

State Machine Replication under Changing Network Conditions

Andreea B. Alexandru^{1*}, Erica Blum^{1*}, Jonathan Katz^{1*}, and Julian Loss^{2**}

¹ University of Maryland, College Park

aandreea,erblum@umd.edu, jkatz2@gmail.com

² CISA Helmholz Center for Information Security
lossjulian@gmail.com

Abstract. Protocols for state machine replication (SMR) are typically designed for synchronous or asynchronous networks, with a lower corruption threshold in the latter case. Recent *network-agnostic* protocols are secure when run in either a synchronous or an asynchronous network. We propose two new constructions of network-agnostic SMR protocols that improve on existing protocols in terms of either the adversarial model or communication complexity:

1. an *adaptively secure* protocol with optimal corruption thresholds and quadratic amortized communication complexity per transaction;
2. a statically secure protocol with near-optimal corruption thresholds and *linear* amortized communication complexity per transaction.

We further explore SMR protocols run in a network that may change between synchronous and asynchronous arbitrarily often; parties can be uncorrupted (as in the proactive model), and the protocol should remain secure as long as the appropriate corruption thresholds are maintained. We show that purely asynchronous proactive secret sharing is impossible without some form of synchronization between the parties, ruling out a natural approach to proactively secure network-agnostic SMR protocols. Motivated by this negative result, we consider a model where the adversary is limited in the total number of parties it can corrupt over the duration of the protocol and show, in this setting, that our SMR protocols remain secure even under arbitrarily changing network conditions.

Keywords: State Machine Replication, Consensus, Proactive Security

1 Introduction

Protocols for *state machine replication* (SMR) allow a set of parties P_1, \dots, P_n to agree on a continuously growing, ordered log of transactions. SMR protocols enable the evolving state of a distributed system to be replicated across multiple

* Work supported by NSF award #1837517.

** Part of this work was done while the author was a postdoctoral researcher at the University of Maryland and at the Carnegie Mellon University.

parties, even when some of them are malicious. SMR lies at the core of many distributed applications and has recently received considerable attention in the context of blockchain protocols. Most of the literature focuses on protocols that are secure in either the *synchronous* or the *asynchronous* model. SMR protocols in the synchronous model can tolerate $t < n/2$ corrupted parties ($t < n$ corrupted parties if external validity is not required [33]), but may fail if the synchrony assumption is violated. On the other hand, asynchronous protocols are secure under arbitrary network conditions, but do not exist when $t \geq n/3$.

Recent work of Blum, Katz, and Loss [6] introduced the *network-agnostic* model in which a single protocol is required to be secure regardless of whether it is run in a synchronous or an asynchronous network, for different corruption thresholds. In subsequent work [7], they show that for any thresholds $t_a \leq t_s$ with $2t_s + t_a < n$, there is an SMR protocol that tolerates t_a corrupted parties if the network is asynchronous and simultaneously tolerates t_s corrupted parties if the network is synchronous. A major benefit of network-agnostic protocols over classical ones is that t_a, t_s can be chosen arbitrarily subject to the above constraints. This allows a protocol designer to flexibly choose t_a, t_s so as to minimize the probability of failure based on assumed properties of the environment.

Although network-agnostic protocols have recently received significant attention [6,8,7,28,4,16], several open questions regarding network-agnostic SMR remain. For one, existing results are primarily concerned with feasibility rather than efficiency; this is especially true when considering protocols secure against an adaptive adversary who can choose which parties to corrupt during the execution of the protocol. Perhaps the most significant limitation of prior work is that it either requires the network to be synchronous for the lifetime of the protocol, or else guarantees security only if the attacker never exceeds the corruption threshold of t_a . Providing a more elegant treatment of networks that can change arbitrarily often between synchronous and asynchronous was left as an explicit open question in prior work.

1.1 Challenges and State-of-the-Art

We begin with a brief overview of network-agnostic SMR, and then explain how existing solutions (do not) deal with the issues raised above.

Network-agnostic SMR. The goal of an SMR protocol is to impose order on transactions that arrive in parties' buffers in an arbitrary fashion. An SMR protocol must ensure *consistency*, which means that all parties agree on the order in which transactions are committed to some log, and *liveness*, which means that any transactions in the buffers of honest parties are eventually appended to the log. SMR is significantly more challenging than the related problem of Byzantine agreement, where parties agree on only a single value.

A network-agnostic SMR protocol must remain secure if the network is synchronous and there are at most t_s corruptions, or if the network is asynchronous and there are at most t_a corruptions. As a key building block for SMR in this setting, Blum et al. [7] introduced a novel protocol for asynchronous common

subset (ACS) that allows parties to agree on a subset of $n - t_a$ inputs in the presence of t_a corrupted parties in an asynchronous network. Their protocol has the property that if all honest parties supply *the same input* B to the protocol, then honest parties include B in their output even when t_s parties are corrupted. This facilitates the following strategy: parties first attempt to agree on an input B using a synchronous protocol. If the network is synchronous, this step will succeed even in the presence of t_s corrupted parties; thus, parties all use the same input B to ACS which outputs this block even if there are t_s corrupted parties. On the other hand, if the network is asynchronous, t_a -security of ACS ensures that all parties can agree on B without relying on the synchronous protocol.

Problems with existing solutions. Blum et al. [7] present two SMR protocols, **Tardigrade** and **Upgrade**. **Tardigrade** is secure against an adaptive adversary and requires $O(n^4)$ bits of communication for n transactions. **Upgrade** gives a more efficient alternative against a static adversary that requires only $O(n^3)$ bits of communication for n^2 transactions. However, **Upgrade** relies on random subcommittees to execute the most expensive steps of the protocol. Such protocols are not adaptively secure and require very large committees in order to provide meaningful corruption bounds. This arguably offsets the communication improvements made by **Upgrade**, as it only offers an asymptotic improvement if the total number of parties in the system is in the order of hundreds of thousands.

Moreover, their work only considers non-switching networks, i.e., the network is either synchronous or asynchronous for the entire duration of the protocol. Thus, if at any point in the lifetime of the protocol the adversary surpasses t_a corrupted parties, their protocols might be insecure if the network is ever asynchronous. We are interested in a more flexible model that tolerates repeated transitions of the network between synchronous and asynchronous behavior, and even in the presence of an adaptive, mobile adversary.

1.2 Our Contributions

We study protocols in a more realistic model where network conditions can arbitrarily change over time, and parties can also recover from corruptions. Such recovery is necessary if we want to allow more than t_a corruptions when the network is synchronous, but then restrict the adversary to fewer than t_a corruptions when the network becomes asynchronous.

Modeling recovery from key exposure. Modeling parties that are temporarily corrupted (sometimes referred to in the literature as *transient faults*) is non-trivial when parties have long-term keys. To model the process of uncorruption, we endow parties with a mechanism to forcibly “flush out” the adversary. (This could be achieved, for example, by having parties restart their computer in safe mode at the onset of a new protocol epoch.) The adaptive adversary can then choose to re-corrupt those parties or new ones. However, without additional measures in place, the internal state of the previously corrupted parties (including their long-term secret keys) remains known to the adversary. Proactive secret sharing is the main technique to refresh parties’ keys for threshold signatures

and related primitives commonly used in communication-efficient randomized SMR protocols. We prove that without further restrictions, secure proactive secret sharing protocols in the pure asynchronous and network-agnostic setting are impossible. While this may seem to be a folklore result, modeling and proving such a result is non-trivial. One of our contributions is to formalize this result and provide a rigorous proof.

To address the above impossibility in the context of SMR protocols, we consider a model in which the attacker is limited to corrupting a set S of at most t_s parties *for the lifetime of the protocol*. (It may corrupt this entire set of parties when the network is synchronous, and must uncorrupt at least $t_s - t_a$ of them when the network becomes asynchronous.) Since transient corruptions are rarely considered in the context of SMR, limiting the total number of faults to t_s seems like a reasonable assumption which is in line with most of the existing literature.

Practical network-agnostic SMR. We propose two new efficient protocols for SMR, **Update** and **Upstate**. **Update** is adaptively secure for optimal corruption thresholds and has $O(n^3)$ communication complexity for committing a block of $O(n)$. This is an $O(n)$ improvement over **Tardigrade** [7], which requires $O(n^4)$ communication to commit blocks of $O(n)$ transactions. We obtain the improvement by carefully applying error-correcting codes in a new ACS protocol.

Upstate is statically secure for near-optimal corruption thresholds and has $O(n^2)$ communication complexity to commit blocks of $O(n)$ transactions. **Upstate** achieves its improved communication complexity by using committees. **Upstate** compares favorably to **Upgrade** [7]: while **Upgrade** requires $O(n^3)$ communication to commit blocks of $O(n^2)$ transactions, **Upstate** commits blocks of $O(n)$ transactions and requires $O(n^2)$ communication.

SMR tolerating key exposure. We show that our protocols are also secure when the network can transition between synchronous and asynchronous behavior and the adversary can be mobile across epochs, but is limited to corrupting at most t_s unique parties. Adding reboots at the beginning of each protocol epoch to flush the adversary out helps **Update** and **Upstate** to withstand the key exposures caused by the adversary’s mobility. Security in this case follows naturally from the structure of network-agnostic protocols. In order to be secure under a higher number of corruptions during the synchronous phase, some parts of the protocol have to use high thresholds for message collection. Although the adversary can know up to t_s keys/key shares during an asynchronous phase following a transition from a synchronous phase, it can only actively corrupt t_a parties and is not able to break security even if it forges or erases keys.

Open questions. We leave open the question of designing an adaptively secure SMR protocol in our setting with quadratic communication complexity per committed block. We also leave open to explore communication-efficient proactive network-agnostic SMR protocols that bypass the impossibility result of network-agnostic proactive secret sharing. We remark that although our protocols use threshold cryptosystems to boost efficiency and censorship resilience, these may not be necessary. Thus, it is plausible that a solution for key refresh could be

achieved without limiting the adversary to corrupting a set of t_s parties. One could then hope to use a network-agnostic ACS protocol to agree on a new list of valid public keys obtained from distributed key generation.

1.3 Related work

Network-agnostic protocols were introduced by Blum et al. in the context of Byzantine agreement [6], and was later extended to multi-party computation [8] and SMR [7]. The latter presents two network-agnostic SMR protocols. *Tardigrade* achieves total communication $O(n^4 + n^3\ell)$ against adaptive adversaries, for n the number of parties and ℓ the block size. *Upgrade* uses committees to achieve total communication $O(n^3 + n\ell)$ against static adversaries (but tolerates fewer corruptions). Appan et al. [4] proposed a protocol for network-agnostic perfectly secure multi-party computation; their protocol uses a novel network-agnostic perfectly secure verifiable secret sharing protocol (but not proactive).

Since our protocols need to support both synchronous and asynchronous networks, and asynchronous SMR protocols are less communication efficient compared to their synchronous counterparts [2,3], we focus here on asynchronous SMR protocols tolerating $t < n/3$ corruptions. Canonical constructions for SMR and atomic broadcast are based on multi-value validated asynchronous Byzantine agreement or asynchronous common subset [23,10,27,14,20] with cubic communication complexity for input sizes linear in n . Only a few existing protocols in the asynchronous setting tolerate *adaptive* corruptions. EPIC [24] and DAG-Rider [22] achieve adaptive security with cubic total communication complexity; Dumbo2 [20] can be modified to achieve adaptive security by using the MVBA from [25]. Neither can be easily adapted to the network-agnostic setting.

A final group of related works concerns secret sharing and distributed key generation (DKG) where parties may crash and then recover or where the set of participants may change. In the *proactive model* [30], the adversary can be mobile across the corrupted parties over time. *Proactive secret sharing* (PSS) was introduced by Herzberg et al. [21]. Canetti et al. [11] and Frankel et al. [15] gave solutions for synchronous DKG against adaptive proactive adversaries using verifiable secret sharing schemes. Benhamouda et al. [5] introduced a secret-sharing protocol for passing secrets from one anonymous committee to another, while Groth [19] proposed a DKG scheme based on publicly verifiable secret sharing that allows refreshing key shares to a new committee. In the asynchronous case, Cachin et al. [9] presented a proactive refresh protocol assuming clock ticks that define epochs, based on [12] which recovers state in an SMR protocol. Schulze et al. [32] proposed a mobile PSS protocol in a partially synchronous network. Recently, several works [26,34,31] have proposed more efficient dynamic/mobile PSS protocols assuming eventual synchrony, short periods of synchrony at the end of an epoch, or synchronized epochs. Subsequent to our work, Yurek et al. [35] construct an asynchronous dynamic PSS protocol (circumventing our impossibility result) but with respect to different definitions than ours.

A related notion of security in the presence of exposed parties was considered in [18], which studied synchronous authenticated broadcast with both corrupted parties and parties who are honest but whose keys have been exposed.

Paper organization. We describe our model in Section 2, and provide definitions in Section 3. In Section 4, we present an ACS protocol that uses error-correcting codes in order to achieve $O(n^3)$ communication against an adaptive adversary, and prove its special properties. This ACS protocol is used as a building block in the Update SMR protocol presented in Section 5, which achieves optimal corruption thresholds in a network-agnostic setting. In Section 6, we describe an asymptotically more efficient SMR protocol, Upstate, that is secure under near optimal thresholds against a static adversary. In Section 7, we prove that under a restricted adversarial model, the SMR protocols discussed so far remain secure under arbitrary network transitions. In Section 8, we model and provide an impossibility proof for proactive asynchronous verifiable secret sharing. This result motivates our restricted mobile adversarial model.

2 Model

Network. We consider n parties P_1, \dots, P_n that are connected via pairwise authenticated channels and have access to a public key infrastructure. During the protocol’s execution, transactions are delivered to parties’ local buffers. We are not concerned with how these transactions originate; in practice, there is an external mechanism where clients gossip these transactions in the network.

When the network is *synchronous*, messages between parties are delivered with a finite, known delay Δ , and the local clocks of the parties are synchronized. When the network is *asynchronous*, messages between parties are eventually delivered to their intended recipient, but may be adversarially delayed or reordered. The local clocks of parties are only assumed to be monotonically increasing and are not necessarily synchronized anymore. If an asynchronous phase is followed by a synchronous phase, all messages sent during the asynchronous phase of the network are delivered by the beginning of the synchronous phase. Transitions between synchronous and asynchronous behaviors can happen arbitrarily.

An SMR protocol operates in logical intervals called *epochs*, which are measured and incremented locally. Another concept is that of a *round of communication*. In the synchronous setting, a round r refers to the time between $(r - 1)\Delta$ and $r\Delta$. In the asynchronous case, the round number will describe some particular send actions that are performed by a party.

We assume that parties perform *atomic send operations*, i.e., parties can send a message to multiple parties simultaneously in such a way that the adversary cannot corrupt them in between individual sends. Moreover, we assume that the adversary cannot perform *after the fact removal*, i.e., the adversary cannot indefinitely prevent a message from being delivered once it is sent by an honest party, even if the adversary corrupts it at some point after the send action.

Threat model. We consider a *Byzantine fault* model, in which some fraction of the parties may be corrupted by an adversary. The adversary controls the

local computations, messages, and current state of any corrupted party, and can coordinate the actions of all corrupted parties. Uncorrupted parties are called *honest*. For any honestly-initiated communication, the adversary receives the epoch τ , the sender identity S , the receiver identity R and the message m (which can be encrypted, in which case the adversary does not see its contents). The adversary determines when to deliver each message.

We assume that the adversary is (t_a, t_s) -limited, i.e., for some fixed thresholds t_s, t_a ($t_a \leq t_s$), up to $t_s < n/2$ parties may be corrupted if the network is synchronous and up to $t_a < n/3$ parties may be corrupted if the network is asynchronous. (The optimal trade-off between t_s, t_a is known to be $2t_s + t_a < n$ [7]). In Sections 4–5 we consider an *adaptive and rushing* adversary that adaptively corrupts parties over the course of a protocol execution; in Section 6, we consider a *static* adversary who corrupts parties prior to the start of an epoch.

Further, we address a mobile adversary. In Section 8, we consider an *epoch-wise mobile adaptive adversary* that can move freely between parties from epoch to epoch as long as it does not exceed more than t_s adaptive corruptions in the synchronous case and t_a adaptive corruptions in the asynchronous case at a given moment in time or in a given epoch. In Section 7, we consider a slightly different adversary who adaptively corrupts at most t_s parties *over the lifetime of the protocol*, and is only permitted to move between those t_s parties between epochs. We will explicitly mention the adversary’s capabilities in each section.

