

Byzantine Fault-Tolerant Deferred Update Replication

Fernando Pedone
University of Lugano (USI)
Switzerland
fernando.pedone@usi.ch

Nicolas Schiper
University of Lugano (USI)
Switzerland
nicolas.schiper@usi.ch

Abstract—Replication is a well-established approach to increasing database availability. Many database replication protocols have been proposed for the crash-stop failure model, in which servers fail silently. Fewer database replication protocols have been proposed for the byzantine failure model, in which servers may fail arbitrarily. This paper considers deferred update replication, a popular database replication technique, under byzantine failures. The paper makes three contributions. First, it shows that making deferred update replication tolerate byzantine failures is quite simple. Second, the paper presents a byzantine-tolerant mechanism to execute read-only transactions at a single server. Third, we consider byzantine client attacks against deferred update replication and discuss effective countermeasures against these attacks.

Keywords—Database replication, byzantine fault-tolerance, dependable systems.

I. INTRODUCTION

Replication is a well-established approach to increasing database availability. By replicating data items in multiple servers, the failure of some servers does not prevent clients from executing transactions against the system. Database replication in the context of crash-stop failures has been largely studied in the past years (e.g., [4], [11], [14], [19], [20]). When a crash-stop server fails, it silently stops its execution. More recently, a few works have considered database replication under byzantine failures (e.g., [21], [24]). Byzantine failures are more severe than crash-stop failures since failed servers can present arbitrary behavior.

Several protocols for the crash-stop failure model are based on deferred update replication. According to deferred update replication, to execute a transaction, a client first picks a server and submits to this server its transaction commands. The execution of a transaction does not cause any communication among servers until after the client requests the transaction’s commit, at which point the transaction enters the termination phase and is propagated to all servers. As part of termination, each server certifies the transaction and commits it, if doing so induces a *serializable* execution, i.e., one in which transactions appear to have been executed in some serial order.

Deferred update replication scales better than state-machine replication and primary-backup replication. With

state-machine replication, every update transaction must be executed by all servers. Thus, adding servers does not increase the throughput of update transactions. With primary-backup replication, the primary first executes update transactions and then propagates the database changes to the backups, which apply them without re-executing the transactions. The throughput of update transactions is limited by the capacity of the primary, not by the number of replicas. Deferred update replication scales better because it allows all servers to act as “primaries”, locally executing transactions and then propagating the modifications to the other servers. As applying transaction modifications to the database is usually cheaper than executing transactions, the technique provides better throughput and scalability.

Ensuring strong consistency despite multiple co-existing primaries requires servers to synchronize. This is typically done by means of an atomic broadcast protocol to order transactions and a certification test to ensure the consistency criterion of interest. One of the key properties of deferred update replication is that read-only transactions can be executed by a single server, without communication across servers. This property has two implications. First, in geographically distributed networks it can substantially reduce the latency of read-only transactions. Second, it enables read-only transactions to scale perfectly with the number of servers in the system.

This paper considers deferred update replication under byzantine failures. It proposes the first byzantine fault-tolerant deferred update replication protocol that is faithful to its crash-stop counterpart: (i) the execution of a transaction does not require communication across servers, only its termination does, and (ii) only one server executes the transaction commands, but all correct servers apply the updates of a committing transaction. Our protocol is surprisingly simple and similar to a typical crash-stop deferred update replication protocol, although based on a more strict certification procedure to guarantee that transactions only commit if they do not violate consistency and read valid data (i.e., data that was not fabricated by a byzantine server).

Our most significant result is a mechanism to execute read-only transactions at a single server under byzantine failures. Some protocols in the crash-stop model achieve this property by carefully scheduling transactions so that they

observe a consistent database view. In the byzantine failure model, however, clients may inadvertently execute a read-only transaction against a byzantine server that fabricates a bogus database view. In brief, our solution to the problem consists in providing enough information for clients to efficiently tell whether the data items read form a valid and consistent view of the database. Clients are still subject to malicious servers executing read-only transactions against old, but consistent, database views. We briefly discuss in the paper the extent of the problem and remedies to such attacks.

Finally, we also consider byzantine client attacks against deferred update replication. Since concurrency control is optimistic in deferred update replication, byzantine clients can launch attacks against honest clients by trying to maximize the abort rate of the transactions submitted by the latter. We discuss different such attacks, present countermeasures to them, and evaluate through simulation byzantine client attacks and our countermeasures.

The remainder of this paper is organized as follows. Section II describes the system model. Sections III and IV discuss deferred update replication in the crash-stop failure model and in the byzantine failure models respectively. Section V considers byzantine clients and Section VI assesses different byzantine client attacks and the effectiveness of our countermeasures. Section VII discusses related work. Section VIII concludes the paper.

II. SYSTEM MODEL AND DEFINITIONS

In this section, we detail the system model and assumptions common to both the crash-stop failure model and the byzantine failure model. Further assumptions, specific to each model, are detailed in Sections III and IV.

A. Clients, servers and communication

Let $C = \{c_1, c_2, \dots\}$ be the set of *client* processes and $S = \{s_1, s_2, \dots, s_n\}$ the set of *server* processes. Processes are either *correct*, if they follow their specification and never fail, or *faulty*, otherwise. We distinguish two classes of faulty processes: *crash-stop* and *byzantine*. Crash-stop processes eventually stop their execution but never misbehave; byzantine processes may present arbitrary behavior.

We study deferred update replication in two models: in one model faulty servers are *crash-stop*; in the other model faulty servers are *byzantine*. In either case, there are at most f faulty servers and an unbounded number of clients. We initially consider crash-stop clients (Sections III and IV) and then byzantine clients (Sections V and VI).

Processes communicate by message passing, using either one-to-one or one-to-many communication. One-to-one communication is through primitives $\text{send}(m)$ and $\text{receive}(m)$, where m is a message. If sender and receiver are correct, then every message sent is eventually received. One-to-many communication is based on atomic broadcast,

through the primitives $\text{abcast}(m)$ and $\text{deliver}(m)$, and used by clients to propagate messages to the group of servers.

In the crash-stop model, atomic broadcast ensures that (i) if one server delivers a broadcast message, then all correct servers also deliver the message; and (ii) no two servers deliver any two messages in different orders. In the byzantine model, atomic broadcast ensures that (i) if one correct server delivers a broadcast message, then all correct servers also deliver the message; and (ii) no two correct servers deliver any two messages in different orders.

B. Transactions and serializability

Let $X = \{x_1, x_2, \dots\}$ be the set of data items, i.e., the *database*, $\text{Cmd} = \{\text{commit}, \text{abort}\} \cup (\{r, w\} \times X \times V)$ the set of *commands*, where V is the set of possible values of a data item, and $S = C \times \text{Cmd}$ the set of *statements*. Statement $(c, (r, x, v))$ means that client c has read item x with value v ; statement $(c, (w, x, v))$ means that c has modified the state of x to v .

We define a *history* h as a finite sequence of statements in S . We define the projection $h|_c$ of history h on $c \in C$ as the longest subsequence h' of h such that every statement in h' is in $c \times \text{Cmd}$. In a projection $h|_c = \sigma_0 \dots \sigma_m$, statement σ_i is *finishing* in $h|_c$ if it is a **commit** or an **abort**; σ_i is *initiating* if it is the first statement in $h|_c$ or the previous statement σ_{i-1} is a finishing statement.

A sequence of commands $t = \sigma_0 \dots \sigma_m$ in $h|_c$ is a *transaction* issued by c if (i) σ_0 is initiating in $h|_c$ and (ii) σ_m is either finishing in $h|_c$ or it is the last statement in $h|_c$. Transaction t is *committing* if σ_m is a **commit** statement. We denote as $\text{com}(h)$ the longest subsequence h' of h such that every statement in h' is part of a committing transaction in h . In other words, $\text{com}(h)$ is the committed projection of h , with all statements of all committed transactions in h .

