How to Tolerate Half Less One Byzantine Nodes in Practical Distributed Systems

Miguel Correia, Nuno Ferreira Neves,
Paulo Veríssimo

DI–FCUL TR–04–6

July 2004

Departamento de Informática
Faculdade de Ciências da Universidade de Lisboa
Campo Grande, 1749–016 Lisboa
Portugal

How to Tolerate Half Less One Byzantine Nodes in Practical Distributed Systems[*]

Miguel Correia Nuno Ferreira Neves Paulo Veríssimo
Faculdade de Ciências da Universidade de Lisboa[1]
Bloco C6, Piso 3, Campo Grande, 1749-016 Lisboa - Portugal
{mpc,nuno,pjv}@di.fc.ul.pt

Abstract

The application of dependability concepts and approaches to the design of secure distributed systems is raising a considerable amount of interest in both communities under the designation of intrusion tolerance*. However, practical intrusion-tolerant replicated systems based on the state machine approach can handle at most f Byzantine components out of a total of $n = 3f + 1$, which is the maximum resilience in asynchronous systems.*

This paper extends the normal asynchronous system with a special distributed oracle called TTCB. Using this extended system we manage to implement an intrusion-tolerant service, based on the state machine approach (SMA), with $2f + 1$ replicas only. Albeit a few other papers in the literature present intrusion-tolerant services based on the SMA, this is the first time the number of replicas is reduced from $3f+1$ to $2f+1$. Another interesting characteristic of the described service is a low time complexity.

1 Introduction

The application of dependability concepts and approaches to the design of secure distributed systems is raising a considerable amount of interest in both communities under the designation of *intrusion tolerance* [29, 37]. The idea is that security concepts like vulnerability, attack and intrusion are contained in the dependability notion of fault, therefore it is possible to build secure systems based, to some extent, on dependability mechanisms. This idea has been used to design several protocols and systems in recent years [3, 6, 9, 14, 17, 23, 26, 27, 30, 31, 38].

The *state machine approach* provides a general solution for the implementation of distributed fault-tolerant services [34]. The idea is to implement a service using a set of server replicas in such a way that the overall service can continue to behave as specified even if a number of servers is faulty. If the service is designed to tolerate arbitrary faults, which include attacks and intrusions, then the service can be said to be *intrusion-tolerant*,

[*] This work was partially supported by the FCT through project POSI/1999/CHS/33996 (DEFEATS), project POSI/CHS/39815/2001 (COPE) and the Large-Scale Informatic Systems Laboratory (LASIGE).

or Byzantine-resilient, since these faults are often called Byzantine in the literature [21][1].

This paper presents a solution for the implementation of *state machine replication services (SMR)* in practical distributed systems. The word *practical* is used in this context to signify open distributed systems with networks that provide weak quality of service guarantees, like the Internet, Ethernet LANs and other common network technologies. This kind of systems is often modelled using the *asynchronous model*, which makes no assumptions about processing times, communication delays or clock drift rates. The asynchronous model is extensively used mainly because it is hard to identify realistic bounds for these delays in practical systems. Moreover, for intrusion-tolerant systems, there is an additional motivation: protocols based on time assumptions frequently have subtle vulnerabilities, which can be exploited in order to cause their failure [4, 2]. We are aware of four asynchronous intrusion-tolerant SMR services in the literature: Rampart [31], BFT [4], SINTRA [3] and FS-NewTOP [27].

The *resilience* of a protocol can be defined as the maximum number of faults in the presence of which the protocol still behaves according to its specification. Notwithstanding the advantages of the asynchronous model discussed above, the optimal resilience for an SMR service based on this model is $\lfloor \frac{n-1}{3} \rfloor$, since the problem essentially boils down to atomic multicast [31, 4], which is equivalent to consensus [16]. A proof of the maximum resilience for asynchronous Byzantine consensus can be found in [1]. This means that the service needs $n > 3f$ replicas to tolerate f faults: four replicas to tolerate one fault, seven to tolerate two faults, etc. Each additional fault the system has to tolerate has a significative cost since it requires three additional machines. Moreover, the whole approach is based on the assumption that replicas fail independently, but this is true only if they do not have the same vulnerabilities [4]. This involves using different replicas, i.e., different code running in different operating systems. To summarize, each additional replica has two costs: (1) the cost of its hardware and software; and (2) the cost of its design, since it has to be different from the other replicas. Notice that the number of faults that can be tolerated can be improved either by detecting and removing faulty replicas [34], or by proactively recovering the state of the replicas [5]. However, in a window of time between detection and removal or between recoveries, the resilience remains $\lfloor \frac{n-1}{3} \rfloor$.

This paper presents a solution that reduces the cost of intrusion-tolerant SMR services by decreasing the number of replicas required to tolerate a number of faults/intrusions. More precisely, the presented SMR service has a resilience of $\lfloor \frac{n-1}{2} \rfloor$, i.e., it requires only a majority of correct replicas ($n > 2f$ servers to tolerate f faults). This means a reduction from 25% to 33% on the number of machines to tolerate the same number of faults: three replicas to tolerate one fault, five to tolerate two faults, seven to tolerate three faults, etc. Detection and removal, or proactive recovery of replicas, can also be used to improve the maximum number of faulty replicas.

How is it possible to improve the resilience from $f = \lfloor \frac{n-1}{3} \rfloor$ to $f = \lfloor \frac{n-1}{2} \rfloor$? The solution has something in common with the approach several protocols in the literature use to circumvent the Fischer, Lynch and Paterson (FLP) impossibility result [15]. FLP says that no deterministic protocol can solve the problem of consensus in an asynchronous system if a single process can crash. One of the most common approaches to circumvent this

[1] Throughout the paper we also use the expression *malicious faults* to emphasize that the cause of the fault is an intelligent attacker that has the purpose of violating some property of the system.

result is to extend the asynchronous system with some kind of oracle, like an unreliable failure detector [7, 22, 18] or an ordering oracle [28]. These oracles allow the protocols to circumvent FLP because they encompass some degree of synchrony, e.g., enough synchrony to detect when a process crashed. The solution in this paper also relies on an oracle, but this particular oracle provides two advantages, instead of a single one: circumventing FLP and increasing the resilience.

In the past few years, we have been exploring a type of oracles called *wormholes* [35], to deal with the uncertainty (or lack of coverage) of assumptions as time bounds [36], or intruder resistance [11, 9]. This paper extends the asynchronous system with a wormhole oracle called *Trusted Timely Computing Base (TTCB)*, already introduced elsewhere [11]. This oracle provides a novel ordering service that allow us to implement an atomic multicast protocol with a resilience of $\lfloor \frac{n-1}{2} \rfloor$. This service is the main building block of our SMR solution.

The paper provides the following main contributions:

- it presents an SMR service implemented mostly on a Byzantine asynchronous systems, but that uses the services provided by a wormhole oracle with stronger properties;
- the SMR service has a resilience of $\lfloor \frac{n-1}{2} \rfloor$ instead of the optimal resilience in asynchronous systems of $\lfloor \frac{n-1}{3} \rfloor$;
- the SMR service circumvents the FLP impossibility result without any synchrony assumptions on the asynchronous part of the system; all synchrony necessary to circumvent FLP is in the wormhole oracle;
- the service arguably exhibits good performance since it has a low time complexity.

The paper is organized as follows. Section 2 presents the system model and the TTCB wormhole. Section 3 defines the main TTCB service used in the paper, the Trusted Multicast Ordering Service. Section 4 defines the state machine approach and describes the solution we propose. Section 5 discusses the performance of the service. Section 6 reviews some related work and Section 7 concludes the paper. The appendix presents a correctness proof of the algorithm.

