John von Neumann Institute for Computing



Computer Science and Numerical Mathematics

Michael Griebel

published in

NIC Symposium 2006 ,
G. Münster, D. Wolf, M. Kremer (Editors),
John von Neumann Institute for Computing, Jülich,
NIC Series, Vol. **32**, ISBN 3-00-017351-X, pp. 297-298, 2006.

© 2006 by John von Neumann Institute for Computing
Permission to make digital or hard copies of portions of this work for personal or classroom use is granted provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise requires prior specific permission by the publisher mentioned above.

http://www.fz-juelich.de/nic-series/volume32

Computer Science and Numerical Mathematics

Michael Griebel

Institut für Numerische Simulation
Wegelerstr. 6, 53115 Bonn, Germany
E-mail: griebel@ins.uni-bonn.de

Computer science aspects and numerical simulation on modern parallel hardware platforms such as the JUMP are discussed in the following section. Here, we put the emphasis on numerical software aspects and their interplay with modern hardware, on the question on how to obtain a cheap crossbar switch with full bisectional bandwidth and on tools which help the programmer of parallel codes on multiprocessor systems.
In their contribution *FEAST: Development of HPC technologies for FEM applications*, Ch. Becker, S. Buijssen, H. Wobker and S. Turek discuss current trends in software developments for PDEs for hierarchical adaptive finite elemente methods. Here, the usually employed data structures and algorithms are not tuned to the specific performance facilities of modern hardware platforms and thus the obtained computational efficiency is far away from the peak rate of the respective processors. To this end hardware-oriented numerics is necessary, i.e. cache-oriented techniques must be employed and locally structured data must be exploited. The authors present such techniques for the FEM package FEAST ("Finite Element Analysis & Solution Tools") which numerically solves fluid dynamics and structure mechanic problems. They present recent results on the JUMP for their hardware oriented multigrid/domain decomposition solvers and discuss scalability, clustering and recursion for their approach. Altogether such developments are important to obtain a good portion of the peak performance on modern parallel hardware platforms for advanced numerical techniques which involve multiscale methods, adaptivity and domain decomposition.

In their contribution *A Scalable Ethernet Clos-Switch*, Norbert Eicker and Thomas Lippert discuss techniques to realize large crossbar-switches which are built from standard Gigabit-Ethernet switches as components. They follow an idea of Charles Clos from 1953 who proposed a similar technique in the context of telephone switching networks. Here, a special topology of cascaded crossbar-switches results in full bisectional bandwidth. The authors show how to build such a large crossbar-switch with up to 1152 ports which achieves full bandwidth at a cost of about 125 Euro per port and has a point-to-point throughput of 220 MB/s where the latency is less than 10 μsec. This approach turns out to be superior to any existing monolithic switch on the market.

In their contribution *The IBM eServer pSeries 690 as a Research Instrument for Computer Scientists*, G. Juckeland, M. Kluge, R. Müller-Pfefferkorn, W.E. Nagel and B. Trenkler discuss some features of the IBM p690 system at the Forschungszentrum Jülich which they tested and analysed with the PARbench, BenchIT and EPcache environment. Their results might help users and administrators to optimize performance and system behavior in their specific applications. Furthermore, they allow in an automatic way to detect MPI overhead which assists a programmer to create efficient parallel code on the IBM p690.