Preface

Tarek S. Abdelrahman
Edward S. Rogers Sr. Department of Electrical and Computer Engineering,
University of Toronto, Toronto, Ontario, Canada M5S 3G4
E-mail: tsa@eecg.toronto.edu

Biographical notes: Tarek Abdelrahman received his PhD degree in Computer Science and Engineering from the University of Michigan at Ann Arbor in 1989. He is currently Professor of Electrical and Computer Engineering and of Computer Science at the University of Toronto, Toronto, Ontario, Canada. His current research interests are in compiler optimisations for parallelism and virtual machine design. Dr. Abdelrahman is a member of ACM and USENIX.

The 31st International Conference on Parallel Processing (ICPP-2002) was held in Vancouver, BC, Canada on August 18–21, 2002. The conference received over 180 submissions of which 67 papers were selected for inclusion in the conference's proceedings. These papers presented state-of-the-art work in both traditional and emerging areas of parallel processing, including architecture, compilers, scheduling, parallel programming, grid computing, mobile computing and middleware systems.

This special issue presents a sample of the outstanding papers presented at ICPP-2002. These papers reflect state-of-the-art contributions to the field of parallel processing, particularly to improving performance in parallel system. The papers come from different areas of parallel processing, reflecting the maturity and diversity of research in this field.

Monreal et al. describe a technique for improving instruction-level parallelism in processors. Ferschweiler et al. describe a repository of performance data (in the form of tracefiles) to aid researchers in testing parallel tools, and in analysing and predicting performance. Similarly, Prodan and Fahringer describe and evaluate a directive-based language that allows researchers to specify experiment parameters on parallel systems. Kettimuthu et al. improve performance through pre-emptive job-scheduling strategies. Finally, Wu and Majumdar exploit the limited heterogeneity of parallel/distributed systems to improve the performance of CORBA-based middleware applications.

The following is a brief overview of each contribution:

- Monreal et al. describe a novel register-renaming scheme that improves instruction-level parallelism in processors. Instead of conservatively releasing registers after instructions commit, registers are released as soon as the processor knows that they are no longer used. This results in speedups up to 9% for some applications.
- Ferschweiler et al. describe a unique repository of tracefiles that serve as input to performance analysis and performance prediction tools used in parallel computing. The repository, called the *Tracefile Testbed*, facilitates searching and retrieval of tracefiles based on a variety of characteristics. It has a web-based interface to allow the browsing, downloading and uploading of collections of tarcefiles.
- Prodan and T. Fahringer present a directive-based language for the specification of synthetic experiments for studying the performance of parallel and distributed scientific applications. The language, called ZEN, introduces substitute and assign directives to specify arbitrary problem, parallelisation, or machine parameters, including program variables, file names, compiler options, target machines, machine sizes, scheduling strategies and/or data distributions. They demonstrate the utility of the language in the performance analysis of an ocean simulation application and for the parameter study of a computational finance code.

- Kettimuthu et al. investigate pre-emptive job-scheduling strategies for scheduling parallel jobs at supercomputer centres. They introduce a new and superior tunable selective-suspension strategy for scheduling, and demonstrate its effectiveness. They further analyse the performance of the proposed strategy under different load conditions.
- Wu and Majumdar present and evaluate a technique called *flyover* that improves the performance of CORBA-based middleware applications executing on parallel/distributed systems with limited heterogeneity. The technique is implemented in a prototype and is shown to lead to significant performance gains compared to pure CORBA-based systems, using commercial middleware products, under various workload and system parameters.

It is my hope that you will find the work contained in this issue useful and interesting.