INFORMATION ACCESS INTENSIVE SYSTEMS

Essential Performance Requirements

D L Childs
Integrated Information Systems
Ann Arbor, Michigan

ABSTRACT

Performance demands of information access systems differ greatly from those of traditional record retrieval systems. The underlying technology supporting database management systems is antithetical to the needs and requirements of information access systems. This short paper will provide a summary of the necessary performance requirements for information access systems, why and how they differ from those of record retrieval systems, and will provide a proposed commercial off-the-shelf, COTS, platform and test environment that can be used to verify (or refute) these assertions.

1. INTRODUCTION

This paper is intended as a summary of conclusions based on recent research. The conclusions are not expected to be popular nor well received. For those curious enough to explore the theoretical foundations of the assertions, ample references are available. For those outside the database community who require more than just record retrieval and who suspect there may be some value in having high performance information access systems, there is a hardware specification and available software that can be used for experimentation.

The paper is divided into two sections. The first examines the distinctions between record retrieval systems and information access systems and why database technology is ill suited for information access. The distinctions are not intended to be judgmental and should not be read as if they were. Record retrieval objectives established for databases are well met by existing technologies. Information access objectives are not.

The second section summarizes recent research resulting from a performance analysis of a well known industry recognized database benchmark oriented for record retrieval, but used for performance comparison in an information access performance analysis, deriving the same query results from the same data source. A COTS hardware platform is configured for those interested in verifying the claimed performance difference between record retrieval and information access systems.

2. INFORMATION & DATA

There is a subtle but critical distinction between information and data. Information is the relationship between objects represented by data items. Data is the structured representation of data items. Simplistic implementations use the structure of the data to reflect data item relationships. There are distinct differences between technologies oriented toward accessing and processing data and technologies oriented toward accessing and processing information. Data access systems, that retrieve records, are ubiquitous while information access systems, that derive relationships, are almost non-existent. *Record retrieval* characterizes the technology for data access, while *set processing* characterizes the technology for information access. Records are located, while sets are derived. The technologies are very different, though both have value.

Record Retrieval Systems: Designed for locating predefined records in pre-structured data arrays.

Information Access Systems: Designed for deriving data relationships from collections of well-defined data.

These are very different design objectives and require very different underlying support technologies. The most dramatic distinction is that record retrieval systems require substantive application investment in database management. While information access system I/O management requires familiarity with set-processing optimization strategies.

2.1 Performance Considerations

Given any mix of applications on any given hardware platform then, ideally, the best system performance can be achieved by providing applications with just the right data, in just the right format, at just the right time. Starting from this ideal and assuming the data is not already available in memory, there is only one performance issue - optimizing the transfer of relevant data from wherever it is to wherever it should be. For record retrieval systems, where every byte of every record is considered relevant data, existing record access strategies are quite effective. For information access systems where data location is subservient to data relationships, record access strategies are grossly inadequate.

2.2 I/O Performance Potential

To most database devotees *I/O Performance Potential* must sound like an oxymoron. It is, in fact, the key component to high performance information access systems. Early developers of data processing applications observed that accessing records in memory was magnitudes faster than accessing records from secondary storage devices. The mantra established at the time was to avoid I/O at all costs. Unfortunately, the unchallenged acceptance of this mantra has inhibited exploration into the discovery and utilization of the I/O performance potential.

XSP Technology Research & Development Series [IAIS_3: 04/13/11]

A close examination of currently available hardware reveals that data can be transfered between disk[1] and RAM at over 512 MB/s per I/O port. With 35 parallel I/O ports a Terabyte of data can be transfered from disk to RAM in less than 60 seconds. Many platforms are available today sporting over 100 I/O ports, yet none of them are taking advantage of the performance potential offered by the available data transfer rates. The question is why? The answer is that traditional record retrieval technologies use performance crippling I/O strategies.

2.3 Performance Crippling I/O Buffers

Record retrieval technology has developed over the last thirty years based on the assumption that indexed record access was the key to high performance data processing. Though it has always been questionable whether or not the assumption is valid for accessing large numbers of records, it is certainly not valid for supporting high performance information access. The reason is simple - indexed record access imposes grossly inefficient use of I/O buffers.

