

# Adaptive Algorithms for Join Processing in Distributed Database Systems

PETER SCHEUERMANN\*

*Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL 60208*

peters@ece.nwu.edu

EUGENE INSEOK CHONG†

*New England R&D Center, Oracle Corporation, 1 Oracle Drive, Nashua, NH 03062*

echong@us.oracle.com

*Received May 10, 1996; Accepted February 19, 1997*

**Recommended by:** Clement Yu

**Abstract.** Distributed query processing algorithms usually perform data reduction by using a semijoin program, but the problem with these approaches is that they still require an explicit join of the reduced relations in the final phase. We introduce an efficient algorithm for join processing in distributed database systems that makes use of bipartite graphs in order to reduce data communication costs and local processing costs. The bipartite graphs represent the tuples that can be joined in two relations taking also into account the reduction state of the relations. This algorithm fully reduces the relations at each site. We then present an adaptive algorithm for response time optimization that takes into account the system configuration, i.e., the additional resources available and the data characteristics, in order to select the best strategy for response time minimization. We also report on the results of a set of experiments which show that our algorithms outperform a number of the recently proposed methods for total processing time and response time minimization.

**Keywords:** distributed query processing, join algorithms, adaptive algorithms, bipartite graphs

## 1. Introduction

Query processing in distributed database systems often requires the transmission of relations and/or temporary results among different sites via a computer network. Until recently it has been assumed that communication costs were the predominant costs and hence much of the research on query optimization in distributed database systems has concentrated on producing optimal or near-optimal strategies with regard to data transmission costs [1, 3, 5, 6, 8, 12-14].

In order to reduce the communication time, most algorithms for distributed query processing involve three phases [20, 30]: 1) a local processing phase which includes all local operations such as selections and projections, 2) a reduction phase in which an effective sequence of semi-join (and join) operations is further used to reduce the sizes of the relations,

\*The work of this author was partially supported by NSF grant IRI-9303583 and NASA-Ames grant NAG2-846.

†The work of this author was performed while he was with Northwestern University.

