

Peer-to-Peer Computing

Adriana Iamnitchi¹, Paolo Trunfio¹, Jonathan Ledlie², and Florian Schintke²

¹ Topic Chairs

² Members

After several years of intensive investigation, peer-to-peer computing has established itself as an accepted research topic in the general area of distributed systems. Going beyond the initial file sharing applications that spurred productive research and development, peer-to-peer computing is associated with inherently decentralized, self-organizing, and self-coordinating large-scale systems. Performance requirements include adaptivity to churn, high resilience to failures, tolerance to network performance variations, and scalability to huge numbers of peers (tens of thousands to millions), but also stronger consistency and security.

This year, fourteen papers were submitted to the Peer-to-Peer Computing track. Each paper was reviewed by at least three reviewers and, finally, we were able to select five regular papers.

Paper “Dynamic publish/subscribe to meet subscriber-defined delay and bandwidth constraints” by M. A. Tariq, G. G. Koch, B. Koldehofe, I. Khan and K. Roethermel, proposes a protocol for publish/subscribe systems in which subscribers are served according to their specified message delay tolerance. In such system, subscribers with tight delay bounds are served first and then forward messages to those with lesser requirements.

In “Combining Hilbert SFC and Bruijn Graphs for Searching Computing Markets in a P2P System” by D. Castellà, H. Blanco, F. Giné and F. Solsona, a resource discovery system for multiple-attribute, range queries in a P2P system is proposed. The solution combines Bruijn graphs with N-trees and Hilbert multi-dimensional space-filling curve functions.

In “Overlay Management for Fully Distributed User-based Collaborative Filtering” by R. Ormándi, I. Hegedus and M. Jelasity, the problem of constructing an overlay network for performing collaborative filtering is addressed. The proposed system builds a network in which similar peers are connected via a directed edge based on a similarity metric.

Paper “A Formal Credit-Based Incentive Model for Sharing Computer Resources” by J. Rius, I. Barri, F. Cores and F. Solsona, proposes an incentive mechanism for increasing peer participation. The mechanism includes a reinvestment policy that distributes credits non-uniformly to participating peers based on their contribution to the network.

Finally, in “Sampling Bias in BitTorrent Measurements” by B. Zhang, A. Iosup, J. A. Pouwelse, D. Epema and H. Sips, a large study on different BitTorrent traces gathered over the last six years is presented. The paper analyzes different strategies of obtaining the traces and whether the methods will lead to a bias in the gathered data.

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 Referees with their highly useful comments, whose efforts have made this conference and this topic possible.