Supercomputing Around the World

Allen D. Malony

Department of Computer and Information Science
University of Oregon
Eugene, Oregon 97403-1202

Abstract

Supercomputing is rapidly becoming a global phenomenon. In keeping with the "Voyages of Discovery" theme of the Supercomputing '92 conference, representatives of supercomputing endeavors from around the world meet in this mini-symposium to speak on national and international supercomputing activities.

1 Introduction

The potential of supercomputing technology to address computationally challenging scientific and engineering problems knows no national bounds, nor are the application areas unique to selected regions of the world (e.g., climate and ocean modelling). Although the United States and Japan have historically been the leaders in the development and marketing of supercomputing technology, the high-performance computing and communications needs of the world community has given rise to international supercomputing endeavors. Some efforts involve the use of supercomputing technology to solve problems of national importance, while others are developing supercomputer hardware and software products to compete in global markets. Whatever the goals of supercomputing initiatives throughout the world - scientific, economic, even political - the exchange of information about these activities is an important step to insuring the success of supercomputing on a global scale.

This mini-symposium brings together international representatives from five areas of the world - Europe, Taiwan, Australia, India, and Brazil - to discuss supercomputing activities in countries that have been under-represented at the Supercomputing conferences in the past. It is hoped that this mini-symposium will help to broaden international participation and exposure in Supercomputing conferences in the future. The abstract of each speaker's talk is presented below.

High Performance Computing and Networking in Europe: A Perspective

David J. Wallace

Director
Edinburgh Parallel Computing Centre
University of Edinburgh
Scotland, United Kingdom
d.wallace@edinburgh.ac.uk

The Commission of the European Communities has supported and is developing a number of Programmes of relevance to high performance computing and networking (HPCN). Within Esprit, a sustained progamme of research and development of components and systems has stimulated industrial and scientific activity. In 1991, the Rubbia committee established by the Commission to advise on HPCN presented a report recommending an annual investment of more than $1 billion per annum. The principles of this report have been accepted and the High Performance Computing Advisory Committee reconvened to recommend on how the first report should be implemented. In a complementary initiative, scientists self-organised to develop the European Teraflops Initiative. An industrial grouping, EI3, has also developed substantial proposals for an HPCN programme. Within the Science Directorate DG XII, the Human Capital and Mobility Programme is also a source of support for the education and training aspects central to the future successful exploitation of HPCN. This talk will review these developments and their relationships, from the personal standpoint of the speaker; it is not an official overview of policy.

1063-9535/92 $3.00 © 1992 IEEE

Supercomputing Environment in Taiwan, R.O.C.

Prof. San-Cheng Chang

Director
National Center for High-Performance Computing
National Science Council
National Taiwan University
4-1 Roosevelt Road
Taipei, Taiwan, Rep. of China
ascchang@nchc.ntu.edu.tw

This presentation describes the planning philosophy, current status, and future plans on the supercomputing environment in Taiwan, Republic of China. Here the national supercomputing environment has gone through a preparation phase of 3-years, and the facility is expected to begin operation in late 1992. This project is carried out by the the National Center for High-Performance Computing (NCHC) under the National Science Council of the government.

In NCHC, state-of-art computing systems and major application software will be installed, and universities and research institutions will access this facility via existing TANet (Taiwan Academic Network). This TANet will be linked via 256K-baud line to NSFNET of U.S in the near future. What makes NCHC unique is the body of research staff. NCHC users are supported by NCHC staff with advanced degrees and expertise in various area of high-performace computing. NCHC is well-acquainted with popular software systems and is capable of developing special codes when users' requirement is beyond the reach of commercial systems. This type of well-structured layered expert support goes far beyond the conventional consulting services. NCHC will be staffed to over 100 employees, among which over 60 will be dedicated to application supports and developmental research.

