Performance Evaluation and Analysis Consortium (PEAC) End Station

Presented by

Patrick H. Worley

Future Technologies Group
Computer Science and Mathematics Division

Overview

The PEAC End Station provides the performance evaluation and performance tool developer communities access to the Leadership Computing Facility (LCF) systems

Consortium goals	
System evaluation	• Evaluate the performance of LCF systems using standard and custom micro-, kernel, and application benchmarks • Determine efficient usage techniques and identify and characterize system performance issues
Performance tools	• Port performance-related tools and system infrastructure to LCF systems and make them available to Oak Ridge Leadership Computing Facility (OLCF) and Argonne Leadership Class Facility (ALCF) users • Further develop the tools to take into account the scale and unique features of LCF systems
Performance modeling	• Validate the effectiveness of and further develop performance characterization and prediction methodologies for LCF systems

Overview (continued)

Consortium goals (continued)	
Application analysis and optimization	• Analyze and help optimize performance of current and candidate LCF application codes
Performance and application community support	• Provide access to other performance researchers who are interested in contributing to the performance evaluation of the LCF systems or in porting complementary performance tools of use to the OLCF user community • Provide access to application developers who wish to evaluate the performance of their codes on LCF systems

All of this must be accomplished while adhering to the "Golden Rules" of the PEAC community:

- Low visibility (no production runs!)
- Open and fair evaluations
- Timely reporting of results

Status as of September 1, 2011

23 active ALCF users, 40 active OLCF users

Working on	Consuming	Contributing to (at least)
• Application performance analysis and optimization (18) • Application performance modeling (10) • Algorithm development (7) • Performance tool development (37) • Performance-related system infrastructure development (16) • System performance evaluation (10)	• XT5: 7,000,000 processor hours out of an allocation of 20,000,000 hours • BG/P: 2,000,000 processor hours out of an allocation of 10,000,000 hours *since January 1, 2011*	• Two book chapters • Four journal articles • Eighteen proceedings papers, including: 2 CUG, 1 Euro-Par, 1 ICPP, 3 ICS, 3 IPDPS, 5 SC11 • Two oral presentations • Two code releases *appearing since January 1, 2011*

many working in multiple areas

System evaluation

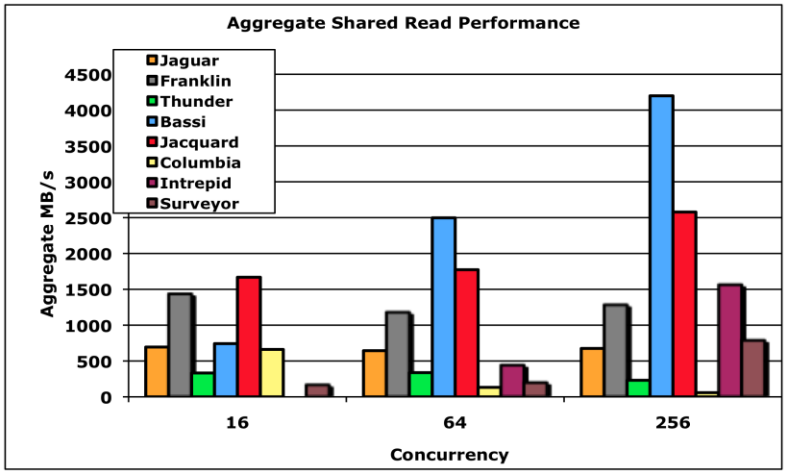

I/O performance characterization [LBL]

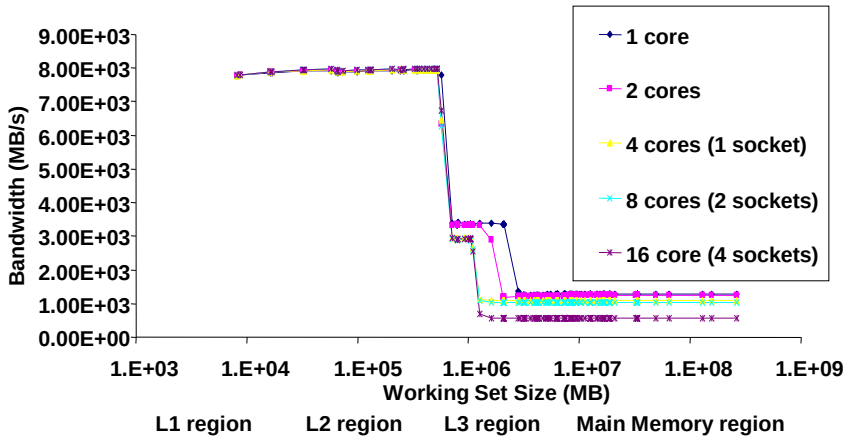

Effect of dual- to quad-core upgrade on memory hierarchy bandwidths [SDSC]

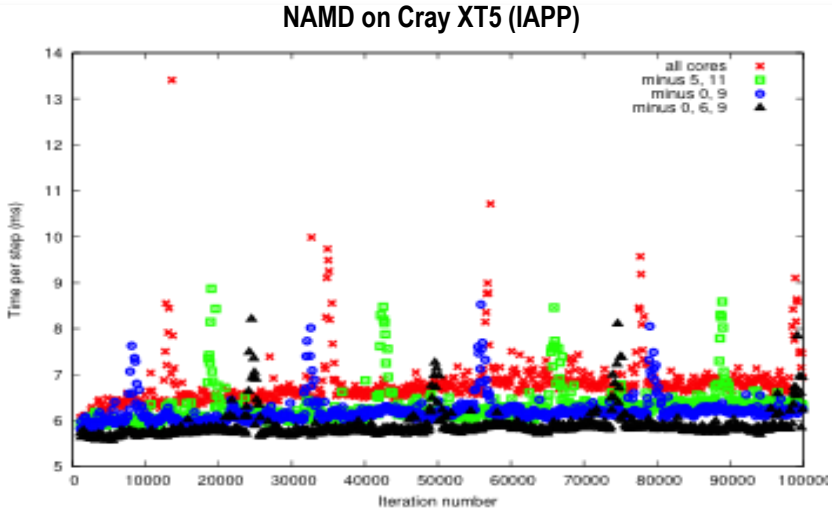

Measuring impact of system noise [UIUC]

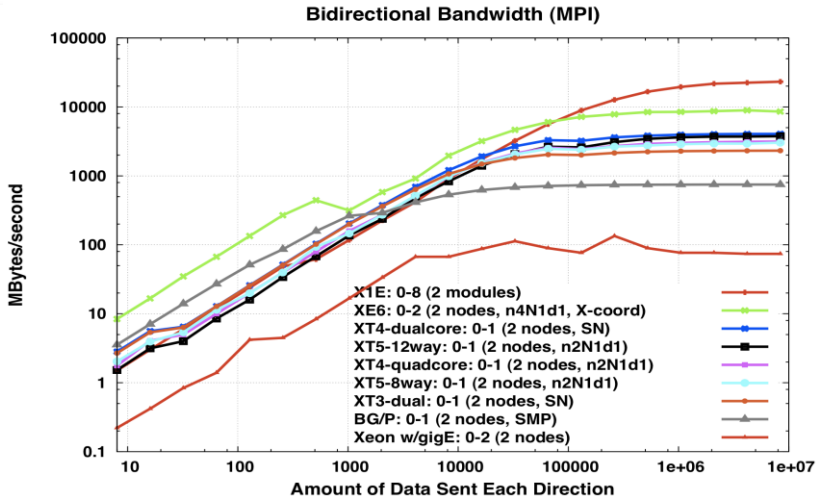

MPI performance evaluation [ORNL]

Application analysis and evaluation

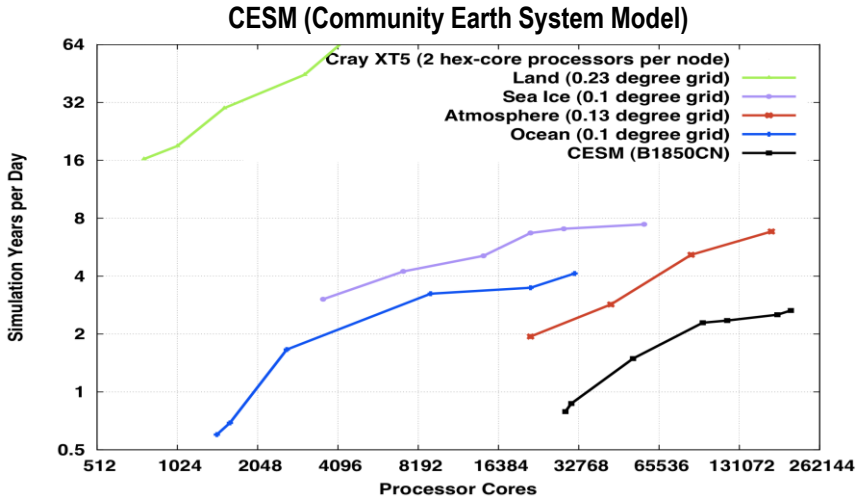

Strong scaling (high resolution) [ORNL/SNL]

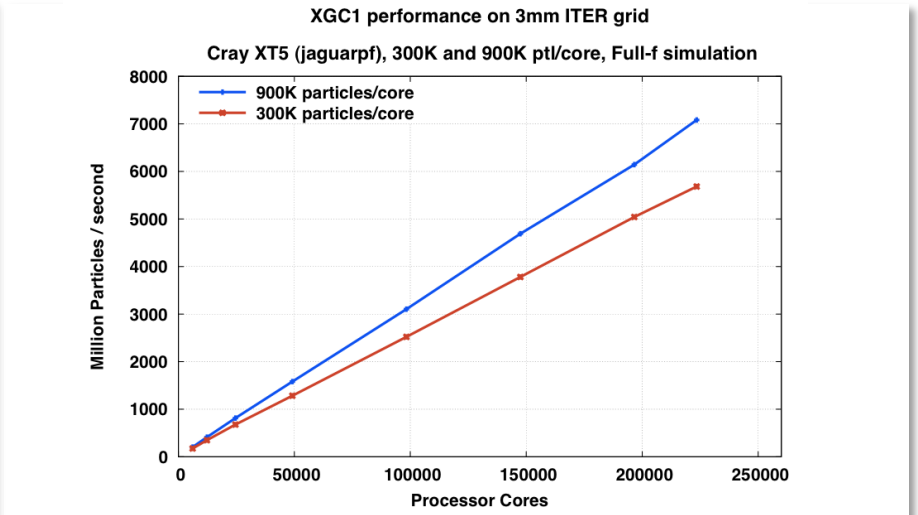

Full system XGC1 weak scaling [ORNL]