2 System Model and the TTCB

The system is essentially composed by a set of hosts interconnected by a network, called payload network. This environment is asynchronous, i.e., there are no assumptions about processing delays or message delivery delays. The hosts have clocks but there are no assumptions, either about local clock drift rates, or about the reliability of the readings they provide.

The asynchronous environment is extended with a TTCB wormhole, a distributed component with local parts in some of the hosts (local TTCBs) and its own communication channel (TTCB control channel). The system is depicted in Figure 1. Besides being distributed, the TTCB has three important features:

- it is assumed to be secure, i.e., resistant to any possible attacks; it can only fail by crashing;

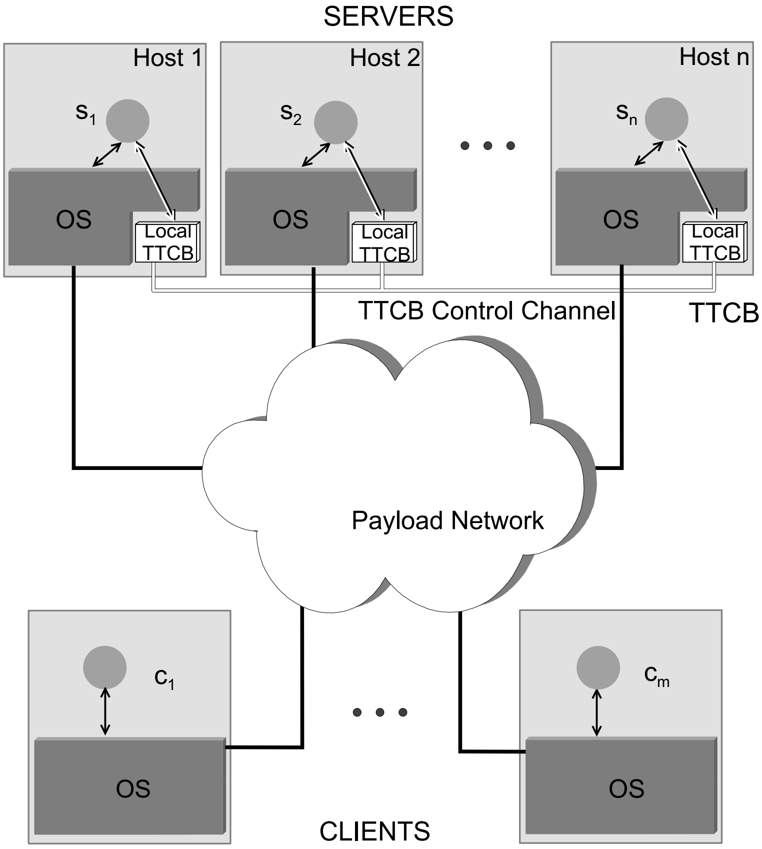

Figure 1. Architecture of the system.

- it is real-time, capable of executing certain operations with a bounded delay;
- it provides a limited set of services, which cannot be possibly affected by malicious faults, since the TTCB is secure.

The TTCB provides a very simple and limited set of services, so that the security of its implementation can be verified. This paper uses only two of these services. The first is the Local Authentication Service, which establishes a trusted path between the server and its local TTCB, i.e., a channel that guarantees the integrity of their communication [11]. The second is the Trusted Multicast Ordering service (TMO), which is the core of our solution and will be described in Section 3.

In relation to the real-time property mentioned above, it is important to make clear that the single consequence of this property for this paper is that the Trusted Multicast Ordering service is not bound by the FLP impossibility result. Otherwise there is no need for the TTCB to be synchronous in the context of this paper. Moreover, the TMO can be implemented in a non-real-time wormhole if another solution is used to circumvent FLP, e.g., randomization or failure detectors.

The approach presented in the paper makes sense only if it is possible to implement the TTCB. There are several possible solutions, which were presented in another paper [11]. Moreover, an implementation based on COTS

(Commercial Off-The-Shelf) components is currently available for free noncommercial use [2]. Let us describe this implementation briefly for the reader to have an idea on how it works.

The local TTCBs have to be secure and real-time. The current TTCB implementation relies on an real-time engineering of Linux called RTAI [8] and is protected by hardening the kernel, since its code is executed inside the kernel. Another solution to protect the local TTCB would be to execute it inside a hardware module inserted in the computer, like a PC/104 board. In relation to the control channel, the current TTCB implementation relies on a Fast-Ethernet network, which is completely independent of the payload network (each host has two network adapters). The control-channel can be assumed to be secure for an inside premises system. Wide-area solutions could be based on virtual private networks over ISDN or Frame Relay. The real-time behavior is ensured by RTAI and by an admission control mechanisms that forces the control channel traffic to be limited and the communication delay bounded. This is a very brief idea and the reader is referred to [11] for a longer discussion on all these issues.

The SMR service is executed by a set of *servers* $S = \{s_1, s_2, ...s_n\}$. The service can be invoked by a set of *clients* $C = \{c_1, c_2, ...c_m\}$. The servers and clients are connected by a fully connected network, although their communication can be delayed arbitrarily, e.g., in consequence of an attack. Every host with a server needs a local TTCB, but not the hosts with clients (see Figure 1). We use the word *processes* to denote both servers and clients. Each server s_i is uniquely identified by eid_i, a number obtained by calling the TTCB Local Authentication Service [11].

A process is *correct* if it follows the protocol it is supposed to execute. We assume that any number of clients can fail, but the number of servers that can fail is limited to $f = \lfloor \frac{n-1}{2} \rfloor$. The failures can be Byzantine or arbitrary, meaning that the processes can simply stop, omit messages, send incorrect messages, send several messages with the same identifier, etc. Faulty processes can pursue their goal of breaking the properties of the protocol alone or in collusion with other corrupt processes. A process is also considered to be faulty if one of the secret keys discussed below is disclosed, or if it is not able to communicate with the local TTCB in its host for some reason (e.g., a local denial of service attack).

The communication among clients and servers is done exclusively through the payload network. The communication among servers is also, to most extent, done through the payload network. We assume that each client-server pair $\{c_i, s_j\}$ and each pair of servers $\{s_i, s_j\}$ is connected by a *reliable channel* with two properties: if the sender and the recipient of a message are both correct then (1) the message is eventually received and (2) the message is not modified in the channel. In practice, these properties have to be obtained with retransmissions and using cryptography. Message Authentication Codes (MACs) are cryptographic checksums that serve our purpose, and only use symmetric cryptography [25]. The processes have to share symmetric keys in order to use MACs. In the paper we assume these keys are distributed before the protocol is executed. In practice, this can be solved using key distribution protocols available in the literature [25]. This issue is out of the scope of the present paper.

Wrapping up, the system is essentially "asynchronous Byzantine": there are no bounds on the processing and communication delays; and the processes can fail arbitrarily. This system is extended with the TTCB wormhole,

[2] The TTCB implementation is available for download at http://www.navigators.di.fc.ul.pt/software/ttcb/

which is synchronous and secure, therefore it provides some “well-behaved” services that the processes can use to perform some steps of their protocols.

3 Trusted Multicast Ordering Service

The SMR service proposed in the paper uses a new TTCB service called *Trusted Multicast Ordering Service (TMO)*. The TMO service is implemented inside the TTCB, therefore its execution cannot be affected by malicious faults.

