A New Paradigm for Verifiable Secret Sharing

Sourav Das*, Zhuolun Xiang†, Alin Tomescu†, Alexander Spiegelman†, Benny Pinkas†, and Ling Ren*
*University of Illinois at Urbana-Champaign, †Aptos
{souravd2, renling}@illinois.edu, {xiangzhuolun, tomescu.alin, sasha.spiegelman}@gmail.com, benny@pinkas.net

***Abstract*—Verifiable Secret Sharing (VSS) is a fundamental building block in cryptography. Despite its importance and extensive studies, existing VSS protocols are often complex and inefficient. Many of them do not support dual threads, are not publicly verifiable, or do not properly terminate in asynchronous networks. This paper presents a new and simple paradigm for designing VSS protocols in synchronous and asynchronous networks. Our VSS protocols are optimally fault-tolerant, i.e., they tolerate a $1/2$ and a $1/3$ fraction of malicious nodes in synchronous and asynchronous networks, respectively. They only require a public key infrastructure and the hardness of discrete logarithms. Our protocols support dual thresholds and their transcripts are publicly verifiable. We implement our VSS protocols and measure their computation and communication costs with up to 1024 nodes. Our evaluation illustrates that our VSS protocols provide asynchronous termination and public verifiability with minimum performance overhead. Compared to the existing VSS protocol with similar guarantees, our protocols are 5-15× and 8-13× better in computation and communication cost, respectively.**

1. Introduction

A Verifiable Secret Sharing (VSS) scheme lets a party holding a secret, commonly referred to as a dealer, share the secret in a verifiable manner among a set of nodes where a fraction of the nodes, including the dealer, could be malicious [26], [34], [56]. The secret sharing process is verifiable in the sense that each node can verify the validity and correctness of its share. VSS is a fundamental building block for secure-multiparty computation (MPC) [7], [51], threshold cryptography [60], Byzantine fault tolerant algorithms [17], distributed key generation (DKG) [37], [33], randomness beacon [25], and so on.

Over the years, numerous works have studied VSS with different properties and in different settings, such as different cryptographic assumptions, network conditions, fault-tolerance, and so on [26], [34], [56], [37], [48], [49], [10], [32], [4], [64], [66], [65], [61]. In this paper, we focus on VSS protocols that use Shamir secret sharing [60], are secure against a computationally bounded adversary, and have optimal fault tolerance in both synchronous and asynchronous networks. We also seek to achieve a few extra nice properties of VSS that we briefly go over next.

A desirable property of VSS protocols is *completeness* which ensures every honest node receives its share of the secret. Applications such as DKG, MPC, and proactive secret sharing crucially rely on the completeness property.

Another desirable property of VSS, especially asynchronous VSS (AVSS), is the support for dual thresholds [16]. Briefly, in an asynchronous network of $n \geq 3t+1$ nodes where at most t nodes are malicious, a dual-threshold AVSS scheme with parameter $\ell \in [t, n-t)$ guarantees secrecy against any coalition of up to ℓ nodes. Dual-threshold AVSS with $\ell = n-t-1$ is used to design high-threshold asynchronous DKG [33], which is in turn used to achieve better secrecy in threshold cryptosystems [62] and better efficiency in Byzantine fault tolerant (BFT) algorithms [17], [18], [63]. Dual-threshold VSS is also useful in designing optimal fault-tolerant BFT systems that rely on sampling for scalability, an approach that is getting wide adoption in recent proof-of-stake blockchains [38], [24], [5].

Finally, some randomness beacon [10], [30] and DKG protocols [44], [41], [47] also require the VSS transcript to be *publicly verifiable* by any external entity. A VSS scheme with a publicly verifiable transcript is also called a Publicly Verifiable Secret Sharing (PVSS) scheme.

In this paper, unless stated otherwise, we always consider VSS protocols with the completeness property, and we primarily study VSS protocols that support dual thresholds in asynchrony and provide publicly verifiable transcripts.

Existing works. Despite years of efforts, there are no VSS schemes that satisfy all our requirements (see §2 for a detailed discussion). For example, the historically dominant approach of designing synchronous VSS protocols relies on interactive complaints [34], [56], [36], [37], [46], [10], [8]. This approach incurs high latencies, is fairly complex, and is not publicly verifiable. Moreover, when extended to asynchronous networks, this approach suffers from a subtle termination issue* [64], [32], [43], [61], and does not support dual thresholds. Several recent asynchronous VSS designs deviate from the interactive complaint framework. But these schemes rely on trusted setups and bilinear pairing for efficiency [49], [4], [3], [66], and they also do not support dual thresholds or public verifiability. On the other hand, existing publicly verifiable VSS uses *verifiable encryption* schemes to let the dealer prove statements over encrypted data, making them expensive [35], [41], [33], [47] or suitable only for limited applications [59], [22], [23].

*In these protocols, parties may never terminate the sharing phase, even if they already output. Malicious nodes can prevent honest nodes from terminating by not sending acknowledgments or complaints. More details can be found in §2.

Our contributions. In this paper, we present a new paradigm for designing VSS protocols for synchronous and asynchronous networks. Our VSS protocols are optimally fault-tolerant, i.e., they tolerate $1/2$ and $1/3$ fractions of malicious nodes in synchronous and asynchronous networks, respectively. Our VSS protocols guarantee completeness and have efficient publicly verifiable transcripts. Our asynchronous protocol also guarantees asynchronous termination without relying on additional cryptographic setups or bilinear pairings and only assumes public key infrastructure.

Our VSS protocols achieve the above-mentioned properties while maintaining the same asymptotic communication and computation costs of best-known VSS protocols. More precisely, in a synchronous network with n nodes, our VSS protocol incurs a communication cost of $O(\kappa n^2 + C_{\mathrm{BB}}(\kappa n))$. Here κ is a computational security parameter, and $C_{\mathrm{BB}}(x)$ is the communication cost of broadcasting a message of size x via a Byzantine broadcast channel. Similarly, our asynchronous VSS (AVSS) protocol incurs a communication cost of $O(\kappa n^2)$.

We then augment our AVSS to support *dual thresholds* for any secrecy threshold $\ell \in [t, n-t)$. Our augmented AVSS protocol maintains the total communication cost of $O(\kappa n^2)$ without relying on a trusted setup. Our dual-threshold AVSS protocol has the following nice properties: (i) The best-case performance with any ℓ is the same as our low-threshold AVSS, where the best-case is when the network is synchronous and the number of malicious nodes is less than $2t - \ell$; and, (ii) the worst-case performance degrades gradually with ℓ where it is the same as our low-threshold AVSS for $\ell = t$. In contrast, existing dual-threshold AVSS protocols [33], [42], [47] incur a high cost regardless of ℓ, and their performance does not improve even under the best-case scenario.

Another useful property of our VSS scheme is that nodes only need to communicate with the dealer. This also means that assuming the presence of a broadcast channel, the synchronous timing assumption needs to apply only between the dealer and other nodes. This assumption is less stringent than requiring bounded communication delays between all pairs of nodes. This property also makes the implementation simpler, as only the dealer needs to establish communication with the other nodes.

As an independent contribution, we design an efficient verifiable encryption scheme for Pedersen commitments. Existing verifiable encryption schemes are designed for the non-hiding Feldman commitment scheme and can not be used to encrypt messages with low entropy [35], [19], [42], [47]. Our verifiable encryption scheme addresses this limitation and supports arbitrary message distribution and is thus more suitable for general applications, including VSS.

Evaluation. We implement our VSS protocol in Rust and measure its computation and communication costs with up to 1024 nodes. Our evaluation illustrates that our AVSS protocol has comparable performance to [64], [32] while additionally achieving asynchronous termination and public verifiability. Compared to the existing VSS protocol with these properties [42], our protocol requires 5-15× and 8-13× less computation and communication.

Paper organization. The rest of the paper is organized as follows. First, we review related work in more detail in §2. In §3, we formally define the problem of Verifiable Secret Sharing and provide an overview of our new VSS paradigm. We describe the required preliminaries in §4. We then describe our synchronous VSS in §5, asynchronous VSS in §6, and dual-threshold asynchronous VSS protocol in §7. We then present our implementation and evaluation results in §8. Finally, we conclude with a discussion in §9.

2. Related Work

VSS protocols consist of two phases: *Sharing* and *Reconstruction*. During the sharing phase, nodes along with the dealer run a protocol so that each node receives its share of the secret at the end of the sharing phase. In the reconstruction phase, nodes interact to recover the shared secret. We categorize existing VSS schemes into three approaches based on the design of their sharing phase. We describe each approach and outline its core idea, advantages, and disadvantages below.

Complaint-based VSS. Historically, the most common approach to designing VSS protocols is to rely on interactive complaints [34], [56], [36], [37], [46], [10], [8]. Briefly, in these protocols, the dealer embeds the secret into a univariate low degree polynomial and publishes a commitment to the polynomial via a broadcast channel. The dealer additionally sends each node its share using a private channel. Upon receiving its share and the commitment, each node validates them for correctness. Nodes that receive no share or invalid shares from the dealer publish complaints against the dealer using the broadcast channel. The dealer responds to the complaints by revealing the share of each complaining node. Intuitively, these protocols rely on complaints to ensure completeness, i.e., prevent malicious dealers from sending valid shares to a subset of honest nodes and not to others.

While this approach provides reasonable efficiency in synchronous networks, they do not extend well to the more realistic partially synchronous and asynchronous networks. Asynchronous VSS (AVSS) protocols that rely on complaints to provide completeness [64], [32], [43], [61] suffer from a subtle termination issue that prevents honest nodes from terminating the protocol, even after outputting their shares. More concretely, these protocols have a step where honest nodes (after outputting their share) wait for either acknowledgments or complaints from all other nodes before terminating. This step is crucial because, in the case of complaints, nodes must assist the complaining nodes in recovering their shares. This allows malicious nodes to prevent honest nodes from terminating by simply not sending acknowledgments or complaints.

Finally, complaint-based VSS protocols have other limitations: they are not publicly verifiable, and they do not support dual thresholds.

Verifiable Encryption-based VSS. One approach to VSS design that addresses the above issues is to use *verifiable*

encryption (VE). Briefly, in a VE-based VSS scheme, the dealer locally generates a transcript that includes encryptions of the shares of all nodes, each under the public key of the corresponding node, along with a non-interactive zero-knowledge (NIZK) proof of the correctness of the encrypted shares. The dealer then publishes the transcript to all the nodes using a broadcast channel. Upon receiving the transcript over the broadcast channel, each node validates the correctness of all encrypted shares using the NIZK proof and recovers its own share by decrypting its encrypted share.

Existing VE-based schemes achieve several nice properties. First, they are non-interactive, i.e., only the dealer broadcasts a single message in the entire protocol. Second, they are also publicly verifiable. Third, the same protocol paradigm, with appropriate instantiations of the broadcast channel, works in both synchronous and asynchronous networks. However, VE-based protocols are generally inefficient, particularly due to their reliance on NIZK over encrypted data [35], [41], [42], [33], [47]. Some works [59], [22], [23] bypass this efficiency issue by weakening the VSS functionality. More precisely, these schemes require the VSS secret to be an elliptic curve group element. Hence, they are not compatible with off-the-self threshold cryptosystems whose keys are field elements [37], [15].

Bivariate polynomial-based AVSS. A more recent approach to designing AVSS is to rely on a bivariate polynomial [49], [4], [3], [66]. In these schemes, the dealer embeds its secret as the constant term of a random low-degree bivariate polynomial. The dealer then publishes a commitment to the bivariate polynomial using reliable broadcast. Additionally, the dealer privately sends partial evaluations of the polynomials to each node. Each node, upon receiving its partial evaluation, communicates with others to recover its share of the secret. Intuitively, the sharing phase terminates only when the dealer sends valid partial evaluations to a majority of the honest nodes. By sending valid partial evaluations to the majority of the honest nodes, the dealer provides these nodes with sufficient information to assist each other in recovering their shares.

Unlike complaint-based AVSS schemes, this approach guarantees asynchronous termination, i.e., a node can terminate the protocol after outputting its share. However, these approaches require the dealer to perform $O(n^2)$ group exponentiations. Moreover, these protocols require a trusted setup and strong cryptographic assumptions in the Algebraic Group Model for efficient communication. More precisely, Haven [4] and Bingo [3] assume hardness of q-SDH in a pairing-friendly group and require a powers-of-tau setup [48] to achieve $O(\kappa n^2)$ total communication. Without the setup, the state-of-the-art protocol Haven incurs $O(\kappa n^2 \log n)$ total communication cost and has $O(n^2)$ per-node computation cost. Lastly, these protocols are not publicly verifiable.

Other related works. A number of works have studied VSS protocols with information-theoretic security [29], [45], [7], [21], [20], [54], [55], [27], [40], in both synchronous and asynchronous networks. However, these have high worst-case communication costs, only guarantee security with abort, or have sub-optimal fault tolerance. A series of works [16], [6], [32] study VSS protocols without completeness, and the latest among them achieve [32] a communication cost of $O(\kappa n^2)$ assuming collision resistance hash functions and hardness of discrete logarithm.

3. Definition and Overview

3.1. Definition of Verifiable Secret Sharing

Definition 1 (Verifiable Secret Sharing)**.** A verifiable secret sharing (VSS) protocol consists of two phases: *Sharing* and *Reconstruction*. During the sharing phase, a dealer L shares a secret $s \in \mathbb{F}$. During the reconstruction phase, nodes interact to recover the secret. We say that a VSS protocol is t-resilient if the following properties hold with probability $1 - \mathsf{negl}(\kappa)$ against any probabilistic polynomial time (PPT) adversary $\mathcal{A}$ that corrupts up to t nodes:

- **Correctness.** If L is honest and has a secret s, then the sharing phase will result in all honest nodes eventually outputting a share of s. Once the sharing phase finishes, if all honest nodes start the reconstruction phase, they will output s.
- **Completeness:** If any honest node outputs in the sharing phase, then there exists a secret $\tilde{s} \in \mathbb{F}$ such that all honest nodes eventually output a share of $\tilde{s}$. Also, $\tilde{s}$ is guaranteed to be reconstructed during the reconstruction phase.
- **Secrecy.** If L is honest, there exists a PPT simulator $\mathcal{S}$ which interacts with an ideal functionality $\mathcal{F}_{\text{VSS}}$ and outputs a view of $\mathcal{A}$, such that the $\mathcal{A}$'s view in the real-world protocol and the simulated protocol are indistinguishable.
- **Termination.** All honest nodes will eventually terminate the Sharing phase.

We will define the functionality $\mathcal{F}_{\text{VSS}}$ and its variants when we analyze its Secrecy property.

