CONCURRENCY AND COMPUTATION: PRACTICE AND EXPERIENCE
Concurrency Computat.: Pract. Exper. 2009; **21**:1757–1759
Published online 18 May 2009 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/cpe.1451

Special Issue: Euro-Par 2008

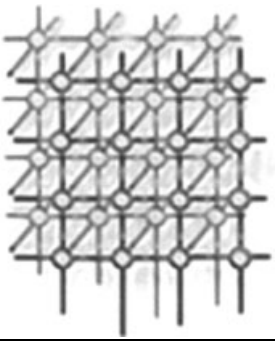

SUMMARY

Preface for the special issue on Euro-Par 2008. Copyright © 2009 John Wiley & Sons, Ltd.

KEY WORDS: distributed computing; parallel computing

This special issue of *Concurrency and Computation*: *Practice and Experience* contains revised and extended versions of selected papers presented at the Euro-Par 2008 conference.

Euro-Par—the European Conference on Parallel Processing—is an annual series of international conferences dedicated to the promotion and advancement of all aspects of parallel and distributed computing. The major themes can be divided into the broad categories of hardware, software, algorithms and applications for parallel and distributed computing.

Euro-Par 2008, the 14th conference in the Euro-Par series, was organized in Las Palmas, Gran Canaria, Spain, by the Computer Architecture & Operating Systems Group (CAOS) of the Universitat Autònoma de Barcelona and the Department of Computer Science and Engineering (DIS) of the Universidad de Las Palmas de Gran Canaria.

Fourteen broad topics were defined and advertised, covering a large variety of aspects of parallel and distributed computing. The call for papers attracted a total of 264 submissions. The submitted papers were reviewed at least three and, in many cases, four times. A total of 89 papers were finally accepted for publication. This makes a global acceptance rate of 34%. The authors of accepted papers come from 25 countries, with the four main contributing countries—Spain, France, Germany and the U.S.A.—accounting for more than 60%.

Based on the results of the reviews and a majority opinion of the topic programme committees, several papers were recommended for a special journal issue. The authors were contacted at the conference and invited to submit revised and extended versions of their papers. These new versions were reviewed independently by three reviewers; two had also reviewed the conference version, the third had not. Eventually, four papers were accepted for publication, all from different Euro-Par topics.

Like last year's best Euro-Par papers in volume 21, number 7 of this journal, this year's selection attests to the variety of Euro-Par research activities. This year, the emphasis of the best papers is on issues close to the hardware—of driving the hardware efficiently. It has been known for

a long time that this is a very difficult problem. As the hardware gets ever more complex, and as the software stack on top of it gets ever higher, rules of thumb have no hope of working any longer. Elaborate conceptual models are needed to reduce the complexity to a level that can be tackled conveniently by humans. This is precisely the business of Euro-Par: bringing together researchers with experience in these major challenges and with the ability to work out suitable abstract models.

Topic 1 on *Support Tools and Environments* is represented by the paper *Clock Synchronization in Cell B.E. Traces*, authored by M. Biberstein, Y. Harel and A. Heilper from Israel. This work is an example that theoretical methods of program analysis can compensate for weaknesses of an architecture. The widely used IBM Cell hardware lacks a global clock. The paper offers a method of assigning, off-line, global time stamps to trace records of a parallel execution on the Cell. Remarkably, this method is based on elaborate theoretical notions, rooted in Lamport's theory of logical locks, and graph theory. Advanced experiments confirm the accuracy and practical impact of the results.

Topic 2 on *Performance Prediction and Evaluation* is represented by the paper *Dynamic File System Semantics to Enable Metadata Optimizations in PVFS*, authored by M. Kuhn, J. Kunkel and T. Ludwig from Germany. This paper describes a modification of the distributed file system PVFS that improves performance by removing a layer of metadata, thus, reducing storage distribution for a file. As a consequence, the data of small files will be stored on the same node instead of being scattered across the network. The role of metadata in enhancing performance and the idea of metadata minimization are not new, but the contribution of this paper is to apply it to the well-known PVFS parallel cluster file system, exploring the right trade-off between performance and functionality for a range of practical applications.

Topic 8 on *Distributed Systems and Algorithms* is represented by the paper *A Meta-predictor Framework for Prefetching in Object-based DSMs*, authored by J. C. Beyler, M. Klemm, P. Clauss and M. Philippsen from France and Germany. The paper describes a dynamic optimizer for object-based distributed shared memory (DSM) systems. The optimizer profiles program execution and, based on the information gathered, inserts code dynamically to prefetch objects from remote nodes at specific program locations. Though profiling and dynamic binary rewriting are all part of the normal execution of the application, the induced overhead is small enough to make the approach efficient, as counter-intuitive as it may seem. Thanks to the use of sophisticated behavior predictors, performance gains of up to 80% can be observed on the benchmarks.

Topic 9 on *Parallel and Distributed Programming* is represented by the paper *Optimizing Multiple Conjugate Gradient Solvers for Large-Scale Systems*, authored by J. C. Sancho and D. J. Kerbyson from the U.S.A. Conjugate gradient solvers are used in many scientific codes, and this domain has received much attention in the past. However, this paper brings in a novel idea for applications running multiple independent instances of solvers. These instances can be interleaved, allowing for the overlapping of communication with other communication and computation, and thereby increasing performance. The performance improvements may be particularly significant in very large-scale systems. Interestingly enough, this idea is closely related to the theoretical work developed by Valiant on the BSP model [1]: overlapping two independent communication statements at a synchronization barrier should result in a reduced start-up time and a more even distribution of messages across a network. Now, it is being applied successfully in practice by the scientific research community.

DOI: 10.1002/cpe

Concluding this preface, we would like to thank Prof. Geoffrey Fox and Prof. Luc Moreau, editors of this journal, for their support of this special issue. Last but not least, we would like to thank our peers who assisted us in reviewing the papers and helped to strengthen the final versions.

REFERENCE

1. Valiant LG. A bridging model for parallel computation. *Communications of the ACM* 1990; **33**(8):103–111.

LUC BOUGÉ
ENS Cachan/Bretagne,
IRISA, *INRIA*,
Campus Ker Lann,
F-35170 Bruz,
France

CHRISTIAN LENGAUER
University of Passau,
D-94030 Passau,
Germany

DOI: 10.1002/cpe