Performance Evaluation and Analysis Consortium (PEAC) End Station

Patrick H. Worley - Oak Ridge National Laboratory

Co-Principal Investigators

Argonne National Laboratory	W. Gropp	*Lawrence Berkeley Natl. Lab.*	D. Bailey & L. Oliker	*Lawrence Livermore Natl. Lab.*	B. de Supinksi
Oak Ridge National Laboratory	J. Vetter & P. Worley	*Rice University*	J. Mellor-Crummey	*Univ. of California - Berkeley*	K. Yelick
Univ. of California - San Diego	A. Snavely	*University of Maryland*	J. Hollingsworth	*University of North Carolina*	D. Reed
University of Oregon	A. Malony	*Univ. of Tennessee - Knoxville*	J. Dongarra	*University of Wisconsin*	B. Miller

Overview

The PEAC End Station provides the performance evaluation and performance tool developer communities access to the National Leadership Computing Facility (NLCF) systems. The consortium goals are:

1. System evaluation
   - Evaluate the performance of NLCF systems using standard and custom micro-, kernel, and application benchmarks.
2. Performance tools
   - Port performance tools to NLCF systems, making these available to National Center for Computational Sciences (NCCS) users, and further develop the tools so as to take into account the scale and unique features of the NLCF systems.
3. Performance modeling
   - Validate the effectiveness of performance modeling methodologies, modifying them as necessary to improve their utility for predicting resource requirements for production runs on the NLCF systems.
4. Application analysis and optimization
   - Analyze performance and help optimize current and candidate NLCF application codes.
5. Performance and application community support
   - Provide access to other performance researchers who are interested in contributing to the performance evaluation of the NLCF systems or in porting complementary performance tools of use to the NCCS user community. Also provide access to application developers who wish to evaluate the performance of their codes on the NLCF systems.

all while adhering to the “Golden Rules” of the performance community:

i. Low visibility (no production runs!)

ii. Open and fair evaluations

iii. Timely reporting (using the PEAC wiki for posting and maintaining results)

Status (10/15/06)

- 48 users and 24 projects:
  - 12 tool development, 1 algorithm development, 5 benchmarking, 4 application characterization or modeling, 2 application code evaluation)
- consuming:
  - XT3: 364,000 processor-hours
  - X1E: 10,000 processor-hours
  - Altix: 2,000 processor-hours
- contributing to (at least)
  - 7 proceedings papers
  - 3 journal papers
  - numerous oral presentations, including at the Workshop on Petascale Computation for the Geosciences

Evaluation Projects

- **ANL:** Nek5000 (Fluid Dynamics) and PFLOTRAN (Subsurface Flow) application code benchmarks
- **LBNL:** Cactus (Astrophysics), ELB3D (Fluid Dynamics), LBMHD3D (Plasma Physics), BeamBeam3D (High Energy Physics), PARATEC (Material Science), HyperClaw (AMR Gas Dynamics), and GTC (Fusion) application code benchmarks. Also APEX-MAP system characterization benchmark.
- **ORNL:** Memory, interprocess communication, and I/O microbenchmarks; CAM/CLM and POP (Climate), AORSA, GTC, GYRO, and XGC (Fusion), S3D (Combustion), VH1 (Astrophysics), Amber (Molecular Dynamics), PFLOTRAN (Subsurface Flow) application code benchmarks.
- **PMAC:** Subsystem probes for system characterization needed for convolution-based performance modeling.

Tool Projects

- **Active Harmony:** runtime optimization framework
- **GASNet:** runtime networking layer for UPC and Titanium compilers
- **HPC Challenge Benchmark**
- **Logistical Networks:** infrastructure for flexible coscheduling of I/O and computation
- **Modeling Assertions:** performance model specification and verification framework
- **mpiP:** MPI profiling infrastructure
- **PAPI:** performance data collection infrastructure
- **PETSc:** toolset for numerical solution of PDEs
- **Sca/LAPACK:** numerical linear algebra library
- **SvPablo:** performance analysis system
- **TAU and KOJAK:** performance analysis system
- **Titanium compiler:** explicitly parallel dialect of Java
- **UPC compiler:** extension of C designed for high performance computing on large-scale parallel systems

Acknowledgements

- Research of presenter sponsored by the Office of Mathematical, Information, and Computational Sciences, Office of Science, U.S. Department of Energy under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC.
- This research used resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725, and of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.