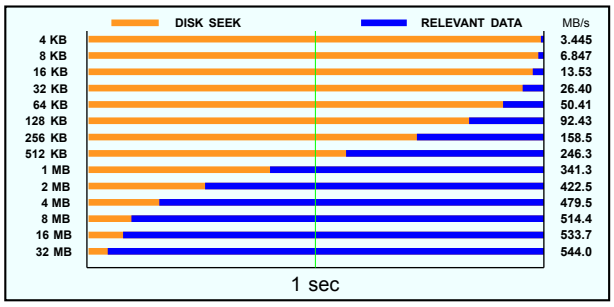

Buffer Size Impact On Relevant Data Transfered

2.3.1 I/O Buffers & SDTR

Sustained Data Transfer Rate, SDTR, is the rate at which a continuous flow of data can be transfered *once the flow has started*. Even with disks having a SDTR of over 512 MB/s record retrieval systems ignore the I/O performance potential by dominate use of small I/O buffers. Just by switching from 4 KB to 8 MB I/O buffers, system performance can improve by a factor of 150, a two magnitude improvement.

2.3.2 Relevant Data Density

Since not every byte of every record retrieved is relevant data, and since indexed access requires partially full I/O buffers, the actual relevant data density of an I/O buffer in practice may be less than 10%. This gives a combined I/O performance potential factor of 1500 times, or a three magnitude improvement.

2.4 Set-Processing Architectures

Record retrieval architectures require applications to be responsible for the representation, organization, and access optimization strategies for secondary storage data. This is as backwards as looking through the wrong end of a telescope. Since each application sharing the same data resource is likely to have very different relevant data access requirements, it is very unlikely that one data storage specification will best serve all applications.

[1]Hitachi Ultrastar C10K300

2.4.1 Relevant Data I/O

As was postulated earlier, the ideal I/O support system would be to provide each application exactly the right data, in exactly the right format, at exactly the right time. For optimized I/O data transfers a set-processing I/O strategy has to recognize and access those *pieces* of files that support each application's relevant data needs.

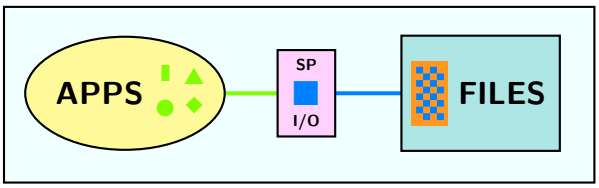

Set-Processing I/O for Relevant Data I/O Transfers.

2.4.2 Set-Processing I/O

Shifting the responsibility of repository data representation, organization, and access optimization from multiple applications to a single system support facility not only improves overall system performance, but also allows applications access to larger volumes of data (thus providing more access to information), substantially reduces application overhead, and ensures inviolate data integrity.

3. SET-PROCESSING I/O SYSTEMS

Set-processing I/O can support system data access by accepting formal (set-theoretically expressed) requests for data and returning just the right data, in just the right format, at just the right time.

Recent research and new advances in disk drive technology show that intensely parallel I/O data access strategies can dramatically outperform traditional record retrieval systems and at a fraction of the cost.

3.1 Intensely Parallel I/O

In theory, the idea is quite simple. If one I/O port can transfer data at 512 MB/s, then two I/O ports, working in parallel can transfer data at 1 GB/s. Pushing the logic, 2048 I/O ports working in parallel can transfer data at 1 TB/s. Somewhere in between is a practical reality.

3.2 Intensely Parallel COTS Platform

One such reality, currently under construction, is a $20,000 48 parallel I/O port COTS platform projected to load 1,000 GB of TPC-H benchmark[2] data in under ten minutes and to exercise all twenty-two queries, as set-processing application relevant data extractions, in under two minutes.

4. CONCLUSION

Since this paper is just a summary of research conclusions, the supporting details are intentionally sparse. However, extensive documentation is available at the XSP Technology website[3] For those actually interested in constructing their own set-processing I/O storage subsystem, the COTS hardware description referenced and the required set-processing operations along with data load and query strategies are available upon request.

[2] A respected record retrieval benchmark. http://www.tpc.org/tpch/

[3] XSP Technology http://xsp.xegesis.org/