EDITORIAL MESSAGE
Special Track on Engineering Large-Scale Distributed Systems

Fabrizio Falchi, ISTI-CNR, Italy
Claudio Lucchese, ISTI-CNR, Italy

Social phenomena like YouTube and Flickr are incontrovertible evidence of users' migration to a new Web overwhelmed by multimedia. In fact, images, videos, music, and other kinds of multimedia objects today constitute about 99% of the Web. Nonetheless, users' chances of a successful search in such a large portion of information are not proportionally supported. Web search is dominated by giants like Google, Yahoo! and Microsoft that exploit centralized text-only indices enriched by an endless toolbox of smart ranking algorithms. Their interest in riding this new tide is witnessed by the acquisition of both Flickr and YouTube. On the other hand, content-based search for image, music and videos has been deeply studied in the last years but it is not yet adopted by the industry because of its cost.

At the same time users are becoming more and more active. There is no doubt that the so called "Web 2.0" marked a new approach of generating and distributing Web content itself, characterized by communities, decentralization of authority, freedom to share and re-use. It is fascinating the opportunity to develop distributed search engines taking advantage of distributed resources controlled by communities, interest groups, single users or providers.

In the last few years Peer-to-Peer systems have been widely used to overcome scalability issues of centralized solutions. P2P algorithms offer robust, scalable and highly available ways of exploiting large pools of storage and computational resources. The approach has been shown to be effective for tasks such as file sharing. In addition, academic research started contributing new insights and paving new promising outlooks for P2P networks. P2P seems the solution that can make a large multimedia information retrieval system able to scale to repositories as large as the Web is. Also, P2P is supposed to lead users to freedom from advertisement-generating, commercial and centralized web search engines. Users interests are not the same of web search engines.

We organized a Special Track of SAC 2008 about this topics called Engineering Large-Scale Distributed Systems. The program committee, comprised of experts from several disciplines, selected 4 papers after a rigorous review process in which every paper was reviewed by 4 reviewers. To ensure high quality of accepted papers, the review process was double blind, i.e., both the reviewers identity was hidden from the authors and the authors' identity was hidden from the reviewers. The accepted papers are:

- *Extending Peer-to-Peer Networks for Approximate Search*
  Alain Mowat, Michael Schumacher, Ion Constantinescu

- *Processing Complex Similarity Queries in Peer-to-Peer*
  Claudio Gennaro, Matteo Mordacchini, Salvatore Orlando, Fausto Rabitti

- *A pull-based e-mail architecture*
  Edson Kageyama, Carlos Maziero, Altair Santin

- *Knowledge-based Semantic Clustering*
  Dominic Jones, John Keeney, Dave Lewis, Dominik Roblek, Declan O'Sullivan

We would like to thank all the program committee members for their fine work reviewing the submitted papers: