

Haven++: Batched and Packed Dual-Threshold Asynchronous Complete Secret Sharing with Applications

Nicolas Alhaddad¹, Mayank Varia¹ and Ziling Yang²

¹ Boston University, United States

² University of Illinois Urbana-Champaign, United States

Abstract. Asynchronous complete secret sharing (ACSS) is a foundational primitive in the design of distributed algorithms and cryptosystems that require confidentiality. ACSS permits a dealer to distribute a secret to a collection of n servers so that everyone holds shares of a polynomial containing the dealer’s secret.

This work contributes a new ACSS protocol, called HAVEN++, that uses packing and batching to make asymptotic and concrete advances in the design and application of ACSS for large secrets. HAVEN++ allows the dealer to pack multiple secrets in a single sharing phase, and to reconstruct either one or all of them later. For even larger secrets, we contribute a batching technique to amortize the cost of proof generation and verification across multiple invocations of our protocol.

The result is an asymptotic improvement in the worst-case amortized communication and computation complexity, both for ACSS itself and for its application to asynchronous distributed key generation. Our ADKG based on HAVEN++ achieves, for the first time, an optimal worst case amortized communication complexity of $O(\kappa n)$ without a trusted setup. To show the practicality of HAVEN++, we implement it and find that it outperforms the work of Yurek et al. (NDSS 2022) by more than an order of magnitude when there are malicious, faulty parties.

1 Introduction

Many cryptographic protocols involving multiple parties begin with the distribution of data related to the parties’ cryptographic secrets. This is the first step in group secure messaging to agree upon shared symmetric keys, in threshold cryptosystems to distribute each party’s public key material to the group [CKLS02], in secure multi-party computation protocols to distribute correlated randomness during preprocessing [AJM⁺23a], and more. Rather than broadcasting an (ephemeral) secret from one party to the rest, instead these protocols require each party to disseminate *shares* of their ephemeral secrets that are subsequently used to agree upon common state.

These applications can be constructed using an *asynchronous complete secret sharing* protocol, or ACSS. Concretely, ACSS is a distributed protocol that protects a secret that a *dealer* distributes among n parties. An ACSS protocol has a sharing phase in which a dealer distributes shares of her secret; later, in the reconstruction phase, the parties can collectively recover the secret. After the initial sharing phase is completed, the parties can leverage ACSS’ confidentiality, integrity, and availability guarantees to reliably use and compute over the disseminated secrets.

Specifically, an ACSS protocol among n parties provides three security guarantees:

E-mail: nhaddad@bu.edu (Nicolas Alhaddad), varia@bu.edu (Mayank Varia), zilingy2@illinois.edu (Ziling Yang)

Table 1: Our HAVEN++ ACSS protocol, compared with other AVSS or ACSS protocols with optimal resilience. A green background indicates the optimal setting for that cell. Many constructions support batching of sufficiently-large messages, including ours. The table only shows worst-case complexity. Here, n = the number of parties and κ is the security parameter.

Works	Comm complexity		limiting setup		number of rounds	dual threshold	crypto assumption
	amortized	batch size	no trust?	no PKI?			
Kokoris-Kogias et al. [KMS20]	$O(\kappa n^3)$	$O(n)$	✓	✗	4	✓	DL
HAVEN [AVZ21]	$O(\kappa n^2 \log(n))$	$O(1)$	✓	✓	3	✗	DL + RO
e-AVSS [BDK13]	$O(\kappa n^2)$	$O(1)$	✗	✓	3	✗	t-SDH [KZG10]
Das et al. [DXR21]	$O(\kappa n^2)$	$O(1)$	✓	✗	3	✗	DDH + RO
Shoup et al. [SS24]	$O(\kappa n^2)^\dagger$	$\Omega(n^2)$	✓	✓	7+	✗	RO
hbACSS2 [YLF ⁺ 22]	$O(\kappa n \log n)$	$O(n^2)$	✓	✗	7+	✗	DL + RO
Bingo [AJM ⁺ 23a]	$O(\kappa n)$	$O(n)$	✗	✓	7	✓	t-SDH [KZG10]
HAVEN++	$O(\kappa n)$	$O(n \log n)$	✓	✓	3	✓	DL + RO

Table 2: ADKG protocols proposed in this paper, compared with the prior state of the art. We show worst-case amortized word complexity, along with the smallest batch size required to reach this amortized cost. The work of Groth et al. [GS24] includes an optimistic path that achieves $O(n)$ word complexity.

