SUPER NEWSLETTER

April 2014

Upcoming Events:

- IPDPS 2014, Phoenix, Arizona, May 19-23 (paper authors Kathryn Mohror, Todd Gamblin, Bronis de Supinski, and Sam Williams)
- ICS 2014, Munich, Germany, June 10-13 (workshop organizers Jeff Hollingsworth and Allen Malony, paper authors Gamblin, de Supinski, Williams, and Phil Roth)
- SC'14, New Orleans, LA, November 16-21, poster deadline July 31, 2014

From the Director

The SUPER project is now in the middle of its third year. As highlighted in previous newsletters, we have made significant progress in our performance, energy, and resilience research areas. Of course, these can all be traded off against each other. For example, reducing voltage and frequency lowers performance, but perhaps not as much as it improves energy efficiency. Or, one can gain resilience with redundancy, but at a cost in energy and performance. This leads us to the need for multi-objective optimization, the focus for this newsletter.

Multi-objective optimization seeks to collectively optimize multiple objectives from SUPER's other research areas. Multi-objective optimization is a difficult problem, and SUPER researchers are applying cutting-edge research results from the field of mathematical optimization to attack the problem. The results will allow users of leadership class systems to weigh the tradeoffs between alternative strategies so as to achieve the objectives most important to their applications. Our research will also point the way toward architecture-aware solutions whereby hardware and systems software designers can expose tuning knobs that can be used by the SUPER auto-tuning framework to optimize code for multiple objectives simultaneously.

In addition to this newsletter, the SUPER website contains up-to-date information on our research results and accomplishments, as well as publications and available software. Our mature performance measurement, analysis, and optimization tools are freely available and are also installed on DOE leadership class systems. We are working through our application partnerships on applying these tools as well as our more experimental research tools to improve all aspects of mapping the application codes effectively to leadership class systems.

- Bob Lucas

Framework for Optimizing Power, Energy, and Performance

When a single objective, such as execution time, is of interest, the autotuning search problem can be posed as a numerical optimization problem. Increasingly, multiple metrics (such as execution time, energy consumption, resilience to errors, power demands, and memory footprint) are of interest simultaneously. When the relative weights or constraints on these objectives are not known at search time, autotuning becomes a multi-objective optimization problem. We are developing a formalism for multi-objective optimization studies of broad applicability in autotuning, architecture design, and other areas of HPC. With a focus on time, power, and energy, our initial work illustrates that a multi-objective analysis provides richer insight than do constrained and single-objective formulations.

Framework for Optimizing Power, Energy, and Performance (cont.)

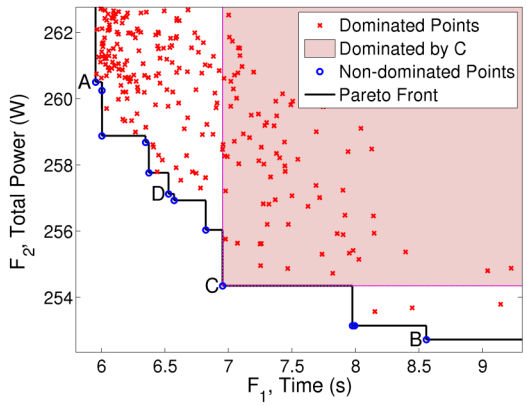

Figure 1. Illustration of Pareto fronts when minimizing two objectives (fdtd kernel, input size 512, Intel Xeon E5530). The points A, B, C, and D are nondominated and hence belong to the Pareto front. The shaded area represents the region in time ($\mathbf{F_1}$) and power ($\mathbf{F_2}$) space that is **dominated** by the point **C**; all points in this region are inferior to **C** in both objectives. The set of nondominated points form the Pareto front. If the objective $\mathbf{F_1}$ ($\mathbf{F_2}$) is minimized in isolation, then we obtain the point **A** (**B**), which necessarily belongs on the Pareto front. However, not all points on the Pareto front necessarily correspond to minimizers of a linear combination of the objectives (see, e.g., **D**). Hence, the Pareto front contains significantly richer information than one obtains from simple single-objective formulations. For example, if one were to minimize time subject to a constraint on power, $\mathbf{F_2(x) \leq P}$, the Pareto front provides the solution for *all possible* values of the cap **P**.

Because of the relationship between power and energy, we have a simple relationship between the two objective spaces. Mathematically, we prove that all points on the energy-time Pareto front have a corresponding point on the power-time Pareto front. Consequently, the number of nondominated points for energy-time is bounded by the number of nondominated points for power-time. Furthermore, we establish a necessary condition to observe such tradeoffs: the power savings of a slower code variant must outpace the product of idle power and the relative slow-down between the slower and faster code variants.