Reboot. To enable protocols to withstand network changes, we assume a reboot mechanism that causes a party to restart its device, thereby flushing out the adversary. Reboots occur at specified times during the protocols, not necessarily simultaneously. The adversary can immediately corrupt a party after rebooting, as long as it does not exceed the allowed threshold at that time. The restart is performed via code written in untamperable memory. Importantly, rebooting does not remove the previous state of a corrupted party from the adversary’s view; in particular, the adversary still knows the secret state of a party, including any secret keys that were held by that party during corruption. Furthermore, the internal state of a corrupted party that has restarted may have been arbitrarily modified by the adversary. For clarity, we call a party *actively corrupted* when the adversary actively controls that party’s behavior and *passively corrupted* or *exposed* if the party was uncorrupted either by the adversary or by reboot.

Keys. Every party P_i holds a private key sk_i of a threshold signature scheme with individual public signature key pk_i and public key pk . Further, every party P_i holds a private key dk_i of a threshold encryption scheme with individual public verification key vk_i and public key ek . The threshold for both schemes is t_s+1 . We assume a trusted dealer that generates $PK = (pk_1, \dots, pk_n, pk, vk_1, \dots, vk_n, ek)$ and $sk_1, \dots, sk_n, dk_1, \dots, dk_n$ and outputs a signature and encryption private keys sk_i, dk_i and the public key PK to each party P_i .

A party P_i can use its signature key sk_i to generate a signature share σ_i on a message m . The signature share σ_i can be verified using the message m and the public verification key pk_i , and is called *valid* if the verification is successful. As a shorthand notation for legibility, we use $\langle m \rangle_i$ for a threshold signature σ_i

of message m under secret key sk_i . A set of $t_s + 1$ valid signature shares on the same message m can be used to compute a signature σ for that message, which can be verified using the public key pk and m .

A party P_i can encrypt a message m using the public encryption key ek to generate a ciphertext c , and can use its decryption key dk_i to obtain a decryption share c_i of c . A decryption share c_i can be verified with respect to c , ek and vk_i and is called *correct* if the verification is successful. A set of $t_s + 1$ correct decryption shares can be used to obtain the decryption m of the ciphertext c .

We assume adaptively secure idealized threshold signature scheme and threshold encryption scheme. For a parameter κ , a signature share and a full signature have length $O(\kappa)$. We implicitly assume that parties use domain separation when constructing signatures to ensure only local context validity. An encryption of a message m of length $|m|$ has length $|m| + O(\kappa)$, and a decryption share has length $O(\kappa)$; these criteria can be met using standard KEM/DEM mechanisms.

3 Preliminaries

State machine replication protocols enable a set of parties to emulate a single server by agreeing on an ever-growing, ordered log of transactions.³ Given that SMR protocols usually continue indefinitely, we opt for a definition that clearly states how the logs are constructed and committed, and their relation order depending on epochs. A party maintains an ever-growing append-only log consisting of *blocks* of transactions: $\text{blocks}_i = (\text{block}_i[1], \text{block}_i[2], \dots)$, where the notation $\text{block}_i[e]$ refers to the block output by party P_i in epoch e . Each $\text{block}_i[e]$ is initialized with a special character \perp and populated by a set of transactions by P_i in epoch e . A party's epoch number is incremented after it outputs a block.

Definition 1 (State Machine Replication (SMR)). *Let Π be a protocol executed by n parties P_1, \dots, P_n . Let pp be some public parameters (e.g., PKI). Parties receive transactions as input, locally maintain arrays blocks , and output blocks and a publicly verifiable proof $\pi_i[e]$ for each $\text{block}_i[e]$ in blocks . Π is a secure SMR protocol tolerating t corruptions if the following properties hold:*

- (*t-Consistency*) *If an honest party outputs a block B in epoch e then all honest parties output B in epoch e .*
- (*t-Completeness*) *Every honest party outputs a block in all epochs.*
- (*t-Liveness*) *If a transaction tx is input to at least $n - t$ honest parties, then all honest parties eventually output a block containing tx .*
- (*t-External validity*) *If an honest party outputs (B, π) , then for a fixed public Boolean function Verify it holds that $\text{Verify}(\text{pp}, B, \pi) = 1$.*

Definition 2 (Binary Byzantine Agreement (BA)). *Let Π be a protocol executed by n parties P_1, \dots, P_n , where each party P_i begins holding input $x_i \in \{0, 1\}$ and parties terminate upon generating output. Π is a secure BA protocol tolerating t corruptions if the following properties hold:*

³ Following [28], we distinguish between SMR and atomic broadcast in that the former explicitly requires an externally verifiable proof of output validity.

- (*t*-Validity) If every honest party's input is equal to the same value x , then every honest party outputs x .
- (*t*-Consistency) All honest parties output the same message x .
- (*t*-Termination) Every honest party eventually terminates with output x .

Definition 3 (Asynchronous Common Subset (ACS)). Let Π be a protocol executed by n parties P_1, \dots, P_n , where each party P_i begins holding input $x_i \in \{0, 1\}^*$ and parties output sets of cardinality at most n . Π is a secure ACS protocol tolerating t corruptions if the following properties hold:

- (*t*-Validity) If every honest party's input is equal to the same value x , then every honest party outputs the value $\{x\}$.
- (*t*-Validity with termination) If every honest party's input is equal to the same value x , then every honest party outputs the value $\{x\}$ and terminates.
- (*t*-Consistency) If an honest party outputs S , all honest parties output S .
- (*t*-Set quality) If an honest party outputs a set S , then S contains the input of at least one honest party.
- (*t*-Termination) Every honest party generates output and terminates.

Block agreement (introduced in [7]) is a validated agreement on objects called *pre-blocks*. A pre-block is a vector of length n where the i th entry is either \perp or a message with a valid signature attached. The *quality* of a pre-block is defined as the number of entries that are not \perp ; a *k-quality pre-block* has quality at least k .

Definition 4 (Block Agreement (BLA)). Let Π be a protocol executed by n parties P_1, \dots, P_n , where each party P_i begins holding input $x_i \in \{0, 1\}^*$ and terminates upon generating output. Π is a secure BLA protocol tolerating t corruptions if the following properties hold:

- (*t*-Validity) If every honest party has input an $(n - t_s)$ -quality pre-block, then every honest party outputs an $(n - t_s)$ -quality pre-block.
- (*t*-Consistency) Every honest party outputs the same pre-block B .

Next, we briefly introduce some standard cryptographic primitives we use throughout. Further details can be found in Appendix A.

Threshold signature schemes. A (t, n) -threshold signature scheme is a signature scheme allowing $t + 1$ parties out of n to compute a signature on a message, with up to $t < n$ corruptions. It is *non-interactive* if parties can non-interactively compute signature shares that can be combined in the signature on a message, using protocols TS.Setup , TS.KeyGen , TS.Sign , TS.ShVer , TS.Verify for setup, key generation, partial signing, share verification and signature verification. The desired properties are correctness, security (unforgeability under chosen-message attack) and robustness (any number $\geq t + 1$ of signature shares can be combined to yield a signature) against a probabilistic polynomial-time adversary.

Linear error correcting codes. We adopt from [29] the description of error correcting codes, in particular, the Reed-Solomon (RS) code. An (n, b) -RS

code encodes b data symbols into codewords of n symbols, and can decode the codewords to recover the original data.

Given inputs m_1, \dots, m_b , the encoding function ENC computes codewords s_1, \dots, s_n . Knowledge of any b elements of the codeword uniquely determines the input message and the remaining of the codeword.

The decoding function DEC computes (m_1, \dots, m_b) , and is capable of tolerating up to c errors and d erasures in codewords (s_1, \dots, s_n) , if and only if $n - b \geq 2c + d$.

Committee election. A first method to elect a committee uses threshold signatures to produce an unpredictable coin. The coin is used to determine an ordering of parties by computing the hash $H(\text{coin}, i)$ and to order the parties accordingly. To elect a size κ committee, one simply takes the first κ parties in the ordering. The second method, known as *cryptographic sortition*, uses verifiable random functions (VRF) to allow each party to individually determine whether they are part of a committee, and then prove their membership to others [17,1]. During the protocol, parties are elected to a committee if and only if the output of the VRF on a specific string is less than a parameter b .

Throughout the paper, we deal with several security parameters. The signature size and the hash output size depend on a parameter that ensures computational security. The committee sizes depend on a parameter that ensures a negligible failure probability. To streamline notation, we denote all these by κ .

4 Asynchronous Common Subset

The protocol proceeds as outlined in Figure 1. Each party P_1, \dots, P_n , starts with an input of size ℓ and splits it into b blocks. These b blocks are then encoded into n codewords of size ℓ/b using a linear error correcting code. Each party P_i forms a message containing the j -th codeword and a hash of the input, signs it and sends it to party P_j . Upon receiving a validly signed message, each party multicasts it, along with the associated signature which will serve as a proof of the codeword validity. We refer to this procedure of input distribution as INDI, and present it in Figure 2. INDI is performed before the agreement on whose messages to output, and ensures that all parties are eventually able to reconstruct the selected inputs despite an adaptive adversary.

Upon receiving $n - t_s$ messages containing codewords, parties attempt to reconstruct the input. Instructions related to reconstruction (referred to as RECON) are shown in Figure 2. Upon reconstructing a valid input from some party P_j , parties multicast a signed *vote* message. Upon receiving $t_s + 1$ votes for P_j , parties assemble a certificate of validity for the reconstructed value of P_j , which consists of $t_s + 1$ signatures on h_j , used to form a full signature. The parties multicast a *commit* message carrying this certificate and the combined signature. We note that recently, Das et al. [13] proposed an asynchronous reliable broadcast protocol using error correcting codes (but without digital signatures) that is related to this step. Finally, upon receiving a unique *commit* message

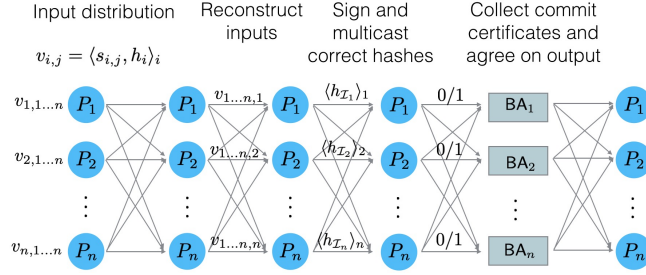


Fig. 1. Diagram of the steps in the Π_{ACS} protocol. BA stands for Byzantine Agreement. \mathcal{I}_i is the set of indices j for which party P_i reconstructed the initial message of party P_j .

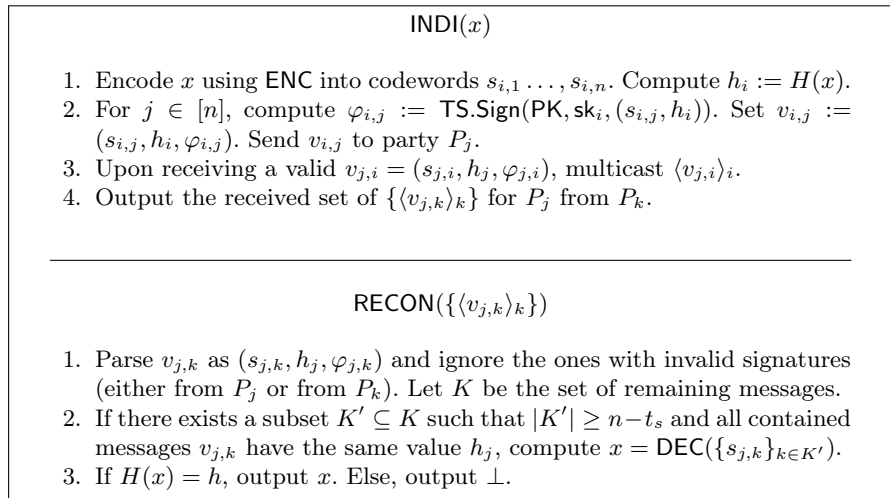


Fig. 2. Input distribution and reconstruction from the perspective of party $P_i \in \{1, \dots, n\}$.

for party P_j , parties input 1 to the corresponding BA_j instance. We implicitly assume that if honest parties receive conflicting commit messages, they do not input 1 to the respective BA.

Protocol Π_{Term} (Figure 4) assembles an output certificate that allows parties to output and terminate (OC 0), ensuring no honest parties are “left behind”.

Across the protocols, we use PK as the public keys output by TS.KeyGen and sk_i the secret key associated to P_i . For simplicity, in Π_{ACS} and the corresponding functionalities, we use $\varphi_{i,j}$ as both the signature of P_i over $s_{i,j}$, and over h_i , sent to party P_j . In this section (and all sections but Section 7), we use a binary BA protocol with t_a -validity, t_a -consistency, and t_a -termination in the presence of $t_a < n/3$ adaptive corruptions, and communication complexity of $O(n^2)$.

Encoding and reconstruction. ENC and DEC are associated to a (n, b) -RS code (Section 3). In the reconstruct procedure RECON, before feeding the code-

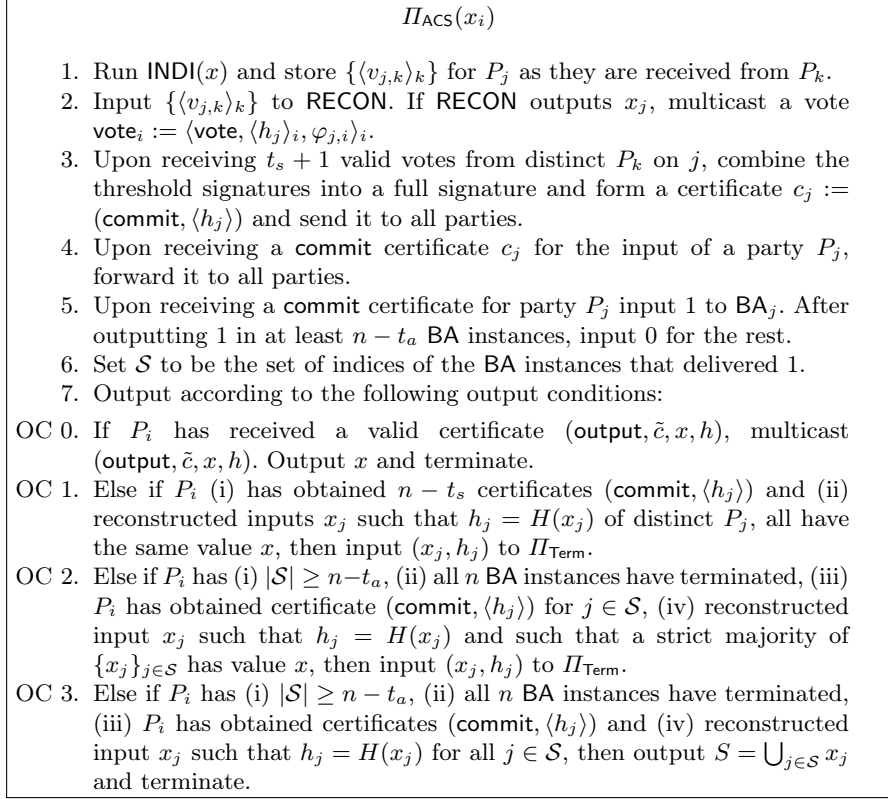


Fig. 3. ACS protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$.




Fig. 4. Termination helper protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$.

words into the DEC algorithm, parties first check that the corresponding signatures are correct. Then, parties check whether at least $n - t_s$ of the messages have the same associated hash value. If an honest party has not managed to reconstruct an input yet, it waits for more messages, then calls RECON again. Thus, each party feeds at least $n - t_s$ valid codewords in DEC. The (n, b) -RS code

allows a party to split an input in b blocks and encode them into n codewords. In order to tolerate d erasures, it must be possible to reconstruct the b blocks from $n - d$ correct codewords. Furthermore, to tolerate c errors among $n - d$ codewords, it must hold that $n - b \geq 2c + d$.

If we let b be equal to t_s , we can tolerate either $t_s + t_a$ erasures, or tolerate t_a errors along with $t_s - t_a$ erasures (since $n > 2t_s + t_a$). This means we need to wait for $n - t_s + t_a$ codewords in total in order to guarantee correct reconstruction in the asynchronous case when t_a parties are corrupted. Thus, a gain in communication efficiency, obtained from using codewords to achieve agreement on length κ hashes instead of length ℓ inputs and from not multicasting the reconstructed output, leads to potentially having to wait for $n - t_s + t_a$ messages in order to reconstruct the correct output if the adversary delivered t_a bad codewords.

