(52)(53)

ELSEVIER

Future Generation Computer Systems 13 (1997/98) 247–249

FGCS
FUTURE GENERATION COMPUTER SYSTEMS

Guest editorial

High Performance Computing and Networking Europe 1997

Alfons Hoekstra [a,*], Arndt Bode [b]

[a] *University of Amsterdam, Netherlands*
[b] *Technical University, Münich, Germany*

It is not easy to set a starting date for the field of High Performance Computing and Networking (HPCN). For more than two decades many researchers and engineers have dedicated their time and energy to parallel computer architectures, parallel computing, and parallelization techniques, all with very impressive results. In this sense, HPCN has been a well accepted topic in computer science research for a long time. The real impact of HPCN on the development of modern society has only been recognized in the beginning of the 1990s. Indeed, 1991 has been an important date for the political support of HPCN activities. In that year the US congress passed the High Performance Computing and Communications Act, and the first Rubia report was presented to the European Commission. This report laid the foundation for funding of HPCN activities within the European Esprit program. Comparable initiatives were taken in Japan with the Real World Computing program.

An important argument to start these HPCN programs was a strong application pull. A large number of so-called grand challenge applications were identified, based on their utmost importance for improving economic welfare of nations and their inhabitants. Without exception, these applications need computing power and memory storage that is far beyond the capabilities of supercomputers that were around in the beginning of the 1990s. To set the stage, sustained computing speeds of 1 Tflop/s using memories of 100 Gbyte were marked almost as a holy grail for HPCN research.

We have come a long way since then, as once more became clear during the final session of the High Performance Computing and Networking Europe 1997 conference, which was held from April 28 to 30, 1997 in Vienna, Austria. Mr. Alex Larzelere, director of the office of the Accelerated Strategic Computing Initiative (ASCI), reported among others on the sustained performance of the Linpeak benchmark of just over 1 Tflop/s, measured in December 1996 on the Intel ASCI Red computer. He also presented growth paths of the ASCI Red and ASCI Blue systems, and predicted that very soon we will witness sustained performance of 1 Tflop/s and more for real production codes.

With these spectacular results in the back of our head, will we in 10 years from now mark 1997 as the year that HPCN was finalized? Definitely not, as was demonstrated by the large number of new, exiting, and innovative results that were presented during the HPCN Europe 1997 conference! The field of HPCN has matured, but has also grown beyond its original paradigms.

The increasing maturity of HPCN results in an ever faster uptake of HPCN technology in industry and

* Corresponding author. Address: Faculteit der Wiskunde Informatica, Natuurkunde en Sterrenkunde, Vakgroep Informatica, Kruislaan 403, 1098 SJ Amsterdam, Netherlands. Tel.: +31 020 535 7463; fax: +31 020 525 7490; e-mail: alfons@wins.uva.nl.

PII S0167-739X(98)00026-5

commerce. The "Industrial End User track" of the HPCN Europe 1997 conference presented many examples of successful use of HPCN in industry, ranging from industrial inspection, via telecommunications applications, to medical imaging and virtual reality. Furthermore, the "General End User track" presented applications of HPCN in the domains of physics, chemistry, and engineering, but also in the field of finance, transaction processing, and data mining.

The "Computational Science" and "Computer Science" tracks presented new scientific results and emerging HPCN technologies. The latter track focused on a wide range of "traditional" computer science subjects, such as (simulation of) parallel architectures, interconnection networks, cache coherency, message passing, data parallel environments, and parallelizing compilers.

The "Computational Science" track focused on new results on modeling and simulation of systems, be it natural systems (from physics, biology, etc.) or man-made systems (cars, computer systems, etc.). New solving methodologies and strategies to map application models to HPC systems were presented. Among others, new optimized preconditioners for parallel iterative solvers, problem solving environments for PDE problems, but also new grid-based methods to solve fluid flow (the Lattice Boltzmann approach) or to model soil contamination were presented. Furthermore, a number of contributions concentrated on mapping of parallel programs and resource management for optimized (static or dynamic) load balancing of heterogeneous distributed computing systems. Finally, a number of contributions on metacomputing environments were presented.

This special issue of FGCS is dedicated to the HPCN Europe 1997 conference, which was embedded in the overall HPCN event, also hosting a large exhibition. The HPCN Europe 1997 conference contained a total of 93 presentations, divided over four tracks as mentioned above. Furthermore, another 38 papers were presented during poster sessions. The series of yearly HPCN conferences started in 1993 in Amsterdam. They always focused on bringing together HPCN end users from industry and commerce with researchers in the field of HPCN, in order to strengthen and speed up the uptake and use of HPCN in real applications.

This special issue contains a selection of 15 papers taken from the conference. We have tried to produce a balanced issue, reflecting the overall composition of the conference, and highlighting a number of new and original applications of HPCN and new research directions in HPCN.

The first three contributions present applications of HPCN. Hancock and Hubbold present a distributed parallel volume renderer server, which is efficiently operated over relative slow network links. Seinstra et al. present a parallel simulation of ion recombination. As the amount of work decreases during the simulation, they have applied the idea of process reduction in order to achieve efficient execution. Breitler et al. show how a computational intensive problem from finance, a portfolio choice problem, can benefit from HPCN.

The next two papers describe application of the inherently parallel, cellular automaton computational model. Guidec et al. introduce their Parflow++ environment to simulate radio wave propagation in cities. Spezzano and Talia present the CARPET language, which is a high level language to express cellular automata models, and which is used to program their parallel cellular automaton environment CAMEL. As an example, modeling and simulation of soil contamination is described.

Issman and Degrez introduce two non-overlapping preconditioners for parallel implicit Navier–Stokes solvers and analyze their performance. De Sturler and Loher show how efficient parallel iterative solvers for irregular sparse matrices can be implemented in High Performance Fortran.

Two papers deal with the BSP paradigm. First, Hill and Skillicorn analyze two frequently encountered criticisms on BSP, and argue that both do not hold. Next, Huang and McColl introduce a two-way algorithm for tridiagonal systems, and analyze its performance using the BSP model. Their predictions are compared with measurements on Sun networks.

Sato et al. present their on-going work on communication libraries for global and cluster computing. Ninf allows users to access global resources. PM is

a high performance communication library for workstation clusters connected through Myrinet. Barak and La'adan describe their MOSIX multicomputer operating system (an enhancement of BSD with adaptive resource sharing). They present performance of the resource sharing algorithms on PCs connected by Myrinet, and of a number of large-scale applications. Cordsen and Schröder-Preikstat present VOTE, a communication support system that provides coexistence of shared-memory and message passing communication.

The final three papers deal with automatic parallelization techniques and compiler technology. Fahringer presents a number of new symbolic analysis techniques to examine non-linear array index functions, complex loop bounds, and to deal with unknown problem, array, or machine sizes. Stöhr and O'Boyle describe a new compiler algorithm to reduce the number of barrier synchronizations in parallellized programs. The new algorithm was implemented in a prototype compiler and next applied to three benchmarks. Karkowski and Corporaal investigate the potential of dynamic loop analyzing techniques for coarse grain loop parallelization for "embedded" applications. Their technique is applied to a set of Audio Signaling Processing programs and on a MPEG-2 video encoder program.

We would like to thank everybody who made this special issue possible: the authors, the organizing committee and the program committee of HPCN'97.