Topic 8: Distributed Systems and Algorithms (Introduction)

Achour Mostefaoui, Andreas Polze, Carlos Baquero,
Paul Ezhilchelvan, and Lars Lundberg

Topic Committee

Distributed Computing is becoming more and more led by technological and application advances. Many works consider new computing models compared to the classical closed model with a fixed number of participants and strong hypothesis on communication and structuration. Indeed, it is hard to imagine some application or computational activity and process that falls outside Distributed Computing. Internet and the web (e.g. social networks, clouds) are becoming the main application field for distributed computing. In addition to the classical challenges that developers have to face (asynchrony and failures) they have to deal with load balancing, malicious and selfish behaviors, mobility, heterogeneity and the dynamic nature of participating processes.

Topic 8 of Euro-Par (Distributed Systems and Algorithms) makes a good mix between research and development. Papers submitted to Topic 8 gave a good overview of the spectrum of Distributed Systems. They focussed on a range of interesting research areas, such as web oriented applications, data managements (data bases) and fault-tolerance. The accepted papers also represent this diverse research landscape, thus making Topic 8 of Euro-Par a good forum to discuss both novel approaches and connections between sub-areas of research in Distributed Systems.

This year five papers were accepted. The paper "Gunther: Search-Based Autotuning of MapReduce", by Guangdeng Liao, Kushal Datta and Theodore Willke proposes a novel approach to parameter optimization in Hadoop clusters based on global search algorithms. Namely, the authors present a tool (Gunther) for automatically finding suitable values for some of the more than 200 configurable parameters in Hadoop. Their approach uses trial execution of Hadoop applications and based on these trial executions they use Genetic Algorithms for finding suitable parameter values.

The paper "Multi-criteria checkpointing strategies: optimizing response-time versus resource utilization", by Aurelien Bouteiller, Franck Cappello, Jack Dongarra, Amina Guermouche, Thomas Herault and Yves Robert, discusses the optimization of system utilization during exhaustive checkpoint rollback operations, with a specific focus on Exascale computing. In uncoordinated checkpointing, when one process rolls-back due to exception, the other cooperating processes need not; however they can be blocked from making progress until the rollback and subsequent restoration are complete. The presented work prevents this by running another application in the meantime so that resource usage is maximized.

F. Wolf, B. Mohr, and D. an Mey (Eds.): Euro-Par 2013, LNCS 8097, pp. 367–368, 2013.

The paper "Efficient event prewarning for sensor networks with multi microenvironments", by Yinglong Li, describes an approach for structuring sensor networks into communication clusters for better performance behavior. To this respect, it proposes an algorithm that eliminates erroneous data by identifying outliers. The paper also proposes a spatial correlation algorithm for generating prewarnings as a function of the location where it was sensed and assesses its efficiency by comparing it with an idealized protocol. In order to improve the quality of these prewarnings the authors suggest spacial correlation analysis.

The paper "On the Scalability of Snapshot Isolation", by Masoud Saeida Ardekani, Pierre Sutra, Nuno Preguica and Marc Shapiro, presents an impossibility result, namely, the proof that genuine partial replication (GPR) cannot be achieved whilst guaranteeing snapshot isolation (SI) consistency as strategies for transaction performance improvements. For this purpose, the authors prove that snapshot isolation (a consistency property) can be broken down into four properties. Then, they show that two of these properties conflict with the property genuine partial replication (GPR). This means that SI and GPR can not be obtained in the same system.

Finally, the paper "Efficient Parallel Block-Max WAND Algorithm", by Veronica Gil Costa, Oscar Rojas and Mauricio Marin, proposes different ways of increasing the performance of web searches (top-k service) through parallel execution on a Master/Slave architecture. The authors consider three different approaches: a distributed two-stage algorithm, and two multi-threaded algorithms using different synchronization schemes (called local and shared heap). The two-stage algorithm is targeted for clusters of distributed computers and the multi-threaded versions are targeted for multi-core computers. This means that the distributed approach and the multi-threaded approaches could be used together if one is using clusters of multiprocessors.

We would like to take this opportunity to thank all authors who submitted their work to Topic 8 of Euro-Par 2013, all external referees who assisted us, and all people involved in organizing the review process for their hard work.