VSS protocols in synchronous and asynchronous networks can tolerate up to $1/2$ and $1/3$ fractions of failures, respectively [2]. It is well known that the standard Termination property is impossible in asynchronous networks since it is impossible to tell apart a slow dealer from a malicious one. Thus, AVSS protocols instead guarantee the asynchronous termination property, similar to that of reliable broadcast [14].

- **Asynchronous termination.** If any honest node outputs in the sharing phase, then all honest nodes will eventually terminate the sharing phase.

Many applications of VSS additionally require the VSS scheme to be publicly verifiable, as defined below.

Definition 2 (Publicly verifiable)**.** A publicly verifiable secret sharing (PVSS) protocol outputs a transcript that enables any third party, not just the original nodes, to verify that the dealer has ensured each node receives its share.

Another desirable property of AVSS protocol is dual-threshold, as defined below.

Definition 3 (Dual-threshold AVSS)**.** A (n, ℓ, t) dual-threshold AVSS for $n \geq 3t+1$ is a t-resilient AVSS scheme where for any given $\ell \in [t, n-t)$, the secrecy of the secret holds against any coalition of up to ℓ nodes. We refer to ℓ as the *secrecy* threshold.

Remark. The dual-threshold guarantees achieved by some VSS and DKG protocols [4], [49], [3], [31] are weaker than Definition 3. Those schemes achieve a secrecy threshold of $\ell > t$ only after the protocol terminates. During the protocol execution, their secrecy threshold is t. In contrast, Definition 3 require a secrecy threshold of ℓ even during the protocol execution.

3.2. Overview of Our Approach

Our starting point is the classical complaint-based synchronous VSS schemes described in §2. In those schemes, nodes publish complaints if they receive an invalid share or no share from the dealer. The dealer responds to complaints by publishing the shares of the complaining nodes. If the dealer fails to do so, it is considered malicious, and nodes output default values. This approach prevents a malicious dealer from violating completeness while still ensuring secrecy. This is because honest nodes will not complain against an honest dealer, thereby safeguarding the shares of honest nodes. Moreover, when the dealer is malicious, secrecy is vacuous.

Note from §2 that the conflict is always between achieving completeness and ensuring secrecy. Without secrecy, achieving completeness is trivial: the dealer simply broadcasts shares of everyone (or even the secret) to all. With this in mind, let us take another look at the complaint-based schemes. Here, the dealer reveals shares of a subset of parties, and the protocol ensures that an honest dealer only reveals shares of malicious nodes. Our new paradigm achieves a similar property but uses a different approach, as we describe next.

The first crucial change we introduce is that, instead of sending explicit complaints, we send explicit acknowledgments instead. The absence of an acknowledgment is in some way a complaint. Specifically, the dealer computes the shares of its secret using a low-degree polynomial, along with a commitment to the polynomial. The dealer, instead of publishing the commitment, first privately sends each node i the commitment along with the share of node i. Each node, upon receiving its share of the secret, validates it for correctness. Upon successful validation, the node responds to the dealer with a signed acknowledgment. This acknowledgment can serve as proof that node i has received its valid share corresponding to the commitment.

The dealer waits to receive an appropriate number of signed acknowledgments. (The dealer cannot wait for acknowledgments from all nodes because malicious nodes may never send acknowledgments.) Next is where our second crucial change comes in. The dealer then publishes, using a broadcast channel, the VSS transcript, which consists of the commitment to the polynomial, the signed acknowledgments it has received, and the shares of nodes who did not respond with a signed acknowledgment. Looking ahead, we will argue that despite the dealer publicly revealing shares of a subset of nodes, $\mathcal{A}$ does not learn enough points on an honest dealer's polynomial, so secrecy is maintained.

Upon receiving the transcript over the broadcast channel, nodes validate it by checking that, for each node $i \in [n]$, either its signature or its share of the secret is included in the transcript. Upon successful validation, each node outputs the commitment and its share and terminates the sharing phase. If the validation fails, a node outputs a default value. Intuitively, completeness is satisfied because a node either explicitly acknowledges receiving its share or will receive its share from the validated transcript.

It is easy to see that the transcript the dealer broadcasts is publicly verifiable. The public verification check of the transcript is precisely the verification check each node performs on the transcript before terminating the sharing phase.

Based on these insights, designing a synchronous VSS protocol is straightforward. In a synchronous network of $n = 2t+1$ nodes, with pair-wise latency Δ, the dealer shares its secret using a degree t polynomial. The dealer then waits for 2Δ time units to receive signed acknowledgments from all honest nodes and reveal the remaining shares using a broadcast channel.

However, this approach fails in asynchronous networks with $n = 3t + 1$. Under asynchrony, the dealer needs to make progress upon receiving $n - t = 2t + 1$ signed acknowledgments. Note that t of these $2t+1$ acknowledgments could be from malicious parties. Now, if the dealer reveals the remaining t honest shares, it would reveal a total of $2t$ shares to $\mathcal{A}$, which is sufficient to reconstruct the degree t polynomial the dealer uses to share its secret. We address this issue by requiring the dealer to share its secret using a degree $2t$ polynomial. This prevents $\mathcal{A}$ from learning the secret even after learning $2t$ shares.

Finally, to construct a dual-threshold AVSS with secrecy threshold ℓ for $\ell \in [t, n-t)$, we combine ideas from verifiable encryption-based VSS with our low-threshold AVSS, i.e., AVSS with $\ell = t$. More precisely, for any ℓ, the dealer still uses a degree $2t$ polynomial to share its secret, but crucially does not reveal all remaining t shares after receiving $2t+1$ signed acknowledgments. Instead, the dealer publicly reveals only $2t - \ell$ of the remaining t shares, encrypts, and broadcasts the remaining $t - (2t - \ell) = \ell - t$ shares using a verifiable encryption scheme. Intuitively, this ensures that any coalition of at most ℓ nodes learns at most $2t$ points on the polynomial. The protocol still ensures completeness because the nodes whose shares are not revealed by the dealer will receive their share from the verifiable encryptions revealed by the dealer. Since the leader broadcasts $\ell - t$ shares using verifiable encryption, the performance degrades gradually with ℓ. And if the leader receives more than $2t + 1$ signed acknowledgments (e.g., in the best case with a synchronous network and few malicious parties), the performance will further improve.

4. Threat Model and Preliminaries

Let $\mathbb{G}$ be an elliptic curve group of order q with $\mathbb{F}$ as its scalar field. Let $g, h \in \mathbb{G}$ be two uniformly random and independent generators. We use κ to denote the security parameter. For example, when we use a signature scheme, κ denotes the size of the secret key. Similarly, we also use κ to denote the size of an element in $\mathbb{F}$ or $\mathbb{G}$. For any integer a, we use $[a]$ to denote the ordered set $\{1, 2, \ldots, a\}$. Also, for two integers a and b where $a < b$, we use $[a, b]$ to denote the ordered set $\{a, a+1, \ldots, b\}$.

4.1. Threat Model

We consider a network of n nodes denoted by $\{1, 2, \ldots, n\}$, where each node is connected with the dealer via a pairwise private and authenticated channel. We assume nodes have access to a broadcast channel with which the dealer can send a value to all nodes. A broadcast channel ensures that the dealer cannot send inconsistent values to different nodes. We can efficiently realize such optimal fault-tolerant broadcast channels in synchronous and asynchronous networks by running a Byzantine broadcast [50], [53] and reliable broadcast [14], [32], respectively. We will give their interfaces in Appendix A.

We consider a *static* adversary $\mathcal{A}$ that can corrupt a threshold fraction of the nodes upfront. For our synchronous VSS protocol, we assume that $\mathcal{A}$ can corrupt less than half of the nodes, i.e., at most t out of $n \geq 2t + 1$ nodes. Also, let Δ be the upper bound on the delay between the honest dealer and any honest node. For our AVSS and dual-threshold AVSS protocols, we assume that for $n \geq 3t + 1$, at most t nodes are malicious. Each node i has its private signing key sk_i and the corresponding public verification key pk_i. We also assume a public key infrastructure (PKI), i.e., all nodes have access to $\{\mathsf{pk}_j\}_{j \in [n]}$.

4.2. Threshold Secret Sharing

A $(n, d+1)$ threshold secret sharing scheme allows a secret $s \in \mathbb{F}$ to be shared into n shares such that any set of $d + 1$ shares are sufficient to recover the original secret, but any set of d shares give no information about the original secret [60], [11]. We use the common Shamir secret sharing [60] scheme, where the secret is embedded in a random degree d polynomial in the field $\mathbb{F}$. Specifically, to share a secret $s \in \mathbb{F}$, a polynomial $p(\cdot)$ of degree d is chosen such that $s = p(0)$ and other coefficients of $p(\cdot)$, $a_1, a_2, \cdots, a_t$ are chosen uniformly randomly from $\mathbb{F}$:

$$p(x) = s + a_1 x + a_2 x^2 + \cdots + a_d x^d$$

The i-th share of the secret is then $p(i)$, i.e., the polynomial evaluated at i. Given $d+1$ points on the polynomial $p(\cdot)$, one can efficiently reconstruct the polynomial using Lagrange interpolation. Also note that s is information-theoretically hidden from an adversary that knows d or fewer evaluation points on the polynomial other than $p(0)$ [60].

4.3. Polynomial Commitment Scheme

The dealer in our VSS scheme commits to its secret by committing to a degree d polynomial $p(\cdot)$. Let $\mathsf{PC} = (\mathsf{PC.Setup}, \mathsf{PC.Commit}, \mathsf{PC.DegCheck}, \mathsf{PC.Open}, \mathsf{PC.Verify})$ be a polynomial commitment scheme.

- $\mathsf{PC.Setup}(1^\kappa) \rightarrow pp$. On input the security parameter κ, outputs the public parameters for the polynomial commitment scheme.
- $\mathsf{PC.Commit}(pp, p(\cdot), n) \rightarrow (\boldsymbol{v}, \boldsymbol{w})$. On input the public parameters pp, number of evaluations n, and the polynomial $p(\cdot)$, outputs the commitment $\boldsymbol{v}$ of the polynomial $p(\cdot)$ and witness $\boldsymbol{w}$.
- $\mathsf{PC.Open}(pp, \boldsymbol{w}, p(\cdot), i) \rightarrow (p(i), \pi)$. On input the index i and the polynomial $p(\cdot)$, outputs $p(i)$, and a valid opening proof π.
- $\mathsf{PC.DegCheck}(pp, \boldsymbol{v}, d) \rightarrow 0/1$. On input the polynomial commitment $\boldsymbol{v}$ and a degree d, outputs 1 if $\boldsymbol{v}$ is a commitment to a polynomial of degree at most d, and outputs 0 otherwise.
- $\mathsf{PC.Verify}(pp, \boldsymbol{v}, i, u, \pi) \rightarrow 0/1$. On input the polynomial commitment $\boldsymbol{v}$ to a polynomial $p(i)$, outputs 1 if $u = p(i)$, and outputs 0 otherwise.

Batch interfaces. As we briefly describe in §3.2, the dealer in our VSS protocols provides opening proofs for a batch of indices and each node verifying them locally. Thus, we use the batched interfaces PC.BatchOpen and PC.BatchVerify for better exposition. Briefly, PC.BatchOpen takes a set I of indices along with $(\boldsymbol{v}, \boldsymbol{w})$ and outputs $(\boldsymbol{s}, \boldsymbol{\pi})$. Here $\boldsymbol{s}$ is the vector of openings for each index in I, and $\boldsymbol{\pi}$ consists of corresponding opening proofs. Similarly, PC.BatchVerify takes as input a set I of indices along with $(\boldsymbol{s}, \boldsymbol{\pi})$, and outputs 1 if all the opening proofs are valid. We formally define these interfaces in Appendix A.2 and present mechanisms to verify a batch of polynomial evaluations more efficiently than verifying each evaluation independently.

A polynomial commitment scheme PC is secure if it satisfies the *Completeness, Evaluation binding*, and *Hiding* [48]. Intuitively, the Completeness property ensures that verification of honestly generated commitments and opening proofs are always successful. The Evaluation binding property prevents $\mathcal{A}$ from successfully opening to two different values at the same index. Lastly, the Hiding property guarantees that the commitment $\boldsymbol{v}$ reveals no information about the polynomial.

Constructions. Figure 1 gives a concrete polynomial commitment scheme. The described scheme combines ideas from the classic Pedersen's polynomial commitment and SCRAPE's low-degree test [22]. The resulting scheme has a linear-sized commitment and constant-sized opening proof. The commitment includes n values of the polynomials in the exponent, and the low degree is verified by multiplying these values in the exponent with a random word from the dual code and checking that the result is $1_\mathbb{G}$, i.e., the identity element of $\mathbb{G}$. The scheme is information-theoretically hiding and evaluation binding assuming hardness of discrete logarithm [56].

$\mathsf{PC.Setup}(1^\lambda)$: Output $pp = (\mathbb{G}, \mathbb{F}, g, h)$, for an elliptic curve group $\mathbb{G}$ with scalar field $\mathbb{F}$, and uniformly random and independent generators $g, h \in \mathbb{G}$.

$\mathsf{PC.Commit}(pp, p(\cdot), d, n)$: Let $p(\cdot)$ be the polynomial of degree d. Sample a random polynomial $r(\cdot)$ of degree d. Let $\boldsymbol{v}$ be the commitment to $p(\cdot)$ where

$$\boldsymbol{v} = \left[g^{p(1)}h^{r(1)}, g^{p(2)}h^{r(2)}, \ldots, g^{p(n)}h^{r(n)}\right]$$

Output $(\boldsymbol{v}, \boldsymbol{w}) = (\boldsymbol{v}, r(\cdot))$.

$\mathsf{PC.Open}(pp, i, p(\cdot), r(\cdot))$: Output $(u, \pi) = (p(i), r(i))$.

$\mathsf{PC.DegCheck}(pp, \boldsymbol{v}, d)$: Sample a random degree $d = n - t - 2$ polynomial $z(\cdot)$ in $\mathbb{F}$. Output 1 if

$$\prod_{i\in[n]} \boldsymbol{v}[i]^{z(i)\cdot\lambda_i} = 1_\mathbb{G} \tag{1}$$

for $\lambda_i = \prod_{j\in[n],j\neq i} 1/(i-j)$; otherwise output 0.

$\mathsf{PC.Verify}(pp, \boldsymbol{v}, i, u, \pi)$: Output 1 if $\boldsymbol{v}[i] = g^z h^\pi$; otherwise output 0.

Figure 1: Pedersen's polynomial commitment scheme combined with SCRAPE's low degree test.

Remark. An alternative approach would have been to commit to $d+1$ coefficients of the polynomials in the exponent. This would have made the commitment shorter and would have eliminated the need for low-degree verification. On the other hand, the opening phase would have become more costly: verifying each opened value would have required $O(d)$ exponentiations instead of one.

