Network-Agnostic State Machine Replication

Erica Blum[1], Jonathan Katz[2], and Julian Loss[1]

[1] University of Maryland
[2] George Mason University

Abstract. We study the problem of *state machine replication* (SMR)—the underlying problem addressed by blockchain protocols—in the presence of a malicious adversary who can corrupt some fraction of the parties running the protocol. Existing protocols for this task assume either a *synchronous* network (where all messages are delivered within some known time Δ) or an *asynchronous* network (where messages can be delayed arbitrarily). Although protocols for the latter case give seemingly stronger guarantees, this is not the case since they (inherently) tolerate a lower fraction of corrupted parties.

We design an SMR protocol that is *network-agnostic* in the following sense: if it is run in a synchronous network, it tolerates t_s corrupted parties; if the network happens to be asynchronous it is resilient to $t_a \leq t_s$ faults. Our protocol achieves optimal tradeoffs between t_s and t_a.

1 Introduction

State machine replication (SMR) is a fundamental problem in distributed computing [18,19,31] that can be viewed as a generalization of *Byzantine agreement* (BA) [20, 30]. Roughly speaking, a BA protocol allows a set of n parties to agree on a value *once*, whereas SMR allows those parties to agree on an infinitely long *sequence* of values with the additional guarantee that values input to honest parties are eventually included in the sequence. (See Section 3 for formal definitions. Note that SMR is not obtained by simply repeating a BA protocol multiple times; see further discussion in Section 1.1.) The desired properties should hold even in the presence of some fraction of corrupted parties who may behave arbitrarily. SMR protocols are deployed in real-world distributed data centers, and the problem has received renewed attention in the context of *blockchain protocols* used for cryptocurrencies and other applications.

Existing SMR protocols assume either a *synchronous network*, where all messages are delivered within some publicly known time bound Δ, or an *asynchronous network*, where messages can be delayed arbitrarily. Although it may appear that protocols designed for the latter setting are strictly more secure, this is not the case because they also (inherently)

tolerate a lower fraction of corrupted parties. Specifically, assuming a public-key infrastructure (PKI) is available to the parties, SMR protocols tolerating up to $t_s < n/2$ adversarial corruptions are possible in a synchronous network, but in an asynchronous network SMR is achievable only for $t_a < n/3$ faults (see [8]).

We study here so-called *network-agnostic* SMR protocols that offer meaningful guarantees regardless of the network in which they are run. That is, fix thresholds t_a, t_s with $0 \leq t_a < n/3$ and $t_a \leq t_s < n/2$. We seek to answer the following question: Assuming a PKI, is it possible to have an SMR protocol that tolerates (1) t_s (adaptive) corruptions if the network is synchronous and (2) t_a (adaptive) corruptions even if the network is asynchronous? We show that the answer is positive iff $t_a + 2t_s < n$.

Our work is directly inspired by recent results of Blum et al. [5], who study the same problem but for the simpler case of Byzantine agreement. We match their bounds on t_a, t_s and, as in their work, show that these bounds are optimal in our setting.[1] While the high-level structure of our SMR protocol resembles the high-level structure of their BA protocol, in constructing our protocol we need to address several technical challenges (mainly due to the stronger liveness property required for SMR; see the following section) that do not arise in their work.

1.1 Related Work

There is extensive prior work on designing both Byzantine agreement and SMR/blockchain protocols; we do not provide an exhaustive survey, but instead focus only on the most relevant prior work.

As argued by Miller et al. [25], many well-known SMR protocols that tolerate malicious faults (e.g., [7,16]) require at least partial synchrony in order to achieve liveness. Their HoneyBadger protocol [25] was designed specifically for asynchronous networks, but can only handle $t < n/3$ faults even if run in a synchronous network. Blockchain protocols are typically analyzed assuming synchrony [12, 26]; Nakamoto consensus, in particular, assumes that messages will be delivered much faster than the time required to solve proof-of-work puzzles.

We emphasize that SMR is *not* realized by simply repeating a (multi-valued) BA protocol multiple times. In particular, the validity property of BA only guarantees that if a value is input by all honest parties then

[1] It is not clear that SMR implies BA in the network-agnostic setting when $t_a + 2t_s \geq n$. Thus, impossibility of SMR when $t_a + 2t_s \geq n$ does not follow from the impossibility result for BA shown by Blum et al. [5].

that value will be output by all honest parties. In the context of SMR the parties each hold multiple inputs in a local buffer (where those inputs may arrive at arbitrary times), and there is no way to ensure that all honest parties will select the same value as input to some execution of an underlying BA protocol. Although generic techniques for compiling a BA protocol into an SMR protocol are known [8], those compilers are not network-agnostic and so do not suffice to solve our problem.

Our work focuses on protocols being run in a network that may be either synchronous or fully asynchronous. Other work looking at similar problems includes that of Malkhi et al. [24], who consider networks that may be either synchronous or *partially* synchronous; Liu et al. [21], who design a protocol that tolerates a minority of malicious faults in a synchronous network, and a minority of *fail-stop* faults in an asynchronous network; and Guo et al. [13] and Abraham et al. [2], who consider temporary disconnections between two synchronous network components.

A slightly different line of work [22, 23, 27, 28] looks at designing protocols with good *responsiveness*. Roughly speaking, such protocols still require the network to be synchronous, but terminate more quickly if the actual message-delivery time is lower than the known upper bound Δ. Kursawe [17] designed a protocol for an asynchronous network that terminates more quickly if the network is synchronous, but does not tolerate more faults in the latter case. Finally, other work [3, 9, 10, 29] considers a model where synchrony is available for some (known) limited period of time, but the network is asynchronous afterward.

1.2 Paper Organization

We define our model in Section 2, before giving definitions for the various tasks we consider in Section 3. In Section 4 we describe a network-agnostic protocol for the asynchronous common subset (ACS) problem. The ACS protocol is used as a subprotocol of our main result, a network-agnostic SMR protocol, that is described and analyzed in Section 5. In Section 6 we prove a lower bound showing that the thresholds we achieve are tight for network-agnostic SMR protocols. As discussed, Blum et al. [5] show an analogous result for BA that does not directly apply to our setting.

2 Model

Setup assumptions and notation. We consider a network of n parties $P_1, \ldots, P_n$ who communicate over point-to-point authenticated chan-

nels. We assume that the parties have established a public-key infrastructure prior to the protocol execution. That is, we assume that all parties hold the same vector $(pk_1, \ldots, pk_n)$ of public keys for a digital-signature scheme, and each honest party P_i holds the honestly generated secret key sk_i associated with pk_i. A *valid signature* σ *on* m *from* P_i is one for which $\mathsf{Vrfy}_{pk_i}(m, \sigma) = 1$. For readability, we use $\langle m \rangle_i$ to denote a tuple (i, m, σ) such that σ is a valid signature on message m signed using P_i's secret key.

For simplicity, we treat signatures as ideal (i.e., perfectly unforgeable); we also implicitly assume that parties use some form of domain separation when signing (e.g., by using unique session IDs) to ensure that signatures are valid only in the context in which they are generated.

Where applicable, we use κ to denote a statistical security parameter.

Adversarial model. We consider the security of our protocols in the presence of an adversary who can *adaptively* corrupt some number of parties. The adversary may coordinate the behavior of corrupted parties and cause them to deviate arbitrarily from the protocol. Note, however, that our claims about adaptive security are only with respect to the property-based definitions found in Section 3, not with respect to a simulation-based definition (cf. [11, 14]).

Network model. We consider two possible settings for the network. In the *synchronous* case, all messages are delivered within some known time Δ after they are sent, but the adversary can reorder and delay messages subject to this bound. (As a consequence, the adversary can potentially be *rushing*, i.e., it can wait to receive all incoming messages in a round before sending its own messages.) In this setting, we also assume all parties begin the protocol at the same time, and parties' clocks progress at the same rate. When we say the network is *asynchronous*, we mean that the adversary can delay messages for an arbitrarily long period of time, though messages must eventually be delivered. We do not make any assumptions on parties' local clocks in the asynchronous case.

We view the network as being either synchronous or asynchronous for the lifetime of the protocol (although we stress that the honest parties do not know which is the case).

3 Definitions

Although we are ultimately interested in state machine replication, our main protocol relies on various subprotocols for different tasks. We therefore provide relevant definitions here. Throughout, when we say a pro-

tocol achieves some property, we include the case where it achieves that property with overwhelming probability (in the implicit parameter κ).