The TMO service was designed with the purpose of assisting the execution of an intrusion-tolerant atomic multicast (or total-order multicast) protocol. The service does *not* implement the atomic multicast protocol, but simply assigns an order number to the messages. The messages, however, are sent through the payload network, not through the TTCB. This is important since the TTCB has limited processing and communication capacities. Let us introduce briefly how an atomic multicast based on the TMO service can be implemented (the full protocol is introduced later in Section 4.2.1). When a process p wants to send a message to a set of recipients, it makes two operations: (1) it gives the TMO a cryptographic hash of the message and (2) it multicasts the message through the payload network reliable channels. Then, when another process q receives the message, it also gives the TMO a hash of the message it received. When a certain number of processes gave the hash of the message, the TMO service assigns an order number to the message and gives that number to the processes. The processes deliver the messages in that order. Figure 2 illustrates the procedure.

The cryptographic hash mentioned above has to be obtained using a *hash function* h defined by the following properties [25]: *HF1 Compression:* h maps an input x of arbitrary finite length, to an output $h(x)$ of fixed length. *HF2 One way:* for all pre-specified outputs, it is computationally infeasible to find an input that hashes to that output. *HF3 Weak collision resistance:* it is computationally infeasible to find any second input that has the same output as a specified input[3]. *HF4 Strong collision resistance:* it is computationally infeasible to find two different inputs that hash to the same output.

The interface of the TMO service contains three functions: TTCB_TMO_send, TTCB_TMO_receive and TTCB_TMO_decide:

error, tag ← TTCB_TMO_send(eid, elist, threshold, msg_id, msg_hash)

error, tag ← TTCB_TMO_receive(eid, elist, threshold, msg_id, msg_hash, sender_eid)

error, order_n, hash, prop_mask ← TTCB_TMO_decide(tag)

A process is said to start *an execution of the TMO service*, or simply to start *a TMO*, when it calls TTCB_TMO_send. The parameters of this function have the following meanings. The first, *eid*, is the identifier of the sender

[3] A guessing attack is expected to break the property HF3 in 2^m hashing operations, where m is the number of bits of the hash. A birthday attack can be expected to break property HF4 in $2^{m/2}$ hashing operations. In a practical setting, a hashing function with 128 bits like MD5, or 160 bits like SHA-1, can be considered secure enough for our protocol. Nevertheless, we consider HF2, HF3 and HF4 to be assumptions.

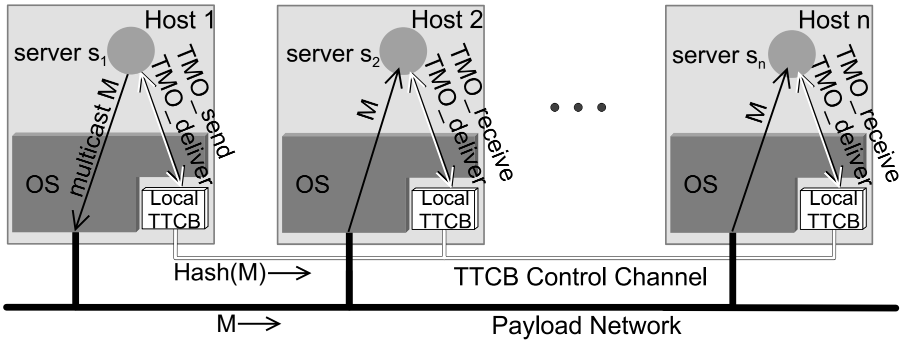

Figure 2. Atomic multicast using the TTCB TMO service.

before the TTCB (see Section 2). *elist* is an array with the identifiers of all processes involved in the set of atomic multicast executions to be ordered. *threshold* is the number of processes in *elist* that must give the TTCB the hash of the message (*msg_hash*) for an order number to be assigned to the message. This parameter will be further discussed in Section 4.2.1. *msg_id* is a message number that has to be unique for the sender. *msg_hash* is a cryptographic hash of the message. The function returns a *tag*, which identifies the TMO execution when the process later calls TTCB_TMO_decide, and an error code.

When a process receives a message it has to call TTCB_TMO_receive. The parameters are the same as for TTCB_TMO_send, except for the eid of the sender, *sender_eid*. How does the TTCB knows that a call to TTCB_TMO_receive corresponds to a certain TMO, which was started by a call to TTCB_TMO_send? It knows by looking at a set of parameters that together uniquely identify a TMO service execution: *(elist, threshold, msg_id, sender_eid)*. The reader should notice that this last sentence has an important implication: if an attacker attempts to break the behavior of the TMO by calling TTCB_TMO_receive with any of these parameters modified, the TTCB will simply consider it to be a call to a different TMO, so the attack would be ineffective.

TTCB_TMO_receive returns a *tag* that is used by TTCB_TMO_decide to identify the TMO. When TTCB_TMO_receive is called and the local TTCB has no data about the TMO, a TMO_UNKNOWN error is returned. If there is data about the TMO but *msg_hash* is different from the hash provided by the sender, a WRONG_HASH error is returned. If there is a TMO_UNKNOWN error, no *tag* is returned; on the contrary, if there is a WRONG_HASH error, the *tag* is returned.

To get the result of the TMO – the order number of the message – a process calls TTCB_TMO_decide. If *threshold* processes did not propose the hash yet, a THRESHOLD_NOT_REACHED error is returned. If there is no error, the function returns the order number *order_n*, the hash of the message *hash*, and a mask with one bit per process, indicating the processes that proposed the correct hash, *prop_mask*. For each TMO execution, the order number returned to all processes must be the same, since the TTCB is assumed to be secure and reliable.

The purpose of the TMO service is to assign consecutive numbers (1, 2, 3, . . .) to a set of TMO executions. At this stage the reader might ask: does the TTCB orders *all* TMO executions with a single sequence of numbers? Or can there be several sets of TMO executions being ordered simultaneously by the TTCB? The answer is related to the purpose of the TMO service: to assist the execution of an atomic multicast protocol; there can be several atomic multicast channels in the system, therefore the TTCB has also to order several sets of TMO executions simultaneously. So, what TMO executions does the TTCB order? The TTCB orders independently each set of TMO executions that belong to the same sequence, defined as follows:

Two TMO executions, identified respectively by (elist$_i$, threshold$_i$, msg_id$_i$, sender_eid$_i$) and (elist$_j$, threshold$_j$, msg_id$_j$, sender_eid$_j$), are said to belong to the same sequence of TMO executions *iff elist$_i$ $=$ elist$_j$.*

TMO Service Implementation

A brief discussion of the implementation we envisage for the TMO service can give a sense of the semantics of the service. The protocol that implements the service is executed by all local TTCBs, which communicate using the TTCB control channel. The protocol can be simple because the TTCB is real-time, local TTCBs can only fail by crashing (they are secure) and they have synchronized clocks. The protocol is implemented on the top of the (crash-tolerant) reliable broadcast protocol presented in [10].

The protocol is based on a fixed coordinator. When a process calls TTCB_TMO_send or TTCB_TMO_receive in a local TTCB, the information about the call is broadcasted to all local TTCBs. When the coordinator gets information about *threshold* calls for a TMO execution, it assigns the next order number to the TMO, defines the mask *prop_mask* with the processes that proposed the correct hash, and broadcasts this information to all local TTCBs. Then, when a process calls TTCB_TMO_decide the order number is returned. If the coordinator crashes, another local TTCB takes over in a consistent manner, since it is aware of the (reliable) broadcasts made by the coordinator.