5. Synchronous VSS

Our synchronous VSS protocol is given in Algorithm 1. We assume $n = 2t + 1$. The CRS $(\mathbb{G}, \mathbb{F}, g, h)$ is the output of the polynomial commitment's setup phase $\mathsf{PC.Setup}(1^\kappa)$. Let Δ be the upper bound on the delay between the honest dealer and any honest node. For any node $i \in [n]$, let $\mathsf{sk}_i, \mathsf{pk}_i$ be its private signing key and public verification key.

5.1. Design

Sharing phase. Let $m \in \mathbb{F}$ be the message the dealer L wants to share. L samples a degree-t polynomial

$$s(x) = m + s_1 x + s_2 x^2 + \cdots + s_t x^t \tag{2}$$

with uniformly random $s_i \in \mathbb{F}$ for each $i \in [n]$. L then computes the commitment of $s(\cdot)$ along with the commitment witness as $\boldsymbol{v}, \boldsymbol{w} \leftarrow \mathsf{PC.Commit}(s(\cdot), n)$.

At time $\tau = 0$, L computes the opening proof $\pi_i = \mathsf{PC.Open}(s(\cdot), i, \boldsymbol{w})$ for each $i \in [n]$ and sends the tuple $\langle \mathtt{SHARE}, \boldsymbol{v}, s(i), \pi_i\rangle$ to node i. Node i, upon receiving the SHARE message from L, validates that $\boldsymbol{v}$ is a polynomial of degree t by checking that $\mathsf{PC.DegCheck}(\boldsymbol{v}, t) = 1$, and checks that its share is valid using $\mathsf{PC.Verify}(\boldsymbol{v}, i, s(i), \pi_i)$.

Algorithm 1 Synchronous VSS

PUBLIC PARAMETERS: $n \geq 2t + 1$, $\{\mathsf{pk}_i\}_{i\in[n]}$, maximum network latency Δ, and public parameters $(\mathbb{G}, \mathbb{F}, g, h)$ of the polynomial commitment scheme.

PRIVATE INPUT: Signing key sk_i.

SHARING PHASE:

// Dealer L at time $\tau = 0$ and with input m:
Sample a t-degree random polynomial $s(\cdot)$ with $s(0) = m$
$\boldsymbol{v}, \boldsymbol{w} \leftarrow \mathsf{PC.Commit}(s(\cdot), n)$
for $i = 1, 2, ..., n$ **do**
  Let $\pi_i = \mathsf{PC.Open}(s(\cdot), i, \boldsymbol{w})$
  send $\langle \mathtt{SHARE}, \boldsymbol{v}, s(i), \pi_i\rangle$ to node i

// Each node i
upon receiving $\langle \mathtt{SHARE}, \boldsymbol{v}, s(i), \pi_i\rangle$ from dealer L **do**
  Check $\mathsf{PC.DegCheck}(\boldsymbol{v}, t) = 1$
  Check $\mathsf{PC.Verify}(\boldsymbol{v}, i, s(i), \pi_i) = 1$
  if both checks are succesful **then**
    Let $\sigma_i = \mathsf{sign}(\mathsf{sk}_i, \boldsymbol{v})$
    send $\langle \mathtt{ACK}, \sigma_i\rangle$ to L

// Dealer L at time $\tau = 2\Delta$
Let $\boldsymbol{\sigma}$ be the set of received valid signatures on $\boldsymbol{v}$.
Let I be the indices of nodes with *missing* signatures.
Let $\boldsymbol{s}, \boldsymbol{\pi} = \mathsf{PC.BatchOpen}(p(\cdot), I, \boldsymbol{w})$.
send $(\boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi})$ using the broadcast channel.

// Each node i once the broadcast outputs $(\boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi})$.
Check if each $\sigma \in \boldsymbol{\sigma}$ is valid and $|\boldsymbol{\sigma}| \geq t + 1$.
Check if $\mathsf{PC.BatchVerify}(\boldsymbol{v}, I, \boldsymbol{s}, \boldsymbol{\pi})$.
Check that I includes all nodes with missing signatures.
if all the checks pass **then**
  Output $(\boldsymbol{v}, s(i), \pi_i)$; **return**
else
  Output 0 as the default share.

RECONSTRUCTION PHASE:

// every node i after finishing the sharing phase
send $\langle \mathtt{RECON}, s(i), \pi_i\rangle$ to all.
upon receiving $\langle \mathtt{RECON}, s(j), \pi_j\rangle$ from node j **do**
  if $\mathsf{PC.Verify}(\boldsymbol{v}, s(j), \pi_j)$ **then**
    $T = T \cup \{s_j\}$
    if $|T| \geq t + 1$ **then**
      output $s(0)$ using Lagrange interpolation; **return**

If both these checks are successful, node i sends a message $\langle \mathtt{ACK}, \sigma_i\rangle$ where σ_i is its signature on $\boldsymbol{v}$.

L waits for 2Δ units of time to collect ACK messages. Here, for ease of exposition, we assume Δ has accounted for the time required to validate $\boldsymbol{v}$ and check the validity of a share. At time $\tau = 2\Delta$, let $\boldsymbol{\sigma}$ be the set of valid signatures L receives and let I be the set of nodes from whom L does not receive valid signatures. L then computes $(\boldsymbol{s}, \boldsymbol{\pi}) = \mathsf{PC.BatchOpen}(s(\cdot), I, \boldsymbol{w})$ where $\boldsymbol{s}$ is of size $|I|$ and consists of $s(k)$ for each $k \in I$, and $\boldsymbol{\pi}$ is the opening proof. Then, L sends the message $\langle \boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi}\rangle$ using the broadcast channel.

When the broadcast channel outputs $\langle \boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi}\rangle$, each node locally checks that: (i) $\boldsymbol{\sigma}$ is a valid set of signatures on $\boldsymbol{v}$ and $|\boldsymbol{\sigma}| \geq t + 1$; (ii) I includes all nodes

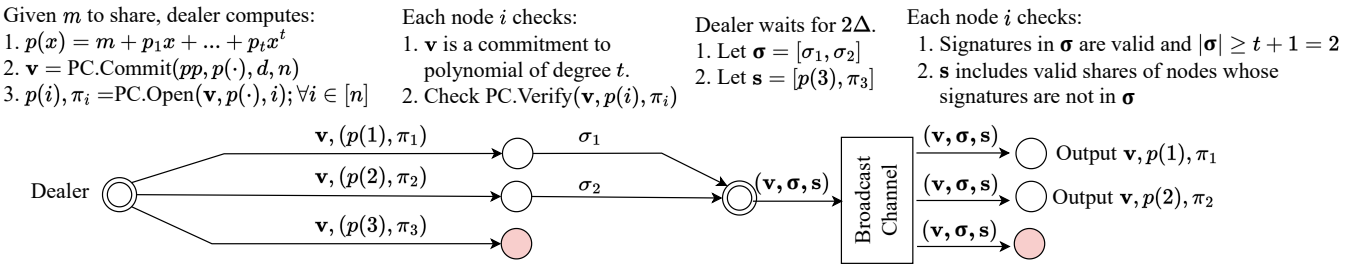

Figure 2: Our synchronous VSS protocol involves three nodes, one of which is malicious (shaded red in the diagram).

whose signatures are not included in $\boldsymbol{\sigma}$; and (iii) $\boldsymbol{s}$ includes valid shares of nodes in I with respect to $\boldsymbol{v}$, i.e., PC.BatchVerify$(\boldsymbol{v}, I, \boldsymbol{s}, \boldsymbol{\pi})$. If all these checks are successful, node i outputs its share $s(i)$, the commitment $\boldsymbol{v}$, and the opening proof π_i. A node gets these from either the broadcast message or the SHARE message it received from the dealer.

Using multisignatures. One simple concrete optimization is to have each node sign its ACK message using a multisignature scheme. More precisely, the ACK message from node i includes its partial signature on $\boldsymbol{v}$. L then broadcast the multisignature σ on $\boldsymbol{v}$ instead of broadcasting a list of signatures. Concretely, in our experiments, we use the BLS multisignature scheme from [12].

Reconstruction phase. Let T be a set of $t + 1$ nodes (including itself) from which node i receives valid shares $s(j)$. Upon receiving $t + 1$ such valid shares, node i computes the secret m using Lagrange interpolation as $m = \sum_{k \in T} \mu_k s(k)$, where $\mu_k = \prod_{j \neq k} \frac{j}{j-k}$ are the Lagrange coefficients.

Optimized reconstruction. In certain situations, it is possible to optimize the reconstruction phase. A node may not need to always wait for $t + 1$ RECON messages. If, during the sharing phase, the dealer has already revealed k shares as part of $\boldsymbol{s}$, a node only needs to wait for $t + 1 - k$ RECON messages for shares not included in $\boldsymbol{s}$.

5.2. Analysis

Correctness. An honest dealer, L, will receive signed ACK messages from all honest nodes within 2Δ time due to the synchrony assumption. Since there are at least $t + 1$ honest nodes, $|\boldsymbol{\sigma}| \geq t+1$. Let $\langle \boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi} \rangle$ be the transcript broadcast by L. Then, by the Validity property of Byzantine agreement, each honest node will output $\langle \boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi} \rangle$. Then, by the Correctness property of the signature scheme and the Completeness property of the polynomial commitment scheme, every honest node will accept the VSS transcript and output its share.

Finally, during the reconstruction protocol, each honest node will multicast a valid RECON message. Thus, every honest node will receive at least $t + 1$ valid shares, which is sufficient to reconstruct the degree t polynomial $s(\cdot)$, and hence $s(0)$. Moreover, the Soundness of the polynomial commitment ensures that honest nodes only accept valid shares on the committed polynomial. This implies that all honest nodes output the same unique secret $s(0)$.

Termination. Follows directly from the Termination property of the Byzantine broadcast scheme (cf. Definition 5).

Completeness. An honest party outputs its share only upon receiving a valid transcript $(\boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi})$ over the Byzantine broadcast channel. The Agreement property of the Byzantine broadcast guarantees that every honest node outputs the same transcript, and hence the same polynomial commitment. Successful validation of the transcript implies that at least $t + 1$ node, hence at least one honest node, signed the commitment. This implies with $1 - \mathsf{negl}(\kappa)$ probability, $\boldsymbol{v}$ is a commitment to a polynomial of degree at most t. Also, for each node $i \in [n]$, either a signature of i or its valid share is included in $\boldsymbol{\sigma}$. In the former case, assuming the existential unforgeability of the signature scheme, node i already received its share. In the latter case, node i will receive its valid share from $\boldsymbol{s}$.

During the reconstruction phase, each honest node will reconstruct the degree t polynomial $p(\cdot)$ corresponding to the commitment $\boldsymbol{v}$. Hence, each node will output the unique secret $s(0)$.

Secrecy. We prove Secrecy using *simulatability*: for every probabilistic polynomial-time (PPT) adversary $\mathcal{A}$ that corrupts up to t nodes, there exists an ideal world PPT simulator $\mathcal{S}_{\text{VSS}}$ that interacts with the ideal functionality $\mathcal{F}_{\text{VSS}}$ (cf. Figure 4) and produces a view such that $\mathcal{A}$'s view in the simulated world is identical to a run of the Sharing phase. We formally prove Secrecy in Appendix C.

Performance. We will analyze our performance using Figure 1 as the polynomial commitment scheme. The dealer performs $O(n \log n)$ field operations to compute shares of each node (using FFT). The dealer then performs $O(n)$ group exponentiations to compute the commitments and $O(n)$ signature verifications. Since group exponentiation is more expensive than $\log n$ field operations, we treat the dealer's computation cost as $O(n)$ group exponentiations. The running time of each node is as follows. Each node performs $O(n)$ group exponentiations to verify the polynomial commitment, signatures of $O(n)$ nodes, and shares of $O(n)$ nodes. Finally, the dealer privately sends an $O(\kappa n)$-bit commitment to each node and broadcasts an $O(\kappa n)$-bit transcript. Hence, the total communication cost of our VSS protocol is $O(\kappa n^2 + C_{BB}(\kappa n))$ where $C_{BB}(a)$ is the communication cost of broadcasting a message of length a.

Combining all the above, we get the following theorem.

Algorithm 2 Asynchronous VSS

PUBLIC PARAMETERS: $n \geq 3t+1$, $\{\mathsf{pk}_i\}_{i\in[n]}$, and public parameters of the polynomial commitment scheme pp.
PRIVATE INPUT: Signing key sk_i.

SHARING PHASE:

// Dealer L with input m:
Sample a $2t$-degree random polynomial $s(\cdot)$ with $s(0) = m$
$\boldsymbol{v}, \boldsymbol{w} \leftarrow \mathsf{PC.Commit}(s(\cdot), n)$
for $i = 1, 2, ..., n$ **do**
  Let $\pi_i \leftarrow \mathsf{PC.Open}(s(\cdot), i, \boldsymbol{w})$
  send $\langle \mathtt{SHARE}, \boldsymbol{v}, s(i), \pi_i \rangle$ to node i

// Each node i
upon receiving $\langle \mathtt{SHARE}, \boldsymbol{v}, s(i), \pi_i \rangle$ from dealer L **do**
  Check $\mathsf{PC.DegCheck}(\boldsymbol{v}, 2t) = 1$
  Check $\mathsf{PC.Verify}(\boldsymbol{v}, i, s(i), \pi_i) = 1$
  if both the checks pass **then**
    Let $\sigma_i = \mathsf{sign}(\mathsf{sk}_i, \boldsymbol{v})$
    send $\langle \mathtt{ACK}, \sigma_i \rangle$ to L

// Dealer L waits for $2t+1$ valid signatures on $\boldsymbol{v}$
Let $\boldsymbol{\sigma}$ be the set of valid signatures on $\boldsymbol{v}$.
Let I be the indices of nodes with missing signatures.
Let $\boldsymbol{s}, \boldsymbol{\pi} = \mathsf{PC.BatchOpen}(p(\cdot), I, \boldsymbol{w})$.
send $(\boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi})$ using a reliable broadcast channel.