Beyond the hardware and software platform, the overall infrastructure of the research environment is a key focus for NCHC. NCHC will set up additional international exchange channels for operation and application experiences with major international institutes. This includes introducing intensive short courses in selected applications, provide residence research positions locally and abroad, and work with vendors on advanced software.

NCHC also hopes to demonstrate how new computer architecture may be integrated with existing computing platform. Pilot systems such as networked high-end workstation and highly to massively parallel systems are included in the plans for 1993. Currently NCHC already has a basic set of distributed programming tools, and Taiwan has more MPP systems installed than other Asia-Pacific countries except Japan. This provides a sound base for entrance into computing environments for next decade.

In summary, NCHC opens the door of high-performance computing for Taiwan and leads the computing environment with a global view in terms of space and time. Through NCHC, Taiwan becomes a more active member in the community of high-performance computing.

Supercomputing in Australia

Prof. Michael A. McRobbie

Executive Director
Centre for Information Science Research
The Australian National University
Canberra, Australia
michael.a.mcrobbie@arp.anu.edu.au

Though Australia has a modern scientific infrastructure in a wide variety of areas (e.g. astronomy, biotechnology, marine sciences and medical research), it has only recently begun to develop a significant supercomputer infrastructure. However, unlike many other industrialised countries, it does not have a national coordinated inter-agency policy for the development of such infrastructure, nor has it yet made an attempt to identify computationally-intensive problems of particular importance to Australia's economic and social well-being.

The first supercomputer in Australia was a Cyber 205 obtained in 1983 by the Commonwealth Scientific and Industrial Research Organisation - Australia's national laboratories. After a somewhat chequered career, this machine was decommissioned around 1988.

In 1987 the Australian National University (ANU) established a research and development agreement with Fujitsu as part of its acquisition of a VP-100 supercomputer - the first supercomputer acquired by any university in Australia - since upgraded to a VP-2200/10. This research agreement has grown to be one of the largest that Fujitsu has with any organisation in the world including in Japan. Under this agreement Fujitsu also placed at ANU its first AP-1000 experimental Sparc-based distributed memory parallel supercomputer as part of a research program in parallel computing. This remains the only AP-1000 outside of

Fujitsu Laboratories. In 1990 ANU acquired the first and still only CM-2 in Australia, and in 1992 acquired the first CM-5 outside of the USA. ANU remains the largest supercomputer site in Australia and one of the largest in the region.

Since ANU acquired its first supercomputer, other universities, research organisations and recently a number of companies, have acquired a variety of supercomputers and mini-supercomputers including a number of small to medium sized machines from Convex, Cray, Fujitsu, IBM and MasPar. All of these sites are interconnected with other research organisations and industry in Australia and overseas by the Australian Academic Research Network (AARNet), which has backbone speeds of 2 Mbits per second. Recent federal funding has been directed towards the enhancement of network services in Australia and there is an expectation that some high speed testbeds, on the model of the US gigabit networks, may be put in place in the near future.

As the unique supercomputing environment at ANU shows, Australia is able to take advantage of being a "neutral" meeting ground for some of the best American and Japanese supercomputing technology.

India's National Initiative in Massively Parallel Supercomputing

Vijay P. Bhatkar

Centre for Development of Advanced Computing
Pune University Campus, Ganeshkhind
Pune 411007, India
bhatkar@parcom.ernet.in

Establishing the Centre for Development of Advanced Computing (C-DAC) in August 1988 marks the launching of India's national initiative in parallel supercomputing. Launching of C-DAC marks the beginning of a new phase in the development of computational science and technology in India. C-DAC's three year first mission goal was to design, develop, and bring into commercial production a state-of-the-art parallel supercomputer with peak computing power exceeding 1,000 Mflops, and demonstrate applications of national importance on the target machine. The project had three main components - Technology, Applications, and Research - which represented, respectively, "height", "weight", and "depth" of the project implementation.

