

Understanding Replication in Databases and Distributed Systems*

M. Wiesmann, F. Pedone, A. Schiper
Swiss Federal Institute of Technology (EPFL)
Operation Systems Laboratory
IN-F Ecublens, CH-1015 Lausanne

B. Kemme, G. Alonso
Swiss Federal Institute of Technology (ETHZ)
Institute of Information Systems
ETH Zentrum, CH-8092 Zürich

Abstract

Replication is an area of interest to both distributed systems and databases. The solutions developed from these two perspectives are conceptually similar but differ in many aspects: model, assumptions, mechanisms, guarantees provided, and implementation. In this paper, we provide an abstract and “neutral” framework to compare replication techniques from both communities. The framework has been designed to emphasize the role played by different mechanisms and to facilitate comparisons. The paper describes the replication techniques used in both communities, compares them, and points out ways in which they can be integrated to arrive to better, more robust replication protocols.

1. Introduction

Replication has been studied in many areas, especially in distributed systems (mainly for fault tolerance purposes) and in databases (mainly for performance reasons). In these two fields, the techniques and mechanisms used are similar, and yet, comparing the protocols developed in the two communities is a frustrating exercise. Due to the many subtleties involved, mechanisms that are conceptually identical, end up being very different in practice. So, it is very difficult to take results from one area and apply them in the other. In the last few years, as part of the DRAGON project [16], we have devoted our efforts to enhance database replication mechanisms by taking advantage of some of the properties of group communication primitives. We have shown how group communication can be embedded into a database [1, 22, 23] and used as part of the transaction manager to guarantee serialisable execution of transactions over replicated data [17]. We have also shown how some of the overheads associated with group communication can be hidden

*Research supported by EPFL-ETHZ DRAGON project and OFES under contract number 95.0830, as part of the ESPRIT BROADCAST-WG (number 22455).

behind the cost of executing transactions, thereby greatly enhancing performance and removing one of the serious limitations of group communication primitives [18]. This work has proven the importance of and the need for a common understanding of the replication protocols used by the two communities.

In this paper, we present a model that allows to compare and distinguish existing replication protocols in databases and distributed systems. We start by introducing a very abstract replication protocol representing what we consider to be the key phases of any replication strategy. Using this abstract protocol as the base line, we analyse a variety of replication protocols from both databases and distributed systems, and show their similarities and differences. With these ideas, we parameterise the protocols and provide an accurate view of the problems addressed by each one of them. Providing such a classification permits to systematically explore the solution space and give a good baseline for the development of new protocols. While such work is conceptual in nature, we believe it is a valuable contribution since it provides a much needed perspective on replication protocols. However, the contribution is not only a didactic one but also eminently practical. In recent years, and in addition to our work, many researchers have started to explore the combination of database and distributed system solutions [25, 29, 21, 15]. The results of this paper will help to show which protocols complement each other and how they can be combined.

The paper is organised as follows. Section 2 introduces our functional model and discusses some basis for our comparison. Section 3 and Section 4 present replication protocols in distributed systems and databases, respectively. Section 5 refines the discussion presented in Section 4 for more complex transaction models. Section 6 discusses the different aspects of the paper. Section 7 concludes this paper.

2. Replication as an Abstract Problem

Replication in databases and distributed systems rely on different assumptions and offer different guarantees to the

clients. In this section, we discuss the context of replication in databases and distributed systems, and introduce a functional model of replication to allow us to treat replication as an abstract problem.

2.1. Replication Context

Hereafter, we assume that the system is composed of a set of *replicas* over which operations must be performed. The operations are issued by *clients*. Communication between different system components (clients and replicas) takes place by exchanging messages.

In this context, distributed systems distinguish between the *synchronous* and the *asynchronous* system model. In the synchronous model there is a known bound on the relative process speed and on the message transmission delay, while no such bounds exist in the asynchronous model. The key difference is that the synchronous system allows *correct crash detection*, while the asynchronous system does not (i.e., in an asynchronous system, when some process p thinks that some other process q has crashed, q might in fact not have crashed). Incorrect crash detection makes the development of replication algorithm more difficult. Fortunately, much of the complexity can be hidden behind the so called *group communication primitives*. This is the approach we have taken in the paper (see Section 3.1).

Databases do not distinguish synchronous and asynchronous systems since they accept to live with *blocking* protocols (a protocol is said to be blocking if the crash of some process may prevent the protocol from terminating). Distributed systems usually look for non-blocking protocols.

This reflects another fundamental difference between distributed systems and database replication protocols. It has been shown that the specification of every problem can be decomposed into *safety* and *liveness* properties [3].¹ Database protocols do not treat liveness issues formally, as part of the protocol specification. Indeed, the properties ensured by transactions (Atomicity, Consistency, Isolation, Durability) [11] are all *safety* properties. However, because databases accept to live with blocking protocols, liveness is not an issue. For the purpose of this paper, we concentrate on safety properties.

Database replication protocols may admit, in some cases, operator intervention to solve abnormal cases, like the failure of a server and the appointment of another one (a way to circumvent blocking). This is usually not done in distributed system protocols, where the replacement of a replica by another is integrated into the protocol (non-blocking protocols).

¹A safety property says that nothing bad ever happens, while a liveness property says that something good eventually happens.

