Topic 7: Peer-to-Peer Computing

Dick Epema[1], Márk Jelasity[2], Josep Jorba[3], and Alberto Montresor[4,⋆]

[1] Delft University of Technology, The Netherlands
[2] Szeged University, Hungary
[3] Universitat Oberta de Catalunya, Spain
[4] University of Trento, Italy

After a decade of intensive investigation, peer-to-peer computing has established itself as an accepted research field in the general area of distributed systems. Peer-to-peer computing can be seen as the democratization of computing—overthrowing the old regime of hierarchies as in client-server systems—largely brought about by last-mile network improvements which have made individual PCs first-class citizens in the computer world. Initially, much of the focus in peer-to-peer systems was on best-effort file sharing. However, over the last few years, research has also been directed at trying to make peer-to-peer systems satisfy properties and have functionality as exhibited by more traditional forms of distributed systems. Examples of such properties and functionality, which in principle run counter to the disorganized nature of peer-to-peer systems, are certain levels of security and providing reliable distributed storage such as required by databases.

Eighteen papers were submitted to this year's Peer-to-Peer Computing track, out of which only four were accepted. Two of these consider problems in the area of security in peer-to-peer systems, one deals with mapping optimization problems to peer-to-peer systems, while the fourth is in the area of query processing in peer-to-peer systems.

In "Scalable Byzantine Fault Tolerant Public Key Authentication for Peer-to-Peer Networks", a group-based scheme for authenticating peers based on public keys is presented that is resilient to Byzantine behavior. When a peer wants to authenticate another peer, it instructs the members of its trusted group to challenge that peer and to send back to it the results of their challenges, based on which it can or cannot authenticate the other peer. An algorithm for maintaining the trusted groups based on a small subset of pre-trusted peers is part of the scheme.

In "Secure Forwarding in DHTs—Is Redundancy the Key to Robustness?", the problem of routing messages to the right node(s) in the presence of malicious peers who may simply drop messages is studied. The two approaches of using reduncancy by having multiple paths and of improving routes by routing messages through selected peers with a high reputation are compared. A new algorithm for either approach is proposed, with a reputation system assumed to be present in the latter case.

In "P2P Evolutionary Algorithms: A Suitable Approach for Tackling Large Instances in Hard Optimization Problems", a population-based method for implementing evolutionary algorithms for solving optimization problems is presented. The population of peers executing the algorithms is structured using a gossiping protocol. With simulations, an assessment is made of such metrics as the computing time of the proposed

⋆ Topic Chairs.

E. Luque, T. Margalef, and D. Benítez (Eds.): Euro-Par 2008, LNCS 5168, pp. 599–600, 2008.

method in comparison to that of a sequential genetic algorithm and the required population size.

In ”Efficient Processing of Continuous Join Queries Using Distributed Hash Tables”, the execution of join queries on continuous data streams flowing into a distributed hash table is studied by means of a newly proposed algorithm, the performance of which is assessed with simulations. The queries are disseminated with a gossiping protocol, the new tuples entering the system are only indexed into the DHT when a query needs them, and the queries are executed by the peers using a distributed sliding window.

Finally, we would like to thank all authors and all reviewers of the papers submitted to this track for their work.