// Each node i once the broadcast outputs $(\boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi})$.
Check if each $\sigma \in \boldsymbol{\sigma}$ is valid and $|\boldsymbol{\sigma}| \geq 2t+1$.
Check that I includes all nodes with missing signatures.
Check if $\mathsf{PC.BatchVerify}(\boldsymbol{v}, I, \boldsymbol{s}, \boldsymbol{\pi})$.
if all the checks pass **then**
  Output $(\boldsymbol{v}, s(i), \pi_i)$; **return**

RECONSTRUCTION PHASE:

// every node i after finishing the sharing phase
send $\langle \mathtt{RECON}, s(i), \pi_i \rangle$ to all.
upon receiving $\langle \mathtt{RECON}, s(j), \pi_j \rangle$ from node j **do**
  if $\mathsf{PC.Verify}(\boldsymbol{v}, s(j), \pi_j)$ **then**
    $T = T \cup \{s_j\}$
    if $|T| \geq 2t+1$ **then**
      output $s(0)$ using Lagrange interpolation; **return**

Theorem 1 (Synchronous VSS)**.** *In a synchronous network of $n \geq 2t+1$ nodes among which at most t nodes are malicious, assuming a polynomial commitment scheme, a signature scheme, and a Byzantine broadcast channel, Algorithm 1 implements a t-resilient publicly verifiable VSS protocol with $O(\kappa n^2 + C_{BB}(\kappa n))$ communication cost. Here κ is the security parameter, and $C_{BB}(a)$ is the communication cost of broadcasting a message of length a using the broadcast channel.*

6. Asynchronous VSS

In this section, we will describe the modifications to make our protocol in an asynchronous or partial synchronous network. As we mention in §3, we seek to design an AVSS with the Completeness property. Since AVSS with completeness implies an asynchronous RBC, $n/3$ is the maximum number of failures any AVSS protocol can tolerate [14]. Throughout this section, we will assume $n = 3t+1$. We summarize our protocol in Algorithm 2, where we highlight in gray the changes on top of Algorithm 1.

Protocol intuition. The natural attempt to adapt the synchronous VSS in an asynchronous network of $n = 3t+1$ is to let the dealer share its secret using a degree t polynomial and keep the rest of the protocol as is. However, as we briefly mention in §3.2, this approach will not work. In asynchrony, there is no fixed upper bound on the message delays, so the dealer cannot wait to receive acknowledgments from all honest nodes. Instead, the dealer must move on upon receiving only $n - t = 2t + 1$ signed acknowledgments. But t of these $n - t$ signed acknowledgments could be from malicious nodes, and the missing t acknowledgments correspond to honest but slow nodes. In this case, an honest dealer would reveal to $\mathcal{A}$ a total of $2t$ shares on a degree t polynomial, which is sufficient for $\mathcal{A}$ to recover the secret.

We address this issue with the following key observation. We let the dealer share the secret using a degree $2t$ polynomial (instead of degree t). The rest of the protocol follows a similar structure, with the following natural changes. The dealer waits for $n-t$ valid signed acknowledgments instead of a pre-specified time bound, and publishes the t shares from the t slow nodes. Intuitively, by using a degree $2t$ polynomial, we ensure that $\mathcal{A}$ does not learn the secret even after learning $2t$ shares on it. Finally, using a degree $2t$ degree polynomial does not affect the reconstructability of the secret as $n-t \geq 2t$, i.e., there are enough honest nodes to reconstruct the secret.

We want to note that although the dealer in Algorithm 2 shares its secret using a degree $2t$ polynomial, the protocol is not dual-threshold. This is because $\mathcal{A}$ learns up to $2t$ points on the polynomial by corrupting only t nodes.

6.1. Design

Sharing phase. During the sharing phase, the dealer L embeds the secret m in a polynomial of degree $2t$ instead of degree t. Let $s(\cdot)$ be the polynomial with $s(0) = m$. Similarly to Algorithm 1, the dealer then computes the commitment $\boldsymbol{v}$ along with witness $\boldsymbol{w}$ and an opening proof π_i for each node $i \in [n]$, and sends $\langle \mathtt{SHARE}, \boldsymbol{v}, \pi_i \rangle$ to each node i.

Each node, upon receiving the SHARE message from L, checks that $\boldsymbol{v}$ is a commitment to a polynomial of degree $2t$, instead of degree t. Also, instead of waiting for 2Δ time to receive ACK messages, L waits until it receives $2t+1$ valid ACK messages. L then broadcasts $\langle \boldsymbol{v}, I, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi} \rangle$ using an asynchronous reliable broadcast (RBC) channel. Here, I consists of the indices of nodes whose signatures are missing, $\boldsymbol{s}$ consists of the shares of nodes in I, and $\boldsymbol{\pi}$ consists of batch opening proofs of shares in $\boldsymbol{s}$.

Once the RBC outputs, a node accepts the transcript only if: (i) $\boldsymbol{\sigma}$ consists of at least $2t+1$ valid signatures, (ii) $\boldsymbol{v}$ is a commitment to a polynomial of degree $2t$, (iii) $\boldsymbol{s}$ includes

valid shares of all nodes with missing signatures. Finally, upon successful validation, each node i outputs $\boldsymbol{v}, s(i), \pi_i$.

Reconstruction phase. The only change from the synchronous scheme is that each node waits for $2t+1$ valid RECON messages because the secret is shared using a degree $2t$ polynomial. Similar to the synchronous scheme, nodes can utilize the shares revealed as a part of $\boldsymbol{s}$ to speed up the reconstruction phase.

Reducing the storage costs. In the AVSS scheme in Algorithm 2, each node stores the entire $\boldsymbol{v}$, which is $O(\kappa n)$ for Pedersen polynomial commitment. We can reduce the storage cost to $O(\kappa)$, using error-correcting code [57] and online error correction [21], similar to AVSS protocols such as [64], [32]. More specifically, each node encodes $\boldsymbol{v}$ using a $[n, t, n-t]$ Reed-Solomon code. Let $\hat{\boldsymbol{v}}$ be the encoded commitment. Each node i then stores $\hat{\boldsymbol{v}}[i]$ and deletes the rest of $\hat{\boldsymbol{v}}$. During the reconstruction phase, each node i sends $\langle \texttt{RECON}, \hat{\boldsymbol{v}}[i], s(i), \pi_i \rangle$ to all. Upon receiving RECON messages, nodes first recover $\boldsymbol{v}$ using online error correction. Then, nodes use the reconstruction protocol described in §6.1 to recover the polynomial.

6.2. Analysis

Correctness. Since $n - t \geq 2t+1$, an honest dealer L will eventually receive $2t+1$ signed acknowledgments. Then, using a similar argument as our synchronous VSS, each honest node will eventually output and accept the transcript broadcast by the honest dealer. Similarly, during the reconstruction phase, each node will eventually receive $2t+1$ valid shares, which is sufficient to reconstruct the degree $2t$ polynomial $p(\cdot)$, and hence $s(0)$. Also, honest nodes will accept only valid shares and hence will output the same unique secret shared by the dealer.

Asynchronous Termination. Follows directly from the Totality property of the Byzantine RBC (cf. Definition 6).

Completeness. Follows using a similar argument as the synchronous VSS protocol.

Secrecy. We will prove the Secrecy in Appendix C.

Performance. The computation cost of the dealer and nodes are similar to that of the synchronous VSS protocol, except the dealer uses a degree $2t$ degree polynomial to share its secret. Precisely, both the dealer and nodes need to perform $O(n)$ group exponentiations. In terms of the bandwidth cost, the dealer sends $O(\kappa n)$ length private message to each node and $O(\kappa n)$ bit long message using a broadcast channel. Thus, using the broadcast channel from [32], the total communication cost is $O(\kappa n^2)$.

Combining all the above, we get the following theorem.

Theorem 2 (Asynchronous VSS)**.** *In an asynchronous network of $n \geq 3t+1$ nodes among which at most t nodes are malicious, assuming a polynomial commitment scheme, a signature scheme, and a Byzantine reliable broadcast channel, Algorithm 2 implements a t-resilient publicly verifiable asynchronous VSS protocol with $O(\kappa n^2)$ communication costs. Here κ is the security parameter.*

7. Dual-threshold AVSS

In this section, we use our paradigm to design an (n, ℓ, t) dual-threshold AVSS scheme.

Protocol intuition. For any given ℓ, the dealer in our dual-threshold AVSS shares its secret using a degree $2t$ polynomial and follows the AVSS protocol until it receives $2t+1$ signed acknowledgments. Then, unlike the AVSS scheme, the dealer does not reveal all remaining t shares. Instead, the dealer publicly reveals only $2t-\ell$ of the remaining t shares and shares the remaining $t-(2t-\ell) = \ell - t$ shares using a verifiable encryption scheme. More precisely, for each of the remaining $t-(2t-\ell) = \ell - t$ shares, the dealer encrypts it with the public key of the corresponding recipient node and computes a NIZK proof of its correctness. Intuitively, by publicly revealing only $2t-\ell$ shares, we ensure that any coalition of ℓ nodes learns at most $2t$ points on the polynomial. The protocol still ensures Completeness, as the nodes whose shares are not revealed by the dealer will receive their share from the verifiable encryptions.

7.1. Verifiable Encryption of Committed Messages

Our dual-threshold AVSS scheme relies on verifiable encryptions for the Pedersen commitment scheme, as defined below.

Definition 4 (Verifiable Encryption of a Committed Message)**.** Verifiable encryption (VE) of a committed message involves three parties: a prover $\mathcal{P}$, a verifier $\mathcal{V}$, and a receiver $\mathcal{R}$. The receiver $\mathcal{R}$ has a public-private key pair $(\mathsf{pk}, \mathsf{sk})$. Let Cm be a commitment scheme. Given (v, c, pk), $\mathcal{P}$ wants to convince $\mathcal{V}$ that c is a public key encryption of a message s under public key pk, and that v is an commitment to s and $\mathcal{P}$ knows s. A verifiable encryption scheme provides the following interfaces.

- $\mathsf{VE.Setup}(1^\kappa, \mathsf{Cm}) \to pp_{\mathrm{VE}}$. On input the security parameter κ, and the commitment scheme Cm, the algorithm outputs the public parameters pp_{VE}.
- $\mathsf{VE.KeyGen}(pp_{\mathrm{VE}}) \to (\mathsf{pk}, \mathsf{sk})$. The algorithm outputs a public-private key pair for the encryption scheme.
- $\mathsf{VE.EncProve}(pp_{\mathrm{VE}}, \mathsf{pk}, s, v, w) \to (c, \pi_{\mathrm{VE}})$: The algorithm takes as input the message s, commitment v with witness w, where $v, w \leftarrow \mathsf{Cm.Commit}(s)$. It outputs an encryption c of the tuple $(s, \pi = \mathsf{Cm.Open}(v, s, w))$ along with a NIZK proof π_{VE} of their correct encryptions.
- $\mathsf{VE.Verify}(pp_{\mathrm{VE}}, \mathsf{pk}, v, c, \pi_{\mathrm{VE}}) \to 0/1$. The algorithm outputs 1, if π_{VE} is a valid proof that there exists α, π such that $\alpha, \pi = \mathsf{VE.Dec}(\mathsf{sk}, c)$ and $\mathsf{Cm.Verify}(v, \alpha, \pi) = 1$. Note that π_{VE} needs to be verifiable without access to the secret key or the underlying message α.
- $\mathsf{VE.Dec}(\mathsf{sk}, c) \to s, \pi$: Given the ciphertext c and a secret key sk, the algorithm outputs a decryption of c using sk.

A verifiable encryption scheme is secure if it satisfies the standard *Completeness, Soundness*, and *Zero-knowledge* properties of verifiable computation schemes [39]. Intuitively, the Completeness property ensures that verification

of an honestly generated π_{VE} is always successful, even if a malicious node generates the public key. The Soundness property prevents a malicious prover from convincing an honest node about the correctness of an incorrectly generated ciphertext. Stating differently, if VE.Verify is successful for a ciphertext c and public key pk, then a node with secret key sk will always be able to recover its share and the opening proof. Lastly, the Zero-knowledge property guarantees that the ciphertext c and the proof π_{VE} reveal no information about the share other than whatever is revealed by the polynomial commitment scheme.

Batch verifiable encryptions. Looking ahead, the dealer in our dual-threshold VSS computes the verifiable encryptions for a batch of shares. Thus, we define the VE scheme to additionally support batched interfaces VE.BatchEncProve and VE.BatchVerify. Trivially, every VE can be modified to support VE.BatchEncProve and VE.BatchVerify by internally invoking the VE.EncProve and VE.Verify for each index in the batch, respectively. The main reason for defining this additional interface is to support the design of batch encryption and verification that are more efficient than the trivial approach.

- $\mathsf{VE.BatchEncProve}(pp_{\text{VE}}, I, \mathsf{pk}_I, \boldsymbol{s}, \boldsymbol{v}, \boldsymbol{w}) \rightarrow (\boldsymbol{c}, \pi_{\text{VE}})$. On input a vector $\boldsymbol{s}$ of messages, their commitments $\boldsymbol{v}$, corresponding witness $\boldsymbol{w}$, the algorithm outputs encryptions $\boldsymbol{c}$ for each $s \in \boldsymbol{s}$, along with a NIZK proof π_{VE} that satisfy VE.BatchVerify.
- $\mathsf{VE.BatchVerify}(pp_{\text{VE}}, I, \mathsf{pk}_I, \boldsymbol{v}, \boldsymbol{c}, \pi_{\mathsf{VE}}) \rightarrow 0/1$. The algorithm outputs 1 if π_{VE} is a valid proof that, for each $i \in I$ there exists (α_i, π_i) such that $\alpha_i, \pi_i = \mathsf{VE.Dec}(\mathsf{sk}_i, \boldsymbol{c}_i)$ and $\mathsf{PC.Verify}(\boldsymbol{v}, \alpha_i, \pi_i) = 1$

Constructions. Only a few VE schemes are known for discrete logarithm-based commitment schemes [35], [19], [42], [47]. These VE schemes are designed to work with the Feldman commitment scheme, where the dealer commits to a secret s as g^s. Note that the Feldman commitment scheme is not hiding. For instance, if the secret has low entropy, an adversary can recover the committed message by running a brute-force search on possible messages. As a result, these VE schemes cannot be directly used in general VSS schemes with arbitrary message distributions. Indeed, these VE schemes were designed for VSS schemes for Distributed Key Generation (DKG) protocols [42], [33], [47], where the shared secret is a random element from a large field.

Our dual-threshold AVSS requires a VE scheme for the Pedersen commitment scheme, where commitments are $g^s h^r$. To our knowledge, no such VE scheme has been described. We present modifications to Groth's VE [42] to make it compatible with the Pedersen commitment scheme in Appendix B.

Remark. If our dual-threshold VSS scheme is used to share secrets with high entropy, we can also employ existing VE schemes, such as those mentioned in [35], [19], [42], [47].

Algorithm 3 Dual-threshold AVSS

PUBLIC PARAMETERS: $n \geq 3t+1$, $\ell \geq t$, $\{pk_i\}_{i\in[n]}$, polynomial commitment PC and verifiable encryption VE.

PRIVATE INPUT: Signing key sk_i.