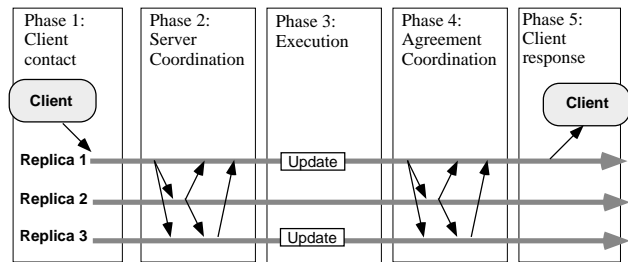


Figure 1. Functional model with the five phases

Finally, distributed systems distinguish between deterministic and non-deterministic replica behaviour. Deterministic replica behaviour assumes that when presented with the same operations in the same order, replicas will produce the same results. Such an assumption is very difficult to make in a database. Thus, if the different replicas have to communicate anyway in order to agree on a result, they can as well exchange the actual operation. By shifting the burden of broadcasting the request to the server, the logic necessary at the client side is greatly simplified at the price of (theoretically) reducing fault tolerance.

2.2. Functional Model

A replication protocol can be described using five generic phases. As we will later show, some replication techniques may skip some phases, order them in a different manner, iterate over some of them, or merge them into a simpler sequence. Thus, the protocols can be compared by the way they implement each one of the phases and how they combine the different phases. In this regard, an abstract replication protocol can be described as a sequence of the following five phases (see Figure 1).

Request (RE): the client submits an operation to one (or more) replicas.

Server coordination (SC): the replica servers coordinate with each other to synchronise the execution of the operation (ordering of concurrent operations).

Execution (EX): the operation is executed on the replica servers.

Agreement coordination (AC): the replica servers agree on the result of the execution (e.g., to guarantee atomicity).

Response (END): the outcome of the operation is transmitted back to the client.

The differences between protocols arise due to the different approaches used in each phase which, in some cases, obviate the need for some other phase (e.g., when messages are ordered based on an atomic broadcast primitive, the agreement coordination phase is not necessary since it

is already performed as part of the process or ordering the messages).

Within this framework, we will first consider transactions composed of a single operation. This can be a single read or write operation, a more complex operation with multiple parameters, or an invocation on a method. A more advanced transaction model will be considered in Section 5.

Request Phase. During the request phase, a client submits an operation to the system. This can be done in two ways: the client can directly send the operation to all replicas or the client can send the operation to one replica which will then send the operation to all others as part of the server coordination phase.

This distinction, although apparently simple, already introduces some significant differences between databases and distributed systems. In databases, clients never contact all replicas, and always send the operation to one copy. The reason is very simple: replication should be transparent to the client. Being able to send an operation to all replicas will imply the client has knowledge about the data location, schema, and distribution which is not practical for any database of average size. This is knowledge intrinsically tied to the database nodes, thus, clients must always submit the operation to one node which will then send it to all others. In distributed systems, however, a clear distinction is made between replication techniques depending on whether the client sends the operation directly to all copies (e.g., active replication) or to one copy (e.g., passive replication).

It could be argued that in both cases, the request mechanisms can be seen as contacting a *proxy* (a database node in one case, or a communication module in the other), in which case there are no significant differences between the two approaches. Conceptually this is true. Practically, it is not a very helpful abstraction because of its implications as it will be discussed below when the different protocols are compared. For the moment being, note that distributed systems deal with *processes* while database deal with *relational schemas*. A list of processes is simpler to handle than a database schema, i.e., a communication module can be expected to be able to handle a list of processes but it is not realistic to assume it can handle a database schema. In particular, database replication requires to understand the operation that is going to be performed while in distributed systems, operation semantics usually play no role.

Server Coordination Phase. During the server coordination phase, the different replicas try to find an order in which the operations need to be performed. This is the point where protocols differ the most in terms of ordering strategies, ordering mechanisms, and correctness criteria.

In terms of ordering strategies, databases order operations according to data dependencies. That is, all operations

must have the same data dependencies at all replicas. It is because of this reason that operation semantics play an important role in database replication: an operation that only reads a data item is not the same as an operation that modifies that data item since the data dependencies introduced are not the same in the two cases. If there are no direct or indirect dependencies between two operations, they do not need to be ordered because the order does not matter. Distributed systems, on the other hand, are commonly based on very strict notions of ordering. From causality, which is based on *potential* dependencies without looking at the operation semantics, to total order (either causal or not) in which all operations are ordered regardless of what they are.

In terms of correctness, database protocols use serialisability adapted to replicated scenarios: one-copy serialisability [6]. It is possible to use other correctness criteria [17] but, in all cases, the basis for correctness are data dependencies. Distributed systems use *linearisability* and *sequential consistency* [5]. Linearisability is strictly stronger than sequential consistency. Linearisability is based on *real-time* dependencies, while sequential consistency only considers the order in which operations are performed on every individual process. Sequential consistency allows, under some conditions, to read *old values*. In this respect, sequential consistency has similarities with one-copy serialisability, but strictly speaking, the two consistency criteria are different. The distributed system replication techniques presented in this paper all ensure linearisability.

Execution Phase. The execution phase represents the actual performing of the operation. It does not introduce many differences between protocols, but it is a good indicator of how each approach treats and distributes the operations. This phase only represents the actual execution of the operation, the applying of the update is typically done in the Agreement Coordination Phase.