Let t and u be transactions in a history h . We say that t precedes u in h , $t <_h u$, if the finishing statement of t occurs before the initiating statement of u in h . A history h is *serial* if for every pair (t, u) of transactions in h , either $t <_h u$ or $u <_h t$. History h is *legal* if in h (i) every read statement $(c_i, (r, x, v_j))$ is preceded by a write statement $(c_j, (w, x, v_j))$ and (ii) in between the two there is no statement $(c_k, (w, x, v_k))$, $v_j \neq v_k$.

History h is *serializable* if there is a serial permutation h' of $\text{com}(h)$ such that for each data item x , $h'|_x$ is legal, where $h'|_x$ is the projection of h' on x . *Serializability* is the set of all serializable histories.

III. DEFERRED UPDATE REPLICATION

In this section, we review deferred update replication in the crash-stop failure model. In this model, some atomic broadcast protocols require a majority of correct processes [12]. Thus, we assume the existence of $2f + 1$ correct servers in the system. The database is fully replicated, that is, every server has a complete copy of the database.

A. Overview

In deferred update replication, transactions pass through two phases in their lifetime: the *execution phase* and the *termination phase*. The execution phase starts when the client issues the first transaction command; it finishes with a client's request to commit or abort the transaction, when the termination phase starts. The termination phase finishes when the transaction is committed or aborted.

Before starting a transaction t , a client c must select the server s that will receive and execute t 's commands; other servers will not be involved in t 's execution. Each data item in the server's database is a tuple (x, v, i) , where x is the item's unique identifier, v is x 's value and i is the value's version. We assume that read and write commands on database tuples are atomic operations. When s receives a read command for x from c , it returns the current value of x (or the most up-to-date value if the database is multiversion) and its corresponding version. Write commands are locally stored by c . It is only during transaction termination that updates are propagated to the servers.

In the termination phase, the client atomically broadcasts t 's *readset* and *writeset*, denoted respectively by $t.rs$ and $t.ws$ —for simplicity, we say that “ c broadcasts t ”. The readset of t is the set of all tuples (x, i) , where x is a data item read by t and i is the version of the value read; the writeset of t is the set of all tuples (x, v) , where x is a data item written by t and v is x 's new value. Notice that the readset does not contain the values read.

Upon delivering t 's termination request, s *certifies* t . Certification ensures a serializable execution; it essentially checks whether t 's read commands have seen values that are still up-to-date when t is certified. If t passes certification, then s executes t 's writes against the database and assigns each new value the same version number k , reflecting the fact that t is the k -th committed transaction at s .

To certify t , s maintains a set CT of tuples (i, up) , where up is a set with the data items written by the i -th committed transaction at s . We state the certification test of t , $C_{cs}(t.rs, CT)$, more formally with the predicate below.

$$C_{cs}(t.rs, CT) \equiv \forall (x, i) \in t.rs : \nexists (j, up) \in CT \text{ s.t. } (x \in up) \text{ and } (j > i) \quad (1)$$

If t passes certification, then s updates the database and CT . Certifying a transaction and creating new database items is an atomic operation. When a new version of x is created, the server can decide to keep older versions of x or not. If multiple versions of a data item exist, then we say the database is *multi-version*; if not it is *single-version*.

B. Algorithm in detail

Algorithms 1 and 2 are high level descriptions of the client's and server's protocol. Notice that to determine the

outcome of a commit request, the client waits for a reply from a single server. For brevity, we do not show in Algorithm 1 the case in which the client decides to abort a transaction. In Algorithm 2, $items(t.ws)$ returns the set of data items in $t.ws$, without the values written by t , that is, $items(t.ws) = \{x \mid (x, v) \in t.ws\}$.

1: Algorithm 1: Deferred update replication, client c 's code:

```

2: Command  $(c, (r, x, v))$  occurs as follows:
3:   if first command for  $t$  then
4:      $t.rs \leftarrow \emptyset$ 
5:      $t.ws \leftarrow \emptyset$ 
6:     choose some server  $s \in S$ 
7:   if  $(x, v) \notin t.ws$  then
8:     send  $(c, (r, x))$  to  $s$ 
9:     wait until receive  $(c, (x, v, i))$  from  $s$ 
10:     $t.rs \leftarrow t.rs \cup \{(x, i)\}$ 
11:    return  $v$ 
12:   else
13:     return  $v$  such that  $(x, v) \in t.ws$ 
14: Command  $(c, (w, x, v))$  occurs as follows:
15:    $t.ws \leftarrow t.ws \cup \{(x, v)\}$ 
16: A commit command is executed as follows:
17:   if  $t.ws \neq \emptyset$  then
18:     abcast( $c$ , commit-req,  $t.rs, t.ws$ ) to all servers
19:     wait until receive  $(c, outcome)$  from  $s$ 

```

1: Algorithm 2: Deferred update replication, server s 's code:

```

2: Initialization
3:    $CT \leftarrow \emptyset$ 
4:    $lastCommitted \leftarrow 0$ 
5: upon receiving  $(c, (r, x))$  from  $c$ 
6:   retrieve  $(x, v, i)$  from database
7:   send  $(c, (x, v, i))$  to  $c$ 
8: upon delivering  $(c, \text{commit-req}, t.rs, t.ws)$ 
9:   if  $C_{cs}(t.rs, CT)$  then
10:     $lastCommitted \leftarrow lastCommitted + 1$ 
11:     $CT \leftarrow CT \cup \{(lastCommitted, items(t.ws))\}$ 
12:    for each  $(x, v) \in t.ws$  do
13:      create database entry  $(x, v, lastCommitted)$ 
14:      send  $(c, \text{commit})$  to  $c$ 
15:   else
16:     send  $(c, \text{abort})$  to  $c$ 

```

C. Read-only transactions

We describe two mechanisms to allow read-only transactions to be executed by a single server only. One mechanism is based on multi-version databases and does not require updates from committing transactions to be synchronized with on-going read-only transactions at a server; the other mechanism assumes a single-version database but synchronizes the updates of committing transactions with read-only transactions at servers.

With multi-version databases, each server stores multiple versions of each data item (limited by a system parameter).

When a transaction t issues its first read command, the client takes the version of the value returned as a reference for future read commands. This version number specifies the *view* of the database that the transaction will see, $t.view$. Every future read command must contain $t.view$. Upon receiving a read command with $t.view$, the server returns the most recent value of the item read whose version is equal to or smaller than $t.view$. If no such a value is available, the server tells the client that t must be aborted. This technique is sometimes called *multiversion timestamps* [10].

If a single version of each data item exists, then read commands must be synchronized (e.g., through two-phase locking [2]) with the updates of committing transactions. During the execution of a transaction t , each read command of t must first acquire a read lock on the data item. If a transaction u passes certification, then the server must acquire write locks on all data items in u 's writeset and then commit u . Since read and write locks cannot be acquired simultaneously, this technique may block transactions. This mechanism has been used by some protocols based on the deferred update replication model (e.g., [18]).

D. Correctness

To reason about correctness, we must define when read, write and commit commands of a transaction t take place. For read and write commands this is simple: a read happens when the client receives the corresponding reply from the server that executes t ; a write happens after the client updates t 's writeset. It seems natural to assume that the commit of t also happens when the client receives the first commit reply from a server. In our model, however, clients are allowed to crash and so, the commit may never take place, despite the fact that some databases consider the transaction committed. To avoid such cases, without assuming correct clients, we define the commit event of t as taking place when the first server applies t 's updates to its database.

We initially argue that Algorithm 1 is correct for update transactions only. We then extend our argument to include read-only transactions.

1) *Update transactions*: Let h_0 be a history created by Algorithm 1. We must show that we can permute the statements of committed transactions in h_0 such that the resulting history, h_s , is serial and legal. Our strategy is to create a series of histories $h_0, \dots, h_i, \dots, h_s$ where h_i is created after swapping two statements in h_{i-1} .

Let t and u be two committing transactions in h_i such that t commits before u , which implies that t 's commit request was delivered before u 's. We show that all of t 's statements that succeed u 's statements in h_i can be placed before u 's statements. There are two cases to consider.

