

Guest Editorial: Selected Papers of European Symposium of Algorithms

Leah Epstein · Paolo Ferragina

Received: 5 June 2014 / Accepted: 21 June 2014 / Published online: 15 July 2014
© Springer Science+Business Media New York 2014

This Special Issue contains a selection of expanded, peer-reviewed papers originally presented at the 20th Annual European Symposium on Algorithms (ESA 2012) held in Ljubljana, Slovenia, on September 10–12, 2012. The ESA symposia are devoted to fostering and disseminating the results of high quality research on the design and evaluation of algorithms and data structures. The forum seeks original algorithmic contributions for problems with relevant theoretical and/or practical applications and aims at bringing together researchers in the computer science and operations research communities. The conference received 285 submissions from 40 countries and the PC selected 69 papers among which we decided to pick seven excellent papers composing this Special Issue. These journal papers, which have been rigorously reviewed according to the journal standards, cover a wide range of algorithmic topics, thereby appealing to both the experts in the field as well as to those who wish a snapshot of the current breadth of research in this area.

The paper on “A Model for Minimizing Active Processor Time” considers an elementary scheduling problem over a collection of n jobs, each having an integer length and a set of time intervals in which it can be feasibly scheduled. Given a parameter B , the processor can schedule up to B jobs at a timeslot t so long as it is “active” at t . The goal is to schedule all the jobs in the fewest number of active timeslots. The machine consumes a fixed amount of energy per active timeslot, regardless of the number of jobs scheduled in that slot (as long as the number of jobs is non-zero). The authors present a linear-time algorithm for the case where jobs are unit length and have a

L. Epstein (✉)
Department of Mathematics, University of Haifa, Haifa, Israel
e-mail: lea@math.haifa.ac.il

P. Ferragina
Dipartimento di Informatica, University of Pisa, Pisa, Italy
e-mail: ferragina@di.unipi.it

single intervals as their feasible regions; for general feasible regions, they show that $B = 2$ admits an efficient solutions whereas the problem is NP-complete starting from $B = 3$. Other variants of the problem are also addressed, thus advancing the existing literature on energy-aware scheduling.

The paper on “The Simplex Tree: An Efficient Data Structure for General Simplicial Complexes” got one of the two Best Paper Awards assigned by the European Association for Theoretical Computer Science (EATCS). It introduces a new data structure, to represent abstract simplicial complexes of any dimension. All faces of the simplicial complex are explicitly stored in a trie whose nodes are in bijection with the faces of the complex. This data structure allows to efficiently implement a large range of basic operations on simplicial complexes. The paper offers a theoretical as well as a detailed experimental analysis, together with a study of Rips and witness complexes.

Another paper on data structures is “Explicit and Efficient Hash Families Suffice for Cuckoo Hashing with a Stash”, which shows that for cuckoo hashing with a stash, very simple families of hash functions can be used, still maintaining the favorable performance guarantees of constant stash size. More specifically, instead of the full randomness needed by previous results, the new approach works with 2-wise independent hash families as building blocks. Moreover these results can be generalized to situations where the stash size is non-constant.

The paper entitled “Polynomial-Time Algorithms for Energy Games with Special Weight Structures” deals with turn-based two-player infinite-duration games played on a weighted directed graph. These games are one of the rare and intriguing combinatorial problems that lie in $NP \cap co-NP$, but are not known to be in P. The existence of polynomial-time algorithms has been a major open problem for decades and apart from pseudo-polynomial algorithms there is no algorithm that solves any non-trivial subclass in polynomial time. In this paper, the authors provide several results based on the weight structures of the graph and a novel technique which converts a non-trivial approximation algorithm into a polynomial time exact algorithm for the problem. Other corollaries referring to specific distributions of weights, ergodic structure of the graph, or boundedness of the clique-widths are addressed.

In “Resource Buying Games” the authors study the existence and computational complexity of pure Nash equilibria for those kinds of problems. The cost of a resource depends on its load. In contrast to classical congestion games for which equilibria are guaranteed to exist, the existence of equilibria in resource buying games strongly depends on the underlying structure of the subsets of resources and the behavior of the cost functions. Authors show that for marginally non-increasing cost functions, matroids are exactly the right structure to consider, and that resource buying games with marginally non-decreasing cost functions always admit a pure Nash equilibria.

The paper entitled “Parameterized Complexity of Induced Graph Matching on Claw-Free Graphs” studies the problem of finding k disjoint induced subgraphs isomorphic to a given graph H within a given graph G and such that there are no edges between vertices of different subgraphs. This problem generalizes the classical Independent Set and Induced Matching problems. The authors show that Induced Graph Matching is fixed-parameter tractable in k on claw-free graphs when H is a fixed connected graph, and admits a polynomial kernel when H is a complete graph. Both results rely on a novel algorithmic structure theorem for claw-free graphs. Other nega-

tive results complement the above positive results whenever H includes some special subgraphs.

We conclude this foreword with the paper “Routing Regardless of Network Stability”, which studies how effective at routing packets are inter-domain routing protocols, such as the Border Gateway Protocol. Theoretical analysis has attempted to answer this question by focusing upon protocol stability; in the present paper authors analyze the effectiveness of packet routing by modeling the interactions between the control plane and the forwarding plane (which determines where packets are forwarded), for the broad class of next-hop preferences with filtering. The paper offers negative (strong inapproximability) results, as well as positive results that guarantee the possibility for all packets to get routed without stability. These positive results are tight—for the general case of filtering lists of cardinality one, it is not possible to ensure that every packet will eventually route.

June 2014

Leah Epstein and Paolo Ferragina