Agreement Coordination Phase. During this phase, the different replicas make sure that they all do the same thing. This phase is interesting because it brings up some of the fundamental differences between protocols. In databases, this phase usually corresponds to a Two Phase Commit Protocol (2PC) during which it is decided whether the operation will be committed or aborted. This phase is necessary because in databases, the Server Coordination Phase takes care only of ordering operations. Once the ordering has been agreed upon, the replicas need to ensure everybody agrees to actually commit the operation. Note that being able to order the operations does not necessarily mean the operation will succeed. In a database, there can be many reasons why an operation succeeds at one site and not at another (load, consistency constraints, interactions with local operations). This is a fundamental difference with distributed systems where once an operation has been success-

fully ordered (in the Server Coordinator Phase) it will be delivered (i.e., “performed”) and there is no need to do any further checking.

Client Response Phase. The client response phase represents the moment in time when the client receives a response from the system. There are two possibilities: either the response is sent only after everything has been settled and the operation has been executed, or the response is sent right away and the propagation of changes and coordination among all replicas is done afterwards. In the case of databases, this distinction leads to (1) eager or synchronous (no response until everything has been done) and (2) lazy or asynchronous (immediate response, propagation of changes is done afterwards) protocols. In the distributed systems case, the response usually takes place only after the protocol has been executed and no discrepancies may arise.

The client response phase is of increasing importance given the proliferation of applications for *mobile* users, where a copy is not always connected to the rest of the system and it does not make sense to wait until updates are applied in the entire system to let the user see the changes made.

3. Distributed Systems Replication

In this section, we describe the model and the communications abstractions used by replication protocols in distributed systems, and present four replication techniques that have been proposed in the literature in the context of distributed systems.

3.1. Replication Model and Abstractions

We consider a distributed system modelled as a set of services implemented by server processes and invoked by client processes. Each server process has a local state that is modified through invocations. We consider that invocations modify the state of a server in an atomic way, that is, the state changes resulting from an invocation are not applied partially. The isolation between concurrent invocations is the responsibility of the server, and is typically achieved using some local synchronisation mechanism. This model is similar to “one operation” transactions in databases (e.g., stored procedures). In order to tolerate faults, services are implemented by multiple server processes or replicas.

To cope with the complexity of replication, the notion of *group* (of servers) and *group communication primitives* have been introduced [7]. The notion of *group* acts as a logical addressing mechanism, allowing the client to ignore the degree of replication and the identity of the individual server processes of a replicated service. *Group communication primitives* provide one-to-many communication with

various powerful semantics. These semantics hide much of the complexity of maintaining the consistency of replicated servers. The two main group communication primitives are *Atomic Broadcast* (or ABCAST) and *View Synchronous Broadcast* (or VSCAST). We give here an informal definition of these primitives. A more formal definition of ABCAST can be found in [14] and of VSCAST can be found in [27] (see also [8, 9]). Group communication properties can also feature FIFO order guarantees.

Atomic Broadcast (ABCAST). Atomic Broadcast provides *atomicity* and *total order*. Let m and m' be two messages that are ABCAST to the same group g of servers. The atomicity property ensures that if one member of g delivers m (resp. m'), then all (not crashed) members of g eventually deliver m (resp. m'). The order property ensures that if two members of g deliver both m and m' , they deliver them in the same order.

View Synchronous Broadcast (VSCAST). The definition of View Synchronous Broadcast is more complex. It is defined in the context of a group g , and is based on the notion of *a sequence of views* $v_0(g), v_1(g), \dots, v_i(g), \dots$ of group g . Each view $v_i(g)$ defines the composition of the group at some time t , i.e. the members of the group that are perceived as being correct at time t . Whenever a process p in some view $v_i(g)$ is suspected to have crashed, or some process q wants to join, a new view $v_{i+1}(g)$ is installed, which reflects the membership change.

Roughly speaking, VSCAST of message m by some member of the group g currently in view $v_i(g)$ ensures the following property: if one process p in $v_i(g)$ delivers m before installing view $v_{i+1}(g)$, then no process installs view $v_{i+1}(g)$ before having first delivered m .

3.2. Active Replication

Active replication, also called the state machine approach [28], is a non-centralised replication technique. Its key concept is that all replicas receive and process the same sequence of client requests. Consistency is guaranteed by assuming that, when provided with the same input in the same order, replicas will produce the same output. This assumption implies that servers process requests in a *deterministic* way.

Clients do not contact one particular server, but address servers as a group. In order for servers to receive the same input in the same order, client requests can be propagated to servers using an Atomic Broadcast. Weaker communication primitives can also be used if semantic information about the operation is known (e.g., two requests that commute do not have to be delivered at all servers in the same order).

The main advantage of active replication is its simplicity (e.g., same code everywhere) and failure transparency.

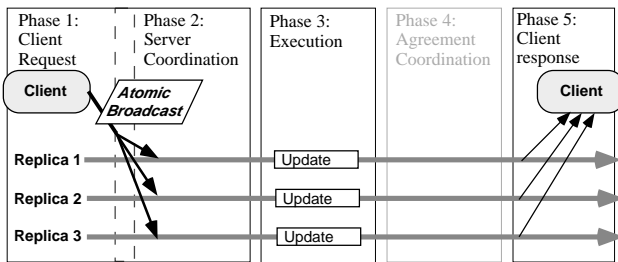


Figure 2. Active replication




Figure 3. Passive replication

Failures are fully hidden from the clients, since if a replica fails, the requests are still processed by the other replicas.