Scheme	Amortized Word Complexity	Batch size	Crypto assumption	Setup	Reference
Low threshold ADKG	$O(\kappa n^4)$	1	DDH + RO	PKI	Kate et al. [KHG12]
	$O(\kappa n^3)$	1	DDH + RO	PKI	Das et al. [DXR21]
	$O(\kappa n^3)$	1	RO + SXDH [BGdMM05]	PKI	Abraham et al. [AJM ⁺ 23b]
	$O(\kappa n^2)^\dagger$	$\Omega(n^2)$	RO	None	Groth et al. [GS24]
	$O(\kappa n)$	$O(n)$	DL + RO	None	This work
High threshold ADKG	$O(\kappa n^4)$	1	DL + RO	None	Kokoris-Kogias et al. [KMS20]
	$O(\kappa n^3)$	1	DL + RO	PKI	Das et al. [DXKKR23]
	$O(\kappa n^3)$	1	t-SDH [KZG10]	Trusted	Abraham et al. [AJM ⁺ 23a]
	$O(\kappa n^2)$	$O(n)$	DL + RO	None	This work

- All honest parties possess a share of a common secret S after the sharing phase even in the presence of t malicious parties (potentially including the dealer).
- The secret S remains confidential if up to p parties attempt to reconstruct the secret (where $p \geq t$, and often $p = t$).
- Reconstruction correctly recovers S if the dealer is honest.

To avoid making assumptions about an upper bound on network latency, ACSS protocols ensure that all three security guarantees hold even if the network reorders or delays messages arbitrarily. The only assumption made is that messages between honest parties must eventually be delivered.

In an asynchronous setting with $n = 3t + 1$, it is challenging to ensure that all honest parties in ACSS receive shares. Imagine a scenario where a party seeks confirmation from all others regarding their received shares' consistency. An honest party can only wait for messages from up to $n - t$ parties to avoid indefinite waiting if t parties fail. Of these messages, t of them might be from the attacker. Consequently, the party can only trust that $n - 2t = t + 1$ responses are from honest parties, leaving the possibility of t honest parties without shares. Recent works have proposed three principal approaches to address this issue and ensure that all honest parties in ACSS possess shares of one or more common secrets.

1. Approaches that rely on reliable broadcast and publicly verifiable secret sharing, as seen in [DXR21]. This method uses public key encryption and zero-knowledge proofs. The dealer verifies that all ciphertexts are shares of the same polynomial and broadcasts both the ciphertexts and proofs to all parties. However, this method does

not scale because it grows linearly with the number of secrets (each secret requires the dealer to generate n proofs, and every party has to verify all of them).

2. Approaches based on two layered secret sharing as seen in [AVZ21, AJM⁺23a, KMS20]. These methods split the secret into shares and then split each share into sub-shares, such that any $t + 1$ sub-shares of a share i can reconstruct back share i . The idea is that if at least $t + 1$ honest parties possess consistent sub-shares of each share and agree on termination, then all honest parties will receive their shares by being handed the proper sub-shares from the $t + 1$ honest parties with proofs of consistency. For large secrets, and specifically for layered secret sharing that are based on bivariate polynomials [AJM⁺23a, KMS20], this approach benefits from using packed secret sharing. Even so, prior works require each party to receive and verify $O(n)$ proofs per secret, so this approach also isn't scalable.
3. Approaches that utilize Asynchronous Verifiable Information Dispersal (AVID) [CT05], public key encryption, and dispute resolution processes as per [SS24, YLF⁺22]. Here, shares are not sent directly; they are encrypted and dispersed via AVID with a proof of consistency with the original polynomial. Discrepancies in shares prompt complaints, with the aggrieved party revealing their key for others to verify the contested encrypted message. In optimistic scenarios, the dealer typically prepares $O(n)$ batched proofs, with each party verifying only their one assigned proof. In contrast, a pessimistic scenario with $O(n)$ complaints requires each honest party to engage in extensive verification and dispute resolution, leading to quadratic communication complexity and expensive runtime.

