Theoretical Computer Science 608 (2015) 199–200

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Preface

Guest editorial: Structural Information and Communication Complexity

This issue is dedicated to a selection of papers from the 21st International Colloquium on Structural Information and Communication Complexity (SIROCCO 2014), held in July 2014, in Takayama, Japan. The proceedings of this colloquium were published by Springer-Verlag as volume 8576 of the Lecture Notes in Computer Science (LNCS).

SIROCCO is an annual meeting devoted to the study of the interplay and trade-offs between the efficiency of algorithms and systems and the availability of information. This basic issue is common to many fields of multi-party systems, and the SIROCCO meetings attract leading researchers from the areas of distributed computing, agent-based computing, network optimization, and multi-party communication complexity.

SIROCCO also emphasizes innovative approaches that expand the collective investigative horizon by formulating new problems and by identifying new research areas, which may even be unorthodox or outside the mainstream. This openness to new research directions and approaches also motivates the Prize for Innovation in Distributed Computing, which is awarded annually at the SIROCCO meetings.

The presentations at the 21st SIROCCO included 5 invited talks, 1 SIROCCO Award talk, and 24 regular presentations selected by the Program Committee from 50 conference submissions. The 5 papers selected for this special issue were accepted after a careful refereeing process according to the normal TCS standards. They cover a range of topics, reflecting both the general scope of SIROCCO and the main themes exhibited in the conference presentations. The following is a summary of the contributions of these papers.

Mobile robots or agents have traditionally been a popular topic at SIROCCO and they feature in four of the papers in this issue, each focusing on different aspects. In the paper “The Beachcombers’ Problem: Walking and searching with mobile robots” by Jurek Czyzowicz, Leszek Gąsieniec, Konstantinos Georgiou, Evangelos Kranakis and Fraser MacQuarrie, robots cooperate to scan quickly a 1-dimensional geometric domain, such as a beach. The twist is that they can choose between two speeds of movement: moving slowly while scanning, or moving quickly without scanning. They give an optimal offline algorithm and a constant-competitive online algorithm when the length of the segment is not known.

In the paper “Rendezvous of heterogeneous mobile agents in edge-weighted networks” by Dariusz Dereniowski, Ralf Klasing, Adrian Kosowski and Lukasz Kuszner, two mobile agents that want to meet are located on a known graph at known locations, but they have different edge-traversal times and are unaware of the traversal time of the other agent. They compare the time required by deterministic algorithm for this problem to the corresponding time for the full-information problem, and show that it differs only by a constant factor under various assumptions or available communication.

In the paper “Improved periodic data retrieval in asynchronous rings with a faulty host” by Evangelos Bampas, Nikos Leonardos, Euripides Markou, Aris Pagourtzis and Matoula Petrolia, the mobile agents are to explore a graph that contains malicious nodes, or *black holes*, that can devour visiting agents. The initially co-located agents are to visit all non-malicious nodes infinitely often. The authors give the first nontrivial lower bound for problems of this kind on the number of agents needed accomplish the task. They also obtain a tight upper bound on the number of agents needed when the malicious node can choose whether to kill the agent or let it pass through.

The paper “How many ants does it take to find the food?” by Yuval Emek, Tobias Langner, David Stolz, Jara Uitto and Roger Wattenhofer, the mobile agents collaborate to find a treasure that is adversarially located on the infinite grid. The basic question of how many agents are needed turns out to depend crucially on the computational power of the agents as well as on the synchronization aspects.

In the final paper, “Lessons from the Congested Clique applied to MapReduce”, James Hegeman and Sriram Pemmaraju relate the power of two prominent computational models. They give a simulation method that allows us to translate algo-

http://dx.doi.org/10.1016/j.tcs.2015.11.019

rithms for the Congested Clique model (the fully connected variant of the CONGEST model) into the popular MapReduce model. They then use it to obtain efficient MapReduce algorithms for several problems, including coloring.

We would like to thank the authors for contributing to this special issue and the referees for their attentive and helpful work in the reviewing process.

Guest Editor
Magnús M. Halldórsson