3.1 Useful Subprotocols

In some cases we consider protocols where parties may not terminate (even upon generating output); for this reason, we mention termination explicitly in some definitions. Honest parties are those who are not corrupted by the end of the execution.

Reliable broadcast. A *reliable broadcast* protocol allows parties to agree on a value chosen by a designated sender. In contrast to the stronger notion of *broadcast*, here honest parties might not terminate (but, if so, then none of them terminate).

Definition 1 (Reliable broadcast). *Let Π be a protocol executed by parties $P_1, \ldots, P_n$, where a designated sender $P^* \in \{P_1, \ldots, P_n\}$ begins holding input v^* and parties terminate upon generating output.*

- **Validity:** *Π is t-*valid *if the following holds whenever at most t parties are corrupted: if P^* is honest, then every honest party outputs v^*.*
- **Consistency:** *Π is t-*consistent *if the following holds whenever at most t parties are corrupted: either no honest party outputs anything, or all honest parties output the same value $v \in \{0,1\}$.*

*If Π is t-valid and t-consistent, then we say it is t-*secure*.*

Byzantine agreement. A *Byzantine agreement* protocol allows parties who each hold some initial value to agree on an output value.

Definition 2 (Byzantine agreement). *Let Π be a protocol executed by parties $P_1, \ldots, P_n$, where each party P_i begins holding input $v_i \in \{0,1\}$.*

- **Validity:** *Π is t-*valid *if the following holds whenever at most t of the parties are corrupted: if every honest party's input is equal to the same value v, then every honest party outputs v.*
- **Consistency:** *Π is t-*consistent *if the following holds whenever at most t of the parties are corrupted: every honest party outputs the same value $v \in \{0,1\}$.*
- **Termination:** *Π is t-*terminating *if whenever at most t parties are corrupted, every honest party terminates with some output in $\{0,1\}$.*

*If Π is t-valid, t-consistent, and t-terminating, then we say it is t-*secure*.*

Asynchronous common subset (ACS). Informally, a protocol for the *asynchronous common subset* problem [4] allows n parties, each with some input, to agree on a subset of those inputs. (The term "asynchronous" in the name is historical, and one can also consider protocols for this task in the synchronous setting.)

Definition 3 (ACS). *Let Π be a protocol executed by parties $P_1, \ldots, P_n$, where each P_i begins holding input $v_i \in \{0,1\}^*$, and parties output sets of size at most n.*

- **Validity:** *Π is t-*valid *if the following holds whenever at most t parties are corrupted: if every honest party's input is equal to the same value v, then every honest party outputs $\{v\}$.*
- **Liveness:** *Π is t-*live *if whenever at most t of the parties are corrupted, every honest party produces output.*
- **Consistency:** *Π is t-*consistent *if whenever at most t parties are corrupted, all honest parties output the same set S.*
- **Set quality:** *Π has t-*set quality *if the following holds whenever at most t parties are corrupted: if an honest party outputs a set S, then S contains the inputs of at least $t+1$ honest parties.*

3.2 State Machine Replication

Protocols for *state machine replication* (SMR) allow parties to maintain agreement on an ever-growing, ordered sequence of *blocks*, where a block is a set of values called *transactions*. An SMR protocol does not terminate but instead continues indefinitely. We model the sequence of blocks output by a party P_i via a write-once array $\mathsf{Blocks}_i = \mathsf{Blocks}_i[1], \mathsf{Blocks}_i[2], \ldots$ maintained by P_i, each entry (or *slot*) of which is initially equal to $\perp$. We say that P_i *outputs a block in slot j* when P_i writes a block to $\mathsf{Blocks}_i[j]$; if $\mathsf{Blocks}_i[j] \neq \perp$ then we refer to $\mathsf{Blocks}_i[j]$ as the *block output by P_i in slot j*. We do not require that honest parties output a block in slot $j-1$ before outputting a block in slot j.

It is useful to define a notion of *epochs* for each party. (We stress that these are not global epochs; instead, each party maintains a local view of its current epoch.) Formally, we assume that each party P_i maintains a write-once array $\mathsf{Epochs}_i = \mathsf{Epochs}_i[1], \mathsf{Epochs}_i[2], \ldots$, each entry of which is initialized to 0. We say P_i *enters epoch j* when it sets $\mathsf{Epochs}_i[j] := 1$, and require:

- For $j > 1$, P_i enters epoch $j-1$ before entering epoch j.
- P_i enters epoch j before outputting a block in slot j.

An SMR protocol is run in a setting where parties asynchronously receive inputs (i.e., transactions) as the protocol is being executed; each party P_i stores transactions it receives in a local buffer buf_i. We imagine these transactions as being provided to parties by some mechanism external to the protocol (which could involve a gossip protocol run among the parties themselves), and make no assumptions about the arrival times of these transactions at any of the parties.

Definition 4 (State machine replication). *Let Π be a protocol executed by parties $P_1, \ldots, P_n$ who are provided with transactions as input and locally maintain arrays* Blocks *and* Epochs *as described above.*

- **Consistency:** *Π is t-*consistent *if the following holds whenever at most t parties are corrupted: if an honest party outputs a block B in slot j then all parties that remain honest output B in slot j.*
- **Strong liveness:** *Π is t-*live *if the following holds whenever at most t parties are corrupted: for any transaction* tx *for which every honest party received* tx *before entering epoch j, every party that remains honest outputs a block that contains* tx *in some slot $j' \leq j$.*
- **Completeness:** *Π is t-*complete *if the following holds whenever at most t parties are corrupted: for all $j \geq 0$, every party that remains honest outputs some block in slot j.*

*If Π is t-consistent, t-live, and t-complete, then we say it is t-*secure*.*

Our liveness definition is stronger than usual, in that we require a transaction tx that appears in all honest parties' buffers by epoch j to be included in a block output by each honest party in some slot $j' \leq j$. (Typically, liveness only requires that each honest party eventually output a block containing tx.) This stronger notion of liveness is useful for showing that SMR implies Byzantine agreement (cf. Appendix A) and is achieved by our protocol.

In our definition, a transaction tx is only guaranteed to be contained in a block output by an honest party if *all* honest parties receive tx as input. A stronger definition would be to require this to hold even if only a *single* honest party receives tx as input. It is easy to achieve the latter from the former, however, by simply having honest parties gossip all transactions they receive to the rest of the network.

4 An ACS Protocol with Higher Validity Threshold

Throughout this section, we assume an asynchronous network. We construct an ACS protocol that is secure when the number of corrupted

parties is below one threshold, and provides validity even for some higher corruption threshold. That is, fix $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$. We show an ACS protocol that is t_a-secure, and achieves validity even for t_s corruptions. This protocol will be a key ingredient in our SMR protocol.

Our construction follows the high-level approach taken by Miller et al. [25], who devise an ACS protocol based on subprotocols for reliable broadcast and Byzantine agreement. In our case we need a reliable broadcast protocol that achieves validity for $t_s \geq n/3$ faults, and in Section 4.1 we show such a protocol. We then describe and analyze our ACS protocol in Section 4.2.

4.1 Reliable Broadcast with Higher Validity Threshold

In Figure 1, we present a variant of Bracha's (asynchronous) reliable broadcast protocol [6] that allows for a more general tradeoff between consistency and validity. Specifically, the protocol is parameterized by a threshold t_s; for any $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$, the protocol achieves t_a-consistency and t_s-validity.

Protocol $\Pi_{\mathsf{BB}}^{t_s}$

The sender P^* sends its input v^* to all parties. Then each party does:

- Upon receiving v^* from P^*, send (echo, v^*) to all parties.
- Upon receiving (echo, v^*) messages on the same value v^* from $n - t_s$ distinct parties, do: if (ready, v^*) was not yet sent, then send (ready, v^*) to all parties.
- Upon receiving (ready, v^*) messages on the same value v^* from $t_s + 1$ distinct parties, do: if (ready, v^*) was not yet sent, then send (ready, v^*) to all parties.
- Upon receiving (ready, v^*) messages on the same value v^* from $n - t_s$ distinct parties, output v^* and terminate.

Fig. 1. Bracha's reliable broadcast protocol, parameterized by t_s.

Lemma 1. *If $t_s < n/2$ then $\Pi_{\mathsf{BB}}^{t_s}$ is t_s-valid.*

Proof. Assume there are at most t_s corrupted parties, and the sender is honest. All honest parties receive the same value v^* from the sender, and consequently send (echo, v^*) to all other parties. Since there are at least $n - t_s$ honest parties, all honest parties receive (echo, v^*) from at least $n - t_s$ different parties, and as a result send (ready, v^*) to all other parties.