The determinism constraint is the major drawback of this approach. Although one might also argue that having all the processing done on all replicas consumes too many resources. Notice however, that the alternative, that is, processing a request at only one replica and transmitting the state changes to the others (see next section), in some cases may be much more complex and expensive than simply executing the invocation on all sites.

Figure 2 depicts the active replication technique using an Atomic Broadcast as communication primitive. In active replication, phases **RE** and **SC** are merged and phase **AC** is not used. The following steps are involved in the processing of an update request in the Active Replication, according to our functional model.

1. The client sends the request to the servers using an Atomic Broadcast.
2. Server coordination is given by the total order property of the Atomic Broadcast.
3. All replicas execute the request in the order they are delivered.
4. No coordination is necessary, as all replica process the same request in the same order. Because replica are deterministic, they all produce the same results.
5. All replica send back their result to the client, and the client typically only waits for the first answer.

3.3. Passive Replication

The basic principle of passive replication, also called *Primary Backup* replication, is that clients send their requests to a primary, which executes the requests and sends update messages to the backups (see Figure 3). The backups do not execute the invocation, but apply the changes produced by the invocation execution at the primary (i.e., updates). By doing this, no determinism constraint is necessary on the execution of invocations.

Communication between the primary and the backups has to guarantee that updates are received and then processed in the same order, which is the case if primary

backup communication is based on FIFO channels. However, FIFO channels are not enough to ensure correct execution in case of failure of the primary. For example, consider that the primary fails before all backups receive the updates for a certain request, and another replica takes over as a new primary. Some mechanism has to ensure that updates sent by the new primary will be “properly” ordered with regard to the updates sent by the faulty primary. VSCAST is a mechanism that guarantees these constraints, and can usually be used to implement the primary backup replication technique [13].

Passive replication can tolerate non-deterministic servers (e.g., multi-threaded servers) and uses little processing power when compared to other replication techniques. However, passive replication suffers from a high reconfiguration cost when the primary fails. The five steps of our framework are the following:

1. The client sends the request to the primary.
2. There is no initial coordination.
3. The primary executes the request.
4. The primary coordinates with the other replicas by sending the update information to the backups.
5. The primary sends the answer to the client.

3.4. Semi-Active Replication

Semi-active replication is an intermediate solution between active and passive replication. Semi-active replication does not require that replicas process service invocation in a deterministic manner. The protocol was originally proposed in a synchronous model [24]. We present it here in a more general system model.

The main difference between semi-active replication and active replication is that each time replicas have to make a non-deterministic decision, a process, called the *leader*, makes the choice and sends it to the *followers*. Figure 4 depicts Semi-active replication. Phases **EX** and **AC** are repeated for each non deterministic choice.

The following steps characterise semi-active replication, according to our framework.

1. The client sends the request to the servers using an Atomic Broadcast.

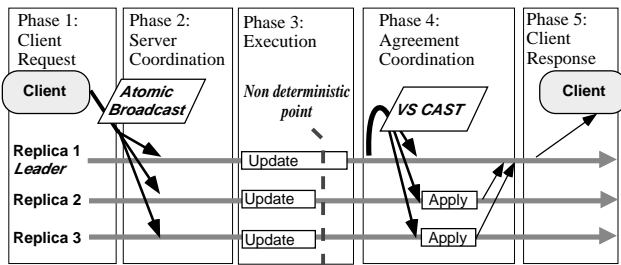


Figure 4. Semi-active replication

	Server Determinism Needed	Server Determinism Not Needed
Server Failure Not Transparent for the Client	Active	Passive
Server Failure Transparent for the Client	Active	semi-Active semi-Passive

Figure 5. Replication in distributed systems

2. The servers coordinate using the order given by this Atomic Broadcast.
3. All replicas execute the request in the order they are delivered.
4. In case of a non deterministic choice, the leader informs the followers using the View Synchronous Broadcast.
5. The servers send back the response to the client.

3.5. Semi-Passive Replication

Semi-passive replication [10] is a variant of passive replication which can be implemented in the asynchronous model without requiring any notion of views. The main advantage over passive replication is to allow for aggressive time-outs values and suspecting crashed processes without incurring too high a cost for incorrect failure suspicions. Because this technique has no equivalence in the context of database replication, we do not discuss it in detail. Roughly speaking, in semi-passive replication the Server Coordination (phase 2) and the Agreement Coordination (phase 4) are part of one single coordination protocol called *Consensus with Deferred Initial Values*.

3.6. Summary

Figure 5 summarises the different replication approaches in distributed systems, grouped according the following two dimensions: (1) failure transparency for clients, and (2) server determinism.

4. Database Replication

Replication in database systems is done mainly for performance reasons. The objective is to access data locally in

order to improve response times and eliminate the overhead of communicating with remote sites. This is usually possible when an operation only reads the data, while write operations require some form of coordination among the replicas. Fault tolerance is an issue but it is solved using backup mechanisms which, even being a form of replication, are entirely transparent to the clients.

