

## Guest-editorial

---

# Best Papers from SC2000

Rod Oldehoeft<sup>1</sup>

*Computer and Computational Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA*

*E-mail: rro@lanl.gov*

The SC2000 Conference on High Performance Networking and Computing was held in Dallas, TX, USA on November 4–10, 2000. Sponsored by the IEEE Computer Society and the ACM Special Interest Group on Computer Architecture (SIGARCH), SC 2000 attracted 5,000 people, including over 2,000 Technical Program registrants.

I was the Technical Papers Committee Chair for SC 2000, and 60 computing professionals from all over the world served on the Committee. In addition, four leaders in high-performance computing presided over disjoint subcommittees of the 60:

- Randy Bramley, Indiana University, Applications, Numerics, and Visualization;
- Bob Lucas, Lawrence Berkeley National Laboratory, Architecture, Networks, and Distributed Computing;
- Dan Reed, University of Illinois, Performance Optimization;
- Ron Perrott, Queens University, Belfast, Programming Models and Tools.

The 179 submissions we received were strong and diverse, and the Committee selected 62 for inclusion. These were organized into 21 sessions for the Conference.

At the Committee Selection Meeting, the subcommittees nominated papers for the Best Paper Award. Before the Conference itself, Professor Perrott also chaired a “subcommittee of the whole” to select the Award winner. This was both difficult and pleasant:

the choice of best paper was closely decided, but all the candidates were of excellent quality.

Six of the seven nominated “Best Papers” are presented in this special issue of *Scientific Programming*. They are not necessarily the same as in the Conference proceedings, for most contain new and extended results. They reflect the diversity found in the Conference as a whole, with two papers each in the currently “hot” areas of Grid computing and OpenMP, while the two application papers represent an always-strong component of the technical program at SC Conferences.

### Grid Computing

Many people are actively engaged in developing various technologies to use the global Internet to build infrastructure for distributed computing and communication. Collectively, these technologies are often referred to as “the Computational Grid” or just “the Grid”. By analogy with the electric power grid, the future is envisioned as one in which pervasive access to advanced computational resources will have a similar transforming effect. These two papers contribute to the Grid infrastructure and provide a modeling and analysis tool.

“The AppLeS parameter sweep template: User-level middleware for the Grid”, by Henri Casanova, Graziano Obertelli, Francine Berman, and Richard Wolski, describes user-level middleware that supports “parameter sweep applications”. Such applications explore a parameter space with independent runs. The “AppLeS Parameter Sweep Template” supports the efficient deployment of these runs across a Grid of computing resources, addressing issues of dynamic and adaptive scheduling, sharing or replicating large input files, and handling large output files. Results are reported for different scheduling strategies, and several issues for future work are delineated.

In “The MicroGrid: A scientific tool for modeling computational grids”, by H.J. Song, X. Liu, D. Jacobsen, R. Bhagwan, X. Zhang, K. Taura, and A. Chien, a tool for simulating the execution of Globus distributed

---

<sup>1</sup>This work was supported by the US Department of Energy through Los Alamos Laboratory Contract W-7405-ENG-36. LA-UR 01-1618.

applications is presented. The Grid is a dynamic place to compute in terms of heterogeneity, availability, and reliability of resources. Supporting middleware must respond flexibly to achieve satisfactory performance for such applications. The MicroGrid allows broad experimentation as applications are simulated in a virtual Grid environment. Results with micro-kernels and the NAS parallel benchmarks show that the MicroGrid can closely match actual experiments.

## OpenMP

OpenMP is a set of compiler directives, library routines, and environment variables that can be used to specify shared-memory parallelism in Fortran, C, and C++ programs. The OpenMP specification has been adopted by many system and parallelizing compiler vendors, so it supports the crafting of portable parallel programs for SMPs. In today's large-scale NUMA SMPs, some memory is closer to a particular processor than other memory. Users have learned that data placement relative to the processor accessing the data is important for high performance. These two papers examine the question of extending OpenMP for making data placement decisions. The first introduces an automatic data migration facility and concludes that such extension is unnecessary, while the second describes extensions and their implementation that allow users to specify placement. In both, performance is improved.

"A transparent runtime data distribution engine for OpenMP", by Dimitrios S. Nikolopoulos, Theodore S. Papatheodorou, Constantine D. Polychronopoulos, Jesús Labarta, and Eduard Ayguadé, won the SC2000 Best Paper Award. First, the authors examined the performance implications of using NUMA hardware, and observe that, using the NAS parallel benchmarks "tuned specifically for the Origin2000", only modest performance losses are incurred. Second, a transparent, user-level page migration engine is presented to improve locality for iterative OpenMP code, with impressive results. The authors conclude that proposed data distribution directives for OpenMP are unneeded.

In "Extending OpenMP for NUMA machines", by John Bircsak, Peter Craig, Raelyn Crowell, Zarka Cvetanovic, Jonathan Harris, C. Alexander Nelson, and Carl D. Offner, the approach is to augment OpenMP with additional directives that allow a user to control data layout, and to specify when a data page is to be

migrated to be near a processor. The implementation of these extensions is detailed, and preliminary results are provided.

## Applications

Papers about scientific applications have always held a prominent place in the technical program of the SC conferences. These two papers reflect a more recent broadening of the scientific scope for such application papers. One presents a breakthrough in computational support for adaptive statistical design of experiments, and the other demonstrates new computational power available for biomolecular simulations.

"Scalable algorithms for adaptive statistical designs", by Robert Oehmke, Janis Hardwick, and Quentin F. Stout, describes an efficient, highly scalable approach to optimizing and analyzing adaptive statistical design of experiments. While static designs define sampling decisions in advance, adaptive designs learn as the experiment proceeds. They have significant ethical (when human subjects are involved) and cost advantages, but are difficult computationally: huge memory requirements, few operations per memory access, and multiply nested loops with dynamic indices. The algorithm described here can solve much larger problems than examples which had recently been deemed impractical. Mr. Oehmke also presented this paper at SC2000 in competition for the Best Student Paper Award.

"Scalable molecular dynamics for large biomolecular systems", by Robert K. Brunner, James C. Phillips and Laxmikant V. Kalé, presents an optimized parallel method for molecular dynamics simulations of large biomolecular systems. Such simulations have become practical with the experimental discovery of structures for an growing set of proteins, but previous parallelization efforts have only scaled well using modest numbers of processors. This C++ program shows sequential performance that compares well with a Fortran counterpart, and has achieved better than 60% efficiency on 2,048 processors of the ASCI Red system. Various techniques for achieving this level of performance are discussed.

These six papers represent the best among technical papers in the Technical Program at the 2000 SC Conference. For information about the SC Conferences, see [www.supercomp.org](http://www.supercomp.org).




**Hindawi**

Submit your manuscripts at  
<http://www.hindawi.com>