By the same argument, all honest parties receive (ready, v^*) from at least $n - t_s$ parties, and so can output v^* (and terminate).

Fix any $v \neq v^*$. To complete the proof, we argue that no honest party will output v. Note first that no honest party will send (echo, v). Thus, any honest party will receive (echo, v) from at most t_s other parties. Since $t_s < n - t_s$, no honest party will ever send (ready, v). By the same argument, this shows that honest parties will receive (ready, v) from at most t_s other parties, and hence will not output v. □

Lemma 2. *Fix $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{BB}}^{t_s}$ is t_a-consistent.*

Proof. Suppose at most t_a parties are corrupted, and that an honest party P_i outputs v. Then P_i must have received (ready, v) messages from at least $n - t_s$ distinct parties, at least $n - t_s - t_a \geq t_s + 1$ of whom are honest. Thus, all honest parties receive (ready, v) messages from at least $t_s + 1$ distinct parties, and so all honest parties send (ready, v) messages to everyone. It follows that all honest parties receive (ready, v) messages from at least $n - t_a \geq n - t_s$ parties, and so can output v as well.

To complete the proof, we argue that honest parties cannot output $v' \neq v$. We argued above that all honest parties send (ready, v) to everyone. Let P be the first honest party to do so. Since $t_a < t_s + 1$, that party must have sent (ready, v) in response to receiving (echo, v) messages from at least $n - t_s$ distinct parties. If some honest P_j outputs v' then, arguing similarly, some honest party P' must have received (echo, v') messages from at least $n - t_s$ distinct parties. But this is a contradiction, since honest parties send only a single echo message but $2 \cdot (n - t_s) - t_a > n$. □

4.2 ACS with Higher Validity Threshold

In Figure 2 we describe an ACS protocol $\Pi_{\mathsf{ACS}}^{t_a, t_s}$ that is parameterized by thresholds t_a, t_s, where $t_a \leq t_s$ and $t_a + 2 \cdot t_s < n$. Our protocol relies on two subprotocols: a reliable broadcast protocol Bcast that is t_s-valid and t_a-consistent (such as the protocol $\Pi_{\mathsf{BB}}^{t_s}$ from the previous section), and a Byzantine agreement protocol BA that is t_a-secure (since $t_a < n/3$, any asynchronous BA protocol secure for that threshold can be used). Our ACS protocol runs several executions of these protocols as sub-routines, so to distinguish between them we denote the ith execution by Bcast_i, resp., BA_i, and say that these executions *correspond to party P_i*.

Lemma 3. *If $t_a + 2 \cdot t_s < n$, then $\Pi_{\mathsf{ACS}}^{t_a, t_s}$ is t_s-valid.*

Protocol $\Pi_{\mathsf{ACS}}^{t_a,t_s}$

At any point during a party's execution of the protocol, let $S^* \stackrel{\text{def}}{=} \{i : \mathsf{BA}_i \text{ output } 1\}$ and let $s = |S^*|$. Define the following boolean conditions:

- $C_1(v)$: at least $n - t_s$ executions $\{\mathsf{Bcast}_i\}_{i\in[n]}$ have output v.
- C_1: $\exists v$ for which $C_1(v)$ is true.
- $C_2(v)$: $s \geq n - t_a$, all executions $\{\mathsf{BA}_i\}_{i\in[n]}$ have terminated, and a majority of the executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ have output v.
- C_2: $\exists v$ for which $C_2(v)$ is true.
- C_3: $s \geq n - t_a$, all executions $\{\mathsf{BA}_i\}_{i\in[n]}$ have terminated, and all executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ have terminated.

Each party does:

- For all i: run Bcast_i with P_i as the sender, where P_i uses input v_i.
- When Bcast_i terminates with output v_i' do: if execution of BA_i has not yet begun, run BA_i using input 1.
- When $s \geq n - t_a$, run any executions $\{\mathsf{BA}_i\}_{i\in[n]}$ that have not yet begun, using input 0.
- **(Exit 1:)** If at any point $C_1(v)$ for some v, output $\{v\}$.
- **(Exit 2:)** If at any point $\neg C_1 \wedge C_2(v)$ for some v, output $\{v\}$.
- **(Exit 3:)** If at any point $\neg C_1 \wedge \neg C_2 \wedge C_3$, output $S := \{v_i'\}_{i\in S^*}$.

After outputting:

- Continue to participate in any ongoing Bcast executions.
- Once $C_1 = \mathsf{true}$, stop participating in any ongoing BA executions.

Fig. 2. An ACS protocol, parameterized by t_a and t_s.

Proof. Note that $t_s < n/2$. Say at most t_s parties are dishonest, and all honest parties have the same input v. By t_s-validity of Bcast, at least $n-t_s$ executions of $\{\mathsf{Bcast}_i\}$ (namely, those for which P_i is honest) will result in v as output, and so all honest parties can take Exit 1 and output $\{v\}$. It is not possible for an honest party to take Exit 1 and output something other than $\{v\}$, since $t_s < n - t_s$. Thus, it only remains to show that if an honest party takes some other exit then it must also output $\{v\}$. Consider the two possibilities:

Exit 2: Suppose some honest party P takes Exit 2 and outputs $\{v'\}$. Then, for that party, $C_2(v')$ is true, and so P must have seen at least $\lfloor \frac{s}{2} \rfloor + 1$ of the $\{\mathsf{Bcast}_i\}_{i\in S^*}$ terminate with output v'. Moreover, P must have $s \geq n - t_a$. Together, these imply that P has seen at least

$$\left\lfloor \frac{n - t_a}{2} \right\rfloor + 1 \geq \left\lfloor \frac{2t_s}{2} \right\rfloor + 1 > t_s$$

executions of $\{\mathsf{Bcast}_i\}$ terminate with output v'. At least one of those executions must correspond to an honest party. But then t_s-validity of Bcast implies that $v' = v$.

Exit 3: Assume an honest party P takes Exit 3. Then P must have $s \geq n - t_a$, must have seen all executions $\{\mathsf{BA}_i\}_{i\in[n]}$ terminate, and must also have seen all executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ terminate. Because

$$|S^*| = s \geq n - t_a > 2t_s,$$

a majority of the executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ that P has seen terminate must correspond to honest parties. By t_s-validity of Bcast, all those executions must have resulted in output v. But then $C_2(v)$ must be true for P, and it would not have taken Exit 3. □

Lemma 4. *Fix $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$, and say at most t_a parties are corrupted. If honest parties P_1, P_2 output sets S_1, S_2, then $S_1 = S_2$.*

Proof. We consider different cases based on the possible exits taken by P_1 and P_2, and show that in all cases their outputs agree.

Case 1: *Either P_1 or P_2 takes Exit 1.* Without loss of generality, assume P_1 takes Exit 1 and outputs $\{v_1\}$. We consider different sub-cases:

- *P_2 takes Exit 1:* Say P_2 outputs $\{v_2\}$. Then P_1 and P_2 must have each seen at least $n - t_s$ executions of $\{\mathsf{Bcast}_i\}$ output v_1 and v_2, respectively. Since $t_s < n/2$, at least one of those executions must be the same. But then t_a-consistency of Bcast implies that $v_1 = v_2$.
- *P_2 takes Exit 2:* Say P_2 outputs $\{v_2\}$. For $C_2(v_2)$ to be satisfied, P_2 must have $s \geq n - t_a$, and must have seen at least

  $$\left\lfloor \frac{s}{2} \right\rfloor + 1 \geq \left\lfloor \frac{n - t_a}{2} \right\rfloor + 1$$

  executions of $\{\mathsf{Bcast}_i\}$ output v_2. As above, P_1 must have seen at least $n - t_s$ executions of $\{\mathsf{Bcast}_i\}$ output v_1. But since

  $$(n - t_s) + \left\lfloor \frac{n - t_a}{2} \right\rfloor + 1 \geq n - t_s + \left\lfloor \frac{2t_s}{2} \right\rfloor + 1 > n,$$

  at least one of those executions must be the same. But then t_a-consistency of Bcast implies that $v_1 = v_2$.

