

Introduction to Partitioning and Scheduling Minitrack

Ishfaq Ahmad

Horst D. Simon

Tao Yang

Dept. of Computer Science
Hong Kong Univ. of Sci.&Tech.
Kowloon, Hong Kong

Silicon Graphics
Mail Stop 7L-580
Mt. View, CA 94043

Dept. of Computer Science
University of California
Santa Barbara, CA 93106

High performance computing in massively parallel processors and a network of workstations has been a major research focus in recent years. The goal of partitioning and scheduling is to efficiently decompose and execute application programs on such an environment. We have received a number of high quality papers submitted for this minitrack and we are only able to accept one-third of the submissions as regular papers and select three extended abstracts from the submissions. The papers contributed to this minitrack address many important issues in partitioning and scheduling. A summary of these papers is listed below.

Lee and Chen propose a dynamic programming algorithm to automatically determine a distribution of data at time compile-time.

Pande and Bali propose a multi-phase partitioning and scheduling scheme for exploiting DOALL loop parallelism on distributed memory machines.

Sivaraman and Raghavendra present a method for automatically generating a dynamic data distribution and alignments on a distributed memory machine.

Hendrickson, Leland and Van Driessche present new mesh partitioning algorithms, ensuring that communication only occurs between architecturally-near processors.

Minyard, Kallinderis, and Schulz present a partitioning method for complex 3-D hybrid (prismatic/tetrahedral) computational grids, using orthogonal recursive bisection of a special octree.

Taylor, Holmer, Schwabe and Hribar propose a domain decomposition scheme to minimize total execution time by considering the trade-off between load balancing and communication minimization.

Barnard describes the design of a parallel implementation of multilevel recursive spectral bisection.

Parashar and Browne propose a run-time partitioning and load balancing scheme for distributed adaptive grid hierarchies.

Erhel, Hahad, and Priol present techniques for par-

tioning and scheduling of irregular loops arising in finite element computations on unstructured meshes.

Raje, Pease and Guy describe a prototype, OFFERS, for the static analysis of programs written in a sequential object-oriented language and suggests class partitioning.

Chen and King study a new model for dynamic processor allocation in large scale multicomputer systems.

Milojicic and Paindaveine discuss a process migration method used in OSF/1 AD operating system.

Dutra presents techniques for distributing and-work and or-work in a parallel logic programming system.

Aubry, Guernic and Machard present a data-flow oriented language called SIGNAL that allows the user to design safe real-time applications.

Vemulapalli and Ali present a two-step mapping approach to scheduling tasks based on the interval ordered task graph theory.

Cherkasova, Kotov, and Rokicki discuss the impact of message scheduling on a packet switching interconnect fabric.

These papers reflect a range of important and innovative work in the area of partitioning and scheduling for parallel and distributed computing. The technical program of this minitrack was made possible through the efforts of many people. We would like to thank all authors who contributed and/or submitted papers and our referees for their insightful comments. We thank Hesham El-Rewini for his help in organizing this minitrack.