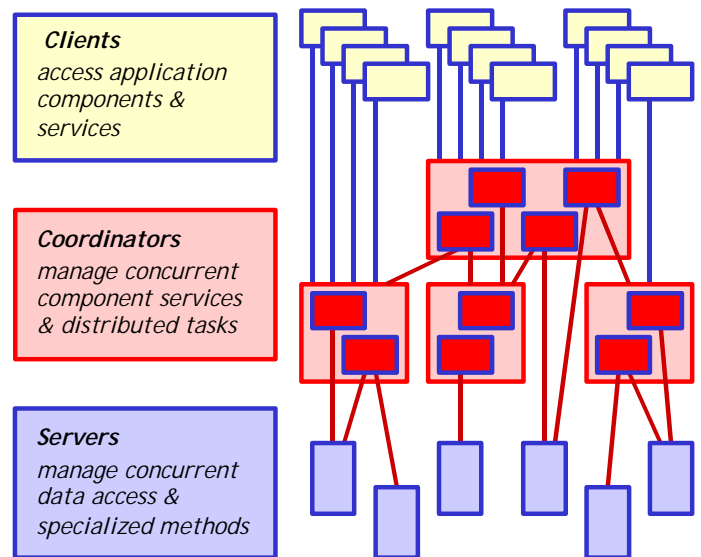


Realizing the Hyperdatabase Vision

Hans-Jörg Schek
DBS Group
ETH Zurich

Figure 1
The Hyperdatabase
Vision



The Vision

The amount of stored information is exploding as a consequence of the immense progress in computer and communication technology during the last decades. However tools for accessing relevant information and processing globally distributed information in a convenient manner are under-developed. The infrastructure for the "Information Space" needs higher attention. In order to improve this situation, we envision the concept of a hyperdatabase that provides database functionality at a much higher level of abstraction, i.e., at the level of complete information components in an n-tier architecture. In analogy to traditional database systems that manage shared data and transactions, a hyperdatabase manages shared information components and transactional processes. It provides "higher-order data independence" by guaranteeing the immunity of applications not only against changes in the data storage and access structures but also against changes in the application components and services, e.g., with respect to the location, implementation, workload, and number of replica of components. In our vision, a hyperdatabase will be the key infrastructure for developing and managing future information systems. At the interface, it will support component and service definition and deployment, specification of transactional

processes encompassing multiple application service invocations, service publication and subscription (see Figure 1). Under the cover, it will perform metadata management, scheduling, optimal routing of service requests, monitoring, flexible failure treatment, availability, and scalability. As illustrated in Figure 2, we have established a number of broad research directions tackling various problems of hyperdatabases:

Transactional coordination in composite systems continues our tradition in transaction research. The two areas below, *database clusters* and *multimedia information management* benefit from transaction research and are two examples of large scale information systems where we explore and realize the hyperdatabase vision by various prototype systems. In the vertical axis, we have established a new area *information dynamics and mobilities* that complements the foundation work in transactions by asynchronous decentralized "coordination". In the following we present a short description of the four areas.

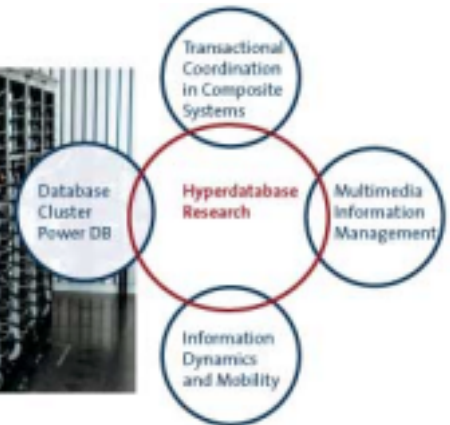
Transactional Coordination in Composite Systems

We have studied the problem of ensuring correctness of concurrent executions in composite n-tier systems. Every coordinator in the composite system performs its transaction management ensuring (local) correctness and (local) recovery. The problem is how global

Figure 3
Physical View on the
database cluster of 128
DBMS nodes



Figure 2
Research areas of the
Database Group



correctness and global recovery is ensured. In the past we have extensively studied this problem from a foundational point of view, and performed several evaluations. In transactional processes, our more recent research activities, we go beyond transactions in that we not only specify conflicts between invocations but we also know about compensation and about retriability. We allow to specify alternative executions and based on these we generalize the “all-or nothing” atomicity of transactions to a notion called guaranteed termination. It means that a single process will eventually terminate along a well-defined path even in case of failure and under concurrency. We have started new investigations aiming at decentralized coordination exploring mobile agent technology, and exploiting cost information of service invocations to optimize scheduling without sacrificing correctness. Work on transactions and transactional processes is a foundation and a basis for transaction implementations in the other main research areas described below.