- 1) Some u 's statement σ_u precedes t 's read statement $(c_t, (r, x, v))$. If σ_u is a read on x , or a read or a write on an item different than x , then we can trivially swap the two. Assume now that σ_u is a write on x . Since

t is delivered before u and only values of a delivered transaction can be read by other transactions, we know that t 's read statement did not read the value written by σ_u , and thus, the two can be swapped.

- 2) Some u 's statement σ_u precedes t 's write statement $(c_t, (w, x, v))$. If σ_u is a statement on an item different than x , then the two can be obviously swapped. Assume then that σ_u is a statement on x . If σ_u is a write, then we can swap the two as no read statement from other transactions have seen them; they only take effect after a transaction is delivered and t is delivered before u . Finally, let σ_u be a read. We show by way of contradiction that $(c_u, (r, x, v_u))$ cannot precede $(c_t, (w, x, v_t))$. Since u is certified after t and u passes certification, from the certification test, it must be that version read by u is still up-to-date. But since t modifies x , it updates x 's version, a contradiction that concludes our argument.

2) *Read-only transactions*: We consider first read-only transactions in the presence of multi-version databases. Let t be a read-only transaction in some history h_i of the system. We extend the argument presented in the previous section to show that all of t 's read commands can be placed in between two update transactions, namely, after the transaction u that created the first item read by t and before any other update transaction. From Section III-C, every future read of t will return a version that is equal to or precedes $t.view$. Therefore, every read command issued by t can be placed before any command that succeeds u 's write commands.

We claim now that read-only transactions are also serializable with single-version databases. In this case, the local execution at a server follows the *two-phase locking* protocol, which is serializable [2]. Notice that although read and write commands are synchronized in a server, certification of update transactions is still needed since two update transactions executing on different servers may see inconsistent reads. For example, transaction t may read data item x and write y , while transaction u , executing at a different server, may read y and write x . In such a case, certification will abort one of the transactions.

IV. BFT DEFERRED UPDATE REPLICATION

To adapt our protocol to the byzantine failure model, we make a few extra assumptions. We first assume that the number of servers is at least $3f + 1$. This is the minimum to solve atomic broadcast with malicious faults [13]. We make use of message digests produced by collision-resistant hash functions to ensure data integrity [22], and public-key signatures to ensure communication authenticity [23]. We follow the common practice of signing message digests instead of signing messages directly. We also assume that each client-server pair and each pair of servers communicate using private channels that can be implemented using symmetric

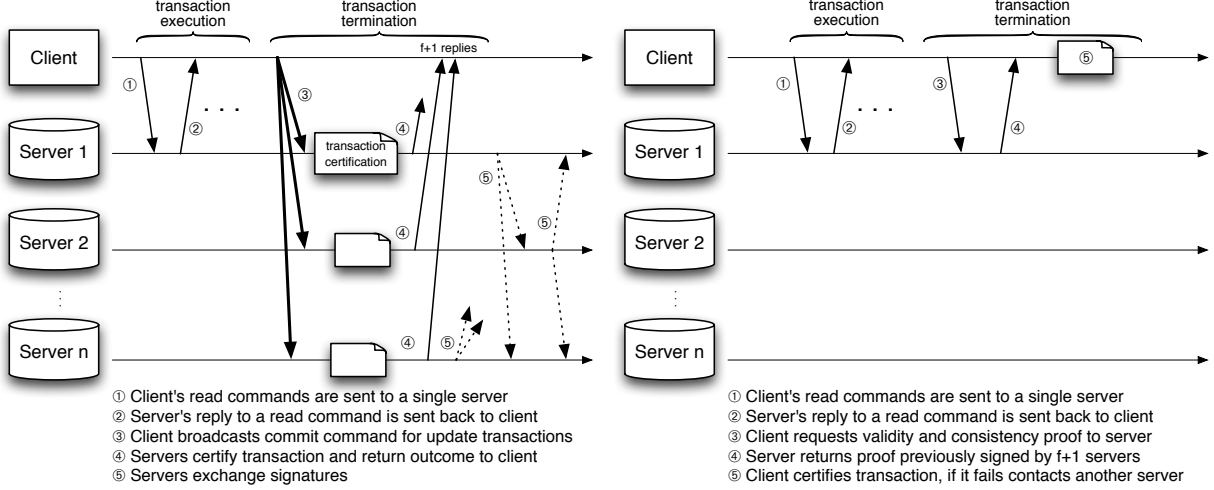


Figure 1. BFT deferred update replication: update (left) and read-only transactions (right)

keys [5]. Finally, clients are authenticated and servers enforce access control. This forbids unauthorized clients from accessing the database, and prevents potentially byzantine servers from issuing transactions that may compromise the integrity of the database. Obviously, a byzantine server can compromise its local copy of the database, but this behaviour is handled by our protocols.

A. Overview

A byzantine server could easily corrupt the execution of the algorithm presented in the previous section. For example, it could hamper the execution of the atomic broadcast protocol or answer client commands with incorrect data. While the first problem can be solved by replacing the atomic broadcast algorithm for crash-stop failures with one that tolerates byzantine failures [17], the second problem is less obvious to address.

A byzantine server may return two types of “incorrect data”: *invalid* or *stale*. Invalid data, as opposed to *valid* data, is fabricated by the server and does not correspond to any value created by a committed transaction. Stale data is too old, although it may be valid. We address the problem of stale data with the certification test by checking that the values read are still up-to-date. To guarantee that transactions read valid data, each database tuple is redefined as (x, v, i, d) , where in addition to x 's value v and version i , it contains a digest d of v . A read command returns the value read and its digest, and readsets include the digest of values.

We decompose the certification test into a component that checks the validity of data read, $C_b^v(t.rs, CT)$ ¹, and another component that checks whether the items read are up to date,

¹The original version of this manuscript contained an error in $C_b^v(t.rs, CT)$. We replaced variable j with i in the second line of this test to ensure that the correct value of x was read.

$C_b^u(t.rs, CT)$.² Moreover, tuples (i, up) in CT contain in set up elements (x, d) , that is, the items written and the digest of the values written.

$$C_b^u(t.rs, CT) \equiv \forall (x, i, d) \in t.rs : \quad (2)$$

$$\nexists (j, up) \in CT \text{ s.t. } ((x, *) \in up) \text{ and } (j > i)$$

$$C_b^v(t.rs, CT) \equiv \forall (x, i, d) \in t.rs : \quad (3)$$

$$\exists (i, up) \in CT \text{ s.t. } ((x, d') \in up) \text{ and } (d = d')$$

Additionally, to shield clients from byzantine servers that would commit a transaction t that violates serializability, clients wait for more than one server's reply before concluding the outcome of transactions. It turns out that at commit time, by waiting for a set of $f + 1$ identical replies, t 's outcome can be safely determined by the clients since only up to f servers can be compromised.

The left side of Figure 1 illustrates the protocol for update transactions. Note that step 5 in the illustration will be explained as part of the protocol for read-only transactions (cf. Section IV-C).

B. Algorithm in detail

Algorithms 3 and 4 present the client's and server's protocol. To execute a read command, a client c contacts a server s and stores the version, the value, and the digest in $t.rs$. Write operations are buffered in $t.ws$ as before. When c wishes to commit t , c atomically broadcasts a signed message to all servers. This is denoted by $\langle m \rangle_{\sigma_c}$, where m is a message and σ_c is c 's signature. Signing messages guarantees that only authenticated clients issue

²The decomposition of the certification test into two components is done for clarity purposes. An implementation would probably combine the two to speed up the execution.

commit requests. After receiving $f+1$ identical replies from servers, c can determine t 's outcome.

Besides a change in the certification test and in the data sent back to the client when answering read requests, the server's code is similar to the crash-stop case. In Algorithm 4, we must instantiate $items(t.ws)$ (see line 11) as $items(t.ws) = \{(x, v's \text{ digest}) \mid (x, v) \in t.ws\}$.

