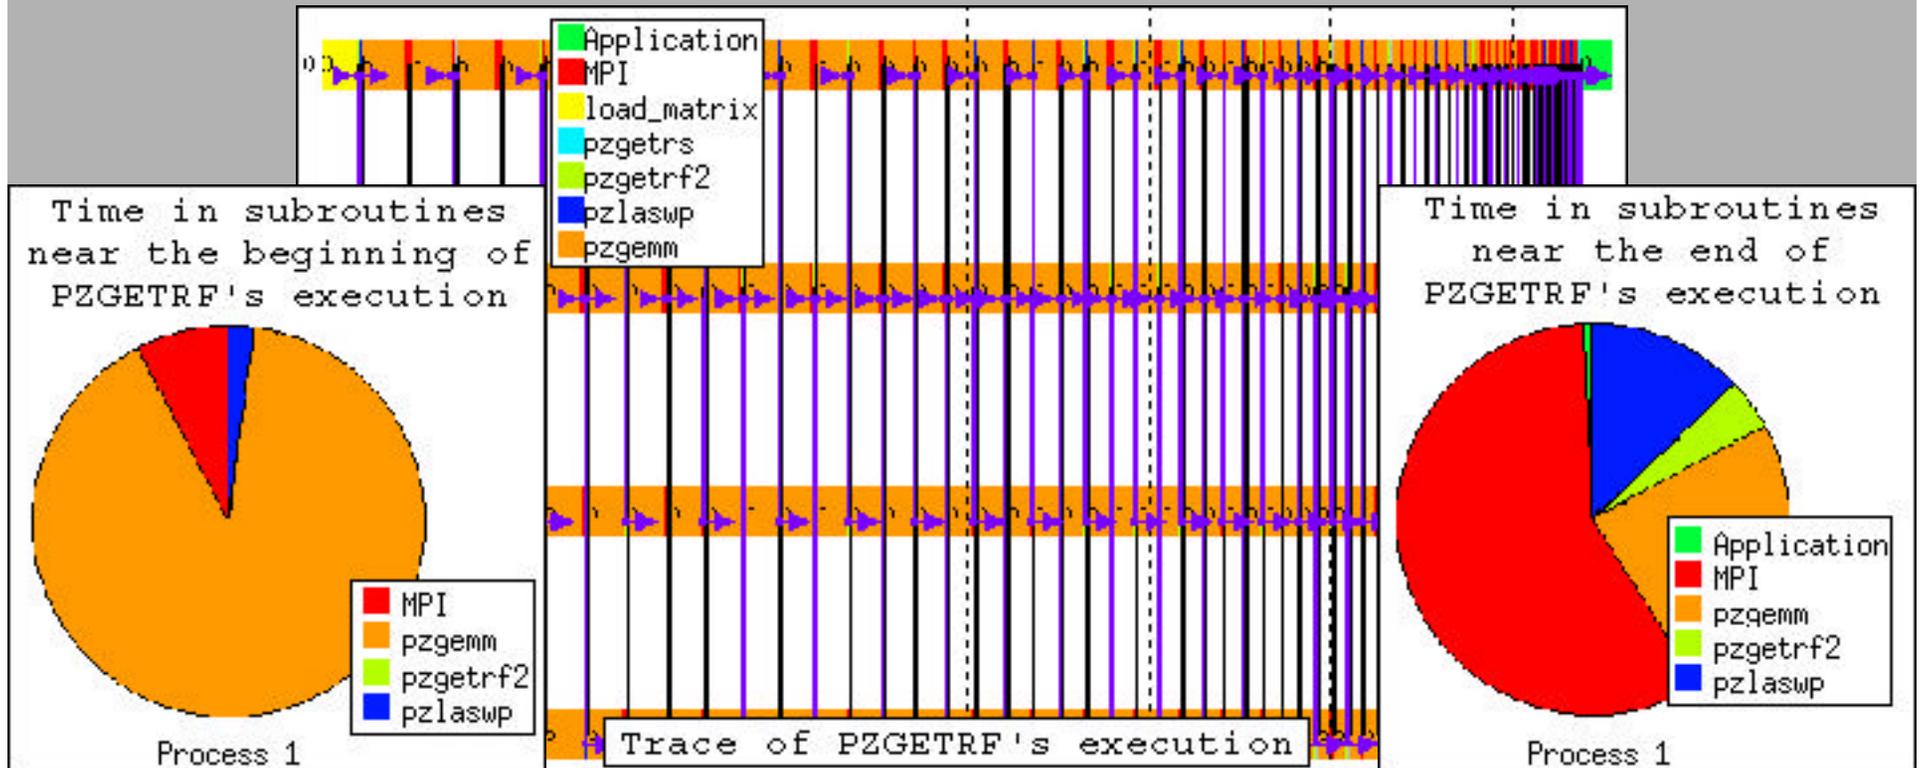
Performance Measurements for AORSA3D – a fusion code



Performance measurements and trace data for ScaLAPACK LU factorization routine PZGETRF.



Floating Point Composition of PZGETRF subroutines





SciDAC Interactions

Codes have been acquired from these projects:

- ✍ Terascale Simulation of Neutrino-Driven Supernovae
- ✍ Advanced Computing for 21st Century Accelerators
- ✍ National Computational Infrastructure for Lattice Gauge Theory
- ✍ Collaborate Design and Development of Community Climate System Model for Terascale Computers
- ✍ Numerical Computation of Wave-Plasma Interactions
- ✍ Accurate Properties for Open-Shell States of Large Molecules
- ✍ Terascale Optimal PDE Solvers
- ✍ An Algorithmic and Software Framework for PDEs



Working with PERC

Benchmarking

-  Application group works with PERC to specify relevant benchmark codes and problems.
-  PERC characterizes performance, generates performance models, and suggests optimizations.

