

The goal of this workshop is to enlight the contribution of complex systems research towards the goal of a new generation of information architectures built on a self-organizing P2P architectures.

The Web bears the potential of being the world's largest encyclopedia and knowledge base, but we are very far from being able to exploit this potential. Search-engine technologies provide support for organizing and querying information. Yet, the current systems mainly support the simple mass-user queries. While advanced information queries too often require excessive manual pre- and postprocessing.

Approaches based on Collaborative Web information search in an Internet-scale peer-to-peer (P2P) system bear the potential of overcoming the shortcomings of today's Google-style search engine technology. In the approach proposed within the DELIS project every user (peer) has a full-fledged search engine that indexes a small portion of the Web, according to the interest profile of the user. Thus the requirements to the peer-to-peer overlay network are higher than in the usual key lookup or file sharing scenario. Consistent Hashing, aka. Distributed Hash Tables (DHT), is the state of the art data structure for such peer-to-peer networks. Because of the more involved task of evaluating documents and various locality issues new designs are necessary. Therefore work on the design of peer-to-peer networks is an important part of the scope of this workshop.

From a different perspective, Delis is studying large-scale IT infrastructures which, in order to be managed or used, pose challenging data management requirements. Such infrastructures include GRIDs, large networks, peer-to-peer applications etc. Their management requires (at least) the ability of building and querying large resource directories, of collecting and processing (summarising) lots of usage and performance data, also mining such data (in real time or near-real time) for detection of significant episodes (e.g. routing failures, denial of service attacks etc).

Access to management data should be decentralised, thus allowing many components in the system to efficiently query or subscribe to information. Also, management information should be organized and indexed to allow users (or applications) to issue queries at a higher level than supported by current frameworks (e.g. SNMP, HTML/XML), using query languages that support taxonomical reasoning and eventually the full power of ontology languages such as Owl. Again, the peer-to-peer approach based on DHTs with its promises for load balancing is a starting point for our research, to be explored for building large-scale, self-managing data repositories, however requiring a number of extensions.

Emphasis in the workshop will be given to understanding the dynamics and behavior of such a network requires analyses at different levels and scales of the overall network.