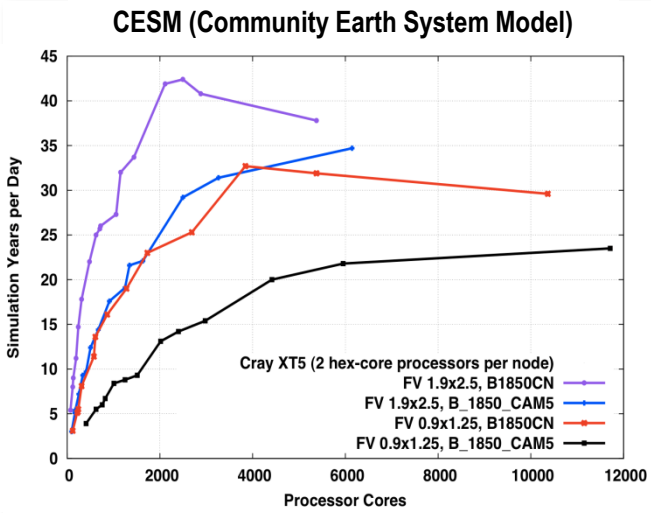

Strong scaling (low resolution) [ORNL]

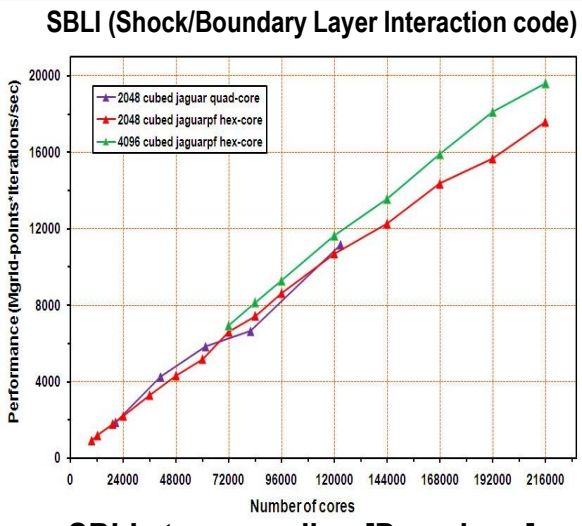

SBLI strong scaling [Daresbury]

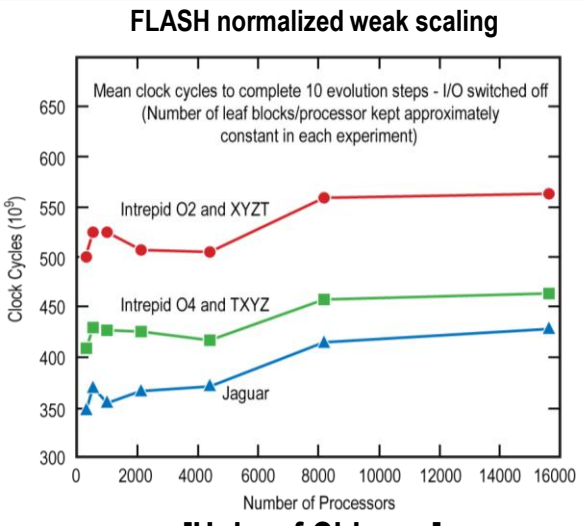

[Univ. of Chicago]

Algorithm analysis and comparison

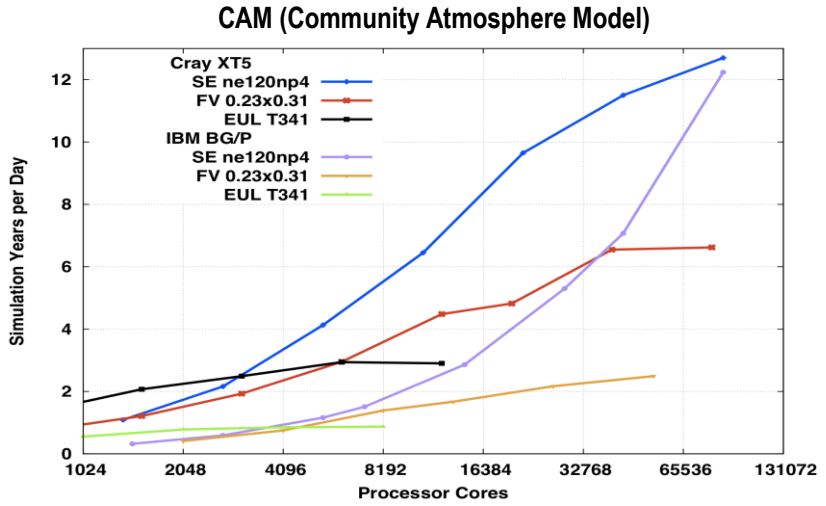

Algorithm and platform comparison [LLNL/ORNL/SNL]

