

Preface

The growth of technical and commercial applications for high-performance computing, combined with the decline in the price of computers overall and the need to deliver new capabilities quickly, has led to the emergence of parallel computers and parallel computing. The newest of the desired applications include digital libraries, transaction processing, decision support, advanced multimedia, high-bandwidth networks, and data mining. Current high-performance applications from business and science also need a path to effective use of parallel computing. IBM has addressed this broad new territory with solutions that are general purpose and scalable and that operate readily within existing technical and commercial computing environments. These solutions are built on the firm foundation of RISC System/6000* workstation microprocessors with the AIX* operating system or on the well-established System/390* enterprise servers. The former solution—the IBM Scalable POWERparallel Systems*—is the focus of this special issue.

The availability of this new family of parallel systems has led to the resurgence of existing applications that previously ran only on mainframes and to the development of new applications in science and business that were not practical before. This issue contains an introductory essay and eight papers on the hardware, software, and systems that comprise the POWERparallel* SP1* and SP2* environments. We are indebted to Z. Barzilai and J. Moulic of the Thomas J. Watson Research Center, IBM Research Division, in Yorktown Heights, New York, for their efforts in laying the groundwork for this issue. We are also indebted to M. Snir and G. S. Almasi, also of the Thomas J. Watson Research Center, for their coordination and development of these papers.

The creation of new technologies, approaches, systems, and user opportunities is a mix of technical

advances and customer needs, some of which are available or anticipated and some not. In an essay intended to serve as an introduction, Wladawsky-Berger relates the recent history of the technology and user needs that made the IBM Scalable POWERparallel Systems both practical for IBM to build on a short schedule and advantageous for customers to have. Fundamentally, the approach augments the family of systems based on RISC System/6000 and AIX, so as to build on and add to the powerful technologies already developed and provide a scalable suite of systems solutions for customers, whether or not they already use those technologies.

The technology and architecture that form the more recent and advanced SP2 scalable parallel system are described in a paper by Agerwala et al. Seven principles for development of the SP2 are presented as a focus for understanding the evolution from scientifically oriented parallel systems to ones that deal effectively with both scientific and commercial applications. The key systems alternatives, competitive choices, SP2 choices, and rationale are explained. The authors also provide a system overview, explain the system components, and briefly discuss system performance using well-known benchmarks.

In order to bind a collection of processors that share work and communicate effectively at high speeds, there must be a high-performance switch. As parallel systems have become more powerful and flexible, such switches have evolved into sophisticated components, with high bandwidth and reliability concerns of their own, plus responsibility for relieving the processors of as many communication tasks as possible. The paper by Stunkel et al. presents the SP2 High-Performance Switch and its architecture, topology, adapter approach, performance, and support software.

The effective use of parallel computing systems depends on software that takes advantage of the multiple, communicating processors. Snir et al. describe the software development and execution capabilities created for the SP2 that allow programmers to exploit the power and flexibility of the parallel programming model with minimal or no code dependence on the specific SP2 system used for execution. In particular, the authors present the new Message-Passing Library that supports the communications paradigm of the SP2 and the new IBM Parallel Operating Environment that supports parallel program development and execution on the SP2.

Software developers for parallel computing also require a file system that allows effective use of resources in parallel and provides file services for the entire parallel system. The file system for the SP family, known as the AIX Parallel I/O File System, provides parallel access to partitioned files through portability, large file support, ease of use, and reliability. The authors, Corbett et al., provide the technical rationale and the developmental history for this file system, largely built on the Vesta file system project.

Gropp and Lusk relate their experiences migrating to and using an early 128-node SP1 at the Argonne National Laboratory and the subsequent migration to SP2 software running on SP1 nodes. Their conclusions provide useful information for other users and feedback for SP developers. They stress the speed with which the migrations can be accomplished—in days—and the effective coexistence of high performance, parallelism, and portability.

The next two papers focus on the effectiveness of the SP2 as demonstrated through implementations of the Numerical Aerodynamic Simulation (NAS) benchmarks. These benchmarks were developed by the National Aeronautics and Space Administration (NASA) specifically to test parallel supercomputing environments. In the first paper, Agarwal et al. show the parallel implementation of the five NAS kernel benchmarks. The results are compared with those of three other scalable parallel systems, revealing that the SP2 outperforms the others, often by wide margins.

In the second paper on NAS benchmarks for the SP2, Naik describes the parallel implementation of the NAS Parallel Benchmark BT (Block Tridiagonal), which is designed to test the effectiveness of

distributed memory systems. These tests, compared with those for two other systems, show results at the top of the range and almost always exceeding results for the others. They also show the delivered power in various situations, measured in billions of floating-point instructions per second (GFLOPS).

One aspect of high-performance systems that is especially important for commercial applications is the availability of equally high-performance database management systems. In this paper, Baru et al. describe DB2* Parallel Edition, the DB2 product that answers this requirement, executing all vital functions in parallel. The architecture and execution model of DB2 Parallel Edition provide processor and disk storage scalability, large database capacity, query optimization, execution time optimization, favorable transaction properties, parallel utilities, and reorganization for load balancing. Performance figures are shown for various complex queries within SP2 operating environments.

We are pleased to announce that, as a service to authors, we now provide an informational page for authors just ahead of the inside back cover. Also, as a service to readers, we recently expanded and reorganized our information about subscriptions and electronic access, and moved it to the inside back cover.

The next issue of the Journal will be a special issue on computer networking.

Gene F. Hoffnagle
Editor

*Trademark or registered trademark of International Business Machines Corporation.