

# 1 How robust are synchronous consensus protocols?

2 Nenad Milošević ✉

3 Università della Svizzera italiana (USI), Lugano, Switzerland

4 Daniel Cason ✉

5 Informal Systems, Toronto, Canada

6 Zarko Milošević ✉

7 Informal Systems, Toronto, Canada

8 Fernando Pedone ✉

9 Università della Svizzera italiana (USI), Lugano, Switzerland

## 10 — Abstract —

11 Synchronous Byzantine fault-tolerant (BFT) protocols have long been a reality in an academic  
12 setting, yet their practicality remains debated. The main concern of skeptics of synchronous systems  
13 is that the correctness of these protocols depends on the timely delivery of all messages within a  
14 predefined synchronous bound,  $\Delta$ . This dependency creates a challenging tradeoff between protocol  
15 correctness and performance, as  $\Delta$  directly impacts both. In this paper, we examine this tradeoff  
16 in detail. Specifically, we introduce BoundBFT, a new synchronous BFT consensus protocol. We  
17 analyze how BoundBFT's correctness can be compromised and use this analysis to design and  
18 implement the most effective attack strategies that malicious processes could employ. Furthermore, we  
19 experimentally determine the synchronous bound  $\Delta$  that provides sufficient confidence in maintaining  
20 protocol correctness even in the presence of malicious replicas. Finally, we apply this discovered  
21 bound to BoundBFT, evaluate its performance, and compare it to state-of-the-art synchronous and  
22 partially synchronous protocols.

23 **2012 ACM Subject Classification** Computer systems organization → Reliability; Computer sys-  
24 tems organization → Availability; Computer systems organization → Redundancy; Computing  
25 methodologies → Distributed algorithms

26 **Keywords and phrases** Synchronous Consensus, Byzantine Failures, Blockchain

27 **Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2024.28

28 **Funding** This work was partially supported by the Swiss National Science Foundation (grant 175717).

## 29 **1** Introduction

30 Synchronous consensus protocols have long been a topic of debate in robust distributed  
31 systems. On the one hand, synchronous consensus protocols can tolerate  $f < n/2$  Byzantine or  
32 malicious processes out of  $n$  processes [17, 15, 26], an improvement over partially synchronous  
33 consensus protocols, which require  $f < n/3$  [13]. On the other hand, the correctness of a  
34 synchronous protocol hinges on the timely delivery of messages within a fixed time bound,  
35  $\Delta$ . To ensure synchrony is not violated (i.e., messages are delivered within  $\Delta$ ), existing  
36 synchronous consensus protocols set  $\Delta$  conservatively, as the 99.99-th percentile of sampled  
37 communication [30] or as a 10-time factor of average latency [2]. This reliance on  $\Delta$  presents  
38 a critical tradeoff: a more conservative  $\Delta$  minimizes the risk of synchrony violations, thus  
39 favoring correctness, but comes at the cost of reduced protocol performance, since synchronous  
40 protocols execute at the pace of  $\Delta$ .


41 This paper offers a new perspective on synchronous systems by starting with the observa-  
42 tion that some synchronous consensus protocols are resilient to synchrony violations; that is,  
43 even if some messages are not delivered within  $\Delta$ , correctness is not compromised. As we  
44 now illustrate, resilience to synchrony violations happens due to communication diversity



© Nenad Milošević, Daniel Cason, Zarko Milošević, and Fernando Pedone;  
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Principles of Distributed Systems (OPODIS 2024).

Editors: Silvia Bonomi, Letterio Galletta, Etienne Rivière, and Valerio Schiavoni; Article No. 28; pp. 28:1–28:26




Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

## 28:2 How robust are synchronous consensus protocols?

45 and redundancy in a protocol. In Figure 1 (left), process  $p$  sends request  $m_A$  to process  $q$   
 46 and sets a  $2\Delta$  timeout for the answer from  $q$ . Even if  $p$ 's request violates synchrony (i.e.,  $m_A$   
 47 takes longer than  $\Delta$  to arrive at  $q$ ),  $q$ 's response ( $m_B$ ) makes up for the delay and arrives at  
 48  $p$  within the expected  $2\Delta$ . In Figure 1 (right),  $p$  sends request  $m_A$  to  $q$  and  $r$  and sets a  
 49  $3\Delta$  timeout for their answer. Process  $q$  receives  $m_A$  timely, replies to  $p$  ( $m_B$ ) and relays  $m_A$   
 50 to  $r$ . Although  $r$  receives  $m_A$  from  $p$  after  $\Delta$ , it receives  $m_A$  from  $q$  timely and responds  
 51 to  $p$  ( $m_C$ ). As a result,  $p$  receives responses from  $q$  and  $r$  within the expected  $3\Delta$ . These  
 52 communication patterns are at the core of BoundBFT, a novel Byzantine fault-tolerant  
 53 synchronous consensus protocol introduced in the paper.



■ **Figure 1** BoundBFT execution patterns where gray messages violate synchronous bound  $\Delta$  but do not compromise protocol correctness.

54 Tolerating synchrony violations can provide substantial improvement in performance. To  
 55 understand why, consider Table 1, reproduced from [30], which compares the 99.99-th and  
 56 99.999-th percentile of communication across Amazon EC2 datacenters. While a protocol  
 57 that can tolerate one synchrony violation in every one hundred thousand messages exchanged  
 58 between US West and US East must set  $\Delta$  to at least 82190 milliseconds, a protocol that  
 59 tolerates a synchrony violation in every ten thousand transmitted messages can set  $\Delta$  to 1097  
 60 milliseconds. Since synchronous consensus protocols run at the pace of  $\Delta$ , this represents a  
 61  $75\times$  improvement in performance!

62 Since BoundBFT tolerates Byzantine failures, synchrony violations should not introduce  
 63 vulnerabilities that could be exploited by malicious processes. In leader-based consensus  
 64 protocols that tolerate Byzantine failures, such as BoundBFT, the leader is the most  
 65 advantageous role for a malicious process as it can induce honest processes into inconsistent  
 66 decisions, possibly with help from other malicious processes. In a synchronous protocol, the  
 67 malicious leader can hope to get “additional help” from synchrony violations, as some honest  
 68 processes may be delayed with respect to other processes. The attack will work as long as  
 69 the deceived honest process does not find out about the trickery before deciding. But honest  
 70 processes communicate with many other honest processes, so there is ample opportunity to  
 71 find out about the attack even if some messages are delayed.

72 In this paper, we assess the robustness of BoundBFT, that is, its ability to maintain  
 73 correctness under synchrony violations, both in the presence and absence of malicious  
 74 processes. Leveraging BoundBFT’s leader-based execution model with signed messages, we  
 75 first characterize the range of potential attacks and examine their effects when combined  
 76 with synchrony violations. We then design and implement specific Byzantine attacks to  
 77 rigorously evaluate BoundBFT’s robustness. Namely, we conduct experiments to determine  
 78 an appropriate synchrony bound,  $\Delta$ , that provides high confidence in preserving protocol  
 79 correctness, even under the attack. Finally, we apply this bound to evaluate BoundBFT’s  
 80 performance, enabling meaningful comparison with partially synchronous protocols.

	US West (CA)		Europe (EU)		Tokyo (JP)		Sydney (AU)		Sao Paolo (BR)	
	99.99%	99.999%	99.99%	99.999%	99.99%	99.999%	99.99%	99.999%	99.99%	99.999%
US East (VA)	1097	82190	1112	85649	1226	81177	1372	95074	1214	85434
US West (CA)			1184	1974	1133	1180	1209	6354	1252	90980
Europe (EU)					1310	1397	1375	3154	1257	1382
Tokyo (JP)							1149	1414	2496	11399
Sydney (AU)									1496	2134

■ **Table 1** Round-trip latency (in milliseconds) of hping3 across Amazon EC2 datacenters, collected during three months [30].

81 We have implemented BoundBFT and compared it to state-of-the-art synchronous (Sync  
82 HotStuff [2, 3]) and partially synchronous consensus protocols (Tendermint [5] and HotStuff-  
83 2 [32]). Our evaluation in an emulated geographically distributed system showed that  
84 BoundBFT’s synchrony bounds can be in some cases more than one order of magnitude  
85 smaller than usual conservative synchronous bounds [30]. As a result, BoundBFT achieves  
86 latency comparable to partially synchronous consensus protocols, while offering higher  
87 throughput, reliability, and availability.

88 The remainder of the paper is structured as follows. Section 2 defines the system model  
89 and introduces background information on blockchain. Section 3 presents BoundBFT, a  
90 new Byzantine fault-tolerant consensus algorithm designed for the synchronous system  
91 model. Section 4 analyzes BoundBFT under synchrony violations and attacks. Section 5  
92 experimentally evaluates BoundBFT and competing approaches. Section 6 overviews related  
93 work and Section 7 concludes. The Appendix contains BoundBFT’s proof of correctness and  
94 the full data of our experimental evaluation.

## 95 2 Background

### 96 2.1 System model

97 We consider a message-passing geographically distributed system consisting of a set of  
98 processes (or replicas) that do not have access to a shared memory or a global clock. Each  
99 process has its own local (hardware) clock, and while these clocks are not synchronized, they  
100 all run at the same speed. Processes can be *honest* or *faulty*. An honest process follows its  
101 specification; a faulty or Byzantine process presents arbitrary behavior. There are  $f$  faulty  
102 processes out of  $n$  processes. Processes communicate using point-to-point reliable links: every  
103 message an honest sender sends to an honest receiver is received.

104 We assume a *synchronous* system: there exists a known bound  $\Delta$  on maximal network  
105 transmission delay in communication between honest processes. We do not assume lock-step  
106 execution (e.g., [28, 12]); instead, we assume that all honest replicas start the execution  
107 within  $\Delta$  time. We compare our proposed synchronous protocol to protocols that assume  
108 *partial synchrony*: the system is initially asynchronous, without communication bounds, and  
109 eventually becomes synchronous.

110 We use cryptographic techniques for authentication, and digest calculation. We assume  
111 that adversaries (and Byzantine processes under their control) are computationally bound  
112 so that they are unable to subvert the cryptographic techniques used. Adversaries can  
113 coordinate Byzantine processes but cannot delay honest processes.

### 114 2.2 Blockchain

115 A blockchain is a distributed append-only log of transactions implemented by geographically  
116 distributed processes. A blockchain (consensus) protocol forms a chain of blocks, where a

117 block's position in the chain is the block's *height*. A block  $B_k$  at height  $k$  has the following  
 118 format  $B_k := (b_k, H(B_{k-1}))$  where  $b_k$  denotes a proposed value (i.e., a set of transactions)  
 119 and  $H(B_{k-1})$  is a hash digest of the predecessor block. The first block,  $B_1 = (b_1, \perp)$ , has no  
 120 predecessor. Every subsequent block  $B_k$  must specify a predecessor block  $B_{k-1}$  by including  
 121 a hash of it. A block is valid if (i) its predecessor is valid or  $\perp$ , and (ii) its proposed value  
 122 meets application-level validity conditions and is consistent with its chain of ancestors (e.g.,  
 123 there are no double-spending transactions). If block  $B_k$  is an ancestor of block  $B_l$  (i.e.,  
 124  $l \geq k$ ), we say  $B_l$  *extends*  $B_k$ . We say blocks  $B_l$  and  $B'_l$  *equivocate* each other if they do not  
 125 extend one another.

126 We assume that a blockchain (consensus) protocol must satisfy the following properties:

- 127 ■ *Agreement*: No two honest replicas commit different blocks at the same height.
- 128 ■ *Progress*: All honest replicas keep committing new blocks.
- 129 ■ *External validity*: Every committed block satisfies the predefined *valid()* predicate.

### 130 **3 BoundBFT**

#### 131 **3.1 The protocol**

132 BoundBFT is a synchronous BFT consensus protocol with a rotating leader [3]. It tolerates  
 133 up to  $f < n/2$  Byzantine replicas. BoundBFT adopts the good-case execution of the  
 134 rotating-leader version of Sync HotStuff [3, 2], achieving optimal latency and responsive  
 135 leader rotations [3]. However, it introduces a different epoch synchronization mechanism  
 136 that reduces the waiting time for a new leader to propose from  $9\Delta$  to  $5\Delta$  when the previous  
 137 leader is silent.

138 Algorithm 1 presents BoundBFT's pseudo-code that covers executions when leaders are  
 139 honest. BoundBFT's execution evolves as a sequence of epochs, numbered  $0, 1, 2, \dots$ , with  
 140 each replica tracking the last epoch it started, denoted as  $e_p$ . Each epoch  $e$  has a designated  
 141 leader, computed using a deterministic function  $leader(e)$ .

142 At the start of an epoch, the leader  $l$  broadcasts the proposal containing a new block  
 143  $b$  that extends the most recently certified block it knows of,  $validBlock_l$  (lines 14–18 in  
 144 Algorithm 1). Along with the new block, the leader includes the certificate for  $validBlock_l$ ,  
 145  $validBC_l$ .

