

dsmDB: A Distributed Shared Memory Approach for Building Replicated Database Systems

[Extended Abstract]

Nelson Duarte Filho
FURG, Brazil

Fernando Pedone
USI, Switzerland

ABSTRACT

Current trends in main memory capacity and cost indicate that in a few years most performance-critical applications will have all (or most of) their data stored in the main memory of the nodes of a small-size cluster. A few recent research papers have pointed this out and proposed architectures taking advantage of clustered environments aggregating powerful processors equipped with large main memories. This position paper proposes yet another approach, which builds on Distributed Shared Memory systems (DSMs) introduced in the early 80's. We introduce the idea of the dsmDB, discuss how its architecture could be organized, and elaborate on some of its algorithms. We conclude the paper with a discussion of some of its advantages and drawbacks.

1. INTRODUCTION

Many current high-end database applications have stringent performance and availability requirements. Commercial solutions to these applications typically rely upon specialized hardware or proprietary software or both, often inheriting their design from early centralized databases [10]. Increasing demand from high-end applications together with widespread availability of powerful clusters, built out of commodity components, have recently led to alternative approaches to implementing efficient and highly available data management systems.

In modern environments, in-memory architectures are more appropriate than on-disk solutions. Main memory capacity has been growing at a steady pace, approximately doubling every 18 months. There is evidence that most future performance-critical applications will have all (or most of) their data stored in the main memory of the nodes of a small cluster. These trends invalidate many fundamental design decisions of current systems and require a re-evaluation of data structures and algorithms adapted to the new environment.

In-memory database systems (IMDBs) have exploited main memory trends focusing mostly on standalone architectures.

Permission to make digital or hard copies of all or part 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 and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WDDDM '08, March 31, 2008, Glasgow, UK

Copyright 2008 ACM 978-1-60558-121-7/08/03 ...\$5.00.

IMDBs provide high throughput and low response time by avoiding disk I/O. IMDBs were originally designed for specific classes of applications (e.g., telecommunication), but have been recently used in more general contexts (e.g., web servers, trading systems, content caches). In most cases, applications are limited by the memory capacity of the server running the IMDB (e.g., [7]).

In a clustered environment, the database can be partitioned into segments and stored in the main memory of several nodes. Applications are then limited by the aggregated memory capacity of the cluster, and not by the capacity of a single node. This approach has been exploited in recent research [1, 4, 10]. In these works, the database is partitioned and replicated either at the granularity of rows (e.g., horizontal partitioning) or tables. Queries either have to be broken into subqueries for execution at the appropriate node, or the workload has to be special enough so that the original query can be submitted to the node containing all needed information.

In this paper we argue for a different in-memory approach that does not suffer from the shortcomings of previous proposals: there are no limitations on the granularity of partitioning and replication, and queries do not have to be broken up for execution. Our architecture is based on Distributed Shared Memory systems (DSMs) of the early 80's. DSMs extend the notion of virtual memory to different nodes of a cluster. Instead of bringing a page from the local disk upon a page-fault, pages are brought from the main memory of remote nodes. Early DSM systems were expensive to implement, as they offered strong consistency. Later proposals weakened the consistency semantics, improving performance, but required changes in the application programs. Interestingly, semantics that were considered weak in DSMs are enough to implement strong database isolation. In this paper we substantiate this claim with an implementation of one-copy serializability on top of a storage manager that ensures very weak consistency, even for DSM standards.

According to Gray et al.'s terminology [3], the dsmDB belongs to the RAPS category, Reliable Array of Partitioned Services. RAPS can be implemented in shared disk (e.g., Oracle's Real Application Cluster [9]) and shared nothing (e.g., MySQL Cluster [5]) environments. While implementations in the former class require specialized hardware, that is not the case for implementations in the latter class, and therefore for the dsmDB. As pointed out in [3], ideally, data in a RAPS should be automatically repartitioned when nodes are added to the system. In reality, however, this is difficult to achieve—MySQL cluster, for example, requires the en-

tire cluster to be shutdown and all nodes to be synchronized offline when adding a new node [2]. As we claim in this paper, automatic reconfiguration is relatively simple and can be done on-the-fly in the dsmDB.

In Section 2 we review the concept of DSMs. In Section 3 we show how it can be used in the context of replicated databases. We conclude in Section 4 with a discussion of advantages and disadvantages of the approach.

2. DISTRIBUTED SHARED MEMORY

Distributed shared memory systems (DSMs) extend the notion of virtual memory to different nodes of a cluster [6]. With virtual memory, if an accessed page is not loaded in main memory, a page fault is triggered and the page is brought from disk. DSMs allow pages to be fetched from the main memory of other nodes in the cluster.