For real-world scenarios and for long-living systems, one might opt to go with method #3 because in the best-case scenario, this approach achieves practical and near-optimal amortized communication complexity. In such systems, a cheating dealer can be detected and subsequently barred from future participation. However, this method may not be feasible when party identities are either unknown or transient. Moreover, it complicates the integration of the ACSS protocol with other protocols and might necessitate the unnecessary retention of all exchanged messages in case of complaints.

An ongoing challenge is to design a practical ACSS in the worst case (when t malicious parties are present including the dealer). The ideal protocol would minimize amortized communication complexity while simultaneously ensuring a practical runtime. Additionally, an advantageous feature of such a system would be its independence from any trusted setup. This is the focus of our work.

1.1 Related Work

AVSS and ACSS. The problem of asynchronous verifiable secret sharing, or AVSS, dates back to at least the 1990s. Early works (e.g., [CR93a, BCG93, Can96]) showed the feasibility of AVSS with unconditional security (i.e., without any cryptographic assumptions), but they had a large communication complexity. In the early 2000s, Cachin et al. [CKLS02] made two important advances: designing an AVSS with optimal message complexity against $t < n/3$ malicious parties, and a dual-threshold AVSS where correctness holds against $t < n/4$ parties yet secrecy holds against $t < n/2$ parties. However, both constructions suffer from suboptimal $O(\kappa n^3)$ communication complexity.

The last few years have seen a renaissance of work in this field, with several works that improve asymptotic and concrete performance, reduce computational assumptions and the need for trusted setup, and increase the thresholds for correctness and secrecy. Many recent works (though not all) leverage recent innovations in the design of polynomial commitments, vector commitments, succinct zero-knowledge proofs, Verifiable Information Dispersal and Reliable Broadcast.

Kokoris-Kogias et al. [KMS20] constructed the first “high-threshold” AVSS protocols that maintain secrecy for up to $p < 2n/3$ parties at a cost of $O(\kappa n^3)$ using polynomial commitments of linear size. Subsequently, Alhaddad et al. [AVZ21] showed how to reduce that to $O(\kappa n^2)$ using constant size polynomial commitments and vector commitments. They also showed how to achieve the first AVSS with worst case amortized communication complexity of $O(\kappa n)$ using erasure coding techniques and encryption. However, since the shares were part of ciphertexts, they weren’t linear and couldn’t be used directly for multiparty computation. It was left as an open problem whether one can achieve an amortized communication complexity $O(\kappa n)$ with linear shares. The subsequent work of Das et al. [DXR21] showed how to achieve a high threshold AVSS using Public Verifiable Secret Sharing (PVSS) and zero knowledge proofs. They leveraged the advancement in reliable broadcast to broadcast a linear polynomial commitment (in the number of parties) and an encryption of all shares, with a proof of consistency to all parties in $O(\kappa n^2)$ communication complexity. However, their work requires a PKI and doesn’t scale to large secrets.

Afterward, the hbACSS protocol of Yurek et al. [YLF⁺22] showed how to achieve $O(\kappa n \log n)$ in the worst case, using verifiable information dispersal in their construction hbACSS2. This work highlighted the need for *complete* secret sharing rather than AVSS. In contrast to AVSS, complete secret sharing guarantees that all honest parties have shares. However, their most concretely efficient hbACSS0 variant suffered from sub-optimal $O(\kappa n^2)$ worst case amortized communication complexity, and their work relies on a public key infrastructure.

Recently, the work of Shoup et al. [SS24] focused on building practical ACSS with large secrets, with the aim of using them for Schnorr Signatures. However, their construction requires many rounds of communication and has a worst-case communication complexity of $O(\kappa n^2)$. Very recently, the Bingo protocol of Abraham et al. [AJM⁺23a] achieved an optimal amortized linear complexity with linear secrets. Their insight was to update the construction of Haven [AVZ21] to use a bivariate polynomial. That way, they can use packed Shamir secret sharing to pack $t + 1$ secrets and make sure that every party has exactly one share of every packed secret. However, Bingo requires a trusted setup, only works with KZG polynomial commitments [KZG10], doesn’t support batching, has seven rounds and relies on a black box reliable broadcast that requires online error correcting code [DXR21].