1: **Algorithm 3: Deferred update replication**, client c 's code:

2: Command $(c, (r, x, v))$ occurs as follows:

3: **if** first command for t **then**

4: $t.rs \leftarrow \emptyset$

5: $t.ws \leftarrow \emptyset$

6: choose some server $s \in S$

7: **if** $(x, v) \notin t.ws$ **then**

8: send $(c, (r, x))$ to s

9: wait until receive $(c, (x, v, i, d))$ from s

10: $t.rs \leftarrow t.rs \cup \{(x, i, d)\}$

11: return v

12: **else**

13: return v such that $(x, v) \in t.ws$

14: Command $(c, (w, x, v))$ occurs as follows:

15: $t.ws \leftarrow t.ws \cup \{(x, v)\}$

16: A commit command is executed as follows:

17: **if** $t.ws \neq \emptyset$ **then**

18: abcast $\langle c, \text{commit-req}, t.rs, t.ws \rangle_{\sigma_c}$ to all servers

19: wait for identical $(c, \text{outcome})$ from $f+1$ servers

x : the data item's key v : the data item's value
 i : the data item's version d : the data item's digest

1: **Algorithm 4: Deferred update replication**, server s 's code:

2: Initialization

3: $CT \leftarrow \emptyset$

4: $lastCommitted \leftarrow 0$

5: **upon** receiving $(c, (r, x))$ from c

6: retrieve (x, v, i, d) from database

7: send $(c, (x, v, i, d))$ to c

8: **upon** delivering $\langle c, \text{commit-req}, t.rs, t.ws \rangle_{\sigma_c}$

9: **if** $C_b^u(t.rs, CT)$ and $C_b^v(t.rs, CT)$ **then**

10: $lastCommitted \leftarrow lastCommitted + 1$

11: $CT \leftarrow CT \cup \{(lastCommitted, items(t.ws))\}$

12: **for each** $(x, v) \in t.ws$ **do**

13: create database entry...

14: ... $(x, v, lastCommitted, v's \text{ digest})$

15: send (c, commit) to c

16: **else**

17: send (c, abort) to c

x : the data item's key v : the data item's value
 i : the data item's version d : the data item's digest

C. Read-Only Transactions

A simple way to handle read-only transactions is to execute and terminate them in the same way as update transactions. This leads to a simple solution but increases the latency of read-only transactions since they need to

be atomically broadcast and certified by servers. In the following, we describe a mechanism that allows read-only transactions to be executed locally to a server only, just like deferred update replication in the crash-stop failure model.

To allow read-only transactions to execute at a single server, without inter-server communication, clients must be able to tell unilaterally whether (i) a value returned by a server as a response to a read command is valid (cf. Section IV-A) and (ii) any set of valid values read by the client belongs to a consistent view of the database. If the client determines that a value returned by the server is invalid or inconsistent, it aborts the transaction and retries using another server.

A set of values read by a client is a *consistent view* of the database if the values could be the result of a serial execution of the committed transactions. For example, assume that transactions t and u modify the values of data items x and y . Any transaction that reads x and y must see either the values created by t or the ones created by u or none, but not a mix of the two (e.g., x from t and y from u).

We ensure proper transaction execution by letting the client ask the server, at the end of the transaction execution, for a proof that the values read are valid and consistent. A *validity and consistency proof* for a transaction t , denoted as $vcp(t)$, consists of all elements of CT whose version lies between the lowest and highest data item version read by t . Moreover, to ensure that byzantine servers do not fabricate data, every element (i, up) in $vcp(t)$ must be signed by $f+1$ servers, denoted $\langle i, up \rangle_{\Sigma_{f+1}}$. Given a validity and consistency proof $vcp(t)$, the client decides to commit t if the following conditions hold:

- 1) The proof $vcp(t)$ is valid: If i_{min} and i_{max} are, respectively, the minimum and maximum data item versions read by t , then $vcp(t)$ contains all tuples (i, up) such that $i_{min} \leq i \leq i_{max}$. Moreover, each element of $vcp(t)$ is signed by $f+1$ servers, which guarantees that at least one correct server abides by this element.
- 2) The values read by t are valid: Each data item with version i read by t matches its corresponding digest in $vcp(t)$ with version i .
- 3) The values read by t are consistent: For each item x read with version i , no newer version $i' > i$ of x exists in $vcp(t)$.

Conditions 2 and 3 can be stated more precisely with predicates (4) and (5), respectively. Note that these predicates are similar to those used to certify update transactions (cf. Section IV-A). The main difference is that read-only transactions do not need to be certified against elements in CT whose version is newer than the highest data item version t read.

$$C_r^c(t.rs, vcp(t)) \equiv \forall (x, i, d) \in t.rs : \quad (4)$$

$$\nexists \langle j, up \rangle_{\Sigma_{f+1}} \in vcp(t) \text{ s.t. } ((x, *) \in up) \text{ and } (j > i)$$

$$C_r^v(t.rs, vcp(t)) \equiv \forall (x, i, d) \in t.rs : \quad (5)$$

$$\exists \langle i, up \rangle_{\Sigma_{f+1}} \in vcp(t) \text{ s.t. } ((x, d') \in up) \text{ and } (d = d')$$

To build a validity and consistency proof from CT , we add to each CT entry a certificate of $f+1$ server signatures. Database servers build certificates asynchronously. When the i -th update transaction t commits on s , the value, version, and digest of each data item x written by t are updated. Periodically, server s signs new tuples (i, up) and sends this information to all servers. When s gathers $f+1$ signatures of the tuple (i, up) , s inserts this new element in CT . This asynchronous scheme does not add communication overhead to update transactions. However, a read-only transaction t may stall until the server answering t 's requests gathers enough signatures to provide a validity and consistency proof for t . The protocol for read-only transactions is illustrated on the right side of Figure 1.

Algorithms 5 and 6 present the client and server protocols to execute read-only transactions. In the algorithms, given a transaction t and a validity and consistency proof $vcp(t)$ for t , read validity and consistency are expressed by predicates $C_r^v(t.rs, vcp(t))$ and $C_r^c(t.rs, vcp(t))$ respectively.

1: **Algorithm 5: Read-only transactions**, client c 's code:

```

2: Command  $(c, (r, x, v))$  occurs as follows:
3:   if first command for  $t$  then
4:      $t.rs \leftarrow \emptyset$ 
5:     choose some server  $s \in S$ 
6:     send  $(c, (r, x))$  to  $s$ 
7:     wait until receive  $(c, (x, v, i, d))$  from  $s$ 
8:      $t.rs \leftarrow t.rs \cup \{(x, i, d)\}$ 
9:   return  $v$ 
10: Command  $(c, \text{commit})/(c, \text{abort})$  occurs as follows:
11:    $i_{max} \leftarrow \max(\{i \mid (x, i, *) \in t.rs\})$ 
12:    $i_{min} \leftarrow \min(\{i \mid (x, i, *) \in t.rs\})$ 
13:   send  $(c, \text{commit-req}, i_{min}, i_{max})$  to  $s$ 
14:   wait until receive  $(c, vcp(t))$ 
15:   if  $vcp(t)$  is valid and  $C_r^c(t.rs, vcp(t))$  and  $C_r^v(t.rs, vcp(t))$  then
16:      $outcome \leftarrow \text{commit}$ 
17:   else
18:      $outcome \leftarrow \text{abort}$ 

```

x : the data item's key v : the data item's value
 i : the data item's version d : the data item's digest

Table I summarizes the costs of the proposed protocols for the crash-stop and byzantine failure models. To compute these costs, we consider a transaction t that performs r reads and present the latency and number of messages sent for the execution and termination phases, when t is an update and a read-only transaction.

