

DBSim: A Simulation Tool for Predicting Database Performance

Mike Lefler

PRC Inc

lefler_mike@prc.com

Mark Stokrp

PRC Inc

stokrp_mark@prc.com

Craig Wong

PRC Inc

wong_craig@prc.com

PRC is a systems integrator, a company which builds information systems (based on relational databases) for federal, state, and local governments. Our database designers generally use rules of thumb, supplemented by readily available tuning guides, to tune a physical schema. Prudent database designers also perform benchmarks prior to operational turnover to monitor performance and identify bottlenecks in the database before the users do it for them. However, performance questions can't be addressed by benchmarking until after the hardware has been purchased and installed, and testing has begun. Consequently, it is often the case that the new information system can initially suffer serious performance problems when contrasted with the legacy system being replaced. Moreover any denormalization introduced this late in the development cycle will force software applications to be modified and regression-tested as well. This is expensive and time-consuming, and may require a hardware upgrade — affecting both the program budget and the customer's satisfaction with the system integrator!

Under this Independent Research and Development (IRAD) project PRC is developing the Database Performance Simulator (DBSim) to give systems engineers, database designers and database administrators an easy-to-use, PC-based tool to model both the transaction throughput (tps) and performance (seconds per transaction) for a specific database configuration across a

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMOD '96 6/96 Montreal, Canada
© 1996 ACM 0-89791-794-4/96/0006...\$3 50

spectrum of popular RDBMSs. Ultimately, DBSim will let knowledgeable users evaluate a system architecture and physical database schema before the hardware is purchased and DBMS installed to significantly reduce both schedule and performance risk.

DBSim's approach to simulating commercial RDBMSs is highly modular and consists of a hierarchy of submodels:

- The query parser/optimizer, which parses the transaction command and compiles it into a query execution plan (QEP).
- The query execution layer, which sequentially executes the operations of a QEP.
- The buffer pool manager, which manages the RDBMS cache and models locking.
- An I/O layer, which simulates the reading and writing of pages to and from disk.

The simulation generates transactions instantiated from pre-defined classes. When a transaction arrives at the system, a "CONTROL" process drives the transaction through the RDBMS layers and tracks the time spent in each layer. During the simulation, DBSim lets users visually track resource utilization.

At the conclusion of a simulation run DBSim generates reports that summarize throughput and response times by transaction class and submodel, and provides utilization and queuing statistics for the processor, disk and buffer pool resources. However our experience to date has demonstrated that the performance visualization aspects of DBSim are the most beneficial to the database designers.