

# Topic 12: Theory and Algorithms for Parallel Computation (Introduction)

Giuseppe F. Italiano, Henning Meyerhenke, Guy Blelloch, and Philippas Tsigas

Topic Committee

Parallelism permeates all levels of current computing systems, from single CPU machines, to large server farms, to geographically dispersed "volunteers" who collaborate over the Internet. The effective use of parallelism depends crucially on the availability of faithful, yet tractable, computational models for algorithm design and analysis and models of efficient strategies for solving key computational problems on prominent classes of computing platforms. Equally important are good algorithmic models of the way the different system components are interconnected. With the development of new genres of computing platforms, such as multicore parallel machines, desktop grids, clouds, and hybrid GPU/CPU-based systems, new computational models and paradigms are needed that will allow parallel programming to advance into mainstream computing. Topic 12 focuses on contributions providing new results on foundational issues regarding parallelism in computing and/or proposing improved approaches to the solution of specific algorithmic problems.

This year, papers submitted to Topic 12 covered nearly all subjects indicated in the call for papers. Among others, subjects included computation and/or communication complexity issues in various computational models, parallel algorithms and data structures for fundamental problems from various domains, and energy considerations in multiprocessor systems. Submissions indicated a significant interest of the parallel computing community towards developing new sound and solid methods for parallel problem solving in the presence of new technological challenges such as increasing core numbers per chip, deep memory hierarchies, complex distributed parallelism, and heterogeneity. Limitations and correctness of parallelism were also under investigation, for example in case of using accelerators.

Among all submissions, three high-quality papers were selected for presentation at the conference. The first paper, "Efficient Parallel and External Matching", by Marcel Birn, Vitaly Osipov, Peter Sanders, Christian Schulz, and Nodari Sitchinava, investigates the local-max algorithm for approximating a maximum weight matching. The algorithm is shown to run in a logarithmic number of phases, incurring work linear in the input size. Also, as their main contribution on parallel aspects of the topic, the authors provide efficient implementations in several models of computation and show that parallel local-max performs well in practice.

The second paper, "Model and complexity results for tree traversals on hybrid platforms", by Julien Herrmann, Loris Marchal, and Yves Robert, analyzes

scheduling of DAGs in a streaming-like model on hybrid systems with a focus on memory usage. The employed model assumes that there are two types of execution units, where each of those has its own (limited) amount of memory. The authors provide complexity results for the question whether a given task graph can be executed given two limited memories, e.g. they show that it is impossible to approximate the optimal consumption of both memories simultaneously by any pair of constant factors.

The third paper, "Splittable Single Source-Sink Routing on CMP Grids: A Sublinear Number of Paths Suffice", by Adrian Kosowski and Przemysław Uznański, addresses the problem of power consumption for routing communication messages between cores of a single-chip multiprocessor assembled in a grid topology. It is shown that (and how) optimal power consumption (within constant factors) can be achieved by splitting requests into a sublinear number of paths.