Topic 05
Parallel and Distributed Databases and Applications

Bernhard Mitschang

Local Chair

Parallel and distributed database technology is critical for many application domains. This is especially true for conventional high-performance transaction systems, but also for novel and intensive data consuming applications like data warehousing, data mining, decision support, and e-commerce. Future database systems must support flexible and adaptive approaches for data allocation, load balancing, and parallel query processing, both at the DML level and at the transaction level.

This year's Euro-Par topic "Parallel and Distributed Databases and Applications" reflects these trends by focussing on replication management and query evaluation; both topics being viewed as indispensable for modern information systems. In our session we have two papers dealing with replica mamangemnt in a direct fashion looking at algorithms and system realization issues. There is yet another paper indirectly dealing with this topic, in that this technology is among others an indispensable means to built up distributed and parallel application systems. The other remaining paper in our session focusses on issues for a novel communication infrastructure to efficiently support parallel and distributed query processing for distributed relational database management systems.

The first two papers deal with synchronous replica management. The paper by Holliday, Agrawal, and Abbadi explores the benefits of epidemic communication for replica management ensuring serializability. A detailed database simulation is used to explore the performance of the proposed protocol. The paper by Böhm, Grabs, Röhm, and Schek investigates the coordination overhead by means of an experimental assessment. Several setups that compare commercial TP-middleware-based solutions to more or less handcrafted ones are discussed. The third paper of our session by Stillger, Scheffner, and Freytag refers to another important topic for parallel and distributed database technology. The design and implementation of a communication infrastructure for an agent-based distributed query evaluation system is described. Whilst the first three papers adhere to the system and research track as mentioned in the call for papers, the fourth and last paper in our session authored by Peinl stresses the experience and application track. A case study of a large-scale online and real-time information system for foreign exchange trading is presented. Distribution, parallelism, and replica management are discussed as the underlying criteria to system efficiency assessment. In perticular it is shown how much the specific requirements of data replication and parallel processing matched with the paradigms and features

A. Bode et al. (Eds.): Euro-Par 2000, LNCS 1900, pp. 425–426, 2000.

of common off-the-shelf components and why proprietary solutions sometimes seemed inevitably.

All in all we can expect in the near future continued interest in research on parallel and distributed database technology and further interesting in-the-field studies on application experiences.