

Synchronous Byzantine Agreement with Expected $O(1)$ Rounds, Expected $O(n^2)$ Communication, and Optimal Resilience

Ittai Abraham¹, Srinivas Devadas², Danny Dolev³, Kartik Nayak^{1,4}, and Ling Ren^{1,5}

¹ VMware Research – {iabraham,nkartik,lingren}@vmware.com

² MIT – devadas@mit.edu

³ Hebrew University of Jerusalem – danny.dolev@mail.huji.ac.il

⁴ Duke University

⁵ University of Illinois at Urbana-Champaign

Abstract. We present new protocols for Byzantine agreement in the synchronous and authenticated setting, tolerating the optimal number of f faults among $n = 2f + 1$ parties. Our protocols achieve an expected $O(1)$ round complexity and an expected $O(n^2)$ communication complexity. The exact round complexity in expectation is 10 for a static adversary and 16 for a strongly rushing adaptive adversary. For comparison, previous protocols in the same setting require expected 29 rounds and expected $\Omega(n^3)$ communication. We also give a lower bound showing that expected $\Omega(f^2)$ communication is necessary against a strongly rushing adaptive adversary.

1 Introduction

Byzantine agreement [26] is a fundamental problem in distributed computing and cryptography. It has been used to build fault tolerant distributed systems [4, 9, 24, 37], secure multi-party computation [6, 19], and more recently cryptocurrencies [3, 23, 30, 32]. In Byzantine agreement, a group n parties, each holding an initial input value, hope to commit on a common value; up to f parties can have Byzantine faults and deviate from the protocol arbitrarily. In a closely related problem called Byzantine broadcast, instead of each party holding an input value, there is one designated *sender* who tries to broadcast a value; To rule out trivial solutions, both problems have additional validity requirements.

Byzantine agreement and Byzantine broadcast have been studied under various combinations of assumptions, most notably timing assumptions – synchrony, asynchrony or partial synchrony, and setup assumptions – cryptography and public-key infrastructure (PKI). It is now well understood that these assumptions drastically affect the fault tolerance bounds. In particular, Byzantine broadcast and Byzantine agreement both require $f < n/3$ under partial synchrony or asynchrony. But under synchrony with digital signatures and PKI, Byzantine agreement can be solved with $f < n/2$ while Byzantine broadcast can be solved with $f < n - 1$.

In this paper, we consider Byzantine agreement in the synchronous and authenticated (i.e., assuming digital signatures and PKI) setting. The efficiency metrics we consider are (1) round complexity, i.e., the number of rounds of communication before the protocol terminates, and (2) communication complexity, i.e., the amount of information exchanged between parties during the protocol. For convenience, we measure communication complexity using the number of signatures exchanged between parties. Assuming each signature has λ bits, multiplying our communication complexity by λ yields the asymptotic communication complexity in bits.

In the synchronous and authenticated setting, Dolev and Strong gave a deterministic Byzantine broadcast protocol for $f < n - 1$ [12]. Their protocol achieves $f + 1$ round complexity and $O(n^2 f)$ communication complexity. The $f + 1$ round complexity matches the lower bound for deterministic protocols [12, 15]. To

A preliminary draft of the paper appeared on ePrint in 2017 [1]. The current version improves and subsumes the Byzantine agreement part of the preliminary draft. In addition, it introduces a lower bound for communication complexity of Byzantine broadcast against a strongly rushing adaptive adversary.

further improve round complexity, randomized protocols have been introduced [5, 14, 17, 35]. The most efficient protocol to our knowledge is proposed by Katz and Koo [21], which solves Byzantine agreement for $f < n/2$ using expected $O(n^3)$ communication and expected $O(1)$ rounds, and to be precise, expected 29 rounds.

In this work, we improve communication complexity to expected $O(n^2)$ and round complexity to expected 16. To the best of our knowledge, our protocol is the first to simultaneously achieve (i) expected $O(1)$ rounds, (ii) expected $O(n^2)$ communication, and (iii) optimal fault tolerance of $n = 2f + 1$ for synchronous and authenticated Byzantine agreement. In fact, we do not know of any existing protocol that achieves both (ii) and (iii).

Our protocols use threshold signatures [8, 36] to reduce communication complexity and a random leader election subroutine to reduce round complexity. The random leader election subroutine can be constructed using common-coin protocols [8, 29], and there exist constructions with a single round and $O(n^2)$ communication in the literature. The protocol by Cachin et al. [8] is secure against a static adversary whereas the protocol by Loss and Moran [29] are secure against an adaptive adversary. With these, we achieve the following result.

Theorem 1. *Synchronous authenticated Byzantine agreement can be solved for $f < n/2$ with*

- *expected 10 rounds and expected $O(n^2)$ communication against a static adversary assuming a single-round common-coin protocol,*
- *expected 16 rounds and expected $O(n^2)$ communication against a strongly rushing adaptive adversary assuming an adaptively secure single-round common-coin protocol.*

It is worth noting that our protocols work even in the presence of a very powerful adversary, which we call a *strongly rushing adaptive adversary*. The adversary can adaptively decide which f parties to corrupt and when to corrupt them. And by “strongly rushing”, we mean that if the adversary decides to corrupt a party h after observing messages sent from h to any other party in round r , it can remove h ’s round- r messages from the network before they reach other honest parties. In comparison, a standard rushing adversary can decide its own round- r messages after learning honest parties’ round- r messages, but if it corrupts h in round r , it cannot “take back” or alter h ’s round- r messages to other parties. The Dolev-Strong and Katz-Koo protocols also work against such a strongly rushing adaptive adversary. Alternatively, random leader election can also be achieved using verifiable random functions [30], albeit to achieve expected $O(1)$ rounds only against a rushing (but not strongly rushing) adaptive adversary.

The $O(1)$ expected round complexity is clearly asymptotically optimal. A natural question is whether or not the expected quadratic communication can be further improved. Building on a work by Dolev and Reischuk [11], we show that $\Omega(f^2)$ expected messages are necessary against a *strongly rushing* adaptive adversary. Specifically, we show the following lower bound:

Theorem 2. *If a protocol solves Byzantine broadcast with $\frac{1}{2} + \epsilon$ probability against a strongly rushing adaptive adversary, then in expectation, honest parties need to send at least $(\epsilon f/2)^2$ messages.*

Note that even though the lower bound is stated for Byzantine broadcast, it applies to Byzantine agreement as well, because Byzantine broadcast can be achieved by having the sender broadcast its value and then running one instance of Byzantine agreement.