SHARING PHASE:

// Dealer L with input m:
Sample a $2t$-degree random polynomial $s(\cdot)$ with $s(0) = m$
$\boldsymbol{v}, r(\cdot) \leftarrow \mathsf{PC.Commit}(s(\cdot), n)$
for $i = 1, 2, ..., n$ **do**
  send $\langle \texttt{SHARE}, \boldsymbol{v}, s(i), r(i)\rangle$ to node i

// Each node i
upon receiving $\langle \texttt{SHARE}, \boldsymbol{v}, s(i), r(i)\rangle$ from dealer L **do**
  Check $\mathsf{PC.DegCheck}(\boldsymbol{v}, 2t) = 1$
  Check $\mathsf{PC.Verify}(\boldsymbol{v}, i, s(i), r(i)) = 1$
  if both the checks pass **then**
    Let $\sigma_i = \mathsf{sign}(sk_i, \boldsymbol{v})$
    send $\langle \texttt{ACK}, \sigma_i\rangle$ to L

// Dealer L waits for $2t+1$ valid signatures
Let $\boldsymbol{\sigma}$ be the set of valid signatures on $\boldsymbol{v}$.
Let I be the indices of nodes with *missing* valid signatures.
Partition I into subsets I_R and I_{VE} with $|I_R| = 2t - \ell$
$\boldsymbol{s}, \boldsymbol{\pi} \leftarrow \mathsf{PC.BatchOpen}(\boldsymbol{v}, I_R, p(\cdot), \boldsymbol{w})$
Let $\boldsymbol{s}_{I_{\mathsf{VE}}} \leftarrow \{p(i)\}$ for all $i \in I_{\mathsf{VE}}$.
$\boldsymbol{c}, \pi_{\text{VE}} \leftarrow \mathsf{VE.BatchEncProve}(I_{\mathsf{VE}}, \mathsf{pk}_{I_{\mathsf{VE}}}, \boldsymbol{s}_{I_{\mathsf{VE}}}, \boldsymbol{v}_{I_{\mathsf{VE}}}, \boldsymbol{w}_{I_{\mathsf{VE}}})$
send $(\boldsymbol{v}, I_R, I_{\mathsf{VE}}, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi}, \boldsymbol{c}, \pi_{\text{VE}})$ using a reliable broadcast.

// Node i upon broadcast outputs $(\boldsymbol{v}, I_R, I_{\mathsf{VE}}, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi}, \boldsymbol{c}, \pi_{\text{VE}})$.
Check if each $\sigma \in \boldsymbol{\sigma}$ is valid and $|\boldsymbol{\sigma}| \geq 2t+1$.
Check $I_R \cup I_{\mathsf{VE}}$ includes all nodes with missing signatures.
Check if $\mathsf{PC.BatchVerify}(\boldsymbol{v}, I_R, \boldsymbol{s}, \boldsymbol{\pi})$.
Check if $\mathsf{VE.BatchVerify}(I_{\mathsf{VE}}, \mathsf{pk}_{I_{\mathsf{VE}}}, \boldsymbol{v}, \boldsymbol{c}, \pi_{\text{VE}})$.
if all the checks pass **then**
  if received no valid SHARE message and $(p(i), \pi_i) \notin \boldsymbol{s}$ **then**
    Let $p(i), \pi_i \leftarrow \mathsf{VE.Dec}(\boldsymbol{c}[i], sk_i)$
  output $(\boldsymbol{v}, p(i), \pi_i)$; **return**

RECONSTRUCTION PHASE:

// every node i after finishing the sharing phase
send $\langle \texttt{RECON}, p(i), \pi_i\rangle$ to all.
upon receiving $\langle \texttt{RECON}, p(j), \pi_j\rangle$ from node j **do**
  if $\mathsf{PC.Verify}(\boldsymbol{v}, j, p(j), \pi_j)$ **then**
    $T = T \cup \{p(j)\}$
    if $|T| \geq 2t+1$ **then**
      output $p(0)$ using Lagrange interpolation; **return**

7.2. Dual-threshold AVSS Design

Let L be the dealer of the (n, ℓ, t) dual-threshold AVSS scheme (cf. Definition 3). Let PC and VE be the polynomial commitment and verifiable encryption scheme, respectively. We summarize our scheme in Algorithm 3 where we highlight the changes with respect to Algorithm 2 in gray.

Sharing phase. The first part of the Sharing phase is the same as the AVSS protocol in Algorithm 2. L shares its secret using a degree $2t$ polynomial $p(\cdot)$, computes its commitment $\boldsymbol{v}, \boldsymbol{w} \leftarrow \mathsf{PC.Commit}(p(\cdot), n)$, and then sends

$\langle \texttt{SHARE}, \boldsymbol{v}, p(i) \rangle$ to each node. Each node i upon receiving the SHARE message, validates it as in Algorithm 2, computes $\sigma_i = \mathsf{sign}(\mathsf{sk}_i, \boldsymbol{v})$, and responds to L with $\langle \texttt{ACK}, \sigma_i \rangle$.

L waits for $2t+1$ valid signed acknowledgements. Let $\boldsymbol{\sigma}$ be the set of valid acknowledgments, and let $I \subset [n]$ be the set of nodes from whom L does not receive ACK messages. Note that these include nodes who sent invalid ACK messages as well as nodes whose messages have not arrived. Next, L arbitrarily partitions I into two disjoint subsets I_R and I_N, such that $|I_R| = 2t - \ell$ and $|I_{\mathsf{VE}}| = \ell - t$. L then computes $\boldsymbol{s}, \boldsymbol{\pi} \leftarrow \mathsf{PC.BatchOpen}(p(\cdot), I_R, \boldsymbol{w})$ and $\boldsymbol{c}, \pi_{\mathrm{VE}} \leftarrow \mathsf{VE.BatchEncProve}(p(\cdot), \boldsymbol{v}, I_{\mathsf{VE}})$.

L then reliably broadcast the dual-threshold AVSS transcript $(\boldsymbol{v}, I_R, I_{\mathsf{VE}}, \boldsymbol{\sigma}, \boldsymbol{s}, \boldsymbol{\pi}, \boldsymbol{c}, \pi_{\mathrm{VE}})$ to all nodes. Upon receiving the transcript, nodes validate it by checking that: (i) $\boldsymbol{\sigma}$ is a valid set of signatures on $\boldsymbol{v}$ and $|\boldsymbol{\sigma}| \geq 2t+1$; (ii) $I_R \cup I_{\mathsf{VE}}$ includes all nodes whose signatures are not included $\boldsymbol{\sigma}$; (iii) $\boldsymbol{s}$ includes of valid shares of nodes in I_R with respect to $\boldsymbol{v}$ i.e., $\mathsf{PC.BatchVerify}(\boldsymbol{v}, I_R, \boldsymbol{s}, \boldsymbol{\pi})$; (iv) $\boldsymbol{c}$ includes verifiable ciphertexts using $\mathsf{VE.BatchVerify}$.

Upon successful verification, each node i locally outputs the commitment $\boldsymbol{v}$, its share $p(i)$, along with the commitment opening proof π_i to be used during the reconstruction phase. Node i either receives $p(i), \pi_i$ from SHARE message, or computes $p(i), \pi_i \leftarrow \mathsf{VE.Dec}(\boldsymbol{c}[i], \mathsf{sk}_i)$.

Reconstruction phase. The reconstruction phase is identical to the reconstruction phase of our AVSS scheme.

7.3. Optimization for Common Case Execution

In the dual-threshold AVSS we have described so far, the dealer L always verifiably encrypts $\ell - t$ of the remaining shares, which can be expensive for both L and other nodes. The following optimizations can significantly lower the number of shares L needs to encrypt in the common case: when the number of active failures is low and the network between L and most honest nodes is synchronous.

In the optimized design, in addition to waiting for $2t+1$ signed acknowledgments, the dealer L also waits for the network latency 2Δ, whichever occurs later. Let $2t+1+k$ for $k \geq 0$ be the number of signed acknowledgments the dealer receives. L then verifiably encrypts shares of $\max\{0, \ell - (t+k)\}$ nodes. This implies that with more signed acknowledgments, L needs to verifiably encrypt fewer shares. In the best-case scenario, i.e., when L receives $\ell - t$ additional signed acknowledgments, it need not compute any verifiable encryptions. Thus, in the best case, we get the dual-threshold property for free.

Remark. The optimization we describe above is also applicable to the AVSS scheme in §6. Also, the storage cost optimization we describe in §6.1 also applies to our dual-threshold AVSS scheme.

7.4. Analysis

Correctness and Asynchronous termination. Follows from similar arguments as the AVSS protocol.

Completeness. The soundness guarantees of the VE scheme ensure that nodes whose signature or share is not included in the VSS transcript will still receive its valid share upon decryption. This, combined with an argument similar to the synchronous VSS protocol, guarantees Completeness.

Secrecy. We prove Secrecy in Appendix C.

Performance. The computation cost of the dealer and nodes is similar to that of the AVSS protocol, except the transcript includes verifiable encryptions for a subset of nodes. Since the (amortized) computation cost of both computing verifiable encryptions and verifying them is linear in the number of encrypted shares [42], [47], both the dealer and nodes need to perform $O(n)$ group exponentiations. Additionally, the dealer sends a private message of length $O(\kappa n)$ to each node and a broadcast channel message of length $O(\kappa n)$ bits. Thus, using the broadcast channel from [32], the total communication cost is $O(\kappa n^2)$.

Combining all the above, we get the following theorem.

Theorem 3 (Dual-threshold AVSS)**.** *In an asynchronous network of $n \geq 3t+1$ nodes among which at most t nodes are malicious, assuming a polynomial commitment scheme, a signature scheme, a Byzantine reliable broadcast channel, and a Verifiable Encryption scheme, Algorithm 3 implements a t-resilient publicly verifiable (n, ℓ, t) dual-threshold AVSS protocol for any $\ell \in [t, n-t)$ with $O(\kappa n^2)$ communication costs. Here κ is the security parameter.*

8. Implementation and Evaluation

We evaluate our VSS schemes by implementing them in Rust. Our implementation is publicly available at https://github.com/sourav1547/vss . Our implementation uses the `blstrs` library [1], which implements efficient finite field and elliptic curve arithmetic. We also use (for both our implementation and the baselines) the multi-exponentiation of group elements using Pippenger's method [9, §4] for efficiency. For our dual threshold AVSS, we implement the verifiable encryption scheme we describe in Appendix B. Our experiments focus on the computation component, excluding any networking aspects. We will separately calculate the bandwidth usage of our scheme and the baselines. Our implementation supports two different signature schemes: the Schnorr signature using Ed25519 elliptic curve [58] scheme and the BLS signature using the BLS12-381 elliptic curve [13]. The Schnorr signature has faster signature verification time but requires interactive aggregation. The BLS signature scheme supports non-interactive aggregation but requires two pairings per signature verification. In the context of our VSS, the dealer needs to perform $2n$ pairings to verify the signatures.

8.1. Evaluation Setup

We evaluate VSS schemes using four key metrics: *dealing time*, *verification time*, *bandwidth usage*, and *reconstruction time*. We explain each of these metrics below:

Table 1: Evaluation of AVSS schemes. Dealing runtime measures the computation cost of the dealer. Verification time refers to the per-node computation cost. Bandwidth usage is the amount of data the dealer sends over the broadcast channel. Our low-threshold AVSS ($\ell = t$) is a special case of our dual-threshold AVSS with $\ell = t$, so they share the same performance numbers. The worst-case performance of our dual-threshold AVSS with $\ell \in [t, 2t]$ degrades linearly as ℓ increases, and the best-case performance is similar to the $\ell = t$ case.

		Dealing time (in ms)			Verification time (in ms)			Bandwidth usage (in KBytes)		
Scheme	ℓ	$n = 256$	512	1024	$n = 256$	512	1024	$n = 256$	512	1024
Yurek et al. [64] (best case)†	t	48.44	96.04	191.78	1.91	3.80	8.11	28	56	112
VE-based VSS [42]+§B	$[t, 2t]$	845.37	1685.20	3419.30	316.31	611.94	1224.01	243.43	482.50	963.43
Ours (w/o multisig)	t	48.70	104.66	235.97	48.70	104.66	235.97	28.03	55.95	112.12
Ours (w/ BLS multisig)	t	156.84	311.43	630.78	6.16	11.30	22.42	17.39	34.68	69.48
Ours (w/o multisig) (worst case)	$2t$	319.22	650.05	1294.90	117.65	236.08	490.25	102.08	200.61	398.83
Ours (w/ BLS multisig) (worst case)	$2t$	419.01	836.41	1676.00	113.89	221.03	430.95	91.42	179.32	356.17

† In the worst-case scenario, each node is required to verify shares from $t + 1$ other nodes and also reconstruct the secret during the sharing phase. This leads to a significant increase in the per-node verification time. The Dealing time and the bandwidth usage remain unchanged.

Dealing time. The dealing time measures the computation cost of the dealer in preparing the transcript. Specifically, it refers to the time dealer takes to compute the polynomial commitment, the opening proofs for each node, verify the signed acknowledgments from all nodes, and aggregate them into a BLS multisignature or a list of Schnorr signatures. For our dual-threshold AVSS, the dealing time also includes the computation time required to generate verifiable encryptions of a subset of shares. The dealing time does *not* include the computation cost of sending messages, such as broadcasting the VSS transcript.

Verification time. This metric measures the computation cost experienced by the nodes. It refers to the time a node takes to verify the degree of the committed polynomial, sign the polynomial commitment, verify the signatures on the polynomial commitment, validate the revealed shares (including its own), and the verifiable encryptions (applicable only to dual-threshold AVSS) provided by the dealer. The verification time does not include the computation cost associated with networking.

Bandwidth usage. We measure bandwidth usage as the amount of data the dealer sends over the broadcast channel. We only include the data sent over the broadcast channel as broadcasting is more expensive than sending private messages. Furthermore, the data the dealer sends privately to each node and the responses of the nodes are smaller than the data sent over the broadcast channel.

Reconstruction time. The reconstruction time measures the computation cost of reconstructing a secret from its shares. This consists of the cost of verifying shares from each node, computing appropriate Lagrange coefficients, and the final inner product. Note that the reconstruction time of a VSS scheme depends on the degree of the polynomial used to share the secret and the cost of verifying each share.

Baselines. The first baseline VSS protocol we compare with is the AVSS protocol of [64] with optimizations from [32]. Recall from §2, this scheme relies on complaints and does not terminate even with a single faulty node. Thus, as our baseline, we measure the dealing and verification time as the computation cost in the best case, i.e., without any faulty nodes. Similarly, we also measure the bandwidth usage of this scheme as the amount of data the dealer sends using an RBC in the best case. We choose this as one of our baselines as it is the most efficient AVSS scheme, and by comparing it with this scheme, we seek to demonstrate that our AVSS guarantees asynchronous Termination and public verifiability with minimal overhead. For the Yurek et al. [64] scheme, we implement the polynomial commitment scheme in Figure 1 instead of standard Pedersen commitment to coefficients. Although committing to the evaluation points increases the dealing time, we adopt this approach as it lowers the computation cost during the complaint and reconstruction phase.

