

Introduction

Amitabha Bagchi, Olivier Beaumont, Pascal Felber, and Alberto Montresor

Topic chairs

Peer-to-peer (P2P) systems enable computers to share information and other resources with their networked peers in large-scale distributed computing environments. The resulting overlay networks are inherently decentralized, self-organizing, and self-coordinating. Well-designed P2P systems should be adaptive to peer arrivals and departures, resilient to failures, tolerant to network performance variations, and scalable to huge numbers of peers (tens of thousands to millions). As P2P research becomes more mature, new challenges emerge to support complex and heterogeneous decentralized environments for sharing and managing data, resources, and knowledge with highly dynamic and unpredictable usage patterns. 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, three papers have been accepted for publication in the P2P track. The general trend among submitted papers was on the study of the properties of P2P networks and their extensions to new services, rather than on the design of new overlays. Each paper was evaluated by four referees.

The paper *Asynchronous Peer-to-Peer Data Mining with Stochastic Gradient Descent* by Róbert Ormándi, István Hegedűs and Mark Jelasity from the University of Szeged, Hungary, proposes a method based on stochastic gradient search to support fully decentralized data mining, with no assumptions on the reliability or synchrony of communication. The idea of applying stochastic gradient descent for SVMs to P2P platforms is particularly original and, in keeping with current trends in computing, opens out the possibility of using a P2P system as a decentralized and dynamic database, thereby creating interesting perspectives for future applications.

The other two papers accepted speak directly to the notion of dynamism that underlies P2P networks by addressing the problem of churn. The first of these, *Evaluation of P2P Systems Under Different Churn Models: Why Should We Bother?* by Marc Sánchez-Artigas and Enrique Fernández-Casado from the Universitat Rovira i Virgili, Spain, helps place the evaluation of P2P systems on a more rigorous basis by investigating the relationships between four different models suggested in the literature for churn in P2P systems. The authors study statistical properties of these models and highlight their similarities and differences. The purpose of this work is to determine if there are significant variations between the models, and hence if they provide different insights when used to study P2P systems.

The paper *ChurnDetect: Gossip-based Churn Estimator for Large-Scale Dynamic Networks* by Andrei Pruteanu, Venkat Iyer and Stefan Dulman from Delft University of Technology, the Netherlands, presents an algorithm to detect in a

distributed way the rate of nodes joining/leaving in multi-hop large scale networks, even in presence of nodes behind firewalls. The algorithm relies on gossip-based communications along with a periodic reset mechanism, and mixes ideas coming from the P2P and ad-hoc communities.

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.