1.1 Related Work

Byzantine agreement and broadcast were first introduced by Lamport, Shostak and Pease [26, 34]. They presented protocols and fault tolerance bounds for two settings (both synchronous). Without cryptographic assumptions (the unauthenticated setting), Byzantine broadcast and agreement can be solved if and only if $f < n/3$. Assuming digital signatures (the authenticated setting), Byzantine broadcast can be solved if $f < n - 1$ and Byzantine agreement can be solved if and only if $f < n/2$. The protocols from [26, 34] had exponential message complexities. Fully polynomial protocols were later shown for both the authenticated ($f < n/2$) [12] and the unauthenticated ($f < n/3$) [18] settings. Both protocols require $f + 1$ rounds of

Table 1: **Comparison with closely related works.** Here, κ is a security parameter and ϵ is a positive constant. We assume all protocols have access to ideal digital signatures.

Protocol	Exp. Comm. Complexity	Exp. Round Complexity	Resilience	Failure Probability	Adversarial Model
Dolev-Strong [12]	$O(n^2 f)$	$f + 1$	$f < n - 1$	0	strongly rushing adaptive
Katz-Koo [21]	$O(n^3)$	29	$f < n/2$	0	strongly rushing adaptive
Micali-Vaikuntanathan [31]	$O(n^3 \kappa)$	$O(\kappa)$	$f < n/2$	$2^{-\kappa}$	strongly rushing adaptive
King-Saia [22]	$n^{1.5} \cdot \text{poly log}(n)$	$\text{poly log}(n)$	$f < (\frac{1}{3} - \epsilon)n$	$1/\text{poly}(n)$	rushing adaptive
Chan-Pass-Shi [10]	$n \cdot \text{poly log}(\kappa)$	$O(1)$	$f < (\frac{1}{2} - \epsilon)n$	$\text{negl}(\kappa)$	rushing adaptive
This work (upper bound)	$O(n^2)$	16	$f < n/2$	0	strongly rushing adaptive
This work (lower bound)	$> (\epsilon f/2)^2$	N/A	any f	$\frac{1}{2} - \epsilon$	strongly rushing adaptive

communication, which matches the lower bound on round complexity for deterministic protocols [12, 15]. To circumvent the $f + 1$ round lower bound, a line of work explored the use of randomization [5, 35] which eventually led to expected constant-round protocols for both the authenticated ($f < n/2$) [21] and the unauthenticated ($f < n/3$) [14] settings. In the asynchronous setting, the FLP impossibility [16] rules out any deterministic solution. Some works use randomization [5, 7, 8] or partial synchrony [13] to circumvent the impossibility.

We have already compared with closely related works by Dolev-Strong and Katz-Koo. Recently, Micali and Vaikuntanathan gave a randomized protocol that solves Byzantine broadcast for $f < n/2$ with $1 - 2^{-\kappa}$ probability using $O(\kappa)$ rounds and $O(n^3 \kappa)$ communication [31]. Their protocol does not support early termination, so all parties always have to run for $O(\kappa)$ rounds and incur the round/communication complexity above. Table 1 summarizes related works and compares with our results. The numbers in Table 1 have considered adding threshold signatures to those works. For Dolev-Strong [12], there is no obvious way to utilize threshold signatures because each round uses a different threshold: a valid message in round k must be signed by k replicas. Katz-Koo [21] and Micali-Vaikuntanathan [31] both invoke n instances of gradecasts in parallel. Threshold signatures can improve one gradecast instance to $O(n^2)$, so n parallel gradecast instances result in $O(n^3)$ communication complexity.

The King-Saia protocol [22] solves synchronous Byzantine agreement for $f < (\frac{1}{3} - \epsilon)n$ with $1/\text{polyn}$ failure probability and uses $n^{1.5} \cdot \text{poly log}(n)$ communication. Recently, Chan, Pass and Shi [10] have shown a protocol that solves *binary* Byzantine agreement for $f < (\frac{1}{2} - \epsilon)n$ with $\text{negl}(\kappa)$ failure probability in expected $O(1)$ rounds and $n \cdot \text{poly log}(\kappa)$ communication [10]. Intuitively, both protocols save communication by down-sampling the population into a smaller committee and running a Byzantine agreement protocol within the committee. Their techniques are orthogonal and complementary to our work since they still need a protocol for the committee. Interestingly, the protocol run by the committee in Chan-Pass-Shi is inspired by the protocol described in this work. We remark that, given Theorem 2, it is not surprising that these sub-quadratic solutions only work against a standard rushing adaptive adversary but not a strongly rushing adaptive one.

1.2 Technical Overview

We first describe our core protocol, which ensures agreement (safety)⁶ and termination as required by Byzantine broadcast/agreement, but provides a weak notion of validity. Specifically, it achieves

- **Termination:** all honest parties eventually commit,
- **Agreement/safety:** all honest parties commit on the same value, and
- **Validity:** if all honest parties start with certificates for the same value v , and no Byzantine party starts with a certificate for a contradictory value, then all honest parties commit on v .

⁶ We will refer to the agreement requirement as “safety” for the rest of the paper.

In Section 4 we will describe how to obtain these certificates to solve Byzantine broadcast or Byzantine agreement.

The core protocol runs in iterations. In each iteration, a unique leader is elected. Each new leader picks up the state left by previous leaders and proposes a value in its iteration. Parties then cast votes on the leader’s value v . In more detail, each iteration consists of 4 rounds. The first three rounds are conceptually similar to Paxos and PBFT: (1) the leader learns the states of the system, (2) the leader proposes a value, and (3) parties vote on the value. If a party receives $f + 1$ votes for the same value and does not detect leader equivocation, it commits on that value. We then add another round: (4) if a party commits, it notifies all other parties about the commit; upon receiving a notification, other parties *accept* the committed value and will vouch for that value to future leaders.

Ideally, if the leader is honest, all honest parties commit v upon receiving $f + 1$ votes for v at the end of that iteration. A Byzantine leader can easily waste its iteration by not proposing. But it can also perform the following more subtle attacks: (1) send contradicting proposals to different honest parties, or (2) send a proposal to some but not all honest parties. We must ensure these Byzantine behaviors do not violate safety.

The need for equivocation checks. To ensure safety in the first attack, parties engage in an all-to-all round of communication to forward the leader’s proposal to each other for an equivocation check. If a party detects leader equivocation, i.e., sees two conflicting signed proposals from the leader, it does not commit even if it receives $f + 1$ votes.