D. Liveness issues

Byzantine servers may compromise the progress of the above protocol by being non-responsive or slow. Besides attacks that would slow down the delivery of atomic broadcast

1: **Algorithm 6: Read-only transactions**, server s 's code:

```

2: Initialization
3:    $forwarded \leftarrow \emptyset$ 
4: upon receiving  $(c, (r, x))$  from  $c$ 
5:   retrieve  $(x, v, i, d)$  from database
6:   send  $(c, (x, v, i, d))$  to  $c$ 
7: upon receiving  $(c, \text{commit-req}, i_{min}, i_{max})$ 
8:    $vcp \leftarrow \emptyset$ 
9:   for  $i$  in  $i_{min}$  to  $i_{max}$  do
10:    wait until element  $e = \langle i, up \rangle_{\Sigma_{f+1}}$  is in  $CT$ 
11:     $vcp \leftarrow vcp \cup \{e\}$ 
12:    send  $(c, vcp)$  to  $c$ 
13: periodically do
14:   for each  $(i, up) \in CT \setminus forwarded$  do
15:     send  $\langle i, up \rangle_{\sigma_s}$  to all servers
16:      $forwarded \leftarrow forwarded \cup \{(i, up)\}$ 
17: upon receiving  $\langle i, up \rangle_*$  from  $f+1$  servers
18:    $CT \leftarrow CT \cup \{\langle i, up \rangle_{\Sigma_{f+1}}\}$ 

```

x : the data item's key v : the data item's value
 i : the data item's version d : the data item's digest

messages [1], byzantine servers may also not answer client read requests or slow down their execution. The first case can be treated as in the crash-stop case, that is, the client may simply consider that the server has crashed and restart executing the transaction on another server. The second case is more problematic since it may not be possible to distinguish between a slow *honest* server and a *malicious* one. To avoid such an attack, the client can execute the transaction on two (or more) servers and abort the transaction on the slower server as soon as the faster server is ready to commit.

A more subtle attack is for a byzantine server to provide read-only transactions with old, but valid and consistent, database views. Although serializability allows old database views to be seen by transactions (i.e., strictly speaking it is not an attack), useful implementations try to reduce the staleness of the views provided to transactions. There are (at least) two ways to confront such server misbehaviour. First, clients can broadcast read-only transactions and ask servers to certify them, just like update transactions. If the transaction fails certification, the client can retry using a different server. Second, clients may submit a read command to more than one server and compare their versions. Submitting read commands to $f+1$ servers ensures the “freshness” of reads, but may be an overkill. More appropriate policies would be to send multiple read commands when suspecting a server misbehavior, and possibly try first with a small subset of servers.

E. Optimizations

Our protocols can be optimized in many ways. In the following, we briefly present three optimizations.

Client caches: To increase scalability, clients can cache data item values, versions, digests, and elements of CT . In

Failure model	Number of servers	Transaction type	Execution		Termination	
			latency	messages	latency	messages
crash-stop	$2f + 1$	update	$r \times 2$	$O(r)$	$\delta(abcast_{cs}) + 1$	$\text{msgs}(abcast_{cs}) + O(n)$
		read-only	$r \times 2$	$O(r)$	-	-
byzantine	$3f + 1$	update	$r \times 2$	$O(r)$	$\delta(abcast_{byz}) + 1$	$\text{msgs}(abcast_{byz}) + O(n^2)$
		read-only	$(r + 1) \times 2$	$O(r)$	-	-

Table I

THE COST OF THE PROPOSED PROTOCOLS (n IS THE NUMBER OF SERVERS, f IS THE MAXIMUM NUMBER OF FAULTY SERVERS, r DENOTES THE NUMBER OF ITEMS THE TRANSACTION READS, $abcast_{cs}$ DENOTES AN ATOMIC BROADCAST ALGORITHM FOR CRASH-STOP FAILURES, AND $abcast_{byz}$ IS AN ATOMIC BROADCAST ALGORITHM FOR BYZANTINE FAILURES).

doing so, clients can execute queries without contacting any server, provided that the necessary items are in the cache. Before inserting a tuple (x, v, i) in the cache, where v and i are x 's value and version respectively, we verify the validity of v and make sure that version i of x has value v by using the appropriate element of CT . We proceed similarly with elements of CT by verifying their signatures before inserting them in the cache. At the end of the execution of a read-only transaction t , the consistency of the values read can be checked using cached elements of CT . If some elements of CT are missing, they are retrieved from a server. If t is an update transaction, the consistency of the values read by t are performed by the certification test. To avoid reading arbitrary old values, cache entries are evicted after some time (e.g., a few hundreds of milliseconds).

Limiting the size of CT : We limit the number of element in CT by some value K to reduce the space overhead of the set of committed transactions on the servers. After the k -th transaction commits and server s inserts tuple $\langle k, up \rangle_{\Sigma_{f+1}}$ into CT , s checks whether CT contains more than K elements. If so, the element of CT with the lowest timestamp is removed. This scheme may force servers to unnecessarily abort transactions due to missing versions in CT . Choosing K must thus be done carefully.

Efficient message authentication: In the above protocol, public-key signatures are used by servers to sign elements of the set of committed transactions and by clients to authenticate their commit messages. Public-key signatures are expensive however. In particular, it is orders of magnitude slower than message authentication codes (MACs). In contrast to public-key signatures, a MAC cannot prove the authenticity of a message to a third party. We thus replace signatures by vectors of MACs [3]. This vector contains one entry per machine of the system, that is, clients and servers. In doing so, any vector of MACs can be verified by any client or server.

F. Correctness

Correctness of the protocol for update transactions relies on the fact that at certification, we use value digests to check for data integrity and versions to check for data staleness. The rest of the correctness argument is based on the same principles as the crash-stop case and we thus omit it.

Proving the correctness of read-only transactions is more subtle. To be serializable, read-only transactions must read values from a consistent view of the database that is the result of a serial execution of some finite sequence of transactions. Let i_{max} and i_{min} respectively be the highest and lowest data item versions read by a transaction t . We claim that if t commits, then the view of the database t observes is the result of the execution, in version order, of all transactions whose corresponding tuple in CT has a version that is smaller than, or equal to i_{max} . Let h_t be the sequence of such transactions sorted in ascending version order.

We first note that since t commits, the server s on which t executed behaved as prescribed by the protocol. This is because each element of the validity and consistency proof $vcp(t)$ of t is signed by $f + 1$ servers, and thus ensures that the values and versions read by t are those that would be returned by a correct server.

Since the client checks that (i) $vcp(t)$ contains all versions between i_{min} and i_{max} and (ii) for any data item x read by t , no newer version of x exists in $vcp(t)$, t reads a consistent view of the database that is the result of the serial execution of transactions in h_t .

V. TOLERATING BYZANTINE CLIENTS

In this section, we discuss serializability in light of byzantine clients and consider some attacks that could be perpetrated by byzantine clients against the deferred update replication technique. We conclude the section with countermeasures to these attacks.

A. Consistency issues

In order to accommodate byzantine clients in our failure model, we must reconsider our consistency criterion since with the current definition, it would be easy for a byzantine client to generate histories that are not serializable. One simple attack is a byzantine client who issues read operations that return bogus information. Another attack is a byzantine client who executes read operations (e.g., against a byzantine server) and does not include them in the transaction's readset. As a consequence, honest servers would certify the transaction using a subset of the transaction's read operations, which could violate serializability.

Our proposal to accomodate byzantine clients in our model is similar to [16], developed in the context of linearizability: Correctness must take into consideration correct clients only. More precisely, (1) correct clients in a history must observe values that could be generated in a serializable history where all clients are correct, and (2) if at some point all byzantine clients stop execution (e.g., the access rights of byzantine clients are revoked), then the system eventually recovers, that is, the number of fictitious writes needed to justify the values read by correct clients is bounded.³

Interestingly, the algorithm presented in Section IV ensures the first property of the modified serializability commented above, but does not guarantee the second property: byzantine clients can leave an unbounded number of “bogus transactions” (denoted “lurking writes” in [15]) in the system after they stop. To see why, notice that a byzantine client can atomically broadcast an unbounded number of transactions before it stops, and it is perfectly possible for these transactions to be delivered after the client stops. In Section V-C we describe a mechanism to guarantee the second property presented above.

B. Byzantine client attacks

Byzantine clients can launch a number of attacks to disrupt the execution of correct clients. The attacks we discuss next target certification by increasing the chances of aborting transactions issued by correct clients.

The most obvious attack to be attempted by a byzantine client is to submit transactions with many write operations and few read operations or no read operation at all (i.e., a “blind transaction”). A reduced number of read operations would increase the chances of the byzantine transaction to commit, while at the same time increasing the probability that legitimate transactions abort. Moreover, as commented in the previous section, a byzantine client can submit an unbounded number of such transactions concurrently.

