

## Topic 5: Parallel and Distributed Databases

Djoerd Hiemstra, Alfons Kemper, Manuel Prieto, and Alex Szalay

U. Twente, TU Munich, U. Complutense Madrid, Johns Hopkins

Euro-Par Topic 5 addresses data management issues in parallel and distributed computing. Advances in data management (storage, access, querying, retrieval, mining) are inherent to current and future information systems. Today, accessing large volumes of information is a reality: Data-intensive applications enable huge user communities to transparently access multiple pre-existing autonomous, distributed and heterogeneous resources (data, documents, images, services, etc.). Data management solutions need efficient techniques for exploiting and mining large datasets available in clusters, peer to peer and Grid architectures. Parallel and distributed file systems, databases, data warehouses, and digital libraries are a key element for achieving scalable, efficient systems that will cost-effectively manage and extract data from huge amounts of highly distributed and heterogeneous digital data repositories.

Each paper submitted to Euro-Par's topic Parallel and Distributed Databases was reviewed by at least three reviewers. Of 11 papers submitted to the topic this year, 3 were accepted, which makes an acceptance rate of 27 %. The three accepted papers discuss diverse issues: database transactions, efficient and reliable structured peer-to-peer systems, and selective replicated declustering.

In their paper *Unifying Memory and Database Transactions*, Ricardo Dias, and João Lourenço present a simple but powerful idea: Combining software transactional memory with database transactions. The paper proposes to provide unified memory and database transactions by integrating the database transaction control with a software framework for transactional memory. Experimental results show that the overhead for unified transactions is low. It is likely that the approach lowers the burden on the application developer.

The paper by Hao Gong, Guangyu Shi, Jian Chen, and Lingyuan Fan, *A DHT Key-Value Storage System with Carrier Grade Performance*, tries to achieve reliability and efficiency in peer-to-peer systems in order to support Telecom services. The proposed design is based on: Adopting a two-layer distributed hash table, embedding location information into peer IDs; Providing one-hop routing by enhancing each peer with an additional one-hop routing table, where super-peers are in charge of updating and synchronizing this routing information. Finally, the approach replicates subscriber data on multiple peers.

Finally, Kerim Oktay, Ata Turk, and Cevdet Aykanat present a paper on *Selective Replicated Declustering for Arbitrary Queries*. The authors present a new algorithm for selective replicated declustering for arbitrary queries. The algorithm makes use of query information available in order to decide on the data assignment to different disks and on which data to replicate respecting space

constraints. Further, it is described how to apply the proposed algorithm in a recursive way for obtaining a multi-way replicated declustering. Experiments show the algorithm outperforms existing replicated declustering schemes, especially for low replication constraints.

We would like to thank the external reviewers for this Euro-Par topic for their helpful reviews, as well as the Euro-Par 2009 organizing committee for excellent coordination of the work.