Concurrent work has explored linear communication complexity in the asynchronous setting. [AAPP24] focuses on AMPC with perfect security and optimal resilience; they introduce a primitive called *Weak-Binding Secret Sharing* with $t < n/4$. This primitive is distinct from ACSS and does not guarantee full reconstruction of the dealer’s shares. [JLS24] considers ACSS in the information-theoretic setting for $t < n/3$ and achieves linear communication complexity. However, their communication complexity is on the order of $\mathcal{O}(Nn\kappa + n^{12}\kappa^2)$, where N denotes the number of secrets, n denotes the number of parties, and κ denotes the security parameter) and it also requires many communication rounds. This renders their approach impractical even for small n and κ . In fact, [JLS24] needs hundreds of millions of secrets to benefit from its linear communication complexity (even for small n such as 4).

ADKG. Asynchronous distributed key generation enables robust, fault-tolerant communication over an unreliable network. As such, it is a valuable building block toward many distributed protocols, including those used for threshold cryptography and blockchains/state machine replication. Several of the works cited above also consider ADKG, as do standalone works like Das et al. [DXKKR23].

Polynomial and vector commitments. Our construction uses polynomial and vector commitments as a building block to achieve consensus on the bivariate polynomial ϕ . For our polynomial commitment scheme, we consider Bulletproofs [BBB⁺18], which have transparent setup but require log-sized proofs. Our HAVEN++ protocol uses these commitments in a black-box manner, so we could alternatively use any polynomial commitment scheme that is deterministic and homomorphic (e.g., [KZG10, BG18, BFS20]). We also use the related idea of vector commitments, which were initially introduced by Libert-Yung [LY10] and Catalano-Fiore [CF13].

1.2 Our Contributions

This work presents asymptotic and tangible improvements in the design and application of Asynchronous Complete Secret Sharing (ACSS) for large secrets. Specifically, we build upon method #2 (layered secret sharing) and introduce HAVEN++: the first practical ACSS that achieves an $O(\kappa n)$ worst-case amortized communication complexity. Compared to Yurek et al. [YLF⁺22] that uses method #3, our benchmarks show a $3\times$ performance improvement when everyone is honest, and orders of magnitude improvement with malicious parties.

Efficient Batched and Packed Two Layered ACSS. In §3, we contribute a new ACSS construction called HAVEN++. This protocol has optimal amortized communication complexity and rounds, and it avoids the need for a trusted setup or PKI when combined with bulletproofs [BBB⁺18]. HAVEN++ is also a *dual-threshold* scheme, which means that it has two different reconstruction protocols with different thresholds: either $t + 1$ parties or $p + 1$ parties are required, for any choice of $p \in (t, n - t)$. HAVEN++ supports reconstruction of all packed secrets under the higher threshold, and reconstruction of a single secret with the smaller threshold.

To improve efficiency on large secrets, we leverage the dual-threshold property in order to *pack* $p - t + 1$ secrets inside of a single ACSS invocation. Additionally, we show how to *batch* multiple invocations of HAVEN++ and only disseminate n^2 proofs independently of how many invocations are being batched. This further improves communication complexity by a $\log(n)$ factor. See Table 1 for a detailed comparison between HAVEN++ and prior works.

Implementation and experimental evaluation. To demonstrate that these techniques improve concrete efficiency in addition to asymptotic efficiency, in §4 we implement our constructions on top of the open-source framework of hbACSS [YLF⁺22], which is the only prior work with near optimal amortized communication complexity that has an open-source implementation. Our implementation reuses all of their low-level field arithmetic and crypto primitives in order to provide an apples-to-apples comparison.

Our experiments show that HAVEN++ substantially reduces computation time compared to hbACSS by a factor of at least $3\times$ when all parties are honest, and often by more than an order of magnitude when some parties are malicious as shown in Figure 1. The exact performance improvement depends on the number of parties and the batch size; see Figures 3-5 for details.