A more subtle strategy is for multiple byzantine clients to collude and coordinate their attack by minimizing the chances that one byzantine transaction would abort another byzantine transaction, but together they would maximize the chances of aborting legitimate transactions. More precisely, this can be obtained as follow. Let \mathcal{T}_B be a set of byzantine transactions created by colluded byzantine clients. For any two transactions t_i and t_j in \mathcal{T}_B , the readset and the writeset of t_i do not intersect the writeset of t_j .

At the very extreme, coordinated byzantine clients could divide the database into disjoint sets, each set associated with one coordinated byzantine client. Then, each byzantine client would permanently submit blind transactions that modify all items in its set.

³Notice that a Byzantine client may issue commands that violate the integrity of the database. Defining mechanisms that preserve database integrity despite malicious clients is out of the scope of this work.

C. Countermeasures against byzantine clients

In this section we discuss mechanisms to address the attacks presented in the previous section. Our first countermeasure limits the number of concurrent transactions that a client can submit to the system to K and addresses bogus transactions. This countermeasure also reduces the negative impact byzantine clients can have on the abort rate of legitimate transactions.

- *CM 1. At most K concurrent transactions per client.* Two transactions t and t' are concurrent if t starts before t' has terminated (i.e., committed or aborted).

To implement CM 1, we assume that clients are authenticated. Thus, an honest server can unmistakably tell which client has submitted a transaction. Each client c appends a sequence number to each of its submitted transaction. This sequence number is provided together with the result of the previous transaction submitted by c and is signed by $f + 1$ servers. The emission of such sequence numbers is used to limit the number of concurrent transactions clients can submit to K . We now explain in details the implementation of CM 1.

Initially, each server s sends signed sequence numbers 1 through K to each client c , and sets the next sequence number to emit, N_s^c , to $K + 1$. Whenever s accepts to certify a new transaction, s increments N_s^c and sends N_s^c signed to c . Server s only accepts to certify transactions whose sequence numbers belong to set V_s^c . Initially, this set contains sequence numbers 1 to K . Each time s accepts to certify a transaction with sequence number seq , seq is removed from V_s^c and N_s^c is added to the set after being incremented.

Countermeasure 1 avoids bogus transactions since if a byzantine client stops, it can leave at most K transactions, which will be eventually processed. Another way to address bogus transactions is to blacklist clients that were identified as byzantine. A violation to one of the countermeasures proposed here is a way to identify malicious behavior. For instance, a violation of CM1 happens when $f + 1$ servers detect that a client c submitted a transaction with an invalid sequence number. A proof of this violation is obtained by receiving $f + 1$ signed messages asserting that c misbehaved.

We now consider two strategies to reduce the number of aborts due to byzantine clients. Notice that, as commented earlier, aborts due to concurrent transactions are inherent to the deferred update replication technique. This fact makes it impossible to completely countermeasure attacks targeting certification. The techniques we describe for such attacks are best effort, that is, they will reduce the effectiveness of an attack, but they will not avoid it altogether.

- *CM 2. At most L writes per transaction.* The number of transaction writes is upper bounded by simply aborting transactions whose number of writes exceed some constant L at certification time.

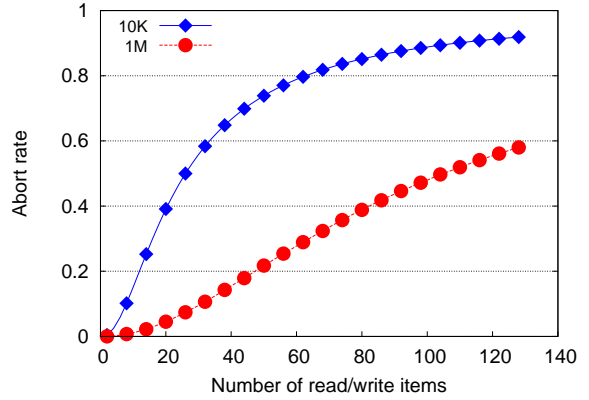



Figure 2. Abort rate versus number of clients (left) and number of items in a transaction (right) in the absence of byzantine clients.

- *CM 3. No blind transactions.* Servers enforce no blind transactions by aborting transactions whose writeset is not included in their readset.

The countermeasures proposed here limit the impact that each individual byzantine client can have on the abort rate but do not address the byzantine client collusion attack described above. In fact, it is unclear how such an attack could be dealt with. Luckily, Section VI shows that this attack has little impact on the abort rate. Determining whether there exist other collusion attacks that would negatively impact the performance of byzantine deferred update replication protocols is an open question.

VI. PERFORMANCE ASSESSMENT

We now assess the attacks described in Section V-B and the effectiveness of the countermeasures proposed in Section V-C through a simple simulation model. We start by describing the simulator and then evaluate the various attacks individually and combined.

A. The simulation model

A run of the simulator is a sequence of steps, where a step can be a read operation or a certification operation. Our goal is to evaluate variations in the abort rates of deferred update replication caused by byzantine clients under data contention, not due to communication and processing delays. Therefore, for simplicity, we assume that each step takes a single time unit.

Table II shows the parameters that can be configured in the simulator and the ranges of values used in our experiments. We consider a large and a small database in the experiments. All read and write operations are uniformly distributed throughout the database, unless noted otherwise in the text. Clients submit a new transaction as soon as the previous one has terminated.

Parameter	Value range
Database size	1 million and 10 thousand entries
Number of honest clients	1..1024
Percentage of byzantine clients	0%, 5%, 10%, 25%, 33%
Concurrent transactions per client	1..10
Reads in honest transactions	2..128
Writes in honest transactions	2..128
Reads in byzantine transactions	2..128
Writes in byzantine transactions	2..128

Table II
SIMULATION PARAMETERS

In all experiments reported, there are no blind honest transactions, that is, every write operation on item x is preceded by a read operation on x . Each point in the graphs is the result of a run of the simulator in which one million honest transactions were executed.

B. Abort rates under normal conditions

Our first set of experiments assesses abort rates in the absence of byzantine clients (see Figure 2). In these experiments, each client submits one transaction at a time (i.e., $K = 1$) and transactions have the same number of read and write operations. In the experiments reported on the left of Figure 2 each transaction has 8 read and 8 write operations in the small database configuration, and 32 reads and 32 writes in the large database configuration. In the experiments reported on the right of the figure, there are 44 and 256 clients, respectively, in the small and in the large database configurations.

As shown in the graphs, deferred update replication is sensitive to the number of clients and the size of transactions. Intuitively, as the number of clients and the number of reads and writes in a transaction increase, the probability of conflicts augments due to data contention, resulting in more aborts. In the space of parameters searched, contention is more important in the small database than in the large database.




Figure 3. Abort rate versus concurrent-transactions attack (left) and write-set attack (right).




Figure 4. Effectiveness of countermeasures CM1 (left, $K = 1$) and CM2 (right, $L = 8$ in small database and $L = 32$ in large database).

C. The “concurrent-transactions” and “write-set” attacks

We initially look at the concurrent-transactions and the write-set attacks. The former attack corresponds to the case in which byzantine clients submit multiple concurrent transactions. The latter attack corresponds to the case in which byzantine clients submit transactions larger than the transactions submitted by honest clients. In both cases, honest and byzantine clients do not submit blind transactions. The abort rates reported in Figure 3 are for honest transactions only.

In both experiments, we set the number of operations in honest transactions to 8 reads and 8 writes in the small database and 32 reads and 32 writes in the large database. There are 44 and 256 clients accessing the small and the large databases, respectively. In the absence of byzantine clients, these configurations lead to abort rates of 10% for both database sizes. In these experiments, we configured 25% of clients to act maliciously.

The number of concurrent transactions submitted by byzantine clients, graph on the left of Figure 3, has almost identical effect on both database setups. Notice, however, that with 25% of byzantine clients, there are 11 byzantine clients in the small database setup and 64 byzantine clients in the large database setup. Thus, increasing the number of concurrent transactions per byzantine client represents a larger number of byzantine transactions in the large database than in the small database.