The need for a notify round. Using the second attack, a Byzantine leader can make some, but not all, honest parties commit on a value v . If the other honest parties do not know that v has been committed, they may commit $v' \neq v$ in a subsequent iteration. Therefore, whenever an honest party h commits on a value v , h needs to *notify* all other honest parties of its commit. h can do this by broadcasting the $f + 1$ votes it received. When another party h' receives such a notification, it “accepts” the value v . If a party has accepted v and receives a proposal $v' \neq v$ in a later iteration, it will not vote for v' unless it is shown a proof that voting for v' is safe. The details can be found in Section 3.

Safety, termination, and validity. Safety is preserved because when an honest party commits, (1) no other party can commit a different value in the same iteration (due to equivocation checks), and (2) no other value can gather enough votes in subsequent iterations (due to *notify* by the honest party). Validity follows from a similar argument: if all honest parties start the protocol with the same certified (i.e., accepted) value v and Byzantine parties do not have a different certified value, only v can gather enough votes. Termination is achieved when some honest party h receives $f + 1$ notify messages. At this point, h sends these $f + 1$ notifications to all other parties and terminates. The $f + 1$ notifications h sends will ensure termination of all other parties in the next round. If an honest leader emerges, all parties terminate in its iteration.

Round complexity and communication complexity. Since there are $f + 1$ honest out of $2f + 1$ parties, by electing a random leader in every iteration, the protocol terminates in 2 iterations in expectation. Depending on the adversarial model, each iteration ranges from 4 to 7 rounds. Each round uses $O(n^2)$ messages (all-to-all) and each message is either a single signature or a single $(f + 1)$ -out-of- n threshold signature. Thus, the protocol runs in expected $O(1)$ rounds and uses expected $O(n^2)$ communication.

Paxos, PBFT, XPaxos, and our protocol. Abstractly, this core protocol resembles the synod algorithm in Paxos [25] but is adapted to the synchronous and Byzantine setting. The main idea of the synod algorithm is to ensure *quorum* intersection [25] at one *honest* party. The core idea of Paxos is to form a quorum of size $f + 1$ before committing a value. With $n = 2f + 1$, two quorums always intersect at one party, which is honest in Paxos. This honest party in the intersection will force a future leader to respect the committed value. In order to tolerate f Byzantine faults, PBFT [9] uses quorums of size $2f + 1$ out of $n = 3f + 1$, so that two quorums intersect at $f + 1$ parties, among which one is guaranteed to be honest. Similar to PBFT, we also need to ensure quorum intersection at $f + 1$ parties. But this requires new techniques with $n = 2f + 1$ parties in total. On the one hand, an intersection of size $f + 1$ seems to require quorums of size $1.5f + 1$. (An subsequent work called Thunderella [33] uses this quorums size to improve the optimistic case.) On the other hand, a quorum size larger than $f + 1$ (the number of honest parties) seems to require participation from

Byzantine parties and thus loses liveness. As described in the core protocol, our synchronous *notify* round forms a *post-commit quorum* of size $2f + 1$, which intersects with any *pre-commit quorum* of size $f + 1$ at $f + 1$ parties. This satisfies the requirement of one honest party in the intersection. Moreover, since parties in the post-commit quorum only receive messages, liveness is not affected.

Our protocol also shares some similarity to XPaxos [28]. In XPaxos, a view-change involves changing a set of $f + 1$ active replicas (instead of only changing the leader). So far as all the active replicas in the old view notify all the active replicas in the new view, there will be one honest replica in the new view that can carry state across views. However, XPaxos makes progress only if all $f + 1$ active replicas are honest. In comparison, our protocol only requires the leader to be honest to make progress.

Achieving Byzantine broadcast and Byzantine agreement. The core protocol already ensures safety and termination, so we only need some technique to boost its weaker validity to what Byzantine broadcast/agreement require. Our protocol achieves this using a single round of all-to-all communication before invoking the protocol. This allows us to avoid the standard transformation of composing n parallel Byzantine broadcasts to achieve Byzantine agreement. As a result, our Byzantine agreement protocol has the same asymptotic round/communication complexity as the core protocol.

2 Model

We assume synchrony. If an honest party i sends a message to another honest party j at the beginning of a round, the message is guaranteed to reach by the end of that round. We describe the protocol assuming lock-step execution, i.e., parties enter and exit each round simultaneously. Later in Section ??, we will present a clock synchronization protocol to bootstrap lock-step execution from bounded message delay.

We assume digital signatures and trusted setup. In the trusted setup phase, a trusted dealer generates public/private key pairs for digital signatures and other cryptographic primitives for each party, and certifies each party’s public keys. We use $\langle x \rangle_i$ to denote a message x signed by party i , i.e., $\langle x \rangle_i = (x, \sigma)$ where σ is a signature of message x produced by party i using its private signing key. For efficiency, it is customary to sign the hash digest of a message. A message can be signed by multiple parties (or the same party) in layers, i.e., $\langle \langle x \rangle_i \rangle_j = \langle x, \sigma_i \rangle_j = (x, \sigma_i, \sigma_j)$ where σ_i is a signature of x and σ_j is a signature of $x \parallel \sigma_i$ (\parallel denotes concatenation). When the context is clear, we omit the signer and simply write $\langle x \rangle$ or $\langle \langle x \rangle \rangle$.

We require a random leader election subroutine. As mentioned, this subroutine can be instantiated using common-coin protocols [8, 29] or verifiable random functions [30]. It may also be left to higher level protocols. For example, a cryptocurrency may elect leaders based on proof of work.

We assume a strongly rushing adaptive adversary. After the trusted setup phase, the adversary can adaptively decide which f parties to corrupt and when to corrupt each of them as the protocol executes. Note, however, that the adversary is not *mobile*: it cannot un-corrupt a Byzantine party to restore its corruption budget. The adversary is also strongly rushing. In each round, the adversary observes any party i ’s message to any other party j . If the adversary decides to corrupt i at this point, it controls which other honest parties (if any) i sends messages to and what messages i sends them *in that round*.

3 A Synchronous Byzantine Synod Protocol

3.1 Core Protocol

Our core protocol is a synchronous Byzantine synod protocol with $n = 2f + 1$ parties. The goal of the core synod protocol is to guarantee that all honest parties eventually commit (termination) on the same value (agreement). In addition, it achieves the following notion of validity: if (1) all honest parties start with the same value and have a *certificate* for this value, and (2) the adversary does not start with a certificate for a contradictory value, then all honest parties commit on this value. In Section 4, we show how to obtain these certificates using a single pre-round to achieve Byzantine broadcast and Byzantine agreement. For ease of exposition, we will temporarily assume a static adversary in Section 3.1 while presenting the core protocol. A

static adversary has to decide which parties to corrupt after the trusted setup phase and before the protocol starts.