Application to ADKG. In §5, we contribute a new, non-black-box application of our batched and packed ACSS protocol HAVEN++ to improve the asymptotic complexity of *asynchronous distributed key generation*, or ADKG. Our ADKG achieves for the first time an optimal $O(\kappa n)$ communication complexity in the worst case (where κ denotes the security parameter), and it also does not require trusted setup or a PKI. This new result is only possible because our ACSS allows parties to have independent shares of each packed secret share. An ADKG protocol allows a collection of parties to agree on a public key




Figure 1: The computational cost per party for HAVEN++, compared against hbACSS0 and hbACSS2 from Yurek et al. [YLF⁺22] (lower is better). For sufficiently many parties n , HAVEN++ wins by a factor of $3\times$ or more, for any batch size.

and each possess a share of the corresponding secret key. It is a critical component of distributed protocols like Byzantine agreement [CKS05] and randomness beacons [HMW18], and of cryptographic protocols like threshold signatures [Bol03, GJKR07] and multiparty computation [YLF⁺22, GS24]. Our ADKG protocol based on HAVEN++ is the first to achieve optimal amortized word complexity even in the worst case; see Table 2 for details.

Application to AMPC. In §5, we contribute a new protocol to generate pre-processing material for secure multi-party computation. Specifically, we provide the first dual secret sharing (also known as random double sharing) in the asynchronous setting on top of HAVEN++ with a worst-case amortized word complexity of $O(\kappa n^2)$ with $t < n/3$. It was only previously known how to achieve similar results in the synchronous setting [DN07].

1.3 Technical Overview of Our Constructions

In this section, we provide a high-level overview of the main techniques used in our ACSS and ADKG constructions.

ACSS construction, against a DoS adversary. We describe our HAVEN++ construction here and in Figure 2. For simplicity, we begin by describing the protocol in the (unrealistic) scenario that the t faulty parties will be truthful in any message that they send, i.e., they are only allowed to drop messages.

The HAVEN++ construction incorporates several design elements of Haven [AVZ21] (hence the name), and it also uses a bivariate polynomial ϕ in a related (but not identical) manner as in prior dual-threshold ACSS protocols [KMS20, AJM⁺23a]. Concretely, ϕ is a polynomial of degree p in the horizontal direction and t in the vertical direction. To construct ϕ , the dealer packs $p - t + 1$ secrets on the points of the x -axis to the left of the y -axis; that is, at locations $(-k + 1, 0)$ for $k \in [1, p - t + 1]$. The dealer then randomly chooses sufficiently many points (as shown in the pink shaded region in Fig. 2) to uniquely determine a *degree* (p, t) bivariate polynomial. The dealer commits to ϕ in a manner that we specify later, and reliably broadcasts this commitment during the 3 rounds of the sharing phase.




Figure 2: The 2d input space to a bivariate polynomial $\phi(x, y)$ used in our ACSS construction; this example uses $t = 2$, $n = 7$, and $p = 3$. The dealer creates a bivariate polynomial ϕ by packing $p - t + 1 = 2$ secrets s_1 and s_2 on the x -axis (shown in red) and randomly choosing the remaining points in the pink shaded region. In the sharing phase, the parties receive points in the yellow region: specifically, party i learns the row polynomial $\phi(i, -)$ (of degree p , shown in brown) and column polynomial $\phi(-, i)$ (of degree t , shown in blue). Then, the parties can reconstruct all secrets (via the x -axis) or one secret (via the corresponding column, shown in orange).

In the sharing phase, each party learns both the i^{th} row and column of ϕ ; as a result, the parties collectively learn the points in the yellow shaded region of Fig. 2 (which is disjoint from the pink region). The sharing phase proceeds in three rounds as follows:

1. The dealer sends each party P_i the i^{th} row polynomial $\phi(i, -)$, as shown in brown in Fig. 2. Looking ahead, the parties will reach agreement if at least $n - 2t$ honest parties receive row polynomials that are consistent with ϕ , but the t malicious parties and up to t honest parties might receive nothing from the dealer.
2. Any honest party i that received a row polynomial in step 1 now sends the point $\phi(i, j)$ to party j . As a result, *every* honest party receives enough points to interpolate their column polynomial (shown in blue in Fig. 2).
3. Each party i sends the point $\phi(j, i)$ to party j . After this round, every honest party receives enough points to interpolate their row polynomial, whether or not they received it previously from the dealer in step 1 above.