146 Upon receiving a proposal (lines 19–25 in Algorithm 1), a replica verifies the proposal's  
 147 validity and votes for it if the leader's block certificate is at least as recent as the replica's  
 148  $lockedBC_p$ . The replica votes by sending a signed vote message to all replicas. A vote  
 149 contains the current epoch number and the hash of the block,  $id(b)$ .

150 When a replica receives a proposal and  $f + 1$  votes for it, it forms a block certificate for  
 151 the proposed block. If the replica has no proof of leader  $l$  misbehaving, it locks on  $b$  and  
 152 triggers  $timeoutCommit(e, b)$  (lines 26–32 in Algorithm 1). The replica then updates its  
 153  $validBlock_p$  and  $validBC_p$  variables (lines 33–36 in Algorithm 1) and starts epoch  $e + 1$ . To  
 154 ensure all honest replicas receive the proposal and its certificate, the replica forwards them  
 155 (lines 25 and 35 in Algorithm 1), allowing all honest replicas to start epoch  $e + 1$  within  $\Delta$   
 156 time.

157 When  $timeoutCommit(e, b)$  expires and the replica has no evidence of leader misbehavior  
 158 for epoch  $e$ , it commits block  $b$  and all blocks  $b$  extends (lines 37–40 in Algorithm 1). In  
 159 other words, it directly commits block  $b$  and indirectly commits all its uncommitted ancestor  
 160 blocks.

161 The replica does not wait for  $timeoutCommit(e, b)$  to expire before starting the next  
 162 epoch. Instead, it begins epoch  $e + 1$  immediately after receiving a block certificate in epoch

■ **Algorithm 1** BoundBFT consensus algorithm: normal case

---

```

1: Initialization:
2:  $e_p := 0$  ▷ the current epoch
3:  $hasVoted_p := false$  ▷ has the replica voted in the current epoch?
4:  $validBC_p := nil$  ▷ the most recent block certificate the replica is aware of and...
5:  $validBlock_p := nil$  ▷ the block certified by  $validBC_p$ 
6:  $lockedBC_p := nil$  ▷ the block certificate the replica is locked on and...
7:  $lockedBlock_p := nil$  ▷ the block certified by  $lockedBC_p$ 
8:  $epochsState_p[] := nil$  ▷ an epoch can be in one of the states: ACTIVE, COMMITTED, NOT-COMMITTED
9: when bootstrapping do  $StartEpoch(0)$  ▷ the execution starts in epoch 0
10: Procedure  $StartEpoch(e)$  : ▷ upon starting an epoch:
11:  $e_p \leftarrow e$  ▷ the replica resets the current epoch variables
12:  $epochsState_p[e_p] \leftarrow ACTIVE$ 
13:  $hasVoted_p \leftarrow false$ 
14: if  $leader(e_p) = p$  then ▷ if the replica is the leader in the current epoch...
15:    $block.txs \leftarrow GetTxs()$  ▷ it gets new transactions to include in the new block
16:   if  $validBlock_p \neq nil$  then ▷ then, if it knows of a previously certified block...
17:      $block.prev \leftarrow id(validBlock_p)$  ▷ it links the new block with that block
18:   broadcast  $\langle PROPOSE, e_p, block, validBC_p \rangle_p$  ▷ lastly, it broadcasts the proposal with the new block...
▷ and the certificate for the block it is extending,  $validBC_p$ 
19: when receive  $\langle PROPOSE, e, b, BC \rangle_l$  where  $valid(b)$  and ▷ upon receiving the valid proposal...
20:  $l = leader(e)$  and  $e = e_p$  do ▷ from the leader of the current epoch:
21:   if  $epochsState_p[e] = ACTIVE \wedge hasVoted_p = false \wedge$  ▷ if the epoch is still active, replica has not voted yet, and...
22:    $BC.epoch \geq lockedBC_p.epoch$  then ▷ proposal's  $BC$  is at least as recent as replica's  $lockedBC$ ...
23:     broadcast  $\langle VOTE, e_p, id(b) \rangle_p$  ▷ the replica votes for a proposal, VOTE message contains block's hash
24:      $hasVoted_p = true$  ▷ then, the replica sets  $hasVoted_p$  so it does not vote twice, and...
25:     forward  $\langle PROPOSE, e, b, BC \rangle_l$  ▷ forwards the proposal message
26: when receive  $\langle PROPOSE, e, b, BC \rangle_l$  and  $f + 1$  distinct  $\langle VOTE, e, id(b) \rangle_*$  ▷ when the replica receives a proposal and...
27: where  $e = e_p$  do ▷  $f + 1$  votes from the current epoch:
28:    $cert \leftarrow NewCert$  from  $f + 1$   $\langle VOTE, e, id(b) \rangle_*$  ▷ it forms a block certificate
29:   if  $epochsState[e] = ACTIVE$  then ▷ if no misbehavior is noticed in the current epoch...
30:      $lockedBC_p \leftarrow cert$  ▷ the replica locks on this block by setting  $lockedBC_p$  to  $cert$  and...
31:      $lockedBlock_p \leftarrow b$  ▷  $lockedBlock_p$  to  $b$ , and...
32:     start  $timeoutCommit(e_p, b)$  ▷ starts  $timeoutCommit$ 
33:      $validBC_p \leftarrow cert$  ▷ the replica always updates its  $validBC_p$  and  $validBlock_p$ ...
34:      $validBlock_p \leftarrow b$  ▷ to the most recent block
35:     forward  $messages$  from  $cert$  ▷ lastly, the replica forwards the votes to other replicas and...
36:      $StartEpoch(e + 1)$  ▷ starts the next epoch
37: when  $timeoutCommit(e, b)$  expires do ▷ when  $timeoutCommit$  expires and...
38:   if  $epochsState[e] = ACTIVE$  then ▷ the replica did not observe any proof of misbehavior,
39:      $epochsState[e] \leftarrow COMMITTED$  ▷ the replica commits the block  $b$  and...
40:      $CommitBlockAndItsAncestors(b)$  ▷ all its already uncommitted ancestor blocks

```

---

163  $e$  (line 36 in Algorithm 1). This approach allows BoundBFT to change leaders without  
164 waiting for the conservative network delay  $\Delta$  when we have a sequence of honest leaders,  
165 a property known as *optimistic responsiveness* [32]. Additionally, BoundBFT implements  
166 *pipelining* [39], enabling replicas to start working on the next block before committing the  
167 previous one. Specifically, the leader in epoch  $e + 1$  will propose a new block once it receives  
168 a block certificate for a block in epoch  $e$ .

169 Algorithm 2 presents BoundBFT's pseudo-code responsible for handling Byzantine leaders.  
170 To detect a malicious leader, a replica  $r$  starts a timer,  $timeoutCertificate(e)$ , when it enters  
171 epoch  $e$  (line 2 in Algorithm 2). If  $timeoutCertificate(e)$  expires and  $r$  is still in epoch  $e$ ,  
172 it indicates that  $r$  did not receive a block certificate, which can only occur if the leader is  
173 Byzantine. Consequently, replica  $r$  blames the leader and broadcasts a message  $\langle BLAME, e \rangle_r$   
174 (lines 3–5 in Algorithm 2).

175 When a replica receives  $f + 1$  blame messages for epoch  $e$  from distinct replicas, it  
176 has proof that at least one honest replica blamed the leader and forms a blame certificate  
177  $C_e(BLAME)$  (lines 6–7 in Algorithm 2).

178 Additionally, if an honest replica receives proposals for two distinct blocks signed by the  
179 leader in the same epoch  $e$ , it has proof that the leader is misbehaving. The replica then

■ **Algorithm 2** BoundBFT consensus algorithm: handling malicious leaders

---

1:	<b>upon</b> starting the epoch $e$ <b>do</b>	▷ when a replica enters a new epoch:
2:	<b>start</b> <i>timeoutCertificate</i> ( $e_p$ )	▷ it starts the timer used to detect a malicious leader
3:	<b>when</b> <i>timeoutCertificate</i> ( $e$ ) expires <b>do</b>	▷ when <i>timeoutCertificate</i> expires...
4:	<b>if</b> $e = e_p \wedge \text{epochsState}[e] = \text{ACTIVE}$ <b>then</b>	▷ in the current epoch that is still ACTIVE
5:	<b>broadcast</b> $\langle \text{BLAME}, e_p \rangle_p$	▷ the replica blames the leader by broadcasting a BLAME message
6:	<b>when</b> receive $f + 1$ distinct $\langle \text{BLAME}, e \rangle_*$ <b>do</b>	▷ when receiving $f + 1$ distinct BLAME messages from an epoch:
7:	$\text{cert} \leftarrow \text{NewCert}$ from $f + 1 \langle \text{BLAME}, e \rangle_*$	▷ the replica forms a blame certificate and...
8:	<i>MissbehaviorDetected</i> ( $\text{cert}, e$ )	▷ calls <i>MissbehaviorDetected</i> with the certificate and epoch as parameters
9:	<b>when</b> receive $\langle \text{PROPOSE}, e, b, BC \rangle_p$ <b>and</b> $\langle \text{PROPOSE}, e, b', BC' \rangle_p$	▷ when replica receives two proposals...
10:	<b>where</b> $p = \text{leader}(e)$ <b>and</b> $b \neq b'$ <b>do</b>	▷ from leader for two distinct blocks:
11:	$\text{cert} \leftarrow \text{NewCert}$ from $\langle \text{PROPOSE}, e, b, BC \rangle_p$ <b>and</b> $\langle \text{PROPOSE}, e, b', BC' \rangle_p$	▷ the replica forms...
12:	<i>MissbehaviorDetected</i> ( $\text{cert}, e$ )	▷ an equivocation certificate and calls <i>MissbehaviorDetected</i>
13:	<b>Procedure</b> <i>MissbehaviorDetected</i> ( $\text{cert}, e$ ) :	▷ when misbehavior is detected in an epoch:
14:	<b>if</b> $\text{epochsState}[e] = \text{ACTIVE}$ <b>then</b>	▷ if the epoch is still active
15:	$\text{epochsState}[e] \leftarrow \text{NOT-COMMITTED}$	▷ the replica sets state to NOT-COMMITTED
16:	<b>if</b> $e = e_p$ <b>then</b>	▷ if $\text{cert}$ is the first certificate for the current epoch,
17:	<b>forward</b> messages from $\text{cert}$	▷ the replica forwards the messages from certificate and...
18:	<b>start</b> <i>timeoutEpochChange</i> ( $e_p$ )	▷ triggers <i>timeoutEpochChange</i>
19:	<b>when</b> <i>timeoutEpochChange</i> ( $e$ ) expires <b>do</b>	▷ when <i>timeoutEpochChange</i> expires:
20:	<b>if</b> $e = e_p$ <b>then</b>	▷ if the replica is in epoch $e$
21:	$\text{StartEpoch}(e_p + 1)$	▷ the replica starts the next epoch

---

180 constructs an equivocation certificate  $C_e(\text{EQUIV})$  (lines 9–11 in Algorithm 2).

181 Whenever a replica has proof of the leader’s misbehavior (i.e., a blame or equivocation  
 182 certificate), it calls the function *MissbehaviorDetected*( $\text{cert}, e$ ) and forwards the certificate  
 183 and epoch number to it (lines 8 and 12 in Algorithm 2). If a block is not committed in epoch  
 184  $e$ , the replica marks the epoch state as NOT-COMMITTED (line 15 in Algorithm 2). Moreover,  
 185 if  $\text{cert}$  is the first certificate in epoch  $e$ , the replica forwards the certificate and triggers  
 186 *timeoutEpochChange*( $e$ ) (lines 16–18 in Algorithm 2). Forwarding the certificate ensures  
 187 that all honest replicas learn that the leader is Byzantine within  $\Delta$  time. Additionally, the  
 188 extra timeout allows the replica to learn if an honest replica  $r$  moved to the next epoch  
 189 before detecting leader misbehavior, i.e.,  $r$  received a block certificate in epoch  $e$ , locked on  
 190 it, and moved to the next epoch.

191 When *timeoutEpochChange*( $e$ ) expires and the replica is still in epoch  $e$ , it moves to  
 192 epoch  $e + 1$  (lines 19–21 in Algorithm 2). Replicas wait for *timeoutEpochChange* before  
 193 moving to the next epoch only in the case of a Byzantine leader. If the leader is honest,  
 194 replicas form a block certificate and move to the next epoch without waiting for any timeouts.

## 195 3.2 BoundBFT’s correctness

196 In this section, we provide the intuition behind BoundBFT’s correctness. The appendix  
 197 contains a detailed correctness proof of BoundBFT.

### 198 3.2.1 Intuition behind epoch synchronization

199 The epoch synchronization mechanism guarantees that honest replicas progress through each  
 200 epoch in a coordinated manner. Specifically, all honest replicas initiate each epoch within  $\Delta$   
 201 time, ensuring synchronization. Additionally, Byzantine replicas cannot disrupt or halt the  
 202 protocol during any epoch.

