

The Query Optimizer in Tandem's new ServerWare SQL Product

Pedro Celis
Tandem Computers
celis@mpd.tandem.com

Tandem has re-written its SQL compiler and its query execution engine into a new product that will be available in multiple operating systems. The new product uses a very different query optimization engine and the talk will highlight the unique aspects of the new query optimizer focusing primarily on the search engine.

A brief overview of NonStop SQL product will be followed by a description of what is different in ServerWare SQL. The current NonStop SQL optimizer uses a traditional bottom-up dynamic programming optimizer. This is the same type of optimization algorithm used in System R and many commercial products. The optimizer in the new product is a top-down branch and bound rule-driven cost based optimizer similar to work done on the Volcano optimizer.

The time and space complexity of several common SQL optimization search spaces will be presented. We will then contrast the space performance of the top down optimizer search engine with the more traditional bottom up search engine used in NonStop SQL and other commercial products showing where dynamic programming is more efficient.

The problems in controlling the time complexity of the top down search engine will then be discussed. The time complexity of rule driven optimizers is very sensitive to the set of rules used to describe the search space. Different rule-sets that describe the same space can produce very different time complexities for the enumeration of that space.

The top down optimization engine does provide very significant benefits in the extensibility and ease of adding new query transformations that should be done on the basis of cost. Dynamic programming optimizers like Starburst have dealt with this problem only by modifying the query graph and re-starting the optimization process. The talk will conclude with a brief comparison of results.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

**Proceedings of the 22nd VLDB Conference
Mumbai(Bombay), India, 1996**