– P_2 *takes Exit 3:* We claim this cannot occur. Indeed, if P_2 takes Exit 3 then P_2 must have $s \geq n - t_a$, and must have seen all executions $\{\mathsf{BA}_i\}_{i\in[n]}$ terminate and all executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ terminate. Because P_1 took Exit 1, P_1 must have seen at least $n - t_s$ executions $\{\mathsf{Bcast}_i\}_{i\in[n]}$ output v_1, and therefore (by t_a-consistency of Bcast) there are at most t_s executions $\{\mathsf{Bcast}_i\}_{i\in[n]}$ that P_2 has seen terminate with a value other than v_1. The number of executions of $\{\mathsf{Bcast}_i\}_{i\in S^*}$ that P_2 has seen terminate with output v_1 is therefore at least $(n-t_a)-t_s > t_s$, which is strictly greater than the number of executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ that P_2 has seen terminate with a value other than v_1. But then $C_2(v_1)$ is true for P_2, and it would not take Exit 3.

Case 2: *Neither P_1 nor P_2 takes Exit 1.* We consider two sub-cases:

– P_1 *and P_2 both take Exit 2.* Say P_1 outputs $\{v_1\}$ and P_2 outputs $\{v_2\}$. Both P_1 and P_2 must have seen all executions $\{\mathsf{BA}_i\}_{i\in[n]}$ terminate; by t_a-consistency of BA they must therefore hold the same S^*. Since $C_2(v_1)$ holds for P_1, it must have seen a majority of the executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ output v_1; similarly, P_2 must have seen a majority of the executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ output v_2. Then t_a-consistency of Bcast implies $v_1 = v_2$.
– *Either P_1 or P_2 takes Exit 3.* Say P_1 takes Exit 3. (The case where P_2 takes Exit 3 is symmetric.) As above, P_1 and P_2 agree on S^* (this holds regardless of whether P_2 takes Exit 2 or Exit 3). Since C_3 holds for P_1 but C_2 does not, P_1 must have seen all executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ terminate but without any value being output by a majority of those executions. But then t_a-consistency of Bcast implies that P_2 also does not see any value being output by a majority of those executions, and so will not take Exit 2. Since P_2 instead must take Exit 3, it must have seen all executions $\{\mathsf{Bcast}_i\}_{i\in S^*}$ terminate; t_a-consistency of Bcast then implies that P_2 outputs the same set as P_1.

This completes the proof. □

Lemma 5. *Fix $t_a \leq t_s$ and $t_a + 2 \cdot t_s < n$. Then $\Pi^{t_a,t_s}_{\mathsf{ACS}}$ is t_a-live.*

Proof. If some honest party P takes Exit 1 during an execution of $\Pi^{t_a,t_s}_{\mathsf{ACS}}$, then P must have seen at least $n - t_s$ executions $\{\mathsf{Bcast}_i\}_{i\in[n]}$ with the same output v. By t_a-consistency of Bcast, all other honest parties will eventually see at least those $n-t_s$ executions output v, and will generate output (if they have not already generated output via another exit).

It remains to consider the case where no honest parties take Exit 1. Let H be the indices of parties who remain honest, with $|H| \geq n - t_a$. By t_s-validity of Bcast, all honest parties see the executions $\{\mathsf{Bcast}_i\}_{i \in H}$ terminate, and so all honest parties initiate the executions $\{\mathsf{BA}_i\}_{i \in H}$. Since no honest party takes Exit 1, all honest parties continue to participate in all those executions. Consider some execution BA_i being run by all honest parties. As long as no honest party has $s \geq n - t_a$, each honest party must be running BA_i using input 1. By t_a-validity of BA, this means that all honest parties will eventually output 1 from that execution. We conclude from this that some honest party will eventually have $s \geq n - t_a$; furthermore, t_a-consistency of BA then implies that all honest parties will eventually have $s \geq n - t_a$. This means that all honest parties execute all $\{\mathsf{BA}_i\}_{i \in [n]}$, and by t_a-security of BA all those executions eventually terminate. Define $\hat{S}^* \stackrel{\text{def}}{=} \{i : \text{some honest player outputs 1 in } \mathsf{BA}_i\}$. We claim that all executions $\{\mathsf{Bcast}_i\}_{i \in \hat{S}^*}$ eventually terminate. To see this, fix $i \in \hat{S}^*$. Then by t_a-validity of BA, some honest party P must have used input 1 to BA_i. But that implies that Bcast_i must have terminated for P. So t_a-consistency of Bcast implies that Bcast_i will terminate for all honest parties. It follows that any honest party can take Exit 2 or 3. □

Lemma 6. *Fix $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{ACS}}^{t_a, t_s}$ has t_a-set quality.*

Proof. Consider some honest party P. Say P takes Exit 1 and outputs $S = \{v\}$. Then P has seen at least $n - t_s$ executions $\{\mathsf{Bcast}_i\}$ terminate with output v. Of these, at least $n - t_s - t_a > t_s \geq t_a$ must correspond to honest parties. By t_s-validity of Bcast, those honest parties all had input v. This means that S contains the inputs of at least $t_a + 1$ honest parties.

Alternatively, say P takes Exit 2 or 3 and outputs a set S. Then P must have $|S^*| \geq n - t_a$. At least

$$n - 2 \cdot t_a > \max\{(n - t_a)/2,\ t_a\}$$

of the indices in S^* correspond to honest parties, and t_s-validity of Bcast implies that for each of those parties the corresponding output value v_i' that P holds is equal to that party's input. Thus, regardless of whether P takes Exit 2 (and S contains the majority value output by $\{\mathsf{Bcast}_i\}_{i \in S^*}$) or Exit 3 (and S contains every value output by $\{\mathsf{Bcast}_i\}_{i \in S^*}$), the set S output by P contains the inputs of at least $t_a + 1$ honest parties. □

Theorem 1. *Fix t_a, t_s with $t_a \leq t_s$ and $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{ACS}}^{t_a, t_s}$ is t_a-secure and t_s-valid.*

Proof. Lemma 3 proves t_s-validity. Lemmas 4 and 5 together prove t_a-consistency, and Lemma 6 shows t_a-set quality. □

We note that $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ does not guarantee termination. Given that any SMR protocol itself runs indefinitely, it is reasonable for an SMR protocol to rely on an ACS protocol that runs forever so long as the ACS protocol has bounded communication complexity. We prove that $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ has bounded communication complexity in Lemma 7. (The local state that parties maintain in $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ may not be bounded, but since the state for any SMR protocol is also unbounded we consider this acceptable.)

Lemma 7. *Fix $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ has bounded communication complexity under either of the following conditions:*

1. *At most t_a parties are corrupted.*
2. *At most t_s parties are corrupted and all honest parties have the same input.*

Proof. Because Bcast from the previous section has bounded communication complexity, we only need to show that all honest parties eventually stop participating in all BA executions. (This can occur either because those executions all terminate, or because honest parties all set $C_1 = \mathsf{true}$ and stop participating in any still-running executions.)

Case 1: *At most t_a parties are corrupted.* If some honest party P takes Exit 1 during an execution of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$, then P must have seen at least $n - t_s$ executions $\{\mathsf{Bcast}_i\}_{i \in [n]}$ with the same output value. By t_a-consistency of Bcast, all honest parties eventually see those executions output the same value, and thus set $C_1 = \mathsf{true}$ and stop participating in any still-running BA executions.

On the other hand, if no honest parties take Exit 1 during an execution of BA, then all honest parties continue to participate in all BA executions. By t_a-termination of BA, each of those executions will terminate.

Case 2: *At most t_s parties are corrupted and all honest parties have the same input v.* Because all honest parties have input v, t_s-validity of Bcast implies that all honest parties receive output v from at least $n - t_s$ executions of $\{\mathsf{Bcast}_i\}$. So all honest parties will eventually set $C_1 = \mathsf{true}$ and thus stop participating in any still-running BA executions. □

5 A Network-Agnostic SMR Protocol

In this section, we show our main result: an SMR protocol that is t_s-secure in a synchronous network and t_a-secure in an asynchronous network. We

begin in Section 5.1 by briefly introducing a useful primitive called *block agreement*. In Appendix B, we construct a block-agreement protocol secure against $t < n/2$ parties in a synchronous network. We then use our block-agreement protocol to construct an SMR protocol in Section 5.2.

5.1 Block Agreement

Block agreement is a form of agreement where (1) in addition to an input, parties provide signatures (in a particular format) on those inputs, and (2) a stronger notion of validity is required. Specifically, consider pairs consisting of a block B along with a set Σ of signed buffers $\langle \mathsf{buf}_j \rangle_j$. (Recall that $\langle m \rangle_i$ denotes a tuple (i, m, σ) such that σ is a valid signature on m with respect to P_i's secret key.) We say a pair (B, Σ) is *t-valid* if:

- Σ contains signed buffers from strictly more than t distinct parties.
- For each $\langle \mathsf{buf}_j \rangle_j \in \Sigma$, we have $\mathsf{buf}_j \subseteq B$.