If we let b be equal to t_a , we can tolerate either t_s errors and no erasures, or $2t_s$ erasures. This corresponds to the synchronous case when t_s parties are corrupted, and honest parties receive all messages that were sent after at most Δ time. Therefore, if an honest party only receives $n - t_s$ codewords, they are all correct. However, we will show below that there is no need to tolerate t_s errors in the synchronous case. Briefly, we can use extra information—the hash value—in order to detect an incorrect reconstruction, and there will be sufficiently many inputs of the honest parties correctly reconstructed in order to achieve termination. Therefore it suffices to let $b = t_s$ throughout.

Lemma 1. *Suppose there are at most t_a corruptions. Given a certificate $(\text{commit}, \langle h \rangle)$ for a party P , all honest parties can eventually reconstruct the same output in a run of Π_{ACS} .*

Proof. If P is honest, then all honest parties will eventually receive $n - t_s$ valid codewords of the true input (since we assume unforgeable signatures), allowing them to correctly reconstruct x .

Assume P is dishonest. To obtain a valid commit certificate on P 's hash $\langle h \rangle$, $t_s - t_a + 1$ honest parties need to have seen $n - t_s$ valid messages, all with the same $h = H(x)$. Of these $n - t_s$ messages, t_a could have been sent by corrupted parties in the multicast round. In the worst case, in the first round when P sent codewords, it could have sent only $n - t_s - t_a$ codewords (but all valid) to distinct honest parties. Eventually, all honest parties receive the $n - t_s - t_a$ codewords and can reconstruct the same input x if the code tolerates $t_s + t_a$ erasures.

On the other hand, the adversary might send t_a malicious codewords which will prevent correct reconstruction from $n - t_s$ codewords. However, assuming H is a collision-resistant hash function, except with negligible probability, there do not exist inputs $x \neq x'$ reconstructed by different sets of codewords such that $h = H(x) = H(x')$. Therefore, if after inputting $n - t_s$ codewords to RECON and not obtaining a valid output with respect to h , the honest parties wait until they receive sufficient codewords in order to be able to correctly reconstruct.

As stated above, each input of size ℓ is split into to $b = t_s$ blocks: $n - t_s > t_a + t_s = 2t_a + t_s - t_a$. This means that the code can tolerate either $t_a + t_s$ erasures, or $t_s - t_a$ erasures and t_a errors if parties wait for $n - t_s + t_a$ messages to honest parties. \square

Lemma 2. *If there are at most t_a -corruptions, there cannot be two valid certificates $(\text{commit}, \langle h \rangle)$, $(\text{commit}, \langle h' \rangle)$, associated with P , and $h \neq h'$.*

Proof. If P is honest, then all honest parties eventually receive $n - t_s$ valid messages containing codewords and the same hash h of the true input, so they can correctly reconstruct x . Therefore, assuming unforgeable signatures, no valid commit message $(\text{commit}, \langle h' \rangle)$ for $h' \neq h$ can exist.

Now suppose P is dishonest. Since there is a certificate $(\text{commit}, \langle h \rangle)$ constructed from at least $t_s + 1$ signatures, and $t_s + 1 > t_a$, at least one honest party P_j signed h . This implies P_j reconstructed an input x such that $h = H(x)$ and saw $n - t_s$ distinct valid messages $v_{*,l} = (s_{*,l}, h)$. At most t_a messages could have originated from malicious parties, so $n - t_s - t_a > t_s + 1$ were messages that honest parties relayed honestly. Assume there is a different honest party P_i that participated in a different commit certificate on h' for P . Then that party also saw $n - t_s$ distinct valid messages $v_{*,l'} = (s_{*,l'}, h')$, out of which $n - t_s - t_a > t_s + 1$ were messages that honest parties relayed honestly. These sets of honest parties should not intersect, so $2(n - t_s - t_a) < n - t_a$, but this contradicts our assumption that $n > 2t_s + t_a$. \square

Note that if the network is synchronous and $t_s = \lfloor n/2 \rfloor, t_a = 0$, different honest parties could receive commit certificates on different hashes of the same malicious party (honest parties always multicast the received certificates). In such a case, honest parties detect equivocation and do not input 1 in the associated BA. However, if the network is asynchronous equivocation is not necessarily detected. Nevertheless, as we see below, validity will still hold.

Lemma 3. Π_{ACS} *satisfies t_s -validity with termination.*

Proof. Suppose all honest parties have the same input x and up to t_s parties are corrupted. At most $t_s < \lfloor \frac{n-t_a}{2} \rfloor + 1 < n - t_s$ reconstructed values can be different than x , so there cannot exist an output certificate on a value $x' \neq x$ even if two honest parties accept different commit certificates for the same corrupted party.

Honest parties will eventually be able to obtain valid commit certificates for the inputs of at least $n - t_s$ honest parties, and therefore (by assumption) eventually obtain at least $n - t_s$ valid certificates for x . At this point, if an honest party has not yet output, it will input $\{x\}$ to Π_{Term} (in OC 1). If at least $t_s + 1$ parties call Π_{Term} via OC 1, then eventually, each party will receive an honest output certificate on $\{x\}$, output and terminate. Below we handle the case in which some honest parties output before the above conditions were satisfied.

Assume party P output before the above could occur. If P called Π_{Term} via OC 2, then despite t_s corruptions that could break security of the t_a -secure BA, it saw x' reconstructed in a strict majority of valid values associated with $n - t_a$ BA terminated instances. Any set of BA instances constituting a strict majority must contain at least one instance corresponding to honest party, since $\lfloor \frac{n-t_a}{2} \rfloor + 1 > t_s + 1$, and so $\{x'\} = \{x\}$ by assumption. Furthermore, in this case P would have input (x, h) to Π_{Term} , and so all parties eventually receive an output certificate on $\{x\}$. Since $n - t_s > \lfloor \frac{n-t_a}{2} \rfloor + 1$, and honest parties' inputs

can always eventually be reconstructed, each honest party will be eventually able to output due to OC 0, even if it was not able to finish the reconstruction of the corrupted parties' inputs.

Finally, if P output S as a result of OC 3, then P did not observe a strict majority of BA instances in \mathcal{S} corresponding to the same value. By assumption, the honest parties have the same input x , so this implies a strict majority of values \mathcal{S} correspond to corrupted parties. However, this contradicts the assumption that only t_s parties are corrupted, because $\lfloor \frac{|\mathcal{S}|}{2} \rfloor \geq t_s$. Therefore, no honest party outputs via OC 3 when all honest parties have the same input. \square

Lemma 4. Π_{ACS} satisfies t_a -set quality.

Proof. Suppose an honest party P_i output a set S .

If P_i output $S = \{x\}$ due to OC 0, then P_i must have obtained a valid output certificate of at least $t_s + 1$ signatures on x , which requires that at least one honest party (call it P_j) input (x, h) to $\Pi_{Term}(x, h)$ in OC 1 or OC 2. Consider each case. If P_j input (x, h) due to OC 1, then it gathered a valid certificate on at least $n - t_s$ values equal to x . At least $n - t_s - t_a \geq t_s + 1$ of the parties associated to these values are honest, so RECON returns their correct original input value. Otherwise, if P_j input (x, h) due to OC 2, then it output 1 in at least $n - t_a$ BA instances and it saw a strict majority of the reconstructed corresponding inputs reconstruct to the value x . Because $n \geq n - t_s + \lfloor \frac{n - t_a}{2} \rfloor + 1$, x was input by some honest party. Thus, in either case some honest party input x .

If P output S due to OC 3, then it output 1 in at least $n - t_a$ BA instances but without the majority condition satisfied. At least one of these instances corresponds to an honest party, so S contains some honest party's input. \square

Lemma 5. Π_{ACS} is t_a -terminating.

Proof. Assume no honest party has output yet. Eventually, all honest parties will obtain at least $n - t_a$ valid commit certificates, since there are at least $n - t_a$ honest parties. Moreover, by Lemma 2, even on malicious inputs, honest parties cannot obtain multiple valid certificates. By the t_a -terminating property of BA, all parties terminate all n BA instances eventually. By the t_a -consistency of BA, all honest parties will agree on the set \mathcal{S} of BA instances that output 1. Finally, by Lemma 1, all honest parties reconstruct the same inputs associated to \mathcal{S} . This allows some honest party to output and terminate.

It remains to show that once some honest party P_i has terminated, all honest parties eventually terminate. If P_i output due to OC 0 (implying it received a valid output certificate from OC 1 or OC 2), then eventually all honest parties receive the certificate multicast by P_i and terminate (if they have not already).

If P_i output due to condition OC 3, then it must have terminated all BA instances, obtained commit certificates and reconstructed all inputs corresponding to $\mathcal{S} = \{i | \text{BA}_i \text{ output } 1\}$ for some $|\mathcal{S}| \geq n - t_a$. Then, t_a -termination and consistency of BA ensure that each other honest party P_j eventually observes parts (i) and (ii) of OC 3 to be true. Furthermore, each honest party eventually reconstructs each $\{x_j\}_{j \in \mathcal{S}}$ and receives the certificates needed to terminate, since P_i must have sent these certificates to all other parties during ACS. \square

Lemma 6. Π_{ACS} satisfies t_a -consistency.

Proof. Assume an honest party P_i has output S . By Lemma 5, each other honest party eventually outputs some set S' . It remains to show that for each possible combination of output conditions, $S = S'$.

Suppose $S = \{x\}$ was output via OC 0, i.e., upon receiving a valid output certificate. There are two subcases.

First, suppose P_j output $S' = \{x'\}$ via OC 0. The existence of an output certificate for x implies that there exists an honest party P who contributed a share via either OC 1 or OC 2; likewise, some honest party P' contributed a share for x' . If both P and P' contributed shares via OC 1, then quorum intersection among the two sets of $n - t_s$ certificates implies $x = x'$. If (say) P and P' contributed shares by OC 1 and OC 2, respectively, then any set of $n - t_s$ BA instances and any set of $\lfloor \frac{n-t_a}{2} \rfloor + 1$ BA instances must intersect at an honest party, and so $x = x'$. Finally, if both P and P' contributed shares via OC 2, then they agree on S , and once again $x = x'$.

Second, suppose towards a contradiction that P_j output $S = \cup_{j \in \mathcal{S}} x_j$ for reconstructed values x_j via OC 3. Of those $n - t_a$ values, at most t_s can have a value $x' \neq x$. But this means that P_j saw at least $n - t_a - t_s \geq t_s + 1$ reconstructed values equal to x , in which case the order of else-if clauses would have caused P_j to output via OC 2, a contradiction.

Third, say P_i outputs S as a result of OC 3. The case in which P_j output $\{x'\}$ via OC 0 is equivalent to the second subcase above. Suppose P_j also output a set S' via OC 3. Both P_i and P_j must have seen all BA instances terminate and agree on the set of BA instances \mathcal{S} that output 1. By Lemma 1, we have $S' = S$. \square

Communication complexity. The Π_{ACS} protocol has a communication complexity of $O(n^2\ell + \kappa n^3)$ per input of size ℓ .

5 The Update SMR Protocol

In this section, we consider an adaptive adversary without mobility, which can actively corrupt at most t_s parties if the network is synchronous, and can corrupt at most t_a parties if the network is asynchronous, in any given epoch. Protocol 5 describes our construction for a network-agnostic SMR protocol.

Apart from the ACS protocol described in Section 4, we also use a block agreement protocol (BLA), whose role is to make parties agree on the input to ACS if the network is synchronous. Honest parties input $(n - t_s)$ -quality pre-blocks of length L to the BLA and ignore any pre-blocks with quality less than $n - t_s$.

We use the adaptively secure block agreement protocol from [7], which we call Π_{BLA} . The protocol has a total complexity of $O(\kappa n^3 + \kappa n^2 L)$ per pre-block of size L . Π_{BLA} has R inner rounds and guarantees t_s -validity, t_s -consistency and t_s -termination in a synchronous network when up to t_s parties are corrupted. We cannot guarantee these in an asynchronous network. However, even if the network is asynchronous, any honest party who terminates Π_{BLA} does so with output that is a valid $n - t_s$ -quality pre-block.

The logical flow of the network-agnostic SMR is the following. In every epoch, each honest party first selects a random sample of L/n transactions from its buffer of transactions. The selected transactions are then threshold encrypted. Next, the parties multicast their encrypted samples and start to assemble a $(n - t_s)$ -quality pre-block. If an honest party succeeds in assembling such a pre-block within the allotted time, it inputs it to Π_{BLA} , which is guaranteed to terminate with consistent output B^* if the network is synchronous. Regardless, honest parties will then input either B^* if obtained from Π_{BLA} or a $(n - t_s)$ -quality pre-block to Π_{ACS} . Recall that Π_{ACS} is guaranteed to terminate regardless of the network condition. Lastly, honest parties participate in constructing the final block: they jointly decrypt the output value of Π_{ACS} , populate the block with the unique transactions, assemble a validity certificate on the hash of the obtained block, and remove the posted transactions from their buffer.

We consider that epoch e starts for a party at time $T_e = \mu(e - 1)$ as measured by the local clock. The parameter μ is a spacing parameter that should be heuristically tuned by the network designers to improve throughput, i.e., not have too much overlap or separation between epochs. If the network is synchronous, then epochs start at the same time for all parties. If the network is asynchronous, parties might start the epochs at different times and might not output a block until they have to start the next epoch. We implicitly assume parties can distinguish between messages from different epochs, e.g. by tagging messages with e .

Below we give our main results on `Update`. The proofs use the results on Π_{ACS} and Π_{BLA} discussed so far, and are provided in Appendix D.

Condition (*). Assume $t_a \leq t_s$, $2t_s + t_a < n$, and $t_a \leq n/3$, $t_s \leq n/2$.

Theorem 1. *Under condition (*), Π_{SMR} is (1) t_s -consistent and t_s -complete if the network is synchronous and (2) t_a -consistent and t_a -complete if the network is asynchronous.*

Theorem 2. *Under condition (*), Π_{SMR} is (1) t_s -externally valid if the network is synchronous and (2) t_a -externally valid if the network is asynchronous.*

Theorem 3. *Under condition (*), Π_{SMR} is (1) t_s -live if the network is synchronous and (2) t_a -live if the network is asynchronous.*

Communication complexity. In Π_{SMR} , the parties select a batch of L/n transactions, construct a pre-block of size $O(L|\text{tx}|)$, and input the pre-block to Π_{BLA} . If Π_{BLA} outputs, it also outputs a pre-block of size $O(L|\text{tx}|)$. The input to Π_{ACS} is of size $O(L|\text{tx}|)$, and if the network is synchronous, the output is of size $O(L|\text{tx}|)$. Conversely, if the network is asynchronous, the output is of size $O(nL|\text{tx}|)$. Since the transactions were randomly selected from honest parties' buffers, with high probability there will be $O(nL)$ transactions in the output block after decryption, assuming that throughput is not limited by a lack of transactions.

Step 1 of Π_{SMR} incurs $O(nL|\text{tx}| + n^2\kappa)$ total communication. In step 2, Π_{BLA} incurs $O(\kappa n^3 + \kappa n^2 L|\text{tx}|)$ total communication and Π_{ACS} incurs $O(\kappa n^3 + n^2 L|\text{tx}|)$ total communication. Finally, in step 3, the parties assemble an output block

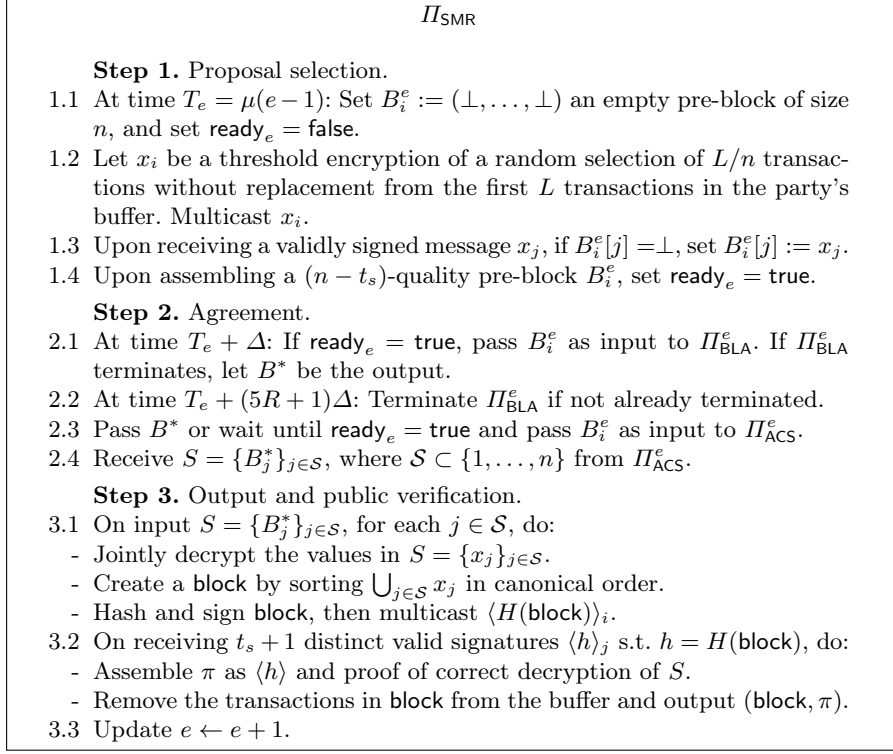


Fig. 5. Update SMR protocol with adaptive security for party $P_{i \in \{1, \dots, n\}}$.

and then multicast the signatures of the hash of the block to construct a proof, incurring $O(\kappa n^2)$ communication.