203 The epoch synchronization mechanism in BoundBFT relies on certificates: to start a new  
 204 epoch, a certificate (i.e., block, blame, or equivocation certificate) must be formed in the  
 205 previous epoch. BoundBFT ensures that, regardless of Byzantine behavior, a certificate is  
 206 created in each epoch.

207 The mechanism BoundBFT employs to guarantee the existence of a certificate is as  
 208 follows: honest replicas initiate a *timeoutCertificate* upon entering a new epoch (line 2  
 209 in Algorithm 1). If the timeout expires without receiving a certificate, the replica blames  
 210 the leader. This results in two possible outcomes: (i) an honest replica forms one of the  
 211 certificates, or (ii) no honest replica receives a certificate before the *timeoutCertificate*  
 212 expires. In case (ii), all  $f + 1$  honest replicas will blame the leader, resulting in the formation  
 213 of a blame certificate.

214 Once a certificate is ensured for an epoch, synchronizing replicas becomes straightforward:  
 215 each replica forwards the received certificate (lines 25 and 35 in Algorithm 1 and line 17 in  
 216 Algorithm 2). Within  $\Delta$  time, all honest replicas receive the certificate and start the next  
 217 epoch if they have not already done so.

### 218 3.2.2 Intuition behind agreement

219 BoundBFT ensures that no two honest replicas commit different blocks in the same blockchain  
 220 height. Consequently, the resulting blockchain remains consistent and does not have forks.

221 In epochs with a Byzantine leader, multiple certificates can be created. As a result,  
 222 different honest replicas may start the next epoch receiving different certificates. For instance,  
 223 one honest replica may receive a block certificate, while another may receive an equivocation  
 224 certificate. To account for this scenario, an honest replica commits a proposed block  $b$  in  
 225 epoch  $e$  only if it knows that the first certificate received by all honest replicas in  $e$  is a  
 226 certificate for  $b$ . This guarantees two properties: (i) all honest replicas vote for block  $b$ , and  
 227 (ii) all honest replicas lock on block  $b$  in epoch  $e$ . Property (i) ensures that no other block  
 228 can be certified and afterward committed in epoch  $e$ . Property (ii) guarantees that honest  
 229 replicas vote only for blocks extending  $b$  in the following epochs. As a result, only  $b$  and  
 230 blocks extending  $b$  will be certified and committed in epochs  $e' \geq e$ , and the agreement  
 231 property will be satisfied.

232 The mechanism BoundBFT uses to verify the commit condition is as follows. Upon  
 233 receiving a certificate for block  $b$ ,  $C_e(b)$ , as the first certificate in epoch  $e$ ,  $r$  forwards  
 234 the certificate and triggers *timeoutCommit*( $e$ ) at time  $t$  (lines 32 and 35 in Algorithm 1).  
 235 Consequently,  $r$  knows that all honest replicas will receive  $C_e(b)$  by time  $t + \Delta$ . If an  
 236 honest replica  $p$  received a different certificate before  $C_e(b)$ , it must have received it at time  
 237  $t_1 < t + \Delta$ . Since  $p$  also forwards its certificate,  $r$  will receive it by time  $t_1 + \Delta < t + 2\Delta$ .  
 238 Therefore, setting *timeoutCommit*( $e$ ) to  $2\Delta$  ensures that  $r$  receives  $p$ 's certificate on time.  
 239 Ultimately, if *timeoutCommit*( $e$ ) expires and  $r$  has not heard about any other certificates,  $r$   
 240 can be sure that the first certificate received by all honest replicas in  $e$  is  $C_e(b)$ . In this case,  
 241  $r$  commits block  $b$ .

### 242 3.2.3 Intuition behind progress

243 BoundBFT ensures that all honest replicas commit a new block in every epoch with an honest  
 244 leader. It does so by: (i) ensuring that all honest replicas vote for the leader's proposal, and  
 245 (ii) preventing the creation of blame or equivocation certificates. Property (i) guarantees the  
 246 creation of a unique block certificate, while property (ii) guarantees that all honest replicas  
 247 must receive the block certificate, trigger *timeoutCommit*, and, when it expires, commit the  
 248 proposed block.

249 An honest leader of an epoch proposes a new block that extends its *validBlock* and sends  
 250 *validBC* together with the new block. Other honest replicas will vote for the new proposal  
 251 only if the *validBC* sent by the leader is at least as recent as their *lockedBC*. Consequently,

## 28:8 How robust are synchronous consensus protocols?

252 BoundBFT ensures that whenever an honest replica locks on a block in epoch  $e$ , all honest  
253 replicas update the *validBlock* and *validBC* to the block certified in epoch  $e$ . Therefore,  
254 the *validBCs* on all honest replicas are always at least as recent as *lockedBCs* on all honest  
255 replicas.

256 BoundBFT ensures that *validBlock* and *validBC* are always up to date by relying on a  
257 mechanism that uses *timeoutEpochChange*. Namely, an honest replica  $r$  cannot start the  
258 next epoch immediately if the first certificate it receives in the current epoch is a blame or  
259 equivocation certificate. Instead, it must ensure no other honest replica locks on a block in  
260 this epoch. Consequently,  $r$  forwards its certificate (line 17 in Algorithm 2), knowing that  
261 in  $\Delta$  time, all honest replicas will receive it. If any honest replica  $p$  locked on a block, it  
262 must have done so before receiving the forwarded certificate. As a result, upon forwarding  
263 its certificate,  $r$  sets *timeoutEpochChange*( $e$ ) to expire in  $2\Delta$  time (line 18 in Algorithm 2).  
264 Moreover,  $r$  starts the next epoch only when this timeout expires, or it receives the block  
265 certificate for the current epoch. Since  $p$  also forwards the certificate after locking (line 35  
266 in Algorithm 1),  $r$  knows it will receive it before *timeoutEpochChange* expires. Notably,  $r$   
267 will not lock on a block certificate if it receives the certificate after *timeoutEpochChange* is  
268 initiated.

269 Lastly, BoundBFT must ensure that no equivocation or blame certificates are possible in  
270 epochs with honest leaders. An equivocation certificate will not be formed since the honest  
271 leader will not propose two different blocks. However, ensuring that no blame certificate  
272 is possible requires that no honest replica blames the leader. In other words, every honest  
273 replica must receive the block certificate before *timeoutCertificate*( $e$ ) expires. Consequently,  
274 honest replicas set *timeoutCertificate*( $e$ ) to  $3\Delta$ . The first  $\Delta$  accounts for epoch drift time,  
275 the second  $\Delta$  for the time it takes for the leader's proposal to reach all honest replicas, and  
276 the last  $\Delta$  is for the reception of the votes broadcast by honest replicas. Since we already  
277 showed that all honest replicas will vote for the honest leader, the block certificate is formed  
278 on all honest replicas before the *timeoutCertificate*( $e$ ) expires, and no honest replica blames  
279 the leader.

## 280 **4 Debunking synchrony violations**

281 BoundBFT relies on synchrony every time it uses one of its timeouts, expressed as a multiple  
282 of a synchrony bound  $\Delta$ . A synchrony violation may result in a scenario where the timeout  
283 expires before a replica receives an expected message from some honest replica. We refer  
284 to this phenomenon as a *timeout violation*. In this section, we first examine how malicious  
285 replicas may attempt to compromise BoundBFT. We then consider the consequences of  
286 timeout violations in the presence and absence of malicious replicas. We present a detailed  
287 Byzantine protocol in the appendix.

### 288 **4.1 Byzantine behavior**

289 Listing possible faulty behaviors of Byzantine replicas is unusual since, by definition, a  
290 Byzantine replica can behave arbitrarily. In the context of leader-based protocols where  
291 messages are signed, however, the scope for deviation is limited, as we now explain.

292 In the case of a Byzantine leader, these are the possible faulty behaviors:

- 293 ■ *SILENCE*: The leader does not send a proposal to a subset of replicas (possibly all).
- 294 ■ *EQUIVOCATION*: The leader proposes multiple blocks in the same epoch.

295 ■ *AMNESIA*: The leader does not extend the blockchain with a new block and proposes an  
 296 alternative block for one of the previously committed blocks.

297 In addition, non-leader Byzantine replicas may misbehave as follows.

298 ■ *MULTI-VOTE*: A Byzantine replica can choose to vote for any proposal it likes. It can  
 299 also vote for multiple proposals in the same epoch.

300 ■ *BLAME*: A Byzantine replica can blame the leader by broadcasting a blame message at  
 301 any point in the execution, even if the leader is honest.

302 In addition, Byzantine replicas can always remain silent or discard messages selectively.

## 303 4.2 Timeout *timeoutCommit*

304 The *timeoutCommit* is the only timeout responsible for BoundBFT's agreement. An honest  
 305 replica sets this timeout after locking on a block in an epoch (line 32 in Algorithm 1). The  
 306 replica then forwards the block certificate and waits for this timeout to receive certificates  
 307 from all other honest replicas.

308 If *timeoutCommit* is violated, an honest replica may miss a certificate from some honest  
 309 replicas. As a result, it may commit block  $b$  in epoch  $e$  thinking all honest replicas locked  
 310 on  $b$ , while in reality, some honest replicas received a different certificate and moved to  
 311 epoch  $e + 1$  without locking on  $b$ . If enough honest replicas did not lock on  $b$ , the agreement  
 312 might be compromised as honest replicas may vote for an alternative block  $b'$ , create a block  
 313 certificate, and commit  $b'$ .

314 The likelihood of this situation in the absence of Byzantine replicas is low, as the following  
 315 conditions must be fulfilled:

- 316 1. A blame certificate must be formed in epoch  $e$ , meaning a majority of replicas must  
 317 have blamed the leader in epoch  $e$  (i.e., *timeoutCertificate* was violated in all of these  
 318 replicas).
- 319 2. A majority of replicas did not lock on  $b$  in epoch  $e$ , receiving a blame certificate before  
 320 receiving a block certificate for  $b$ .
- 321 3. The leader of epoch  $e + 1$  did not receive  $b$ 's block certificate in epoch  $e$ , thus not updating  
 322 its *validBlock* and *validBC* to  $b$  and  $b$ 's certificate (i.e., its *timeoutEpochChange* was  
 323 violated in epoch  $e$ ).

324 Even though *timeoutCertificate* and *timeoutEpochChange* are responsible for BoundBFT's  
 325 progress, in this scenario, they also play a role in guarding the protocol's agreement.

326 Byzantine replicas can exploit *timeoutCommit* violations and potentially compromise  
 327 BoundBFT's agreement through the following attacks:

328 ■ *AMNESIA-ATTACK*: The Byzantine leader ignores the algorithm (line 17 in Algorithm  
 329 1) and does not propose a block that extends its *validBlock*. Instead, it proposes an  
 330 alternative block  $b$  for its *validBlock* (i.e.,  $b.prev = validBlock.prev$ ). Byzantine replicas  
 331 vote for this proposal. The agreement can be violated if an honest replica committed  
 332 *validBlock* while some honest replicas, due to *timeoutCommit* violations, did not lock  
 333 on *validBlock*. As a result, these replicas will vote for block  $b$ , and if their votes, together  
 334 with Byzantine votes, form a majority, block  $b$  will be certified and committed. To  
 335 increase the probability of this scenario, in epochs with an honest leader, Byzantine  
 336 replicas send votes for the block proposed by the honest leader to one subset of honest  
 337 replicas to help them form the block certificate faster and commit a block. At the same  
 338 time, they send BLAME messages to a different subset of honest replicas to help them  
 339 form a blame certificate before receiving a block certificate.

## 28:10 How robust are synchronous consensus protocols?

340 ■ *EQUIVOCATION-ATTACK*: The Byzantine leader proposes two distinct proposals in  
341 the same epoch, and Byzantine replicas vote for both. The first proposal and its votes are  
342 sent to one subset of honest replicas, and the second proposal and its votes to another  
343 subset. As a result, two honest replicas may vote for, receive block certificates, and  
344 commit different blocks if their *timeoutCommits* expire before they learn about the  
345 equivocated proposal. In epochs when the leader is honest, Byzantine replicas remain  
346 silent and update *validBlock* and *validBC* to stay aware of the most recently certified  
347 block. This is important so that when they become the leader, Byzantine replicas can  
348 generate new blocks that honest replicas will vote for.

### 349 4.3 Timeout *timeoutCertificate*

350 Honest replicas initiate this timeout upon starting an epoch (line 2 in Algorithm 2). Its  
351 purpose is twofold. First, it ensures that an honest replica does not wait indefinitely for  
352 a silent malicious leader. Second, when the leader is honest and proposes a block that all  
353 honest replicas will vote for, *timeoutCertificate* should not expire before all honest replicas  
354 receive the proposal and the votes from all other honest replicas. In other words, before they  
355 receive a block certificate. Consequently, no honest replicas will blame an honest leader.

356 If *timeoutCertificate* is violated, an honest replica will incorrectly blame the honest  
357 leader. If a majority of honest replicas blame the leader, a blame certificate can be formed,  
358 and the decision might not be reached in the current epoch. However, the block will be  
359 committed when the next honest leader proposes a block and other honest replicas receive  
360 the block certificate on time.