The DSM model improves performance since, with current technology, bringing a page from the main memory of a remote node in a cluster is faster than bringing the page from the local disk. It also allows programs running on different nodes to easily share pages. However, multiple nodes concurrently reading and writing shared pages raises the problem of consistency.

Many consistency criteria have been defined in the literature. The most intuitive one is that a read should always return the last value written. However, in a distributed system the notion of the last value written is not well defined. Sequential consistency solves this problem by requiring the memory to appear to all nodes as if they were executing on a single multiprogrammed processor.

Although intuitive and well defined, sequential consistency is expensive to implement. To see why, consider that node N wants to write page P . Before the page can be written, N has to invalidate any remote copies of P . The write will be processed once all nodes have replied to N confirming that any future access to P will see the newest value written. Even though there are mechanisms to keep multiple copies of a page on several nodes, they involve some level of cooperation between nodes, and multiple communication steps until the operations can be processed.

Alternatively, weaker consistency criteria have been proposed, which require less synchronization and data movement, resulting in better performance [6]. Weaker consistency can be obtained, for example, by explicitly specifying, through synchronization operators, which parts of the application require strong consistency. Obviously, synchronization is no longer transparent, and to work properly some programs may have to be modified.

In the following sections we argue that the DSM model (with weak consistency) is an appropriate paradigm for building high performance and high availability database systems.

3. THE DSM DATABASE APPROACH

In this section we briefly discuss the architecture of the DSM database approach (dsmDB), explain how database consistency criteria can be built on top of weak shared memory guarantees, and present some of the algorithms involved.

3.1 The dsmDB Architecture

We assume a simplified database architecture composed of a Distributed Storage Manager, a Consistency Manager, and a Query Manager (see Figure 1).

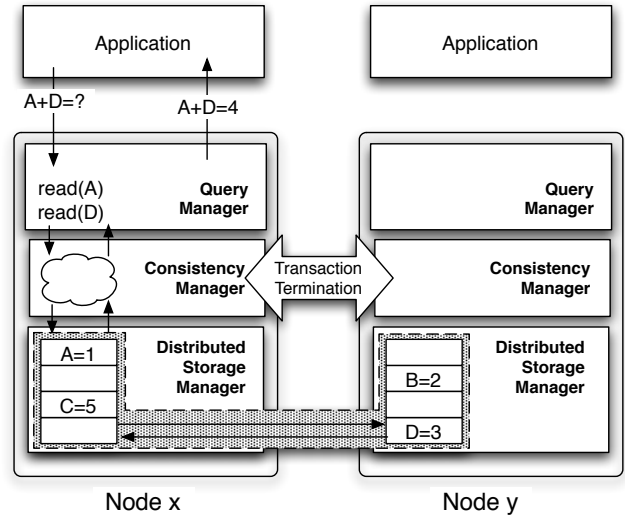


Figure 1: Architecture of the dsmDB

The *Storage Manager* executes read and write operations on the data, providing transactional access to a virtual storage that consists of the complete database space. This illusion is implemented by local and remote memory. Storage Managers may synchronize local access to the virtual database space, in which case two transactions executing on the same node would be mutually consistent. However, transactions executing on different nodes are not synchronized during their execution by Storage Managers.

The *Consistency Manager* ensures that the execution is globally consistent according to some correctness criteria. It uses an optimistic mechanism to guarantee global consistency, similarly to the deferred update replication model [8]: Transactions execute locally on some node and, as part of the transaction termination protocol, are validated. If the transaction passes the validation phase, it is locally committed by the Storage Manager; if it does not pass validation it is aborted. From the perspective of the Consistency Managers, the execution of each transaction is local. In reality, some data may be fetched from remote nodes on demand.

The *Query Manager* is responsible for receiving high level client requests (e.g., SQL statements), pre-processing them and possibly breaking them up into lower level operations to be executed by the Storage Manager. It receives the results from the Consistency and Storage Managers and passes them to the clients, possibly after some post-processing.

In the example in Figure 1, an application program has submitted a request to compute the sum of items A and D to the Query Manager of node x . This request is parsed and results in two read operations, passed to the Storage Manager. The value of A is stored locally and is available immediately. Reading D triggers an exception since it is not stored locally. After D is fetched from a remote node, the Storage Manager replies to the Query Manager, which computes the result and returns it to the application program.