Summing over all steps, we see that **Update** incurs a total communication of $O(\kappa n^3 + \kappa n^2 L |\text{tx}|)$. Choosing a proposal sample size L that is $O(n)$ yields an asymptotic total communication of $O(\kappa n^3)$ per block of transactions and an amortized communication complexity of $O(\kappa n^2)$ per transaction.

6 The Upstate SMR Protocol

We consider a static adversary that is able to corrupt up to $\hat{t}_a = (1 - \epsilon)t_a$ parties in the asynchronous case and up to $\hat{t}_s = (1 - \epsilon)t_s$ parties in the synchronous case, for a small $\epsilon > 0$. Informally, the ϵ slack in the corruption thresholds ensures that with high probability the fraction of corruptions in a smaller committee chosen at random is close to the fraction of corruptions in the pool of n parties.

Figure 6 describes the input selection mechanism INSE^κ that handles input encoding and primary committee election. The input of size $\ell = L/\kappa$ is split as before into b blocks, which are then encoded into n codewords of size ℓ/b

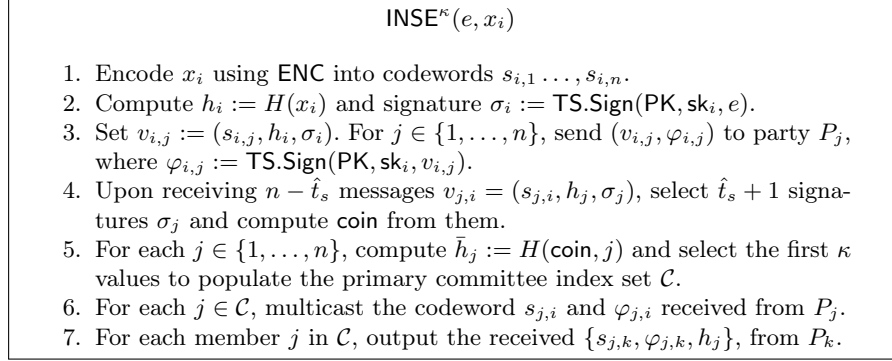


Fig. 6. Input selection—input encoding and primary committee election—from the perspective of party $P_{i \in \{1, \dots, n\}}$ in epoch e .

(Section 3). Each party sends the i -th codeword with a hash and a threshold signature over the epoch number to party P_i . Combining $\hat{t}_s + 1$ threshold signatures yields an unpredictable value that is used to select a committee of κ parties whose inputs will form the output.

Protocol 7 describes our construction for a network-agnostic committee-based SMR protocol. At the start of each epoch, parties choose a random sample of L/κ transactions from their buffers. The parties then run an input selection procedure, called INSE, to select κ committee members. Inputs from committee members are gathered into pre-blocks, which are passed to committee-based versions of BLA and ACS in the same way as in Update. Because the committee is of size κ , the pre-blocks are $(1 - t_s/n)\kappa$ -quality. The committee-based ACS and BLA protocols are described at the end of the section, with additional details in Appendix D. After running BLA and ACS, the parties construct the final block by jointly decrypting the output value of Π_{ACS}^κ .

Condition ().** Assume $t_a \leq t_s$, $2t_s + t_a < n$, $t_a \leq n/3$, $t_s \leq n/2$ and $\hat{t}_a := (1 - \epsilon)t_a$, $\hat{t}_s := (1 - \epsilon)t_s$ for $\epsilon > 0$.

Theorem 4. *Under condition (**) except with negligible probability, Π_{SMR}^κ is (1) \hat{t}_s -consistent, \hat{t}_s -complete, \hat{t}_s -externally valid and \hat{t}_s -live if the network is synchronous and (2) \hat{t}_a -consistent, \hat{t}_a -complete, \hat{t}_a -externally valid and \hat{t}_a -live if the network is asynchronous.*

The proof follows along the same lines as the proofs of Theorems 1–3, using the properties of the committee-based protocols Π_{ACS}^κ and Π_{BLA}^κ .

Committee-based asynchronous common subset. We now present an ACS protocol Π_{ACS}^κ in a network-agnostic setting with static corruptions.

An overview of the protocol appears in Figure 8. Inputs of size ℓ are passed to the input selection procedure INSE (Figure 6), which determines the *primary committee* \mathcal{C} . Next, each party multicasts the codewords they received from the

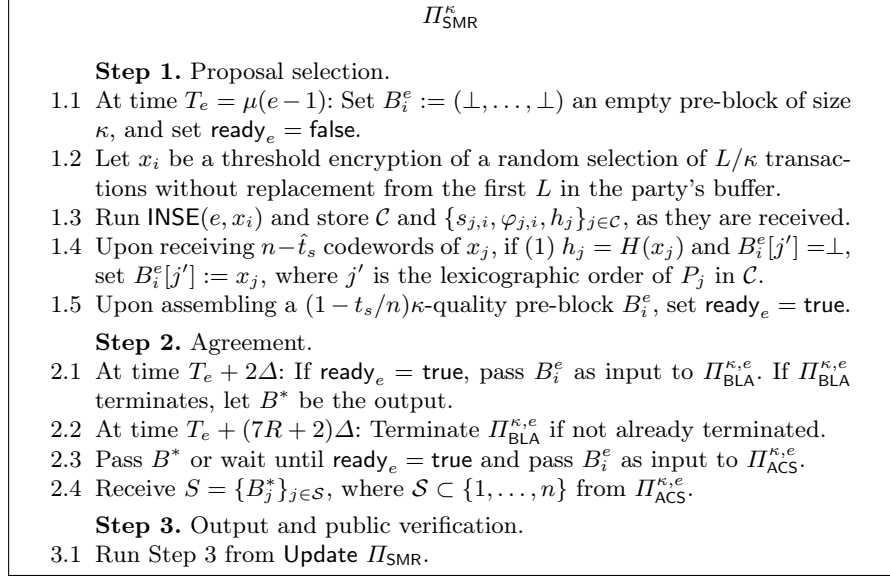


Fig. 7. SMR protocol with static security for party $P_{i \in \{1, \dots, n\}}$.

members of the primary committee. To reduce communication, one *secondary* committee is elected for each member of the primary committee. The secondary committee is responsible for constructing certificates of correctness for the reconstructed values of the primary committee. The secondary committees are self-elected as described in Section 3. Finally, parties agree on which primary committee members' values to output by running κ parallel BA instances.

Inputs are split into $b = \hat{t}_s$ blocks using an error correcting code that tolerates either \hat{t}_s erasures or \hat{t}_a errors and $\hat{t}_s - \hat{t}_a$ erasures. For simplicity, in $\Pi_{\text{ACS}}^{\kappa}$, we use $\varphi_{i,j}$ as both the signature of P_i over $s_{i,j}$ and over h_i , sent to P_j . Across the protocols, H denotes a collision-resistant hash function and b a bound ensuring committees of size κ in expectation.

Lemma 7. $\Pi_{\text{ACS}}^{\kappa}$ is \hat{t}_a -consistent, \hat{t}_a -terminating, has \hat{t}_s -validity with termination and \hat{t}_a -set quality except with negligible probability.

Committee-based block agreement protocol. Throughout the remainder of the section, we consider a network that is synchronous with up to $\hat{t}_s = (1 - \epsilon)t_s$ corruptions, such that with high probability a committee of size κ will have up to $t_s\kappa/n$ corrupted members. Honest parties are assumed to input $(1 - t_s/n)\kappa$ -quality pre-blocks of total length κ to the block agreement protocol.

We construct a protocol BLA^{κ} , based on the BA protocol from [2,1] and the block agreement protocol from [7], with several changes to achieve security against adaptive adversaries at a quadratic communication per pre-block. The high-level idea is to elect a leader who proposes an input among the ones sent




Fig. 8. Diagram of the steps in the Π_{ACS}^{κ} protocol. CE stands for committee election and BA for Byzantine Agreement.

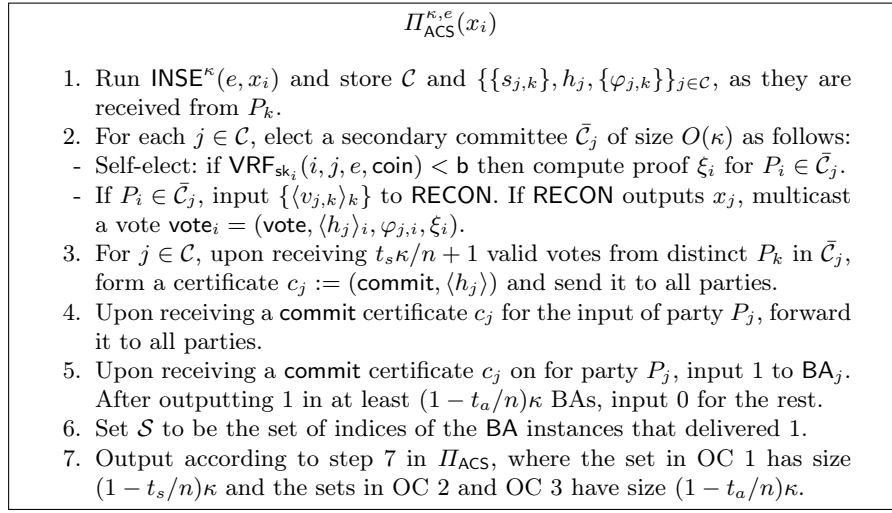


Fig. 9. ACS protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$ in epoch e .

by the parties, such that honest parties will commit on the same value. In our protocol, the proposal of inputs is performed before the leader election. Due to the forward secure signatures, the adversary cannot later corrupt the leader and cause them to equivocate. The construction is given in Appendix D.

Parties encode their pre-blocks into codewords and distribute them, along with the hash, for future reconstruction and verification. The protocol is run for multiple rounds, and a leader is elected at each round. The parties commit on a value when they receive sufficient votes on that value, prioritizing votes with higher round numbers. In each round, a different committee is tasked with assembling a certificate. In a given round, only votes from the current committee are considered valid. Π_{BLA}^{κ} makes calls to a graded consensus protocol Π_{GC}^{κ} , which makes a call to a Propose protocol $\Pi_{Propose}^{\kappa}$.

Communication complexity. Π_{ACS}^κ has communication complexity $O(\kappa n \ell + \kappa^2 n^2)$ communication and Π_{BLA}^κ has communication complexity $O(R\kappa^2 n^2 + \kappa n \ell)$, per input of size ℓ . In Π_{SMR}^κ , Π_{BLA}^κ and Π_{ACS}^κ are run on pre-blocks of size $O(L|\text{tx}|)$. If the network is synchronous, the output is of size $O(L|\text{tx}|)$, while if the network is asynchronous, the output is of size $O(\kappa L|\text{tx}|)$. After decryption, since the transactions were randomly selected from honest parties buffers, with high probability, there will be $O(\kappa L)$ transactions in the output block.

For simplicity, we omit the $|\text{tx}|$ factor in the following paragraph. Π_{SMR}^κ incurs $O(n^2 L/(\kappa b) + n^2 \kappa)$ total communication for step 1.3 and $O(n^2 \kappa L/b + n^2 \kappa^2)$ total communication in step 1.4. In step 2, Π_{BLA}^κ incurs $O(R\kappa^2 n^2 + \kappa n^2 L/b + \kappa n L)$ total communication and Π_{ACS}^κ incurs $O(\kappa n^2 L/b + \kappa n L + \kappa^2 n^2)$ total communication.

Since $b = \hat{t}_s = O(n)$, Upstate incurs a total communication of $O(R\kappa^2 n^2 + \kappa n L|\text{tx}|)$. This allows us to select a proposal sample size of $L = O(R\kappa n)$ and obtain a total communication of $O(R\kappa^2 n^2)$ per transaction and an amortized communication complexity of $O(\kappa n)$ per block L .

7 SMR under Arbitrary Network Changes

We now consider a network that can arbitrarily transition between synchronous and asynchronous behaviors and a *constrained epoch-mobile adaptive adversary*, who can corrupt at most t_s unique parties over the duration of the protocol, and can move between those t_s parties from epoch to epoch, as long as it does not exceed the t_a or t_s limit in any epoch or at any moment in time. In this model, parties' local machines may reboot to flush the adversary out. Importantly, the state of the parties is not removed from the adversary's view after uncorruptions.

Adding a reboot step at the beginning of each epoch to the network-agnostic protocols discussed so far, Update and Upstate, as well as Tardigrade, results in protocols that are secure under arbitrary network changes, as long as rebooting ensures that $n > 2t_s + t_a$, $t_a \leq t_s$, with at most $t_s - t_a$ exposed keys in the asynchronous case, in the restricted epoch-mobile model. For simplicity, we assume the reboot is instantaneous; otherwise we can adjust the timings of the steps.

Theorem 5. *Protocols Update, Upstate, and Tardigrade [7] with reboots are secure under arbitrary network changes against a constrained epoch-mobile adaptive adversary, where $n > 2t_s + t_a$, $t_a \leq t_s$.*

We prove the first part of Theorem 5 below, after some technical observations. Proofs of the rest of Theorem 5 and of the Lemmata are given in Appendix E.

Throughout, we use threshold cryptographic primitives with a threshold of $t_s + 1$. Although the adversary has access to up to t_s keys/key shares, it cannot create full signatures or certificates on its own because these require at least $t_s + 1$ valid contributions; likewise, it cannot decrypt independently of the honest parties. Moreover, while forming commit or output certificates, honest parties only sign messages that they locally verified, such as a hash value whose associate input was correctly reconstructed, or the output of the Π_{ACS} protocol.

In all protocols in this section, we use the binary BA protocol from [6], which is also designed for a network-agnostic setting with $n > 2t_s + t_a$. It is signature-free, apart from a threshold cryptosystem with high threshold of $t_s + 1$ to compute the common coin and ensure termination. This ensures that even with t_s key exposures (but only t_a active corruptions), the protocol remains t_a -valid, t_a -consistent and t_a -terminating against an adaptive adversary.

Lemma 8. *In a Π_{ACS} execution, if there are at most t_a corruptions and $t_s - t_a$ exposed parties, then at least $n - t_a$ BA instances will terminate with output 1.*

Lemma 9. *Suppose there are at most t_a corruptions and $t_s - t_a$ exposed parties during an execution of Π_{ACS} . Given a certificate for a party P , $(\text{commit}, \langle h \rangle)$, all honest parties eventually reconstruct the same output.*

Lemma 10. *If there are at most t_a corruptions, there cannot be two valid certificates $(\text{commit}, \langle h \rangle)$, $(\text{commit}, \langle h' \rangle)$ associated with P such that $h \neq h'$.*

Proof. (Theorem 5, Update) When the network is only synchronous or only asynchronous, or there is a single asynchronous to synchronous transition, the proof follows directly from the security proof of Update in Section 5.

Suppose the network has undergone a transition from synchronous to asynchronous. The adversary actively controls at most t_a parties, but may have exposed up to t_s parties. This means that each pre-block created by an actively corrupted party may contain up to t_s validly signed adversarial ciphertexts. However, exposed parties still act honestly, so each pre-block created by an honest party contains at most t_a malicious ciphertexts. Because pre-block entries are received directly from the corresponding party, an honest party's $(n - t_s)$ -quality pre-block will have at least $n - t_s - t_a$ honestly created and signed ciphertexts.