This construction has $O(\kappa n^2)$ word complexity, where κ denotes both the security parameter and field size of an individual secret. After the sharing phase is complete: $p + 1$ honest parties can reconstruct all secrets by revealing the points on the x -axis, or $t + 1$ honest parties can reconstruct a single secret by revealing the points on the corresponding column (shown in orange in Fig. 2).

ACSS construction, against a malicious adversary. To account for malicious parties (including the dealer), we use polynomial and vector commitments so that everyone can prove that they are sending points on the polynomial ϕ . We do not use bivariate polynomial commitments directly in HAVEN++; instead, the dealer produces a polynomial

commitment to each *column* polynomial $\hat{\phi}_i = \text{pCom}(\phi(-, i))$ and broadcasts a vector $\vec{v} = [\hat{\phi}_i]_{i \in [n]}$ of these commitments. Anyone who receives \vec{v} is eventually assured that these polynomial commitments collectively correspond to a bivariate polynomial because: (i) they can check directly that each commitment corresponds to some column polynomial of degree at most t , and (ii) collectively they know that at least $t + 1$ parties have checked that $t + 1$ rows are of degree p and by lemma 1 this vector commitment must have been for a bivariate polynomial commitment.

Whenever the parties send a point $\phi(i, j)$, they always attach a proof of inclusion on the corresponding column polynomial ϕ_j . We emphasize that even in steps 1 and 3 that involve sending points on a row, the corresponding proofs are still with respect to the respective column polynomials instead.

These proofs become the dominant cost in the HAVEN++ protocol. For this reason, the protocol benefits even more by batching $O(n)$ executions in parallel, which collectively contain a total of $O(n^2)$ secrets. We construct a batch proof in Algorithm 4 that efficiently proves correctness of all $O(n)$ row or column polynomials at once. Batching improves the amortized communication complexity by a factor of n , as shown in Table 1.

ADKG construction. Asynchronous distributed key generation allows n parties to agree on a public key such that each party holds one share of the corresponding secret key. In this work, we construct an ADKG protocol by combining n instances of ACSS (one per party) with a multi-valued Byzantine agreement or MVBA protocol. However, it’s worth noting that the way we combine the n instances of our ACSS is novel and relies on the dual threshold property of our ACSS. The full construction is shown in Figure 6 and Algorithm 5; we provide only a brief overview here.

The high-level idea of our construction is simple: each of the n parties acts as the dealer and disperses a bivariate polynomial, each packed with $t + 1$ secrets, and then we “mix and match” columns from everyone to form $t + 1$ new bivariate polynomials that the adversary does not fully know. This is only possible because our ACSS construction guarantees that every party (except the dealer) has exactly one share of every packed secret on every column being mixed.

One challenge here is that some parties might act maliciously as the dealer and disperse shares that will never reach agreement. This is where the multi-valued Byzantine agreement protocol comes in: it allows the n parties to ‘vote’ on which $n - t$ of the dispersed bivariate polynomials to use in the mix-and-match stage. Confidentiality is maintained even if the adversary knows t of these polynomials.

Another challenge is that ADKG requires polynomial arithmetic “in the exponent,” since it is used to determine a public/secret key pair in a group where the discrete logarithm is assumed to be hard. Fortunately, all of our polynomial operations like evaluation and interpolation are linear, so they can be performed in the exponent. See §5 for details.

2 Preliminaries

In this section, we introduce our model and some of building blocks that are going to be used in this paper for our constructions.

2.1 Network and Adversary Model

We study a network of n parties, each pair interconnected via an authenticated and private channel. A malicious adversary, denoted by \mathcal{A} , can corrupt up to t parties. Our network is asynchronous: \mathcal{A} can delay but must eventually deliver messages between honest parties. To evaluate the latency of asynchronous protocols, we adopt the standard concept of asynchronous rounds [CR93a]. A protocol operates in R asynchronous rounds

if its execution time is at most R times the longest message delay between honest parties during the protocol run.