4.1. Replication Model in Databases

A database is a collection of data items controlled by a database management system. A replicated database is thus a collection of databases that store copies of the same data items (for simplicity, we assume full replication). Hence, we distinguish a logical data item X and its physical copies X_i on the different sites. The basic unit of replication is the data item.

Clients access the data by submitting transactions. An operation, $o_i(X)$, of a transaction, T_i , can be either a read or a write access to a logical data item, X , in the database. Moreover, a transaction is a unit of work that executes atomically, i.e., a transaction either commits or aborts its results on all participating sites. Furthermore, if transactions run concurrently they must be isolated from each other if they conflict. Two operations of different transactions conflict if both access the same data item and one of them is a write. Isolation is provided by concurrency control mechanisms such as locking protocols [6] which guarantee serialisability. These protocols are extended to work in replicated scenarios and to provide 1-copy serialisability, the accepted correctness criterion for database replication [6].

A client submits its transactions to only one database and, in general, it is connected only to this database. If a database server fails, active transactions (not yet committed or aborted) running on that server are aborted. Clients can then be connected to another database server and re-submit the transaction. The failure is seen by the client but, in return, the client's logic is much simpler. From a practical point of view, in any working system, failures are the exception so it makes sense to optimise for the situation when failures do not occur.

In this section, we will use a very simple form of transaction that consists of a single operation. This allows us to concentrate on the coordination and interaction steps and makes it possible to directly compare with distributed system approaches. The next section will refine this model to extend it to normal transactions. Although the *single operation* approach may seem restrictive, it is actually used by many commercial systems in the form of *stored procedures*. A stored procedure resembles a procedure call and contains all the operations of one transaction. By invoking the stored procedure, the client invokes a transaction.

Note that the use of quorums is orthogonal to the follow-

update propagation	
Eager Primary Copy	Lazy Primary Copy
Eager Update Everywhere	Lazy Update Everywhere

Figure 6. Replication in database systems

ing discussion. Quorums only determine how many sites and which of them need to be contacted in order to execute an operation. Independently of which sites participate, the phases of the different protocols are the same. In an extreme case, read operations only access the local copy (read-one/write-all approach [6]), while write operations require coordination in any case.

4.2. Replication Strategies

Gray et. al [12] have categorised database replication protocols using two parameters (see Figure 6). One is when update propagation takes place (eager vs. lazy) and the second is who can perform updates (primary vs. update-everywhere). In eager replication schemes, updates are propagated within the boundaries of a transaction, i.e., the user does not receive the commit notification until sufficient copies in the system have been updated. Lazy schemes, on the other hand, update a local copy, commit and only some time after the commit, the propagation of the changes takes place. The first approach provides consistency in a straightforward way but it is expensive in terms of message overhead and response time. Lazy replication allows a wide variety of optimisations, however, since copies are allowed to diverge, inconsistencies might occur.

In regard to who is allowed to perform updates, the primary copy approach requires all updates to be performed first at one copy (the primary or master copy) and then at the other copies. This simplifies replica control at the price of introducing a single point of failure and a potential bottleneck. The update everywhere approach allows any copy to be updated, thereby speeding up access but at the price of making coordination more complex.

4.3. Eager Primary Copy Replication

In an eager primary copy approach, an update operation is first performed at a primary master copy and then propagated from this master copy to the secondary copies. When the primary has the confirmation that the secondary copies have performed the update, it commits and returns a notification to the user. Ordering of conflicting operations is determined by the primary site and must be obeyed by the secondary copies. Reading transactions can be performed




Figure 7. Eager primary copy

on any site and reading transactions will always see the latest version of each object. Early solutions, e.g., distributed INGRES [4, 30], used this approach. Currently, it is only used for fault-tolerance in order to implement a hot-standby backup mechanism where a primary site executed all operations and a secondary site is ready to immediately take over in case the primary fails [11, 2].²

Figure 7 shows the steps of the protocol in terms of the functional model as it would be used in a hot stand-by backup mechanism. The server coordination phase disappears since execution takes place only at the primary. The execution phase involves performing the transaction to generate the corresponding log records which are then sent to the secondary and applied. Then a 2PC protocol is executed during the agreement coordination phase. Once this finishes, a response is returned to the client.

From here, it is easy to see that eager primary copy replication is functionally equivalent to passive replication with VSCAST. The only differences are internal to the Agreement Coordination phase (2PC in the case of databases and VSCAST in the case of distributed systems). This difference can be explained by the use of transactions in databases. As explained, VSCAST is used to guarantee that operations are ordered correctly even after a failure occurs. In a database environment, the use of 2PC guarantees that if the primary fails, all active transactions will be aborted. Therefore, there is no need to order operations from “before the failure” and “after the failure” since there is only one source and the different views cannot overlap with each other.

4.4. Eager Update Everywhere Replication

From a functional point of view there are two types of protocols to consider depending on whether they use distributed locking or atomic broadcast to order conflicting operations.

Distributed Locking Approach When using distributed locking, a item can only be accessed after it has been locked

²Note that the primary is still a single point of failure, such an approach assumes that a human operator can reconfigure the system so that the backup is the new primary




Figure 8. Eager update everywhere with distributed locking

at all sites. For transactions with one operation, the replication control runs as follows (see Figure 8). The client sends the request to its local database server. This server sends a lock request to all other servers which grant or do not grant the lock. The lock request acts as the Server Coordination phase. If the lock is granted by all sites, we can proceed. If not, the transaction can be delayed and the request repeated some time afterwards. When all the locks are granted, the operation is executed at all sites. During the Agreement Coordination phase, a 2PC protocol is used to make sure that all sites commit the transaction. Afterwards, the client gets a response.