In the following, we first examine the security of the building blocks and then the security of the overall protocol.

ACS. In Π_{ACS} , parties need to be able to reconstruct all values corresponding to the at least $n - t_a$ BA instances that terminated with output 1. The use of codewords makes the analysis slightly subtler, since the adversary can forge valid but bad codewords and distribute them in the multicast round of INDI as if they originated from the exposed parties. By Lemma 8, at least $n - t_a$ BA instances will still terminate, despite exposures. Coupled with Lemmata 9 and 10, which show there cannot be conflicting certificates and all honest parties are able to eventually correctly reconstruct the same input, it follows that Π_{ACS} achieves t_a -termination, t_a -set quality and t_a -consistency. Finally, t_s -validity with termination has the same proof as in Lemma 3.

BLA. There is a Leader mechanism in Π_{BLA} [7], that is obtained using a strict majority of parties. Hence it is still unpredictable in the presence of t_s exposed parties. The property required of Π_{BLA} in the asynchronous case is the following: if an honest party does output in Π_{BLA} , its output is a $(n - t_s)$ -quality pre-block. Honest parties only validate and multicast $(n - t_s)$ -quality blocks, so this property still holds.

SMR. A corrupted party can forge the signature of an exposed party when assembling its own $(n - t_s)$ -quality pre-block. Therefore, up to t_a pre-blocks input to Π_{BLA} could have only $n - 2t_s$ entries originating from honest parties. If such a block is output by Π_{BLA} , then the same holds for the the output of Π_{ACS} .

By t_a -consistency and t_a -validity with termination of Π_{ACS} , all honest parties output the same set of pre-blocks. As a result, at least $n - t_a > t_s$ parties contribute valid decryption shares, and so every honest party is able to reconstruct the same block. Therefore, Update SMR is t_a -consistent and t_a -complete.

Next, we argue that t_a -liveness holds. If an adversarial pre-block is output by ACS, only $n - 2t_s$ honest parties are guaranteed to remove L/n transactions in a given epoch. Thus, the presence of key exposures increases the number of epochs needed for tx to move to the front of sufficiently many honest parties' buffers (see Appendix B). However, this still happens and ensures that tx is eventually output, but the probability increases with the number of epochs.

External validity follows from consistency of Π_{ACS} , since a threshold of $t_s + 1$ is used in the validity certificates over the block hashes.

Finally, the adversary cannot break the liveness of the protocol by erasing threshold key shares of the corrupted parties: any $t_s + 1$ shares can be used to reconstruct, so in order to prevent reconstruction, the adversary would need to erase at least $n - t_s - t_a$ shares. But this would require the adversary to corrupt more than t_s parties over the duration of the protocol, since $2t_s + t_a < n$. We conclude that security is preserved even across multiple network transitions. \square

8 Asynchronous Proactive Secret Sharing

We first consider an asynchronous network in the presence of a mobile adaptive adversary. At the end, we extend the analysis to changing network conditions.

In each epoch, the adversary is limited to t_a corruptions, but those t_a corruptions need not target the same parties in each epoch. Thus, over multiple epochs, the adversary could have controlled more than $t_a + 1$ different parties. While a party is corrupted, its current epoch is considered to be undefined, since it can behave arbitrarily. Upon becoming uncorrupted, a party's local epoch number is considered to be the epoch in which it was originally corrupted. We refer to the parties that are not corrupted as *honest* (in that epoch).

Here, we use an additional assumption of secure (authenticated private) channels, implemented using a pairwise shared key inaccessible to the adversary, e.g., stored in secure hardware. We show that even with secure channels, it is impossible to have a proactive asynchronous protocol without making any assumption on epoch length (as in [9] where epochs are defined to take place between clock ticks) but with epochs determined by a successful reshare of the secret (as in [32] but where the network is partially synchronous). While Cachin et al. [9] briefly remark upon this impossibility before making the assumption of clock ticks and "asynchronous proactive channels", we fully model and prove this result.

Definition 5. A (t_a+1) -out-of- n proactive verifiable secret sharing scheme with reshare is defined by an algorithm *Share* and protocols *Reshare*, *Reconstruct* that satisfy the following:

- *Share* takes as input a secret $s \in \mathbb{F}$ and outputs shares $(s_1^{(0)}, \dots, s_n^{(0)})$. Party P_i , $i = 1, \dots, n$ is given $s_i^{(0)}$ and sets its epoch number to 0.
- *Reshare* is an interactive protocol run by a subset of parties \mathcal{S} of size at least $n - t_a$ that takes as input an epoch number τ , a set of shares associated to that epoch number consisting of the share of each of the parties in \mathcal{S} : $(s_{i_1}^{(\tau)}, \dots, s_{i_{|\mathcal{S}|}}^{(\tau)})$ and outputs to every party P_i , $i \in [n]$ a new share $s_i^{(\tau+1)}$ or an error symbol \perp . A party P_i that receives output from *Reshare* with associated epoch τ sets its epoch number to $\tau + 1$.
- *Reconstruct* is an interactive protocol run by a subset of parties \mathcal{S} of size at least $n - t_a$, that takes as input a epoch number τ , a set of shares $(s_{i_1}^{(\tau)}, \dots, s_{i_{|\mathcal{S}|}}^{(\tau)})$ and outputs to all parties either a value $s' \in \mathbb{F}$ or an error symbol \perp .

An honest party is said to *complete* *Share*, *Reshare*, or *Reconstruct* in epoch τ when they generate the corresponding output from the algorithm in epoch τ .

In Appendix G, we give a standard privacy game between a challenger and an adversary \mathcal{A} where the goal of the adversary is to learn the secret. The advantage of the adversary is denoted by $\text{Adv}(\mathcal{A})$.

Definition 6. A proactive verifiable secret sharing scheme with reshare is secure against a t_a -limited adversary if it satisfies the following:

- (*Privacy*): $\text{Adv}(\mathcal{A})$ is negligible.
- (*Correctness*): For any $s \in \mathbb{F}$, conditioned on the adversary eventually delivering all messages between honest parties, it holds that: if during any epoch τ , a set \mathcal{S} of least $n - t_a$ honest parties locally call *Reconstruct* on epoch number τ and local shares associated with τ , they obtain the initially shared secret: $\text{Reconstruct}(\tau, \{s_i^{(\tau)}\}_{i \in \mathcal{S}}) = \text{Reconstruct}(\text{Share}(s))$. Furthermore, all parties in \mathcal{S} proceed to epoch $\tau + 1$.
- (*Liveness*): For any epoch number $\tau \geq 0$, if an honest party has reached epoch τ , i.e., has obtained output from the *Reshare* protocol associated to epoch $\tau - 1$, then all honest parties will eventually reach a epoch number $\tau' \geq \tau$, provided the adversary delivers all messages sent between honest parties so far and the responses triggered by these messages.

In verifiable secret sharing, in order to achieve correctness, *Share*, *Reshare* and *Reconstruct* need to implicitly have validation procedures of the inputs. We asked for at least $n - t_a$ instead of $t_a + 1$ parties to participate in *Reconstruct* to guarantee success against t_a malicious parties who could submit t_a invalid shares. Nevertheless, $t_a + 1$ valid shares are sufficient to reconstruct the secret.

Theorem 6. There does not exist a secure asynchronous (t_a+1) -out-of- n proactive verifiable secret sharing scheme with reshare.

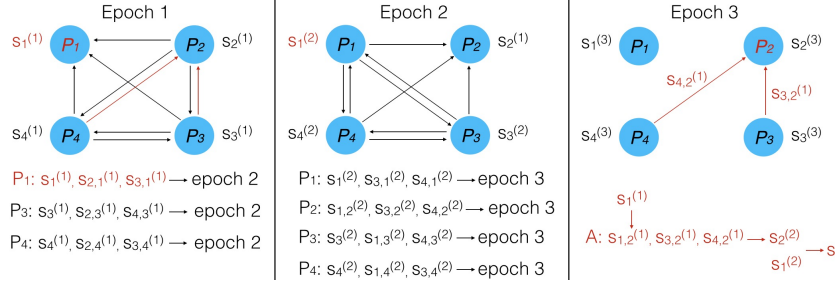


Fig. 10. We denote by $s_{j,i}^{(\tau)}$ the intermediate share obtained by party P_i from party P_j in epoch τ . P_j can construct its share for the next epoch $s_j^{(\tau+1)}$ from $n - t_a$ values $s_{j,i}^{(\tau)}$. The red quantities are in the view of the adversary. The red edges represent delayed messages from epoch 1 delivered in epoch 3.

Proof. We show that an adversary can break privacy by amassing shares corresponding to $t_a + 1$ parties in a single epoch. Then, we prove that protocols which avoid the prior attack do not satisfy liveness. For simplicity we first consider the case of non-interactive reshare protocols, and then handle the general case.

Non-interactive Reshare protocol. Consider $n = 4$ and $t_a = 1$. This counterexample is depicted in Figure 10 and can be extended to arbitrary n and corruption threshold $t_a < n/3$.

The adversary corrupts party P_1 in epoch 1. At this point in time, the adversary knows the state of P_1 , which includes the share $s_1^{(1)}$. Each honest party locally initiates the Reshare protocol at the onset of epoch 1. The adversary instructs P_1 not to deliver any message and delivers all the following messages: from P_2 to all other parties, from P_3 only to P_1 and P_4 , and from P_4 only to P_1 and P_3 . The parties P_3, P_4 thus obtain sufficient information to construct their shares $s_1^{(2)}, s_3^{(2)}, s_4^{(2)}$ and advance to epoch 2. However, P_2 remains in epoch 1. The adversary uncorrupts party P_1 after Reshare was completed. At this point in time, the view of the adversary includes $s_1^{(1)}$ and the intermediate shares for $s_1^{(2)}$. The adversary allows P_1 to also advance to epoch 2.

At the onset of epoch 2, each honest party locally initiates the Reshare protocol. The adversary delivers all messages between parties. This enables all parties to obtain their corresponding share $s_1^{(3)}, s_2^{(3)}, s_3^{(3)}, s_4^{(3)}$, and advance to epoch 3.

At the onset of epoch 3, the adversary corrupts party P_2 and delivers the messages originated in epoch 1 from P_3 and P_4 destined to P_2 . The adversary now has 3 messages, counting $s_1^{(1)}$, and is able to obtain $s_2^{(2)}$. Hence, it reconstructs s from two correct shares in epoch 2: $s_1^{(2)}, s_2^{(2)}$, without corrupting more than $t_a = 1$ party per epoch.

Restarting and flushing the adversary out does not prevent this attack, since there is no synchronizing signal instructing a corrupted party to restart before the first Reshare is completed. This could be addressed using erasures and/or interaction; however, we show that protocols that avoid this attack are not live.

Interactive Reshare protocols. Consider a generic interactive Reshare protocol where two parties, P_i and P_j , start an epoch with $s_i^{(\tau)}$ and $s_j^{(\tau)}$, respectively. After r rounds of communication, P_i obtains $s_{j,i}^{(\tau)}$ and P_j obtains $s_{i,j}^{(\tau)}$.

If only one of the r messages is useful for computing the new share, then the previous attack still applies. If more than one of the r messages are needed for computing the new share, and honest parties erase their previous state when transitioning to a new epoch (implying they do not respond to messages originated from previous epochs), then the above attack does not break privacy. But such an interactive asynchronous protocol where parties can only advance to the next epoch after repeated interactions does not achieve liveness, as shown next.

Consider now that the adversary delays all messages destined to P_1 , hence keeping it in epoch 1, while allowing the rest of the parties to progress an arbitrarily large number of epochs τ . At this point, the adversary delivers all messages that were sent so far, including the messages originated at P_1 as response to the received messages. However, since obtaining the output of any Reshare requires interaction and the other honest parties do not respond to messages originated in previous epochs in order to preserve privacy, a party P_1 cannot reach a subsequent epoch based only on the messages sent so far, breaking liveness. \square

The attack above hinges on the fact that a party can still retrieve in epoch $\tau' > \tau$ the contents of a message sent to it in epoch τ . Both privacy and liveness would be maintained if parties had access to “setup-free asynchronous forward-secure channels” with the following properties: (1) A message sent in epoch τ can only be read in epoch τ ; (2) At the onset of epoch $\tau + 1$, the sender and receiver on that channel have access to the new secret and public key, respectively, i.e., the adversary does not control the delivery of this information (it should not be interactive); (3) Messages in different epochs are encrypted with different keys.

Secure co-processors using forward secure encryption are not sufficient to implement this kind of channel. Say a party P_1 was delayed and is still in epoch τ , and all other parties advanced to epoch $\tau' > \tau$, updating their channel keys. But when honest parties start a new Reshare, they cannot use the key associated to P_1 's epoch τ , because an adversary corrupting P_1 in epoch τ would learn shares from epoch τ' and break privacy. These are points (1) and (3). So until the adversary delivers the messages from epoch τ , P_1 is stuck, but this does not break liveness if the protocol is non-interactive. If point (2) is satisfied, the other parties need to already have the public key in the channel for epoch $\tau + 1$, otherwise the impossibility proof for interactive protocols would apply. But a forward secure with unique public key allows a ciphertext encrypted at epoch $\tau + 1$ to be decrypted at epoch τ , so privacy is broken.

Note that in [9], the transition between epochs is *external*, triggered by a clock tick, and can happen even if a party did not complete the Reshare protocol in the current epoch. This allows parties to rely on the clock tick event to set new channel keys in a synchronized way.

To circumvent the result in Theorem 6, Yurek et al. [35] considered high reconstruction thresholds and defined local epochs such that a party can decide to not pass to a subsequent epoch even if it has all shares to do so, unlike our

definition based on completing a Reshare. Briefly, the impossibility does not hold because (i) a party decides to progress to the next epoch after receiving at least $n - t_a$ epoch τ messages (while in epoch τ), forcing the adversary to deliver at least these many messages to every party per epoch; (ii) combined with a high reconstruction threshold of $n - t_a$, the t_a shares held by the adversary in epoch τ and the at most t_a messages it could have delayed are not sufficient to reconstruct $s^{(\tau)}$, as $n - t_a > 2t_a$. We also mention that the constructions in [35] assume every party has Paillier key pairs that are not refreshed after corruptions.

Proactive secret sharing under network changes. We again consider a network that can arbitrarily switch between synchronous and asynchronous cases and $n > 2t_s + t_a$, $t_a \leq n/3$, $t_s \leq n/2$. Note that in this setting, the Reconstruct threshold is at least $t_s + 1$ and the Reshare threshold is $n - t_s$ in order to satisfy privacy in case the network is synchronous.

Corollary 1. *There does not exist a secure (t_s, t_a) -proactive verifiable secret sharing scheme with reshare under arbitrary network transitions.*

Proof. Assume the network is in an asynchronous state, so the adversary can corrupt up to t_a parties in the same local epoch. The arguments in the proof of Theorem 6 still hold. For the privacy attack, the adversary delays the messages in epoch τ towards $t_s - t_a + 1$ honest parties, until the epoch(s) it corrupts these parties (if $t_s \geq 2t_a$, it needs more epochs to corrupt all $t_s - t_a + 1$ parties), while allowing the rest of the parties to complete the refresh in all epochs, i.e., deliver and receive at least $n - t_s$ share messages. For the interactive liveness attack, the adversary can still cause the parties to be arbitrarily far apart. \square

We remark that the clock ticks used in II_{SMR} (Section 5) to start an epoch are not the same as the ones assumed in [9]. In our model, the epoch started at T_e does not necessarily finish by T_{e+1} , and can continue in the background, so liveness could be lost if all parties would erase their key shares at T_{e+1} .

References

1. I. Abraham, T.-H. H. Chan, D. Dolev, K. Nayak, R. Pass, L. Ren, and E. Shi. Communication complexity of byzantine agreement, revisited. In *Symposium on Principles of Distributed Computing (PODC)*, pages 317–326. ACM, 2019.
2. I. Abraham, S. Devadas, D. Dolev, K. Nayak, and L. Ren. Synchronous byzantine agreement with expected $O(1)$ rounds, expected $O(n^2)$ communication, and optimal resilience. In *Financial Cryptography (FC)*, volume 11598 of *LNCS*, pages 320–334. Springer, Heidelberg, 2019.
3. I. Abraham, D. Malkhi, K. Nayak, L. Ren, and M. Yin. Sync HotStuff: Simple and practical synchronous state machine replication. In *IEEE Symposium on Security and Privacy (SP)*, pages 106–118. IEEE Computer Society Press, 2020.
4. A. Appan, A. Chandramouli, and A. Choudhury. Perfectly-secure synchronous MPC with asynchronous fallback guarantees. Cryptology ePrint Archive, Report 2022/109, 2022. <https://eprint.iacr.org/2022/109>.