We now describe the protocol in detail. When a leader proposes a value v in iteration k , we say the proposal has rank k and write them as a tuple (v, k) . The first iteration has $k = 1$. Each party i internally maintains states $\text{accepted}_i = (v_i, k_i, \mathcal{C}_i)$ across iterations to record its *accepted* proposal. Initially, each party i initializes $\text{accepted}_i := (\perp, 0, \perp)$. If party i later accepts (v, k) , it sets $\text{accepted}_i := (v, k, \mathcal{C})$ such that \mathcal{C} certifies that v is legally accepted in iteration k . \mathcal{C} consists of $f + 1$ commit requests for proposal (v, k) (see the protocol for details). We also say \mathcal{C} certifies, or is a certificate for, (v, k) . Proposals are ranked by the iteration number in which they are made. Namely, (v, k) is ranked higher than, lower than, or equal to (v', k') if $k > k'$, $k < k'$ and $k = k'$, respectively. Certificates are ranked by the proposals they certify. When we say a party “broadcasts” a message, we mean it sends the message to all parties including itself.

Round 0 (elect) All parties participate in the threshold coin-tossing scheme from [8]. Their scheme costs a single round and outputs a random string to all parties. The random string modulo n defines a random leader L_k for the current iteration k . We henceforth write L_k as L for simplicity.

Round 1 (status) Each party i sends a $\langle k, \text{status}, v_i, k_i, \mathcal{C}_i \rangle_i$ message to L to report its current accepted value.

At the end of this round, if party i reports the highest certificate to L (i could be L itself), L sets $\text{accepted}_L = (v_L, k_L, \mathcal{C}_L) := (v_i, k_i, \mathcal{C}_i)$. If no party reports a certificate, L chooses v_L freely and sets $k_L := 0$ and $\mathcal{C}_L := \perp$.

Round 2 (propose) L broadcasts a signed proposal $\langle \langle k, \text{propose}, v_L \rangle_L, k_L, \mathcal{C}_L \rangle_L$.

At the end of this round, party i sets $v_{L \rightarrow i} := v_L$ if the certificate it receives in the above leader proposal is no lower than what i reported to the leader, i.e., if $k_L \geq k_i$. Otherwise (leader is faulty), it sets $v_{L \rightarrow i} := \perp$.

Round 3 (commit) If $v_{L \rightarrow i} \neq \perp$, then party i forwards the proposal $\langle k, \text{propose}, v_{L \rightarrow i} \rangle_L$ to all other parties and broadcasts a $\langle k, \text{commit}, v_{L \rightarrow i} \rangle_i$ request.

At the end of this round, if party i is forwarded a properly signed proposal $\langle k, \text{propose}, v' \rangle_L$ in which $v' \neq v_{L \rightarrow i}$, it does not commit in this iteration (leader has equivocated). Else, if party i receives $f + 1$ $\langle k, \text{commit}, v \rangle_j$ requests in all of which $v = v_{L \rightarrow i}$, it commits on v and sets its internal state \mathcal{C}_i to be these $f + 1$ commit requests concatenated. In other words, party i commits if and only if it receives $f + 1$ matching commit requests and does not detect leader equivocation.

Round 4 (notify) If party i has committed on v at the end of the previous round, it sends a notification $\langle \langle \text{notify}, v \rangle_i, \mathcal{C}_i \rangle_i$ to every other party.

At the end of this round, if party i receives a $\langle \langle \text{notify}, v \rangle_j, \mathcal{C}_j \rangle_j$ message, it accepts v by setting $\text{accepted}_i = (v_i, k_i, \mathcal{C}_i) := (v, k, \mathcal{C})$. If party i receives multiple valid notify messages with different values (how this can happen is explained at the end of Section 3.2), it can accept an arbitrary one. Lastly, party i increments the iteration counter k and enters the next iteration.

Early and non-simultaneous termination. At any point during the protocol, if a party gathers *notification headers* (excluding certificates) $\langle \text{notify}, v \rangle$ from $f + 1$ distinct parties, it sends these $f + 1$ notification headers to all other parties and terminates. This ensures that when the first honest party terminates, all other honest parties receive $f + 1$ notification headers and terminate in the next round.

3.2 Safety, Termination, and Validity

In this section, we prove that the core protocol in Section 3.1 provides safety, termination and a weak notion of validity.

Safety. We first give some intuition to aid understanding. The scenario to consider for safety is when an honest party h commits on a value v^* in iteration k^* . We first show that Byzantine parties cannot hold a certificate for a value other than v^* in iteration k^* . Thus, all other honest parties accept v^* at the end of iteration k^* upon receiving notify from the honest party h . Thus, a value other than v^* cannot gather enough

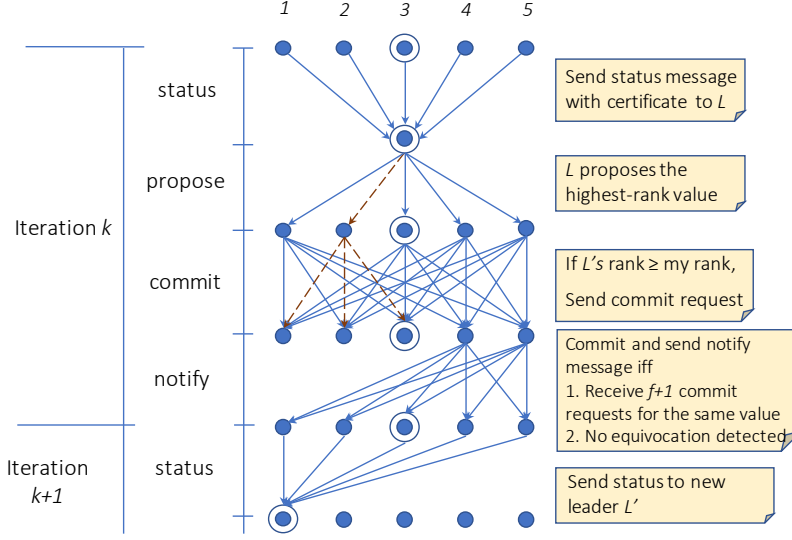


Fig. 1: An example iteration of the core protocol. In this example, $f = 2$, $n = 2f + 1 = 5$, parties 3 and 4 are Byzantine. **1.** (status) Each party sends its current states to $L = 3$. **2.** (propose) No party has committed or accepted any value, so L can propose any value of its choice. L equivocates and sends one proposal to party 4 (shown by dashed red arrow) and a different proposal to honest parties. **3.** (commit) Honest parties forward L 's proposal and send commit requests to all parties. Party 4 only sends to parties $\{3, 4, 5\}$. Parties 1 and 2 receive $f + 1$ commit requests for the blue value and do not detect equivocation, so they commit. Party 5 detects leader equivocation and does not commit despite also receiving $f + 1$ commit requests for the blue value. **4.** (notify) Parties 1 and 2 notify all other parties. On receiving a valid notification, party 5 accepts the blue value. **5.** (status) The parties send status messages to the new leader $L' = 1$ for iteration $k + 1$.

votes in iteration $k^* + 1$, and hence cannot be committed or accepted in iteration $k^* + 1$, and hence cannot gather enough votes in iteration $k^* + 2$, and so on. Safety then holds by induction.