2.2 Definitions and Building Blocks

Dual-threshold Verifiable Secret Sharing. Secret sharing is a method where a secret is divided into shares in such a way that only specific subsets of shares can reconstruct the original secret. Verifiable Secret Sharing (VSS) enhances this by allowing participants to verify that their shares reconstruct to the same secret, even in the presence of a bad dealer. In verifiable secret sharing, an attacker is allowed to control t out of n parties, and if the attacker doesn't corrupt the dealer then the attacker learns nothing about the secret being shared. Moreover, any $t + 1$ parties can reconstruct the secret. VSS schemes have two phases: the sharing phase in which a special party called the *dealer* disperses a secret s among the n parties, and the reconstruction phase in which the n parties collaborate to recover s .

Dual-threshold verifiable secret sharing adds another degree of flexibility: the number of shares p that is insufficient to reconstruct the message is not restricted to t but can be higher. Additionally, we can use packed secret sharing to store multiple secrets within set of shares, and then reconstruct either all secrets or a specific secret (say, in location k). The following definition is adapted from Abraham et al. [AJM⁺23a] to incorporate the dual threshold property.

Definition 1 (AVSS [AJM⁺23a]). A dual-threshold asynchronous verifiable secret sharing protocol contains three protocols **Share**, **Reconstruct**, and **Reconstruct(k)** that satisfy the following three properties, even against an adversary who controls t malicious parties.

- *Termination:* If the dealer is honest, then all honest parties will complete **Share**. Also if one honest party completes **Share**, then all honest parties will. Finally, if all honest parties complete **Share** and invoke **Reconstruct** or **Reconstruct(k)**, then they all will complete reconstruction.
- *Correctness:* All honest parties who complete the partial reconstruction protocol **Reconstruct(k)** should agree on the same secret. The same is true for **Reconstruct**, and moreover it should produce the same secret at location k . Finally, this secret should be the same as the one initially used by the dealer in the **Share** protocol, if the dealer was honest.
- *Secrecy:* An adversary should not be able to learn anything about the k^{th} secret until the point at which some honest party invokes **Reconstruct(k)**. For the full reconstruction protocol **Reconstruct**, an adversary should not be able to learn anything even if it participates in the protocol with up to $p - t$ honest parties.

An asynchronous *complete* secret sharing protocol, or ACSS, additionally satisfies the completeness property.

Definition 2 (Completeness [DXR21]). If some honest party completes **Share**, then there exists a degree- t polynomial p such that $p(0) = s$ and each honest party i will eventually hold a share $s_i = p(i)$. Moreover, when the dealer is honest, then s is the secret that it initially shared.

Reed-Solomon Error Correcting Code. Reed-Solomon error-correcting codes play a fundamental role in state-of-the-art reliable broadcast protocols, verifiable secret sharing schemes, and information-theoretic multi-party computation protocols. In adversarial settings, they empower honest parties to reconstruct the dealer's secret (or the plaintext message in reliable broadcasts) even amidst failures.

Formally, an (m, n) error-correcting code is defined by a pair of algorithms (ECCEnc, ECCDec). The encoding algorithm, denoted by $\text{ECCEnc}(M, m, k)$, ingests a message M comprised of k symbols, interprets it as a polynomial of degree $k - 1$, and emits m evaluations of said polynomial. Conversely, the decoding algorithm, represented by $\text{ECCDec}(k, r, T)$, receives a set of symbols T —some potentially erroneous—and produces a polynomial of degree $k - 1$, or equivalently, k symbols. This is achieved by amending up to r errors (incorrect symbols) within T . It is a well-established fact [MS77] that ECCDec can rectify up to r errors in T and yield the initial message given that $|T| \geq k + 2r$.

Note that in this paper we will only formally call the decoding algorithm. The encoding algorithm will not be called.

