
ELSEVIER

Parallel Computing 26 (2000) 175–177

PARALLEL COMPUTING

www.elsevier.com/locate/parco

Guest editorial

Parallel computing on clusters of workstations

There has been an increasing interest in the use of clusters of workstations connected together by high speed networks for solving large computation-intensive problems. The trend is mainly driven by the cost effectiveness of such systems as compared to large multiprocessor systems with tightly coupled processors and memories. Parallel computing on a cluster of workstations connected together by high speed networks has given rise to a range of hardware and network related issues on any given platform. Load balancing, inter-processor communication, and transport protocol for such machines are being widely studied. With the availability of cheap personal computers, workstations and networking devices, the recent trend is to connect a number of such workstations to solve computation-intensive tasks in parallel on such clusters.

The workstations can be connected using different network technologies such as off the shelf devices like Ethernet to specialized networks. Such networks and the associated software and protocols introduce latency and throughput limitations thereby increasing the execution time of cluster based computation. Researchers are engaged in designing algorithms and protocols to minimize the effect of this latency.

This special issue was planned to bring together the present state of the art in hardware and network requirements for running parallel applications in a network of workstations connected together by high speed networks, as well as to explore future directions of research and development. A very large number of papers were submitted indicating the public interest in the topics; after a careful review the following papers were accepted.

There is a common belief that the load should be evenly distributed among all the workstations in cluster. If a bus is used to connect the workstations in a cluster, uniform distribution of load may lead to performance degradation due to bus conflicts when all the processors complete their tasks at the same time. Lin and Xie have proposed skewing of the task completion times in order to avoid bus contention. The authors have carried out analysis and simulation to show that communication conflicts can be minimized in a bus based system by using a load skewing task assignment strategy.

An Ethernet switch can be used to interconnect a number of workstations in a cluster. Transport protocols normally use some form of acknowledgment to make

PII: S0167-8191(99)00101-5

the transport mechanism reliable. Donaldson, Hill and Skillicorn have proposed a reliable transport protocol which maximizes the switch performance. Their scheme uses a receiver driven protocol and a fine grained acknowledgment reduction technique. It has been argued that performance similar to high performance machines can be obtained at a fraction of the cost by using networks of workstations connected by an Ethernet switch.

The paper by Brightwell, Frisk, Greenberg, Hudson, Levenhagen, Maccabe and Riesen describes a 96-node prototype of a massively parallel computational architecture which has been built at Sandia National Laboratory using off the shelf computing and networking hardware. Performance results of parallel benchmark codes have been presented with indication of problems in building such a massively parallel machine.

In a network of workstations, it is important to utilize the available computing power to its fullest potential, i.e., not to let the workstations idle as long as there are ready tasks. The paper by Melab and Talbi describes a method to automatically distribute tasks to idle workstations in a network of workstations. They have provided experimental results for the efficiency of their scheme by using examples like Gauss Jordan algorithm. In a distributed computing environment, bandwidth and latency of the underlying network are the usual bottlenecks for achieving high throughput. Chu and Dowd, in their paper, propose a cache coherence protocol to reduce the effects of latency. Simulation results have shown better performance of their protocol as compared to previous protocols.

Blum, Wang and Leung have studied the architecture of cluster and message passing algorithms and their effects on the performance of workstation based cluster computing. They have also investigated new network topologies and message passing software systems which could reduce latency. Test results have been presented to demonstrate that the new architecture and software result in very fast systems.

Different types of networks can be used to connect a number of processors in a workstation cluster. The delays introduced by the communication layers in the operating system is usually a bottleneck. Chiola and Ciaccio have described a fast communication layer for 100Base-T based workstation clusters.

A number of primitives such as multi-cast and barrier synchronization have been implemented. One of the challenges in network based computing is the parallel implementation of DO-ACROSS loops. Many compiler optimization techniques have been proposed to parallelize DO-ACROSS loops while keeping the communication overhead low. Liu and King have proposed a methodology to parallelize DO-ACROSS loop at runtime in a network of workstations environment. It uses multi-threading to hide latency. The authors have demonstrated the usefulness of their technique using a number of Alpha workstations connected together using a 10 Mbps Ethernet.

We are grateful to the Regional Editor, Professor Robert Hiromoto and the Editor-in-Chief, Professor Dr. Gerhard R. Joubert, for their guidance through the whole process. We thank the large number of authors who submitted papers for this special issue and we gratefully acknowledge the efforts of the reviewers who spent their valuable time to provide reviews in a timely fashion.

11 October 1999

Mohammed Atiquzzaman
Department of Electrical and Computer Engineering,
University of Dayton,
Dayton, *OH 45469*, *USA*
E-mail address: atiq@ieee.org

Pradip K. Srimani
Department of Computer Science,
Colorado State University,
Ft. Collins, *CO 80523*, *USA*
E-mail address: srimani@cs.colostate.edu