

Topic 7

Peer-to-Peer Computing

Alberto Montresor, Fabrice Le Fessant, Dick Epema, and Spyros Voulgaris

Topic Chairs

Distributed systems have experienced a shift of scale in the past few years. This evolution has generated an interest in peer-to-peer systems and resulted in much interesting work. Peer-to-peer systems are characterized by their potential to scale due to their fully decentralized nature. They are self-organizing, adapting automatically to peer arrivals and departures, and are highly resilient to failures. They rely on a symmetric communication model where peers act both as servers and clients. As the peer-to-peer concepts and technologies become more mature, many distributed services and applications relying on this model are envisaged in the context of large-scale distributed and parallel systems. This topic examines peer-to-peer technologies, applications, and systems, and also identifies key research issues and challenges. Twenty-six papers were submitted to the track and we accepted six. We organized two sessions, the first devoted to the problem of query management in structured and unstructured overlay networks, the second containing a broader selection of topics.

In “Path Query Routing in Unstructured Peer-to-Peer Networks”, a solution is proposed to the problem of storing an XML database over an unstructured peer-to-peer network. The system combines multi-level bloom filters for path queries with exponentially decaying bloom-filters for neighbourhood knowledge, to decide to which neighbours queries should be forwarded in the network. In “Processing Top-k Queries in Distributed Hash Tables”, a new algorithm is presented, to extract only the first k replies matching a given query in a Distributed Hash Table. Multiple parameters can be taken into account in the query, provided each parameter domain can be split uniformly among some of the peers. In “Multi-dimensional Range Queries Over Structured Overlays” the authors suggest SOMA, a CAN-like overlay for executing range queries on multiple attributes. SOMA is based on a virtual d -dimensional Cartesian coordinate space, where each dimension corresponds to one attribute. SOMA processes range queries in $\log(N)$ routing steps, N being the number of nodes.

In “Asynchronous Distributed Power Iteration with Gossip-based Normalization”, the authors have designed a fully distributed and robust algorithm based on gossip for finding the dominant eigenvector of large and sparse matrices. In “Capitalizing on Free Riders in P2P Networks”, instead of fighting free riders by shutting them out, a mechanism is presented to benefit from them by letting them handle the forwarding of search queries. This mechanism shifts the load of search queries to the free riders, turning them into an asset instead of a menace. The paper “Content-Based Publish/Subscribe Using Distributed R-Trees”, introduces a class of distributed R-trees and their applicability to building content-

based publish/subscribe overlays. The study focuses on the properties of the resulting topology and the accuracy of event dissemination.