The graph on the right of Figure 3 shows that growing the size of byzantine transactions leads to more aborted honest transactions. While the impact of this attack seems to be more effective in the small database than in the large database, we note that when byzantine transactions are twice the size of honest transactions (i.e., 16 operations and 64 operations in the small and large database configurations), the abort rate of honest transactions is 12% in both cases.

In both experiments, countermeasures CM1 and CM2 are effective in reducing the abort rate by limiting the number of transactions issued by byzantine clients to one at a time (CM1), as done by honest clients (graph on the left of Figure 4), and by keeping the number of reads and writes in byzantine transactions equal to the number of such operations in honest transactions (CM2) (graph on the right of Figure 4). Effectiveness is the ratio between the abort rate with an attack and without it, after the corresponding countermeasure was applied.

D. The impact of the “blind-transactions” attack

We consider now the blind-transactions attack when combined with the concurrent-transactions attack (graph on the left of Figure 5) and the write-set attack (graph on the right of Figure 5). To evaluate the blind-transactions attack, we ran the experiments presented in the previous section after removing all read operations from byzantine transactions.

Blind transactions, a powerful mechanism to boost byzantine client attacks, benefit from two properties: (a) they never abort and (b) they execute more quickly than honest transactions, which contain read operations. From property (a), every transaction created by a byzantine client may cause honest transactions to abort. Property (b) stems from the fact that without reads, byzantine transactions are executed in a single step. As a consequence, more transactions can be generated by byzantine clients in an execution. We have assessed that property (b) can create more damage to honest transactions than property (a). In our assessment, we configured byzantine transactions to issue read operations that no transaction updates. Thus, such transactions never abort, although they issue read operations. It turned out that in this case, byzantine transactions were much less harmful than blind transactions.

Figure 6 assesses the effectiveness of countermeasures CM1 and CM3 combined, CM2 and CM3 combined, and CM3 alone, targeting blind transactions only. Effectiveness of CM3 alone is calculated as the ratio between the abort rate of the combined attacks (e.g., write-set and blind transactions attacks) and the abort rate without blind transactions (e.g., write-set attack alone).

The effectiveness of CM3 alone decreases with the increase of the number of concurrent byzantine transactions and the number of writes in a byzantine transaction, regardless the database size. This happens because as the abort rate increases, due to more concurrent byzantine transactions per client and augmented writes in byzantine transactions, there are fewer honest transactions to be aborted by blind transactions. Therefore, there is less gain in preventing blind transactions. Combining CM3 with CM1 and CM2 can neutralize both client attacks.

E. The impact of the “colluded-clients” attack

Figure 7 shows the abort rate of the colluded-clients attack combined with the concurrent-transactions and the write-set attacks. In the colluded-clients attack, we assigned each byzantine client to a range of database keys. Ranges were of same size and non-overlapping. A byzantine client creates transactions that read and write data items within its range.

The colluded-clients attack led to a small increase in the abort rate of the two other attacks: the differences between the abort rates reported in Figure 3 and Figure 7 are almost unnoticed. The colluded-clients attack aims to reduce the number of byzantine transactions aborted by other byzantine transactions. Since the overall abort rate is low, this technique leads to few additional committed byzantine transactions, and these transactions are not enough to significantly abort more honest transactions.

We also ran experiments combining the blind-transactions attack and the colluded-clients attack but do not report the results as they provided no new insight. Since blind transactions are never aborted, there was no benefit from avoiding byzantine transactions to overlap.

F. Vulnerability to percentage of byzantine clients

Our last experiments target the effect of the number of byzantine clients. The y -axis in Figure 8 is normalized by the abort rate of the corresponding configuration without byzantine clients. While the effect of a larger number of byzantine clients is more pronounced in the large database configuration than in the small database configuration, in both cases more byzantine clients implies more aborted transactions.

From the graphs, the effect of the number of byzantine clients is limited in two cases: (a) when there are few writes per transaction and (b) when there are approximately 34 writes in the small database and 80 writes in the large database. In case (a), the abort rate is quite low, thus the number of byzantine clients is not important. In case (b), the number of abort is very high, and there is not much more damage that can be caused by adding additional byzantine clients.

G. Summary

To conclude, we have assessed four attacks that can be perpetrated by byzantine clients against certification: (a) write set, (b) concurrent transactions, (c) blind transactions, and (d) colluded clients attacks. We determined that the first three attacks can increase the number of aborts in honest transactions, but not the last. Luckily, the countermeasures proposed in Section V-C target attacks (a)–(c), and can prevent them. Had the colluded byzantine clients attack proved more effective, deferred update replication would suffer from a serious vulnerability since it is not clear what counter measure could be used against collusion.




Figure 5. Abort rate versus blind-transactions attack combined with concurrent-transactions attack (left) and write-set attack (right).

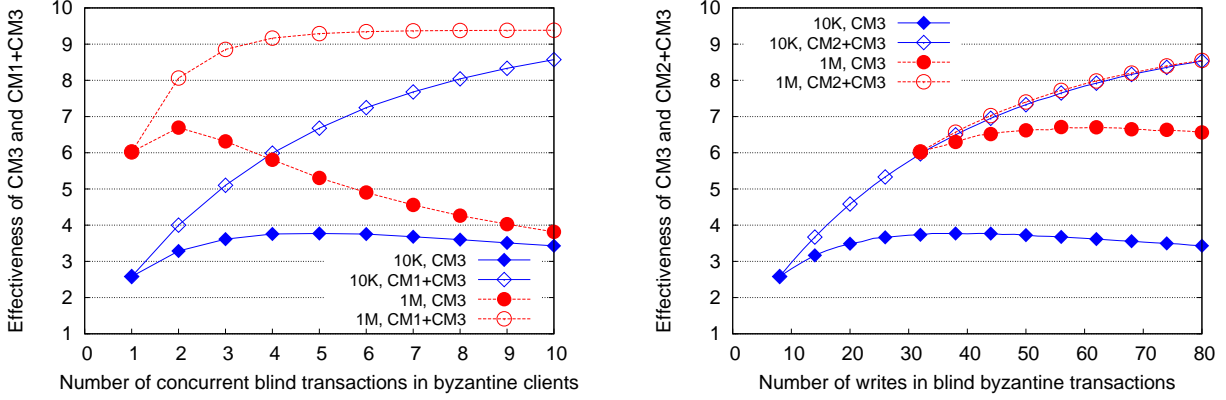


Figure 6. Effectiveness of countermeasures CM3 alone, CM1 combined with CM3 (left), and CM2 combined with CM3 (right).

VII. RELATED WORK

Database replication and deferred update replication have been largely studied under benign faults (e.g., non-byzantine servers subject to crash-stop failures). Few works have considered the effects of byzantine servers on database replication. The first paper to consider the problem is [9], written more than two decades ago. This paper investigates the use of byzantine agreement and state machine replication in the context of databases. It proposes to execute transactions atop serializable databases at the expense of limiting transaction concurrency. The deferred update replication protocol we propose allows concurrency among transactions in that multiple transactions can be simultaneously executed by different servers. Only transaction termination needs to be serialized.

More recently, Vandiver et al. [24] proposed a system that allows more concurrency between transactions. Clients communicate through a central coordinator that chooses a

replica as primary and the rest as secondaries. Transactions are first executed on the primary to determine which transactions can be executed in parallel. When the result of a query is returned to the coordinator, the latter ships this query to the secondaries. A commit barrier, maintained by the coordinator and incremented whenever a transaction commits, is used to determine which transactions can be executed in parallel at the secondaries. That is, two transactions that executed in parallel on the primary will be executed in parallel on the secondaries, provided that no transaction commits in the mean time. This is, roughly speaking, the basics of the Commit Barrier Scheduling protocol proposed in [24]. This approach is similar to ours in the sense that it allows the concurrent execution of transactions, although our protocol allows the execution of transactions at *any replica*. Moreover, in contrast to [24], our protocol does not require a trusted coordinator, a strong assumption.