We now formalize the above intuition by proving the following lemma about certificates: once an honest party commits, all certificates in that iteration and future iterations can only certify its committed value.

Lemma 1. *Suppose party h is the first honest party to commit and it commits on v^* in iteration k^* . If a certificate \mathcal{C} for (v, k^*) exists, then $v = v^*$.*

Proof. \mathcal{C} must consist of $f + 1$ commit requests for v . At least one of these comes from an honest party (call it h_1). Thus, h_1 must have received a proposal for v from the leader, and must have forwarded the proposal to all other parties. If $v \neq v^*$, h would have detected leader equivocation, and would not have committed on v^* in this iteration. So we have $v = v^*$.

Lemma 2. *If at the start of iteration k , (1) every honest party i has a certificate for (v, k_i) , and (2) all conflicting certificates are lower ranked, i.e., any certificate for (v', k') where $v \neq v'$ must have $k' < k_i$ for all honest i , then the above two conditions will hold at the end of iteration k .*

Proof. Suppose for contradiction that some party (honest or Byzantine) acquires a *higher* certificate than what it had previously for $v' \neq v$. Then it must receive from one honest party (call it h) a $\langle k, \text{commit}, v' \rangle_h$ request in iteration k . Note that h has a certificate for (v, k_h) at the start of iteration k . In order for h to send a commit request for v' , the leader L_k must show a certificate for (v', k') such that $k' \geq k_h$, which contradicts condition (2).

A simple induction shows that the above two conditions, if true at the start of an iteration, will hold true forever.

Theorem 3 (Safety). *If two honest parties commit on v and v' respectively, then $v = v'$.*

Proof. Suppose party h is the first honest party to commit, and it commits on v^* in iteration k^* . After the notify round of iteration k^* , every honest party receives a certificate for (v^*, k^*) and accepts v^* . Furthermore, due to Lemma 1, there cannot be a certificate for (v, k^*) in iteration k^* for $v \neq v^*$. Thus, the two conditions in Lemma 2 hold at the end of iteration k^* . So no certificate for a value other than v^* can be formed from this point on. In order for an honest party to commit on v , there must be a certificate for (v, k) where $k \geq k^*$. Therefore, $v = v^*$. Similarly, $v' = v^*$, and we have $v = v'$.

Termination. We now show that an honest leader will guarantee all honest parties terminate by the end of that iteration.

Theorem 4 (Termination). *If the leader L_k in iteration k is honest, then every honest party terminates one round after iteration k (or earlier).*

Proof. The honest leader L_k will send a proposal to all parties. It will propose a value reported by the highest certificate it collects in the **status** round. This certificate will be no lower than any certificate held by honest parties. Additionally, the unforgeability of digital signatures prevents Byzantine parties from falsely accusing L of equivocating. Therefore, all honest parties will send **commit** requests for v , receive $f + 1$ **commit** requests for v , commits on v , send notification headers for v , receive $f + 1$ notification headers for v (this is the end of iteration k), and terminate in the next round. (It is possible that they receive $f + 1$ notification headers and terminate at any earlier time.)

Validity. We now discuss the validity achieved by our core protocol. In the theorem, we assume the existence of initial certificates for $(v, 0)$ that are input to our core protocol. These initial certificates will be provided by higher-level protocols that invoke the core protocol (c.f. Section 4).

Theorem 5 (Validity). *All honest parties will commit on v if (1) every honest party starts with an initial certificate C certifying v , and (2) no Byzantine party has a certificate C' certifying $v' \neq v$.*

Proof. The proof is straightforward from Lemma 2 and Theorem 4. The input constraints satisfy the two conditions for Lemma 2 with each $k_i = 0$. Due to Lemma 2, for all subsequent iterations, only v can have certificates and thus, only v can be committed. By Theorem 4, when an honest leader emerges, all honest parties will commit on v .

Finally, we mention an interesting scenario that does not have to be explicitly addressed in the proofs. Before any honest party commits, Byzantine parties may obtain certificates for multiple values in the same iteration. In particular, the Byzantine leader proposes two values v and v' to all the f Byzantine parties. (An example with more than two values is similar.) Byzantine parties then exchange f **commit** requests for both values among them. Additionally, the Byzantine leader proposes v and v' to different honest parties. Now with one more **commit** request for each value from honest parties, Byzantine parties can obtain certificates for both v and v' , and can make honest parties accept different values by showing them different certificates (notify messages). However, this will not lead to a safety violation because no honest party would have committed in this iteration: the leader has equivocated to honest parties, so all honest parties will detect equivocation from forwarded proposals and thus refuse to commit. This scenario showcases the necessity of both the synchrony assumption and the use of digital signatures for our protocol. Lacking either one, equivocation cannot be reliably detected and any protocol will be subject to the $f < n/3$ bound. For completeness, we note that the above scenario will not lead to a violation of the termination property, either. At the end of the iteration, honest parties may accept either value. But in the next iteration, they can still vote for either value despite having accepted the other, since the two values have the same rank.

3.3 Random Leader Election against an Adaptive Adversary

The protocol presented so far does not achieve expected constant rounds against an adaptive adversary. The adversary learns who the leader L is after the `elect` round in an iteration. It can then immediately corrupt L and prevent it from sending any proposal. This way, the adversary forces the protocol to run for f iterations.

A first modification towards adaptive security is to move the `elect` round after the `propose` round and before the `commit` round. The hope is that, by the time L is corrupted, all honest parties have already received its proposal. This means every party should act as a potential leader before L_k is revealed, i.e., in `status` and `propose` rounds to collect status and make a proposal. From the `commit` round onward, only L 's proposal is relevant.

However, this idea alone is not sufficient. At the end of the `elect` round, after learning the identity of L , the adversary corrupts L , signs an equivocating proposal using L 's secret key and forwards it to all honest parties. Honest parties will detect equivocation from L and will not commit in this iteration. We are again forced to run the protocol for f iterations.

