

Preface

The most powerful computer systems of today—or supercomputers—exploit massive parallelism to achieve their superlative performance. A recent TOP500™ list of the highest-performing computers in the world, compiled by TOP500.Org (www.top500.org) in June of 2007 as this special issue was being prepared, presents some striking performance numbers. The fastest computer in the world, an IBM System Blue Gene/L™ Solution housed at Lawrence Livermore National Laboratory, harnesses the power of 131,072 processors working in parallel to achieve a performance of 280.6 teraflops per second on LINPACK, a linear algebra benchmark that is used to rank supercomputers. We note that a *teraflop* is a trillion floating-point operations such as multiplications or additions. The second, third, and fourth most powerful computers harness the power of between 23,000 and 41,000 processors. Indeed, the most powerful computer with fewer than 10,000 processors appears only in the eighth position on the list.

Massively parallel systems are not solely found in the area of supercomputing. Cluster-based systems often exploit massive parallelism. Google, for example, has reported using clusters with more than 15,000 processors.¹ However, massively parallel clusters are often not recognized as supercomputers either because they perform poorly on the LINPACK benchmark or because their owners have little interest in appearing on the list. Furthermore, grid-based systems have also exploited massive parallelism. The IBM-sponsored World Community Grid (www.worldcommunitygrid.org/index.jsp), for example, harnesses the power of more than 300,000 computers for humanitarian research on a number of important topics—many dealing with global health or climate issues. Clearly, we are in an era in which many computer systems can be described as massively parallel.

The scope of this parallelism is so far beyond the everyday experience of most computer professionals that it can only inspire questions about how massively parallel systems are used. These questions are:

- What applications can benefit from massive parallelism?
- How many applications benefit from massively parallel systems?
- Are there new applications that have been enabled by massive parallelism?

This issue of the *IBM Journal of Research and Development* is intended to provide insight into these questions.

¹L. A. Barrossa, J. Dean, and U. Holze, "Web Search for a Planet: The Google Cluster Architecture," *IEEE Micro*, 23, No. 2, 22–28 (2003).

Statistics from the TOP500 Web site also point to a clear trend: The most powerful computer systems are using increasingly massive parallelism to achieve their increased performance. Ten years ago, the minimum number of processors employed by any of the top ten supercomputers was 167; five years ago, it was 768; in June 2007, it was 9,600. This trend inspires further questions such as Will future technologies enable us to exploit even more massive parallelism, and what future applications can benefit from even more massive parallelism?

The concluding paper in this issue of the *Journal* provides insight into this question, because it presents an overview of the IBM System Blue Gene/P™ Solution. The Blue Gene/P architecture is the successor to the Blue Gene/L architecture, which was used to produce the most powerful and massively parallel supercomputer in the world. It will utilize even more massive parallelism than the Blue Gene/L Solution to achieve unprecedented performance.

We are fortunate that massively parallel systems have been made accessible for application prototyping and experimentation. Indeed, many of the applications described in this issue have been benchmarked at the IBM Blue Gene® Watson facility at the IBM Thomas J. Watson Research Center, which houses the fourth most powerful computer in the world—a Blue Gene/L computer that harnesses the power of 40,960 processors to achieve a LINPACK-benchmark performance of 91.29 teraflops.

One application domain that has been immensely propelled by massively parallel systems is the life sciences, which is represented in the first six papers of this issue. Some of these papers describe groundbreaking research into the processes that govern the operation of biological systems, while others describe improved methods for drug discovery.

Of great importance in life science research is the operation and function of proteins, which are the building blocks of life, as well as the operation of systems of proteins. As proteins form from amino acids, they are quickly shaped by molecular forces into a near-steady-state conformation (which is the native shape of the folded protein) and shape states around it—through a process called *protein folding*. Once in its native state, the free energy of the protein is minimized. While the massively parallel systems of today make simulation of protein folding possible, protein folding can consume months of processing time on a massively parallel system.

The first paper, by Raman et al., describes software that predicts the native shape of the folded protein—without fully simulating protein folding—by intelligently selecting a number of likely candidates and examining their free energy. While this can greatly reduce the

computations needed to determine the structure of the native state, it, too, is computationally demanding and can benefit from implementation on a massively parallel system.

Zhou et al. describe simulation of the structural consequence of a genetic mutation for the protein lysozyme. Such mutations can cause proteins to fold into an alternate native state, or *misfold*. Protein misfolding is a problem of considerable interest, because it is believed to be related to a number of important human disorders that include Alzheimer's and Parkinson's diseases.

Djurfeldt et al. and Kozloski et al. simulate brain functionality with different neuron-level models. Although these simulations involve relatively small areas of brain, they require massively parallel systems. Simulating the entire human brain with a suite of neuron-level models, as aspired to by the Blue Brain Project mentioned by Kozloski et al., will require even more massively parallel future systems.