A pair is *valid* if it is 0-valid (meaning it contains signed buffers from at least one party).

Definition 5 (Block agreement). *Let Π be a protocol executed by parties $P_1, \ldots, P_n$, where each party P_i begins holding input (B_i, Σ_i) and parties terminate upon generating output.*

- **Validity:** *Π is t-*valid *if whenever at most t of the parties are corrupted, every honest party outputs either $\perp$ or a valid pair.*
- **Termination:** *Π is t-*terminating *if whenever at most t of the parties are corrupted, every honest party terminates.*
- **Consistency:** *Π is t-*consistent *if the following holds whenever at most t of the parties are corrupted: for any $s \leq t$, if every honest party inputs an s-valid pair, there is an s-valid (B, Σ) such that every honest party outputs (B, Σ).*

*If Π is t-valid, t-consistent, and t-terminating, then we say it is t-*secure*.*

We prove the following in Appendix B.

Theorem 2. *There is a block-agreement protocol Π_{BLA} that is t-secure for any $t < n/2$ when run in a synchronous network. Moreover, all honest parties terminate with probability $1 - 2^{-O(\kappa)}$ after time $\kappa \cdot \Delta$.*

5.2 State Machine Replication

We now combine our various sub-protocols to realize network-agnostic SMR. At a high level, our SMR protocol $\Pi_{\mathsf{SMR}}^{t_a,t_s}$ (see Figure 3) proceeds as follows. For each slot j, the parties attempt to reach agreement on a block using the block-agreement protocol Π_{BLA}. If that protocol terminates, parties use its output B as input to our ACS protocol $\Pi_{\mathsf{ACS}}^{t_a,t_s}$. If Π_{BLA} fails to terminate after a sufficiently long time, parties abandon it and instead attempt to reach agreement using the ACS protocol directly.

By setting the timeout appropriately, we can ensure that in a synchronous network Π_{BLA} terminates with overwhelming probability. Thus, if the network is synchronous and at most t_s parties are corrupted, all parties agree on their input B to $\Pi_{\mathsf{ACS}}^{t_a,t_s}$, and t_s-validity of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ ensures that all parties output B. On the other hand, if the network is asynchronous and at most t_a parties are corrupted, then t_a-security of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ ensures agreement.

Protocol $\Pi_{\mathsf{SMR}}^{t_a,t_s}$

We describe the protocol from the point of view of party P_i holding a set buf_i that grows asynchronously via some external process.

For $k = 1, \ldots,$ do the following starting at time $T_k := (\Delta + 5\Delta\kappa) \cdot (k-1)$:

1. Set $\mathsf{Epochs}_i[k] := 1$, and initialize $B := \emptyset, \Sigma := \emptyset$.
2. Send $\langle \mathsf{buf}_i \rangle_i$ to every party.
3. While $|\Sigma| \leq t_s$:
   - The first time $\langle \mathsf{buf}_j \rangle_j$ is received from P_j, set $B := B \cup \mathsf{buf}_j$ and $\Sigma := \Sigma \cup \{\langle \mathsf{buf}_j \rangle_j\}$.
4. At time $T_k + \Delta$, run Π_{BLA} on input (B, Σ).
5. If Π_{BLA} produces t_s-valid output, let (B^*, Σ^*) denote that output. Otherwise, at time $T_k + \Delta + \kappa \cdot \Delta$ set $(B^*, \Sigma^*) := (B, \Sigma)$.
6. Run $\mathsf{BlockSet} \leftarrow \Pi_{\mathsf{ACS}}^{t_a,t_s}$ using input B^*.
7. Set $\mathsf{Blocks}_i[k] := \bigcup_{\hat{B} \in \mathsf{BlockSet}} \hat{B}$. Set $\mathsf{buf}_i := \mathsf{buf}_i \setminus \mathsf{Blocks}_i[k]$.

Fig. 3. A protocol for state machine replication.

Theorem 3 (Consistency). *Fix t_a, t_s with $t_a < n/3$ and $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{SMR}}^{t_a,t_s}$ is t_a-consistent when run in an asynchronous network, and t_s-consistent when run in a synchronous network.*

Proof. Assume first that at most t_s parties are dishonest and the network is synchronous. In any slot k, each honest party receives $\langle \mathsf{buf}_j \rangle_j$ from at least the $n - t_s > t_s$ honest parties, and the input (B, Σ) they use to Π_{BLA}

is t_s-valid. Consistency of Π_{BLA} implies that every honest party outputs the same t_s-valid pair (B^*, Σ^*) after running Π_{BLA} for time $\kappa \cdot \Delta$. By t_s-validity of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$, this means every honest party obtains output $\{B^*\}$ from $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ and then sets $\mathsf{Blocks}[k] = B^*$.

If at most t_a parties are dishonest and the network is asynchronous, then t_a-consistency of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ implies that all honest parties agree on the same value $\mathsf{BlockSet}$, and hence set $\mathsf{Blocks}[k]$ to the same value. □

Theorem 4 (Strong liveness). *Fix $t_a \leq t_s$ with $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{SMR}}^{t_a,t_s}$ is t_a-live when run in an asynchronous network, and t_s-live when run in a synchronous network.*

Proof. By consistency of $\Pi_{\mathsf{SMR}}^{t_a,t_s}$, we can refer to the values of $\mathsf{Blocks}[i]$ without specifying any particular party. Consider some transaction tx that every honest party received before entering epoch k. If tx appears in $\mathsf{Blocks}[k']$ for some $k' < k$ then we are done. Otherwise, every honest party has tx in their buffer when they enter epoch k. We show that in this latter case, tx is in $\mathsf{Blocks}[k]$.

Assume at most t_s parties are corrupted and the network is synchronous. Reasoning as in the proof of Theorem 3, every honest party outputs the same t_s-valid pair (B^*, Σ^*) after running Π_{BLA} for time $\kappa \cdot \Delta$, and sets $\mathsf{Blocks}[k] = B^*$. Since (B^*, Σ^*) is t_s-valid, Σ^* must contain a signature on a subset of B^* from at least one honest party. But an honest party would have only signed a subset that includes tx, implying $\mathsf{tx} \in B^*$.

Consider next the case where at most t_a parties are dishonest and the network is asynchronous. Every honest party P_i runs $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ using an input B_i^* for which they have a t_s-valid pair (B_i^*, Σ_i^*). Arguing as above, each B_i^* must contain tx. By t_a-security of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$, all honest parties output the same set $\mathsf{BlockSet}$ that contains B_i^* for some honest party P_i, and hence contains tx. It follows that every honest party includes tx in $\mathsf{Blocks}[k]$. □

Theorem 5 (Completeness). *Fix t_a, t_s with $t_a < n/3$ and $t_a + 2 \cdot t_s < n$. Then $\Pi_{\mathsf{SMR}}^{t_a,t_s}$ is t_a-complete when run in an asynchronous network, and t_s-complete when run in a synchronous network.*

Proof. By inspection of $\Pi_{\mathsf{SMR}}^{t_a,t_s}$, a party outputs a block in slot k iff its execution of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ in iteration k produces output. So if at most t_a parties are corrupted, completeness follows from t_a-liveness of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$. If at most t_s parties are corrupted and the network is synchronous, then consistency of Π_{BLA} implies that all honest parties run $\Pi_{\mathsf{ACS}}^{t_a,t_s}$ using the same input; completeness then follows from t_s-validity of $\Pi_{\mathsf{ACS}}^{t_a,t_s}$. □

6 Optimality of Our Thresholds

In this section we show that the parameters achieved by our SMR protocol are optimal. This extends the analogous result by Blum et al. [5], who consider the case of BA. We remark that, although SMR is generally viewed as a stronger form of consensus than BA, it is unclear whether SMR generically implies BA in a network-agnostic setting, and we were not able to show such a result for the corruption thresholds of interest (namely, when $t_a + 2t_s \geq n$). We thus need to prove impossibility directly.