361 The creation of a blame certificate is easier in the presence of Byzantine replicas:

362 ■ *BLAME-ATTACK*: Byzantine replicas do not vote for the proposal sent by the honest  
363 leader. Instead, they broadcast *BLAME* messages upon starting the epoch with an honest  
364 leader. As a result, the blame certificate can be formed if a single honest replica blames  
365 the leader. In epochs when the leader is Byzantine, the Byzantine replicas just remain  
366 silent.

367 Apart from unconditionally blaming the leader and hoping that one of the honest  
368 replicas will also blame the leader, there is no other way for Byzantine replicas to prevent  
369 honest replicas from committing a block in epochs with honest leaders. *timeoutCertificate*  
370 violations may slow down the execution but will not lead to violations in agreement (i.e.,  
371 when two honest replicas decide on different blocks).

### 372 4.4 Timeout *timeoutEpochChange*

373 An honest replica  $r$  triggers this timeout when it receives a blame or equivocation certificate  
374 as the first certificate in an epoch (line 18 in Algorithm 2). This timeout ensures that if  
375 another honest replica receives a block certificate as the first certificate and locks on it in  
376 the same epoch,  $r$  will receive this block certificate and update its *validBlock* and *validBC*  
377 before starting the next epoch. Consequently,  $r$  starts the next epoch when it receives a  
378 block certificate from the current epoch or when *timeoutEpochChange* expires.

379 If the *timeoutEpochChange* is violated, an honest replica will not hear about the locked  
380 block. Consequently, if the replica is the next epoch leader, it will propose a block that locked  
381 replicas will not accept. If the set of remaining honest replicas that vote for a proposed block  
382 is less than the majority, the block certificate will not be formed, and a decision will not be  
383 reached even though the epoch leader is honest. However, the new block will be committed  
384 when one of the locked honest replicas becomes a leader.

385 Honest replicas rely on this timeout only in epochs when an equivocation or a blame  
 386 certificate is formed. In the absence of attacks, creating an equivocation certificate is  
 387 impossible. As a result, an honest replica uses *timeoutEpochChange* solely in case it receives  
 388 a blame certificate. This can happen only if *timeoutCertificate* is violated on a majority  
 389 of honest replicas, and they blame the honest leader.

390 With Byzantine replicas, however, both equivocation and blame certificates are possible.  
 391 Byzantine replicas can exploit *timeoutEpochChange* violations as follows:

- 392 ■ **EQUIVOCATION-CERTIFICATE-ATTACK**: The Byzantine leader broadcasts a pro-  
 393 posal for block  $b$ . Then, the Byzantine replicas send votes for block  $b$  to one subset of  
 394 honest replicas to help them create a block certificate and lock on  $b$ . At the same time,  
 395 the Byzantine leader sends a second proposal for block  $b' \neq b$  to the other subset of  
 396 honest replicas. These honest replicas will form an equivocation certificate and start  
 397 *timeoutEpochChange*. As a result, they will not lock on block  $b$ .
- 398 ■ **BLAME-CERTIFICATE-ATTACK**: Similarly, Byzantine replicas impose locking on one  
 399 subset of honest replicas by sending a proposal and votes for block  $b$ . Instead of equivoc-  
 400 ating, the Byzantine leader remains silent and does not send any proposal to the other  
 401 subset of honest replicas. Moreover, all Byzantine replicas send blame messages to these  
 402 replicas. If these replicas did not receive a block certificate before the *timeoutCertificate*  
 403 expires,<sup>1</sup> they will blame the leader and, together with blame messages from Byzantine  
 404 replicas, form a blame certificate and start *timeoutEpochChange*, without locking on  
 405 block  $b$ .

406 In both attacks, Byzantine replicas remain silent in epochs with an honest leader. By  
 407 remaining silent, these replicas ensure that if a *timeoutEpochChange* violation happened  
 408 and the next honest leader proposes a block that does not extend block  $b$ , a block certificate  
 409 will not be formed and a decision will not be reached.

410 Similarly to *timeoutCertificate*, violations of *timeoutEpochChange* can slow down the  
 411 execution but cannot lead to violations in agreement.

## 412 **5 Evaluation**

### 413 **5.1 Experimental environment and setup**

414 We conducted our experiments in a cluster with emulated wide-area latencies between 6 AWS  
 415 zones (see Table 1). Latencies between nodes were configured using the Linux Traffic Control  
 416 kernel module [25]. The emulated WAN provided an affordable approximation of the AWS  
 417 environment since our evaluation required hundreds of hours of experiments (see Appendix  
 418 B). The cluster contains 60 nodes divided in two groups: (i) EPYC Zen 2 with two 16-Core  
 419 AMD EPYC 2881 MHz and 32GB of RAM, and (ii) HP SE1102 with two Quad-Core Intel  
 420 Xeon 2.5GHz and 8GB of RAM. We implemented BoundBFT, all competing protocols, and  
 421 all proposed attacks (see the Appendix) in Go. The implementations use SHA256 hashes  
 422 and Ed25519 64-byte digital signatures. We rely on libp2p [1] for communication between  
 423 replicas.

<sup>1</sup> Even though the Byzantine replicas do not send the proposal and votes to these replicas, they can receive the forwarded messages from other honest replicas and form a block certificate before *timeoutCertificate* expires.

424 **5.2 BoundBFT’s synchrony bound**

425 In this section, we experimentally determine the value for BoundBFT’s synchrony bound  $\Delta$ :  
 426 We ran BoundBFT in the presence of malicious replicas and determined  $\Delta$  that gives enough  
 427 confidence that BoundBFT’s correctness will not be compromised. Initially, we set  $\Delta$  to  
 428 1250 ms (99.99%), the synchronous bound from [30], and gradually decreased it to the point  
 429 where we started to observe BoundBFT’s agreement and progress violations. The complete  
 430 data for these experiments can be found in Appendix B. In the following, we comment on  
 431 the main takeaways.

432 When there is a single ( $f = 1$ ) or no ( $f = 0$ ) Byzantine replicas in the system, we did  
 433 not observe any agreement violations, even if we set  $\Delta$  as low as 50 ms—the average latency  
 434 between 80% of replicas in our system is higher than 50 ms. This shows that agreement  
 435 violations are highly unlikely if the number of Byzantine replicas is low, even if many messages  
 436 violate synchronous bounds. When the number of Byzantine replicas is  $f = 19$  (i.e., the  
 437 maximum number of Byzantine replicas partially synchronous protocols can tolerate), the  
 438 agreement violations were observed only when we lowered the  $\Delta$  to 50 ms. Moreover, even  
 439 with  $\Delta = 50$  ms, the BoundBFT’s agreement was violated in less than 10% of epochs in  
 440 which the attack was launched. However, to prevent agreement violations when the number  
 441 of Byzantine replicas is  $f = 29$  (i.e., the maximum BoundBFT can tolerate), we needed to  
 442 increase  $\Delta$  to 150ms and 300ms, for 1KB and 32KB block sizes, respectively. This makes  
 443 sense since to create a block certificate an honest replica needs to receive a vote from itself,  
 444 Byzantine replicas, and one honest replica. Importantly, even in this case, the resulting  $\Delta$   
 445 was 8 and 4 times lower than the initial one. Notice that when  $f = 29$ , partially synchronous  
 446 protocols will halt if Byzantine replicas remain silent.

447 Table 2 shows the  $\Delta$ s that resulted in no agreement violations and less than 5% of progress  
 448 violations for all considered setups. Namely, no two honest replicas committed on different  
 449 blocks for the same height and in less than 5% of epochs with an honest leader, some honest  
 450 replicas did not commit a new block. The progress violations can also be lowered to 0 %,  
 451 but this would require a slight increase in the  $\Delta$  chosen (see Appendix B). We believe this is  
 452 not necessary since progress violations can only lead to a slight decrease in performance and  
 453 do not affect agreement.

Attack type	1KB						32KB					
	$f=29$		$f=19$		$f=1$		$f=29$		$f=19$		$f=1$	
	$k_{min}$	$k_{max}$	$k_{min}$	$k_{max}$	$k_{min}$	$k_{max}$	$k_{min}$	$k_{max}$	$k_{min}$	$k_{max}$	$k_{min}$	$k_{max}$
<i>EQUIVOCATION</i>	150	150	100	100	100	100	150	300	150	150	100	100
<i>AMNESIA</i>	150	150	150	150	100	100	300	300	150	150	100	100
<i>EQUIVOCATION-CERTIFICATE</i>	150	150	150	150	100	100	300	300	150	150	100	100
<i>BLAME-CERTIFICATE</i>	150	150	150	150	100	100	150	150	150	150	100	100
<i>BLAME</i>	150		150		100		300		150		100	
<i>NO ATTACK</i>			100						100			

■ **Table 2** The  $\Delta$  in ms BoundBFT must adopt to achieve 0% of Agreement and < 5% of Progress violations under a specific attack. The table shows data for the setup of 60 replicas, 1KB and 32KB block sizes and different number of Byzantine replicas ( $f$ ). Additionally, for attacks that partition honest replicas in two subsets, it shows results when Byzantine replicas divide honest replicas into the two smallest subsets ( $k = k_{min} = 1$ ) and the two largest subsets ( $k = k_{max} = n - f/2$ ).

454 **5.3 Performance**


455 In this section, we compare BoundBFT to state-of-the-art partially synchronous and syn-  
 456 chronous protocols. In the partially synchronous model, we choose Tendermint [5] and  
 457 HotStuff-2 [32], the most recent protocol of the HotStuff family [39]. In the synchronous  
 458 system model, we consider Sync HotStuff [2, 3]. We limited our evaluation to Byzantine

459 consensus protocols with a rotating leader: a new leader is elected when the protocol changes  
 460 epoch (or round or view) as part of the normal execution, not only when the leader fails.  
 461 These protocols are preferred in blockchain systems since they provide better fairness and  
 462 censorship resistance than protocols with a stable leader (e.g., PBFT) [3].

463 To generate a load in our experiments, we equip every replica with a built-in client that  
 464 generates transactions in advance and stores them in a local pool. When a replica is a leader  
 465 in an epoch, it takes transactions from the pool and forms a block. The block size defines  
 466 the number of transactions taken from the pool. This design leaves the mempool (i.e., the  
 467 part of a blockchain responsible for propagating client transactions across the system) out of  
 468 the discussion as different systems may implement it in different ways. Consequently, the  
 469 latencies we report in the paper represent consensus latencies (i.e., the time the leader of an  
 470 epoch needs to commit a block). Throughput is computed as the rate of committed blocks  
 471 per time unit. Every point in the graphs is an average of 3 runs. We ran each experiment for  
 472 5 minutes.

### 473 5.3.1 Latency

474 BoundBFT and Sync HotStuff wait for  $2\Delta$  (BoundBFT's  $timeoutCommit(e)$ ) before com-  
 475 mitting a block in epoch  $e$ . Consequently, their latency is directly affected by the chosen  
 476 synchrony bound. For BoundBFT, we adopt the synchrony bound based on the experiments  
 477 from the previous section (see Table 2). Conversely, for Sync HotStuff, we use the  $\Delta$  from  
 478 [30].



479 **Figure 2** Latency (left) and throughput (right) comparison for all protocols for 1 KB and 32 KB  
 480 block sizes in a system with 60 replicas.

481 We measured latencies in a system with 60 replicas with 1KB and 32KB block sizes.  
 482 Figure 2 (left) shows the average latency computed by epoch leaders. First, we notice the  
 483 significant impact of BoundBFT's synchrony bound on latency. Namely, BoundBFT achieves  
 484  $5.4\times$  and  $3.4\times$  lower latency than Sync HotStuff with 1KB and 32KB block sizes, respectively.  
 485 Second, BoundBFT's latency is in between the latencies of partially synchronous protocols. It  
 486 is  $1.3\times$  and  $1.8\times$  lower than HotStuff-2's latency and  $1.4\times$  and  $2\times$  higher than Tendermint's  
 487 for small and large blocks, respectively. HotStuff-2 has higher latency due to its linear  
 488 communication pattern, which requires five communication steps, while Tendermint has  
 489 quadratic communication and commits a block in only three communication steps. We can  
 also see that HotStuff-2 has the most significant standard deviation; we attribute this to the  
 lower redundancy due to its linear communication pattern.

### 490 5.3.2 Throughput

491 BoundBFT and Sync HotStuff [2, 3] use *pipelining* to limit the impact of  $\Delta$  on throughput,  
 492 which allows the leader to propose a block  $B_{k+1}$  that extends block  $B_k$  after receiving

## 28:14 How robust are synchronous consensus protocols?

493  $C_e(B_k)$ , i.e., before committing block  $B_k$ . This way, protocols can order multiple blocks in  
494 parallel, and the throughput is unaffected by  $\Delta$ . This technique was initially introduced in  
495 HotStuff [39], and we implemented a pipelined version of HotStuff-2. Adapting pipelining to  
496 Tendermint is more complex and out of the scope of this paper. Moreover, notice that we  
497 compare BoundBFT a pipelined partially synchronous protocol, HotStuff-2.

