Topic 7: Peer-to-Peer Computing
(Introduction)

Damiano Carra, Thorsten Strufe, György Dán, and Marcel Karnstedt

Topic Committee

In Peer-to-peer (P2P) systems computers form an overlay network and share their resources (storage, CPU, bandwidth) to implement a service on top of the Internet. P2P computing has a great potential for creating systems that are self-organizing, efficient, and scalable, but it also faces many challenges: dynamic peer arrivals and departures, which may be correlated (e.g., flash crowd effects, or software failures), high variability of resources, and resource heterogeneity. This topic provides a forum for researchers to present new contributions to P2P systems, technologies, middleware, and applications that address key research issues and challenges.

This year, one paper, which was evaluated by four referees, has been accepted for publication in the peer-to-peer track: "Design and Implementation of a Scalable Membership Service for Supercomputer Resiliency-Aware Runtime", by Yoav Tock, Benjamin Mandler, Jose Moreira and Terry Jones from IBM Haifa Research Laboratory, IBM T.J. Watson Research Center and Oak Ridge National Laboratory. The paper presents the design and implementation of two services for ultra-large HPC systems: a node membership service and an attribute replication service. For such services, the design uses techniques from Peer-to-Peer computing. To deal with the very large number of nodes, the design is based on a hierarchical structure of the nodes and uses an eventual consistency model. The proposed approach also supports versioning.

We would like to take the opportunity of thanking the authors who submitted a contribution, as well as the Euro-Par Organizing Committee, and the external referees with their highly useful comments, whose efforts have made this conference and this topic possible.

F. Wolf, B. Mohr, and D. an Mey (Eds.): Euro-Par 2013, LNCS 8097, p. 353, 2013.