A comparison between Figures 4 and 8 shows that semi-active replication and eager update everywhere using distributed locking are conceptually similar. The differences arise from the mechanisms used during the Server Coordination and Agreement Coordination phases. In databases, Server Coordination takes place using 2 Phase Locking [6] while in distributed systems this is achieved using ABCAST. The 2 Phase Commit mechanism used in the Agreement Coordination phase of the database replication protocol corresponds to the use of a VSCAST mechanism in the distributed systems protocol. If the databases were deterministic, 2PC would not be needed and the protocol would be functionally identical to active replication (see Figure 8).

Data Replication based on Atomic Broadcast It has been suggested to use group communication primitives to implement database replication. However, it has not been until recently that the problem has been tackled with sufficient depth so as to provide realistic solutions [26, 17, 18]. The basic idea behind this approach is to use the total order guaranteed by ABCAST to provide a hint to the transaction manager on how to order conflicting operations. Thus, the client submits its request to one database server which then broadcasts the request to all other database servers. Instead of 2 Phase Locking, the server coordination is done based on the total order guaranteed by ABCAST and using some techniques to obtain the locks in a consistent manner at all sites [17, 19]. It must be guaranteed that two conflicting operations are executed in the order of the ABCAST at




Figure 9. Eager update everywhere based on atomic broadcast

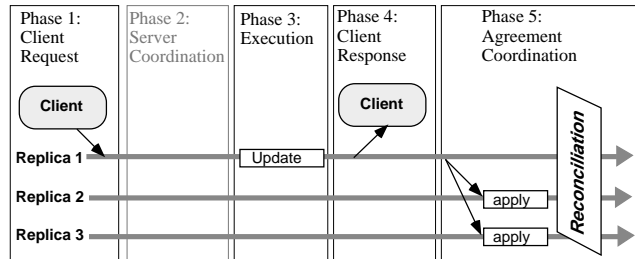


Figure 10. Lazy update everywhere

all sites. Once the local server has executed the operation it sends the response to the client (see Figure 2). The five phases are the following:

1. The client sends the request to the local server
2. The server forwards the request to all servers which coordinate using the Atomic Broadcast.
3. The servers execute the transaction.
4. There is no coordination at this point.
5. The local server sends back the response to the client.

The similarities between active replication and eager update everywhere using ABCAST are obvious when Figures 2 and 9 are compared. The only significant difference is the interaction between the client and the system. Regarding the determinism of the databases, a complete study of the requirements and the conditions under which ABCAST can be used for database replication and when an Agreement Coordination is necessary can be found in [17].

4.5. Lazy Replication

Lazy replication avoids the synchronisation overhead of eager replication by providing a response to the client before there is any coordination between servers. As in eager solutions there exist both primary copy and update everywhere approaches (see Figure 10 for lazy update everywhere). In the case of primary copy, all clients must contact the same server to perform updates while in update everywhere any server can be accessed. Directly after the execution of the transaction the local server sends the response




Figure 11. Eager primary copy approach for transactions

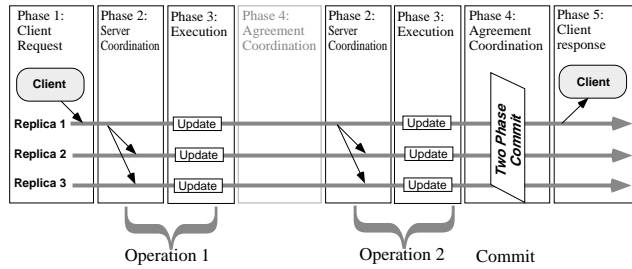


Figure 12. Eager update everywhere approach for transactions

back to the client. Only some time after the commit the updates are propagated to the other sites. This allows to bundle changes of different transactions and propagate updates on an interval basis to reduce communication overhead. In the case of primary copy the Agreement Coordination phase is relatively straightforward in that all ordering takes place at the primary and the replicas need only to apply the propagated changes. In the case of update everywhere, coordination is much more complicated. Since the other sites might have run conflicting transactions at the same time, the copies on the different sites might not only be stale but inconsistent. Reconciliation is needed to decide which updates are the winners and which transactions must be undone.

Note that the concept of laziness, while existing in distributed systems approaches [20], is not widely used. This reflects the fact that those solutions are mainly developed for fault-tolerant purposes, making an eager approach obligatory. Lazy approaches, on the other hand, are a straightforward solution if performance is the main issue.

5. Transactions

In many databases, transactions are not one single operation or are not executed via a stored procedure. Instead, transactions are a partial order of read and write operations which are not necessarily available for processing at the same time. This has important consequences for replica control resulting in protocols which have no equivalent in distributed systems.

The fact that now a transaction has multiple operations and that those operations need to be properly ordered with respect to each other requires to modify the functional model. The modification involves introducing a loop including the Server Coordination and Execution phases or the Execution and Agreement Coordination phases, depending on the protocol used. The loop will be executed once for each operation that needs to be performed.

