

Parallel Index Building in Informix OnLine 6.0

*Wayne Davison
Informix Software, Inc.
921 SW Washington, Suite 670
Portland, OR 97205
(503) 221-6246*

Abstract

Informix Software, in a joint venture with Sequent Computer Systems, is augmenting its OnLine database server to provide a multiple process, multi-threaded parallel architecture for shared memory computing environments. This development project, known internally as PDQ (Parallel Database Query), is creating a server which divides large database queries into multiple parallel tasks for improved throughput. Queries will be optimized into tree-structured evaluation plans of demand-driven dataflow iterators. The control and synchronization of parallel streams of data is implemented by an "exchange iterator", modelled after the Volcano database system [1,2]. The exchange iterator can provide vertical parallelism (pipelining between processes) or horizontal parallelism (process replication to iterate over partitioned data subsets).

The exchange iterator, and the iterator model of data flow, have been applied to the parallel construction of B-tree indexes. If key value distributions are not already known, the index builder performs random sampling of the data tuples to build key value histograms. The data is then scanned in parallel from multiple disks and repartitioned by the exchange iterator into appropriate bins. Each bin is then sorted in parallel and passed to a B-tree appender iterator which builds an index subtree from the bottom up. The root node of each subtree is routed through another exchange iterator to a master index iterator which performs the final assembly of the completed index.

This talk will provide an overview of the portable MT multithreaded package, a description of the exchange iterator, an outline of the algorithm for parallel index building, and will conclude with performance comparisons of parallel vs. non-parallel index building.

1. G. Graefe, Encapsulation of Parallelism in the Volcano Query Processing System, Proceedings of the ACM SIGMOD Conference, Atlantic City, N.J., May 1990.
2. G. Graefe, Volcano, an Extensible and Parallel Dataflow Query Processing System, to appear in IEEE Trans. on Knowledge and Data Eng.