Performance Tools

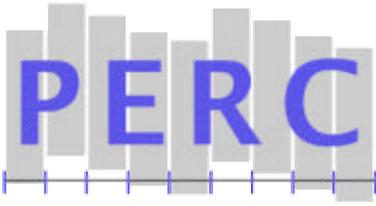
-  PERC trains application developers to use tools.
-  Application group uses tools in development, providing feedback on functionality and future development

For further information: <http://perc.nersc.gov>

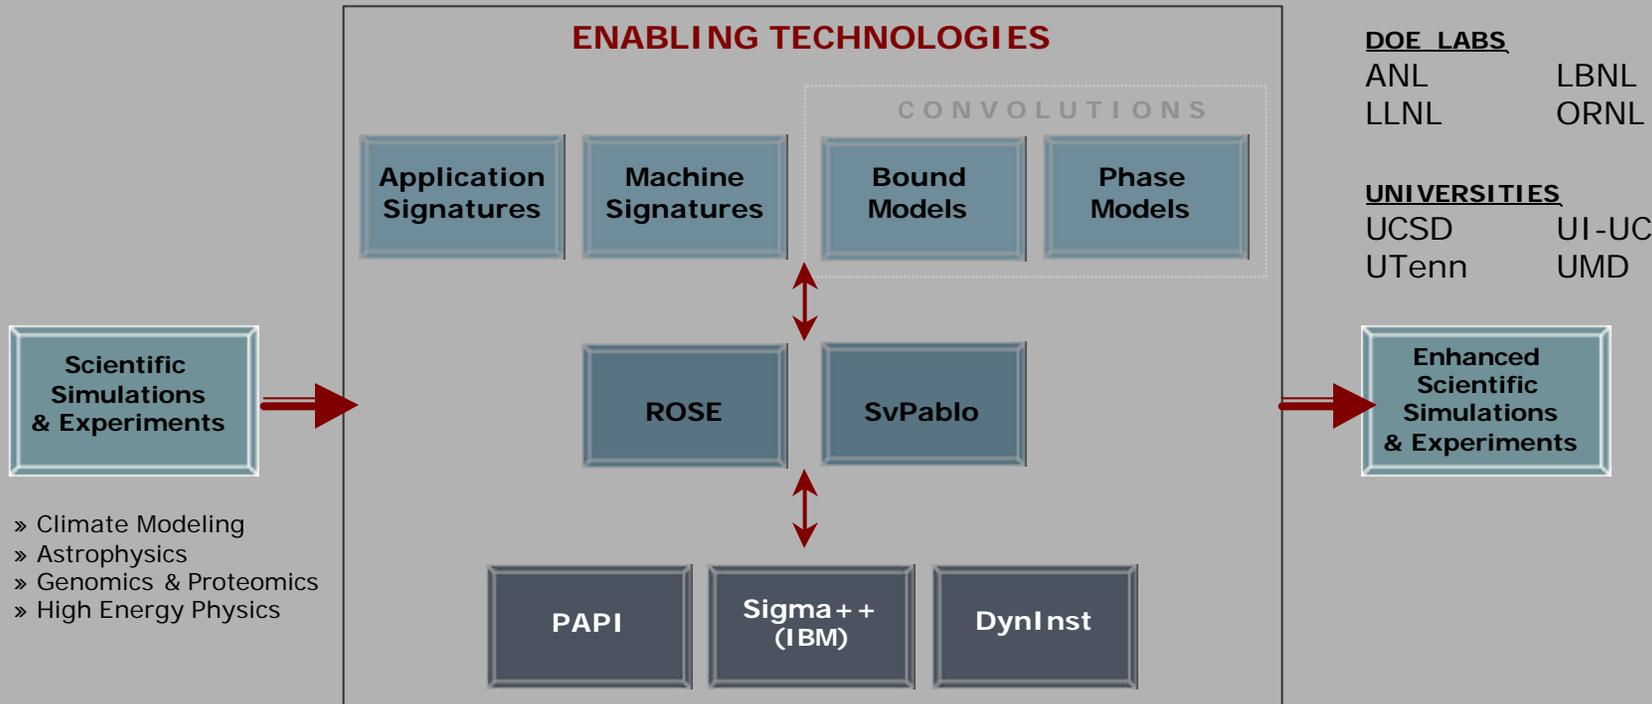


Summary

- ✍ Achieving optimal performance on HPC systems has compelling economic and scientific rationales.
- ✍ Performance is poorly understood – in depth-studies do not exist except in a handful of cases.
- ✍ PERC will pursue “performance science” and “performance engineering”, including improved benchmarks, monitoring tools, modeling techniques, and optimizers.



Developing a *science* for understanding performance of scientific applications on high-end computer systems, and *engineering* strategies for improving performance on these systems.



GOALS

Optimize and Simplify:

- Profiling of real applications
- Measurement of machine capabilities
 (emphasis on memory hierarchy)
- Performance prediction
- Performance monitoring
- Informed tuning

- Understand the key factors in applications that affect performance.
- Understand the key factors in computer systems that affect performance.
- Develop models that accurately predict performance of applications on systems.
- Develop an enabling infrastructure of tools for performance monitoring, modeling and optimization.
- Validate these ideas and infrastructure via close collaboration with DOE SC and other application owners.
- Transfer the technology to end-users.