Theory Comput Syst (2010) 47: 809–810
DOI 10.1007/s00224-010-9284-5

Foreword

Cyril Gavoille · Boaz Patt-Shamir · Christian Scheideler

Published online: 5 August 2010

This special issue contains selected papers whose preliminary versions were presented at the 19th Annual ACM Symposium on Parallelism on Algorithms and Architectures (SPAA), held in San Diego, California, from June 9 to 11, 2007. The conference defines parallelism very broadly and therefore includes papers on various problems in wired and wireless networks as well as parallel and multicore systems, cache-oblivious algorithms, network games and concurrent and parallel programming approaches. This breadth is reflected in the six papers that we, the guest editors, chose to invite for this special issue:

- "Distributed Approximation of Capacitated Dominating Sets" derives a number of results on the distributed complexity and local approximability of the Capacitated Dominating Set Problem.
- "Strong-Diameter Decompositions of Minor Free Graphs" provides the first sparse covers and probabilistic partitions for graphs excluding a fixed minor that have strong diameter bounds in the sense that each set of the cover resp. partition has a small diameter as an induced subgraph.
- "Approximation Algorithms for Multiprocessor Scheduling under Uncertainty" presents polynomial-time approximation algorithms for the multiprocessor scheduling problem in scenarios where there is uncertainty in the successful execution of jobs when assigned to processors.

C. Gavoille
LaBRI—Université de Bordeaux, 351, cours de la Libération, 33405 Talence cedex, France

B. Patt-Shamir
Department of Electrical Engineering, Tel Aviv University, Tel Aviv 69978, Israel

C. Scheideler (✉)
Institut für Informatik, Universität Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
e-mail: scheideler@in.tum.de

- "The Cache-Oblivious Gaussian Elimination Paradigm: Theoretical Framework, Parallelization and Experimental Evaluation" presents parallel variants of the standard Gaussian elimination algorithm that achieve a good speed-up and match the sequential caching performance cache-obliviously for both shared and distributed caches for sufficiently large inputs.
- "Local MST Computation with Short Advice" explores the distributed computing with advice paradigm for the local computation of a minimum spanning tree and shows that a logarithmic-time algorithm is possible with just a constant number of bits of advice per node.
- The paper on "Optimal Sparse Matrix Dense Vector Multiplication in the I/O-Model" investigates the worst-case complexity in terms of the number of I/O-operations in order to multiply a matrix with a vector that are stored in some external memory.

We thank the authors for contributing their work, the referees for ensuring the high quality of the material presented here, and the Editor-in-Chief for his valuable assistance with the process of assembling the special issue.

Guest Editors