498 We evaluated throughput in a system with 60 replicas and various block sizes (see Figure 2  
499 (right)). BoundBFT and Sync HotStuff have similar throughput, as they both start ordering  
500 the next block after receiving a certificate for the previous block. Moreover, they outperform  
501 partially synchronous protocols for all block sizes considered, reaching throughput more than  
502  $2\times$  higher than Tendermint's for all block sizes. The reason behind this is that Tendermint  
503 does not use pipelining. They also perform better, from  $1.4\times$  to  $3\times$ , than HotStuff-2, a  
504 partially synchronous protocol with pipelining. This is because even though both protocols  
505 start ordering the next block after collecting a certificate for the previous block, the certificate  
506 in HotStuff-2 requires votes from a two-third majority of replicas, while in BoundBFT the  
507 votes from the majority are enough.

### 508 5.4 Summary

509 In this section, we summarize the main takeaways of our evaluation.

- 510 ■ We did not observe any agreement violations when the number of Byzantine replicas  
511 was 0 and 1. Synchrony violations combined with a minority or one-third of colluded  
512 Byzantine replicas resulted in agreement violations.
- 513 ■ BoundBFT can use a  $\Delta$  that is  $4\times$  to  $8\times$  smaller than the conservative 99.99%  $\Delta$  [30],  
514 allowing it to improve latency from  $\approx 3.4\times$  to  $5.4\times$ .
- 515 ■ BoundBFT's  $\Delta$  is big enough to ensure correctness with high probability when the system  
516 is under attack.
- 517 ■ BoundBFT achieves from  $1.4\times$  to  $3\times$  higher throughput and comparable latency to  
518 state-of-the-art partially synchronous protocols.

## 519 6 Related work

520 The synchronous system model in its purest form requires that every message sent in the  
521 system be delivered within some known synchrony bound  $\Delta$ . Some protocols have been  
522 designed for this model such as Dfinity [24] and Sync HotStuff [2, 3]. These protocols are  
523 optimal in terms of resilience [15, 26] as they can tolerate up to a minority of faulty processes  
524 (i.e.,  $f < n/2$ , where  $f$  is the number of Byzantine processes out of  $n$ ). However, deploying  
525 these protocols in a blockchain environment usually comes with conservative  $\Delta$  that should  
526 ensure with high probability that the system is synchronous. Consequently, these protocols  
527 perform poorly. BoundBFT belongs to this family of protocols and is similar to the rotating  
528 version of Sync HotStuff [3]. The two protocols share similar behavior in the common-case  
529 but they have different epoch synchronization mechanisms. In this paper, we show that we  
530 do not need to use conservative bounds when deploying BoundBFT and as a consequence it  
531 can deliver reasonably good performance.

532 Many deterministic consensus protocols assume the partially synchronous system model  
533 [13]. The model allows consensus protocols that need synchrony only for liveness but not for  
534 safety. More precisely, the partially synchronous system model ensures that after some point,  
535 usually referred to as GST (Global Stabilization Time), messages exchanged among honest  
536 processes will be delivered within an unknown synchrony bound. A fundamental limitation

537 of the partially synchronous system model is that any consensus protocol designed for this  
538 model can tolerate up to one third of faulty processes (i.e.,  $f < n/3$ ) [13]. The first practical  
539 representative of these protocols is PBFT [7], and more recently, in the blockchain context,  
540 Tendermint [5], HotStuff [39], HotStuff-2 [32] and ICC [6], to name a few.

541 Guo et al. [23] introduced the “weak synchronous model” (called mobile sluggish model  
542 in [2] for consistency with other works in the literature). The idea is to distinguish between  
543 *prompt* processes, those that respect synchrony bounds, from *sluggish* processes, those for  
544 whom messages may violate synchrony bounds. Moreover, the set of sluggish processes may  
545 change over time. The mobile sluggish model weakens the synchronous model and may result  
546 in more practical protocols. However, this is true only in situations when the actual number  
547 of faulty processes in the system is less than the minority (e.g., [2, 8]).

548 XFT [30] is based on the observation that typical BFT consensus protocols assume a  
549 powerful adversary that fully controls malicious processes and the network between honest  
550 processes, which is unrealistic. We share this view, and consider an adversary that cannot  
551 control the network between honest processes. XFT differentiates between three types of  
552 faulty processes: crash, Byzantine, and partitioned (i.e., processes that cannot exchange  
553 messages with other honest processes within the known synchrony bound). It ensures progress  
554 as long as the total number of faulty processes in the system is lower than  $f < n/2$ . In  
555 other words, XFT assumes a majority of honest replicas that can communicate timely. Since  
556 selecting a quorum of  $f + 1$  responsive replicas out of  $n$  replicas requires an exponential  
557 number of attempts, the solution is practical when  $f$  is small.

558 An alternative way to increase the resilience of consensus protocols while assuming  
559 a partially synchronous system is to rely on trusted hardware (e.g., Trusted Execution  
560 Environment). This idea was introduced in A2M [9] and explored in many works (e.g., [10,  
561 29, 37, 29, 38, 11]). These techniques have the disadvantage of requiring blockchain servers  
562 to be equipped with special hardware, which is not the case with BoundBFT.

563 The consensus protocols presented above assume some synchrony to circumvent the FLP  
564 impossibility result [16]. Alternatively, there are many asynchronous consensus protocols that  
565 rely on randomization to solve consensus [4, 34, 14, 20, 36, 33, 22, 18, 21]. Asynchronous  
566 protocols are robust since they do not assume any synchrony, but they provide probabilistic  
567 guarantees and perform worse than partially synchronous and synchronous protocols. Con-  
568 sequently, some protocols use a simpler leader-based deterministic protocol to improve the  
569 latency in good cases [19, 27, 35, 31].

## 570 **7 Conclusion**

571 In this paper, we have shown how we can circumvent major performance drawbacks of  
572 synchronous consensus protocols by choosing protocol-specific time bounds instead of con-  
573 servative model-specific bounds. Instead of ensuring that all messages are received within  
574 synchronous bounds with high probability, we have analyzed protocol semantics and potential  
575 correctness violations in case of synchrony violations and Byzantine attacks, and have shown  
576 how to select a time bound that does not hurt correctness. As a showcase, we designed  
577 BoundBFT, a new Byzantine fault-tolerant synchronous consensus protocol, and have shown  
578 experimentally that BoundBFT withstands synchrony bound violations under attack and  
579 outperforms traditional synchronous consensus protocols. Furthermore, BoundBFT achieves  
580 similar latency and better throughput than state-of-the-art partially synchronous protocols,  
581 offering higher resilience.

## 582 — References

- 583 1 Libp2p. <https://libp2p.io>. [Accessed 2023-01-12].
- 584 2 Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Maofan Yin. Sync HotStuff: Simple and practical synchronous state machine replication. In *2020 IEEE Symposium on Security and Privacy (SP)*, may 2020. doi:10.1109/sp40000.2020.00044.
- 585 3 Ittai Abraham, Kartik Nayak, and Nibesh Shrestha. Optimal good-case latency for rotating leader synchronous BFT. In *25th International Conference on Principles of Distributed Systems (OPODIS 2021)*, pages 27:1–27:19, 2022. doi:10.4230/LIPICS.OPODIS.2021.27.
- 586 4 Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). In Robert L. Probert, Nancy A. Lynch, and Nicola Santoro, editors, *Proceedings of the Second Annual ACM Symposium on Principles of Distributed Computing, Montreal, Quebec, Canada, August 17-19, 1983*, PODC’83, pages 27–30. ACM, aug 1983. doi:10.1145/800221.806707.
- 587 5 Ethan Buchman, Jae Kwon, and Zarko Milosevic. The latest gossip on BFT consensus. jul 2018. URL: <http://arxiv.org/abs/1807.04938>, arXiv:1807.04938[cs.DC], doi:10.48550/arXiv.1807.04938.
- 588 6 Jan Camenisch, Manu Drijvers, Timo Hanke, Yvonne-Anne Pignolet, Victor Shoup, and Dominic Williams. Internet computer consensus. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC’22, pages 81–91, jul 2022. doi:10.1145/3519270.3538430.
- 589 7 Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In Margo I. Seltzer and Paul J. Leach, editors, *Proceedings of the Third USENIX Symposium on Operating Systems Design and Implementation (OSDI), New Orleans, Louisiana, USA, February 22-25, 1999*, OSDI ’99, pages 173–186, USA, 1999. USENIX Association. URL: <https://dl.acm.org/citation.cfm?id=296824>.
- 590 8 T-H. Hubert Chan, Rafael Pass, and Elaine Shi. Pili: An extremely simple synchronous blockchain. Cryptology ePrint Archive, Paper 2018/980, 2018. <https://eprint.iacr.org/2018/980>. URL: <https://eprint.iacr.org/2018/980>.
- 591 9 Byung-Gon Chun, Petros Maniatis, Scott Shenker, and John Kubiatowicz. Attested append-only memory: making adversaries stick to their word. In Thomas C. Bressoud and M. Frans Kaashoek, editors, *Proceedings of the 21st ACM Symposium on Operating Systems Principles 2007, SOSP 2007, Stevenson, Washington, USA, October 14-17, 2007*, volume 41, pages 189–204. ACM, oct 2007. doi:10.1145/1294261.1294280.
- 592 10 Miguel Correia, Nuno Ferreira Neves, and Paulo Veríssimo. How to tolerate half less one byzantine nodes in practical distributed systems. In *23rd International Symposium on Reliable Distributed Systems (SRDS 2004), 18-20 October 2004, Florianopolis, Brazil*, pages 174–183. IEEE Computer Society, 08 2004. doi:10.1109/RELDIS.2004.1353018.
- 593 11 Jérémie Decouchant, David Kozhaya, Vincent Rahli, and Jiangshan Yu. DAMYSUS: streamlined bft consensus leveraging trusted components. In *Proceedings of the Seventeenth European Conference on Computer Systems*. ACM, mar 2022. doi:10.1145/3492321.3519568.
- 594 12 Danny Dolev and H. Strong. Authenticated algorithms for byzantine agreement. *SIAM J. Comput.*, 12:656–666, 11 1983. doi:10.1137/0212045.
- 595 13 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2):288–323, apr 1988. doi:10.1145/42282.42283.
- 596 14 Pesech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM Journal on Computing*, 26(4):873–933, aug 1997. doi:10.1137/s0097539790187084.
- 597 15 Michael J. Fischer, Nancy A. Lynch, and Michael Merritt. Easy impossibility proofs for distributed consensus problems. *Distributed Comput.*, 1(1):26–39, jan 1986. doi:10.1007/BF01843568.

- 632 16 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed  
633 consensus with one faulty process. *Journal of ACM*, 32(2):374–382, apr 1985. doi:10.1145/  
634 3149.214121.
- 635 17 Matthias Fitzi. *Generalized communication and security models in Byzantine agreement*. PhD  
636 thesis, ETH Zurich, Zürich, Switzerland, 3 2003. Reprint as vol. 4 of ETH Series in Information  
637 Security and Cryptography, ISBN 3-89649-853-3, Hartung-Gorre Verlag, Konstanz, 2003. URL:  
638 <https://d-nb.info/967397375>.
- 639 18 Yingzi Gao, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Dumbong:  
640 Fast asynchronous BFT consensus with throughput-oblivious latency. In Heng Yin,  
641 Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC*  
642 *Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA,*  
643 *November 7-11, 2022*, CCS '22, pages 1187–1201, New York, NY, USA, 2022. ACM. doi:  
644 10.1145/3548606.3559379.
- 645 19 Rati Gelashvili, Lefteris Kokoris-Kogias, Alberto Sonnino, Alexander Spiegelman, and Zhuolun  
646 Xiang. Jolteon and ditto: Network-adaptive efficient consensus with asynchronous fallback.  
647 *CoRR*, abs/2106.10362, 06 2021. URL: <https://arxiv.org/abs/2106.10362>, arXiv:2106.  
648 10362, doi:10.48550/arXiv.2106.10362.
- 649 20 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nikolai Zeldovich. Algorand:  
650 Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on*  
651 *Operating Systems Principles, Shanghai, China, October 28-31, 2017*, SOSP '17, pages 51–68,  
652 New York, NY, USA, 2017. ACM. doi:10.1145/3132747.3132757.
- 653 21 Bingyong Guo, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Speeding  
654 dumbong: Pushing asynchronous bft closer to practice. *Cryptology ePrint Archive*, Paper  
655 2022/027, 2022. <https://eprint.iacr.org/2022/027>. URL: [https://eprint.iacr.org/](https://eprint.iacr.org/2022/027)  
656 2022/027.
- 657 22 Bingyong Guo, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Dumbo: Faster  
658 asynchronous BFT protocols. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni  
659 Vigna, editors, *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications*  
660 *Security, Virtual Event, USA, November 9-13, 2020*, CCS '20, pages 803–818, New York, NY,  
661 USA, 2020. ACM. doi:10.1145/3372297.3417262.
- 662 23 Yue Guo, Rafael Pass, and Elaine Shi. Synchronous, with a chance of partition tol-  
663 erance. In *Advances in Cryptology – CRYPTO 2019*, pages 499–529, aug 2019. doi:  
664 10.1007/978-3-030-26948-7\_18.
- 665 24 Timo Hanke, Mahnush Movahedi, and Dominic Williams. DFINITY technology overview  
666 series, consensus system. *CoRR*, abs/1805.04548, may 2018. URL: [http://arxiv.org/abs/](http://arxiv.org/abs/1805.04548)  
667 1805.04548, arXiv:1805.04548, doi:10.48550/arXiv.1805.04548.
- 668 25 Bert Hubert, Gregory Maxwell, Martijn van Oosterhout, Remco van Mook, Paul B. Schroeder,  
669 et al. Linux advanced routing & traffic control HOWTO. <https://lartc.org/lartc.html>,  
670 2002. [Accessed 2024-22-5].
- 671 26 Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byz-  
672 antine agreement. In *Advances in Cryptology - CRYPTO 2006, 26th Annual International*  
673 *Cryptology Conference*, volume 4117 of *Lecture Notes in Computer Science*, pages 445–  
674 462. Springer, 2006. URL: <https://iacr.org/archive/crypto2006/41170440/41170440.pdf>,  
675 doi:10.1007/11818175\_27.
- 676 27 Klaus Kursawe and Victor Shoup. Optimistic asynchronous atomic broadcast. *IACR Cryptol.*  
677 *ePrint Arch.*, page 22, 2001. <https://eprint.iacr.org/2001/022>. URL: [http://eprint.](http://eprint.iacr.org/2001/022)  
678 [iacr.org/2001/022](http://eprint.iacr.org/2001/022).
- 679 28 Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The byzantine generals problem.  
680 *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, jul 1982. doi:10.1145/357172.357176.
- 681 29 Dave Levin, John R Douceur, Jacob R Lorch, and Thomas Moscibroda. TrInc: Small  
682 trusted hardware for large distributed systems. In *NSDI*, volume 9, pages 1–14, 2009. URL:  
683 [http://www.usenix.org/events/nsdi09/tech/full\\_papers/levin/levin.pdf](http://www.usenix.org/events/nsdi09/tech/full_papers/levin/levin.pdf).