3.2 Distributed Storage Manager

The Distributed Storage Manager provides access to a virtual storage that consists of the complete database space, which hereafter we assume to be a set of blocks. There are

four simple operations: $\text{read}(T, B)$, $\text{write}(T, B, V)$, $\text{commit}(T)$, and $\text{abort}(T)$, where T is a transaction, B a unique block id, and V a value. These primitives provide weak guarantees:

(Read committed.) Every transaction can only read committed data or its own uncommitted writes.

(Data freshness.) If a node x executes a committed write on block B in isolation and no other node writes B at a later time, then eventually every node that successively reads B will see x 's write.

Providing weak consistency (w.r.t. traditional consistency guarantees of Distributed Shared Memory systems) is key to the dsmDB: First, it allows applications to tune the level of global consistency needed through the Consistency Manager. Second, weak consistency can be efficiently implemented by the Storage Manager.

We now briefly describe one possible implementation of the Distributed Storage Manager.

The virtual storage abstraction provided by a node is implemented using its local memory and the memory of remote nodes. At any time, each data block B can be stored by any number of nodes, but it is surely stored by all *home* nodes of B , determined by a deterministic function with B 's id as sole input. Thus, every node keeps locally a set of *permanent* blocks, the ones for which it is home, and a set of *temporary* blocks, blocks that are not permanent and can be discarded at the node's will.

Therefore, we divide the local memory in the Distributed Storage Manager of a node into *permanent* and *temporary*. Permanent storage keeps permanent blocks; temporary storage keeps temporary blocks and ongoing writes of executing transactions. For durability purposes, permanent storage can be located both in main memory and on disk. Alternatively, if blocks have multiple home nodes, then only main memory may suffice to guarantee durability.

Read operations are served immediately from the node's local storage (i.e., from the permanent or temporary storage) or fetched from the block's home node. Write operations are executed locally in temporary storage. At commit time, they may become permanent, temporary, or simply be discarded from the node's memory. To guarantee data freshness, from time to time a node invalidates all blocks in its temporary storage.

The Storage Manager can be optimized in a number of ways. For example, it can discard all write operations of blocks for which it is not the home node. This results in faster response for committing transactions, since fewer operations have to be executed, and leaves more space in the memory buffer for ongoing and new transactions.

3.3 Consistency Manager

Storage Managers implement weak consistency semantics. Stronger guarantees can be obtained by a separate mechanism implemented by the Consistency Manager. In the following we will discuss how One-Copy Serializability (ISR) can be implemented. We adopt an approach similar to the validation process used by the Database State Machine (DBSM) [8].

ISR specifies that any concurrent execution of transactions in a possibly replicated setting should be equivalent to a serial execution of the same transactions in a single node. As a consequence, if two conflicting transactions execute concurrently on different nodes, only one transaction can be allowed to commit. Two transactions conflict if they

access the same database block and at least one transaction writes the block.

The Consistency Manager uses an atomic broadcast abstraction, defined by the primitives $\text{broadcast}(m)$ and $\text{deliver}(m)$, where m is a message. Every message broadcast by the Consistency Manager of a node is delivered by all operational (i.e., nonfaulty) nodes in total order. More formally:

1. (Agreement.) If node x delivers message m , then every nonfaulty node also delivers m .
2. (Total order.) If nodes x and y deliver messages m and m' , they do so in the same order.

During the local execution of T , the Consistency Manager captures its read and write operations, that is, its readset and writeset. Notice that readsets and writesets contain an indication of the data blocks read and written by the transactions, not the contents of the blocks. When the commit operation is requested, T 's readsets, writesets, and the values written are broadcast to all nodes.

Upon delivery, each node x proceeds with the validation of T , which checks whether all blocks read by T are still up to date, i.e., they were not written by another transaction during T 's execution. This is performed by checking T 's readset against the writesets of all transactions that executed concurrently with T and have already committed.

If T passes validation then it can be committed against the local Storage Manager. At the node where T executed, it suffices for the Consistency Manager to send the $\text{commit}(T)$ operation to the Storage Manager and reply to the application; at the other nodes, the local Consistency Manager first executes all writes against the local Storage Manager and then commits the transaction.

4. CONCLUDING REMARKS

This short note discusses the dsmDB, a novel approach to building high performance clustered database systems inspired by DSMs. The dsmDB approach enhances transaction processing in two ways: first, it provides in-memory access of executing transactions; second, it allows nodes to reduce the number of updates to be processed when transactions commit, improving performance. This last property is achieved by the Storage Managers, which can drop write requests for data blocks they are not responsible for (i.e., the node is not the block's home node).

