

# Dynamic Extensible Query Processing in Super-Peer Based P2P Systems\*

Christian Wiesner

Alfons Kemper

Stefan Brandl

University of Passau, Germany

## 1. Introduction

P2P applications have been quite successful, e.g., for exchanging music files, where simple attributes describe the resources. A lot of effort has been put into refining topologies and query routing functionalities of these P2P networks. Less effort has been put into extending the query functionality within P2P networks. Query processing techniques for P2P systems are still inefficient: sending (atomic) queries to the appropriate peers clearly fails for queries which need data from more than one peer to be executed. While quite a few database techniques can be re-used, a P2P data management infrastructure poses additional challenges. Because of the dynamic nature of P2P networks, we can neither assume global knowledge about data distribution, nor are static topologies and static query plans suitable. In contrast to traditional distributed database systems, we cannot assume a complete schema but rather work with a distributed allocation schema which directs query processing tasks from one node to the neighbors. Therefore, we utilize distributed routing indexes which contain both metadata and content information and are hosted by a backbone of super-peers.

## 2. Extensible Distributed Query Processing

To enable dynamic, extensible, and distributed query processing in super-peer based P2P networks, where standard query operators and user-defined code can be executed nearby the data, we distribute query processing to (super-)peers. Therefore, super-peers provide functionality for the management of the indices, query optimization, and query processing. Additionally, we expect that peers provide query processing capabilities to be full members of the P2P network. These query processors can be dynamically extended by operators such that query evaluation plans with user-defined operators, e.g., selection predicates, join predicates, etc., can be pushed from the client to the (super-)peers where they are executed. To enable this, super-peers have to provide an optimizer for generating efficient query plans from the queries they receive. The distribution process is guided by the routing index which is dynamic and corresponds to the data allocation schema in traditional distributed DBMSs. However, as the index is dynamic and dispersed, static query optimization is not possible. Thus, query optimization must also be dynamic and based on the allocation schema of the data known at the super-peer.

Figure 1 illustrates such a super-peer based P2P network with extensible distributed query processing capabilities. The fig-

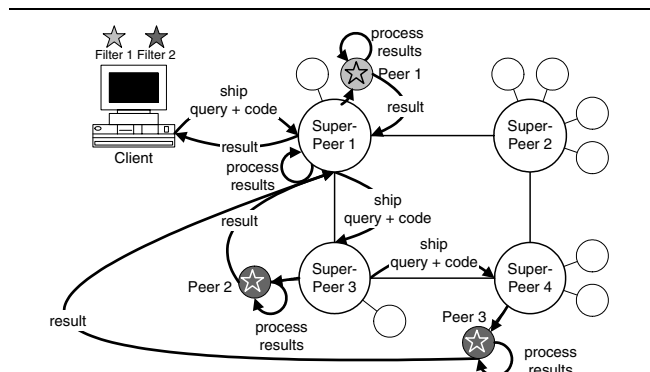


Figure 1. Pushing Code-Carrying Query Plans

ure shows the architecture and the flow of messages in our approach where queries and code are pushed through the network. The client sends the query including user-defined operators to the first super-peer where the local routing indices are consulted and the query is split into two parts. The subplans including the user-defined filters are shipped to the neighbors where the plan optimization continues. This way, operators are pushed to the data sources.

## 3. The Demonstration

We present a P2P architecture based on ObjectGlobe [1] using a super-peer topology where research labs and universities register their projects including description, staff, papers, etc. Users of this network may search for possible research partners, i.e., people who are involved in the same projects and have written many papers, e.g., about *P2P*. As the data is specified using a *semantic* description, these queries could be executed in such a P2P network much more efficiently than by existing search engines. We demonstrate the processing of queries including user-defined operators and the distributed plan generation. Pushing user-defined code towards the data sources will be illustrated using operators for the analysis and filtering of PDF and XML documents.

## References

- [1] R. Braumandl, M. Keidl, A. Kemper, D. Kossmann, A. Kreutz, S. Seltzsam, and K. Stocker. ObjectGlobe: Ubiquitous query processing on the Internet. *The VLDB Journal: Special Issue on E-Services*, 10(3):48–71, August 2001.
- [2] I. Brunkhorst, H. Dhraief, A. Kemper, W. Nejdl, and C. Wiesner. Distributed Queries and Query Optimization in Schema-Based P2P-Systems. In *International Workshop on Databases, Information Systems and Peer-to-Peer Computing*, Berlin, Germany, September 2003. Springer.

\* A detailed description of the proposed techniques is presented in [2].