Shave et al. and Pang et al. describe the application of massively parallel systems to "virtual screening," a process that computes the binding strength between a candidate protein and a large number of simple chemicals, or ligands. The information gathered in this process allows promising candidate drugs to be identified. This allows the drug testing process to focus on the more promising candidates and thus speed drug discovery.

Another application domain that has been energized by massively parallel systems is energy production, which is represented by the next three papers. Each describes work that has the potential to help satisfy the increasing need for energy in the world.

Calandra et al. describe the exploitation of massively parallel systems for *seismic imaging*—a fundamental technique for determining the composition of geological strata to aid petroleum exploration. Here, the ability to exploit massively parallel systems enables seismic imaging to produce more accurate and detailed models of the strata and hence more effective petroleum exploration. Commer et al. apply the power of massively parallel systems to the more-recent controlled-source electromagnetic (CSEM) three-dimensional geophysical imaging, which processes electromagnetic signals to gain further knowledge of the surface strata. The ninth paper of the issue, by Ethier et al., describes a massively parallel simulation of microturbulence in *nuclear fusion*—an energy source of great potential.

Another application domain that can benefit from massive parallelism is climate modeling, which offers us guidance on the management of our global environment. Although this domain has long benefited from large-scale parallelism, Dennis and Tufo demonstrate its ability to exploit massive parallelism. Given the complexity of the climate and the whole-earth scope of the simulation

required, the opportunity for use of even more massively parallel systems should be apparent.

The last application domain represented in this issue is cosmology. Here, Fisher et al. demonstrate that astrophysics can successfully exploit massively parallel systems. In this case, the scope of the region of interest of simulation is so large—the universe—that the need for even more massively parallel systems is unmistakable.

The next group of papers in the issue describe application enablers that are fundamental to large classes of applications. The first four papers in this set describe software bases that are used to simulate *molecular dynamics*—the effect of molecular forces on molecular systems. Although all four techniques implement molecular dynamics, each has application domains that ensure its continued importance.

Gygi describes the Qbox molecular dynamics software that won a prestigious 2006 Gordon Bell Prize in the peak performance category. We note that the Gordon Bell Prizes recognize groundbreaking achievements for performance and scalability in several categories on genuine and specific scientific applications. Fitch et al. describe the Blue Matter molecular dynamics software that has been a fundamental enabler of many biological simulations run at the Blue Gene Watson facility, including the lysozyme misfolding work. Bohm et al. describe Car-Parrinello molecular dynamics, which has many important applications including the simulation of new semiconductor materials. Kumar et al. describe NAMD, which is an important general molecular dynamics modeling application that is also notable for having received a 2002 Gordon Bell Prize.

The next paper, by Vranas et al., describes the implementation of quantum chromodynamics (QCD) on massively parallel systems. QCD simulates subatomic forces and has application to high-energy physics and cosmology. The particular implementation of QCD described in this paper is also notable for having won a 2006 Gordon Bell Prize for Special Achievement.

While the suite of papers appearing in this issue are excellent in content, are representative of current work on applications of massively parallel systems, and demonstrate a diversity of applications, they do not truly demonstrate the scope of applications for massively parallel systems. One of the frustrating tasks faced in preparing this issue was the turning away of many proposals for other papers on excellent work. Among the notable application topics absent from this issue are industrial system design, computational fluid dynamics, financial analysis, and digital content creation. Furthermore, applications on massively parallel clusters and grids are also prominently absent. The reader should appreciate that these omissions are due to limitations in the number of papers that could be accommodated—and

not to the absence of publication-worthy efforts. The application domain of massively parallel systems is indeed broad and increasing.

The last paper of this issue, by the IBM Blue Gene team, describes the IBM System Blue Gene/P Solution, which will enable even more massively parallel systems. A fully configured Blue Gene/P system would utilize more than 450,000 processors, which is more than three times the parallelism of today's most powerful supercomputer. Furthermore, a fully configured Blue Gene/P system would achieve a peak performance in excess of 1.5 petaflops, which is more than five times the performance of the most powerful supercomputer of today.

What will be the impact of this greater scale of parallelism? Certainly, there will be opportunities to reduce the time to solution of many of the important problems of today, such as those mentioned in this issue. Some applications that today represent leading-edge research may enter routine commercial usage. Certainly, there will also be opportunities for solutions to new problems that will be possible only with more massively parallel systems. One of the great trends of the information era has been the substitution of computer simulations for physical simulations, and more massively parallel systems will allow more massively detailed simulations than are currently possible. Certainly, there will be opportunities for inventions that will allow researchers to create new application enablers that will allow extremely demanding applications to utilize even greater parallelism than that which exists today.

Fred Mintzer
Program Director, Blue Gene Watson
Associate Director, Deep Computing Institute
IBM Thomas J. Watson Research Center

Guest Editor