At the top two layers, Query Manager and Consistency Manager, the dsmDB is similar to the Database State Machine (DBSM) [8]. In fact, the Storage Managers implement the illusion of a fully replicated database. Differently from the DBSM, and other fully replicated database protocols, the dsmDB has more flexibility to handle storage.

We enumerate next a number of reasons we believe the dsmDB approach is promising.

- *Location transparency.* The dsmDB offers a transparent mechanism to manage memory and disk, either local or remote. As a consequence, adding nodes to and removing nodes from a cluster can be done on-the-fly, with the resulting performance benefits. It also provides simple system reconfiguration (e.g., changing the unit of partitioning).
- *Concurrency control "à la carte".* By providing weak consistency at the storage level, the dsmDB allows an

efficient implementation. Stronger requirements can be built on top of it in a modular way. We have commented on how ISR can be guaranteed in this context. We conjecture that other strong database consistency criteria (e.g., snapshot isolation) could be also implemented on top of Storage Managers.

- *Simple load balancing.* The dsmDB could be combined with a load balancer in order to group transactions that access common portions of the database on specific nodes, favoring locality. This technique has proved to be quite efficient in similar contexts [4]. The work in [4] assumes that the workload can be divided at the table level, and each table placed on a separate replica. Since the dsmDB provides location transparency, it does not need such an assumption.
- *Flexible data partitioning and replication.* As another consequence of location transparency, the dsmDB allows data to be freely partitioned and replicated. This is an important feature if the database is large but the hot spot portion of the data fits the aggregated main memory of the cluster nodes.
- *Incremental recovery.* Traditional database systems have to undergo a recovery procedure before becoming operational after a crash. This procedure involves operations such as bringing the database state from disk, undoing the effects of unfinished transactions, and making sure that all committed transactions have their effects reflected on the database state. This operation can be sped up by bringing the current state from remote nodes, but in any case new transactions can be only accepted once the complete state has been recovered. The dsmDB allows incremental recovery by simply restarting the crashed node with an empty storage. New transactions can be accepted immediately and data read by these transactions will be fetched from remote nodes on demand.

We have also tried to anticipate the drawbacks of the dsmDB. So far, we identified the following ones.

- *Loss of semantics.* Since the virtual database space is implemented by the Storage Managers, after high level operations are translated into lower level block operations by the Query Manager, some information available for sharing data may be lost. In this sense, moving the virtual shared space to higher modules (e.g., the Query Manager) could be wiser. It's unclear to us though whether the complexity involved in such a design could bring any real benefits.
- *Difficulty to re-engineer into existing database engines.* We are currently building a prototype of the dsmDB by modifying an existing database (Berkeley DB). Despite the relatively simplicity of the Storage Manager, modifying a storage engine to integrate the ideas described in this paper turned out to be more complex than we initially expected. One reason is that although it is easy to identify the portions of the database that handle storage, dependencies across layers make it difficult to introduce the needed modifications.

5. REFERENCES

- [1] L. Camargos, F. Pedone, and M. Wieloch. Sprint: A middleware for highperformance transaction processing. In *Eurosys*, 2007.
- [2] E. Cecchet, G. Candea, and A. Ailamaki. Middleware-based database replication: The gaps between theory and practice. In *SIGMOD '08: Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, 2008.
- [3] B. Devlin, J. Gray, B. Laing, and G. Spix. Scalability terminology: Farms, clones, partitions, and packs: Racs and raps. Technical report, Microsoft, 1999.
- [4] S. Elnikety, S. Dropsho, and W. Zwaenepoel. Tashkent+: Memory-aware load balancing and update filtering in replicated databases. In *Eurosys*, 2007.
- [5] Mysql cluster architecture overview, April 2004. Online Technical White Paper.
- [6] B. Nitzberg and V. Lo. Distributed shared memory: A survey of issues and algorithms. *IEEE Computer*, 24(8):52–60, August 1991.
- [7] Oracle TimesTen products and technologies, February 2007. Online White Paper.
- [8] F. Pedone. *The Database State Machine and Group Communication Issues*. PhD thesis, École Polytechnique Fédérale de Lausanne, Switzerland, 1999. Number 2090.
- [9] Oracle real application cluster 11g, April 2007. Online White Paper.
- [10] M. Stonebraker, S. Madden, D. Abbadi, S. Harizopoulos, N. Hachem, and P. Helland. The end of an architectural era (it's time for a complete rewrite). In *VLDB*, 2007.