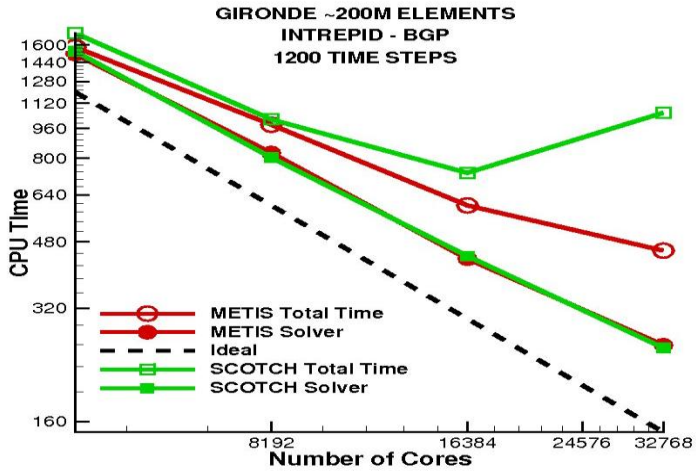

Parallel algorithm comparison [Daresbury]

Processing of genomes into domain maps: need improved load balancing that takes into account scale-free nature of the graphs [RENCI, NCSA]

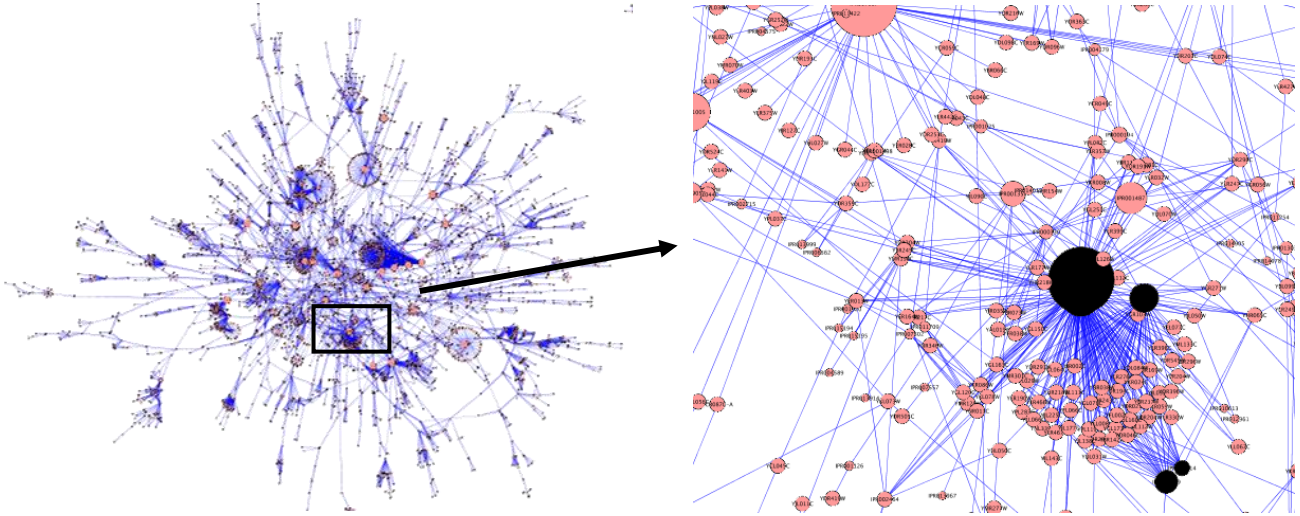

Optimization and optimization technology

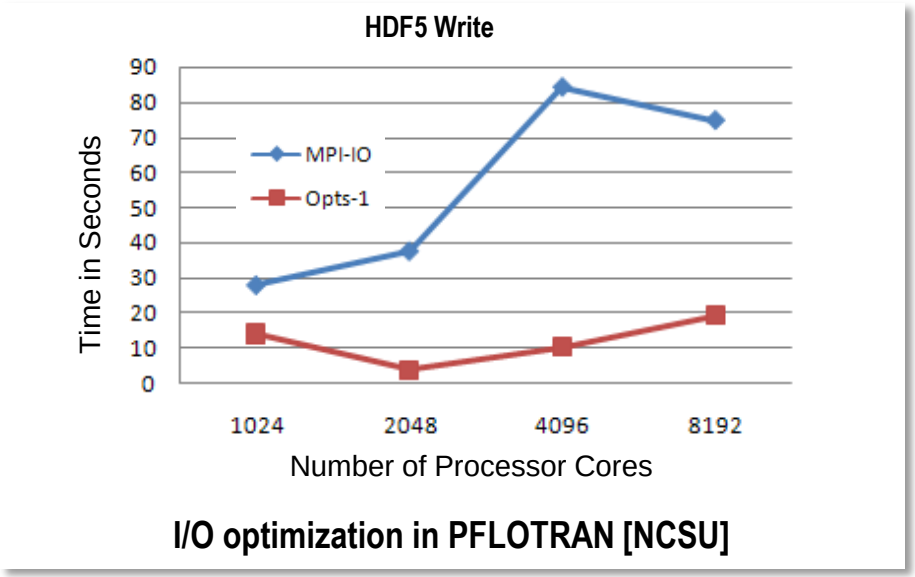

I/O optimization in PFLOTRAN [NCSU]