Lemma 8. *Fix t_a, t_s, n with $t_a + 2t_s \geq n$. If an n-party SMR protocol is t_s-live in a synchronous network, then it cannot also be t_a-consistent in an asynchronous network.*

Proof. Assume $t_a + 2t_s = n$ and fix an SMR protocol Π. Partition the n parties into sets S_0, S_1, S_a where $|S_0| = |S_1| = t_s$ and $|S_a| = t_a$, and consider the following experiment:

- Choose uniform $m_0, m_1 \in \{0,1\}^\kappa$.
- Parties in S_b begin running Π at global time 0 with their buffers containing only m_b. All communication between parties in S_0 and parties in S_1 is blocked (but all other messages are delivered within time Δ).
- Create virtual copies of each party in S_a, call them S_a^0 and S_a^1. Parties in S_a^b begin running Π (at global time 0) with their buffers containing only m_b, and communicate only with each other and parties in S_b.

Consider an execution of Π in a synchronous network where parties in S_1 are corrupted and simply abort. Uniform $m_0, m_1 \in \{0,1\}^\kappa$ are chosen, and the remaining (honest) parties start with their buffers containing only m_0. The views of the honest parties in this execution are distributed identically to the views of $S_0 \cup S_a^0$ in the above experiment. In particular, t_s-liveness of Π implies that, in the above experiment, all parties in S_0 include m_0 in $\mathsf{Blocks}[1]$. Moreover, since parties in S_0 have no information about m_1, they include m_1 in $\mathsf{Blocks}[1]$ with negligible probability. Analogously, all parties in S_1 include m_1 in $\mathsf{Blocks}[1]$ but include m_0 in $\mathsf{Blocks}[1]$ with negligible probability.

Next consider an execution of Π in an asynchronous network where parties in S_a are corrupted, and run Π with their buffers containing m_0 when interacting with S_0 while running Π with their buffers containing m_1 when interacting with S_1. Moreover, all communication between the (honest) parties in S_0 and S_1 is delayed indefinitely. The views of the

honest parties here are distributed identically to the views of $S_0 \cup S_1$ in the above experiment, yet the conclusion of the preceding paragraph shows that t_a-consistency is violated with overwhelming probability. □

Acknowledgments

Work supported in part under financial assistance award 70NANB19H126 from the U.S. Department of Commerce, National Institute of Standards and Technology, and NSF award #1837517.

References

1. Ittai Abraham, Srinivas Devadas, Danny Dolev, Kartik Nayak, and Ling Ren. Efficient synchronous Byzantine consensus, 2017. Available at `https://eprint.iacr.org/2017/307`.
2. Ittai Abraham, Dahlia Malkhi, Kartik Nayak, Ling Ren, and Maofan Yin. Sync HotStuff: Simple and practical synchronous state machine replication, 2019. Available at `http://eprint.iacr.org/2019/270`.
3. Zuzana Beerliová-Trubíniová, Martin Hirt, and Jesper Buus Nielsen. On the theoretical gap between synchronous and asynchronous MPC protocols. In *29th Annual ACM Symp. on Principles of Distributed Computing (PODC)*, pages 211–218. ACM Press, 2010.
4. Michael Ben-Or, Boaz Kelmer, and Tal Rabin. Asynchronous secure computations with optimal resilience. In *13th Annual ACM Symp. on Principles of Distributed Computing (PODC)*, pages 183–192. ACM Press, August 1994.
5. Erica Blum, Jonathan Katz, and Julian Loss. Synchronous consensus with optimal asynchronous fallback guarantees. In *14th Theory of Cryptography Conference—TCC 2019*, volume 11891 of *LNCS*. Springer, 2019. Available at `https://eprint.iacr.org/2019/692`.
6. Gabriel Bracha. An asynchronous $\lfloor (n-1)/3 \rfloor$-resilient consensus protocol. In *3rd Annual ACM Symp. on Principles of Distributed Computing (PODC)*, pages 154–162. ACM Press, 1984.
7. Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Trans. Computer Systems*, 20(4):398–461, 2002.
8. Miguel Correia, Nuno Ferreira Neves, and Paulo Veríssimo. From consensus to atomic broadcast: Time-free Byzantine-resistant protocols without signatures. *The Computer Journal*, 49(1):82–96, 2006.
9. Ivan Damgård, Martin Geisler, Mikkel Krøigaard, and Jesper Buus Nielsen. Asynchronous multiparty computation: Theory and implementation. In *12th Intl. Conference on Theory and Practice of Public Key Cryptography—PKC 2009*, volume 5443 of *LNCS*, pages 160–179. Springer, 2009.
10. Matthias Fitzi and Jesper Buus Nielsen. On the number of synchronous rounds sufficient for authenticated Byzantine agreement. In *23rd Intl. Symp. on Distributed Computing (DISC)*, volume 5805 of *LNCS*, pages 449–463. Springer, 2009.
11. Juan A. Garay, Jonathan Katz, Ranjit Kumaresan, and Hong-Sheng Zhou. Adaptively secure broadcast, revisited. In *30th Annual ACM Symp. on Principles of Distributed Computing (PODC)*, pages 179–186. ACM Press, 2011.

12. Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology—Eurocrypt 2015, Part II*, volume 9057 of *LNCS*, pages 281–310. Springer, 2015.
13. Yue Guo, Rafael Pass, and Elaine Shi. Synchronous, with a chance of partition tolerance. In *Advances in Cryptology—Crypto 2019, Part I*, volume 11692 of *LNCS*, pages 499–529. Springer, 2019.
14. Martin Hirt and Vassilis Zikas. Adaptively secure broadcast. In *Advances in Cryptology—Eurocrypt 2010*, volume 6110 of *LNCS*, pages 466–485. Springer, 2010.
15. Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for Byzantine agreement. *J. Computer and System Sciences*, 75(2):91–112, 2009.
16. Ramakrishna Kotla, Lorenzo Alvisi, Michael Dahlin, Allen Clement, and Edmund L. Wong. Zyzzyva: Speculative Byzantine fault tolerance. *ACM Trans. Computer Systems*, 27(4):7:1–7:39, 2009.
17. Klaus Kursawe. Optimistic Byzantine agreement. In *21st Symposium on Reliable Distributed Systems (SRDS)*, pages 262–267. IEEE Computer Society, 2002.
18. Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7), 1978.
19. Leslie Lamport. The part-time parliament. Technical Report 49, DEC Systems Research Center, 1989.
20. Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The Byzantine generals problem. *ACM Trans. Programming Language Systems*, 4(3):382–401, 1982.
21. Shengyun Liu, Paolo Viotti, Christian Cachin, Vivien Quéma, and Marko Vukolic. XFT: Practical fault tolerance beyond crashes. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 485–500. USENIX Association, 2016.
22. Chen-Da Liu-Zhang, Julian Loss, Tal Moran, Ueli Maurer, and Daniel Tschudi. Robust MPC: Asynchronous responsiveness yet synchronous security. Unpublished manuscript.
23. Julian Loss and Tal Moran. Combining asynchronous and synchronous Byzantine agreement: The best of both worlds, 2018. Available at `http://eprint.iacr.org/2018/235`.
24. Dahlia Malkhi, Kartik Nayak, and Ling Ren. Flexible Byzantine fault tolerance. In *26th ACM Conf. on Computer and Communications Security (CCS)*, pages 1041–1053. ACM Press, 2019. Available at `https://arxiv.org/abs/1904.10067`.
25. Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of BFT protocols. In *23rd ACM Conf. on Computer and Communications Security (CCS)*, pages 31–42. ACM Press, 2016.
26. Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Advances in Cryptology—Eurocrypt 2017, Part II*, volume 10211 of *LNCS*, pages 643–673. Springer, 2017.
27. Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. In *31st International Symposium on Distributed Computing (DISC)*, volume 91 of *LIPIcs*, pages 39:1–39:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
28. Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. In *Advances in Cryptology—Eurocrypt 2018, Part II*, volume 10821 of *LNCS*, pages 3–33. Springer, 2018.
29. Arpita Patra and Divya Ravi. On the power of hybrid networks in multi-party computation. *IEEE Trans. Information Theory*, 64(6):4207–4227, 2018.

30. M. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.
31. Fred Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, 1990.

A SMR Implies Weak BA

We briefly discuss how SMR relates to BA. Specifically, we show that SMR implies *weak* BA. A weak BA protocol Π satisfies validity and consistency as in Definition 2, but instead of termination it achieves a weaker liveness property. Namely, we say that Π is t-live if whenever at most t parties are corrupted, every honest party outputs a value in $\{0, 1\}$ (but may not terminate).

In Figure 4 we show how to use an SMR protocol Π_{SMR} to achieve weak BA.

Protocol $\Pi^{t_s}_{\mathsf{WBA}}$

We describe the protocol from the point of view of a party P_i with input v_i.