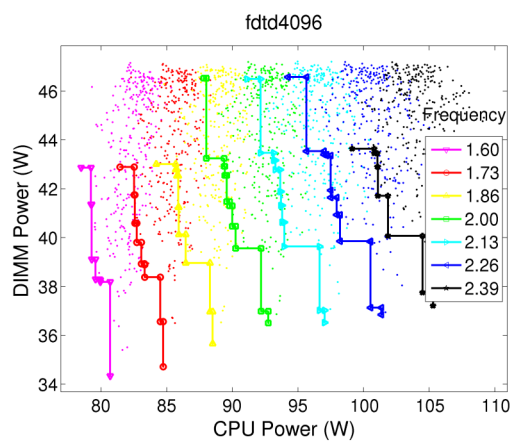

Figure 2. Pareto fronts (for each clock frequency) on Intel Xeon E5530 for component-level power draws of the **fdtd4096** kernel. When we analyze each of the fronts for different clock frequencies in isolation, we see a clear tradeoff between DIMM and CPU power draws for different code variants. We attribute this behavior intuitively to the optimizations that impact data motion. Code variants that better utilize the caches can reduce the stress on DIMMs by lowering the number of data transfers from the main memory, thereby lowering the DIMM power. At the same time, better cache utilization leads to fewer stalls and more compute work for the CPU, thereby raising its power demand. Such tradeoffs are of interest in studies for future architectures where one may consider constraining CPU draw (e.g., for thermal/fault considerations) and/or DIMM draw (e.g., as a proxy for the effective memory footprint or as a simulator of memory-starved systems).

Recent Software Releases:

Active Harmony 4.5, September 2013

http://www.dyninst.org/harmony/

TAU 2.23, November 2013

http://tau.uoregon.edu/

PAPI 5.3.0, December 2013

http://icl.utk.edu/papi/

PaRSEC/DPLASMA 1.2.0, April 2014

http://icl.utk.edu/parsec/

See the SUPER website for more software releases

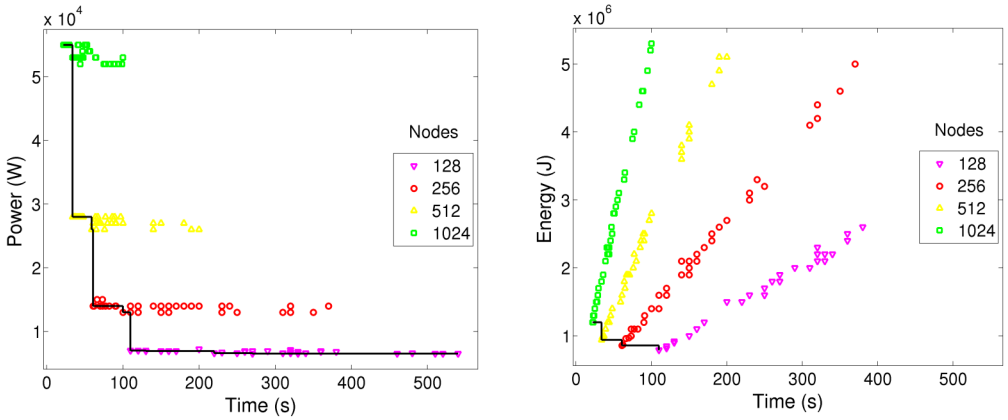

Figure 3. Tradeoffs in time, power, and energy obtained for miniFE, a finite-element mini-application, on the BG/Q system at Argonne. The results show that there are tradeoffs between time to completion and both power and energy. As expected, increasing the node count decreases the time to completion but increases the power draw. Concerning energy, the best parameter configuration within each node count provides a tradeoff between time to completion and energy consumption. Within a given node count, however, the fastest code variant consumes the least energy.

Our findings show that in some settings objectives are strictly correlated and there is a single, ideal decision point; in others, significant tradeoffs exist. The existence of these tradeoffs can motivate hardware designers to expose a richer set of "knobs" (or configuration options) to future administrators and software designers. This framework and our analysis are sufficiently general and can be easily extended to incorporate new hardware - and software-based power and energy knobs as they become available. We expect to develop a tool to guide SciDAC code developers to explore and analysis such tradeoffs for their applications.

[More details can be found in: Prasanna Balaprakash, Ananta Tiwari, and Stefan M. Wild. Multi-objective optimization of HPC kernels for performance, power, and energy. In Proceedings of the 4th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS13). To appear, preprint available at http://www.mcs.anl.gov/papers/P4069-0413.pdf.]