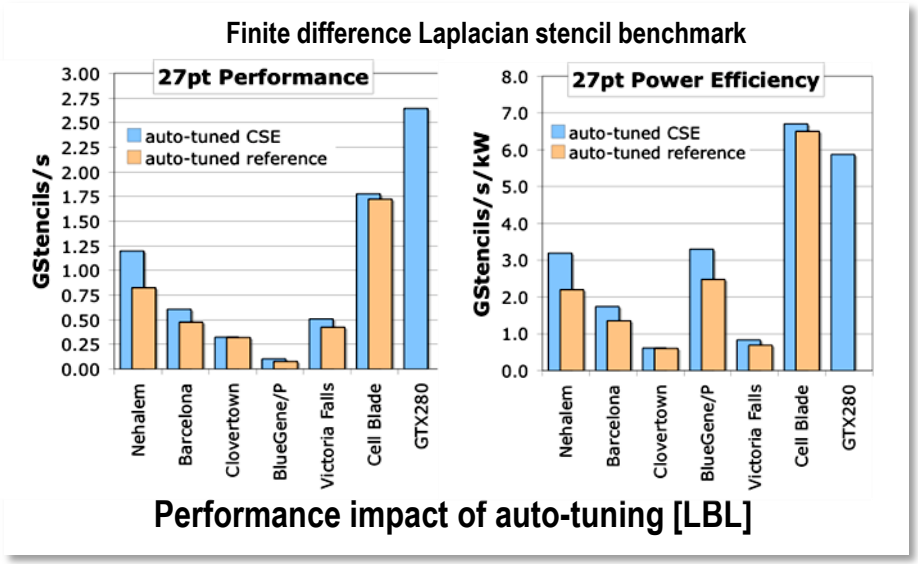

Performance impact of auto-tuning [LBL]

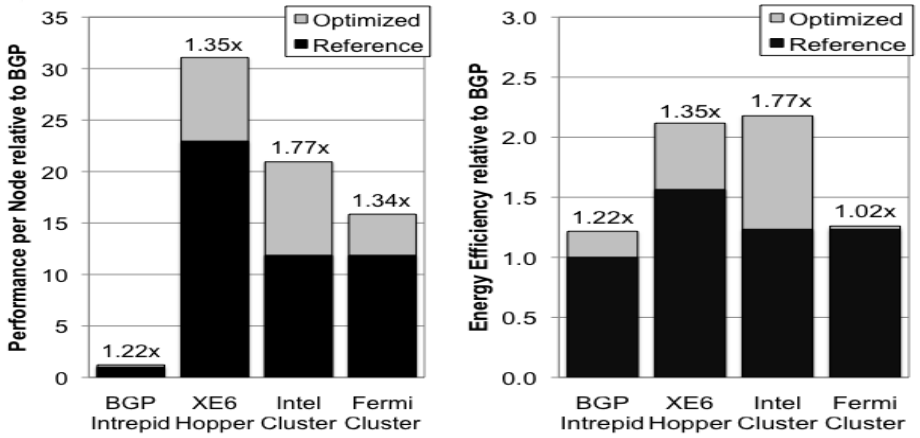

Optimizing for CPUs and GPUs [LBL]

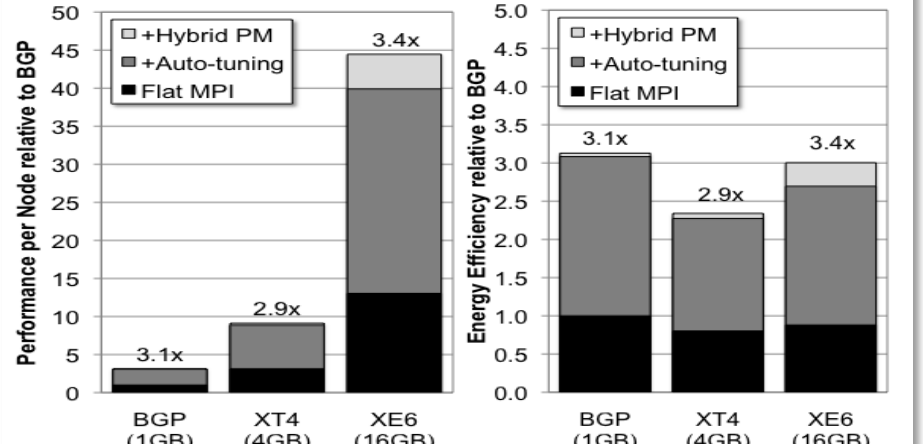

Application of auto-tuning [LBL]

Tool development

TAU graphical display of profile data by thread (Univ. Oregon)