## 28:18 How robust are synchronous consensus protocols?

- 684 30 Shengyun Liu, Paolo Viotti, Christian Cachin, Vivien Quéma, and Marko Vukolic. XFT:  
685 practical fault tolerance beyond crashes. In Kimberly Keeton and Timothy Roscoe, editors,  
686 *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016, Sa-*  
687 *vannah, GA, USA, November 2-4, 2016*, OSDI'16, pages 485–500, USA, 2016. USENIX  
688 Association. URL: [https://www.usenix.org/conference/osdi16/technical-sessions/](https://www.usenix.org/conference/osdi16/technical-sessions/presentation/liu)  
689 [presentation/liu](https://www.usenix.org/conference/osdi16/technical-sessions/presentation/liu).
- 690 31 Yuan Lu, Zhenliang Lu, and Qiang Tang. Bolt-dumbo transformer: Asynchronous consensus  
691 as fast as pipelined BFT. *CoRR*, abs/2103.09425, 2021. URL: [https://arxiv.org/abs/2103.](https://arxiv.org/abs/2103.09425)  
692 [09425](https://arxiv.org/abs/2103.09425), arXiv:2103.09425, doi:10.48550/arXiv.2103.09425.
- 693 32 Dahlia Malkhi and Kartik Nayak. Extended abstract: Hotstuff-2: Optimal two-phase responsive  
694 BFT. *IACR Cryptol. ePrint Arch.*, page 397, 2023. URL: [https://eprint.iacr.org/2023/](https://eprint.iacr.org/2023/397)  
695 [397](https://eprint.iacr.org/2023/397).
- 696 33 Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of BFT  
697 protocols. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers,  
698 and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and*  
699 *Communications Security, Vienna, Austria, October 24-28, 2016*, CCS '16, pages 31–42, New  
700 York, NY, USA, 2016. ACM. doi:10.1145/2976749.2978399.
- 701 34 Michael O. Rabin. Randomized byzantine generals. In *24th Annual Symposium on Foundations*  
702 *of Computer Science (sfcs 1983)*, nov 1983. doi:10.1109/sfcs.1983.48.
- 703 35 HariGovind V. Ramasamy and Christian Cachin. Parsimonious asynchronous byzantine-fault-  
704 tolerant atomic broadcast. *IACR Cryptol. ePrint Arch.*, page 82, 2006. [https://eprint.iacr.](https://eprint.iacr.org/2006/082)  
705 [org/2006/082](https://eprint.iacr.org/2006/082). URL: <http://eprint.iacr.org/2006/082>.
- 706 36 Dmitry Tanana. Avalanche blockchain protocol for distributed computing security. In *2019*  
707 *IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom)*.  
708 IEEE, jun 2019. doi:10.1109/blackseacom.2019.8812863.
- 709 37 Giuliana Santos Veronese, Miguel Correia, Alysson Neves Bessani, Lau Cheuk Lung, and Paulo  
710 Verissimo. Efficient byzantine fault-tolerance. *IEEE Transactions on Computers*, 62(1):16–30,  
711 2013. doi:10.1109/TC.2011.221.
- 712 38 Sravya Yandamuri, Ittai Abraham, Kartik Nayak, and Michael K. Reiter. Communication-  
713 efficient bft protocols using small trusted hardware to tolerate minority corruption. *Cryptology*  
714 *ePrint Archive*, Paper 2021/184, 2021. <https://eprint.iacr.org/2021/184>. URL: <https://eprint.iacr.org/2021/184>.
- 715 39 Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan Gueta, and Ittai Abraham.  
716 HotStuff: BFT consensus with linearity and responsiveness. In *Proceedings of the 2019*  
717 *ACM Symposium on Principles of Distributed Computing, PODC'19*, pages 347–356, jul 2019.  
718 doi:10.1145/3293611.3331591.  
719

## A Appendix: Algorithm correctness

### A.1 Proof of correctness

This section presents the proof that BoundBFT satisfies all the properties of blockchain consensus protocol (Section 2.2).

► **Lemma 1.** *Every honest replica always moves to the next epoch.*

**Proof.** Assume for contradiction that an honest replica  $r$  remains in some epoch  $e$  indefinitely. This would imply that  $r$  did not generate any of the certificates  $C_e(B_k)$ ,  $C_e(\text{BLAME})$ , or  $C_e(\text{EQUIV})$ . However, each honest replica starts the timer,  $\text{timeoutCertificate}(e)$ , upon entering epoch  $e$  (line 2 in Algorithm 2). When this timeout expires, if an honest replica has not received any certificate, it broadcasts the BLAME message (lines 3–5 in Algorithm 2). Consequently, if no certificate is formed before the  $\text{timeoutCertificate}(e)$  expires, all honest replicas will broadcast the BLAME message, leading to the formation of the blame certificate  $C_e(\text{BLAME})$ . This contradicts our assumption, thus proving that every honest replica moves to the next epoch. ◀

► **Lemma 2.** *If an honest replica starts epoch  $e$  at time  $t$ , then all honest replicas start epoch  $e$  by time  $t + \Delta$ .*

**Proof.** Suppose an honest replica  $r$  starts epoch  $e$  at time  $t$ . This implies that  $r$  receives and broadcasts  $C_{e-1}(B_k)$  at time  $t$  (lines 35–36 in Algorithm 1), or at time  $t - \text{timeoutEpochChange}(2\Delta)$ ,  $r$  receives and broadcasts  $C_{e-1}(\text{BLAME})$  or  $C_{e-1}(\text{EQUIV})$  (lines 17 and 21 in Algorithm 2). Messages with certificates will arrive within  $\Delta$  time. Consequently, in the former case, all honest replicas receive  $C_{e-1}(B_k)$  by time  $t + \Delta$  and start epoch  $e$ . In the latter case, all honest replicas receive  $C_{e-1}(\text{BLAME})$  or  $C_{e-1}(\text{EQUIV})$  by time  $t - \Delta$  and within  $2\Delta$  they start epoch  $e$ , ensuring that all honest replicas start epoch  $e$  by time  $t + \Delta$ . ◀

► **Theorem 3.** (*Epoch synchronization*) *All honest replicas continuously move through epochs, with each replica starting a new epoch within  $\Delta$  time of any other honest replica.*

**Proof.** We prove this theorem by combining Lemma 1 and Lemma 2.

First, from Lemma 1, we know that every honest replica always moves to the next epoch. This ensures that no honest replica remains stuck in any epoch indefinitely.

Second, from Lemma 2, we know that if an honest replica starts epoch  $e$  at time  $t$ , then all honest replicas start epoch  $e$  by time  $t + \Delta$ . This guarantees that all honest replicas start each epoch within  $\Delta$  time of each other.

Combining these two results, we can conclude that all honest replicas continuously move through epochs, with each replica initiating a new epoch within  $\Delta$  time of any other honest replica. ◀

► **Lemma 4.** *If an honest replica directly commits block  $B_k$  in epoch  $e$ , then (i) no block different than  $B_k$  can be certified in epoch  $e$ , and (ii) every honest replica locks on block  $B_k$  in epoch  $e$ .*

**Proof.** Suppose an honest replica  $r$  directly commits  $B_k$  in epoch  $e$  at time  $t$  (line 40 in Algorithm 1). This means that at time  $t - 2\Delta$ ,  $r$  received  $C_e(B_k)$ , locked on it, and started  $\text{timeoutCommit}(e)$  (lines 26–32 in Algorithm 1). Moreover, replica  $r$  forwarded all messages representing  $C_e(B_k)$  (lines 25 and 35 in Algorithm 1) so all honest replicas received these messages in  $\Delta$  time, by time  $t - \Delta$ .

763 For part (i), assume for a contradiction that some honest replica  $p$  received and voted  
 764 in epoch  $e$  for the block  $B_l \neq B_k$ . Since every honest replica votes only once,  $p$  must have  
 765 received a proposal for  $B_l$  before receiving a proposal message for  $B_k$ , at some time  $t_1 < t - \Delta$ .  
 766 As a result,  $p$  forwards the propose message for  $B_l$  at time  $t_1$  (line 25 in Algorithm 1). Replica  
 767  $r$  will receive this message by time  $t_1 + \Delta$ , that is, before  $t$ . Since these two propose messages  
 768 lead to a  $C_e(\text{EQUIV})$  certificate,  $p$  would not commit (lines 9–12 and 15 in Algorithm 2), a  
 769 contradiction. Therefore, property (i) holds since no honest replica votes for a block different  
 770 than  $B_k$ ; otherwise replica  $r$  would not commit.

771 For part (ii), we know by (i) that if replica  $r$  directly commits in epoch  $e$ , there is not any  
 772 possible  $C_e(B_l) \neq C_e(B_k)$ . So, we need to prove that every honest replica  $p$  receives  $C_e(B_k)$   
 773 before receiving  $C_e(\text{BLAME})$  or  $C_e(\text{EQUIV})$ . For a contradiction, assume that  $p$  receives  
 774  $C_e(\text{BLAME})$  or  $C_e(\text{EQUIV})$  before receiving  $C_e(B_k)$ . This must happen at time  $t_1 < t - \Delta$  as  
 775  $p$  receives  $C_e(B_k)$  by time  $t - \Delta$ . After receiving  $C_e(\text{BLAME})$  or  $C_e(\text{EQUIV})$ ,  $p$  broadcasts  
 776 them (line 17 in Algorithm 2). So,  $p$  broadcasts  $C_e(\text{BLAME})$  or  $C_e(\text{EQUIV})$  at time  $t_1$  and  $r$   
 777 receives them by time  $t_1 + \Delta$ . Since  $t > t_1 + \Delta$  replica  $r$  will not commit  $B_k$ , a contradiction.  
 778  $\blacktriangleleft$

779 **► Lemma 5.** *If  $B_k$  is the only certified block in epoch  $e$  and  $f + 1$  honest replicas lock on*  
 780 *block  $B_k$  in epoch  $e$  ( $\text{lockedBlock} = B_k$  and  $\text{lockedBC} = C_e(B_k)$ ), then in all epochs  $e' > e$ ,*  
 781 *they vote only for blocks extending  $B_k$ , or they blame the proposer.*

782 **Proof.** The proof proceeds by induction on the epoch number.

783 **Base step** ( $e' = e + 1$ ): Let  $C$  denote the set of  $f + 1$  honest replicas. The replicas in  
 784 set  $C$  do not vote for proposals that do not extend blocks certified in epochs higher than or  
 785 equal to their  $\text{lockedBC}$  (line 22 in Algorithm 1). As a result, when  $\text{timeoutCertificate}(e')$   
 786 expires, no block certificate will be formed since no honest replica has voted, causing honest  
 787 replicas to blame the proposer by sending  $\langle \text{BLAME}, e' \rangle_*$  message. Therefore, the lemma holds  
 788 for the base step since honest replicas vote only for a block if it extends  $\text{lockedBlock}$ .

