

Preface

This group of papers pertains to *deep computing*—high-performance computing (extending to the teraflops range) applied to problems that require significant computing resources in order to facilitate relatively complex decision-making in fields ranging from business to the sciences. The papers capture a few illustrations of deep computing that cut across traditional disciplines. A previous double issue of this journal (Volume 45, Number 3/4, May/July 2001) and a companion issue of the *IBM Systems Journal* (Volume 40, Number 2, 2001) contained papers that covered aspects of deep computing for the life sciences.

Deep computing requires not only massive amounts of computing, but also the processing of vast amounts of data and the use of sophisticated means for the interpretation of results. Most often, scalability is also important—i.e., the use of scalable hardware, software that is tuned to take advantage of that hardware, and algorithms that match applications to the hardware to achieve optimum performance. In order to address societal problems, the systems must also be able to implement extensive computer interactions involving multimedia, immersive visualization, audio input, etc.

New algorithmic approaches are being developed that take into account computer architecture in order to address problems of general interest and answer questions that have impact beyond just scientific inquiry. The paper by Crivelli and Head-Gordon discusses a load-balancing technique to exploit the use of parallel architectures for addressing protein structure problems. In their paper, Fann et al. investigate multiwavelet approaches and apply them to areas of computational chemistry, computational electromagnetics, and fluid dynamics. Newns et al. discuss new algorithmic techniques they have developed for improving magnetic recording simulation.

The applicability of deep computing to astrophysics is illustrated by the paper by Lake et al. The authors apply new algorithms coupled with deep computing to advance our understanding of the *N*-body problem.

Coupling *specialized* hardware (MDGRAPE-2) with algorithmic techniques accelerates the ability to carry out micromagnetic simulations of an array of dipoles, as Elmegreen and colleagues describe in their paper. They coupled MDGRAPE-2 hardware with a commercially available system, using special algorithms to achieve their results.

The paper by Kramer et al. pertains to issues associated with the massive amounts of data found in a deep computing environment and the network infrastructure needed to support large data movement.

Deep computing is not just about a single machine, but also about a system that is becoming a significant computational environment—*grid computing*. In their paper, Grimshaw et al. discuss the grid and its evolving ecosystem.

Kirk E. Jordan

Emerging Solutions Executive
IBM Life Sciences Solutions
Cambridge, Massachusetts

Guest Editor