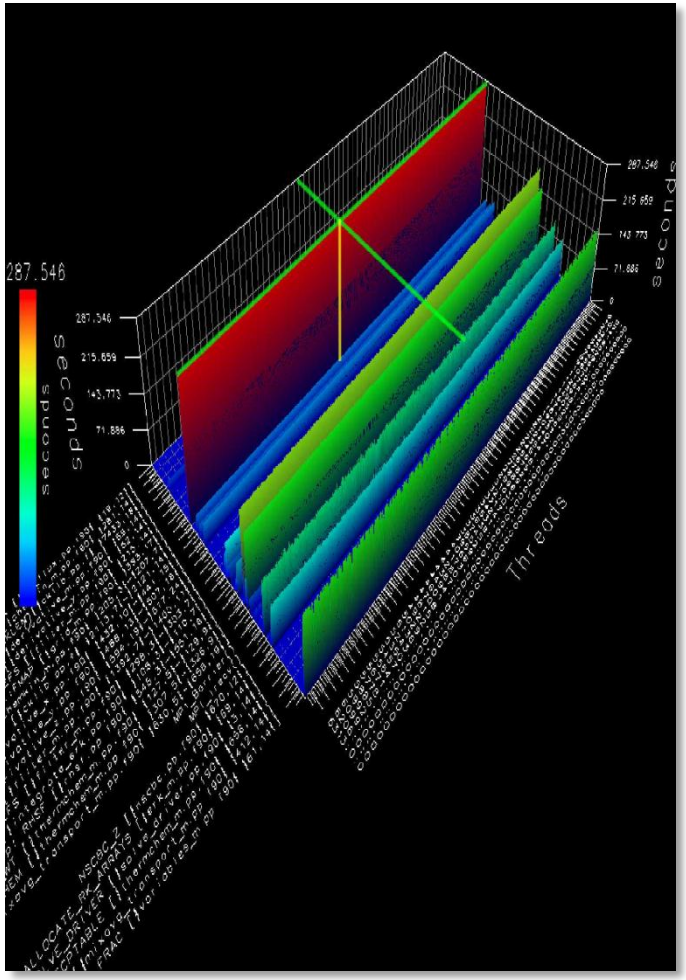

Visualization of point-to-point communication topology collected with mpiP (LLNL, ORNL)

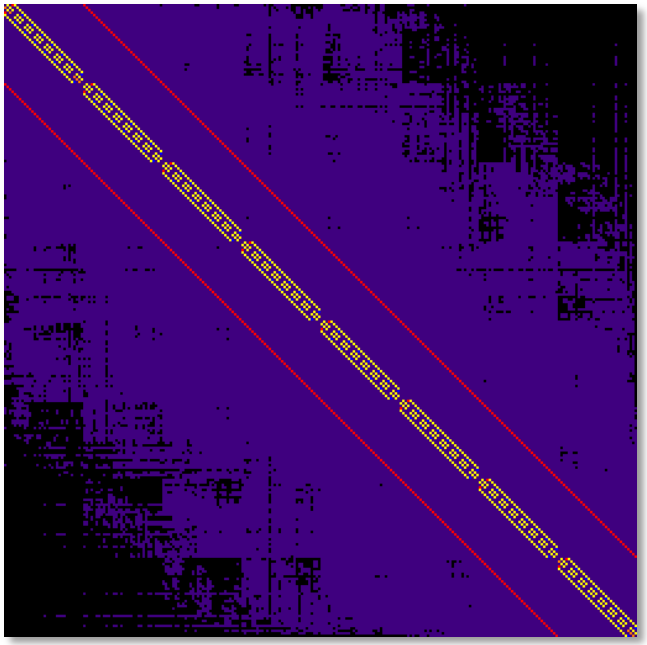

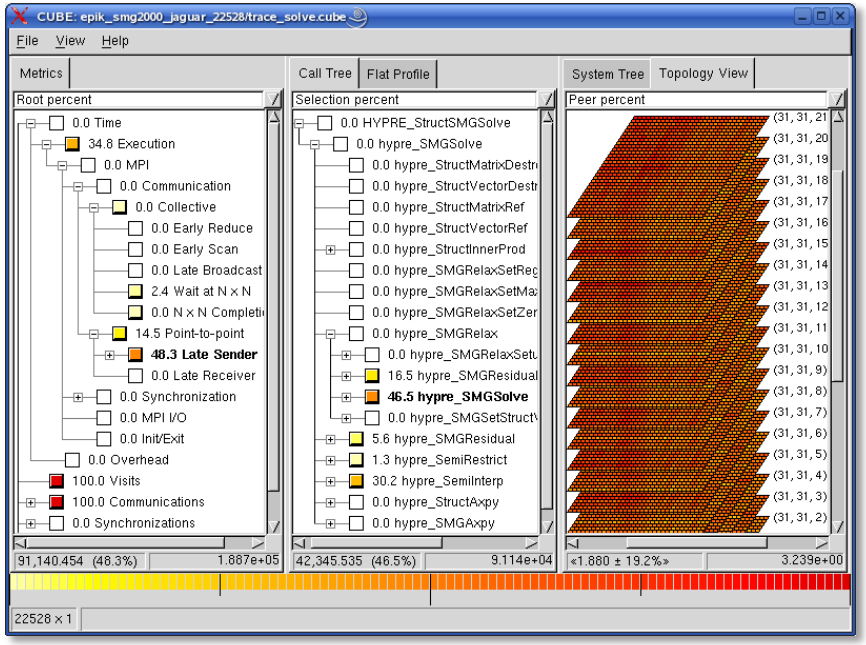

SCALASCA trace-based performance analysis (FZ-Jülich, Univ. Tenn.)

