

RAPID RESPONSE TRANSACTION PROCESSING

Elapsed-Time Critical Transactions

Transaction processing is an act of deriving information from data. Needed information may be confined to a single record or woven among millions of records. Record-access transaction processing architectures address the first, while set-processing architectures are ideal for the second. Much is known about record-accessing architectures, very little is known about set-processing architectures. Neither architecture is best for all transaction processing needs. Systems intending to provide rapid response for all types of transaction processing need to be a hybrid of both record-accessing and set-processing architectures.

ELAPSED TIME Too late is too late! Critical information acquired too late to be of value, even if only too late by a nano-second, is too late to be of value. For applications requiring rapid access to individual records, record-accessing architectures provide the minimal elapsed time. For applications requiring complex derivation of information from multiple large files, set-processing architectures are vastly superior to record-accessing architectures, [see RIAM]. Fortunately, these two architectures are quite compatible and can work synergistically on the same platform.

SET-PROCESSING The difference between record-access and set-processing is just a difference in emphasis in the execution of a transaction. A transaction process has two distinct components: accessing all necessary data and deriving needed information from that data. Record-accessing architectures emphasize the first, while set-processing architectures emphasize the second.


EXECUTION STRATEGIES Record-accessing strategies first organize the physical representation of the data for rapid access, then structure the information derivation process to take advantage of the access organization. Set-processing strategies rely on capturing the *information essence* of a transaction then adapting the physical organization of data to fit the information derivation needs. With both strategies, raw source data has to be preprocessed for acceptance by a transaction.

COMPARISON OF STRATEGIES A well respected industry benchmark provides a collection of 22 transactions for comparing the ‘query processing **power**’ of transaction processing systems at a specified database size. The results of a comparison of ten systems with a database size of 100GB is reflected by Graph-1. Of the 22 transactions, three transactions account for 42 to 54 percent of the total elapsed time. These three transactions are better suited for set-processing execution strategies than they are for record-accessing execution strategies. Graph-2 shows a re-ranking of the 10 systems based on their ability to respond to these three specific time consuming transactions. Graph-2 includes the preprocessing of raw data. Graph-3 is the elapsed time for just the three critical transactions. These three graphs give very different performance pictures for these systems.


HYBRID SYSTEMS Also included in Graphs 2 & 3 is a comparison on how set-processing based software (XSP) might improve the response of each system. The same elapsed time for each (though percentage improvement varies) is a reasonable expectation since set-processing performance is a function of information derivation, not a function of pre-organized data, and independent evaluation of set-processing software has validated these expectations. These results seem to suggest that hybrid systems coupling record-accessing and set-processing technologies together might provide for very capable rapid response transaction processing systems.

ELAPSED TIME COMPARISONS OF TEN COMMERCIAL SYSTEMS


ALL NORMALIZED TO SYSTEM-0 AT 8 HOURS



Graph-1: Elapsed Time Based On Query-per-Hour Metric



Graph-2: Elapsed Time for Transactions Including Preprocessing Time



Graph-3: Elapsed Time for Transactions Without Preprocessing Time