789 **Induction step** ( $e' \rightarrow e' + 1$ ): Assume that no replica in set  $C$  has voted for a block  
 790 not extending  $B_k$  until epoch  $e' + 1$ . We now show that the lemma holds for epoch  $e' + 1$ .  
 791 Since replicas from the set  $C$  vote for blocks extending  $B_k$  or blame the proposer in epochs  
 792  $e \leq e'' \leq e'$ , no block  $B_l$  not extending  $B_k$  can receive  $f + 1$  votes in those epochs. Therefore,  
 793 for all processes in set  $C$ ,  $\text{lockedBlock} = B_{k'}$  and  $\text{lockedBC.epoch} \geq e$ , where  $B_{k'} = B_k$  or  
 794  $B_{k'}$  extends  $B_k$ . Assume, for the sake of contradiction, that a process  $p$  in set  $C$  votes in  
 795 epoch  $e' + 1$  for a block not extending  $B_k$ . An honest replica will not vote for a block not  
 796 extending its  $\text{lockedBlock}$  (line 22 in Algorithm 1), leading to a contradiction. Hence, the  
 797 lemma holds for epoch  $e' + 1$  as well.  
 798  $\blacktriangleleft$

799 **► Lemma 6.** *If an honest replica directly commits block  $B_k$  in epoch  $e$ , then any block  $B_l$*   
 800 *that is certified in epoch  $e' > e$  must extend  $B_k$ .*

801 **Proof.** The proof follows directly from Lemmas 4 and 5. More precisely, if an honest replica  
 802 directly commits block  $B_k$  in epoch  $e$ , by Lemma 4, we know that  $f + 1$  honest replicas (set  
 803  $C$ ) lock on block  $B_k$  in epoch  $e$  and  $B_k$  is the only certified block in epoch  $e$ . Consequently,  
 804 by Lemma 5, replicas from  $C$  vote only for the blocks extending block  $B_k$  in epochs  $e' > e$ .  
 805 Therefore, no block  $B_l$  that does not extend  $B_k$  can collect  $f + 1$  votes and thus cannot be  
 806 certified in any epoch  $e' > e$ .  
 807  $\blacktriangleleft$

807 **► Theorem 7. (Agreement)** *No two honest replicas commit different blocks at the same*  
 808 *height.*

809 **Proof.** Suppose, for the sake of contradiction, that two distinct blocks  $B_k$  and  $B'_k$  are  
 810 committed for the same height  $k$ . Assume that  $B_k$  is committed as a result of  $B_l$  being  
 811 directly committed in epoch  $e$  and  $B'_k$  is committed as a result of  $B_{l'}$  being directly committed  
 812 in epoch  $e'$ . Without loss of generality, assume  $l < l'$ . Note that all directly committed  
 813 blocks are certified. This is true because in order to start  $timeoutCommit(e)$  for block  $B_k$ ,  
 814 a replica needs to receive  $C_e(B_k)$  (lines 26–32 in Algorithm 1). By Lemma 6,  $B_{l'}$  extends  $B_l$ .  
 815 Therefore,  $B_k = B'_k$ , which contradicts the assumption that  $B_k$  and  $B'_k$  are distinct. Hence,  
 816 no two honest replicas can commit different blocks at the same height.

817

818 ► **Lemma 8.** *If an honest replica  $r$  locks on a block  $B_k$  in epoch  $e$ , no honest replica starts  
 819 epoch  $e + 1$  before updating  $validBC$  to  $C_e(B_l)$ , where  $B_l$  does not need to be equal to  $B_k$ .*

820 **Proof.** Assume that an honest replica  $r$  locks on a block  $B_k$  in epoch  $e$  at time  $t$ . This  
 821 implies  $r$  receives  $C_e(B_k)$  at time  $t$  and does not receive  $C_e(BLAME)$  or  $C_e(EQUIV)$  before  
 822 that. Since all messages representing  $C_e(B_k)$  are broadcast (lines 25 and 35 in Algorithm 1),  
 823 all honest replicas receive these messages by time  $t + \Delta$ .

824 Suppose for a contradiction that some honest replica  $p$  starts the epoch  $e + 1$  before  
 825 receiving  $C_e(B_k)$  or some other  $C_e(B_l)$ , in other words at time  $t_1 < t + \Delta$ . This means it  
 826 had received  $C_e(BLAME)$  or  $C_e(EQUIV)$  and broadcast messages representing them at time  
 827  $t_1 - 2\Delta$  (line 17 in Algorithm 2). Consequently, replica  $r$  receives  $C_e(BLAME)$  or  $C_e(EQUIV)$   
 828 by time  $t_1 - \Delta$  and as  $t > t_1 - \Delta$ , it does not lock on  $B_k$ , a contradiction.

829

830 ► **Corollary 9.** *Every honest replica starts epoch  $e$  with  $validBC$  that is at least as recent as  
 831 any certificate any honest replica locks on in any epoch  $e' < e$ .*

832 **Proof.** Suppose that the last epoch in which some honest replica locks on a block is epoch  
 833  $e' < e$ . By Lemma 8, we know that all honest replicas update their  $validBC$  to some  
 834 certificate from the same epoch ( $e'$ ), before starting epoch  $e' + 1$ . From this and the fact  
 835 that no honest replica, in any of the following epochs ( $e' < e'' < e$ ), updates its  $validBC$  to  
 836 an older certificate (lines 33–34 in Algorithm 1), we see that this corollary holds. ◀

837 ► **Theorem 10.** (*Progress*) *All honest replicas keep committing new blocks.*

838 **Proof.** From Lemmas 1 and 2, we see that replicas proceed through epochs, each epoch  
 839 having a dedicated leader. If the leader of an epoch is Byzantine and does not propose any  
 840 block or proposes equivocating blocks, honest replicas will collect  $C_e(BLAME)$  or  $C_e(EQUIV)$   
 841 and move to the next epoch. Due to the round-robin leader election, there will be epochs  
 842 with honest leaders.

843 Consider an epoch  $e$  with an honest leader  $l$ . Let  $t$  be the time when the first honest  
 844 replica starts epoch  $e$ . By Lemma 2, all honest replicas enter epoch  $e$  by the time  $t + \Delta$ .  
 845 Therefore, by the time  $t + \Delta$  at the latest, an honest leader  $l$  broadcasts the proposal  
 846  $\langle PROPOSE, e, B_k, validBC_l \rangle_l$ . All honest replicas receive proposal by time  $t + 2\Delta$ . Since by  
 847 the Corollary 9,  $validBC_l$  is at least as recent as any  $lockedBC$  of any honest replica, all  
 848 honest replicas vote for the proposal. As a result, all honest replicas receive  $C_e(B_k)$  by  
 849 time  $t + 3\Delta$ . Since  $timeoutCertificate(e) > 3\Delta$ , no honest replica will send a  $\langle BLAME, e \rangle_*$   
 850 message in epoch  $e$ , and  $C_e(BLAME)$  cannot be formed. Furthermore, considering that replica  
 851  $l$  is honest, it does not equivocate, so no  $C_e(EQUIV)$  will be formed in epoch  $e$ . Consequently,  
 852 all honest replicas start  $timeoutCommit(e)$ , and when it expires, they commit  $B_k$  and all  
 853 its ancestors. This scenario will occur in every epoch with an honest leader, ensuring that all  
 854 honest replicas consistently commit new blocks across all such epochs. ◀

## 28:22 How robust are synchronous consensus protocols?

855 ► **Theorem 11.** (*External validity*) Every committed block satisfies the predefined `valid()`  
 856 predicate.

857 **Proof.** This follows directly from the requirement that every committed block must first be  
 858 certified (lines 26, 32, and 40 in Algorithm 1). This implies that at least one honest replica  
 859 accepted the block, meaning that `valid()` returned true for this block on at least one honest  
 860 replica (line 19 in Algorithm 1). ◀

## 861 A.2 The Byzantine protocol

### ■ Algorithm 3 The Byzantine protocol

---

```

1: Initialization:
2:    $e_p := 0$  ▷ the current epoch
3:    $validBC_p := nil$  ▷ the most recent block certificate the replica is aware of and...
4:    $validBlock_p := nil$  ▷ the block certified by  $validBC_p$ 
5:    $C := getAllHonestReplicas()$  ▷ the set of honest replicas
6:    $f := getNumberOfByzantineReplicas()$  ▷ the number of Byzantine replicas
7:    $attackType := getAttackType()$  ▷ the attack type the Byzantine replica launches
8:    $k := getTargetSize()$  ▷ the size of the two random sets of honest replicas that are under the attack
9: when bootstrapping do StartEpoch(0) ▷ the execution starts in epoch 0
10: Procedure StartEpoch( $e$ ) : ▷ upon starting the epoch...
11:    $e_p \leftarrow e$  ▷ the replica sets the current epoch, and...
12:   switch  $attackType$  : ▷ invokes the specific attack and pass the necessary arguments to it
13:     case EQUIVOCATION-ATTACK :
14:       LaunchEquivocationAttack( $e, validBlock, validBC, C, k, f$ )
15:     case AMNESIA-ATTACK :
16:       LaunchAmnesiaAttack( $e, validBlock, C, k, f$ )
17:     case BLAME-ATTACK :
18:       LaunchBlameAttack( $e, C, f$ )
19:     case EQUIVOCATION-CERTIFICATE-ATTACK :
20:       LaunchEquivocationCertificateAttack( $e, validBlock, validBC, C, k, f$ )
21:     case BLAME-CERTIFICATE-ATTACK :
22:       LaunchBlameCertificateAttack( $e, validBlock, validBC, C, k, f$ )
23:   when receive  $\langle PROPOSE, e, b, BC \rangle_l$  and  $f + 1$  distinct  $\langle VOTE, e, id(b) \rangle_*$  ▷ when replica receives a proposal...
24:     where  $e = e_p$  do ▷ and  $f + 1$  votes for it in the current epoch...
25:        $cert \leftarrow NewCert$  from  $f + 1$   $\langle VOTE, e, id(b) \rangle_*$  ▷ it forms a block certificate,...
26:        $validBC_p \leftarrow cert$  ▷ updates its  $validBC_p$  and  $validBlock_p$ ...
27:        $validBlock_p \leftarrow b$  ▷ to the most recent block, and...
28:       StartEpoch( $e + 1$ ) ▷ starts immediately the next epoch
29:   when receive  $\langle PROPOSE, e, b, BC \rangle_p$  and  $\langle PROPOSE, e, b', BC' \rangle_p$  ▷ upon receiving two different proposals...
30:     where  $e = e_p$  and  $p = leader(e)$  and  $b \neq b'$  do ▷ from the leader in the current epoch:
31:       start timeoutEpochChange( $e_p$ ) ▷ the replica triggers timeoutEpochChange
32:   when receive  $f + 1$  distinct  $\langle BLAME, e \rangle_*$  where  $e = e_p$  do ▷ upon receiving a blame certificate in the...
33:     start timeoutEpochChange( $e_p$ ) ▷ current epoch the replica triggers timeoutEpochChange
34:   when timeoutEpochChange( $e$ ) expires do ▷ when timeoutEpochChange expires and...
35:     if  $e = e_p$  then ▷ the replica is still in epoch  $e$ ...
36:       StartEpoch( $e_p + 1$ ) ▷ the replica starts the next epoch

```

---

862 Algorithm 3 presents the Byzantine replica protocol. Byzantine replicas proceed through  
 863 epochs in the same way as honest replicas. Namely, if they receive a block certificate, they  
 864 start the next epoch immediately (lines 23–28 in Algorithm 3), while if they receive a blame  
 865 or equivocation certificate they wait for *timeoutEpochChange* before starting the next epoch  
 866 (lines 29–36 in Algorithm 3). They do this to be synchronized with honest replicas so they  
 867 can launch the attack at the moment that maximizes the attack’s effectiveness. Moreover,  
 868 a Byzantine replica waits for *timeoutEpochChange* to update its *validBlock* and *validBC*  
 869 to the most recent values. As a result, when leader in an epoch, a Byzantine replica can  
 870 propose a valid block (i.e., an invalid block would be easily dismissed by honest replicas).

871 Algorithms 4 and 5 present the logic for attacks on BoundBFT’s agreement and progress,  
 872 respectively. We empower the attacks by assuming that Byzantine replicas know each other  
 873 and collude (Section 2.1): each Byzantine replica has private keys of all Byzantine replicas.

874 Therefore, a Byzantine replica can sign and send messages on behalf of other Byzantine  
875 replicas.

876 Upon starting an epoch, a Byzantine replica launches a specific attack, which is a  
877 parameter of an algorithm:

- 878 ■ *EQUIVOCATION-ATTACK*,
- 879 ■ *AMNESIA-ATTACK*,
- 880 ■ *BLAME-ATTACK*,
- 881 ■ *EQUIVOCATION-CERTIFICATE-ATTACK*, or
- 882 ■ *BLAME-CERTIFICATE-ATTACK*.