Database Cluster - PowerDB

In the PowerDB project we explore a hyperdatabase consisting of a set of component databases in a PC cluster as shown in Figure 3. The objective is to explore the limits of scalability and availability of today’s database technology. In every component we have a complete DBMS with its data. Clients access data via the coordinator, i.e., via the hyperdatabase. We explore protocols for high-level transaction management under special consideration of semantic conflicts and of data partitioning and replication. Replication of complete databases contributes to considerable speed-ups in case of read transactions. Due to the second layer transaction management we avoid the disadvantages of traditional commit protocols and of synchronous updates. In addition, query routing aims at detecting components that have sufficiently fresh data and that have the shortest response time due to queries that have been processed before. Replication can be full or be restricted to certain parts of the database. We investigate methods that dynamically allow to add more components to the cluster. We put special emphasis on “Online Analytical Processing in a Cluster of Databases”, and on “XML Document Management with PowerDB”.

Multimedia Information Management

Multimedia information systems consist of many specialized components such as databases, object repositories, special image servers, feature extractors, and indexing components. In several projects in the past, we have developed a prototype system for interactive image similarity search on a PC cluster. The cluster is coordinated by a hyperdatabase that contains descriptions of the components including their actual load. Simple transactional processes for insertion, similarity search, and bulk load can run in parallel and the subtasks are “optimally” and reliably assigned to the components by the hyperdatabase as shown in Figure 4. At any point in time, a new component can be added to the cluster in order to improve response times. Interactive similarity retrieval is based on the VA-File, a simple but efficient approximation of the inherently high-dimensional feature vectors. In order to improve the retrieval effectiveness, we support complex similarity queries consisting of several reference images, several feature types, textual attributes and predicates. In combination with relevance feedback, our similarity search system provides a convenient interface for effective queries, as exemplified in Figure 5. We further apply these techniques to organize, manage, and present the individual information spaces of users in a more natural and efficient way.



Figure 4
Multimedia
components
coordinated by a
hyperdatabase in the
ETHWorld application

Information Dynamics and Mobility

The combination of wireless and wired connectivity along with increasingly small and powerful mobile devices, such as laptops, personal digital assistants, handheld PCs, and smart phones, enables a wide range of new applications that will radically change the way information is managed and processed today. Therefore in this new research area we put strong emphasis on networked information systems where at any point in time nodes may become (partially) disconnected. Nevertheless processing should continue with the objective of afterwards resolving potential conflicts if there are any, i.e., by performing some coordination afterwards when nodes are re-connected. In our vision, information systems will be composed of self-describing and self-organizing mobile information components that are abstractions of both data and application logic



Contact

Prof. Hans-Jörg Schek
Dr. Klemens Böhm
Dr. Heiko Schuldts
Dr. Can Türker
Dr. Roger Weber
Database Research Group
ETH Zürich
Institute of Information Systems
ETH Zentrum
CH-8092 Zürich, Switzerland
Phone +41 1 632 7241
Fax +41 1 632 1172
E-Mail: schek@inf.ethz.ch
<http://www.dbs.ethz.ch>

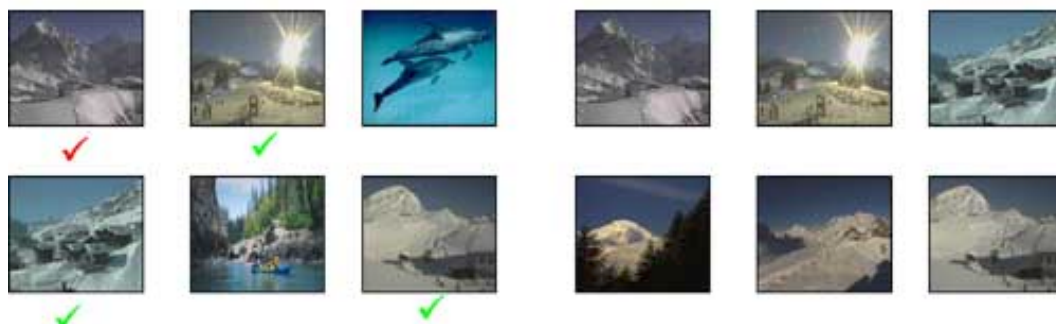


Figure 5
Image Search: Most 5 similar images to the query image (✓) before (left) and after (right) consideration of relevance feedback. Result found in a test collection of 350000 images