Applying Multiple-Query Optimization Techniques to Database Processing

Park, Jooseok

Department of Business Administration, Kyung Hee University

Abstract

The objective of this paper is to develop an extensible query processing system for database applications using multiple queries. Such database applications frequently occur in deductive database, active database, and object-oriented database environments.

The general strategy used for the extensible query optimizer is to identify common subexpressions and to use them in simultaneous optimization of multiple queries. In this paper, our efforts are concentrated on extending operator graphs to include processing related information. Note that most previous works assumed no access to a base table through an index, no other join method than the nested-loop method for a join operator, and no pipelined and sorting operations. Using that information, better cost/benefit analysis can be done with regard to sharing activities, and we can also detect new sharing opportunities (e.g., sharing of direct access paths, sorting operations, or join methods) related to common operations rather than common intermediate results.

In order to capture various cases of sharing opportunities for the multiple-query optimization model, we also try to categorize possible sharing abilities of the query processing system; five cases are given in this paper, depending on the basic sharing unit and its storage location. Particularly, by allowing each tuple to have an attribute for the assignment of plan identifiers to it, Case 5 provides sharing opportunity (e.g., compression of different results into one for the same

subsequent access operations) which are completely different from the other proposals. For each case, we show how to detect sharing activities between two tasks, and provide the general savings functions for each sharing activity. We also illustrate their benefit through computational results.

Finally, we focus on an efficient and effective algorithm transforming multiple queries to an optimal global plan. A set of transformation rules and a state-space search algorithm are presented in order to construct an optimal global plan efficiently. These rules transform access plan instances into an integrated graph with only the necessary information for the algorithm. The state-space search algorithm performs two levels of optimization: (1) task-level optimization for the global sharing among tasks and (2) plan-level optimization for minimization of the global plan's cost.