4 State Machine Replication

A *state machine* is characterized by a set of *state variables*, which define the state of the machine, and a set of *commands* that modify the state variables [34]. Commands have to be atomic in the sense that they cannot interfere with other commands. The *state machine approach* consists of replicating a state machine in n servers $s_i \in S$. The set of servers S implements the *service*. We assume that no more than $f = \lfloor \frac{n-1}{2} \rfloor$ servers fail. All servers follow the same history of states if four properties are satisfied:

- *SM1 Initial state.* All servers start in the same state.
- *SM2 Agreement.* All server execute the same commands.
- *SM3 Total order.* All servers execute the commands in the same order.
- *SM4 Determinism.* The same command executed in the same initial state generates the same final state.

The first property states that each state variable has the same initial value in all servers, something that is usually simple to guarantee. The second and third properties demand that the servers agree in the commands to execute and in their order. This can be guaranteed sending the commands to the servers using an atomic multicast protocol. The fourth property is about the semantics of the commands at the application level, so in this paper we simply make the assumption that the commands are deterministic.

The system works essentially the following way:

1. a client sends a command to one of the servers;
2. the server sends the command to all servers using an atomic multicast protocol;
3. each server executes the command and sends a reply to the client;
4. the client waits for $f + 1$ identical replies from different servers; the result in these replies is the result of the issued command.

This is a simplified description of the process, so let us first delve into the details of the clients, and later we describe the protocol executed by the servers.

4.1 Clients

A client c_i issues a command *cmd* to the service by sending a REQUEST message to one of the servers, s_j. The message is sent through the payload network, since the only communication that is performed inside the TTCB is the one related to the execution of the TMO service. The format of the message is:

$$\langle \text{REQUEST, addr, num, cmd, vec} \rangle$$

where REQUEST is the type of the message, *addr* is the address of the client (e.g., the IP address and the port), *num* is the request number, *cmd* is the command to be executed (including its parameters) and *vec* a vector of MACs (see discussion below). The request number has to be unique, since the SMR service discards requests from the same client with the same number. A solution to generate these numbers is to use a counter incremented for each sent message.

If the client and the server are correct, the REQUEST message is eventually received by s_j, due to the properties of the reliable channels (Section 2). Then, if the server is correct it atomically multicasts the message to all servers in S. Finally, all correct servers in S execute the command and send a reply to the client. The format of the reply message is:

$$\langle \text{REPLY, addr, num, res} \rangle$$

where REPLY is the type of the message, *addr* is the address of the server, *num* is the request number, and *res* is the result of the executed command.

This scheme, albeit simple, is vulnerable to some attacks. A server s_j can be malicious and forward the message only to a subset of S, or discard it altogether. To solve this problem, if c_i does not receive $f + 1$ replies from different servers after T_{resend} units of time read in its local clock, it assumes that s_j did not forward the request, so it multicasts the message to another f servers. If this happens, it sends the message to a total of $f + 1$ servers, therefore at least one must be correct, and the request will eventually be atomically multicasted.

Ideally, T_{resend} should be greater than the maximum round trip delay between any client and a server. However, the payload system is assumed to be asynchronous, i.e., there are no bounds on communication delays, so it is not possible to define such an "ideal" value for T_{resend}. Therefore, the value of T_{resend} involves a tradeoff: if too high, the client can take long to have the command executed; if too low, the client can resend the command without necessity. The value should be selected taking this tradeoff into account. If the command is resent without need, the duplicates are discarded using a mechanism discussed in the next section.

A malicious server might attempt a second attack: to modify the message before multicasting it to the other servers. To tolerate this attack, the request message takes a vector of MACs *vec*. This vector takes a MAC per server, each obtained with the key shared between the client and that server. Therefore, each server can test the integrity of the message by checking if its MAC is valid, and discard the message otherwise[4].

In general, there will be restrictions on the commands that each client can execute. For instance, if the commands are queries on a database, probably not all the clients are allowed to query all registers in the same way. This involves implementing some kind of access control. There are several schemes available in the literature and this issue is orthogonal to the problem we are addressing in the paper, so we do not propose any particular scheme.

4.2 Servers

The protocol executed by the servers is a thin layer on the top of an atomic multicast protocol. A server calls *atomic_mcast(M_{REQ})* to atomically multicast a request M$_{REQ}$ to all servers, and the atomic multicast protocol layer calls *atomic_dlv(M_{REQ})* to deliver M$_{REQ}$ to a server. The basic protocol executed by each server is in Algorithm 1.

Algorithm 1 SMR protocol (for server s_i).

```
When a request M_{REQ} = ⟨REQUEST, addr, num, cmd, vec⟩ is received from a client: if there is no message $M_{REQ}$', with $M_{REQ}$'.addr = $M_{REQ}$.addr and $M_{REQ}$'.num = $M_{REQ}$.num, for which atomic_dlv($M_{REQ}$') has been previously called, then call atomic_mcast($M_{REQ}$); otherwise discard the request.
When atomic_dlv($M_{REQ}$) is called: if there is no message $M_{REQ}$', with $M_{REQ}$'.addr = $M_{REQ}$.addr and $M_{REQ}$'.num = $M_{REQ}$.num, for which atomic_dlv($M_{REQ}$') has been previously called, then execute cmd and send a message ⟨REPLY, addr, num, res⟩ with the result of the command to the client.
```

The objective of checking in both steps if *atomic_dlv(M_{REQ}')*, with *M_{REQ}'.addr* = *M_{REQ}.addr* and *M_{REQ}'.num* = *M_{REQ}.num*, has been previously called, is to guarantee that a request from a client is executed only once. Recall that a client, even if correct, can resend a request (Section 4.1). This condition is implemented using a set

[4] A malicious client might build a vector of MACs with a combination of valid and invalid MACs. This attack would be ineffective: if enough correct servers received the message with the correct MAC the command would be executed, otherwise it would be discarded.

that stores the request number and the client address (*M_{REQ}'.num* and *M_{REQ}'.addr*) for all requests for which *atomic_dlv(M_{REQ}')* has already been called. If the server already received the request, the request is simply discarded (step 1). If several requests with the same command are delivered by the atomic multicast protocol, only the first one causes the execution of the command (step 2). When the command is executed, a reply is sent to the client.

This basic protocol makes at least one atomic multicast for each client request. This cost may be excessive depending on the rate of requests being issued. This cost can be greatly reduced using a *batching mechanism*, i.e., aggregating several requests in a single atomic multicast. The decision about batching requests is left for each server to take; if it assesses that the rate of requests is greater than a given bound, it starts collecting a number of requests before atomic multicasting them together in a single message. This mechanism introduces some delay in the system, so the client's T_{resend} has to take this delay into account.

4.2.1 Atomic Multicast Protocol

The core of the algorithm executed by the servers is the atomic multicast protocol. This protocol guarantees basically two properties: all correct servers deliver the same messages in the same order; if the sender is correct all servers deliver the sent message. More formally, a server is said to (atomically) multicast a message M if it calls *atomic_mcast(M)*, and it is said to (atomically) deliver a message M if *atomic_dlv(M)* is called in the server. The protocol is defined in terms of four properties:

- *AM1 Validity.* If a correct server multicasts a message M with a vector with all MACs valid, then some correct server eventually delivers M.
- *AM2 Agreement.* If a correct server delivers a message M, then all correct servers eventually deliver M.
- *AM3 Integrity.* For any identifier ID, every correct server delivers at most one message M with identifier ID, and if *sender(M)* is correct then M was previously multicast by *sender(M)* [5].
- *AM4 Total order.* If two correct servers deliver two messages M_1 and M_2 then both servers deliver the two messages in the same order.