- Set $V_i := \mathsf{buf}_i := \emptyset$.
- Send $\langle v_i \rangle_i$ to every party. Upon receiving $\langle v_j \rangle_j$ from party P_j, set $\mathsf{buf}_i := \mathsf{buf}_i \cup \{\langle v_j \rangle_j\}$.
- Begin to run Π_{SMR} at time Δ.
- Upon outputting a block $B = \mathsf{Blocks}_i[k]$ do: for all j such that B contains $\langle v_j \rangle_j$ and there is no pair $(\star, j)$ in V, add (v_j, j) to V.
- If at any point during the execution $|V_i| \geq n - t_s$, then output the majority value among all values in V_i.

Fig. 4. A protocol for weak Byzantine agreement, parameterized by t_s.

Lemma 9 (Validity and liveness). *Let $t_a + 2t_s < n$. If Π_{SMR} is t_s-live in a synchronous network (resp., t_a-live in an asynchronous network), then $\Pi^{t_s}_{\mathsf{WBA}}$ is t_s-valid and t_s-live in a synchronous network (resp., t_a-valid and t_a-live in an asynchronous network).*

Proof. Assume all honest parties hold input v. Consider first the case where at most t_s parties are corrupted and the network is synchronous. The initial message from each honest party is received by all other honest parties by time Δ. By t_s-liveness of Π_{SMR}, the block $B = \mathsf{Blocks}[1]$ output by any honest party contains (v, i) for each honest party P_i. At that point,

each honest party will have $|V| \geq n - t_s$, and since $t_s < n/2$ the majority value in V will be v. Thus, all honest parties output v.

Next, consider the case where there are at most t_a corrupted parties and the network is asynchronous. If some honest party has $|V| \geq n - t_s$, then at least $n - t_s - t_a > t_a$ of those values correspond to honest parties, and hence v will be the majority value. Thus, any honest party who outputs anything will output v. It remains to show that all honest parties eventually have $|V| \geq n - t_s$. This follows from the fact that honest parties' initial messages are eventually delivered to all honest parties, along with t_a-liveness of Π_{SMR}. □

Lemma 10 (Consistency). *For all t, if Π_{SMR} is t-consistent in a synchronous (resp., asynchronous) network, then $\Pi^{t_s}_{\mathsf{WBA}}$ is t-consistent in a synchronous (resp., asynchronous) network.*

Proof. The lemma is immediate. □

B A Block-Agreement Protocol

Throughout this section, we assume a synchronous network.

The structure of our block-agreement protocol is inspired by the *synod protocol* of Abraham et al. [1]. We construct our protocol in a modular fashion. We begin by defining a subprotocol $\Pi^{P^*}_{\mathsf{Propose}}$ (see Figure 5) in which a designated party P^* serves as a *proposer*. A tuple (k, B, Σ, C) is called a *k-vote on* (B, Σ) if (B, Σ) is valid and either:

- $k = 0$, or
- $k > 0$ and C is a set of valid signatures from a majority of the parties on messages of the form $(\mathsf{Commit}, k', B, \Sigma)$ with $k' \geq k$ (where possibly different k' can be used in different messages).

When the exact value of k is unimportant, we simply refer to the tuple as a *vote*. A message of the form $\mathsf{status} = \langle \mathsf{Status}, k, B, \Sigma, C \rangle_i$ is a *correctly formed* Status *message (from party P_i)* if (k, B, Σ, C) is a vote. A message $\langle \mathsf{Propose}, \mathsf{status}_1, \ldots \rangle_*$ is a *correctly formed* Propose *message* if it contains correctly formed Status messages from a majority of the parties.

We first show that any two honest parties who generate output in this protocol agree on their output.

Lemma 11. *If honest parties P_i and P_j output $(B_i, \Sigma_i), (B_j, \Sigma_j) \neq \perp$, respectively, in an execution of $\Pi^{P^*}_{\mathsf{Propose}}$, then $(B_i, \Sigma_i) = (B_j, \Sigma_j)$.*

Protocol $\Pi^{P^*}_{\mathsf{Propose}}$

We describe the protocol from the point of view of a party P_i with input a vote (k, B, Σ, C). Let $t = \lceil (n+1)/2 \rceil$.

1. At time 0, send $\mathsf{status}_i := \langle \mathsf{Status}, k, B, \Sigma, C \rangle_i$ to P^*.
2. At time Δ, if P^* has received at least $s \geq t$ correctly formed Status messages $\mathsf{status}_1, \ldots, \mathsf{status}_t$ (from distinct parties), then P^* sets

$$m := (\mathsf{Propose}, \mathsf{status}_1, \ldots, \mathsf{status}_s),$$

   and sends $\langle m \rangle_*$ to all parties.
3. At time 2Δ, if a correctly formed $\mathsf{Propose}$ message $\langle m \rangle_*$ has been received from P^*, then send $\langle m \rangle_*$ to all parties. Otherwise, output $\perp$.
4. At time 3Δ, let $\langle m \rangle^j_*$ be the correctly formed $\mathsf{Propose}$ message received from P_j (if any). If there exists j such that $\langle m \rangle^j_* \neq \langle m \rangle_*$, output $\perp$. Otherwise, let $\mathsf{status}_{\mathsf{max}} = \langle \mathsf{Status}, k', B', \Sigma', C' \rangle$ be the status message in $\langle m \rangle_*$ with maximal k' (picking the lowest index in case of ties). Output (B', Σ').

Fig. 5. A protocol $\Pi^{P^*}_{\mathsf{Propose}}$ with designated proposer P^*.

Proof. If P_i outputs $(B_i, \Sigma_i) \neq \perp$, then P_i must have received a correctly formed $\mathsf{Propose}$ message $\langle m \rangle_*$ by time 2Δ that would cause it to output (B_i, Σ_i). That message is forwarded by P_i to P_j, and hence P_j either outputs $\perp$ (if it detects an inconsistency) or the same value (B_i, Σ_i). □

Assume less than half the parties are corrupted. We show that if there is some (B, Σ) such that the input of each honest party P_i is a vote of the form (k_i, B, Σ, C_i), and no honest party ever receives a vote (k', B', Σ', C') with $k' \geq \min_i\{k_i\}$ and $(B', \Sigma') \neq (B, \Sigma)$, then the only value an honest party can output is (B, Σ).

Lemma 12. *Assume fewer than $n/2$ parties are corrupted, and that the input of each honest party P_i to $\Pi^{P^*}_{\mathsf{Propose}}$ is a k_i-vote on (B, Σ). If no honest party ever receives a k'-vote on $(B', \Sigma') \neq (B, \Sigma)$ with $k' \geq \min_i\{k_i\}$, then every honest party outputs either (B, Σ) or $\perp$.*

Proof. Consider an honest party P who does not output $\perp$. That party must have received a correctly formed $\mathsf{Propose}$ message $\langle m \rangle_*$ from P^*, which in turn must contain a correctly formed Status message from at least one honest party P_i. That Status message contains a vote (k_i, B, Σ, C_i) and, under the assumptions of the lemma, any other vote (k', B', Σ', C') contained in $\langle m \rangle_*$ with $k' \geq k_i$ has $(B', \Sigma') = (B, \Sigma)$. It follows that P outputs (B, Σ). □

Finally, we show that when P^* is honest then all honest parties do indeed generate output.

Lemma 13. *Assume fewer than $n/2$ parties are corrupted. If every honest party's input to $\Pi^{P^*}_{\mathsf{Propose}}$ is a vote and P^* is honest, then every honest party outputs the same valid $(B, \Sigma) \neq \perp$.*

Proof. Since every honest party's input is a vote, P^* will receive at least $\lceil (n+1)/2 \rceil$ correctly formed Status messages, and so sends a correctly formed Propose message to all honest parties. Since P^* is honest, this is the only correctly formed Propose message the honest parties will receive, and so all honest parties will output the same valid $(B, \Sigma) \neq \perp$. □

We now present a protocol Π^k_{GC} that uses $\Pi^{P^*}_{\mathsf{Propose}}$ to achieve a form of graded consensus on a valid pair (B, Σ). (See Figure 6.) As in the protocol of Abraham et al. [1], we rely on an atomic leader-election mechanism Leader with the following properties: On input k from a majority of parties, Leader chooses a uniform leader $\ell \in \{1, \ldots, n\}$ and sends (k, ℓ) to all parties. This ensures that if less than half of all parties are corrupted, then at least one honest party must call Leader with input k before the adversary can learn the identity of ℓ. A leader-election mechanism tolerating any $t < n/2$ faults can be realized (in the synchronous model with a PKI) based on general assumptions [15]; it can also be realized more efficiently using a threshold unique signature scheme.

