Cluster Computing Workshop

organized at Euro-Par´ 2000

European Parallel Computing Conference

August 29th to September 1st, 2000 in Munich, Germany

Recent Advances in Cluster Computing

Cluster computing can be described as a fusion of the fields of parallel, high-performance, distributed, and high-availability computing. It has become a popular topic of research among the academic and industrial communities, including system designers, network developers, algorithm developers, as well as faculty and graduate researchers. The use of clusters as a application platform is not just limited to the scientific and engineering area; there are many business applications, including E-commerce, that are benefiting from the use of clusters.

There are many exciting areas of development in cluster computing. These include new ideas as well as hybrids of old ones that are being deployed in production and research systems. There are attempts to couple multiple clusters, either located within one organisation or situated across multiple organisations forming what is known as a *federated clusters* or *hyperclusters*. The exploitation of federated clusters (clusters of clusters) as an infrastructure can seem to be approaching the area of the increasingly popular GRID infrastructure.

The concept of portals that offers web-based access to applications running on clusters is becoming widely accepted. Such computing portals, offering access to scientific applications online, are known as scientific portals. PAPIA (Parallel Protein Information Analysis system) developed by the Japanese Real World Computing Partnership (RWCP) is an example of one. PAPIA allows scientists to have online access to a Protein Data Bank (PDB) in order to perform protein analysis on clusters. Anyone can perform the analysis of protein molecules and genetic DNA sequence from anywhere, at any time using any platform. In the future, we will see many such applications that exploit clustering and the Internet technologies for scientific discovery.

The TopClusters project is TFCC collaboration with the TOP500 team. Plans are underway to build a database to record the performances of the most powerful cluster systems in different areas. These is great interest in this area as clusters are being used as platforms to host a diverse range applications, including scientific computing, web serving, database, and mission critical systems. Each of these areas has its own specific requirements. For example scientific applications are driven by floating-point performance whereas database applications are driven primarily by system I/O performance. The TopClusters project will measure these parameters (Mflop/s, I/O, TPC, MTBF, etc.) and use this to try and understand more fully the key aspects that need to addressed when building clusters for new and emerging applications.

The IEEE Task Force on Cluster Computing (TFCC) is acting as a focal point and guide to the current cluster computing community. The TFCC has been actively promoting the field of cluster computing with the aid of a number of novel projects, for example we have an educational activity that has a book donation programme, holds forums for informal discussion, helps guide R&D work both in academic and industrial settings through workshops, symposiums and conferences.

The recent developments in high-speed networking, middleware and resource management technologies have pushed clusters into the mainstream as general purpose computing system. This is clearly evident from the use of clusters as a computing platform for solving problems in number of disciplines. It also raises a number of challenges that cluster systems need to address with respect to their ability to support for example:

- System architecture.
- Heterogeneity.
- Single system image.
- System scalability.
- Resource management.
- System administration.
- Performance.
- Reliability.
- Application scalability.
- Management and administration of hyperclusters.

Based on the aspects allready mentioned, a number of challenges listed above are among the issues addressed by the 13 research papers that we have selected from 27 contributions for the EuroPar'2000 Cluster Computing Workshop. The program of the Workshop presents articles which demonstrate both theoretical and practical results of research works and new developments regarding cluster computing. The following papers have been accepted for presentation and discussion in the workshop, they cover:

- F. Rauch, C. Kurmann and T. M. Stricker presents an analytical model that guides an implementation towards an optimal configuration for any given PC cluster. The model is validated by measurements on a cluster using Gigabit- and Fast Ethernet links.
- W. Hu, F. Zhang, H. Liu and G. Chen and of F. Zhang and Z. Zhang describe topics belonging to the SMP and DSMP topics. While the first one introduces an SMP protocol for the home-based software DSM system JIA-JIA, the second paper is the design of a compiler aided software DSM system, called CA-JIAJIA, which retains the ease of programming of shared memory, while exploiting the compile-time collected information to reduce various overhead of run-time consistency maintenance and im-plements write vector based modified blocks prefetching.
- R. Cunniffe and B. A. Coghlan proposes a framework for cluster management which enables a cluster to be more efficiently utilized within a research environment.
- V. Shurbanov, D. Avresky, P. Mehra and W. Watson describe in their paper the performance implications of several end-to-end flow-control schemes clusters based on the ServerNet system-area network.
- The authors of the next paper H. Pedroso and J. G. Silva introduce the system WMPI as the first implementation of the MPI standard for Windows based machines.
- The goal of the paper of F. Solsona, F. Giné, P. Hernández and E. Luque is to build a NOW that runs parallel applications with performance equivalent to a MPP as well as executing sequential tasks as a dedicated uniprocessor with acceptable performance.
- Z. Juhasz and L. Kesmarki investigates the possible use of Jini technology for building Java-based metacomputing environments and gives an overview of Jini and highlights those features that can be used effectively for metacompting.
- The implementation of a skeleton library allowing the C programmer to write parallel programs using skeleton abstractions to structure and exploit parallelism is given by M. Danelutto and M. Stigliani.
- The contributions of the paper of C. Wagner and F. Mueller are twofold. First, a protocol for distributed mutual exclusion is introduced using a token-based decentralized approach, which allows either multiple concurrent readers or a single writer to enter their critical sections. Second, this protocol is evaluated in comparison with another protocol that uses a static structure instead of dynamic path compression.
- W. Schreiner, C. Mittermaier and F. Winkler describe a parallel solution to the problem of reliably plotting a plane algebraic curve based on Distributed Maple, a distributed programming extension written in Java.
- An Application Programming Interface (PCI-DDC) is described by E. Renault, P. David and P. Feautrier which provides different levels of integration in the kernel depending on the security and the performances expected by the administrator.
- A Clustering Approach for Improving Network Performance in Heterogeneous Systems is the topic of the presentation of V. Arnau, J. M. Orduña, S. Moreno, R. Valero and A. Ruiz. They propose on one hand a clustering algorithm that, given a network topology, provides a network partition adapted to the communication requirements of the set of applications running on the machine. On other hand, they propose a criterion to measure the quality of each one of the possible mappings of processes to processors that the provided network partition may generate.

All in all, we think we have an interesting program for this scientific event on cluster computing.

We are grateful to the authors for their contributions and the reviewers who provided many useful comments in a very short time. We are also grateful to the organizers of the Europar Conference 2000 for their helpful support regarding the organization of the Workshop.

June 2000
EuroPar'2000 Cluster Computing Workshop Team

Rajkumar Buyya,
Monash University, Australia (Global Chair).
Dan Hyde,
Bucknell University, USA (Vice Chair).
Mark Baker,
University of Portsmouth, UK (Vice Chair).
Djamshid Tavangarian,
University of Rostock,Germany (Local Chair)