Guest Editorial

Mark de Berg

Received: 1 February 2012 / Accepted: 4 February 2012 / Published online: 10 February 2012
© The Author(s) 2012. This article is published with open access at Springerlink.com

This special issue contains a selection of papers from the *18th European Symposium on Algorithms (ESA)*, which was held September 6–8, 2010 in Liverpool (UK). The European Symposium on Algorithms is one of the premier algorithms conferences worldwide, featuring high-quality papers on efficient algorithms and data structures in computer science, discrete applied mathematics, and operations research. ESA publishes both theory-oriented papers (in the Design and Analysis Track) and experimentation-oriented papers (in the Engineering and Application Track). For this special issue seven excellent papers were selected from the Design and Analysis Track. All papers went through the thorough reviewing process that is standard for papers published in *Algorithmica*.

Bansal, Gupta, Li, Mestre, Nagarajan and Rudra study matchings in random graphs in their paper *When LP is the Cure for Your Matching Woes: Improved Bounds for Stochastic Matchings*. They consider a model where each edge exists with a certain probability and the algorithm can query edges for their existence; when the edge exists, the algorithm must use it in the matching. Using clever LP rounding schemes, they obtain several results that improve and/or generalize many of the existing solutions for stochastic matchings.

The sporadic task model is a model of recurrent processes in real-time systems that has received great attention in the last years. The paper *Feasibility Analysis of Sporadic Real-Time Multiprocessor Task Systems* by Bonifaci and Marchetti-Spaccamela solves a fundamental open problem in the sporadic task model, namely to decide if a system of sporadic tasks is feasible. They also solve two other important problems in the area, concerning online feasibility and concerning schedulability.

M. de Berg (✉)
TU Eindhoven, Eindhoven, The Netherlands
e-mail: mdberg@win.tue.nl

Caching is a classical optimization problem. The status of several offline variants of the problem has been open for years, however. In their paper *Caching is Hard—Even in the Fault Model*, Chrobak, Woeginger, Makino and Xu settle the complexity of these variants by proving NP-hardness using an intricate reduction from Vertex Cover.

In their paper *Budgeted Red-Blue Median and its Generalizations*, Hajiaghayi, Khandekar and Kortsarz consider an interesting variant of the well-studied k -median problem. They show that a surprisingly simple local search algorithm gives a constant factor approximation. In addition, they give improved results for the prize-collecting k -median problem.

Brodal, Davooda and Rao study a fundamental data-structuring problem, namely to preprocess a two-dimensional array such that the smallest element in a rectangular query range can be reported efficiently. In their paper *On Space Efficient Two Dimensional Range Minimum Data Structures* they study data structures that require only few extra bits of storage in addition to the storage needed for the array itself, obtaining improved lower and upper bounds on the trade-off between extra storage and query time.

Pountourakis and Vidali study an algorithmic problem from the field of mechanism design: find a mechanism that, given a set of customers who bid on a certain service, decides which customers will be serviced at what price. The goal is to design the mechanism such that it does not help customers to lie about the price they are willing to pay, even if they team up. Their paper *A complete characterization of group-strategyproof mechanisms of cost-sharing* characterizes all such mechanisms, and opens new perspectives on the study of the famous Cost-Sharing Problem.

In their paper *f -Sensitivity Distance Oracles and Routing Schemes*, Chechik, Langberg, Roditty and Peleg study approximate distance queries in weighted graphs where, besides the source vertex s and target vertex t , the query also specifies a set of f edges that the path from s to t is not allowed to use. They are the first to obtain results for $f > 2$, and they are the first to obtain a compact routing scheme in this setting.

Together these seven beautiful papers show the breadth and the depth of algorithms research in general, and of the European Symposium on Algorithms in particular. I want to thank all authors for submitting their papers to this special issue and all reviewers for their careful reviews, and I want to thank the editor-in-chief Ming-Yang Kao for giving me the opportunity to edit this special issue. Finally, I hope that the readers will find papers presented here interesting and enjoyable.

Eindhoven, January 2012

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.