

PANEL

Database and Transaction Processing Benchmarks

Jim Gray
San Francisco Systems Center
Digital Equipment Corp
Gray@SFbay.enet.dec.com

Panelists: Francois Raab, Information Paradigm
Samuel DeFazio, Sequent Computer Systems
Jim Frew, University of California, Berkeley

This session is a compendium of the emerging performance and price/performance metrics for database systems and transaction processing systems. Each benchmark tries to answer the question: "What computer should I buy?" Clearly, the answer question is "The system that does the job with the lowest cost-of-ownership." Cost-of-ownership includes project risks, programming costs, operations costs, hardware costs, and software costs. It's difficult to quantify project risks, programming costs, and operations costs. In contrast, computer performance can be quantified and compared.

Generic benchmarks are often used in this way as a rough estimate of the relative system performance because the cost of implementing and measuring a specific application on many different systems is usually prohibitive.

Certain generic benchmarks have become so widely recognized that vendors announce the performance of new products in terms of those benchmarks. For example, DEC, HP, and IBM state the relative performance and price/performance of new machines and system releases by stating their ratings on the Transaction Processing Performance Council's benchmark TPC BM A. This practice is becoming increasingly common.

No single metric can measure the performance of computer systems on all applications. System performance varies enormously from one application domain to another. Each system is typically designed for a few problem domains and may be incapable of performing other tasks. For example, most supercomputers lack database and transaction processing software and so are inappropriate for most business applications.

Domain-specific benchmarks are a response to this diversity of computer system use. Each such benchmark specifies a synthetic workload characterizing typical applications in that problem domain. The performance and of this workload on various computer systems then gives a rough estimate of their relative performance on that problem domain. This handbook contains seven domain-specific benchmarks covering database and transaction processing systems. To be useful, a domain-specific benchmark must meet four important criteria. It must be:

Relevant: It must measure the peak performance and price/performance of systems when performing typical operations within that problem domain.

Portable: It should be easy to implement the benchmark on many different systems and architectures.

Scaleable: The benchmark should apply to small and large computer systems. It should be possible to scale the benchmark up to larger systems, and to parallel computer systems as computer performance and architecture evolve.

Simple: The benchmark must be understandable, otherwise it will lack credibility.

This session presents three such benchmarks.

1. The Transaction Processing Performance Council is about to ratify its third benchmark, variously called TPC-C or the Order Entry Benchmark. TPC-C is an online transaction processing benchmark that captures more complex transactions profiles and functionality. It supersedes the simpler TPC-A and TPC-B which are being made obsolete by technology trends.
2. Several vendors of document retrieval systems have devised a benchmark called the Textual Document Retrieval benchmark. It specifies a scaleable text database and a set of queries that characterize the search of the database. This proposed benchmark is being submitted to the TPC for consideration as a standard.
3. A university-industry consortia, called the Sequoia-2000 project, is investigating the application of computers to the measurement, modeling, and prediction of global change. As part of this project, earth scientists, government users, and computer scientists have devised a benchmark that characterizes the database storage and manipulation requirements for the earth-science problem domain. They have also developed a network and a visualization benchmark. The Sequoia-2000 database benchmark is presented here. It characterizes the database needs typical of scientific applications.