J Supercomput (2010) 51: 241–243
DOI 10.1007/s11227-010-0399-x

Preface

Zhiyi Huang · John Hine · Laurent Lefevre · Tony McGregor · Yi Pan · Hong Shen · Zhiwei Xu

Published online: 5 March 2010

This special issue titled "Emerging Research in Parallel and Distributed Computing" reflects the recent research efforts of authors of the Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT 2008). This special issue is divided into two sections: the first addresses the issues in multicore computing, while the second contains the new developments in grid computing.

Parallel computing is facing new challenges with the advent of multicore and many-core architectures. In the past decades, CPU speeds doubled every 18 months by increasing the clock speed. Software has had a "free ride" achieving better performance through increasing hardware speed. However, this free ride is now over due to physical limitations such as heat dissipation and energy consumption in the manufacture of CPU chips. Because of the stalling CPU clock speeds, parallel computing becomes the only way for software to increase performance. Researchers are now working hard to provide suitable tools and environments for parallel computing, from languages to runtime systems, from hardware accelerators to scheduling algorithms for operating systems.

In the invited paper "The Cilk++ Concurrency Platform", Charles Leiserson gives an overview of a state-of-the-art parallel programming environment—Cilk++, which is derived from MIT Cilk. Cilk++ consists of a compiler, a runtime system, and a race-detection tool. It is simple to use and guarantees to load-balance computations within theoretical bounds. One recent prominent feature of Cilk++ is the "hyperobject" library to mitigate the data race problem.

In "Data Race: Tame the Beast", Kai-Cheung Leung, Zhiyi Huang, Qihang Huang, and Paul Werstein propose a data-race prevention scheme in View-Oriented Parallel Programming (VOPP). VOPP is a novel parallel programming model, which uses

Z. Huang (✉)
University of Otago, Dunedin, Otago, New Zealand
e-mail: zhuang@cs.otago.ac.nz

views to bundle mutual exclusion with data access. The proposed data-race prevention scheme is implemented with existing virtual memory protection mechanisms, while traditional solutions are focused on data race detection. Experimental results show that the extra overhead of the scheme is relatively small. This prevention scheme ushers in a new approach to the chronic problem of data races.

In "An Enhancer of Memory and Network for Applications with Large-Capacity Data and Non-Continuous Data Accessing", Noboru Tanabe, Hirotaka Hakozaki, Hiroshi Ando, Yasunori Dohi, Zhengzhe Luo and Hironori Nakajo propose a hardware accelerator to improve memory access of non-contiguous data. With multicore architectures, the performance of memory modules is vital for applications to reduce memory access delay so as to utilize the full power of many cores. However, current memory systems and cache architectures are designed for applications with spatial locality. In order to overcome the problems of current memory systems, this paper proposes a memory enhancer equipped with scatter and gather vector access functions, which can offer high performance for non-contiguous data access and network connectivity.

In "Scheduling Task Graphs Optimally with A^*", Ahmed Zaki Semar Shahul and Oliver Sinnen investigate the task scheduling problem using the A^* search algorithm which is a best-first state space search. Task scheduling is essential to parallel runtime systems. The adaptation of the A^* search algorithm for task scheduling can produce optimal schedules in reasonable time for small to medium sized task graphs. In comparison to previous approaches, the A^* scheduling algorithm presented in this paper has a significantly smaller search space due to a much improved consistent and admissible cost function $f(s)$ and additional pruning techniques.

Grid computing continues to advance including work on cloud computing and virtual organizations. In "A Layered Virtual Organization Architecture for Grid", Yongqiang Zou, Li Zha, Xiaoning Wang, Haojie Zhou, and Peixu Li present a layered architecture to construct virtual organizations. Virtual organizations, defined as dynamic collections of individuals, institutions and resources, are widely accepted in grid computing and other distributed computing environments. This paper identifies four non-functional goals in designing a virtual organization. The proposed three-layer architecture satisfies the functional requirements while achieving the non-functional goals: decentralization, flexibility, simplicity, and efficiency.

With the backdrop of global warming, energy efficiency has become an important issue for grid and cloud computing. Large-scale cloud architectures face increasing energy consumption. In "Designing and Evaluating an Energy-Efficient Cloud", Laurent Lefevre and Anne-Cecile Orgerie propose an energy-efficient framework for cloud architectures. Through experiments on modern multicore computers, they show that the proposed framework can save up to 25% of energy consumption.

Replication placement is an essential but challenging issue in large-scale data grids. In "Adaptive Popularity-Driven Replica Placement in Hierarchical Data Grids", Mohammad Shorfuzzaman, Peter Graham and Rasit Eskicioglu propose a new dynamic replica-placement algorithm for hierarchical data grids. The algorithm is guided by file "popularity", which places replicas close to clients to reduce data access time while still using network and storage resources efficiently.

Agent technology is useful for coordinated problem solving in virtual organizations. It can facilitate semi-automated negotiation between distributed elements and

enables efficient management of distributed systems. In “A Framework for Facilitating Cooperation in Multi-Agent Systems”, Toktam Ebadi, Maryam Purvis and Martin Purvis introduce a multi-agent framework that facilitates cooperation in multi-agent robotic systems. It uses a layered approach based on Coloured Petri Nets for modeling complex, concurrent conversations among agents.

Papers in this special issue were selected from 17 submissions and underwent several rounds of review. We are very grateful to our external reviewers Ian Foster, Stefan Turek, and Naixue Xiong.

Enjoy reading.