

CoreGRID/ERCIM Workshop on Grids, Clouds and P2P Computing – CGWS2012

Frédéric Desprez¹, Domenico Talia², and Ramin Yaghypour³

¹ INRIA and ENS Lyon, France

² University of Calabria, Italia

³ Georg-August University of Göttingen, Germany

CoreGRID is a European research Network of Excellence (NoE) that was initiated in 2004 as part of the EU FP6 research framework. CoreGRID partners, from 44 different countries, developed theoretical foundations and software infrastructures for large-scale, distributed Grid and P2P applications. An ERCIM sponsored CoreGRID Working Group was established to ensure the continuity of the CoreGrid programme after the original funding period of the NoE. The working group extended its interests to include the emerging field of service-based cloud computing due to its great importance to the European software industry. The working group's main goals consist in i) sustaining the operation of the CoreGRID Network, ii) establishing a forum encouraging collaboration between the Grid and P2P Computing research communities, and (iii) encourage research on the role of cloud computing as a new paradigm for distributed computing in e-Science.

In particular, the ERCIM CoreGRID working group managed to organize an annual CoreGRID workshop, traditionally associated to the Euro-Par conference, thus continuing the successful tradition of the annual CoreGRID workshops during the initial Network of Excellence. Past ERCIM CoreGRID workshops have been organized in Delft (2009), Ischia-Naples (2010), Bordeaux (2011). In 2012 the workshop has been organized in Rhodes Island, Greece. The topics of interest included Service Level Agreements, Data & Knowledge Management, Scheduling, Virtual Environments, Network Monitoring, Volunteer Computing Systems, Trust & Security, Self-* and adaptive mechanisms, Advanced programming models, IaaS, PaaS and SaaS, Tools and Environments for Application Development and Execution.

The 2012 ERCIM CoreGRID workshop was organized jointly with the first workshop in Big Data Management (BDMC) and gathered around 40 researchers from the European community, in August 2012. Six papers were presented after a keynote talk from Frédéric Suter (CC-IN2P3, France) about the SimGrid simulation framework.

The first paper presents an evaluation of a distributed storage plugin for Cumulus, an S3-compatible open-source Cloud service over Grid'5000. The results show that the application, while managing big chunks of data, is able to scale with the size of data and the number of processes. The second paper describes the extensions added to FastFlow, a structured parallel programming framework targeting shared memory multi-core architectures, to support the execution of

programs structured as fine grain parallel activities running on a single workstation. Experiments are conducted over state-of-the-art networked multi-core nodes. The third paper studies the management of pipeline workflow applications that are executed on a distributed platform with setup times. Based on a theoretical study of the problem, the authors provide an optimal algorithm for constellations with identical buffer capacities. For the problem with non-fixed buffer sizes, a $b/(b+1)$ -approximation algorithm is presented. The fourth paper presents the design and implementation of a meteorological application using the ASKALON environment comprising graphical workflow modeling and execution in a Cloud computing environment. The experiments show that a good speedup can be obtained when executed in a virtualized Cloud environment with important operational cost reductions. The next paper describes a workload archive acquired at the science-gateway level. Its added value is demonstrated on several case studies related to user accounting, pilot jobs, fine-grained task analysis, bag of tasks, and workflows. Results show that science-gateway workload archives can detect workload wrapped in pilot jobs, improve user identification, give information on distributions of data transfer times, make bag-of-task detection accurate, and retrieve characteristics of workflow executions. The sixth and last paper proposes a novel mechanism for energy adaptation in P2P file sharing protocols to enhance the possibility of a client completing the file download before exhausting its battery. The proposed mechanism can be implemented in BitTorrent. An analysis is given through a comprehensive set of simulations.

We wish to thank all who contributed to the success of the workshop: authors submitting papers, invited speakers, colleagues who refereed the submitted papers and attended the sessions, and Euro-Par 2012 organizers whose invaluable support greatly helped in the organisation of the Workshop.

October 2012

F. Desprez, D. Talia, R. Yayhapour