Our second baseline is the verifiable encryption-based VSS scheme, which works as follows: The dealer computes shares of each node using a polynomial of degree d (typically $d = t$), computes the commitment to the polynomial, and verifiably encrypts shares of each node. The dealer then broadcasts the commitment and ciphertexts to all nodes using a broadcast channel. Concretely, we implement a VSS scheme based on the verifiable encryption scheme we describe in Appendix B, as existing VSS schemes based on verifiable encryption [42], [33], [47] only achieve a weaker Secrecy property. This baseline achieves similar properties to our scheme: it supports dual-threshold, is publicly verifiable, and works in synchronous and asynchronous networks.

We want to note that the VE in Appendix B has a parameter m that indicates the number of chunks we divide a secret into. A smaller value of m results in quicker dealing times but also leads to longer worst-case decryption times. For our evaluations, we opt for $m = 16$ to favor the baseline, i.e., to give it faster dealing and verification time in the absence of failures. However, with $m = 16$, in the worst case, a node would have to perform more than 2^{21} group exponentiations to decrypt its shares.

8.2. Evaluation Results

All experiments are run using a single thread on an Apple M2 Pro device with 16 GB RAM and 12 cores. We report our results in Table 1. Recall that our scheme and VE-based VSS [42] both provide asynchronous termination, public verifiability, and dual-threshold, while Yurek et al. [64] does not. We want to show through the evaluation that our scheme only adds a small overhead compared to Yurek et al. [64]

Table 2: AVSS reconstruction time (in milliseconds). For synchronous VSS, our reconstruction time is the same as the baseline.

Scheme	$n = 256$	$n = 512$	$n = 1024$
Baseline	16.42	32.61	65.49
Ours	32.57	65.16	131.60

to achieve these properties while significantly improving the performance over VE-based VSS [42].

Dealing time. For low threshold, $\ell = t$, compared to Yurek et al. [64], our AVSS dealing time is slightly larger as our dealer additionally needs to validate the signature on the acknowledgment messages. With BLS multisignature scheme, the dealer performs two pairings per signature verification; hence the dealing time is 3 times higher. Compared to VE-based VSS, our shortens the dealing time (without multisignature) by more than $15\times$. For high threshold, $\ell = 2t$ or dual-threshold, compared to VE-based VSS [42]+§B, our AVSS dealing time is better by about $2\times$, with or without BLS multisignature. This is because the dealer in our dual-threshold AVSS scheme verifiably encrypts $\ell - t$ shares instead of all n shares. We reiterate that for $\ell > t$, in the best-case scenario, our dual-threshold AVSS has a dealing time comparable to our low-threshold AVSS. Hence, in the best case, our dual-threshold AVSS also improves the performance by 5-15$\times$.

Verification time. As we report in Table 1, for low threshold $\ell = t$, our verification with BLS multisignature is about $3\times$ larger than the best case verification time of [64]. This is because each node in our schemes needs to additionally validate the signatures and shares of other nodes revealed by the dealer. However, the absolute verification time is very small, e.g., only 22 milliseconds for 1024 nodes. Also, it is $60\times$ smaller than the VE-based VSS scheme. For high threshold $\ell = 2t$, compared to VE-based VSS schemes, our verification time of our protocol is $3\times$ and $60\times$ better in the worst and best case, respectively.

Bandwidth usage. For low threshold $\ell = t$, our scheme has similar (without multisignature) or better (with multisignature) bandwidth cost compared with Yurek et al. [64], and is significantly better than VE based VSS scheme. For dual threshold with $\ell = 2t$, our bandwidth usage is about $2\times$ better (in the worst case) than the VE-based VSS scheme.

Reconstruction time. Recall from §8.1 that the reconstruction time depends only on the degree of the polynomial used to share the secret and the cost of verifying each share. Since our synchronous VSS protocol uses the same polynomial degree and the same share verification procedure as the baseline, its reconstruction time is identical to the baseline protocol. On the other hand, the dealer in our AVSS scheme uses a degree $2t$ polynomial, compared to degree t polynomial used by existing AVSS schemes. We report the reconstruction time (in milliseconds) in Table 2. As expected, our reconstruction time is twice as expensive as the baseline. Nevertheless, the absolute values are very small, e.g., 132 milliseconds for 1024 nodes.

Comparison with Class-group based VSS [47]. Very recently, Kate et al. [47] improved the efficiency of [42] for high-entropy secrets using a non-standard class-group assumption. Since their implementation is not publicly available, we only estimate how it compares to our scheme. The bandwidth usage of [47] is $219(n+1)+48n$ bytes (assuming we commit to evaluation points instead of coefficients), which is approximately $3\times$ higher than our bandwidth usage. Regarding dealing and verification time, [47] reports $2.7\times$ improvement over [42]. Since our dealing time is $15\times$ better than [42], we anticipate that our dealing time is $5\times$ better than that of [47]. Similarly, we expect our verification time to be 2-3$\times$ better. Note that we achieve these improvements while relying on the standard discrete logarithm assumption.

9. Discussion and Conclusion

Interactive vs. non-interactive protocols. In existing VE-based VSS [33], [42], [47] the dealer sends a single message over the broadcast channel. On the other hand, our VSS protocols require interaction between the dealer and the other nodes (but not among the nodes). As a result, our protocols are slightly more complex to implement. Yet, we believe that the substantial performance improvements offered by our protocols outweigh the added complexity. Designing a more efficient non-interactive public verifiable secret sharing scheme remains a fascinating open question.

Applications with polynomials of an arbitrary degree. Although our AVSS scheme shares the secret using a degree $2t$ polynomial, applications of AVSS such as asynchronous DKG, asynchronous proactive secret sharing, etc., need not use a degree $2t$ polynomials. Instead, these applications can share their secret using an arbitrary degree polynomial, using the degree switching trick of [31].

Conclusion. We have presented a simple paradigm to design three efficient verifiable secret sharing protocols for various settings, i.e., under synchrony, asynchrony, and asynchrony with dual-threshold. All our protocols are optimal fault-tolerant, i.e., we can tolerate $n/2$ and $n/3$ failures in synchrony and asynchrony, respectively. Unlike existing schemes, our VSS protocols do not rely on complaints and require only a single broadcast. Our protocols output efficient publicly verifiable transcripts and support dual-threshold in asynchrony. Moreover, our asynchronous VSS protocols ensure natural termination, a shortcoming in many existing asynchronous VSS schemes.

Our VSS protocols maintain the same asymptotic performance as state-of-the-art counterparts while relying on milder cryptographic assumptions and setups. Furthermore, our scheme also results in significant concrete improvements compared to existing VSS protocols with similar properties.

Future research directions. Several recent VSS schemes support batching [64], [43], [8], [3]. Very recently, the Bingo protocol [3] achieved security of the AVSS scheme in the presence of an adaptive adversary [3]. Extending our VSS to support these properties while maintaining its simplicity and performance is a fascinating research direction.

Acknowledgments

This work is funded in part by a VMware early career faculty grant, a Chainlink Labs Ph.D. fellowship, and the National Science Foundation award #2240976.

References

[1] blstrs library. [Online]. Available: https://docs.rs/blstrs/latest/blstrs/

[2] I. Abraham, D. Dolev, and G. Stern, "Revisiting asynchronous fault tolerant computation with optimal resilience," in *Proceedings of the 39th Symposium on Principles of Distributed Computing*, 2020, pp. 139–148.

[3] I. Abraham, P. Jovanovic, M. Maller, S. Meiklejohn, and G. Stern, "Bingo: Adaptively secure packed asynchronous verifiable secret sharing and asynchronous distributed key generation," in *Annual International Cryptology Conference*. Springer, 2023.

[4] N. Alhaddad, M. Varia, and H. Zhang, "High-threshold avss with optimal communication complexity," in *International Conference on Financial Cryptography and Data Security*. Springer, 2021, pp. 479–498.

[5] O. Alpos, C. Cachin, S. H. Kamp, and J. B. Nielsen, "Practical large-scale proof-of-stake asynchronous total-order broadcast," *Cryptology ePrint Archive*, 2023.

[6] M. Backes, A. Datta, and A. Kate, "Asynchronous computational vss with reduced communication complexity," in *Cryptographers' Track at the RSA Conference*. Springer, 2013, pp. 259–276.

[7] M. Ben-Or, S. Goldwasser, and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation," in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, ser. STOC '88, New York, NY, USA, 1988, p. 1–10.

[8] F. Benhamouda, S. Halevi, H. Krawczyk, A. Miao, and T. Rabin, "Threshold cryptography as a service (in the multiserver and yoso models)," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 323–336.

[9] D. J. Bernstein, J. Doumen, T. Lange, and J.-J. Oosterwijk, "Faster batch forgery identification," in *Progress in Cryptology-INDOCRYPT 2012: 13th International Conference on Cryptology in India, Kolkata, India, December 9-12, 2012. Proceedings 13*. Springer, 2012, pp. 454–473.

[10] A. Bhat, N. Shrestha, Z. Luo, A. Kate, and K. Nayak, "Randpiper–reconfiguration-friendly random beacons with quadratic communication," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 3502–3524.

[11] G. R. Blakley, "Safeguarding cryptographic keys," in *1979 International Workshop on Managing Requirements Knowledge (MARK)*. IEEE, 1979, pp. 313–318.

[12] D. Boneh, M. Drijvers, and G. Neven, "Compact multi-signatures for smaller blockchains," in *Advances in Cryptology–ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part II*. Springer, 2018, pp. 435–464.

[13] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the weil pairing," in *Advances in Cryptology—ASIACRYPT 2001: 7th International Conference on the Theory and Application of Cryptology and Information Security Gold Coast, Australia, December 9–13, 2001 Proceedings 7*. Springer, 2001, pp. 514–532.

[14] G. Bracha, "Asynchronous byzantine agreement protocols," *Information and Computation*, vol. 75, no. 2, pp. 130–143, 1987.

[15] L. T. Brandao, L. T. Brandao, M. Davidson, and A. Vassilev, "Nist roadmap toward criteria for threshold schemes for cryptographic primitives," 2020.

[16] C. Cachin, K. Kursawe, A. Lysyanskaya, and R. Strobl, "Asynchronous verifiable secret sharing and proactive cryptosystems," in *Proceedings of the 9th ACM Conference on Computer and Communications Security*, 2002, pp. 88–97.

[17] C. Cachin, K. Kursawe, F. Petzold, and V. Shoup, "Secure and efficient asynchronous broadcast protocols," in *Annual International Cryptology Conference*. Springer, 2001, pp. 524–541.

[18] C. Cachin, K. Kursawe, and V. Shoup, "Random oracles in constantipole: practical asynchronous byzantine agreement using cryptography," in *Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing*, 2000, pp. 123–132.

[19] J. Camenisch and V. Shoup, "Practical verifiable encryption and decryption of discrete logarithms," in *Annual International Cryptology Conference*. Springer, 2003, pp. 126–144.

[20] R. Canetti, "Studies in secure multiparty computation and applications," Ph.D. dissertation, Citeseer, 1996.

[21] R. Canetti and T. Rabin, "Fast asynchronous byzantine agreement with optimal resilience," in *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, 1993, pp. 42–51.

[22] I. Cascudo and B. David, "Scrape: Scalable randomness attested by public entities," in *International Conference on Applied Cryptography and Network Security*. Springer, 2017, pp. 537–556.

[23] ——, "Albatross: publicly attestable batched randomness based on secret sharing," in *Advances in Cryptology–ASIACRYPT 2020: 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7–11, 2020, Proceedings, Part III 26*. Springer, 2020, pp. 311–341.

[24] P. Chaidos and A. Kiayias, "Mithril: Stake-based threshold multisignatures," *Cryptology ePrint Archive*, 2021.

[25] K. Choi, A. Manoj, and J. Bonneau, "Sok: Distributed randomness beacons," in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2023.

[26] B. Chor, S. Goldwasser, S. Micali, and B. Awerbuch, "Verifiable secret sharing and achieving simultaneity in the presence of faults," in *26th Annual Symposium on Foundations of Computer Science (sfcs 1985)*. IEEE, 1985, pp. 383–395.

[27] A. Choudhury, "Optimally-resilient unconditionally-secure asynchronous multi-party computation revisited," Cryptology ePrint Archive, Report 2020/906, 2020. https://eprint. iacr. org ..., Tech. Rep., 2020.

[28] I. Damgård, "On σ-protocols," *Lecture Notes, University of Aarhus, Department for Computer Science*, p. 84, 2002.

[29] I. Damgård and J. B. Nielsen, "Scalable and unconditionally secure multiparty computation," in *Annual International Cryptology Conference*. Springer, 2007, pp. 572–590.

[30] S. Das, V. Krishnan, I. M. Isaac, and L. Ren, "Spurt: Scalable distributed randomness beacon with transparent setup," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 2502–2517.

[31] S. Das, Z. Xiang, L. Kokoris-Kogias, and L. Ren, "Practical asynchronous high-threshold distributed key generation and distributed polynomial sampling," in *32st USENIX Security Symposium (USENIX Security 23)*. USENIX Association, 2023.

[32] S. Das, Z. Xiang, and L. Ren, "Asynchronous data dissemination and its applications," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.

[33] S. Das, T. Yurek, Z. Xiang, A. Miller, L. Kokoris-Kogias, and L. Ren, "Practical asynchronous distributed key generation," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 2518–2534.

[34] P. Feldman, "A practical scheme for non-interactive verifiable secret sharing," in *28th Annual Symposium on Foundations of Computer Science (sfcs 1987)*. IEEE, 1987, pp. 427–438.

[35] P.-A. Fouque and J. Stern, "One round threshold discrete-log key generation without private channels," in *International Workshop on Public Key Cryptography*. Springer, 2001, pp. 300–316.

[36] R. Gennaro, S. Jarecki, H. Krawczyk, and T. Rabin, "Robust threshold dss signatures," in *Advances in Cryptology—EUROCRYPT'96: International Conference on the Theory and Application of Cryptographic Techniques Saragossa, Spain, May 12–16, 1996 Proceedings 15*. Springer, 1996, pp. 354–371.

[37] ——, "Secure distributed key generation for discrete-log based cryptosystems," *Journal of Cryptology*, vol. 20, no. 1, pp. 51–83, 2007.

[38] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," in *Proceedings of the 26th symposium on operating systems principles*, 2017, pp. 51–68.

[39] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof-systems," in *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, ser. STOC '85, p. 291–304.

[40] V. Goyal, Y. Song, and C. Zhu, "Guaranteed output delivery comes free in honest majority mpc," in *Annual International Cryptology Conference*. Springer, 2020, pp. 618–646.