This definition is similar to other definitions found in the literature, e.g., in [16]. However, property AM1 does not guarantee that the message is delivered in case the message does not have a vector filled with valid MACs (i.e., MACs properly obtained using the key shared between the client and each of the servers). Recall that the objective of this vector of MACs if to prevent a malicious server from atomically multicasting a corrupted request (Section 4.1). Albeit the objective is to deal with malicious servers, if the client itself is malicious and sends a message with some invalid MACs, the message may not be delivered by the atomic multicast protocol.

[5] The predicate *sender(M)* gives the sender field of the header of M.

The protocol is shown in Algorithm 2. It has four parts: initialization (lines 1-8), processing of a call to *atomic_mcast(M)* (lines 9-13), processing of the reception of an ACAST message (line 14), and a task that processes the messages stores in a number of buffers (lines 15-34).

The atomic multicast protocol uses a single type of messages:

$$\langle \text{ACAST, addr, mreq, msg\_id, sender\_eid, elist, threshold} \rangle$$

where *ACAST* is the message type, *addr* the address of the sender server, *mreq* the request message (*mreq* = M_{REQ}), *msg_id* a message number unique for the sender, *sender_eid* the eid of the server that atomically multicasted the message, *elist* is the list of *eid*'s of the processes involved in the protocol, and *threshold* is the value $f + 1 = \lfloor \frac{n-1}{2} \rfloor + 1$.

Lines 1-7 initialize several local variables, including three sets used to store messages in different stages of processing: *Wait_tmo*, *Wait_thresh* and *Wait_deliv*. Line 8 starts task T1.

When *atomic_mcast(M_{REQ})* is called, the server calls *verify_mac* to test if in the vector of MACs, the MAC that corresponds to itself (s_i) is valid (line 9). If it is not, the server simply dismisses the message. If the MAC is valid, the request M_{REQ} is enveloped in an ACAST message and multicasted to all servers except the sender (line 10). Then, the server starts the execution of one instance of the TTCB TMO service by calling TTCB_TMO_send (line 11). Each call to *atomic_mcast* causes exactly one execution of the TMO service. After starting the TMO service, the server puts the ACAST message in the set *Wait_thresh*, waiting for the TMO threshold to be achieved (line 13). When an ACAST message is received by a server, it is simply stored in *Wait_tmo* (line 14).

Task T1 is permanently checking if the messages in the three sets can be processed. The messages in *Wait_tmo* are handled in lines 16-22. For each message in *Wait_tmo,* task T1 makes a call to TTCB_TMO_receive (lines 16-18). If the MAC corresponding to s_i is valid, the hash of message is given to TTCB_TMO_receive (lines 17-18). Otherwise, a value out of the range of valid hashes is given, $\perp$ (lines 17-18). If the local TTCB is still now aware of that TMO execution[6], then TTCB_TMO_receive returns the error TMO_UNKNOWN. If the TTCB is aware of the TMO but the hash of the request is wrong, then an error WRONG_HASH is returned. If the TTCB is aware of the TMO and either the hash is correct, or the hash is $\perp$ (the MAC is invalid), the message is removed from *Wait_tmo* and inserted in *Wait_thresh* (lines 19-22). If the TTCB is aware of the TMO but the hash is wrong (but not $\perp$), the message is discarded since it has been corrupted at some stage (lines 19-22).

The set *Wait_thresh* contains messages waiting for the number of calls to their TMO to reach the threshold. These messages are handled in lines 23-30. The purpose of the *threshold* is to guarantee that the servers only decide to deliver a message if they eventually become able to deliver it. In other words, they can only decide to deliver a message if at least one correct server has the message. This is guaranteed if at least $f + 1$ servers prove that they know the hash of the message, therefore the threshold is set to $f + 1$ (line 3). Notice that a server that received a message with an invalid MAC does not contribute to the threshold, since it gives TTCB_TMO_receive

[6] The TMO is started in the local TTCB of the server that atomically multicasts the ACAST message, so the information about the TMO takes a certain time to be broadcasted and received by the other local TTCBs. Therefore, it is not possible to guarantee that the TMO information will be available in a local TTCB when the corresponding ACAST message is received.

Algorithm 2 Atomic multicast protocol (server s_i).

```
INITIALIZATION:
elist ← {all eid's in E in canonical order}                                  {for TMO service}
msg_id_next ← 1                                                   {number of next ACAST to send}
threshold ← $\lfloor \frac{n-1}{2} \rfloor + 1$                  {threshold for TMO service $(f+1)$}
order_next ← 1                                               {number of next request to deliver}
Wait_tmo ← ∅                                        {set w/recvd ACASTs while TMO unknown}
Wait_thresh ← ∅                                      {set w/ACASTs while thresh. not reached}
Wait_deliv ← ∅                                        {set with requests waiting for delivery}
activate task (T1)

WHEN ATOMIC_MCAST(M$_{REQ}$) IS CALLED DO
if verify_mac(M$_{REQ}$.vec[s$_i$]) then   {if the MAC for s$_i$ is valid, handle M$_{REQ}$, otherwise discard it}
    multicast M$_{ACAST}$ = ⟨ACAST, addr$_i$, M$_{REQ}$, msg_id_next, my_eid, elist, threshold⟩ to servers S \ {$s_i$}
    err, tag ← TTCB_TMO_send(eid$_i$, elist, threshold, msg_id_next, Hash(M$_{REQ}$))
    msg_id_next ← msg_id_next + 1
    Wait_thresh ← Wait_thresh ∪ {(M$_{ACAST}$,tag)}

WHEN M$_{ACAST}$ = ⟨ACAST, ADDR, M$_{REQ}$, MSG_ID, SENDER_EID, ELIST, THRESHOLD⟩ IS RECEIVED DO
Wait_tmo ← Wait_tmo ∪ {M$_{ACAST}$}

TASK T1:
loop
    for all M$_{ACAST}$ ∈ Wait_tmo do                        {messages waiting while TMO is unknown}
        if verify_mac(M$_{REQ}$.vec[s$_i$]) then hash ← Hash( M$_{ACAST}$.mreq)  else  hash ← ⊥
        err, tag ←TTCB_TMO_receive(eid$_i$, M$_{ACAST}$.elist, M$_{ACAST}$.threshold, M$_{ACAST}$.msg_id, hash, M$_{ACAST}$.
        sender_eid)
        if err ≠ TMO_UNKNOWN then
            Wait_tmo ← Wait_tmo \ {M$_{ACAST}$}
            if (err ≠ WRONG_HASH) or (hash = ⊥) then
                Wait_thresh ← Wait_thresh ∪ {(M$_{ACAST}$,tag)}
    for all (M$_{ACAST}$,tag)∈Wait_thresh do            {messages waiting while threshold not reached}
        err, n, hash, prop_mask ← TTCB_TMO_decide(tag)
        if err ≠ THRESHOLD_NOT_REACHED then
            Wait_thresh ← Wait_thresh \ {(M$_{ACAST}$,tag)}
            if Hash(M$_{ACAST}$.mreq) = hash then
                Wait_deliv ← Wait_deliv ∪ {(M$_{ACAST}$.mreq,n)}
                if M$_{ACAST}$.addr ≠ addr$_i$ then                                  {if not the sender}
                    multicast M$_{ACAST}$ to {$\forall_{s_j \in S}$ : $s_j$ ∉prop_mask}
    while $\exists_{(M_{REQ},n) \in Wait\_deliv}$ : n = order_next do          {messages waiting to be delivered}
        Wait_deliv ← Wait_deliv \ {(M$_{REQ}$,n)}
        order_next ← order_next + 1
        ATOMIC_DLV(M$_{REQ}$)
```

the value $\perp$ instead of the hash of the message (lines 17-18). When the threshold is reached, the message is removed from *Wait_thresh* (lines 25-26). If the message corresponds to the hash returned, the message is inserted in *Wait_deliv* (lines 27-28). Then, if the server is not the message sender, it resends the message to the servers that did not 'contribute' to the threshold, i.e., to the servers not in *prop_mask* (lines 29-30). The rationale for resending the message is that a malicious sender can send the message only to a subset of the servers; therefore, these servers may not have the message.