5. F. Benhamouda, C. Gentry, S. Gorbunov, S. Halevi, H. Krawczyk, C. Lin, T. Rabin, and L. Reyzin. Can a public blockchain keep a secret? In *Theory of Cryptography Conf. 2020, Part I*, volume 12550 of *LNCS*, pages 260–290. Springer, Heidelberg, 2020.
6. E. Blum, J. Katz, and J. Loss. Synchronous consensus with optimal asynchronous fallback guarantees. In *Theory of Cryptography Conf. 2019, Part I*, volume 11891 of *LNCS*, pages 131–150. Springer, Heidelberg, 2019.
7. E. Blum, J. Katz, and J. Loss. Tardigrade: An atomic broadcast protocol for arbitrary network conditions. In *Advances in Cryptology - Asiacrypt 2021, Part II*, volume 13091 of *LNCS*, pages 547–572. Springer, Heidelberg, 2021.
8. E. Blum, C.-D. L. Zhang, and J. Loss. Always have a backup plan: Fully secure synchronous MPC with asynchronous fallback. In *Advances in Cryptology - Crypto 2020, Part II*, volume 12171 of *LNCS*, pages 707–731. Springer, Heidelberg, 2020.
9. C. Cachin, K. Kursawe, A. Lysyanskaya, and R. Strobl. Asynchronous verifiable secret sharing and proactive cryptosystems. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 88–97. ACM Press, 2002.
10. C. Cachin and J. A. Poritz. Secure intrusion-tolerant replication on the internet. In *Intl. Conf. on Dependable Systems and Networks*, pages 167–176. IEEE, 2002.
11. R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin. Adaptive security for threshold cryptosystems. In *Advances in Cryptology - Crypto '99*, volume 1666 of *LNCS*, pages 98–115. Springer, Heidelberg, 1999.
12. M. Castro and B. Liskov. Proactive recovery in a Byzantine-Fault-Tolerant system. In *Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX Association, 2000.
13. S. Das, Z. Xiang, and L. Ren. Balanced quadratic reliable broadcast and improved asynchronous verifiable information dispersal. Cryptology ePrint Archive, Report 2022/052, 2022. <https://eprint.iacr.org/2022/052>.
14. S. Duan, M. K. Reiter, and H. Zhang. BEAT: Asynchronous BFT made practical. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 2028–2041. ACM Press, 2018.
15. Y. Frankel, P. D. MacKenzie, and M. Yung. Adaptively-secure optimal-resilience proactive RSA. In *Advances in Cryptology - Asiacrypt '99*, volume 1716 of *LNCS*, pages 180–194. Springer, Heidelberg, 1999.
16. D. Ghinea, C.-D. Liu-Zhang, and R. Wattenhofer. Optimal synchronous approximate agreement with asynchronous fallback. Cryptology ePrint Archive, Report 2022/354, 2022. <https://eprint.iacr.org/2022/354>.
17. Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Symposium on Operating Systems Principles (OSDI)*, pages 51–68. ACM, 2017.
18. S. D. Gordon, J. Katz, R. Kumaresan, and A. Yerukhimovich. Authenticated broadcast with a partially compromised public-key infrastructure. *Information and Computation*, 234:17–25. Elsevier, 2014.
19. J. Groth. Non-interactive distributed key generation and key resharing. Cryptology ePrint Archive, Report 2021/339, 2021. <https://eprint.iacr.org/2021/339>.
20. B. Guo, Z. Lu, Q. Tang, J. Xu, and Z. Zhang. Dumbo: Faster asynchronous BFT protocols. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 803–818. ACM Press, 2020.
21. A. Herzberg, S. Jarecki, H. Krawczyk, and M. Yung. Proactive secret sharing or: How to cope with perpetual leakage. In *Advances in Cryptology - Crypto '95*, volume 963 of *LNCS*, pages 339–352. Springer, Heidelberg, 1995.

22. I. Keidar, E. Kokoris-Kogias, O. Naor, and A. Spiegelman. All you need is DAG. In *Symposium on Principles of Distributed Computing (PODC)*, pages 165–175. ACM, 2021.
23. K. Kursawe and V. Shoup. Optimistic asynchronous atomic broadcast. In *Intl. Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3580 of *LNCS*, pages 204–215. Springer, Heidelberg, 2005.
24. C. Liu, S. Duan, and H. Zhang. EPIC: Efficient Asynchronous BFT with Adaptive Security. In *Intl. Conf. on Dependable Systems and Networks (DSN)*, pages 437–451. IEEE, 2020.
25. Y. Lu, Z. Lu, Q. Tang, and G. Wang. Dumbo-MVBA: Optimal multi-valued validated asynchronous byzantine agreement, revisited. In *Symposium on Principles of Distributed Computing (PODC)*, pages 129–138. ACM, 2020.
26. S. K. D. Maram, F. Zhang, L. Wang, A. Low, Y. Zhang, A. Juels, and D. Song. CHURP: Dynamic-committee proactive secret sharing. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 2369–2386. ACM Press, 2019.
27. A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song. The honey badger of BFT protocols. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 31–42. ACM Press, 2016.
28. A. Momose and L. Ren. Multi-threshold byzantine fault tolerance. In *ACM Conf. on Computer and Communications Security (CCS)*, pages 1686–1699. ACM Press, 2021.
29. K. Nayak, L. Ren, E. Shi, N. H. Vaidya, and Z. Xiang. Improved extension protocols for byzantine broadcast and agreement. In *Intl. Symposium on Distributed Computing (DISC)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
30. R. Ostrovsky and M. Yung. How to withstand mobile virus attacks (extended abstract). In *Symposium on Principles of Distributed Computing (PODC)*, pages 51–59. ACM, 1991.
31. M. Rambaud and A. Urban. Asynchronous dynamic proactive secret sharing under honest majority: Refreshing without a consistent view on shares. Cryptology ePrint Archive, Report 2022/619, 2022. <https://eprint.iacr.org/2022/619>.
32. D. A. Schultz, B. Liskov, and M. Liskov. Mobile proactive secret sharing. In *Symposium on Principles of Distributed Computing (PODC)*, page 458. ACM, 2008.
33. E. Shi. Foundations of distributed consensus and blockchains. Book manuscript, 2020.
34. R. Vassantlal, E. Alchieri, B. Ferreira, and A. Bessani. Cobra: Dynamic proactive secret sharing for confidential BFT services. In *IEEE Symposium on Security and Privacy (SP)*, pages 1528–1528. IEEE Computer Society, 2022.
35. T. Yurek, Z. Xiang, Y. Xia, and A. Miller. Long live the honey badger: Robust asynchronous DPSS and its applications. Cryptology ePrint Archive, Report 2022/971, 2022. <https://eprint.iacr.org/2022/971>.

A Preliminaries

A.1 Threshold Signature Scheme

We describe the syntax of a threshold signature scheme:

Setup: $\text{pp} \leftarrow \text{TS.Setup}(1^\kappa, n, t)$. Given the security parameter, the number of participants and threshold, output public parameters pp .

Key generation: $(\text{pk}, \text{SK}, \text{VK}) \leftarrow \text{TS.KeyGen}(\text{pp})$. This is an interactive protocol where the outcome is a public key pk , a vector of private keys $\text{SK} = (\text{sk}_1, \dots, \text{sk}_n)$, where party P_i obtains only sk_i and a public vector of verification keys $\text{VK} = (\text{vk}_1, \dots, \text{vk}_n)$. Denote by $\text{PK} = (\text{pk}, \text{VK})$.

Partial signing: $\sigma_i \leftarrow \text{TS.Sign}(\text{PK}, \text{sk}_i, M)$. This algorithm takes in the public keys PK , a message M and a private key share sk_i and outputs a signature share σ_i .

Signature share verification: $b_i \leftarrow \text{TS.ShVer}(\text{PK}, M, i, \sigma_i)$. This algorithm takes in the public keys PK , a message M and a partial signature on that message and outputs 1 if the partial signature is valid and 0 otherwise.

Share combination: $\sigma \leftarrow \text{TS.Combine}(\text{PK}, M, (i, \sigma_i)_{i \in S})$. On input the public keys PK , the message M and a set of shares of size $|S| = t + 1$, this algorithm outputs a full signature σ if $\text{TS.ShVer}(\text{pk}, \text{VK}, m, i, \sigma_i)$ returned 1 for all shares and outputs \perp otherwise.

Verification: $b \leftarrow \text{TS.Verify}(\text{PK}, M, \sigma)$. Anyone can verify a signature σ for on message M under public keys PK by running the verification algorithm, which returns 1 to indicate that the signature is valid and 0 otherwise.

A.2 Cryptographic Sortition

We describe the syntax of cryptographic sortition:

Key generation: $(vk, sk) \leftarrow \text{VRF.KeyGen}(\text{pp})$. Given input pp , the key generation algorithm outputs a verification key vk and a secret key sk .

Evaluation: $(\rho, \pi) \leftarrow \text{VRF.Eval}(\text{pp}, sk, M)$. Given a secret key sk and message M , the evaluation algorithm outputs a value ρ and proof π .

Verification: $b \leftarrow \text{VRF.Verify}(\text{pp}, vk, x, \rho, \pi)$. Given a verification key vk , input M , output ρ and proof π , the verification algorithm outputs 1 if the output ρ and proof π are valid with respect to vk and x and outputs 0 otherwise.

B Liveness Bound

The following lemma is based on the pigeonhole principle.

Lemma 11. *Assume there are at most t_y corrupted parties. If there are $n - t_s - t_x$ honest entries in a pre-block output in any given epoch, there exists at least*

$$M_\alpha < 1 + (n - t_s - t_x) \frac{1 - \alpha + \frac{n - t_y}{e^{(n - t_x - t_s)}}}{1 - \alpha \frac{n - t_s - t_x}{n - t_y} + \frac{1}{r}},$$

honest parties who output more than $\alpha r \frac{n-t_s-t_x}{n-t_y}$ times over r epochs for $0 < \alpha \leq 1$, and at least

$$M_0 \leq n - t_x - t_s$$

honest parties who output more than once over r epochs.

If the network is synchronous, we have $t_x = 0$ and $t_y = t_s$. If the network is asynchronous, we have $t_x = t_y = t_a$ if there are no key exposures, and $t_x = t_y = t_s$ if there are key exposures.

Proof. Assume this does not hold, i.e., for $0 < \alpha \leq 1$ and M_α in the statement, there do not exist M_α honest parties that output at least $\alpha r \frac{n-t_s-t_x}{n-t_y}$ times over r epochs. Then, in expectation, there exist at least $n - t_y - M_\alpha + 1$ honest parties that output less than $\alpha r \frac{n-t_s-t_x}{n-t_y}$ times over r epochs. Then the number of honest entries output over r epochs will be bounded below and above by:

$$\begin{aligned} r(n - t_s - t_x) &\leq (n - t_y - M_\alpha + 1) \left(\alpha r \frac{n - t_s - t_x}{n - t_y} - 1 \right) + r(M_\alpha - 1) \\ M_\alpha &\geq 1 + (n - t_s - t_x) \frac{1 - \alpha + \frac{n-t_y}{e(n-t_x-t_s)}}{1 - \alpha \frac{n-t_s-t_x}{n-t_y} + \frac{1}{r}}, \end{aligned} \quad (1)$$

which contradicts the assumption at the beginning of the proof.

Consider $\alpha = 0$ and there are only $M_0 - 1$ parties that ever output. This means that there can be at most $r(M_0 - 1)$ outputs. Therefore $r(M_\alpha - 1) \geq r(n - t_s - t_x)$, which contradicts the statement of the lemma. \square

Lemma 11 implies that at least $n - t_x - t_s$ honest parties output at least once over r epochs. Moreover, we always have at least one honest party consistently outputting over r epochs, since

$$\alpha \leq \min \left\{ 1, \frac{1 + (n - t_y - 1)/r(n - t_s - t_x)}{1 - 1/(n - t_y)} \right\} = 1, \quad (2)$$

as long as $n > \max\{t_s + t_x, t_y\}$, which always happens in our settings.

A similar argument follows for the committee-based protocols with $\hat{t}_s = (1 - \epsilon)t_s$ and $\hat{t}_a = (1 - \epsilon)t_a$ corruptions.

C Probability Bounds

Lemma 12 (Chernoff's inequalities). *Let X_1, X_2, \dots, X_n be independent random binary variables such that, for $1 \leq i \leq n$, $\mathbb{P}[X_i = 1] =: p_i$. Then, for $X := \sum_{i=1}^n X_i$ and $\mu := \mathbb{E}[X] = \sum_{i=1}^n p_i$:*

$$\mathbb{P}[X \geq (1 + \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2 + \delta}}, \quad 0 < \delta, \quad (3)$$

$$\mathbb{P}[X \leq (1 - \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2}}, \quad 0 < \delta < 1. \quad (4)$$

Lemma 13 (Hoeffding's inequality). *Let X be a variable sampled from a binomial distribution with n independent trials and probability of success p . Then:*

$$\mathbb{P}[X \geq k + np] \leq e^{-2k^2/n}. \quad (5)$$

Let X_1, \dots, X_n be random binary variables sampled with probability t/n and Y_1, \dots, Y_n be random binary variables sampled with probability $1 - t/n$. We want to bound the probability that $X = \sum_{i=1}^{\kappa} X_i$ is greater than a value s , and the probability that $Y = \sum_{i=1}^{\kappa} Y_i$ is less than a value s .

$$\mathbb{P}[X \geq s] \leq e^{-\frac{(s-t\kappa/n)^2}{s+2-t\kappa/n}}, \text{ if } s > t\kappa/n, \text{ by (3)}. \quad (6)$$

$$\mathbb{P}[Y \leq s] \leq e^{-\frac{\kappa(1-t/n-s/\kappa)^2}{1-t/n}}, \text{ if } 0 < s < (1-t/n)\kappa, \text{ by (4)}. \quad (7)$$

C.1 Unique Committees

In the ACS and BLA protocol, the committee \mathcal{C} of size κ is selected by using a collision-resistant hash function on unpredictable inputs. Hence, the event of choosing a particular party to be part of the committee can be thought of as an independent uniform at random sampling procedure.

Let Y_s denote the number of honest parties among the κ randomly elected committee members when the number of corrupted parties is $(1 - \epsilon)t_s$. The probability that there are fewer than $t_s\kappa/n$ honest parties in the committee is bounded by

$$\mathbb{P}[Y_s \leq \kappa t_s/n] \leq e^{-\frac{\kappa(1-(2-\epsilon)t_s/n)^2}{1-(1-\epsilon)t_s/n}}. \quad (8)$$

Since $t_s/n \leq 1/2$, we get $\mathbb{P}[Y_s \leq \kappa/2] \leq e^{-\frac{\epsilon^2\kappa}{1+\epsilon}}$.

The probability that there are fewer than $(1 - t_s/n)\kappa$ honest parties in the committee is bounded by

$$\mathbb{P}[Y_s \leq (1 - t_s/n)\kappa] \leq e^{-\frac{\kappa(\epsilon t_s/n)^2}{1-(1-\epsilon)t_s/n}}. \quad (9)$$

Since $t_s/n \leq 1/2$, we get $\mathbb{P}[Y_s \leq \kappa/2] \leq e^{-\frac{\epsilon^2\kappa}{2+2\epsilon}}$.

The probability that there are more than $t_a\kappa/n$ actively corrupted parties in the committee when the threshold of corruptions is $(1 - \epsilon)t_a$ is bounded by

$$\mathbb{P}[X_a \geq \kappa t_a/n] \leq e^{-\frac{(\epsilon\kappa t_a/n)^2}{2-\epsilon\kappa t_a/n}}. \quad (10)$$

Since $t_a/n \leq 1/3$, we get $\mathbb{P}[X_a \geq \kappa/3] \leq e^{-\frac{\epsilon^2\kappa^2}{3(6-\epsilon\kappa)}}$.

