



Chapter 6

**An Adaptive Probe-Based
Technique to Optimize Join Queries
in Distributed Internet Databases**

Latifur Khan

University of Texas at Dallas, USA

Dennis McLeod and Cyrus Shahabi

University of Southern California, Los Angeles, USA

An adaptive probe-based optimization technique is developed and demonstrated in the context of an Internet-based distributed database environment. More and more common are database systems, which are distributed across servers communicating via the Internet where a query at a given site might require data from remote sites. Optimizing the response time of such queries is a challenging task due to the unpredictability of server performance and network traffic at the time of data shipment; this may result in the selection of an expensive query plan using a static query optimizer. We constructed an experimental setup consisting of two servers running the same DBMS connected via the Internet. Concentrating on join queries, we demonstrate how a static query optimizer might choose an expensive plan by mistake. This is due to the lack of a priori knowledge of the run-time environment, inaccurate statistical assumptions in size estimation, and neglecting the cost of remote method invocation. These shortcomings are addressed collectively by proposing a probing mechanism. Furthermore, we extend our mechanism with an adaptive technique that detects sub-optimality

Previously Published in the *Journal of Database Management*, vol.12, no.4, Copyright © 2001, Idea Group Publishing.

This chapter appears in the book, *Human Computer Interaction Development and Management* by Tonya Barrier.

Copyright © 2002, Idea Group Publishing.

of a plan during query execution and attempts to switch to the cheapest plan while avoiding redundant work and imposing little overhead. An implementation of our run-time optimization technique for join queries was constructed in the Java language and incorporated into an experimental setup. The results demonstrate the superiority of our probe-based optimization over a static optimization.

A distributed database is a collection of partially independent databases that share a common schema, and coordinates processing of non-local transactions. Processors communicate with one another through a communication network (Silberschatz, Korth, and Sudarshan, 1997; Yu and Meng, 1998). We focus on distributed database systems with sites running homogeneous software (i.e., database management system, DBMS) on heterogeneous hardware (e.g., PC and Unix workstations) connected via the Internet. The Internet databases are appropriate for organizations consisting of a number of almost independent sub-organizations such as a University with many departments or a bank with many branches. The idea is to partition data across multiple geographically or administratively distributed sites where each site runs an almost autonomous database system.

In a distributed database system, some queries require the participation of multiple sites, each processing part of the query as well as transferring data back and forth among themselves. Since usually there is more than one plan to execute such a query, it is crucial to obtain the cost of each plan, which highly depends on the amount of participation by each site as well as the amount of data shipment between the sites. Assuming a private/dedicated network and servers, this cost can be computed *a priori* due to the predictability of servers and network conditions and availability of effective network bandwidth. However, in the Internet environment, which is based on a best effort service, there are a number of unpredictable factors that make the cost computation complicated (Paxson and Floyd; 1997). A *static* query optimizer that does not consider the characteristics of the environment or only considers the *a priori* knowledge on the run-time parameters might end up choosing expensive plans due to these unpredictable factors. In the following paragraph, we explain some of these factors via simple examples.

Participating sites (or servers) of Internet database systems might have different processing powers. One site might be a high-end multiprocessor system while the other is a low-end PC running (say) Windows NT. In addition, since most queries are I/O intensive, a site having faster disk drives might observe a better performance. In an Internet-based environment these sites might be dedicated to a single application or multiple simultaneous applications. For example, one site might only run a database server while the other is a database server, a web server,

22 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/adaptive-probe-based-technique-optimize/22208

Related Content

Smart Phone Activity: Risk-Taking Behaviours and Perceptions on Data Security among Young People in England

Balbir S. Barn, Ravinder Barnand Jo-Pei Tan (2013). *International Journal of Social and Organizational Dynamics in IT* (pp. 43-58).

www.irma-international.org/article/smart-phone-activity/114983

Integrating Digital Technologies in Accounting Preservice Teacher Education: A Case Study in Portugal

Ana Luísa Rodrigues (2022). *International Journal of Technology and Human Interaction* (pp. 1-19).

www.irma-international.org/article/integrating-digital-technologies-in-accounting-preservice-teacher-education/293200

Technological Beyond: Merging of Man and Machine

Enric Bas (2021). *Technological Breakthroughs and Future Business Opportunities in Education, Health, and Outer Space* (pp. 64-81).

www.irma-international.org/chapter/technological-beyond/276252

A Bluetooth User Positioning System for Locating, Informing, and Extracting Information using Data Mining Techniques

John Garofalakis and Christos Mettouris (2011). *Emerging Pervasive and Ubiquitous Aspects of Information Systems: Cross-Disciplinary Advancements* (pp. 132-155).

www.irma-international.org/chapter/bluetooth-user-positioning-system-locating/52435

Giving Voice to Feminist Projects in MIS Research

Lynette Kvasny, Anita Greenhill and Eileen M. Trauth (2005). *International Journal of Technology and Human Interaction* (pp. 1-18).

www.irma-international.org/article/giving-voice-feminist-projects-mis/2857