Tool development

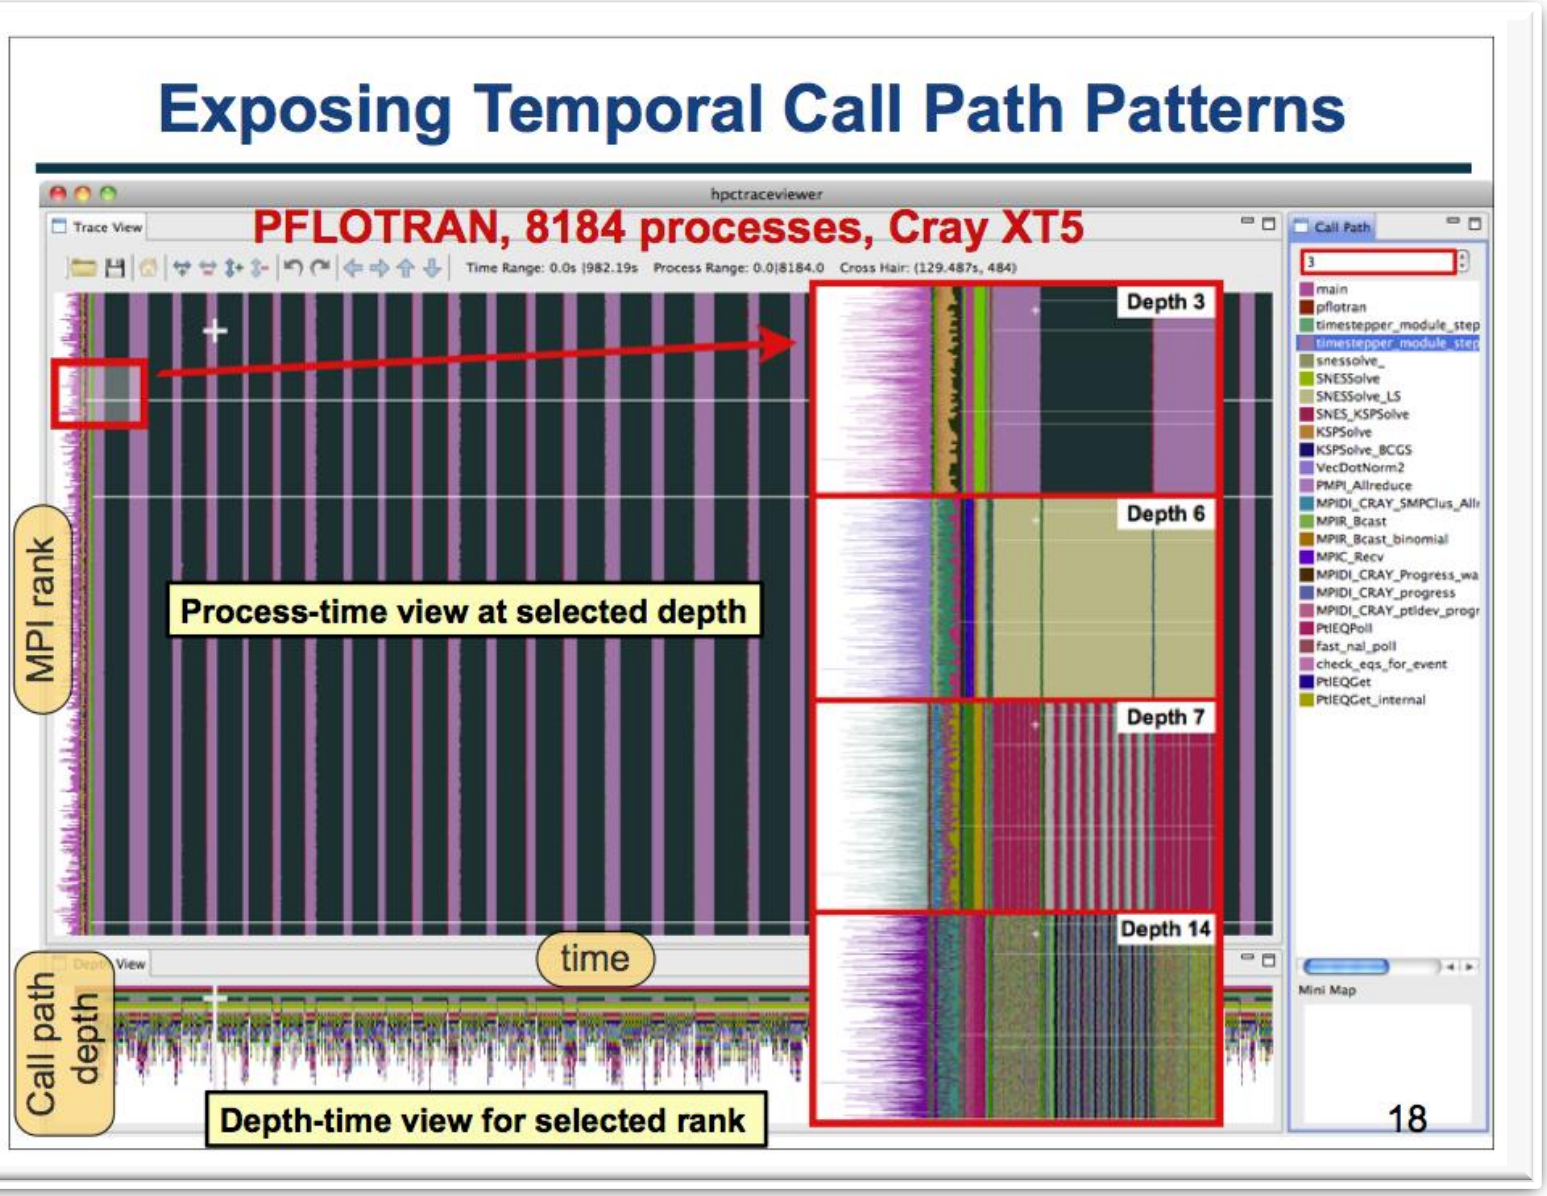

Sample-based tracing and visualization using HPCToolkit [Rice]