The prime delivery of the mission has been the 256-node PARAM parallel supercomputer complete with parallel disk array, advanced programming environment and a range of applications in science and engineering. PARAM has been made available in two series of configurations. The PARAM 8000 series of replicated scalar processor machines can be configured with 16, 32, 64, 128, 256, or, if required, 1024 compute nodes based on transputers. The PARAM 8600 is a PARAM 8000 machine additionally equipped with vector processing capability based on i860 vector nodes. The 256-node PARAM 8000 has a peak computing power exceeding 1 GFlops / 7500 MIPS, aggregate main memory of 1 GBytes and auxiliary memory of over 20 BGytes in the form of parallel disk arrays. The PARAM 8600 coherently integrates vector nodes into the architecture boosting the peak computing power of a 256-node PARAM 8600 to over 5 GFlops.

In order to provide scalable and affordable computing power on industry standard platforms, C-DAC has also developed several application accelerators which bring high-performance computing power to PCs, SUNs, VAXes and VME bus-based systems. These application accelerators have been realized through add-in boards and a TRAM motherboard approach. Application specific TRAMs have been developed for SCSI, graphics, frame grabbing, ADC, signal processing, etc. The range of application accelerators developed is comparable to that available currently in Europe.

PARAS is a comprehensive parallel programming environment that has been developed for PARAM and similar message passing parallel computers. It is comprised of the host servers, compilers, run-time environment, parallel file system, on-line debugger and profiling tools, graphics and visualization support, networking interfaces, off-line parallel processing tools, program restructurers, and libraries. The entire system software, which has been developed from the root level, consists of one and a half million lines of source code and represents the largest system software effort ever done in India. PARAS compares well with similar programming environments that are available internationally.

C-DAC has initiated a significant applications development program in the area of parallel computing, involving over 20 organizations for joint applications development on PARAM. Currently, over 40 application kernels are running on PARAM in the areas of remote sensing, image processing, signal processing, launch vehicle dynamics, computational fluid dynamics, finite element modelling, oil reservoir modelling, seismic data processing, power systems analysis, process control, robotics, computational physics, computational astrophysics, computational mathematics, scheduling and optimization, molecular modelling, bio-technology, and graphics and visualization. The plausible sustained performance obtained has demonstrated the general purpose applicability of PARAM in science and engineering applications.

PARAM and complementary parallel processing products have already been commercialized. Currently, there are 10 installations of the PARAM system in leading academia and research institutions in India and over 100 parallel processing accelerator platforms in the field. A few systems have also been exported to advanced countries. The 256-node PARAM facility at Pune is being operated as a national parallel supercomputing facility.

With 200 scientists and engineers, C-DAC today represents one of the "critical mass" groups in massively parallel supercomputing. C-DAC's perspective plan for the Second Mission is to develop a teraflop, massively parallel supercomputer architecture by the mid-1990s.

Supercomputing in Brazil

Dr. Kepler Oliveira

General Director
Center for Supercomputation
Universidade Federal do Rio Grande do Sol (UFRGS)
Av. Paulo Gama, 110
90049 Porto Allegre, Brazil
kepler@ifl.ufrgs.br

At the Center for Supercomputation, we have just installed the first supercomputer in Brazil, a Cray YMP2E232. The Center is open to all the academic and research institutions in the country, and to a lesser degree, to industries. The Center is connected to all the institutes within the University via an optical fiber, and to the National Research Network, that connectes all academic and research instituions in Brazil, which in turn is connected to Internet, Bitnet and HEPnet. As a part of the Center, we have also installed a Scientific Visualization Laboratory, based on Silicon Graphics workstations.

The projects that are already using the supercomputer include astrophysics, biological sciences, computational science, engineering, and physics. We are still in the process of hiring software development personal, and acquiring the application software needed. This is a necessary step if the Center is to be used by industries in the development and simulation of new products or procedures.

We hope supercomputing will allow Brazil to improve its share of scientific and technological research, necessary to maintain and increase its competitivety in the international market.