[41] J. Groth, "Short pairing-based non-interactive zero-knowledge arguments." in *Asiacrypt*, vol. 6477. Springer, 2010, pp. 321–340.

[42] ——, "Non-interactive distributed key generation and key resharing." *IACR Cryptol. ePrint Arch.*, vol. 2021, p. 339, 2021.

[43] J. Groth and V. Shoup, "Design and analysis of a distributed ecdsa signing service," *Cryptology ePrint Archive*, 2022.

[44] K. Gurkan, P. Jovanovic, M. Maller, S. Meiklejohn, G. Stern, and A. Tomescu, "Aggregatable distributed key generation," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2021, pp. 147–176.

[45] M. Hirt, J. B. Nielsen, and B. Przydatek, "Asynchronous multiparty computation with quadratic communication," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2008, pp. 473–485.

[46] A. Kate and I. Goldberg, "Distributed key generation for the internet," in *2009 29th IEEE International Conference on Distributed Computing Systems*. IEEE, 2009, pp. 119–128.

[47] A. Kate, E. V. Mangipudi, P. Mukherjee, H. Saleem, and S. A. K. Thyagarajan, "Non-interactive vss using class groups and application to dkg," *Cryptology ePrint Archive*, 2023.

[48] A. Kate, G. M. Zaverucha, and I. Goldberg, "Constant-size commitments to polynomials and their applications," in *International conference on the theory and application of cryptology and information security*. Springer, 2010, pp. 177–194.

[49] E. Kokoris Kogias, D. Malkhi, and A. Spiegelman, "Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures." in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020, pp. 1751–1767.

[50] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," in *Concurrency: the works of leslie lamport*, 2019, pp. 203–226.

[51] Y. Lindell, "Secure multiparty computation (mpc)," *Cryptology ePrint Archive*, 2020.

[52] A. Momose and L. Ren, "Optimal communication complexity of authenticated byzantine agreement," in *35th International Symposium on Distributed Computing*, 2021.

[53] K. Nayak, L. Ren, E. Shi, N. H. Vaidya, and Z. Xiang, "Improved extension protocols for byzantine broadcast and agreement," in *34th International Symposium on Distributed Computing, DISC 2020*, 2020.

[54] A. Patra, A. Choudhary, and C. P. Rangan, "Efficient statistical asynchronous verifiable secret sharing with optimal resilience," in *International Conference on Information Theoretic Security*. Springer, 2009, pp. 74–92.

[55] A. Patra, A. Choudhury, and C. P. Rangan, "Efficient asynchronous verifiable secret sharing and multiparty computation," *Journal of Cryptology*, vol. 28, no. 1, pp. 49–109, 2015.

[56] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *Annual international cryptology conference*. Springer, 1991, pp. 129–140.

[57] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of the society for industrial and applied mathematics*, vol. 8, no. 2, pp. 300–304, 1960.

[58] C.-P. Schnorr, "Efficient identification and signatures for smart cards," in *Advances in Cryptology—CRYPTO'89 Proceedings 9*. Springer, 1990, pp. 239–252.

[59] B. Schoenmakers, "A simple publicly verifiable secret sharing scheme and its application to electronic voting," in *Annual International Cryptology Conference*. Springer, 1999, pp. 148–164.

[60] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.

[61] V. Shoup and N. P. Smart, "Lightweight asynchronous verifiable secret sharing with optimal resilience," *Cryptology ePrint Archive*, 2023.

[62] R. Vassantlal, E. Alchieri, B. Ferreira, and A. Bessani, "Cobra: Dynamic proactive secret sharing for confidential bft services," in *2022 IEEE symposium on security and privacy (SP)*. IEEE, 2022, pp. 1335–1353.

[63] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "Hotstuff: Bft consensus with linearity and responsiveness," in *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. ACM, 2019, pp. 347–356.

[64] T. Yurek, L. Luo, J. Fairoze, A. Kate, and A. Miller, "hbacss: How to robustly share many secrets," in *Proceedings of the 29th Annual Network and Distributed System Security Symposium*, 2022.

[65] H. Zhang, S. Duan, C. Liu, B. Zhao, X. Meng, S. Liu, Y. Yu, F. Zhang, and L. Zhu, "Practical asynchronous distributed key generation: Improved efficiency, weaker assumption, and standard model," *Cryptology ePrint Archive*, 2022.

[66] J. Zhang, T. Xie, T. Hoang, E. Shi, and Y. Zhang, "Polynomial commitment with a {One-to-Many} prover and applications," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2965–2982.

Appendix A. Additional Preliminaries

A.1. Broadcast Channel

Our synchronous VSS and AVSS protocols make black box invocations to a Byzantine broadcast and Byzantine reliable broadcast protocol, respectively. We use state-of-the-art broadcast extension protocols, i.e., for long messages [53], [32]. For completeness, we include the definitions of Byzantine (reliable) broadcast below.

Definition 5 (Byzantine Broadcast)**.** A protocol for a set of nodes $\{1, \ldots, n\}$ including a designated broadcaster who holds an initial input, is a Byzantine broadcast protocol if the following properties hold

- **Agreement.** If an honest node outputs a message M and another honest node outputs M', then $M = M'$.

$\mathsf{PC.BatchOpen}(pp, \boldsymbol{w}, p(\cdot), I = \{i_1, \ldots, i_k\})$: Output $(\boldsymbol{s}, \boldsymbol{\pi})$ where

$$\boldsymbol{s} = [p(i_1), \ldots, p(i_k)]; \text{ and } \boldsymbol{\pi} = [r(i_1), \ldots, r(i_k)] \quad (3)$$

$\mathsf{PC.BatchVerify}(pp, \boldsymbol{v}, I = \{i_1, \ldots, i_k\}, \boldsymbol{s}, \boldsymbol{\pi})$: Given a subset $I \subseteq [n]$, let $k = |I|$. Assert $k = |\boldsymbol{v}| = |\boldsymbol{\pi}|$. Sample a uniform random vector $[\gamma_1, \ldots, \gamma_k] \in \mathbb{F}^k$. Let $s = \sum_{j \in [k]} \gamma_j s_j$ and $\pi = \sum_{j \in [k]} \gamma_j \pi_j$. Output 1, if the following holds, otherwise output 0.

$$\prod_{j \in [k]} \boldsymbol{v}[i_j]^{\gamma_j} = g^s h^{\pi} \quad (4)$$

Figure 3: Batched interfaces for the Polynomial commitment.

- **Validity.** If the sender is honest and has input M, all honest nodes output M.
- **Termination.** Every honest node outputs a message.

Definition 6 (Byzantine Reliable Broadcast)**.** A protocol for a set of nodes $\{1, \ldots, n\}$ including a designated broadcaster who holds an initial input, is a Byzantine reliable broadcast protocol if the following properties hold

- **Agreement** and **Validity.** Same as Byzantine broadcast.
- **Totality.** If an honest node outputs a message, then every honest node eventually outputs a message.

The optimal fault-tolerant synchronous Byzantine broadcast [53], [52] achieves $O(n|M|+\kappa n^2)$ communication cost for a message M assuming powers-of-tau [48] and q-SDH, and $O(n|M|+\kappa n^2 \log n)$ communication cost assuming collision resistant hash function. The optimal fault-tolerant asynchronous Byzantine reliable broadcast [32] achieves a communication cost of $O(n|M|+\kappa n^2)$ for a message M assuming collision-resistant hash functions.

A.2. Batched interface for polynomial commitment

As we briefly describe in §3.2, in our VSS schemes, the dealer reveals shares for a list of nodes for everyone to verify. Thus, we introduce the following additional interface for batched opening and verification. Specifically, for any set $I \subseteq [n]$, we require the polynomial commitment to provide the following interfaces.

- $\mathsf{PC.BatchOpen}(pp, \boldsymbol{w}, p(\cdot), I = \{i_1, \ldots, i_k\}) \rightarrow (\boldsymbol{u}, \boldsymbol{\pi})$. On input the set of indices I, the polynomial $p(\cdot)$ and witness $\boldsymbol{w}$, outputs $\boldsymbol{u} = [p(i_1), \ldots, p(i_k)]$ along with batch opening proof $\boldsymbol{\pi} = [\pi_{i_1}, \ldots, \pi_{i_n}]$.
- $\mathsf{PC.BatchVerify}(pp, \boldsymbol{v}, I = \{i_1, \ldots, i_k\}, \boldsymbol{u}, \boldsymbol{\pi}) \rightarrow 0/1$. On input the commitment $\boldsymbol{v}$ to a polynomial $p(i)$, outputs 1 if $\boldsymbol{u}[j] = p(i_j)$ for all $i_j \in I$, and outputs 0 otherwise.

We describe the concrete instantiations of the batched interfaces in Figure 3. Here we use a random linear combination to verify all the openings using a single multi-exponentiations of width k instead of $2k$ exponentiations.

Appendix B. Verifiable Encryptions of Discrete Logarithm

B.1. Verifiable Encryption Scheme of [42]

The VE scheme of Groth [42] works with Feldman commitment where a message $s \in \mathbb{F}$ is committed as g^s for some pre-specified generator $g \in \mathbb{G}$. We can not use it to design a VSS protocol as the Feldman commitment scheme is not hiding for messages with small entropy; an adversary can exhaustively search the message space to derive a matching commitment. Nevertheless, we will use Groth's VE to design a VE that works with the Pedersen commitment scheme. Next, we will briefly describe the relations $\mathcal{P}$ in Groth's VE proves and discuss how our modifications require $\mathcal{P}$ to prove a similar relation.

Let $v = g^s$ be the commitment to the secret s. $\mathcal{P}$ computes, among other things, the ElGamal encryption of v, i.e., $c_v = (c_{v,0}, c_{v,1}) = (g^a, v\mathsf{pk}^a)$. Here $\mathsf{pk} = g^{\mathsf{sk}}$ is the public key of the recipient with secret key sk. $\mathcal{P}$ then computes the NIZK proof in two parts: *Proof of correct sharing* and *Proof of correct chunking*.

Proof of correct sharing. In the first part, for the tuple (v, c_v, pk), $\mathcal{P}$ proves, using a Σ-protocol, that c_v is an ElGamal encryption of v for the public key pk.

Proof of correct chunking. In the second part, $\mathcal{P}$ proves that the ciphertext is decryptable. Let $\boldsymbol{c}_v$ be a vector of ElGamal ciphertexts where each ciphertext encrypts a small number of bits (called chunks) of s. Let $(\mathsf{pk}, c_v, \boldsymbol{c}_v)$ be the entire ciphertext (of commitment and each chunk of s), then $\mathcal{P}$ proves that $s \leftarrow \mathsf{Dec}(\mathsf{sk}, c_v, \boldsymbol{c}_v)$. We refer the reader to [42, §6.5]) for more details.

B.2. VE for Pedersen commitments

Our new VE for Pedersen commitment maintains the two-part structure of Groth's VE. Looking ahead, we provide support for the Pedersen commitment scheme only by changing the protocol for proof of correct sharing. Moreover, our modification adds only two group elements and a single field element to the Groth's VE proof. We discuss our changes next.

Proof of correct sharing. Let $g^s h^r$ be the Pedersen commitment to s. Let $v = g^s$ and $u = h^r$, hence the commitment to s is $v \cdot u$. In our scheme, the ciphertext also contains the ElGamal encryption of u, i.e. $c_u = (c_{u,0}, c_{u,1}) = (g^b, u\mathsf{pk}^b)$, along with $c_v = (c_{v,0}, c_{v,1}) = (g^a, v\mathsf{pk}^a)$. Now, $\mathcal{P}$ and $\mathcal{V}$ locally computes c_{vu}, where,

$$c_{vu} = (c_{v,0} \cdot c_{u,0}, c_{v,1} \cdot c_{u,1}) = (g^{a+b}, vu \cdot \mathsf{pk}^{a+b})$$

$\mathcal{P}$ in our VE then uses the protocol for proof of correct sharing of Groth's VE (with standard modifications [28]) for the tuple $(vu, c_{vu}, \mathsf{pk})$ to prove that c_{uv} is an ElGamal encryption of vu for public key pk.

Proof of correct chunking. Since the ciphertext of our VE remains unchanged (with the exception of one additional ElGamal encryption), a tempting approach is to directly use

the protocol for proof of correct chunking of Groth's VE protocol as the second part of our VE scheme. Intuitively, proof of correct chunking protocol of Groth's VE guarantees that a node with secret sk will be able to decrypt s as $\mathsf{Dec}(sk, c_v, \boldsymbol{c}_v)$. Although it is true, there is one subtle issue.

Eventually, to reconstruct the secret, we require each node to reveal its share along with a opening proof. For Pedersen commitment $g^s h^r$, the natural opening proof is r. This implies that to fully support Pedersen commitments, we need to add additional information $\boldsymbol{c}_u$ to the ciphertext and the NIZK proof such that $(s, r) \leftarrow \mathsf{Dec}(\mathsf{sk}, c_u, c_v, \boldsymbol{c}_v, \boldsymbol{c}_u)$.

The obvious approach is to repeat the protocol to prove the decryptability of c_v for c_u, as well. However, this would increase the computation cost of dealing and verifying the transcript and the transcript size by a factor of 2. Next, we describe our approach that addresses this issue without increasing the ciphertext size, thus avoiding the $2\times$ overhead.

Our key observation is that the opening proof of a Pedersen commitment $g^s h^r$ need not be r. Instead, it can be $(u = h^r, \pi_u)$ where π_u proves that u is correctly computed. Thus, in our VE, we let $\mathcal{R}$ recover (u, π_u), where $\mathcal{R}$ uses π_u to convince others regarding the correctness of u.

Computing u is trivial as it is the ElGamal decryption of c_u using the secret key sk. We define π_u as the tuple $(\mathsf{pk}^b, \pi_{\mathsf{pk}})$ where π_{pk} is a discrete logarithm equality (DLEq) proof for the tuple $(g, \mathsf{pk}, c_{u,0}, \mathsf{pk}^b)$. More precisely, π_{pk} convinces any verifier that $\log_g \mathsf{pk} = \log_{c_{u,0}} \mathsf{pk}^b$.

Each node upon receiving $\pi_u = (\mathsf{pk}^b, \pi_{\mathsf{pk}})$, checks the correctness of the DLEq relation using π_{pk} and $c_{u,0}$. Upon successful validation, the node computes $h^r = c_{u,1}/\mathsf{pk}^b$. Finally, the node checks the correctness of s by checking whether $g^s h^r = vu$.

Appendix C. Secrecy Proofs

We prove Secrecy of our VSS protocols using *simulatability*: for every probabilistic polynomial-time (PPT) adversary $\mathcal{A}$ that statically corrupts up to t nodes, there exists an ideal world PPT simulator that interacts with the ideal functionality and produces a view such that $\mathcal{A}$'s view in the simulated world is indistinguishable to a run of the Sharing phase.