The SUPER team at the Utah all-hands meeting, March 2014

Featured SUPER Researcher: Prasanna Balaprakash

Where do you work and how are you involved with SUPER?

I am an assistant computer scientist with a joint appointment in the Mathematics and Computer Science Division and the Leadership Computing Facility at Argonne National Laboratory. Within SUPER, I focus on two related areas. First, machine learning approaches for performance modeling and architecture exploration. Second, multi-objective mathematical algorithms to optimize several possibly conflicting code performance metrics related to run time, power, energy, and resilience simultaneously.

Can you briefly summarize your educational and work background?

I received a Bachelor's degree in computer science engineering from the Periyar University in Salem, India; a Master's degree in computer science from the Otto-von-Guericke University Magdeburg in Germany; and a PhD in Engineering Sciences from the Artificial Intelligence Laboratory of the Universite Libre de Bruxelles in Belgium. I worked as a chief technology officer at Mentis Sprl., a data analytics startup in Brussels, Belgium for a year before moving to Argonne in late 2010, where I was a postdoc with Stefan Wild until the end of 2013.

Where are you from originally?

I am from Erode, Tamil Nadu, India.

What are your research areas of interest?

My research interests span the areas of machine learning, numerical optimization, and performance engineering. My research focus is on the design, development, and analysis of algorithms for solving large-scale problems that arise in automating the tuning of computer codes and on computationally expensive design-space explorations.

What do you see yourself doing five years from now?

I plan to continue my focus on automatic performance modeling and optimization for high-performance and emerging computing platforms. I believe that deploying automatic performance modeling algorithms, together with appropriate machine learning and mathematical optimization approaches for the compiler and runtime systems, is key to tackling some of the main problems facing performance engineers today, including meeting power and energy constraints; achieving performance goals on diverse architectures; and managing the ever-growing complexity in the design, analysis, and portability of scientific codes on extreme-scale systems. To that end, I will apply a whole-system view and develop relevant modeling and prediction functionalities for both compiler/runtime and compiler/programmer interfaces. I am excited about utilizing my background in performance engineering and machine learning to form collaborations with other research groups. In a nutshell, I hope to advance the state of the art in performance modeling for future extreme-scale systems by increasing the applicability of mathematical modeling, machine learning, and optimization approaches.

What are some things you enjoy doing that don't involve computers?

Road trips, driving, cooking, radio, cricket, and ultimate frisbee with the Argonne club.

Selected Recent Publications

- Chiang, W-F., G. Gopalakrishnan, Z. Rakamaric, **and** A. Solovyev, **"Efficient Search for Inputs Causing High Floating-point Errors", *19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'14)*, Orlando, Florida, February, 2014.**
- Balaprakash, P., K. Rupp, A. Mametjanov, R. B. Gramacy, P. D. Hovland, **and** S. M. Wild, **"Empirical Performance Modeling of GPU Kernels Using Active Learning", *Parallel Computing: Accelerating Computational Science and Engineering (CSE)*, vol. 25: IOS Press, pp. 646–655, 2014.**
- Sato, K., A. Moody, K. Mohror, T. Gamblin, B. R. de Supinski, N. Maruyama, **and** S. Matsuoka, **"FMI: Fault Tolerant Messaging Interface for Fast and Transparent Recovery", *28th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2014)*, 05/2014.**
- Malony, A. D., **and** K. A. Huck, **"General Hybrid Parallel Profiling", *Parallel, Distributed, and Network-Based Processing (PDP 2014)*, Turin, Italy, 02/2014.**
- Behzad, B., S. Byna, S. M. Wild, Prabhat, **and** M. Snir, **Improving Parallel I/O Autotuning with Performance Modeling, , 2014.**
- Pearce, O., T. Gamblin, B. R. de Supinski, T. Arsenlis, **and** N. M. Amato, **"Load Balancing N-Body Simulations with Highly Non-Uniform Density", *International Conference on Supercomputing (ICS'14)*, 06/2014.**
- Aktulga, H.. M., A.. Buluc, S. Williams, **and** C.. Yang, **"Optimizing Sparse Matrix-Multiple Vector Multiplication for Nuclear Configuration Interaction Calculations", *International Parallel and Distributed Processing Symposium (IPDPS 2014)*, 05/2014.**

See the SUPER website for additional recent publications

SUPER website: http://www.super-scidac.org/
Contact: Bob Lucas, rflucas@isi.edu
Newsletter editor: Bonnie Poore, bonnie.e.poore@gmail.com