5.1. Eager Primary Copy Replication

In the case of primary copy, there is no need for server coordination. Hence, the loop will involve the Execution and the Agreement Coordination phases. In this loop an operation is performed at the primary copy and then the changes sent to the replicas. This is done for every operation and, at the end, a new Agreement Coordination phase is executed in order to go through a 2PC protocol that will commit the transaction at all sites (see Figure 11).

Note that the Agreement Coordination phases for each operation and that the one at the end use different mechanisms. If we compare this with the algorithm in Section 4.4, we notice that the last phase is the same. For each operation except the last, it suffices to send the operation. In the final Agreement Coordination phase, a 2PC protocol is used to make sure all sites either commit or abort the transaction.

An alternative approach to this one is to use shadow copies and propagate the changes made by a transaction only after the transaction has completed (note that completed is not the same as committed). If this approach is used, the resulting protocol is identical to that shown in Figure 7.

5.2. Eager update everywhere replication

We will again look at the two different approaches used to implement eager update everywhere replication.

Distributed Locking In this case, a lock must be obtained for every operation in the transaction. This requires repeating the Server Coordination and Execution phases for every operation. At the end, once all operations have been processed in this way, a 2PC protocol is used during the Agreement Coordination phase to commit or abort the transaction at all sites (see Figure 12).

Certification Based Database Replication When using ABCAST to send the operations to all replicas, the resulting total order has no bearing on the serialisation order that needs to be produced. For this reason, it does not make

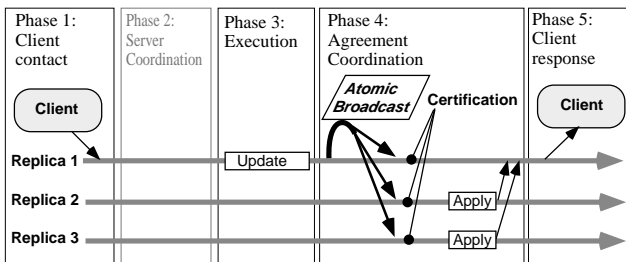


Figure 13. Certification based Database Replication

much sense to use ABCAST to send every operation of a transaction separately. It makes sense, however, to use shadow copies at one site to perform the operations and then, once the transaction is completed, send all the changes in one single message [17]. Due to the fact that now a transaction manager has to unbundle these messages, the agreement coordination phase becomes more complicated since it involves deciding whether the operations can be executed correctly. This can be seen as a *certification* step during which sites make sure they can execute transactions in the order specified by the total order established by ABCAST (see Figure 13).

5.3. Lazy Replication

When using lazy replication, updates are not propagated until the transaction commits. Then all the updates performed by the transaction are sent as a unit. Thus, whether transactions have one or more operations does not make a difference for lazy replication protocols.

6. Discussion

This paper presents a general comparison of replication approaches used in the distributed system and database communities. Our approach was to first characterise replication algorithms using a generic framework. Our generic framework identifies five basic steps, and, although simple, allows us to classify classical replication protocols described in the literature on distributed systems and databases.

Figure 14 summarises the different replication techniques. We see that any replication technique that ensures strong consistency has either an **SC** and/or **AC** step before the **END** step. All techniques have at least one synchronisation step (**SC** or **AC**). If the execution step (**EX**) is deterministic, no synchronisation after **EX** is needed, as the execution will yield the same result on all servers. For the same reason, if only one server does the execution step, there is no need for synchronisation before the execution.




Figure 14. Synthetic view of approaches

Several conclusion can be drawn from this figure. First, primary copy and passive replication schemes share one common feature: they do not have an **SC** phase (since the primary does the processing, there is no need for early synchronisation between replicas). Furthermore, update everywhere replication schemes need the initial **SC** phase before an update can be executed by the replicas. The only exceptions are the Certification based techniques that use Atomic Broadcast (Sect. 5.2). Those techniques are *optimistic* in the sense that they do the processing without initial synchronisation, and abort transactions in order to maintain consistency. Finally, the difference between eager and lazy replication techniques is the ordering of the **AC** and **END** phases: in the eager technique, the **AC** phase comes first, while in the lazy technique, the **END** phase comes first.

7. Conclusion

Despite different models, constraints and terminologies, replication algorithms for distributed systems and databases bear several similarities. These similarities put into evidence the need for stronger cooperation between both communities. For example, replicated databases could benefit from the abstractions of distributed systems. Presently, we are planning a performance study of the different approaches, taking into account different workload and failure assumptions.