The probability that there are more than $t_s\kappa/n$ exposed parties in the committee when the threshold of corruptions is $(1 - \epsilon)t_s$ is bounded by

$$\mathbb{P}[X_s \geq \kappa t_s/n] \leq e^{-\frac{(\epsilon\kappa t_s/n)^2}{2+\epsilon\kappa t_s/n}}. \quad (11)$$

Since $t_s/n \leq 1/2$, we get $\mathbb{P}[X_s \geq \kappa/2] \leq e^{-\frac{\epsilon^2\kappa^2}{2(4+\epsilon\kappa)}}$.

C.2 Self-Elected Committees

The committees obtained through self-election have in expectation κ members, obtained using a VRF with output length of κ bits and bound $\mathbf{b} = \kappa 2^\kappa / n$. Applying the VRF can be idealized as flipping coins for each party to determine whether they (κ of them) are in the committee or not.

Let Z denote the number of parties selected in the secondary committee $\bar{\mathcal{C}}$. The expected value of Z is $\mathbb{E}[Z] = \kappa$. By (4), the probability that committee $\bar{\mathcal{C}}$ has strictly fewer than $(1 - \epsilon)\kappa + 1$ members is:

$$\mathbb{P}[Z \leq (1 - \epsilon)\kappa] \leq e^{-\epsilon^2 \kappa / 2}. \quad (12)$$

By (3), the probability that committee $\bar{\mathcal{C}}$ has more than $(1 + \epsilon)\kappa$ members is:

$$\mathbb{P}[Z \geq (1 + \epsilon)\kappa] \leq e^{-\epsilon^2 \kappa / (2 + \epsilon)}. \quad (13)$$

Let Y_s denote the number of honest parties among the randomly elected secondary committee members. In expectation, the number of honest parties selected will be greater than the initial fraction of honest parties times the committee size: $\mathbb{E}[Y_s] \geq (1 - (1 - \epsilon)t_s/n)\kappa$, so equation (8) holds for the secondary committee as well. Analogously, equation (11) holds for X_a , the number of corrupted parties among the randomly elected secondary committee members when the initial corruption threshold is t_a .

Denote by W the number of committees that have more than $\kappa t_s / n + 1$ corrupted/exposed members. Note that because the selection of the secondary committees is independent (also of the corruption selection), W is a binomial variable with probability of success $p := P[X \geq \kappa t_s / n]$ out of κ independent trials. We are interested in bounding the probability of W being more than $t_a \kappa / n$, which is the cumulative distribution function of a binomial random variable with parameters (p, κ) . Using the Hoeffding inequality (5), we obtain:

$$\mathbb{P}[W \geq t_a \kappa / n] \leq e^{-2\kappa(t_a/n - p)^2}. \quad (14)$$

For $t_a/n \leq 1/3$ we get $\mathbb{P}[W \geq \kappa/3] \approx e^{-2\kappa(1/3 - e^{-\epsilon^2 \kappa^2 / 2(4 + \epsilon \kappa)})^2}$.

D Technical Details

D.1 Update SMR proofs

Proof. (Theorem 1) We start with (1). Say a honest party P output a valid block in epoch e . Then P must have generated output in Π_{ACS} in epoch e , call it B , and at least $t_s + 1$ decryption shares on B were gathered. By t_s -validity with termination of Π_{ACS} , all honest parties will output B if they started Π_{ACS} with a valid pre-block B , so to prove t_s -consistency of Π_{SMR} , it remains to show that all honest parties input the same B to Π_{ACS} . Since the network is synchronous,

by time $T_e + \Delta$, all honest parties have managed to assemble a $(n - t_s)$ -quality pre-block B_i^e and input it to Π_{BLA} , which is t_s -terminating, t_s -valid and t_s -consistent, so it terminates by time $T_e + (5R + 1)\Delta$ with the same valid output B . Finally, t_s -completeness follows from t_s -validity with termination of Π_{ACS} .

We now address (2). Say a honest party P output a valid block in epoch e . P must have generated output in Π_{ACS} in epoch e , call it B , and gathered at least $t_s + 1$ decryption shares on B . By t_a -consistency of Π_{ACS} , all honest parties should have generated B in epoch e , so this proves t_a -consistency of Π_{SMR} . Every honest party will eventually assemble a valid $n - t_s$ -quality pre-block B_i^e , either as an output of Π_{BLA} if it terminates, or by waiting until $n - t_s$ codewords multicast by honest parties are delivered for at least $n - t_s$ parties. By t_a -consistency and t_a -termination of Π_{ACS} , all honest parties will output the same pre-block B in epoch e , and therefore there are at least $t_s + 1 \leq n - t_a$ valid decryption shares (for the same B). This ensures each honest party successfully recovers a block, proving t_a -completeness of Π_{SMR} . \square

Proof. (Theorem 2) By Theorem 1, all honest parties will output the same valid block, obtained by decrypting the output of Π_{ACS} , which means that they have valid certificates of correct decryption. External validity follows from the fact that the adversary cannot generate invalid certificates because it controls fewer than $t_s + 1$ parties. \square

Proof. (Theorem 3) Assume all honest parties (at least $n - t_s$ in (1) and at least $n - t_a$ in (2)) have received a transaction tx . If by some epoch e , tx is not in an honest party's buffer anymore it means it was output in $\text{blocks}[e']$ for $e' < e$. Then, by consistency of Π_{SMR} proven in Theorem 1, tx will not be in any honest party's buffer after epoch e' . Otherwise, suppose tx is still in an honest party P 's buffer at epoch e . By completeness of Π_{SMR} proven in Theorem 1, each party outputs a block in every epoch. This block is obtained by decrypting a pre-block of $(n - t_s)$ -set quality to which at least $n - t_s - t_a$ honest parties contributed L/n transactions, by Lemma 14 proved below. Thus, each honest party that contributes removes in expectation at least L/n transactions from their buffer in each epoch. Assuming parties cannot receive an infinite amount of transactions in a finite number of epochs, there will be a finite number of transactions in P 's buffer alongside tx . By the lower bound in Lemma 11 (Appendix B), honest parties continue to clear transactions from their buffers so that eventually tx appears among the first L transactions of their buffers. Once this has occurred, the probability that tx fails to appear in the output block at the e 'th epoch if at least one of the honest parties that contributes its input to the block has tx among the first L transactions of its buffer is at most $1 - 1/n$. Thus, a transaction tx is included in $\text{blocks}[e : e + r]$ with probability at least $1 - (1 - 1/n)^{r+1}$, which approaches 1 as r goes to infinity. ⁴ \square

⁴ The notion of eventual liveness considered here is standard under asynchrony. In the synchronous case, a more detailed analysis of liveness can be made, e.g. following the approach used in [7].

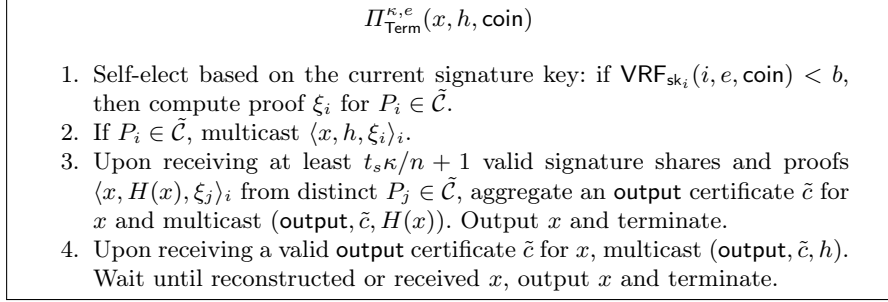


Fig. 11. Termination helper protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$.

Lemma 14. *Under condition (*), at least $n - t_s - t_a$ honest parties have contributed transactions in any block output by a honest party in Π_{SMR} .*

Proof. All honest parties input valid pre-blocks in Π_{BLA} and Π_{ACS} , meaning that they wait to receive at least $n - t_s$ validly signed encrypted entries. By the t_s -security of Π_{BLA} , if Π_{BLA} outputs, it outputs a $(n - t_s)$ -quality pre-block; even if the network is asynchronous, an honest party would not output an invalid pre-block. Therefore, honest parties' inputs to Π_{ACS} are also $(n - t_s)$ -quality.

In case the network is asynchronous and there are at most t_a corrupted parties, $n - t_s - t_a$ entries in the pre-block originate from honest parties. By t_a -set quality of Π_{ACS} (Lemma 4), the output of Π_{ACS} contains at least a pre-block of $(n - t_s)$ -quality, therefore with $(n - t_s - t_a)$ honest entries.

In case the network is synchronous, each honest party has received all messages from all other honest parties upon reaching Step 2 of Π_{SMR} , so the number of honest entries in their pre-blocks is at least $(n - t_s)$. Moreover, all honest parties complete Π_{BLA} with the same output pre-block B containing $(n - t_s)$ honest entries. By the t_s -validity with termination of Π_{ACS} (Lemma 3), the output pre-block of Π_{ACS} is also B . \square

D.2 Committee-Based ACS

Throughout this section, we consider a network that is asynchronous with up to $\hat{t}_a = (1 - \epsilon)t_a$ corruptions.

In $\Pi_{\text{ACS}}^{\kappa}$, the termination conditions OC 1 and OC 2 make calls to $\Pi_{\text{Term}}^{\kappa, e}$.

Sketch of proof. (Lemma 7) We discuss the changes arising from the use of committees in the proofs for the properties of the $\Pi_{\text{ACS}}^{\kappa}$ protocol. The proof will then follow from the proofs of Lemmata 1–6.

The static adversary cannot tamper with the election of the primary committee because it can corrupt only up to \hat{t}_s parties, while the signature aggregation requires $\hat{t}_s + 1$ signatures. The election of the secondary committees is done independently and in parallel, based on the coin computed this epoch. An adversary




Fig. 12. BLA protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$ in epoch e .

cannot tamper with these elections because of the unforgeability of the signature scheme and cannot predict the membership from previous epochs.

A committee election is unpredictable and modeled as a uniformly random sampling of κ parties (in the primary committee) or $O(\kappa)$ parties (in the other committees) from the pool of n parties. In expectation, the fraction of corrupted parties over all parties will be reflected in the committee. We select parameters κ and ϵ such that with high probability, there are at most $t_x \kappa / n$ corrupted parties in the committees and at most $t_x \kappa / n$ secondary committees contain more than $t_x \kappa / n$ corrupted members, where $t_x = t_a$ in the asynchronous case and $t_x = t_s$ in the synchronous case. The failure probabilities are given in Appendix C, using standard arguments based on the Chernoff and Hoeffding bounds. \square

D.3 Committee-Based BLA

Throughout this section, we consider a network that is synchronous with up to $\hat{t}_s = (1 - \epsilon)t_s$ corruptions. We use the following validity definition that depends on the committee size:

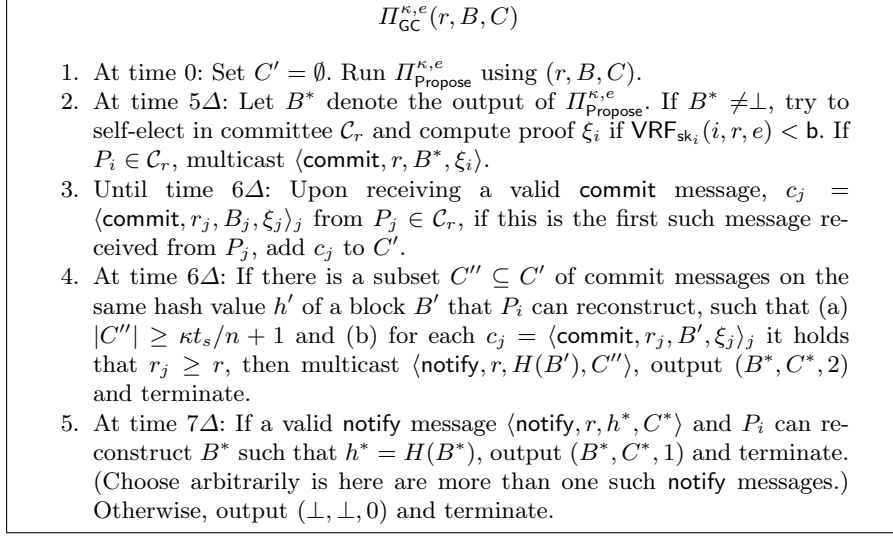
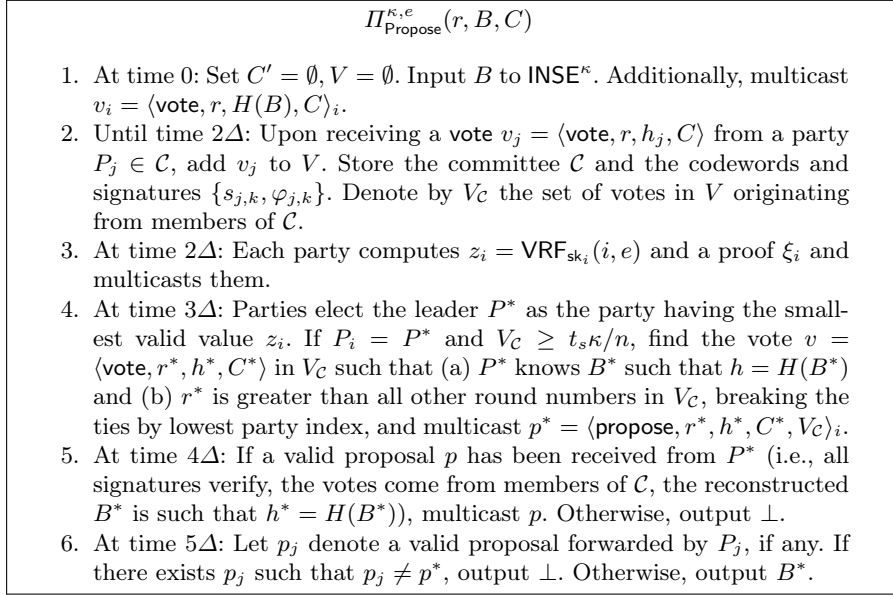
- (*t-Validity*) If every honest party has input an $(1 - t/n)\kappa$ -quality pre-block, then every honest party outputs an $(1 - t/n)\kappa$ -quality pre-block.

Lemma 15. $\Pi_{\text{BLA}}^{\kappa}$ achieves \hat{t}_s -termination, \hat{t}_s -validity and \hat{t}_s -consistency except with negligible probability.

In each call of $\Pi_{\text{Propose}}^{\kappa}$, parties encode their inputs and send the codewords and hashes to the other parties, such that honest parties are able to reconstruct the proposed pre-block of the leader. Since parties might output a different pre-block than the one they started with in the current round, they have to send the hash and the codewords of their new input during the next round.

We sketch the main ideas of the proof below. The full proofs can be obtained by expanding the BLA proofs in [7].

Sketch of proof. (Lemma 15) Parties terminate after participating in all R rounds of the protocol (even if they generated output earlier), so \hat{t}_s -termination is ensured by design.

**Fig. 13.** GC protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$ in round r .**Fig. 14.** Propose protocol from the perspective of party $P_{i \in \{1, \dots, n\}}$ in round r .

We now argue \hat{t}_s -validity. Suppose all honest parties start Π_{BLA}^κ with input B of quality $t_s \kappa/n$. If the leader is honest, all honest parties terminate $\Pi_{Propose}^\kappa$

with a pre-block B^* of quality $t_s\kappa/n$. This is because honest parties agree on the committee and can reconstruct their pre-blocks after INSE^κ since the code tolerates t_s erasures (honest parties distinguish invalid signatures and ignore codewords with invalid signatures). Except with negligible probability, there will be at most $t_s\kappa/n$ corrupted parties in the primary committee, so V_C cannot contain $t_s\kappa/n + 1$ votes of corrupted parties. Thus, the proposal will be valid, ensuring all honest parties receive only valid proposals by time 5Δ and will output B^* . On the other hand, if the leader is dishonest, it can refuse to send a valid proposal by the required time. However, it cannot force honest parties to accept an invalid proposal since the adversary cannot tamper with the election of the primary committee (the network is synchronous and at most) and the corrupted leader cannot forge the signatures of the honest parties that are in the primary committee. In this case, the honest parties may terminate $\Pi_{\text{Propose}}^\kappa$ with output \perp .