Secrecy of Synchronous VSS. We prove Secrecy of our synchronous VSS with respect to $\mathcal{F}_{\text{VSS}}$ ideal functionality (cf. Figure 4). Let $\mathcal{S}_{\text{VSS}}$ be corresponding simulator. $\mathcal{S}_{\text{VSS}}$ simulate $\mathcal{A}$'s view using the Pedersens' polynomial commitment scheme. We summarize $\mathcal{S}_{\text{VSS}}$ in Figure 5, and prove the following theorem.

Lemma 1 (Synchronous VSS Secrecy)**.** *$\mathcal{A}$'s view in its interaction with $\mathcal{S}_{\text{VSS}}$ is identically distributed to its view in the real protocol.*

Proof. Let $h = g^\alpha$ for some non-zero $\alpha \in \mathbb{F}$. For any fixed commitment $\boldsymbol{v}$, consider the probability of outputting $\boldsymbol{v}$ and $s(i)$ for each $i \in \mathcal{C}$ in a real protocol. For a fixed polynomial $s(\cdot)$, there exists a unique polynomial $r(\cdot)$ that outputs $\boldsymbol{v}$ as the commitment. Since the dealer in the honest protocol samples $r(\cdot)$ uniformly at random, in the real protocol $\Pr[\boldsymbol{v}, \{s(i)\}_{i\in\mathcal{C}}]_{\text{real}} = 1/|\mathbb{F}|^{t+1}$.

Now consider the probability of the same event in the simulated view. For a fixed $\hat{s}(\cdot)$, a unique degree t polynomial $\hat{r}(\cdot)$ exists that results in $\boldsymbol{v}$ as the commitment. In particular, the unique $\hat{r}(\cdot)$ is:

$$\hat{r}(x) = r(x) + \frac{s(x) - \hat{s}(x)}{\alpha} \tag{5}$$

Since $\mathcal{S}_{\text{VSS}}$ samples $\hat{r}(\cdot)$ uniformly at random,

$$\begin{aligned}\Pr[\boldsymbol{v}, \{\hat{s}(i)\}_{i\in\mathcal{C}}]_{\text{id}} &= \Pr\left[\hat{r}(x) = r(x) + \frac{s(x) - \hat{s}(x)}{\alpha}\right]_{\text{id}} \\ &= 1/|\mathbb{F}|^{t+1} \end{aligned} \tag{6}$$

Equation (6) implies that the polynomial commitment and shares seen by $\mathcal{A}$ are identically distributed in real and simulated view. Since $\mathcal{S}_{\text{VSS}}$ simulates the rest of the protocol as per protocol specification, the distribution of the remaining messages seen by $\mathcal{A}$ is also identical in both the real and simulated world. □

Secrecy of Asynchronous VSS. We prove Secrecy of our AVSS scheme with respect to $\mathcal{F}_{\text{AVSS}}$ ideal functionality (cf. Figure 6). Let $\mathcal{S}_{\text{AVSS}}$ be corresponding simulator. $\mathcal{S}_{\text{AVSS}}$ also uses the polynomial commitment scheme from Figure 1 to simulate $\mathcal{A}$'s view. We summarize $\mathcal{S}_{\text{AVSS}}$ in Figure 7, and prove the following.

Lemma 2 (Asynchronous VSS Secrecy)**.** *$\mathcal{A}$'s view in its interaction with $\mathcal{S}_{\text{AVSS}}$ is identically distributed to its view in the real protocol.*

Functionality $\mathcal{F}_{\text{VSS}}$

Parameters: Maximum number of malicious nodes t, the total number of nodes $n \geq 2t+1$. Let $\mathbb{G}$ be an elliptic curve group of order q with scalar field $\mathbb{F}$.

1) Wait for secret s from the dealer.
2) Wait for $\mathcal{C}$ with $|\mathcal{C}| \leq t$, the set of nodes $\mathcal{A}$ will corrupt.
3) Compute (n, t) Shamir secret shares of s over the field $\mathbb{F}$. Let $s(x)$ be the degree t polynomial with $s = s(0)$.
4) Send $s(j)$ to each honest node j. Send $\{s(i)\}_{i\in\mathcal{C}}$ to $\mathcal{A}$.

Figure 4: Functionality for the Sharing phase of synchronous VSS.

Inputs. $\mathcal{C}$, $\mathbb{F}$ and $\mathbb{G}$. **Notation.** Let $\mathcal{H} = [n] \setminus \mathcal{C}$

1) Sample signing and public key $(\mathsf{sk}_j, \mathsf{pk}_j)$ for each $j \in \mathcal{H}$. Send the public keys to $\mathcal{A}$.
2) Send $\mathcal{C}$ to $\mathcal{F}_{\text{VSS}}$ and receive $\{s(i)\}_{i\in\mathcal{C}}$.
3) Sample uniformly random generators $g, h \leftarrow \mathbb{G}$.
4) Sample a polynomial $\hat{s}(\cdot)$ of degree t such that $\hat{s}(i) = s(i)$ for each $i \in \mathcal{C}$. Additionally, sample a uniform random polynomial $\hat{r}(\cdot)$ of degree t.
5) Compute the commitment $\boldsymbol{v} = [v_1, v_2 \ldots, v_n]$ where $v_i = g^{\hat{s}(i)} h^{r(i)}$ for each $i \in [n]$.
6) Simulate the dealer by sending $\boldsymbol{v} = [v_1, v_2, \ldots, v_n]$ as the polynomial commitment. Participate in the rest of the protocol on behalf of the honest parties.

Figure 5: Synchronous VSS simulator $\mathcal{S}_{\text{VSS}}$

Functionality $\mathcal{F}_{\text{AVSS}}$

Parameters: Maximum number of malicious nodes t, the total number of nodes $n \geq 3t+1$. Let $\mathbb{G}$ be an elliptic curve group of order q with scalar field $\mathbb{F}$.

1) Wait for secret s from the dealer.
2) Wait for $\mathcal{C}$ and $\mathcal{H}_R$ from $\mathcal{A}$. Here $\mathcal{C}$ is the set of nodes $\mathcal{A}$ will corrupt and $\mathcal{H}_R$ is the set of honest nodes whose shares the the functionality will reveal. Check that $|\mathcal{C}| \leq t$, $|\mathcal{C} \cup \mathcal{H}_R| \leq 2t$. Let $\mathcal{C}_0 = \mathcal{C} \cup \mathcal{H}_R$.
3) Compute $(n, 2t)$ Shamir secret shares of s over the field $\mathbb{F}$. Let $s(x)$ be the degree $2t$ polynomial with $s = s(0)$.
4) Send $s(j)$ to each honest node j. Send $\{s(i)\}_{i \in \mathcal{C}_0}$ to $\mathcal{A}$.

Figure 6: Functionality for the Sharing phase of our AVSS.

Inputs. $\mathcal{C}, \mathcal{H}_R$, $\mathbb{F}$ and $\mathbb{G}$. **Notation.** Let $\mathcal{H} = [n] \setminus \mathcal{C}$ and let $\mathcal{C}_0 = \mathcal{C} \cup \mathcal{H}_R$.

1) Sample signing and public key $(\mathsf{sk}_j, \mathsf{pk}_j)$ for each $j \in \mathcal{H}$. Send the public keys to $\mathcal{A}$.
2) Send $(\mathcal{C}, \mathcal{H}_R)$ to $\mathcal{F}_{\text{AVSS}}$ and receive $\{s(i)\}_{i \in \mathcal{C}_0}$.
3) Sample uniformly random generators $g, h \leftarrow \mathbb{G}$.
4) Sample a polynomial $\hat{s}(\cdot)$ of degree $2t$ such that $\hat{s}(i) = s(i)$ for each $i \in \mathcal{C}_0$. Additionally, sample a uniform random polynomial $\hat{r}(\cdot)$ of degree $2t$.
5) Compute the commitment $\boldsymbol{v} = [v_1, v_2 \ldots, v_n]$ where $v_i = g^{\hat{s}(i)} h^{r(i)}$ for each $i \in [n]$.
6) Simulate the dealer by sending $\boldsymbol{v} = [v_1, v_2, \ldots, v_n]$ as the polynomial commitment. Participate in the rest of the protocol on behalf of the honest parties.

Figure 7: Asynchronous VSS simulator $\mathcal{S}_{\text{AVSS}}$

Proof. Follows using a similar argument as the proof of Lemma 1. □

Secrecy of the dual-thresold AVSS. We prove Secrecy of our dual-threshold AVSS scheme with respect to $\mathcal{F}_{\text{DtAVSS}}$ ideal functionality (cf. Figure 8). Let $\mathcal{S}_{\text{DtVSS}}$ be corresponding simulator. $\mathcal{S}_{\text{DtVSS}}$ also uses Pedersens' polynomial commitment from Figure 1 and the VE scheme from Appendix B to simulate $\mathcal{A}$'s view. We summarize $\mathcal{S}_{\text{DtVSS}}$ in Figure 9, and prove the following.

Lemma 3 (Dual-threshold AVSS)**.** *$\mathcal{A}$'s view in its interaction with $\mathcal{S}_{\text{DtVSS}}$ is computationally indistinguishable from its view in the real protocol.*

Proof. We will prove this using a sequence of hybrids using the verifiable encryption scheme from Appendix B.

Hybrid 0. This corresponds to the real-world execution.

Hybrid 1. Same as Hybrid 0, except we will change the NIZK proof of correct sharing of the VE scheme with a simulated proof. Hybrid 1 is indistinguishable from Hybrid 0 due to the zero-knowledge property of the NIZK scheme.

Hybrid 2 to Hybrid $k+1$. Without loss of generality let $\mathcal{H}_E = 1, 2, \ldots, k$. Hybrid $i+1$ for any $i \in [1, k]$ is the same as Hybrid i, except it swaps out VE of $s(i)$ and with VE of $\hat{s}(i)$. For each $i \in [1, k]$, Hybrid $i+1$ is indistinguishable from Hybrid i due to the CPA security of the VE scheme.

Functionality $\mathcal{F}_{\text{DtAVSS}}$

Parameters: The maximum number of malicious nodes t, the total number of nodes $n \geq 3t+1$, and maximum coalition size $\ell \in [t, n-t)$. Let $\mathbb{G}$ be an elliptic curve group of order q with scalar field $\mathbb{F}$.

1) Wait for secret s from the dealer.
2) Wait for $\mathcal{C}$, $\mathcal{H}_R$, and $\mathcal{H}_C$ from $\mathcal{A}$. Here $\mathcal{C}$ is the set of nodes $\mathcal{A}$ will corrupt and $\mathcal{H}_R$ is the set of honest nodes whose shares the the functionality will reveal. Also, let $\mathcal{H}_C$ is the set of honest nodes who will collude with $\mathcal{A}$ to learn the secret. Let $\mathcal{C}_0 = \mathcal{C} \cup \mathcal{H}_C \cup \mathcal{H}_R$
3) Assert that $|\mathcal{C}| \leq t$, $|\mathcal{C} \cup \mathcal{H}_C| \leq \ell$, and $|\mathcal{C} \cup \mathcal{H}_C \cup \mathcal{H}_R| \leq 2t$.
4) Compute $(n, 2t)$ Shamir secret shares of s over the field $\mathbb{F}$. Let $s(x)$ be the degree $2t$ polynomial with $s = s(0)$.
5) Send $s(j)$ to each honest node j. Send $\{s(i)\}_{i \in \mathcal{C}_0}$ to $\mathcal{A}$.

Figure 8: Dual-threshold AVSS functionality.

Inputs. $\mathcal{C}, \mathcal{H}_C, \mathcal{H}_R$, $\mathbb{F}$ and $\mathbb{G}$.
Notation. Let $\mathcal{H} = [n] \setminus \mathcal{C}$ and let $\mathcal{C}_0 = \mathcal{C} \cup \mathcal{H}_R \cup \mathcal{H}_C$.

1) Sample signing and public key $(\mathsf{sk}_j, \mathsf{pk}_j)$ for each $j \in \mathcal{H}$. Send the public keys of all nodes to $\mathcal{A}$. Additionally, send sk_j for each $j \in \mathcal{H}_C$ to $\mathcal{A}$.
2) Send $(\mathcal{C}, \mathcal{H}_R, \mathcal{H}_C)$ to $\mathcal{F}_{\text{DtAVSS}}$ and receive $\{s(i)\}_{i \in \mathcal{C}_0}$.
3) Sample uniformly random generators $g, h \leftarrow \mathbb{G}$.
4) Sample a polynomial $\hat{s}(\cdot)$ of degree $2t$ such that $\hat{s}(i) = s(i)$ for each $i \in \mathcal{C}_0$. Additionally, sample a uniform random polynomial $\hat{r}(\cdot)$ of degree $2t$.
5) Compute the commitment $\boldsymbol{v} = [v_1, v_2 \ldots, v_n]$ where $v_i = g^{\hat{s}(i)} h^{r(i)}$ for each $i \in [n]$.
6) Simulate the dealer by sending $\boldsymbol{v} = [v_1, v_2, \ldots, v_n]$ as the polynomial commitment. Simulate the dual-threshold VSS protocol on behalf of honest nodes up until receiving $2t+1$ signed acknowledgments.
7) Let $\mathcal{H}_E$ be the set of nodes whose shares $\mathcal{S}_{\text{DtVSS}}$ will verifiably encrypt. For each node $j \in \mathcal{H}_E$, use $\hat{s}(j)$ and $\hat{r}(j)$ as inputs to the VE scheme.

Figure 9: Dual-threshold AVSS simulator $\mathcal{S}_{\text{DtVSS}}$

Hybrid $k+2$ to Hybrid $2k+1$. Hybrid $k+i+1$ for any $i \in [1, k]$ is the same as Hybrid $k+i$, except it swaps out the encryption of $r(i)$ and with encryption of $\hat{r}(i)$. For each $i \in [1, k]$, Hybrid $k+i+1$ is indistinguishable from Hybrid $k+i$ due to the CPA security of the ElGamal encryption scheme

Hybrid $2k+2$. Same as Hybrid $2k+1$, except change the Pedersen commitment $\{g^{s(i)} h^{r(i)}\}_{i \in [n]}$ to $\{g^{\hat{s}(i)} h^{\hat{r}(i)}\}_{i \in [n]}$. Using a similar argument as Proof of Lemma 1, Hybrid $2k+2$ is identically distributed to Hybrid $2k+1$.

Hybrid $2k+3$. Same as Hybrid $2k+2$, except we will change the simulated NIZK proof of correct sharing of the VE scheme with a real NIZK proof. Hybrid $2k+3$ is indistinguishable from Hybrid $2k+2$ due to the zero-knowledge property of the NIZK scheme. Moreover, Hybrid $2k+3$ is the simulated transcript. □