The set *Wait_deliv* keeps messages that already have an order number assigned by the TMO service, therefore they can be delivered. These messages are handled in lines 31-34. The algorithm keeps a number with the next message to be delivered, *order_next*. If the next message to be delivered is stored in *Wait_deliv* (line 31), then task T1 delivers it (lines 32-34). Otherwise, the message has to wait for its turn.

Appendix A gives a proof that the protocol satisfies its specification in terms of properties AM1-AM4.

4.2.2 FLP Impossibility Result

The consensus problem has been proven to be impossible to resolve deterministically in asynchronous systems if a process is allowed to fail, even if only by crashing [15]. This FLP impossibility result also applies to the atomic multicast problem since it is essentially equivalent to consensus [16]. Therefore, it is important to discuss how the atomic multicast protocol proposed in the paper circumvents this result.

A first observation is that our system does not have to be bound by FLP, since it is not strictly asynchronous: it is mostly asynchronous, but includes the TTCB subsystem, which is synchronous. The problem of atomic multicast is essentially equivalent to a consensus about the set of messages to deliver and their order. Our protocol leaves this consensus to the TTCB TMO service, which is executed in a synchronous environment, therefore FLP does not apply.

Another way of reasoning about the problem is given in a paper by Dolev et al. [13]. The TMO service implements a sort of communication with two properties: (1) the communication can be considered to be ordered, since the service service assigns order numbers to the messages and a simple buffering scheme allows the messages to be delivered in that order; (2) the communication is, according to that paper nomenclature, by "broadcast", because all (correct) servers deliver the same messages in the same step. Therefore, the classification of consensus protocols in terms of communication primitives presented in the paper, allow us to conclude that the FLP result does not apply to communication based on a mechanism like the TMO service.

5 Performance

The evaluation of the performance of distributed protocols is usually made in terms of time and message complexities. In asynchronous systems, the *time complexity* is usually measured in terms of the maximum number of *asynchronous rounds* of message exchange. An asynchronous round involves a process sending a message and receiving one or more messages in response. For the Byzantine fault model, only the number of rounds executed by correct processes matter, since malicious processes can behave arbitrarily. We consider separately the number of rounds of TMO execution.

The time complexity is two rounds of message exchange in the payload network, plus one round of TMO executions. Let us justify this complexity by presenting the worst case. The client sends a request to a sender s_j (half round), but s_j is crashed (or is malicious), so s_j does not multicast the message to the other servers. This situation forces the client to resend the request to another f servers, which we count as another half round. Then, all correct servers that received the request, multicast the request in an ACAST message to all other servers (half round) and start one TMO (one round of TMO executions, since all TMOs are executed in parallel). When the first of these TMOs terminates, the command is executed and all correct servers send a reply to the client (half round). Therefore, there are two rounds of message exchange plus one round of TMO executions.

The *message complexity* is measured in number of messages (unicasts) sent. We start by discussing this complexity when the batching mechanism is disabled. The complexity of the SMR service can be divided essentially in three cases:

1. *One request.* For each command, a client sends only one REQUEST message to a single server because the client is correct, the server is correct, and the servers answer in less than T_{resend} units of time measured in the client's clock. A single TMO is executed.

2. $f + 1$ *requests.* For each command, a client sends REQUEST messages to $f + 1$ servers because the servers do not respond before T_{resend}, although both the client and the server for which it first sends the request are correct. $f + 1$ TMOs are executed.

3. n *requests.* For each command, a malicious client sends REQUEST messages to all n servers. n TMOs are executed. A malicious client can issue any number of commands but the SMR protocol prevents it from forcing the execution of more than n TMOs by command (see Algorithm 1).

Table 1 summarizes the message complexities for the three situations. The deduction of these values from the protocol is straight forward.

Requests	Message complexity	TMOs
1	$O(n^2)$	1
$f + 1$	$O(n^3)$	$f + 1$
n	$O(n^3)$	n

Table 1. Message complexity and number of TMOs executed (batching disabled).

The table presents the message complexities and the number of TMOs executed when the batching mechanism is disabled. However, the purpose of this mechanism is precisely to reduce these numbers. How much are they reduced? If we consider that the average number of requests batched in each atomic multicast is B, then the message complexities and the number of TMOs presented in the table have to be divided by B. Therefore, the higher the value of B, the higher is the reduction in the complexity and number of TMOs. Nevertheless, there is a tradeoff involved. To increase B the algorithm has to delay requests until a certain number can be batched in an atomic multicast, therefore increasing the average latency of the algorithm.

6 Related Work

The state machine approach was first introduced by Lamport for systems in which faults were assumed not to occur [19]. Later, Schneider generalized the approach for systems with crash faults [33], and Lamport generalized it for a class of Byzantine faults [20]. The Byzantine faults considered in this latter paper cannot be considered to include malicious faults. The algorithm is essentially synchronous, in the sense that an interaction that exceeds a maximum delay δ is assumed to be a fault, which the algorithm tolerates if its resilience is not exceeded. If the system model assumes malicious faults, then an attacker might purposely delay the communication to force correct processes to be considered failed.

More recently, two Byzantine-resilient or intrusion-tolerant state machine replication systems appeared: Rampart and BFT. Both services have the optimal resilience for asynchronous systems, i.e., $\lfloor \frac{n-1}{3} \rfloor$.

Rampart is an intrusion-tolerant group communication system. It provides a set of communication primitives and a membership service, which handles the joining and leaving of group members [31, 32]. The atomic multicast protocol relies on a reliable multicast protocol that guarantees, essentially, that all correct processes deliver the same messages. When a message is atomically multicast to the group, the reliable multicast protocol is used to send the message. Then, a special process, the sequencer, defines an order for the messages and also reliably multicasts this order to the group. All these protocols use digital signatures to authenticate some messages [25]. Rampart is mostly asynchronous but assumes enough synchrony to detect process failures. The state machine approach is implemented by a set of servers, which form a group [31]. Clients send their requests to a server of their choice, similarly to our algorithm. The output of the service has to be voted so that the results provided by correct servers prevail over those returned by malicious servers. Rampart implemented two solutions. In the first, the client receives individual results from the servers and performs the voting, in the same way as in our approach. In the second, the voting is executed by the servers using a (k,n)-threshold signature scheme [12]. This scheme generates a public key and n shares of the corresponding private key. Each share can be used to obtain a partial signature of a message and any k of those partial signatures form a full signature that can be verified using the public key. Albeit elegant, this scheme has poor performance.

BFT is a Byzantine-resilient state machine replication service. The system is optimized for having good performance, therefore, on the contrary to Rampart, it does not use public-key cryptography most of the time. BFT is not a full-fledged group communication system since it does not have a membership service and does not provide generic group communication primitives, similarly to our algorithm. In BFT, all clients send the requests to the same server, the primary. Then, the primary atomically multicasts the request to the backups (the other servers); all replicas execute the request and send the result to the client; the client waits for $f + 1$ replies with the same result, which is the result of the operation. BFT assumes enough synchrony to detect the failure of the primary. When it fails, a new primary is elected.

