Scaling Science Applications on Blue Gene

William D. Gropp, Wolfgang Frings,
Marc-André Hermanns, Ed Jedlicka, Kirk E. Jordan,
Fred Mintzer, Boris Orth

Scaling Science Applications on Blue Gene

William D. Gropp[1], **Wolfgang Frings**[2], **Marc-André Hermanns**[2], **Ed Jedlicka**[1],
Kirk E. Jordan[3], **Fred Mintzer**[3], **and Boris Orth**[2]

[1] Argonne National Laboratory
Mathematics and Computer Science Division
9700 South Cass Avenue, Argonne, IL 60439, USA
E-mail: {gropp, jedlicka}@mcs.anl.gov

[2] Forschungszentrum Jülich,
John von Neumann Institute for Computing, 52425 Jülich, Germany
E-mail: {w.frings, m.a.hermanns, b.orth}@fz-juelich.de

[3] Deep Computing, IBM Systems and Technology Group
1 Rogers Street, Cambridge, MA 02142, USA
E-mail: {kjordan, mintzer}@us.ibm.com

Massively parallel supercomputers like IBM's Blue Gene/L offer exciting new opportunities for scientific discovery by enabling numerical simulations on an unprecedented scale. However, achieving highly scalable performance is often not straightforward as the system's extraordinary level of parallelism and its specialized nodes present challenges to applications in many areas, including: communication efficiency, memory usage, and I/O.

This mini-symposium aimed to highlight some of the remarkable scaling and performance results achieved, and bring together users and system experts to discuss possible solutions to yet unresolved issues. It featured speakers whose applications have run at large scale on the 8K node system at John von Neumann Institute for Computing at Forschungszentrum Jülich and the 20K node system at IBM's Watson Research Center. The talks provided a sampling of the different applications and algorithms that have successfully run on Blue Gene. The speakers discussed best practices, particularly with respect to scaling to tens of thousands of processes, and challenges faced in using BlueGene/L to massive scale, and they showcased some of the breakthrough science that has already been achieved.

Mini-Symposium Contributions

Jörg Schumacher of Ilmenau Technical University reported on his work on turbulence in laterally extended systems. He has studied three-dimensional turbulence in a very flat Cartesian cell using periodic boundaries in the lateral directions and free-slip boundaries in the vertical one. The pseudospectral simulations were conducted on Blue Gene/L, the memory limitations of which required a two-dimensional parallelization of the numerical scheme. Schumacher reported on performance and scaling tests on up to eight racks and presented first physical results for the case of convective turbulence.

Kevin Stratford from Edinburgh Parallel Computing Centre presented an implementation of Lees-Edwards sliding periodic boundary conditions used to perform simulations of sheared binary fluids. In conjunction with domain decomposition and message passing he was able to perform shear simulations of significant size and duration. Stratford discussed the scaling and performance of these large calculations on the Blue Gene/L architecture.

Ilian T. Todorov talked about DL_POLY_3, a new generation software package for molecular dynamics simulations developed at Daresbury Laboratory. After an overview of

the package he reported on the weak scaling behaviour of DL_POLY_3 and discussed its dependence on force field complexity. Todorov showed benchmark results on massively parallel architectures like BlueGene/L and Cray XT3, and concluded with a review of the challenges and successes in applying DL_POLY_3 in two different simulation studies: (i) radiation damage cascades in minerals and oxides, where the problem size (length scale) is of importance, and (ii) biochemical simulations, where long time scale simulations are required.

Erik Koch from the Institute for Solid State Research at Forschungszentrum Jülich reported on simulations of materials with strong correlations. In these simulations the Lanczos method was used to calculate ground state and dynamical properties of model Hamiltonians approximating the Hamiltonian of the full, non-perturbative many-body problem. Koch presented an efficient implementation for Blue Gene that overcomes the method's principal problem of non-local memory access and scales very well. Besides correlated electronic systems it may also be applied to quantum spin systems, which are relevant for quantum computing.

Jeffrey Fox of Gene Network Sciences reported on massively parallel simulations of electrical wave propagation in cardiac tissue. The software of his group is being used in an on-going study to connect drug-induced modifications of molecular properties of heart cells to changes in tissue properties that might lead to a rhythm disorder. The performance of the code was studied by simulating wave propagation in an anatomically realistic model of the left and right ventricles of a rabbit heart on Blue Gene partitions of up to 4,096 processors. Judging from these results, the Blue Gene architecture seems particularly suited for cardiac simulation, and offers a promising platform for rapidly exploring cardiac electrical wave dynamics in large spatial domains.

Hinnerk Stüben from Zuse Institute Berlin reported on how he achieved, for the first time, a sustained performance of more than 1 TFlop/s in everyday simulations of Lattice QCD on the IBM BlueGene/L system at Jülich and the SGI Altix 4700 at LRZ, Garching. He compared the performance of Fortran/MPI code with assembler code, the latter of which allows for a better overlapping of communication and computation, and prefetching of data from main memory. Lower-level programming techniques resulted in a speed-up of 1.3-2.0 compared to the Fortran code. His program scales up to the full eight Blue Gene/L racks at Jülich, at a remarkable maximum performance of 8.05 Tflop/s measured on the whole machine.

Henry M. Tufo from the National Center for Atmospheric and Climate Research (NCAR) presented HOMME, the High-Order Method Modeling Environment, developed by NCAR in collaboration with the University of Colorado. HOMME is capable of solving the shallow water equations and the dry/moist primitive equations and has been shown to scale to 32,768 processors on BlueGene/L. Tufo reported that the ultimate goal for HOMME was to provide the atmospheric science community a framework upon which to build a new generation of atmospheric general circulation models for CCSM based on high-order numerical methods that efficiently scale to hundreds-of-thousands of processors, achieve scientifically useful integration rates, provide monotonic and mass conserving transport of multiple species, and can easily couple to community physics packages.

Last but not least, Kirk E. Jordan from IBM presented interesting early experiences with Blue Gene/P, the next generation of the IBM Blue Gene system.