To this end, we need to add a step for each party to “prepare” its proposal before the leader is revealed. Afterwards, only “prepared” proposals are considered in equivocation checking. The prepare step should guarantee that, if a party h is honest throughout the prepare process but becomes corrupted afterwards, an adversary cannot construct a “prepared” equivocating proposal on h 's behalf. We achieve the prepare step in two rounds as follows.

Round P1 (`prepare1`) Each party i broadcasts its proposal $\langle v_i, k \rangle_i$.

Round P2 (`prepare2`) If party j receives a proposal $\langle v_i, k \rangle_i$ from party i in the previous round, party j signs the proposal and sends $\langle v_i, k \rangle_j$ back to party i .

We say a proposal (v_i, k) is prepared if it carries $f + 1$ signatures from distinct parties. Each honest party will be able to prepare its proposal. If party i is honest in the two `prepare` rounds and becomes corrupted only afterwards, preparing a conflicting proposal on party i 's behalf requires forging at least one honest party's signature, which a computationally bounded adversary cannot do.

The core protocol against a strongly rushing adaptive adversary now has 7 rounds: `status`, `prepare1`, `prepare2`, `propose`, `elect`, `commit`, and `notify`. Proofs for safety and validity remain unchanged from the static case. Proof of termination and round complexity analysis also hold once we observe that (1) there is a $> 1/2$ chance that each leader L_k is honest up to the point at which it is revealed, (2) if L_k is still honest by the end of the `propose` round of iteration k , all honest parties will consider its proposal valid and terminate one round after iteration k .

We remark that leader election based on verifiable random function [30], when combined with our `prepare` rounds, achieves expected 2 iterations against a (normal) rushing adaptive adversary. But it will run into f iterations against a strongly rushing adaptive adversary, who can prevent a leader from announcing its rank after receiving it.

3.4 Round Complexity and Communication Complexity

The first honest leader will ensure termination. The random leader election subroutine ensures a $(f + 1)/(2f + 1) > 1/2$ probability that each leader is honest, so the core protocol terminates in expected 2 iterations, plus one extra round to forward $f + 1$ `notify`. Thus, if an iteration requires r rounds, our core protocol requires $2r + 1$ rounds to terminate in expectation. If the adversary is adaptive and strongly rushing, each iteration requires $r = 7$ rounds. If the adversary is adaptive and normal rushing, the `elect` round can happen in parallel to `propose`, and each iteration has $r = 6$ rounds. If the adversary is static (rushing or otherwise), we do not need the two `prepare` rounds, and the `elect` round can happen in parallel to either `status` or `propose`, giving $r = 4$ rounds per iteration.

Next, we analyze the communication complexity. We will show that each round consumes $O(n^2)$ communication. Hence, the core protocol requires expected $O(n^2)$ communication (whether the adversary is adaptive or static, rushing or not). First of all, note that although a certificate consists of $f + 1$ signatures, its size can be reduced to a single signature using threshold signatures [8, 20, 27, 36].

1. In the **status** round, every party is reporting its currently accepted certificate to every other party (every party can potentially be the leader since the leader identity has not been revealed).
2. In **prepare₁**, every party sends a signed proposal, which is $O(1)$ in size, to every other party.
3. In **prepare₂**, every party sends back a doubly signed proposal, which is $O(1)$ in size, to every other party.
4. In the **propose** round, every party sends a proposal, which carries a certificate, to every other party. (A proposal need not contain **status** messages, following the suggestion of the HotStuff protocol [2]).
5. In the **elect** round, the common-coin protocol by Loss and Moran [29] requires $O(n^2)$ communication.
6. In the **commit** round, every party sends an $O(1)$ -sized **commit** message to every other party.
7. In the **notify** round, every party sends a **notify** message, which carries a certificate, to every other party.
8. Lastly, before termination, every party sends $f + 1$ notification headers $\langle \text{notify}, v \rangle$, which can be reduced to a single threshold signature, to every other party.

4 Byzantine Broadcast and Agreement

In this section, we describe how to use the core protocol to solve synchronous authenticated Byzantine broadcast and agreement for the $f < n/2$ case. For both problems, we design a “pre-round” to let honest parties obtain initial certificates and then invoke the core protocol.

Byzantine broadcast. In Byzantine broadcast, a designated *sender* tries to broadcast a value to n parties. A solution needs to satisfy three requirements:

- (**termination**) all honest parties eventually commit,
- (**agreement**) all honest parties commit on the same value, and
- (**validity**) if the sender is honest, then all honest parties commit on the value it broadcasts.

Let L_s be the designated sender. In the pre-round, L_s broadcasts a signed value $\langle v_s \rangle_{L_s}$ to every party. Such a signed value by the sender is an initial certificate certifying $(v_s, 0)$. We then invoke the core protocol. Safety and termination are satisfied due to Theorems 3 and 4. If the designated sender is honest, each honest party has a certificate for $(v_s, 0)$ and no conflicting initial certificate can exist, satisfying the condition for Theorem 5. Thus, validity is satisfied.

Byzantine agreement. In Byzantine agreement, every party holds an initial input value. A solution needs to satisfy the same termination and agreement requirements as in Byzantine broadcast. There exist a few different validity notions. We adopt a common one known as strong unanimity [13]:

- (**validity**) if all honest parties hold the same input value v , then they all commit on v .

In the pre-round, every party i broadcasts its value $\langle v_i \rangle_i$. $f + 1$ signatures from distinct parties for the same value v form an initial certificate for $(v, 0)$. We then invoke the core protocol. Safety and termination are satisfied due to Theorems 3 and 4. If all honest parties have the same input value, then they will have an initial certificate for v and no conflicting initial certificate can exist, satisfying the condition for Theorem 5. Thus, validity is satisfied.

The efficiency of the protocols is straightforward given the analysis of the core protocol. Both protocols require one more round than the core protocol and the same $O(n^2)$ communication complexity as the core protocol.

5 Lower Bound on Expected Communication Against a Strongly Rushing Adaptive Adversary

Our proof strategy builds on and extends the classic Dolev and Reischuk [11, Theorem 2] lower bound, which shows that in every deterministic Byzantine broadcast protocol honest parties need to send at least $\Omega(f^2)$ messages. We show that in every randomized protocol solving Byzantine broadcast with ϵ error probability against a strongly rushing adaptive adversary, honest parties need to send at least $(\epsilon f/2)^2$ messages.