References

- [1] D. Agrawal, G. Alonso, A. E. Abbadi, and I. Stanoi. Exploiting atomic broadcast in replicated databases. In *Proceedings of EuroPar (EuroPar'97)*, Passau (Germany), 1997.
- [2] G. Alonso, M. Kamath, D. Agrawal, A. El Abbadi, R. Günthör, and C. Mohan. Advanced transaction models in the workflow contexts. In *Proceedings of the International Conference on Data Engineering*, New Orleans, Feb. 1996.
- [3] B. Alpern and F. Schneider. Recognizing safety and liveness. *Distributed Computing*, 2:117–126, 1987.
- [4] P. Alsberg and J. Day. A principle for resilient sharing of distributed resources. In *Proceedings of the International Conference on Software Engineering*, Oct. 1976.
- [5] H. Attiya and J. Welch. Sequential consistency versus linearizability. *ACM Transactions on Computer Systems*, 12(2):91–122, May 1994.
- [6] P. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987.
- [7] K. P. Birman. The process group approach to reliable distributed computing. *Communications of the ACM*, 36(12):37–53, Dec. 1993.
- [8] K. P. Birman and T. A. Joseph. Exploiting virtual synchrony in distributed systems. In *Proceedings of the 11th ACM Symposium on OS Principles*, pages 123–138, Austin, TX, USA, Nov. 1987. ACM SIGOPS, ACM.
- [9] K. P. Birman, A. Schiper, and P. Stephenson. Lightweight causal and atomic group multicast. *ACM Transactions on Computer Systems*, 9(3):272–314, Aug. 1991.
- [10] X. Défago, A. Schiper, and N. Sergent. Semi-passive replication. In *Proceedings of the 17th IEEE Symposium on Reliable Distributed Systems (SRDS)*, pages 43–50, West Lafayette, IN, USA, Oct. 1998.
- [11] J. Gray and A. Reuter. *Transaction Processing: concepts and techniques*. Data Management Systems. Morgan Kaufmann Publishers, Inc., San Mateo (CA), USA, 1993.
- [12] J. N. Gray, P. Helland, and a. D. S. P. O'Neil. The dangers of replication and a solution. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data*, pages 173–182, Montreal, Canada, June 1996. SIGMOD.
- [13] R. Guerraoui and A. Schiper. Software-based replication for fault tolerance. *IEEE Computer*, 30(4):68–74, Apr. 1997.
- [14] V. Hadzilacos and S. Toueg. Fault-tolerant broadcasts and related problems. In S. Mullender, editor, *Distributed Systems*, chapter 5. adwe, second edition, 1993.
- [15] J. Holliday, D. Agrawal, and A. E. Abbadi. The performance of database replication with group multicast. In *Proceedings of IEEE International Symposium on Fault Tolerant Computing (FTCS29)*, pages 158–165, 1999.
- [16] Information & Communications Systems Research Group, ETH Zürich and Laboratoire de Systèmes d'Exploitation (LSE), EPF Lausanne. *DRAGON: Database Replication Based on Group Communication*, May 1998. <http://www.inf.ethz.ch/departement/IS/iks/research/dragon.html>.
- [17] B. Kemme and G. Alonso. A suite of database replication protocols based on group communication primitives. In *Proceedings of the 18th International Conference on Distributed Computing Systems (ICDCS)*, Amsterdam, The Netherlands, May 1998.
- [18] B. Kemme, F. Pedone, G. Alonso, and A. Schiper. Processing transactions over optimistic atomic broadcast protocols. In *Proceedings of the International Conference on Distributed Computing Systems*, Austin, Texas, June 1999. to appear.
- [19] B. Kemme, F. Pedone, G. Alonso, and A. Schiper. Using optimistic atomic broadcast in transaction processing systems. Technical report, Department of Computer Science, ETH Zürich, Mar. 1999.
- [20] R. Ladin, B. Liskov, and S. Ghemawat. Providing high availability using lazy replication. *ACM Trans. Comput. Syst.*, 10(4):360–391, November 1992.
- [21] E. Pacitti, P. Minet, and E. Simon. Fast algorithms for maintaining replica consistency in lazy master replicated databases. In *Proceedings of the 25th International Conference on Very Large Databases*, Edinburgh - Scotland - UK, 7–10 Sept. 1999.
- [22] F. Pedone, R. Guerraoui, and A. Schiper. Transaction reordering in replicated databases. In *Proceedings of the 16th Symposium on Reliable Distributed Systems (SRDS-16)*, Durham, North Carolina, USA, Oct. 1997.
- [23] F. Pedone, R. Guerraoui, and A. Schiper. Exploiting atomic broadcast in replicated databases. In *Proceedings of EuroPar (EuroPar'98)*, Sept. 1998.
- [24] D. Powell, M. Chéréque, and D. Drackley. Fault-tolerance in Delta-4*. *ACM Operating Systems Review, SIGOPS*, 25(2):122–125, Apr. 1991.
- [25] M. Raynal, G. Thia-Kime, and M. Ahamad. From serializable to causal transactions for collaborative applications. Technical Report 983, Institut de Recherche en Informatique et Systèmes Aléatoires, Feb. 1996.
- [26] A. Schiper and M. Raynal. From group communication to transactions in distributed systems. *Communications of the ACM*, 39(4):84–87, Apr. 1996.
- [27] A. Schiper and A. Sandoz. Uniform reliable multicast in a virtually synchronous environment. In *Proceedings of the 13th International Conference on Distributed Computing Systems (ICDCS-13)*, pages 561–568, Pittsburgh, Pennsylvania, USA, May 1993. IEEE Computer Society Press.
- [28] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, Dec. 1990.
- [29] I. Stanoi, D. Agrawal, and A. E. Abbadi. Using broadcast primitives in replicated databases. In *Proceedings of the International Conference on Distributed Computing Systems ICDCS'98*, pages 148–155, Amsterdam, The Netherlands, May 1998.
- [30] M. Stonebraker. Concurrency control and consistency of multiple copies of data in distributed INGRES. *IEEE Transactions on Software Engineering*, SE-5:188–194, May 1979.