Editorial

Pingzhi Fan
School of Computer and Communications Engineering,
Southwest Jiaotong University,
Chengdu, 610031, China
Fax: +86 28 87600743 E-mail: p.fan@ieee.org

Biographical notes: Pingzhi Fan (Senior Member, IEEE) received his MS in computer engineering from the Southwest Jiaotong University, PRC, and PhD in electronic engineering from the Hull University, UK. He is currently a Professor and Dean of the School of Computer and Communication Engineering, Southwest Jiaotong University, China, and as Guest Professor of the Leeds University, UK and Shanghai Jiaotong University. He was a recipient of the UK ORS Award (1992) and the National Science Foundation of China for Outstanding Young Scientist (1998). He is inventor of 15 patents and author of about 200 research papers and five books.

Aiming at bringing together state-of-the-art contributions ranging from papers that report novel research to papers that describe industrial application experiences, in this special issue, eight papers have been selected to address various aspects of architectures, software systems, algorithms, applications, and technologies in the field of parallel and distributed computing.

All the papers published in this special issue are closely related to the Fourth International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT'03) held in Chengdu, China, August 27–29, 2003. PDCAT is a major forum for scientists, engineers, and practitioners throughout the world to present the latest research, results, ideas, developments, techniques, and applications in all areas of parallel and distributed computing. Following PDCAT'00 in Hong Kong, PDCAT'01 in Taipei, PDCAT'02 in Kanazawa, Japan, PDCAT'03 took place in Chengdu, China. PDCAT'03 received 343 submissions that came from 21 countries worldwide. Through a careful review process, where each submitted manuscript was reviewed by at least two programme committee members or their external referees, the programme committee accepted 211 papers (maximum five pages each), which are included in the proceedings published by IEEE Press. Because of the limited space of this special issue, based on the accepted PDCAT'03 papers with excellent quality as evaluated by reviewers and the excellent presentations recommended by PDCAT'03 session chairs, only prospective authors were invited to extend and enhance their papers for submission. Except for an invited survey paper based on the keynote speech delivered at PDCAT'03 by Professor Timothy Mark Pinkston, at least two review reports are obtained for each paper through a strict double blind process.

The first paper 'Trends toward on-chip networked microsystems' by Timothy Mark Pinkston and Jeonghee Shin addresses some trends in the application, implementation technology, and processor architecture areas. A taxonomy, which captures the influence of these trends on processor microsystems is presented. This paper gives an insightful and systematic description on the trend towards on-chip networked microsystems derived from logically and physically partitioning the processor architecture. Partitioning the architecture physically enables the scaling properties of the underlying implementation technology to continue to provide increasing performance and not be encumbered by chip-crossing wire delay. In addition, the impact on future research directions of this paradigm shift in the way microsystems are designed and intraconnected is also highlighted.

Advances in electro-optic switches have made optical communication a good networking choice, which can satisfy the large channel bandwidth, low communication latency, low error rate, and gigabit transmission requirements of high performance computing and communication applications. The paper, 'Optimal all-to-all personalised exchange in a novel optical multistage interconnection network' by Siu-Cheung Chau and Ada Wai-Chee Fu, investigates the optimal all-to-all personalised exchange, where each node in the network needs to send a different message to each of the other nodes. The authors propose an optimal optical multistage interconnection network (MIN) that requires only one pass to realise a permutation for an all-to-all personalised exchange.

The paper by Haibo Wen et al. deals with the contention resolution (CR) in optical burst switching (OBS) networks. Unlike most previous researches, which focus mainly within the local core switching node, this paper investigates an integrated CR technique in OBS networks based on generalised multi-protocol label switching (GMPLS). Based on a new network model with label switched path (LSP), a novel load-balancing contention resolution (LBCR)

technique, which can significantly reduce burst contention and burst loss probability is proposed and analysed.

Motivated by the fact that conventional threaded programming model exhibits poor performance under high concurrency workloads, and traditional thread-per-request solution is no longer feasible, Yan Chen et al. present a stage-based parallel programming model for high concurrency stateful network services. By making use of the advantages of thread-per-stage, three key design principles for delivering a high performance stage-based design are revealed and a sample SIP proxy server is implemented.

The problem of faulty-node tolerant cycle embedding is to find a cycle in a network with some faulty nodes, and the cycle length depends on the number of faulty nodes. The paper by Yamin Li, Shietung Peng, and Wanming Chu describes a low-degree dual-cube (DC) as an alternative to the hypercubes. A dual-cube uses hypercubes as basic components and each N-node hypercube contains every N-node ring as a subgraph. It is shown in this paper that a fault-free cycle containing at least $2^{2m+1} - 2f$ nodes can be constructed in a DC(m) with $f \leq m$ faulty nodes, which can be used as an interconnection network for large scale parallel computers.

Because efficient utilisation of processing resources in a large, multi-user parallel system depends on the reliable processor management scheme, the paper 'A dynamic processor management strategy on the reconfigurable meshes' by Kyung-Hee Seo and Sung-Chun Kim is concerned with the reconfigurable mesh consisting of a mesh-connected topology augmented by a dynamically reconfigurable bus system. By reconfiguring a faulty mesh system into a maximum convex system, this paper presents a dynamic and reliable processor allocation strategy to increase the performance of mesh-connected parallel systems with faulty processors. It is shown that the proposed strategy performs more efficiently than other strategies in terms of the completion time, the job response time, and the system utilisation.

The paper 'A new strategy for multiprocessor scheduling of cyclic task graphs' by Frode E. Sandnes and Oliver Sinnen covers a modular strategy for scheduling iterative computations. In this paper, iterative computations are represented using cyclic task-graphs that are transformed into acyclic task-graph. The authors experimentally quantify how the task transformation affects the make-span of the schedules and the effectiveness of the approach is compared to other methods including a graph unfolding strategy. In addition, a more intuitive graph unfolding formulation is also presented.

The last paper by Said Boussakta and Mohamed Aziz proposes a fast and efficient 3-D digital filtering method. It is based on the combination of a 3-D hybrid parallel filtering algorithm and the 3-D vector radix fast Hartley transform (3-D VR FHT) for applications in high resolution/high speed multidimensional signal and image processing. The proposed 3-D hybrid parallel algorithm overcomes the computations overhead and performance limitations associated with the block filtering method by eliminating the overlapping blocks and boundary conditions in parallel filtering applications.

It is our sincere hope that this selection of papers, although necessarily limited by space considerations, will serve to give the readership of this journal a sense of recent advances in the exciting field of parallel and distributed computing and applications.

The guest editors are grateful to all the authors for their hard work in preparing and revising their papers, the reviewers who contributed significantly to this special issue, and IJHPCN Editors-in-Chief Laurence Tianruo Yang and Yi Pan for their consistent support and encouragement. This work was supported partly by the National Science Foundation of China (NSFC) and the Research Grants Council of Hong Kong (RGC) joint research scheme (No.60218001/ N_HKUST617-02 and No.60272070), and the University Doctorial Research Foundation under the grant No.20020613020, Ministry of Education, China.