SINTRA (previously called Hydra) is a framework aimed to support the implementation of replicated intrusion-tolerant services [3]. It provides a number of group communication primitives, like reliable, atomic and causal multicast. These primitives are implemented on the top of a randomized Byzantine agreement protocol based on cryptographic primitives like threshold cryptography and coin tossing, therefore they are strictly asynchronous. The resilience is also $\lfloor \frac{n-1}{3} \rfloor$.

Besides Rampart, there are two other intrusion-tolerant group communication systems: SecureRing [17] and SecureGroup [26]. However, in the literature about those systems there is no discussion about their use for the implementation of the state machine approach. Castro and Liskov argue that these systems are slower than BFT [4]. The resilience is the same.

FS-NewTOP is a recent intrusion-tolerant SMR system based on fail-signal (FS) processes, i.e., processes that announce when they fail [27]. Each FS process is implemented by two nodes connected by a synchronous channel. Each node monitors its peer. When one node detects that its peer has misbehaved in some way, it signals the failure to all processes and stops the FS process. The resilience is allegedly $4f + 2$, which is sub-optimal. However, the algorithm does not tolerate the failure of two nodes, if they form one FS process. The assumption that two nodes of the same FS process do not fail simultaneously is hard to substantiate in environments prone to malicious faults.

Quorum systems are an alternative to the state machine replication approach to implement fault-tolerant systems. Malkhi and Reiter were, to the best of our knowledge, the first to present a study of their application to tolerate Byzantine faults and to use them to build a dependable data repository that supports shared data abstractions, Phalanx [23]. The applications for quorum systems are not the same as for SMR. Phalanx provides data stores (read/write operations) and locks, not a generic service. Fleet builds on Phalanx but provides support for generic objects instead of just read/write operations on variables [24].

Pedone et al. used *weak ordering oracles* to solve crash-tolerant agreement problems in asynchronous systems [28]. The oracle basically gives a hint about the order of the messages, which may be right or wrong. The hint is simply the order in which the messages are received from the network, which is often right in a local network. Although our work is completely different from theirs, the ‘TTCB with TMO’ oracle might be considered to be a perfect ordering oracle.

7 Conclusion

This paper proposes a novel state machine approach solution. The algorithm is executed in an asynchronous and Byzantine environment, with the exception of a synchronous and secure distributed subsystem, the Trusted Timely Computing Base wormhole. The algorithm is based on a novel TTCB service, the Trusted Multicast Ordering Service, which defines an order for a set of messages represented by their hashes. Using this service, we managed to design an atomic multicast protocol with a resilience lower than the maximum theoretical bound in asynchronous systems: $\lfloor \frac{n-1}{2} \rfloor$ against $\lfloor \frac{n-1}{3} \rfloor$. The paper also shows how the TTCB can be used to circumvent

FLP.

The performance of the system was assessed in terms of time and message complexities, and number of TMOs executed. The system is currently being implemented using the COTS-based TTCB [11].

References

[1] G. Bracha and S. Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM*, 32(4):824–840, October 1985.

[2] C. Cachin, K. Kursawe, and V. Shoup. Random oracles in Contanstinople: Practical asynchronous Byzantine agreement using cryptography. In *Proceedings of the 19th ACM Symposium on Principles of Distributed Computing*, pages 123–132, July 2000.

[3] C. Cachin and J. A. Poritz. Secure intrusion-tolerant replication on the internet. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 167–176, June 2002.

[4] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, pages 173–186, February 1999.

[5] M. Castro and B. Liskov. Proactive recovery in a Byzantine-fault-tolerant system. In *Proceedings of the Fourth Symposium on Operating Systems Design and Implementation*, pages 273–288, October 2000.

[6] M. Castro and B. Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems*, 20(4):398–461, November 2002.

[7] T. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, 43(2):225–267, March 1996.

[8] P. Cloutier, P. Mantegazza, S. Papacharalambous, I. Soanes, S. Hughes, and K. Yaghmour. DIAPM-RTAI position paper. In *Real-Time Linux Workshop*, November 2000.

[9] M. Correia, L. C. Lung, N. F. Neves, and P. Veríssimo. Efficient Byzantine-resilient reliable multicast on a hybrid failure model. In *Proceedings of the 21st IEEE Symposium on Reliable Distributed Systems*, pages 2–11, October 2002.

[10] M. Correia, P. Veríssimo, and N. F. Neves. The design of a COTS real-time distributed security kernel (extended version). DI/FCUL TR 01–12, Department of Computer Science, University of Lisbon, 2001.

[11] M. Correia, P. Veríssimo, and N. F. Neves. The design of a COTS real-time distributed security kernel. In *Proceedings of the Fourth European Dependable Computing Conference*, pages 234–252, October 2002.

[12] Y. Desmedt and Y. Frankel. Threshold cryptosystems. In G. Brassard, editor, *Advances in Cryptology–Crypto ’89*, volume 435 of *Lecture Notes in Computer Science*, pages 307–315. Springer-Verlag, May 1989.

[13] D. Dolev, C. Dwork, and L. Stockmeyer. On the minimal synchronism needed for distributed consensus. *Journal of the ACM*, 34(1):77–97, January 1987.

[14] B. Dutertre, V. Crettaz, and V. Stavridou. Intrusion-tolerant Enclaves. In *Proceedings of the IEEE International Symposium on Security and Privacy*, pages 216–226, May 2002.

[15] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985.

[16] V. Hadzilacos and S. Toueg. A modular approach to fault-tolerant broadcasts and related problems. Technical Report TR94-1425, Cornell University, Department of Computer Science, May 1994.

[17] K. P. Kihlstrom, L. E. Moser, and P. M. Melliar-Smith. The SecureRing group communication system. *ACM Transactions on Information and System Security*, 4(4):371–406, November 2001.

[18] K. P. Kihlstrom, L. E. Moser, and P. M. Melliar-Smith. Byzantine fault detectors for solving consensus. *The Computer Journal*, 46(1):16–35, January 2003.

[19] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.

[20] L. Lamport. Using time instead of timeout for fault-tolerant distributed systems. *ACM Transactions on Programming Languages and Systems*, 6(2):254–280, April 1984.

[21] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.

[22] D. Malkhi and M. Reiter. Unreliable intrusion detection in distributed computations. In *Proceedings of the 10th Computer Security Foundations Workshop*, pages 116–124, June 1997.

[23] D. Malkhi and M. Reiter. Secure and scalable replication in Phalanx. In *Proceedings of the 17th IEEE Symposium on Reliable Distributed Systems*, October 1998.

[24] D. Malkhi, M. K. Reiter, D. Tulone, and E. Ziskind. Persistent objects in the Fleet system. In *Proceedings of the 2nd DARPA Information Survivability Conference and Exposition (DISCEX II)*, June 2001.

[25] A. J. Menezes, P. C. Van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1997.

[26] L. E. Moser, P. M. Melliar-Smith, and N. Narasimhan. The SecureGroup communication system. In *Proceedings of the IEEE Information Survivability Conference*, pages 507–516, January 2000.

[27] D. Mpoeleng, P. Ezhilchelvan, and N. Speirs. From crash tolerance to authenticated Byzantine tolerance: A structured approach, the cost and benefits. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 227–236, June 2003.

[28] F. Pedone, A. Schiper, P. Urbán, and D. Cavin. Solving agreement problems with weak ordering oracles. In *Proceedings of the Fourth European Dependable Computing Conference*, pages 44–61, October 2002.