In Π_{GC}^κ , with high probability, at least $(1 - t_s/n)\kappa$ members of \mathcal{C}_r will be honest and send B^* in step 2 of Π_{GC}^κ , which will reach all honest parties by the beginning of the next round. Therefore, in step 5 of Π_{GC}^κ , a party receiving a valid notify message will be able to determine to what block B^* the hash h^* corresponds and to output B^* . Moreover, fewer than $t_s\kappa/n$ parties in \mathcal{C}_r are corrupted, so there cannot be sufficient votes in C'' if no honest party participates. (Recall that honest parties do not output invalid B^* from $\Pi_{\text{Propose}}^\kappa$.) Therefore, if the leader is honest in r^* , all honest parties output a valid B^* with grade 2 and thus output in Π_{BLA}^κ . Since $\hat{t}_s/n < 1/2$, a dishonest leader is elected with probability smaller than $1/2$, so the probability no honest leader is elected in R rounds is negligible. This proves \hat{t}_s -validity.

To prove \hat{t}_s -consistency, suppose r^* is the first round in which some honest party P_i has output a pre-block B . P_i must have generated a notify message and output with grade 2 in Π_{GC}^κ . Then no honest party can output with grade 0, and all honest parties must have received that notify message in the same round. Therefore, all honest parties will use B as input in iterations greater than $r^* + 1$. Moreover, no honest party could have sent a commit message on a different pre-block $B' \neq B$ in the same execution of Π_{GC}^κ , and so a corrupted leader cannot construct a valid vote on β' in a subsequent round number. Inductively, we can argue that honest parties will keep inputting B and not voting on other blocks in all subsequent rounds until R , so all honest parties will output B at the end of Π_{BLA}^κ . \square

If the network is asynchronous, note that we cannot guarantee any of the termination, consistency and validity properties. However, in the SMR protocol, the BLA step is followed by ACS, so that ACS compensates for BLA if the network is asynchronous (and likewise BLA compensates for ACS if the network is synchronous). This is the reason it suffices to show that BLA achieves termination, consistency and validity in a synchronous network.

D.4 Adaptive Attack on Network-Agnostic Committee-Based SMR

We describe an attack on Upstate from an adaptive adversary. Designing an adaptively secure SMR protocol in the network-agnostic case that achieves $O(n^2)$ communication per $O(n)$ -block is a challenging open problem.

If the network is purely synchronous, Upstate SMR is adaptively secure. The reason is that all parties start each part of the protocol at the same time and the individual subprotocols are designed such that an adaptive adversary cannot learn information about committee membership before the honest parties have taken their actions.

In a multi-epoch SMR protocol run in a asynchronous network, the adversary can delay a party from proceeding to epoch $e + 1$. Since \hat{t}_a delayed parties are indistinguishable from \hat{t}_a corrupted parties, the protocols should be designed to proceed and terminate while requiring at most $n - \hat{t}_a$ messages. Therefore, the adversary can delay up to \hat{t}_a parties from starting, e.g., the INSE part in Π_{ACS}^κ . The protocol proceeds and generates the coin, and the adversary can learn the primary committee membership before the delayed members sent their messages.

Say the adversary has corrupted $f < \hat{t}_a$ parties so far in this epoch, and has a budget of $\hat{t}_a - f$ remaining corruptions. In expectation, out of the corrupted f parties, $f\kappa/n$ of them will be members, and out of the delayed \hat{t}_a parties, $t_a\kappa/n$ will be members. However, $\hat{t}_a - f \geq t_a\kappa/n$, for a variety of values of f , so the adversary can corrupt $(t_a + f)\kappa/n$ members of the committee. This is more than the $t_a\kappa/n$ allowed committee corruptions and breaks the \hat{t}_a -security of Π_{ACS}^κ .

We note that a self-elected primary committee resolves the attack in the sense that the adversary is not able to corrupt and control the inputs of more than $t_a\kappa/n$ committee members. However, there will be no agreement between honest parties on the membership in the primary committee in the asynchronous case, so we cannot simultaneously achieve \hat{t}_a -consistency and \hat{t}_s -validity in the Π_{ACS}^κ . Designing an adaptively secure SMR protocol in the network-agnostic case that achieves $O(n^2)$ communication per $O(n)$ -block is a challenging open problem.

D.5 Communication Complexity

Let ℓ be the input length and let b be the number of blocks into which the input is divided, so that each codeword has size ℓ/b . Also note that we consider a regime where $\kappa < n$.

Protocol Π_{ACS} (Section 4). Step 2 of INDI involves n^2 messages of $O(\kappa + \ell/b)$ bits, so the total communication is $O(\kappa n^2 + n^2 \ell/b)$. Resending the received codewords and signatures corresponding in step 3 of INDI takes $O(n^3 \ell/b + \kappa n^3)$ total communication. Step 2 of Π_{ACS} where parties send their votes takes $O(\kappa n^3)$. A certificate contains a full signature of size κ , so forwarding the received certificates takes $O(\kappa n^3)$ communication in step 3. Using an instantiation of asynchronous binary BA with quadratic communication, the communication for all n BA instances takes $O(\kappa n^3)$ communication. Π_{Term} is run only on inputs of size ℓ , and requires $O(n^2 \ell + \kappa n^2)$ communication. We consider a regime where $\kappa < n$.

Therefore, the asymptotic communication of Π_{ACS} is $O(n^2l + n^3l/b + \kappa n^3)$ per input of size ℓ .

The parameter b of the error correcting code is chosen to be $b = t_s = O(n)$, leading to a communication of $O(n^2l + \kappa n^3)$. This yields an amortized communication of $O(n^2)$ per input of size $\ell = O(\kappa n)$.

Protocol Π_{ACS}^κ from Section 6 Step 3 of INSE^κ involves n^2 messages of $O(\kappa + l/b)$ bits, so the total communication is $O(\kappa n^2 + n^2l/b)$. Resending the received codewords and signatures corresponding to the κ parties in the primary committee in step 4 of INSE^κ takes $O(\kappa n^2l/b + \kappa^2 n^2)$ total communication. Step 2 of Π_{ACS}^κ where members of κ secondary committees send their signatures on hashes to all other parties takes $O(\kappa^3 n)$. A certificate contains $O(\kappa)$ signatures of size κ , so forwarding the received certificates corresponding to the proposals of the primary committee takes only $O(\kappa^3 n)$ communication in steps 3 and 4. Using an instantiation of asynchronous binary BA with $O(\kappa n^2)$ communication, the communication for all κ BA instances takes $O(\kappa^2 n^2)$ communication. Termination conditions OC 1 and OC 2 incur another $O(\kappa n l + \kappa^2 n^2)$ total communication. Overall, the total communication of Π_{ACS}^κ is $O(\kappa n^2l/b + \kappa n l + \kappa^2 n^2)$ communication per block.

The parameter of the error correcting code $b = \hat{t}_s = O(n)$. To have total communication $O(\kappa^2 n^2)$ per block, we can set $\ell \leq O(\kappa n)$ so $\ell/b \leq O(\kappa)$. This yields an amortized communication of $O(\kappa n)$.

Protocol Π_{BLA}^κ (Section 6). The commit list can have $R\kappa^2$ size. $\Pi_{\text{Propose}}^\kappa$ incurs $O(\kappa^2 n^2 + \kappa n^2l/b)$ total communication from INSE^κ , $O(n^2\kappa^2 R)$ from the vote and propose multicasts, and $n^2\kappa$ for the leader election, for a total of $O(\kappa^2 n^2 R + \kappa n^2l/b)$. Additionally, Π_{GC}^κ contributes $O(\kappa n l + \kappa^2 n^2 R)$ to the communication cost. These steps are run for $R < \kappa$ rounds, so that in total Π_{BLA}^κ incurs $O(R\kappa^2 n^2 + \kappa n^2l/b + \kappa n l)$ communication complexity.

To obtain $O(R\kappa^2 n^2)$ communication complexity, we can set $\ell = O(\kappa n R)$ so that, yielding an amortized communication per block of $O(\kappa n)$.

Upgrade [7] achieves a total communication complexity of $O(\kappa^2 n^3 + \kappa n L |\text{tx}| + \kappa^3 n + L |\text{tx}| \kappa^2)$ per block (the number of rounds R is absorbed), which enables a linear amortized cost for a block size of $L = O(\kappa n^2)$. Such a block size implies a communication complexity of $O(\kappa^2 n^3)$ only to distribute them to a small committee of size κ , which already exceeds the quadratic communication limit we imposed for **Upstate**. Nevertheless, both **Upgrade** and **Upstate** need to be run over hundreds of thousands of parties in order to achieve asymptotic communication improvements over **Tardigrade** and **Update**. In this large-scale context, assuming that parties have $O(n^2)$ transactions in their buffers at every epoch is quite restrictive. Therefore, we believe that providing a protocol with linear amortized communication per $O(n)$ transactions has more relevance in practice.

E Network Changes

Proof. (Lemma 8) By t_a consistency and validity of BA, at least one honest party needs to input 1 in a BA instance in order for it to output 1. This means honest parties need to be able to construct at least $n - t_a$ certificates. Clearly, certificates corresponding to the $n - t_s$ honest and unexposed parties can be eventually reconstructed. Therefore, we focus on the case of building a certificate for an exposed party P .

When multicasting messages in step 3 of INDI, corrupted parties can send erroneous codewords on behalf of P . Therefore, in RECON, up to t_a of the at least $n - t_a$ codewords can have a valid signature but are erroneous (but need to have the same hash h in order to be taken into consideration). While the code cannot tolerate at the same time t_a errors and t_s erasures, with overwhelming probability, the value x output by DEC on erroneous codewords will not satisfy $h = H(x)$. Therefore, honest parties wait for more correct codewords, which are guaranteed to eventually arrive, since $n - t_a$ parties behave honestly, so honest parties can assemble certificates for exposed parties as well.

Proof. (Lemma 9) Since the adversary cannot act on behalf of the exposed parties directly, the arguments for when P is honest and unexposed, and for when P is dishonest are the same as the arguments in the proof of Lemma 1.

Assume P is exposed. The same argument as for a dishonest P that sends codewords in the first round of INDI applies for an exposed P . Assuming H is a collision-resistant hash function, there do not exist values $x \neq x'$ reconstructed by different sets of codewords such that $h = H(x) = H(x')$. Therefore, if after inputting $n - t_s$ codewords to RECON and not obtaining a valid output with respect to h , the honest parties wait until they receive $n - t_s + t_a$ codewords in order to be able to correctly reconstruct. \square

Proof. (Lemma 10) Since the adversary cannot act directly on behalf of the exposed parties, the arguments for when P is honest and unexposed and when P is dishonest can be taken directly from the proof of Lemma 2.

Suppose P is exposed. The same argument as for an exposed P ensures that because $n > 2t_s + t_a$, there needs to be one honest party that would sign both certificates, implying $h = h'$. \square

The adversary can include more malicious entries in the block by forging the signatures of t_s instead of t_a parties. However, asymptotically, the honestly generated throughput remains the same as in the network-agnostic case.

E.1 Upstate SMR under Network Changes

Sketch of proof. (Theorem 5, Upstate) Despite knowing \hat{t}_s keys, a static adversary cannot actively corrupt more than $t_a \kappa / n$ parties in any of the committees with high probability (see Appendix C), because it selects the corrupted parties before the epoch starts and committee membership is unpredictable.

The arguments for t_a -security of Upstate under arbitrary network changes follow from the arguments in the proof of Theorem 5. \square

As a side note, an adaptive adversary knowing \hat{t}_s keys would be able to corrupt up to $t_s \kappa / n$ parties in the secondary committees. Nevertheless, since the secondary committees are only used to construct certificates of $\hat{t}_s + 1$ keys, this would not be a problem.

E.2 Tardigrade SMR under Network Changes

Sketch of proof. (Theorem 5, Tardigrade) Let $\overline{\Pi_{\text{ACS}}}$ denote the asynchronous common subset from Tardigrade.

ACS. The t_s -valid and t_a -consistent reliable broadcast protocol used in $\overline{\Pi_{\text{ACS}}}$ is signature-free, and furthermore, the non-terminating ACS protocol is signature-free (apart from the BA components). Thus, the proofs for t_a -consistency, t_s -validity, t_a -liveness, and t_a -set quality of the ACS protocol from [7, Sec. 4] hold.

The terminating wrapper of the ACS protocol, $\overline{\Pi_{\text{ACS}}}$, uses threshold digital signatures. Thus, the adversary can forge the threshold signatures of up to t_s parties. However, it cannot create acceptable certificates on its own. An honest party will sign an output if and only if it has already terminated ACS with that output. Thus, there can never be a valid certificate for an invalid output. Therefore, $\overline{\Pi_{\text{ACS}}}$ is t_a -consistent, t_a -live, t_a -terminating, t_s -valid with termination and has t_a -set quality [7, Sec. 4].

The arguments for Π_{BLA} and Π_{SMR} are similar to the ones in the proof of Theorem 5. \square

We note that we recover a similar result as [18] concerning synchronous authenticated broadcast for $n > 2t_s + \min(t_s, t_e - t_s)$, where t_s is the number of corrupted parties and t_e is the number of exposed parties.

F Discussion on the Model

Mobile corruptions. Without even mild synchrony assumptions, such as a time signal from an external clock, a mobile adaptive adversary that can corrupt any of the n parties can break either privacy or liveness of a proactive threshold protocol, as we describe in Section 8. This motivates a limiting assumption over the adversary's powers: the overall number of parties that can be corrupted over the duration of the protocol is below the total number of parties, and the adversary can only be mobile in between that smaller set of parties. In order to further characterize the adversarial mobility, we need to granulate the time. We do so by defining local epochs, characterized by the parties' actions, not by a clock. We assume that the adversary can only corrupt a limited number of parties at any given moment in time, and in any given epoch. Without this assumption, an adversary could corrupt the maximum number of parties at an onset of the epoch, take some actions and deliver only those associated messages, then immediately move into a different maximal set of parties, which in the asynchronous case would be equivalent to corrupting twice the allowed threshold.

Channels. An initially agreed upon public key infrastructure (initially) ensures agreement and confidentiality. However, as the adversary corrupts the parties and learns their secret keys, authentication and confidentiality are not ensured anymore, even if the adversary is flushed out from the party, since it can use the learned keys to decrypt communication, forge signatures, and create messages that appear to have been created by the corresponding party. Without an assumption of authenticated channels for all the duration of the protocol, the adversary can spawn other parties in the network to impersonate the uncorrupted parties. This assumption is standard in other proactive protocols [9,32].

G Asynchronous Proactive Secret Sharing

The game proceeds as follows between a t_a -limited adversary and a challenger.

1. The adversary chooses two secrets $s_0, s_1 \in \mathbb{F}$ and gives them to the challenger.
2. The challenger chooses $b \leftarrow \{0, 1\}$ and runs $(s_1^{(0)}, \dots, s_n^{(0)}) \leftarrow \text{Share}(s_b)$.
3. For an epoch number $\tau \geq 0$ the adversary specifies a set of parties \mathcal{C}^τ , where $|\mathcal{C}^\tau| \leq t_a$ and sends it to the challenger. Denote by \mathcal{H}^τ the set of honest parties in epoch τ . The challenger sends to the adversary the corresponding state of the parties, which includes the shares $\{s_i^{(\tau)}\}_{i \in \mathcal{C}^\tau}$.
4. The adversary sends replacing shares $\{s_i^{(\tau)'}\}_{i \in \mathcal{C}^\tau}$ to the challenger, which sets $s_i^{(\tau)} = s_i^{(\tau)'}$ for $i \in \mathcal{C}^\tau$ and specifies the actions of the parties in \mathcal{C}^τ . The adversary specifies the set of messages \mathcal{M} (including other messages than the ones associated to the proactive VSS protocol) to be delivered to the honest parties from the set \mathcal{C}^τ . It also specifies to the challenger which messages between the remaining honest parties are to be delivered over the network; note that this means that the adversary can refuse to deliver a message originated in epoch $\tau' < \tau$ from a party in $\mathcal{C}^{\tau'}$ to a party in $\mathcal{H}^{\tau'}$.
5. For a epoch τ for which the challenger has a set \mathcal{S} of at least $n - t_a$ associated shares, the challenger runs $(s_1^{(\tau+1)}, \dots, s_n^{(\tau+1)}) \leftarrow \text{Reshare}(\tau, \{s_i^{(\tau)}\}_{i \in \mathcal{S}})$.
6. If τ is as large as desired by the adversary, go to step 7. Otherwise, increment the epoch number by 1 and go to step 3.
7. Eventually, the adversary outputs a guess $b' \in \{0, 1\}$ of b .

Denote the adversary's advantage by $\text{Adv}(\mathcal{A}) = |\Pr[b' = b] - \frac{1}{2}|$.