The Dolev-Reischuk lower bound. We first explain the Dolev-Reischuk proof at a high level. Observe that for a deterministic protocol, an execution is completely determined by the input (of the designated sender) and the adversary's strategy. Consider the following adversarial strategy. The adversary corrupts an arbitrary set B of $f/2$ parties that does not include the designated sender. Let A denote the set of remaining parties. All parties in B behave like honest parties, except that (i) they ignore the first $f/2$ messages sent to them, and (ii) they do not send messages to each other. Call the above adversarial strategy **Strat**. Suppose the honest designated sender sends 0 under **Strat**. For validity to hold, then all honest parties commit on 0.

If at most $(f/2)^2$ messages are sent to B in the above execution, then there exists a party $p \in B$ such that p receives at most $f/2$ messages. Let $S(p)$ denote the set of parties that send messages to p . Clearly, $|S(p)| \leq f/2$. We define strategy **Strat'** identically as **Strat** except that: (i) do not corrupt p , (ii) corrupt all parties in $S(p)$ (possibly including the designated sender) and prevent them from sending any messages to p (but let $S(p)$ behave honestly to other parties). Since $|S(p)| \leq f/2$, **Strat'** corrupts at most f parties.

Observe that honest parties in $A \setminus S(p)$ receive identical messages from all other parties under **Strat** and **Strat'**, so they cannot distinguish the two executions. Thus, $A \setminus S(p)$ would still commit on 0 under **Strat'**. However, p does not receive any message and has to commit on an arbitrary value. If this value is 1, safety is violated. If p commits on 0 when receiving no messages, we can let the sender send $v = 1$ under **Strat** and derive a safety violation under **Strat'** following a symmetric argument.

Our lower bound. We now extend the above proof to randomized protocols. In a randomized protocol, there are two sources of randomness that need to be considered carefully. On the one hand, honest parties can use randomization to their advantage. On the other hand, an adaptive adversary can also leverage randomness. Indeed our lower bound uses a randomized adversarial strategy. In addition, our lower bound crucially relies on the adversary being *strongly rushing* – the adversary can observe that a message is sent by an honest party h to any other party in a given round r , decide to adaptively corrupt h , and then remove messages sent by h in round r . We prove the following theorem:

Theorem 6 (Theorem 2, restated). *If a protocol solves Byzantine broadcast with $\frac{1}{2} + \epsilon$ probability against a strongly rushing adaptive adversary, then in expectation, honest parties collectively need to send at least $(\epsilon f/2)^2$ messages.*

Proof. We first define adversarial strategy **Strat** (same as Dolev-Reischuk):

1. Corrupt an arbitrary set B of $f/2$ parties. Let A denote the remaining parties.
2. Each party in B behaves like an honest party except that, it ignores the first $f/2$ messages sent to it, and it does not send messages to other parties in B .

Suppose for contradiction that a protocol solves Byzantine broadcast with less than ϵ probability of error using less than $(\epsilon f/2)^2$ expected messages. For a protocol to have an expected message complexity of $(\epsilon f/2)^2$, honest parties collectively need to send fewer than that many messages in expectation *regardless of* the adversary's strategy. Let z be a random variable denoting the number of messages sent by honest parties to B . We have $E[z] < (\epsilon f/2)^2$. Let X be the event that $z < \epsilon(f/2)^2$. By Markov's inequality, $\Pr[z \geq \frac{1}{\epsilon} E[z]] \leq \epsilon$. Thus, $\Pr[z < \epsilon(f/2)^2] \geq \Pr[z < \frac{1}{\epsilon} E[z]] > 1 - \epsilon$.

Let Y be the event that among the first $\epsilon(f/2)^2$ messages, a party p picked uniformly at random by the adversary receives at most $f/2$ messages. Observe that among the first $\epsilon(f/2)^2 = \epsilon|B|(f/2)$ messages, there exist at most $\epsilon|B|$ parties that receive more than $f/2$ of those. Since p has been picked uniformly at random from B , $\Pr[Y] \geq 1 - \epsilon$.

$$\begin{aligned} \Pr[X \cap Y] &= \Pr[X] + \Pr[Y] - \Pr[X \cup Y] \\ &\geq (1 - \epsilon) + (1 - \epsilon) - 1 \\ &= 1 - 2\epsilon \end{aligned}$$

Define adversarial strategy **Strat'** identically to **Strat** except:

1. The adversary picks a party $p \in B$ uniformly at random. The adversary does not corrupt p .
2. Whenever some party $s \in A$ attempts to send a message to p in a round, if the adversary has not used up its corruption quota f , it immediately corrupts s and removes the message sent by s to p in that round. Once corrupted, s does not send p any messages but otherwise behaves correctly.

Observe that $X \cap Y$ denotes the event where the total number of messages sent by honest parties to B is less than $\epsilon(f/2)^2$ and among those p has received at most $f/2$ messages. Let $S(p)$ be the set of parties that attempt to send p messages (some or all of these attempts are blocked). Thus, we have shown that, $\Pr[S(p) \leq f/2] \geq 1 - 2\epsilon$. This is the probability that the random party p picked by the adversary under **Strat** receives $\leq f/2$ messages, which also means p receives no message at all under **Strat'**. In this case, honest parties in $A \setminus S(p)$ receive identical messages under the two adversarial strategies while p needs to commit on a value without receiving any message under **Strat'**.

Without loss of generality, suppose that when receiving no message p commits on 1 with at least $1/2$ probability (otherwise, p commits on 0 with at least $1/2$ probability that the proof follows from a symmetric argument). Let the designated sender send 0 under **Strat**. Conditioned on $X \cap Y$, validity holds under **Strat** if and only if safety is violated under **Strat'**. Thus, with at least $\frac{1}{2}(1 - 2\epsilon) = \frac{1}{2} - \epsilon$ probability, either validity is violated under **Strat** or safety is violated under **Strat'**.

6 Clock Synchronization

An important question is how practical the synchrony assumption is, which will be the topic of this section. The synchrony assumption essentially states that all honest replicas' messages arrive in time. This requires two properties: (i) a bounded message delay and (ii) locked step execution, i.e., honest replicas enter each round roughly at the same time. The second property is important because, if replica i enters a round much earlier than replica j , then i may end up finishing the round too soon without waiting for j 's message to arrive. In our protocol, for example, this could prevent i from detecting leader equivocation and result in a safety violation.

The XFT paper provided some justification for the bounded message delay assumption in certain applications [28]. But we still need a mechanism to enforce locked step execution. To this end, we will use the following clock synchronization protocol. It will be executed at known time intervals. We call each interval a "day".

Round 0 (*sync*) When party i 's clock reaches the beginning of day X , it sends a $\langle \text{sync}, X \rangle_i$ message to all parties including itself.