Below, we refer to a message $\langle \mathsf{Commit}, k, B, \Sigma \rangle_i$ as a *correctly formed* Commit *message* (*from P_i on (B, Σ)*) if (B, Σ) is valid. We refer to a message $(\mathsf{Notify}, k, B, \Sigma, C)$ as a *correctly formed* Notify *message on* (B, Σ) if (B, Σ) is valid and C is a set of valid signatures on $(\mathsf{Commit}, k, B, \Sigma)$ from more than $n/2$ parties; in that case, C is called a *k-certificate for* (B, Σ). For an output $((B, \Sigma, C), g)$, we refer to g as the *grade* and (B, Σ, C) as the *output*. When a party's output is (B, Σ, C), we may also say that its output is a k-certificate for (B, Σ).

Lemma 14. *Assume fewer than $n/2$ parties are corrupted, and that the input of each honest party P_i to Π^k_{GC} is a k_i-vote on (B, Σ). If no honest party ever receives a k'-vote on $(B', \Sigma') \neq (B, \Sigma)$ with $k' \geq \min_i\{k_i\}$ in step 1 of Π^k_{GC}, then (1) no honest party sends a Commit message on $(B', \Sigma') \neq (B, \Sigma)$ and (2) any honest party who outputs a nonzero grade outputs a k-certificate for (B, Σ).*

Proof. By Lemma 12, every honest party outputs either (B, Σ) or $\perp$ in every execution of Π_{Propose} in step 1. It follows that no honest party P_i

Protocol Π^k_{GC}

We describe the protocol from the point of view of a party P_i with input a vote (k', B, Σ, C'). Let $t = \lceil (n+1)/2 \rceil$.

1. At time 0, run parallel executions of $\Pi^{P_1}_{\mathsf{Propose}}, \ldots, \Pi^{P_n}_{\mathsf{Propose}}$, each using input (k', B, Σ, C'). Let (B_j, Σ_j) be the output from the jth protocol.
2. At time 3Δ, call $\mathsf{Leader}(k)$ to obtain the response ℓ. If $(B_\ell, \Sigma_\ell) \neq \perp$, send $\langle \mathsf{Commit}, k, B_\ell, \Sigma_\ell \rangle_i$ to every party.
3. At time 4Δ, if at least t correctly formed Commit messages $\langle \mathsf{Commit}, k, B_\ell, \Sigma_\ell \rangle_j$ from distinct parties have been received, then form a k-certificate C for (B_ℓ, Σ_ℓ), send $m := (\mathsf{Notify}, k, B_\ell, \Sigma_\ell, C)$ to every party, output $((B_\ell, \Sigma_\ell, C), 2)$, and terminate.
4. At time 5Δ, if a correctly formed Notify message $(\mathsf{Notify}, k, B, \Sigma, C)$ has been received, output $((B, \Sigma, C), 1)$ and terminate. (If there is more than one such message, choose arbitrarily.) Otherwise, output $(\perp, 0)$ and terminate.

Fig. 6. A graded block-consensus protocol Π^k_{GC}, parameterized by k.

sends a Commit message on $(B', \Sigma') \neq (B, \Sigma)$, proving the first part of the lemma. Since less than half the parties are corrupted, this means an honest party will receive fewer than $\lceil (n+1)/2 \rceil$ correctly formed Commit messages on anything other than (B, Σ); it follows that if an honest party outputs grade $g = 2$ then that party outputs (B, Σ, C) with C a k-certificate for (B, Σ).

Arguing similarly, no honest party will receive a correctly formed Notify message on anything other than (B, Σ). Hence any honest party that outputs grade 1 outputs (B, Σ, C) with C a k-certificate for (B, Σ). □

Lemma 15. *Assume fewer than $n/2$ parties are corrupted. If an honest party outputs (B, Σ, C) with a nonzero grade in an execution of Π^k_{GC}, then no honest party sends a* Commit *message on $(B', \Sigma') \neq (B, \Sigma)$.*

Proof. Say an honest party outputs (B, Σ, C) with a nonzero grade. That party must have received a correctly formed Notify message on (B, Σ). Since that Notify message includes a k-certificate C with signatures from more than half the parties, at least one honest party P must have sent a Commit message on (B, Σ). This means that P must have received (B, Σ) as its output from $\Pi^{P_\ell}_{\mathsf{Propose}}$. By Lemma 11, this means the output of any other honest party from $\Pi^{P_\ell}_{\mathsf{Propose}}$ is either (B, Σ) or $\perp$. □

Lemma 16. *Assume fewer than $n/2$ parties are corrupted. If an honest party outputs (B, Σ, C) with grade 2 in an execution of Π^k_{GC}, then every honest party outputs a k-certificate on (B, Σ) with a nonzero grade.*

Proof. Say an honest party P outputs (B, Σ, C) with a grade of 2. By Lemma 15, this means no honest party sent a correctly formed Commit message on $(B', \Sigma') \neq (B, \Sigma)$; it is thus impossible for any honest party to output $(B', \Sigma') \neq (B, \Sigma)$ with a nonzero grade. Since P sends a correctly formed Notify message on (B, Σ) to all honest parties, every honest party will output (B, Σ) with a nonzero grade. □

Lemma 17. *Assume fewer than $n/2$ parties are corrupted. Then with probability at least $1/2$ every honest party outputs a k-certificate on the same valid (B, Σ) with a grade of 2.*

Proof. The leader ℓ chosen in step 2 was honest in step 1 with probability at least $1/2$. We show that whenever this occurs, every honest party outputs grade 2. Agreement on a valid (B, Σ) follows from Lemma 16.

Assume ℓ was honest in step 1. Lemma 13 implies that every honest party holds the same valid $(B_\ell, \Sigma_\ell) \neq \perp$ in step 2, and so sends a correctly formed Commit message on (B_ℓ, Σ_ℓ). Since there are at least $\lceil (n+1)/2 \rceil$ honest parties, the lemma follows. □

In Figure 7 we describe our block-agreement protocol Π_{BLA}.

Protocol Π_{BLA}

We describe the protocol from the point of view of a party P with input a valid pair (B, Σ).

Initialize $(k^*, B^*, \Sigma^*, C^*) := (0, B, \Sigma, \emptyset)$ and $k := 1$. While $k \leq \kappa$ do:

1. At time $(5k-5) \cdot \Delta$, run Π^k_{GC} using input $(k^*, B^*, \Sigma^*, C^*)$ to obtain output $((B, \Sigma, C), g)$.
2. At time $5k \cdot \Delta$ do: If $g > 0$, set $(k^*, B^*, \Sigma^*, C^*) := (k, B, \Sigma, C)$. If $g = 2$, output (B, Σ). Increment k.

Fig. 7. A block-agreement protocol Π_{BLA}.

Lemma 18. *If $t < n/2$, then Π_{BLA} is t-secure.*

Proof. Assume fewer than $n/2$ parties are corrupted. Let k be the first iteration in which some honest party outputs (B, Σ). We first show that

in every subsequent iteration: (1) every honest party P_i uses as its input in step 1 a k_i-vote on (B, Σ); and (2) corrupted parties cannot construct a k'-vote on $(B', \Sigma') \neq (B, \Sigma)$ for any $k' \geq \min_i\{k_i\}$.

Say an honest party outputs (B, Σ) in iteration k. Then that party must have output a k-certificate for (B, Σ) in the execution of Π^k_{GC} in iteration k. By Lemma 16, this means every honest party output a k-certificate on (B, Σ) in the same execution of Π^k_{GC}, and so (1) holds in iteration $k+1$. Moreover, Lemma 15 implies that no honest party sent a Commit message on $(B', \Sigma') \neq (B, \Sigma)$ in the execution of Π^k_{GC}, and so (2) also holds in iteration $k+1$. Lemma 14 implies, inductively, that the stated properties continue to hold in every subsequent iteration.

It follows from Lemma 14 that any other honest party P who generates output in Π_{BLA} also outputs (B, Σ), regardless of whether they generate output in iteration k or a subsequent iteration.

Lemma 17 shows that in each iteration of Π_{BLA}, with probability at least $1/2$ all honest parties output some (the same) valid (B, Σ) in that iteration. Thus, after κ iterations all honest parties have generated valid output with probability at least $1 - 2^{-\kappa}$ (note that all parties terminate after κ iterations). $\square$