■ **Algorithm 4** Byzantine attacks on BoundBFT's agreement

---

```

1: Procedure LaunchEquivocationAttack(                               ▷ ***EQUIVOCATION-ATTACK***
2:   e, validBlock, validBC, honestReplicas, k, f)
3:   if isByzantineLeader(e) then                                     ▷ if the epoch leader is a Byzantine replica:
4:      $p_1, p_2 \leftarrow \text{generateTwoDifferentProposals}(e, \text{validBlock}, \text{validBC})$  ▷ the replica creates two distinct...
                                                                                                       ▷ proposals
5:      $v_1 \leftarrow \text{generateVoteMessages}(p_1, f)$                                ▷ then, the replica generates f VOTE messages for p1, and...
6:      $v_2 \leftarrow \text{generateVoteMessages}(p_2, f)$                                ▷ f VOTE messages for p2, one for each Byzantine replica
7:      $set_1, set_2 \leftarrow \text{getTwoRandomSets}(e, k, \text{honestReplicas})$            ▷ lastly, the replica divide honest replicas...
                                                                                                       ▷ in two random sets of size k
8:     send  $p_1$  and  $v_1$  to  $set_1$                                                ▷ then, it sends the first proposal and votes for it to the first set,  $set_1$ ,
9:     send  $p_2$  and  $v_2$  to  $set_2$                                                ▷ while, sending the second proposal and its votes to the second set,  $set_2$ 
10: Procedure LaunchAmnesiaAttack(                                       ▷ ***AMNESIA-ATTACK***
11:   e, validBlock, honestReplicas, k, f)
12:   if isByzantineLeader(e) then                                       ▷ if the epoch leader is a Byzantine replica:
13:      $p \leftarrow \text{generateAlternativeProposal}(e, \text{validBlock})$              ▷ the replica generates...
                                                                                                       ▷ an alternative proposal for validBlock
14:      $v \leftarrow \text{generateVoteMessages}(p, f)$                                ▷ then, the replica generates f VOTE messages for p,...
                                                                                                       ▷ one for each Byzantine replica
15:     send  $p$  and  $v$  to honestReplicas                                       ▷ the replica sends the proposal and...
                                                                                                       ▷ votes for it to the all honest replicas
16:   else                                                                 ▷ else, if the leader is an honest replica:
17:      $set_1, set_2 \leftarrow \text{getTwoRandomSets}(e, k, \text{honestReplicas})$            ▷ the replica divides honest replicas...
                                                                                                       ▷ in two random sets of size k
18:     upon receiving  $p = \langle \text{PROPOSE}, e, b, BC \rangle_l$                        ▷ then, when it receives the proposal...
19:     where  $l = \text{leader}(e)$  do                                             ▷ from epoch leader...
20:        $v \leftarrow \text{generateVoteMessages}(p, f)$                                ▷ it generates f VOTE messages for received proposal and...
21:        $b \leftarrow \text{generateBlameMessages}(e, f)$                                ▷ f BLAME messages for epoch e, one for each Byzantine replica
22:       send  $v$  to  $set_1$                                                        ▷ then, it sends votes to one subset of honest replicas,  $set_1$ ,...
23:       send  $b$  to  $set_2$                                                        ▷ while sending blames to the second subset of honest replicas,  $set_2$ 

```

---

883 All Byzantine replicas launch the same attack, not only the current epoch leader. This  
884 ensures that messages arrive at their destinations as fast as possible. So, if the malicious  
885 leader is far from some honest replica, the honest replica will receive attack messages from  
886 its closest Byzantine replica. For example, if the attack is *EQUIVOCATION-ATTACK*, each  
887 Byzantine replica will generate two proposals and votes for these proposals and send one  
888 proposal and its votes to one subset of honest replicas and another proposal and its votes to  
889 the other subset of honest replicas.

890 In attacks where Byzantine replicas divide honest replicas into two subsets and send  
891 different messages to them, the subsets are picked randomly. The size of these subsets is a para-  
892 meter of the algorithm, defined by *k*. If *k* is set to 2, function *getTwoRandomSets*(*e, k, set*)  
893 will return two different subsets, each containing two random elements. Byzantine replicas  
894 ensure they have the same subsets by using the current epoch number as a random number  
895 generator seed.

## 28:24 How robust are synchronous consensus protocols?

### Algorithm 5 Byzantine attacks on BoundBFT's progress

---

```

1: Procedure LaunchBlameAttack(▷ ***BLAME-ATTACK***
2:  $e, honestReplicas, f$ )
3: if isHonestLeader( $e$ ) then ▷ if the epoch leader is honest:
4:    $b \leftarrow generateBlameMessages(e, f)$  ▷ the replica generates  $f$  BLAME messages,...
▷ one for each Byzantine replica, and...
▷ send them to all honest replicas
5:   send  $b$  to honestReplicas
6: Procedure LaunchEquivocationCertAttack(▷ ***EQUIVOCATION-CERT-ATTACK***
7:  $e, validBlock, validBC, k, honestReplicas, f$ ) :
8: if isByzantineLeader( $e$ ) then ▷ if the epoch leader is Byzantine:
9:    $p_1, p_2 \leftarrow generateTwoDifferentProposals(e, validBlock, validBC)$  ▷ the replica generates...
▷ two different proposals, and...
10:   $v_1 \leftarrow generateVoteMessages(p_1, f)$  ▷  $f$  VOTE messages only for  $p_1$ , one for each Byzantine replica
11:   $set_1, set_2 \leftarrow getTwoRandomSets(e, k, honestReplicas)$  ▷ then, the replica divides honest replicas...
▷ in two random sets of size  $k$ 
12:  send  $p_1$  and  $v_1$  to  $set_1$  ▷ finally, it sends the first proposal and votes for it...
▷ to the first set of honest replicas,  $set_1, \dots$ 
13:  send  $p_1$  and  $p_2$  to  $set_2$  ▷ while sending both proposals to the second set of honest replicas,  $set_2$ 
14: Procedure LaunchBlameCertAttack(▷ ***BLAME-CERT-ATTACK***
15:  $e, validBlock, validBC, k, honestReplicas, f$ ) :
16: if isByzantineLeader( $e$ ) then ▷ if the epoch leader is Byzantine:
17:   $p \leftarrow generateNewProposal(e, validBlock)$  ▷ the replica generates new proposal for epoch  $e, \dots$ 
18:   $v \leftarrow generateVoteMessages(p, f)$  ▷  $f$  VOTE messages for proposal, and...
19:   $b \leftarrow generateBlameMessages(e, f)$  ▷  $f$  BLAME messages, one for each Byzantine replica
20:   $set_1, set_2 \leftarrow getTwoRandomSets(e, k, honestReplicas)$  ▷ then, replica divides honest replicas...
▷ in two random sets of size  $k$ 
21:  send  $p$  and  $v$  to  $set_1$  ▷ finally, it sends the proposal and votes for it...
▷ to the first set of honest replicas,  $set_1, \dots$ 
22:  send  $b$  to  $set_2$  ▷ while sending blame messages to the second set of honest replicas,  $set_2$ 

```

---

## 896 B Determining BoundBFT's synchrony bound

897 This section presents the complete data of the experiments we used to determine BoundBFT's  
898 synchronous bound (see Section 5.2). Namely, we implemented all proposed attacks (see  
899 Algorithm 3 and Figures 4 and 5). Then, we ran BoundBFT in our cluster while varying  
900 the number of Byzantine replicas  $f$ . As a starting point, we set  $\Delta$  to 1250 ms (99.99%),  
901 the synchronous bound from [30], and gradually decreased it to the point where we started  
902 to observe BoundBFT's agreement and progress violations. More than 300 hours worth of  
903 experiments were conducted in total. Tables 3 and 4 below show data for 1KB and 32KB  
904 block sizes, respectively.

N = 60, block size = 1KB												
Δ (ms)	No attack											
	Agreement						Progress					
1250	0%						0%					
600	0%						0%					
300	0%						0%					
150	0%						0%					
100	0%						0%					
50	0%						65%					
Δ (ms)	Equivocation attack											
	f=29				f=19				f=1			
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>	
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
300	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
150	0%	1%	0%	3%	0%	0%	0%	0%	0%	0%	0%	0%
100	5%	36%	6%	6%	0%	5%	0%	0%	0%	0%	0%	0%
50	19%	80%	31%	60%	2%	73%	8%	54%	0%	65%	0%	65%
Δ (ms)	Amnesia attack											
	f=29				f=19				f=1			
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>	
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
300	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
150	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
100	4%	19%	0%	18%	0%	14%	0%	8%	0%	0%	0%	0%
50	26%	90%	0%	80%	9%	81%	0%	97%	0%	66%	0%	67%
Δ (ms)	Blame attack											
	f=29				f=19				f=1			
	Agreement		Progress		Agreement		Progress		Agreement		Progress	
1250	0%		0%		0%		0%		0%		0%	
600	0%		0%		0%		0%		0%		0%	
300	0%		0%		0%		0%		0%		0%	
150	0%		0%		0%		0%		0%		0%	
100	0%		84%		0%		46%		0%		0%	
50	0%		100%		0%		100%		0%		67%	
Δ (ms)	Equivocation certificate attack											
	f=29				f=19				f=1			
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>	
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
300	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
150	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
100	0%	16%	0%	12%	0%	17%	0%	12%	0%	0%	0%	0%
50	2%	88%	0%	90%	0%	86%	0%	80%	0%	63%	0%	66%
Δ (ms)	Blame certificate attack											
	f=29				f=19				f=1			
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>	
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
300	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
150	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
100	0%	21%	0%	25%	0%	15%	0%	17%	0%	0%	0%	0%
50	0%	90%	0%	92%	0%	92%	0%	99%	0%	63%	0%	64%

Table 3: Percentage of Agreement and Progress violations when running BoundBFT under different attacks while using different values as its Δ. The table shows data for the setup of 60 replicas, 1KB block size and different number of Byzantine replicas (f). Additionally, for attacks that partition honest replicas in two subsets, it shows results when Byzantine replicas divide them in two minimal subsets (k=k<sub>min</sub>=1) and two maximal subsets (k=k<sub>max</sub>=n-f/2).

28:26 How robust are synchronous consensus protocols?

N = 60, block size = 32KB													
Δ (ms)	No attack												
	Agreement						Progress						
1250	0%						0%						
600	0%						0%						
300	0%						0%						
150	0%						0%						
100	0%						0%						
50	0%						66%						
Equivocation attack													
Δ (ms)	f=29				f=19				f=1				
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
300	0%	0%	0%	5%	0%	0%	0%	0%	0%	0%	0%	0%	
150	2%	0%	0%	15%	0%	0%	0%	0%	0%	0%	0%	0%	
100	5%	34%	0%	67%	0%	13%	0%	9%	0%	0%	0%	0%	
50	27%	91%	5%	85%	1%	94%	0%	97%	0%	66%	0%	66%	
Amnesia attack													
Δ (ms)	f=29				f=19				f=1				
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
300	0%	2%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
150	3%	26%	1%	10%	0%	1%	0%	0%	0%	0%	0%	0%	
100	6%	42%	0%	44%	0%	22%	0%	0%	0%	0%	0%	0%	
50	27%	81%	10%	88%	9%	74%	0%	37%	0%	67%	0%	67%	
Blame attack													
Δ (ms)	f=29				f=19				f=1				
	Agreement		Progress		Agreement		Progress		Agreement		Progress		
1250	0%		0%		0%		0%		0%		0%		
600	0%		0%		0%		0%		0%		0%		
300	0%		0%		0%		0%		0%		0%		
150	0%		25%		0%		0%		0%		0%		
100	0%		85%		0%		43%		0%		0%		
50	0%		100%		0%		99%		0%		79%		
Equivocation certificate attack													
Δ (ms)	f=29				f=19				f=1				
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
300	0%	0%	0%	2%	0%	0%	0%	0%	0%	0%	0%	0%	
150	0%	0%	0%	18%	0%	0%	0%	0%	0%	0%	0%	0%	
100	0%	31%	0%	61%	0%	13%	0%	17%	0%	0%	0%	0%	
50	3%	86%	0%	83%	0%	95%	0%	95%	0%	67%	0%	67%	
Blame certificate attack													
Δ (ms)	f=29				f=19				f=1				
	k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		k <sub>min</sub>		k <sub>max</sub>		
	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	Agreement	Progress	
1250	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
600	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
300	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	
150	0%	2%	0%	3%	0%	0%	0%	0%	0%	0%	0%	0%	
100	0%	30%	0%	32%	0%	17%	0%	12%	0%	0%	0%	0%	
50	1%	90%	0%	96%	0%	95%	0%	99%	0%	66%	0%	66%	

Table 4: Percentage of Agreement and Progress violations when running BoundBFT under different attacks while using different values as its Δ. The table shows data for the setup of 60 replicas, 32KB block size and different number of Byzantine replicas (f). Additionally, for attacks that partition honest replicas in two subsets, it shows results when Byzantine replicas divide them in two minimal subsets (k=k<sub>min</sub>=1) and two maximal subsets (k=k<sub>max</sub>=n-f/2).