Round 1 (*new-day*) The first time a party j receives $f + 1$ $\langle \text{sync}, X \rangle$ messages from distinct parties (either as $f + 1$ separate *sync* messages or within a single *new-day* message), it

- sets its clock to the beginning of day X , and
- sends all other parties a *new-day* message, which is the concatenation of $f + 1$ $\langle \text{sync}, X \rangle$ messages from distinct parties.

The above protocol bootstraps lock-step synchrony from the message delay bound Δ and a clock drift bound. Each *sync* message is triggered by a party's own local clock, independent of when day X would start for other parties. The protocol refreshes honest parties' clock difference to at most the message delay bound Δ at the beginning of each day. The first honest party to start a new day will broadcast a *new-day* message, which makes all other honest parties start the new day within Δ time. Obtaining a *new-day* message also means at least one honest party has sent a valid *sync* message, ensuring that roughly one day has indeed passed since the previous day. We can then set the duration of each round to $2\Delta + \phi$ where ϕ is the maximum clock drift between two honest parties in a "day". The clock synchronization protocol may also be useful to synchronous protocols other than Byzantine agreement.

Acknowledgments

We thank Dahlia Malkhi and Benjamin Chan for many useful discussions.

References

1. Ittai Abraham, Srinivas Devadas, Danny Dolev, Kartik Nayak, and Ling Ren. Synchronous byzantine agreement with expected $O(1)$ rounds, expected $O(n^2)$ communication, and optimal resilience. Cryptology ePrint Archive, Report 2018/1028, 2018. <https://eprint.iacr.org/2018/1028>.
2. Ittai Abraham, Guy Gueta, and Dahlia Malkhi. Hot-stuff the linear, optimal-resilience, one-message bft devil. *arXiv preprint arXiv:1803.05069*, 2018.
3. Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Alexander Spiegelman. Solida: A blockchain protocol based on reconfigurable byzantine consensus. OPODIS, 2017.
4. Atul Adya, William J. Bolosky, Miguel Castro, Gerald Cermak, Ronnie Chaiken, John R. Douceur, Jon Howell, Jacob R. Lorch, Marvin Theimer, and Roger P. Wattenhofer. FARSITE: Federated, available, and reliable storage for an incompletely trusted environment. *ACM SIGOPS Operating Systems Review*, 36(SI):1–14, 2002.
5. Michael Ben-Or. Another advantage of free choice (extended abstract): Completely asynchronous agreement protocols. In *Proceedings of the second annual ACM symposium on Principles of distributed computing*, pages 27–30. ACM, 1983.
6. Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the 20th annual ACM symposium on Theory of computing*, pages 1–10. ACM, 1988.
7. Gabriel Bracha. Asynchronous byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987.
8. Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in Constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, 2005.
9. Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
10. TH Chan, Rafael Pass, and Elaine Shi. Communication-efficient byzantine agreement without erasures. *arXiv preprint arXiv:1805.03391*, 2018.
11. Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement. *Journal of the ACM (JACM)*, 32(1):191–204, 1985.
12. Danny Dolev and H. Raymond Strong. Authenticated algorithms for byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.
13. Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, 1988.
14. Pesech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM Journal on Computing*, 26(4):873–933, 1997.
15. Michael J. Fischer and Nancy A. Lynch. A lower bound for the time to assure interactive consistency. *Information processing letters*, 14(4):183–186, 1982.
16. Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, 1985.
17. Matthias Fitzi and Juan A Garay. Efficient player-optimal protocols for strong and differential consensus. In *Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 211–220. ACM, 2003.
18. Juan A. Garay and Yoram Moses. Fully polynomial byzantine agreement for $n > 3f$ processors in $f + 1$ rounds. *SIAM Journal on Computing*, 27(1):247–290, 1998.
19. Shafi Goldwasser, Silvio Micali, and Avi Wigderson. How to play any mental game, or a completeness theorem for protocols with an honest majority. In *Proc. of the 19th Annual ACM STOC*, volume 87, pages 218–229, 1987.
20. Guy Golan Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael K Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. SBFT: a scalable decentralized trust infrastructure for blockchains. *arXiv preprint 1804.01626*, 2018.
21. Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for Byzantine agreement. In *Annual International Cryptology Conference*, volume 4117, pages 445–462. Springer, 2006.
22. Valerie King and Jared Saia. Breaking the $O(n^2)$ bit barrier: scalable byzantine agreement with an adaptive adversary. *Journal of the ACM*, 58(4):18, 2011.
23. Eleftherios Kokoris Kogias, Philipp Jovanovic, Nicolas Gailly, Ismail Khoffi, Linus Gasser, and Bryan Ford. Enhancing Bitcoin security and performance with strong consistency via collective signing. In *25th USENIX Security Symposium*, pages 279–296. USENIX Association, 2016.
24. John Kubiatowicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishan Gummadi, Sean Rhea, Hakim Weatherspoon, Westley Weimer, and Ben Zhao. Oceanstore: An architecture for global-scale persistent storage. *ACM Sigplan Notices*, 35(11):190–201, 2000.

25. Leslie Lamport. The part-time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, 1998.
26. Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
27. Benoît Libert, Marc Joye, and Moti Yung. Born and raised distributively: Fully distributed non-interactive adaptively-secure threshold signatures with short shares. *Theoretical Computer Science*, 645:1–24, 2016.
28. Shengyun Liu, Christian Cachin, Vivien Quéma, and Marko Vukolic. XFT: practical fault tolerance beyond crashes. In *12th USENIX Symposium on Operating Systems Design and Implementation*, pages 485–500. USENIX Association, 2016.
29. Julian Loss and Tal Moran. Combining asynchronous and synchronous Byzantine agreement: The best of both worlds. *Cryptology ePrint Archive 2018/235*, 2018.
30. Silvio Micali. Algorand: The efficient and democratic ledger. arXiv:1607.01341, 2016.
31. Silvio Micali and Vinod Vaikuntanathan. Optimal and player-replaceable consensus with an honest majority, 2017.
32. Rafael Pass and Elaine Shi. Feasibilities and infeasibilities for achieving responsiveness in permissionless consensus. In *International Symposium on Distributed Computing*. Springer, 2017.
33. Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 3–33. Springer, 2018.
34. Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM (JACM)*, 27(2):228–234, 1980.
35. Michael O. Rabin. Randomized Byzantine generals. In *Proceedings of the 24th Annual Symposium on Foundations of Computer Science*, pages 403–409. IEEE, 1983.
36. Victor Shoup. Practical threshold signatures. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 207–220. Springer, 2000.
37. Lidong Zhou, Fred Schneider, and Robbert van Renesse. COCA: A secure distributed online certification authority. *ACM Transactions on Computer Systems*, 20(4):329–368, 2002.