Byzantium [21] considers byzantine failures of servers




Figure 7. Abort rate versus colluded-clients attack combined with concurrent-transactions attack (left) and write-set attack (right).

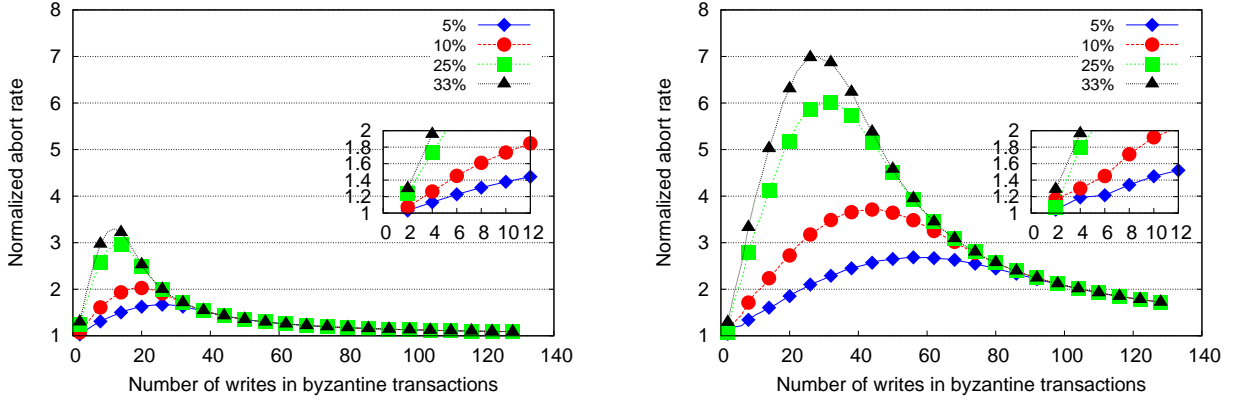


Figure 8. Normalized abort rate versus number of writes in byzantine transactions in small database (left) and large database (right).

that guarantee snapshot isolation, as opposed to serializability. Under *snapshot isolation*, transactions observe a committed instant of the database, which may correspond to the database state when the transaction started or, more likely in a distributed environment, to some earlier state of the database [7]. Transactions that execute concurrently can only commit if they do not modify the same data items, a policy sometimes referred to as *first-committer-wins rule*. In Byzantium, a client first selects a replica that will act as the coordinator for its transaction and then atomically broadcast the *begin* operation to all replicas so that they all use the same database snapshot for the transaction. The transaction is entirely executed by the coordinator. At commit time, the operations along with their results are atomically broadcast to all replicas. If the transaction was executed in a correct coordinator, then a quorum of servers will obtain the same results and the transaction can commit; otherwise, the transaction is aborted and the client is notified

about the byzantine coordinator. Both update and read-only transactions are atomically broadcast.

Byzantium is extended in [8] to a more efficient protocol that allows read-only transactions to be executed on any subset of $f+1$ replicas and propagates operations before commit time for better performance. Their empirical assessment shows that the extended version of Byzantium introduces a moderate performance overhead of roughly 30% over a non-replicated solutions in the TPC-C benchmark.

VIII. FINAL REMARKS

This paper has considered the deferred update replication technique in the byzantine failure model. Deferred update replication has been largely used to implement database replication in the crash-stop failure model. It is more scalable than other replication techniques such as state machine replication and primary-backup since transactions may be executed at any server. Moreover, it allows read-only

transactions to be executed at a single replica. The paper shows that it is surprisingly simple to use deferred update replication under byzantine failures—in fact, our protocol only requires a small modification of the certification procedure and an additional check, performed by clients, to filter out transaction outcomes sent by byzantine servers. The paper also shows that even though some servers may behave maliciously, read-only transactions can be executed at a single server—the execution must be certified by the client at the end of the transaction however. Finally, the paper considers the effects of byzantine client attacks against certification, presents countermeasures against these attacks, and analyzes their performance and effectiveness.

ACKNOWLEDGEMENTS

The authors wish to thank Antonio Carzaniga, Rui Oliveira, Ricardo Padilha, José Orlando Pereira, José Enrique Armendáriz-Iñigo, and the LADC 2011 and JBCS anonymous reviewers for the insightful comments about this work.

REFERENCES

- [1] Y. Amir, B. A. Coan, J. Kirsch, and J. Lane. Byzantine replication under attack. In *DSN*, pages 197–206, 2008.
- [2] P. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987.
- [3] M. Castro and B. Liskov. Practical byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems (TOCS)*, 20(4):398–461, November 2002.
- [4] E. Cecchet, J. Marguerite, and W. Zwaenepoel. C-jdbc: Flexible database clustering middleware. In *USENIX Annual Technical Conference, FREENIX Track*, 2004.
- [5] W. Diffie and M. E. Hellman. Multiuser cryptographic techniques. In *AFIPS '76: Proceedings of the June 7-10, 1976, national computer conference and exposition*, pages 109–112, New York, NY, USA, 1976. ACM.
- [6] S. Elnikety, S. Dropsho, and F. Pedone. Tashkent: Uniting durability with transaction ordering for high-performance scalable database replication. In *Proceedings of EuroSys*, 2006.
- [7] S. Elnikety, F. Pedone, and W. Zwaenepoel. Database replication using generalized snapshot isolation. In *Symposium on Reliable Distributed Systems (SRDS'2005)*, Orlando, USA, 2005.
- [8] R. Garcia, R. Rodrigues, and N. Preguiça. Efficient middleware for byzantine fault tolerant database replication. In *Proceedings of the sixth conference on Computer systems*, EuroSys '11, pages 107–122. ACM, 2011.
- [9] H. Garcia-Molina, F. M. Pittelli, and S. B. Davidson. Applications of byzantine agreement in database systems. *ACM Trans. Database Syst.*, 11(1):27–47, 1986.
- [10] H. Garcia-Molina, J. D. Ullman, and J. Widom. *Database Systems: The Complete Book*. Prentice Hall, 2008.
- [11] B. Kemme and G. Alonso. A new approach to developing and implementing eager database replication protocols. *ACM Transactions on Database Systems (TODS)*, 25(3), September 2000.
- [12] L. Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, May 1998.
- [13] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [14] Y. Lin, B. Kemme, M. Patino-Martinez, and R. Jimenez-Peris. Middleware based data replication providing snapshot isolation. In *International Conference on Management of Data (SIGMOD)*, Baltimore, Maryland, USA, 2005.
- [15] B. Liskov and R. Rodrigues. Byzantine clients rendered harmless. Technical Report MIT-CSAIL-TR-2005-047, MIT, July 2005.
- [16] D. Malkhi, M. Reiter, and N. Lynch. A correctness condition for memory shared by byzantine processes. Unpublished manuscript, Sept. 98.
- [17] J.-P. Martin and L. Alvisi. Fast byzantine consensus. In *DSN'05*, pages 402–411, 2005.
- [18] F. Pedone. *The Database State Machine and Group Communication Issues*. PhD thesis, École Polytechnique Fédérale de Lausanne, Switzerland, 1999. Number 2090.
- [19] F. Pedone, R. Guerraoui, and A. Schiper. Transaction reordering in replicated databases. In *Proceedings of the 16th IEEE Symposium on Reliable Distributed Systems*, Durham (USA), 1997.
- [20] C. Plattner and G. Alonso. Ganymed: scalable replication for transactional web applications. In *Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware*, pages 155–174, 2004.
- [21] N. M. Preguiça, R. Rodrigues, C. Honorato, and J. Lourenço. Byzantium: Byzantine-fault-tolerant database replication providing snapshot isolation. In *HotDep*, 2008.
- [22] R. L. Rivest. The md5 message-digest algorithm. internet rfc-1321. 1992.
- [23] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and Public-Key cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.
- [24] B. Vandiver, H. Balakrishnan, B. Liskov, and S. Madden. Tolerating byzantine faults in transaction processing systems using commit barrier scheduling. In *SOSP*, pages 59–72, 2007.