Performance Characterization of the World's Most Powerful Supercomputers

Guest Editors: Rupak Biswas[1] and Leonid Oliker[2]

[1]NAS Division, NASA Ames Research Center, Moffett Field, CA 94035
[2]CRD/NERSC, Lawrence Berkeley National Laboratory, Berkeley, CA 94720

Over the last decade, the architectural spectrum of large-scale parallel systems has been dominated by the cluster-of-workstation paradigm. These systems are characterized by loose hardware integration between the networking and the CPU/memory components, and lack the efficiency of custom-designed computing platforms. As a result, while peak processor performance increases exponentially according to Moore's Law, improvements in sustained performance lag far behind. The problem is expected to exacerbate by the end of this decade, because traditional uniprocessor scaling trends are slowing down as most of the benefits of instruction level parallelism and pipelining have been mined, and clock frequency increases are limited by power dissipation concerns. In the same time frame, mission-critical applications will have computational requirements that are at least two orders of magnitude larger than current levels. Numerous recent developments in high-end computing (HEC) system design have the potential to overcome these deficiencies, via the use of parallel vector-based platforms; ultra-scale systems built from low-power components; scalable shared-address space architectures; and commodity clusters utilizing custom interconnects.

To understand the tradeoffs between the myriad of architectural design choices, comparative application performance data must be readily available to the HEC community at large, as performance of full-scale scientific applications is the final arbiter of supercomputing platforms. However, evaluating large-scale scientific applications using realistic problem sizes on leading HEC platforms is an extremely complex process, requiring coordination, communication, and management of a team of application scientists coming from highly disparate backgrounds. A detailed application analysis reveals algorithmic limitations to scalability and indicates where a particular code is poorly suited to certain architectural features. Direct comparison of applications across architectures, combined with a careful analysis of the problem and implementation characteristics, provides invaluable insight into the suitability of a given machine type for a particular class of scientific methods.

This special issue is a collection of papers that present and analyze performance results from the latest generation of HEC platforms, representing the world's most powerful supercomputers. Several of these papers were presented at the 2006 SIAM Conference on Parallel Processing for Scientific Computing held in San Francisco.

The first paper by Oliker et al. compare and analyze the performance of four application codes drawn from the areas of magnetic fusion, plasma physics, astrophysics, and materials science on leading vector and superscalar architectures. The platforms include Cray X1, X1E, Earth Simulator, NEC SX-8, and those utilizing IBM Power3, Intel Itanium, and AMD Opteron processors. The authors also present several major implementation innovations, and conclude that modern parallel vector systems have tremendous potential to attain impressive aggregate performance for well-structured computations.

The paper by Shan, Strohmaier, and Qiang analyzes the performance of a particle simulation code on seven different computing architectures: IBM BlueGene/L, Power3, Power5, Cray XT3, X1E, NEC SX-8, and AMD Opteron cluster. The code is actively used to model beam dynamics and thereby optimize the performance of high-energy colliders. The authors examine various performance issues such as workload partitioning to effectively use the computational resources, performance bottlenecks caused by architectural characteristics, and the utility of application codes to guide next-generation HEC architectures.

The IBM BlueGene/L paper by de Supinski et al. demonstrates how a supercomputer of unprecedented concurrency and computational power can enable scientific breakthroughs but also present unique challenges with scaling and optimization. The authors describe their experiences on this system at Lawrence Livermore National Laboratory with eight diverse applications including turbulence simulations, dislocation dynamics, adaptive mesh refinement, and semantic graph searching. Performance results demonstrate excellent scalability to very large processor counts, with a classical molecular dynamics application sustaining more than 107 TFlops/s.

The next paper by Alam et al. is a performance evaluation of the Cray XT3 system at Oak Ridge National Laboratory. The authors utilize micro-benchmarks, HPC Challenge (HPCC) kernels, and full-scale strategic applications of interest to the U.S. Department of Energy to compare and contrast XT3 performance against a suite of other architectures. Results demonstrate that this Cray supercomputer is an extremely viable platform to satisfy the diverse needs of a broad scientific computing community interested in climate, biology, astrophysics, combustion, and fusion.

Labarta, Rodriguez, and Badia report on the performance of the MareNostrum IBM PowerPC cluster located at the Barcelona Supercomputing Center. The paper critically evaluates performance of individual components to aggregate system behavior using micro-benchmarks and compact application kernels. The authors also present performance models that are generated using the Dimemas trace-driven simulation tool. They conclude that such characterizations are extremely useful to understand observed behavior and extrapolate performance, thereby driving designs of future architectures and algorithms.

The paper by Hood et al. benchmarks the Columbia supercomputer, a SGI Altix constellation installed at NASA Ames Research Center. The authors examine floating-point performance, memory bandwidth, and message passing communication speeds using HPCC kernels, NAS Parallel Benchmarks, and the Overflow computational fluid dynamics application. Results demonstrate performance improvements due to changes in interconnect type, processor speed, and cache size across the Altix product line. Substantial scalability improvements are shown if appropriate interconnect fabrics are utilized.

The next paper is by Nakano et al. presents a hierarchical framework for atomistic simulations of materials, and its performance on HEC platforms such as the IBM BG/L, AGI Altix, and AMD Opteron-based cluster. The framework is also tested on a geographically distributed grid of six supercomputer centers in the United States and Japan. Chemically reactive and non-reactive molecular dynamics simulations, with embedded quantum mechanical models, are utilized to understand microscopic mechanisms that drive macroscopic material properties. Simulations at unprecedented scales on BlueGene/L provide invaluable insights.

The paper by Tiyyagura et al. is a comprehensive performance evaluation of the NEC SX-8 system at the Stuttgart High Performance Computing Center (HLRS). The authors consider the HPCC suite of synthetic benchmarks, and five full-scale applications from the fields of fluid mechanics and quantum molecular dynamics developed by various European institutions. Their results demonstrate that this architecture is suitable for large-scale computational science in many different application fields, with sustained performance between 30 and 90 percent of peak.

The final paper in this special issue is by Wehner, Oliker, and Shalf. This work presents an extrapolation of the computational requirements for ultra-high resolution climate simulations, demonstrating that a credible kilometer-scale calculation would require at least ten petaflops of sustained computation. The article argues that impressive power- and cost-efficiency could be attained for a system of this scale via application-driven processor customization, and by leveraging the tremendous resources of the embedded-computing market.

Overall, this special issue contains one of the most comprehensive collections of high-performance application studies on modern HEC systems. These investigations provide invaluable insight in helping identify the suitability of particular architectures for given application classes while revealing performance-limiting system bottlenecks that can aid designers of the next generation systems. We believe that these kinds of studies lead to more efficient use of computational resources in both current installations and future designs.