Tool development

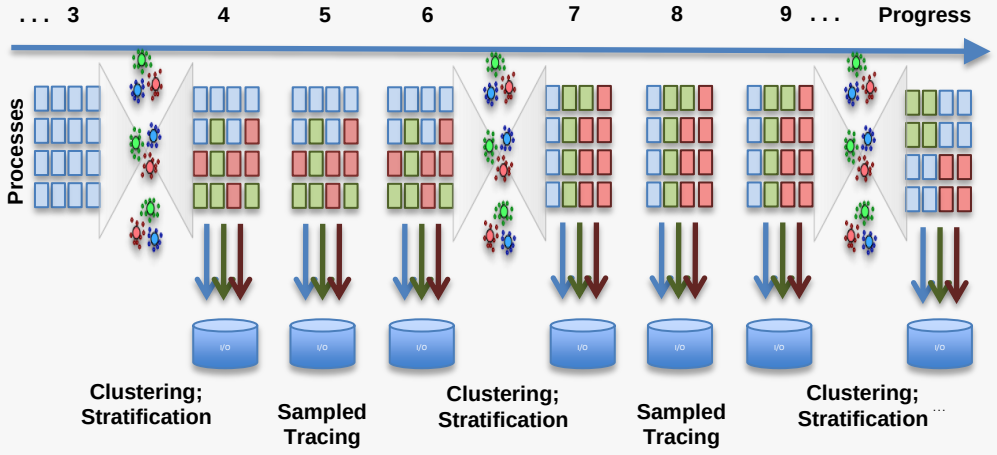

Using Clustering Algorithm with Parallel Extensions to K-Medoids (CAPEK) to adaptively sample performance traces (improving performance tracing at scale) [LLNL, RENCI, Microsoft]

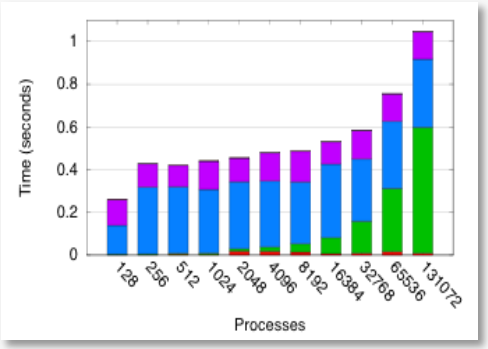

Using performance modeling to predict where performance optimization might be most effective [SDSC, ORNL]

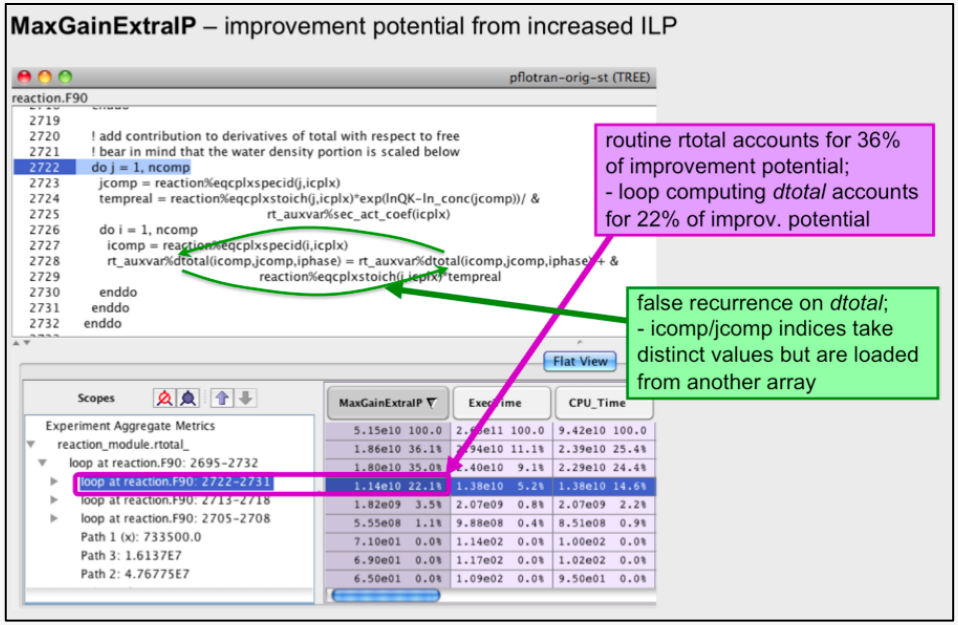

Co-principal investigators

University of Illinois at Urbana-Champaign	Lawrence Berkeley National Laboratory	Lawrence Livermore National Laboratory	Oak Ridge National Laboratory	Rice University	University of California–Berkeley and NERSC
William Gropp	David Bailey Leonid Oliker	Bronis de Supinski	Jeffrey Vetter Patrick Worley (PI)	John Mellor-Crummey	Kathy Yelick

University of California–San Diego	University of Maryland	University of Oregon	University of Tennessee	University of Wisconsin
Allan Snavely	Jeffrey Hollingsworth	Allen Malony	Jack Dongarra	Barton Miller

Contacts

Patrick H. Worley

Future Technologies Group
Computer Science and Mathematics Division
(865) 574-3128
worleyph@ornl.gov

Barbara Helland (OLCF Program Manager)

Betsy Riley (ALCF Program Manager)

Office of Advanced Scientific Computing Research
DOE Office of Science