

Advances in high performance computing

Minyi Guo · Jingling Xue

Published online: 25 May 2007
© Springer Science+Business Media, LLC 2007

This special issue brings together five research papers providing a snapshot of the current state of the art in the broad area of high performance computing. These papers have been selected as extended and revised versions, subject to peer review, from those originally accepted for presentation at the *3rd International Symposium on Parallel and Distributed Processing and Applications (ISPA'2005)*, which was held during 2–5 November, 2005 in Nanjing, China. The selected papers address a range of issues in high performance computing and their contributions are summarised below.

In “An Energy-Efficient Protocol for Data Gathering and Aggregation in Wireless Sensor Networks,” **Liu, Cao, Zheng and Xie** introduce a hierarchical clustering protocol for data gathering and aggregation in wireless sensor networks. Their simulation results show that their protocol outperforms two existing ones, LEACH and PEGASIS, in terms of network lifetime and the amount of data gathered.

More and more of the top 500 super-computers are cluster-based. However, the increased reliance on super-computing clusters based on commercial off-the-shelf parts has created a growing need for lightweight fault tolerance. In their article, “ER-TCP: An Efficient TCP Fault-tolerance Scheme for Cluster Computing,” **Shao, Jin, Cheng and Jiang** present a scheme that combines active replication with a logging mechanism to achieve fault tolerance for the TCP connections on the server-side.

In “Improving the Parallelism of Iterative Methods by Aggressive Loop Fusion,” **Xue, Guo and Dai** present a compiler approach to aggressively fusing loop nests

M. Guo (✉)
School of Computer Science and Engineering, The University of Aizu,
Fukushima 965-8580, Japan
e-mail: minyi@u-aitu.ac.jp

J. Xue
Programming Languages and Compilers Group, School of Computer Science and Engineering,
University of New South Wales, Sydney, NSW 2052, Australia

to improve performance on multiprocessor platforms by improving cache locality. Unlike existing loop fusion algorithms, their algorithm is capable of fusing loop nests even in the presence of fusion-preventing data dependences.

In “Runtime Support for Scalable Programming in Java,” **Lim, Carpenter, Fox** and **Lee** discuss the design and implementation of a communication library useful for supporting high performance computing in Java. They evaluated their library using a multi-grid application and a GUI application.

Redundancy schemes are widely used in distributed systems to achieve high data availability. In “Insight into Redundancy schemes in DHTs,” **Chen, Qiu** and **Wu** present a scheme for achieving good performance by sharing user downloaded files for subsequent accesses while utilising erasure coding to maintain file availability. The authors compared their scheme with two existing ones and their simulation results show that their scheme can save more maintenance bandwidth with an acceptable redundancy factor.

Finally, we would like to thank the authors for their contributions, including the authors of papers that could not be included in this special issue. We would also like to thank the reviewers for their dedication and effort in providing timely and thorough reviews. This special issue would not have been possible without their hard work.