Topic 8: Distributed Systems and Algorithms

Andrzej Goscinski, Marios Mavronicolas, Weisong Shi, and Teo Yong Meng

Topic Committee

The increasing significance of *Distributed Computing* becomes more and more crucial with the prevail of technological advances that make *Global Computing* a reality in modern world. Indeed, it is hard to imagine some application or computational activity and process that falls outside *Distributed Computing.* With the large advent of distributed systems, we are faced with the real challenges of distributed computation: How do we cope with asynchrony and failures? How (and how well) do we achieve load balancing? How do we model and analyze malicious and selfish behavior? How do we address mobility, heterogeneity and the dynamic nature of participating processes? What can we achieve in the presence of disconnecting operations that cause network partitioning?

These and many more are some of the questions that are routinely scrutinized under the light of current research in *Distributed Systems and Algorithms,* the well-known **Topic 8** of *Europar.* This *Europar* topic provides a forum for both research and development, of interest to both academia and industry, to present and discuss novel approaches to *Distributed Computing* and its relation and connection to *Parallel Processing.* The *Europar 2012* Call for Papers encouraged submission of papers accross the whole spectrum of *Distributed Systems and Algorithms,* with emphasis on several classical and currently popular subareas.

This year five papers were accepted. The paper *Towards Load Balanced Distributed Transactional Memory,* by *G. Sharma* and *C. Busch,* considers the problem of implementing transactional memory in d-dimensional mesh networks. It presents and analyzes *multibend*, a novel load balanced directory-based protocol, which is designed for the *data-flow* distributed implementation of software transactional memory. The paper *CUDA-For-Clusters: A System for Efficient Execution of CUDA Kernels on Multi-Core Clusters,* by *R. Prabhakar, G. Ramaswamy* and *M. J. Thazhuthaveetil,* presents and explores *CUDA* as a programming language for multicores and develops in this way *CUDA-For-Clusters (CFC),* a framework that transparently orchestrates execution of *CUDA* kernels on a cluster of multi-core machines. The paper *From a Store-collect Object and Ω to Efficient Asynchronous Consensus,* by *M. Raynal* and *J. Stainer,* presents an efficient algorithm to build a consensus object, which is based on an Ω failure detector (to obtain liveness) and a *store-collect* object (to maintain its safety). The paper *An Investigation into the performance of reduction algorithms under load imbalance,* by *P. Marendic, J. Lemeire, T. Haber*, *D. Vucinic,* and *P. Schelkens*, investigates contexts where it is not guaranteed that all processes start *reduction* at about the same time; this is a common context in practice, where significant load imbalances may occur and affect the performance of algorithms. The paper investigates the impact of such imbalances on the most commonly employed reduction algorithms and propose a new algorithm specifically adapted

C. Kaklamanis et al. (Eds.): Euro-Par 2012, LNCS 7484, pp. 401–402, 2012.

for such contexts. The paper *Achieving Reliability in Master-worker Computing via Evolutionary Dynamics,* by *E. Christoforou, A. Fernández Anta, C. Georgiou, M. A. Mosteiro* and *A. Sánchez,* consider Internet-based Master-Worker Computations where a master process sends tasks, across the Internet, to worker processes; workers execute and report back some result but they are not trustworthy. To this respect, the paper models such computations using evolutionary dynamics and studies the conditions under which the master can reliably obtain tasks results. The paper develops and analyzes an algorithmic mechanism that uses reinforcement learning to provide workers with the necessary incentives to eventually become truthful.

We would like to take this opportunity to thank all authors who submitted their work to Topic 8 of *Europar 2012,* all external referees who assisted us, and all people involved in organizing the review process for their hard work.