[29] D. Powell and R. J. Stroud, editors. *Conceptual Model and Architecture of MAFTIA. Project MAFTIA deliverable D21*. January 2002.

[30] H. Ramasamy, P. Pandey, J. Lyons, M. Cukier, and W. H. Sanders. Quantifying the cost of providing intrusion tolerance in group communication systems. In *Proceedings of the International Conference on Dependable Systems and Networks*, pages 229–238, June 2002.

[31] M. K. Reiter. The Rampart toolkit for building high-integrity services. In *Theory and Practice in Distributed Systems*, volume 938 of *Lecture Notes in Computer Science*, pages 99–110. Springer-Verlag, 1995.

[32] M. K. Reiter. Distributing trust with the Rampart toolkit. *Communications of the ACM*, 39(4):71–74, April 1996.

[33] F. B. Schneider. Synchronization in distributed programs. *ACM Transactions on Programming Languages and Systems*, 4(2):179–195, April 1982.

[34] F. B. Schneider. Implementing faul-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, December 1990.

[35] P. Veríssimo. Uncertainty and predictability: Can they be reconciled? In *Future Directions in Distributed Computing*, volume 2584 of *Lecture Notes in Computer Science*, pages 108–113. Springer-Verlag, 2003.

[36] P. Veríssimo and A. Casimiro. The Timely Computing Base model and architecture. *IEEE Transactions on Computers*, 51(8):916–930, August 2002.

[37] P. Veríssimo, N. F. Neves, and M. Correia. Intrusion-tolerant architectures: Concepts and design. In R. Lemos, C. Gacek, and A. Romanovsky, editors, *Architecting Dependable Systems*, volume 2677 of *Lecture Notes in Computer Science*, pages 3–36. Springer-Verlag, 2003.

[38] L. Zhou, F. Schneider, and R. van Renesse. COCA: A secure distributed on-line certification authority. *ACM Transactions on Computer Systems*, 20(4):329–368, November 2002.

A Correctness Proofs

This appendix presents a proof that the atomic multicast protocol in Algorithm 2 behaves according to its specification in Section 4.2, i.e., that it satisfies properties AM1-AM4. We assume the system and TTCB models in Section 2. We also assume that the TMO service behaves according to its specification in Section 3 and that at most $f = \lfloor\frac{n-1}{2}\rfloor$ servers might fail, out of a total of n servers.

Theorem 1 *If a correct server multicasts a message M with a vector with all MACs valid, then some correct server eventually delivers M (AM1 Validity).*

Proof (sketch): Let us consider M$_{REQ}$ = M. A server s_j is said to multicast a message M$_{REQ}$ when it calls *atomic_multicast(M$_{REQ}$)*. This call is handled by lines 9-13. The server follows the protocol since it is assumed to be correct. M$_{REQ}$ is enveloped in a M$_{ACAST}$ message and multicasted to the other servers (line 10). Then, the server calls TTCB_TMO_send and starts a TMO for the message M$_{REQ}$ (line 11).

Now let us consider another correct server s_k. This server eventually receives a correct copy of message M$_{ACAST}$ attending to the properties of the reliable channels and that both s_j and s_k are correct. When s_k receives the message it saves it in *Wait_tmo* (line 14). This set is processed by task T1, specifically by lines 16-22. TTCB_TMO_receive is called with the same *elist* and *threshold* as the ones passed by s_j to TTCB_TMO_send. The *msg_id* and the hash of the message are also the same since the message M$_{ACAST}$ received is the same as the one sent (all MACs are valid so *hash* is set to the hash of the message in line 17). Attending to the properties of the TTCB and the TMO service, TTCB_TMO_receive eventually returns an error different from TMO_UNKNOWN. The error WRONG_HASH can not occur since the message received is the same as the one sent. This allow us to conclude that the message M$_{ACAST}$ is eventually removed from *Wait_tmo* and inserted in *Wait_thresh* (lines 19-22).

This second set, *Wait_thresh*, is handled by task T1 in lines 23-30. TTCB_TMO_decide returns an error THRESHOLD_NOT_REACHED until (1) s_j has called TTCB_TMO_send and (2) $threshold - 1 = \lfloor\frac{n-1}{2}\rfloor$ other servers in *elist* have called TTCB_TMO_receive with *hash* $\neq \perp$. The first condition is a direct consequence of the sender s_j being correct (assumed by the theorem). In relation to the second condition, all correct servers (except s_j) eventually receive M$_{ACAST}$ with a valid MAC for the same reasons as s_k. There are at least $n - f = n - \lfloor\frac{n-1}{2}\rfloor$ correct servers, so $n - \lfloor\frac{n-1}{2}\rfloor - 1$ eventually call TTCB_TMO_receive with *hash* $\neq \perp$. So we need to have:

$$n - \lfloor\tfrac{n-1}{2}\rfloor - 1 \geq \lfloor\tfrac{n-1}{2}\rfloor \Longleftrightarrow n \geq 2\lfloor\tfrac{n-1}{2}\rfloor + 1$$

This is always true, so the threshold is eventually reached and TTCB_TMO_decide eventually returns an error different from THRESHOLD_NOT_REACHED and the hash of the message M$_{ACAST}$. When this happens, s_k puts the message M$_{REQ}$ in *Wait_deliv* (lines 25-28).

This third set is processed in lines 31-34. A message is delivered when its order number n, returned by TTCB_TMO_decide (line 24), is the next one to be delivered, i.e., equal to *order_next* (line 31). TMO gives these numbers in order, starting with 1. Only messages that reach the threshold are counted and all correct servers receive these messages (as discussed above and attending to the fact that all correct servers, in line 30, multicast the M$_{ACAST}$ messages to all servers that they are not aware of having receive them). Therefore, all messages with numbers lower than n are eventually received and delivered, so also does M$_{REQ}$ = M. □

Theorem 2 *If a correct server delivers a message M, then all correct servers eventually deliver M (AM2 Agreement).*

Proof (sketch): For a correct server to deliver a message M$_{REQ}$ = M, *threshold* servers have to give the hash of M$_{REQ}$ to the TMO service (lines 23-24). Before a correct server delivers a message M$_{REQ}$, it multicasts a message M$_{ACAST}$ containing M$_{REQ}$ to all servers that did not give the correct hash of M$_{REQ}$ to the TMO, given in the mask *prop_mask* (line 30). The theorem assumes one correct server delivers M, therefore all correct servers receive the message either from the sender or from another correct server. If a correct server receives the message then it also delivers it (see the proof for Theorem 1). □

Theorem 3 *For any identifier ID, every correct server delivers at most one message M with identifier ID, and if* sender(M) *is correct then M was previously multicast by* sender(M) *(AM3 Integrity).*

Proof (sketch): The identifier (ID) of a message is *(ACAST, msg_id, sender_eid, elist, threshold)*. The TMO service uniquely identifies one TMO execution by the combination of parameters *(elist, threshold, msg_id, sender_eid)*, so only one TMO can be executed for a message with the identifier ID. Therefore, only one message with that ID can be put in *Wait_tmo* (lines 18-20), *Wait_thresh* (lines 21-22), *Wait_deliv* (line 28), and finally delivered (lines 31-34).

The second part of the theorem is an immediate consequence of the properties of the reliable channels that interconnect the servers. □

Theorem 4 *If two correct servers deliver two messages M_1 and M_2 then both servers deliver the two messages in the same order (AM4 Total order).*

Proof (sketch): The servers deliver the messages in the order indicated by the TMO service. This service gives the same order to all servers. □