Online Error Correcting. Online error correction (OEC) refers to a set of techniques where error detection and correction are performed as data is transmitted or processed. It was first used by Canetti et al. [BCG93] for doing verifiable secret sharing and was later used for asynchronous reliable broadcast [DXR21], [ADD⁺22a] and asynchronous verifiable information dispersal [ADD⁺22b]. In contrast to traditional error-correcting codes that first gather all data before starting the correction procedure, online methods operate as data streams in. This capability is particularly useful to honest parties that are trying to filter out bad shares as they are receiving them. For our use case, when $n = 3t + 1$ and when all honest parties have evaluation of a polynomial of degree t , it allows a receiver to recover the polynomial of degree t after hearing from $2t + 1$ honest parties, even though t parties might send bad evaluations. We refer the reader to the original paper of Canetti et al. [BCG93] for full details.

2.3 Polynomial and Vector Commitments

We consider polynomial commitment scheme that allows a prover to commit non interactively to a polynomial such that, later, the prover can be asked to open the commitment at any particular point and reveal the corresponding value. We follow the same definition as HAVEN; like them, we require the polynomial commitment to be deterministic and additively homomorphic. We re-state the definition for convenience with the addition of three extra optional algorithms AggBatchProof , AggBatchVerify and BatchProof that are all used strictly for the batched variant of our algorithm described in §3.3.

Definition 3. A *polynomial commitment scheme* \mathcal{P} comprises four algorithms Setup , pCom , Eval , Verify and four optional algorithms Hom , AggBatchProof , AggBatchVerify , BatchProof that act as follows:

- $\text{Setup}(1^\kappa, \mathbb{F}, D) \rightarrow \mathbf{pp}$ is given a security parameter κ , a finite field \mathbb{F} , and an upper bound D on the degree of any polynomial to be committed. It generates public parameters \mathbf{pp} that are required for all subsequent operations.
- $\text{pCom}(\mathbf{pp}, \phi(x), d) \rightarrow \hat{\phi}$ is given a polynomial $\phi(x) \in \mathbb{F}[x]$ of degree $d \leq D$. It outputs a commitment string $\hat{\phi}$ (throughout this work, we use the hat notation to denote a commitment to a polynomial).
- $\text{Eval}(\mathbf{pp}, \phi, i) \rightarrow \langle i, \phi(i), w \rangle$ is given a polynomial ϕ as well as an index $i \in \mathbb{F}$. It outputs a 3-tuple containing i , the evaluation $\phi(i)$, and witness string w_i .
- $\text{Verify}(\mathbf{pp}, \hat{\phi}, y, d) \rightarrow \text{True/False}$ takes as input a commitment $\hat{\phi}$, a 3-tuple $y = \langle i, j, w \rangle$, and a degree d . It outputs a Boolean value indicating whether the provided evaluation j corresponds to the committed polynomial $\hat{\phi}$ evaluated at the point determined by i . Specifically, Verify checks whether $\phi(i) = j$. If the evaluation is consistent with the commitment and the degree constraint, the function returns True ; otherwise, it returns False .
- $\text{Hom}(\mathbf{pp}, \hat{\phi}_1, \hat{\phi}_2, a) \rightarrow \widehat{\phi_1 + a\phi_2}$ takes in commitments to two polynomials ϕ_1 and ϕ_2 of degree at most D , as well as a field element $a \in \mathbb{F}$. Outputs the commitment

- $\text{pCom}(\text{pp}, \phi, \max\{d_1, d_2\})$ to the polynomial $\phi = \phi_1 + a\phi_2$.
- $\text{BatchProof}(\text{pp}, \phi, n, d)$ is given a polynomial ϕ of degree d , where n is a positive integer. It outputs n 3-tuples $\langle i, \phi(i), w_i \rangle$ where $1 \leq i \leq n$ and w_i is a proof for $\phi(i)$. Each such proof can be verified with Verify .
- $\text{AggBatchProof}(\text{pp}, [\phi_1, \dots, \phi_\beta], n, d)$ is given a list of polynomials $\phi_1 \dots \phi_\beta$ of the same degree d , and a positive integer n . It outputs n different 3-tuples $\langle i, (\phi_1(i), \dots, \phi_\beta(i)), w_i \rangle$ for $1 \leq i \leq n$. That is, each tuple contains evaluations of all β polynomials together with a single proof string w_i .
- $\text{AggBatchVerify}(\text{pp}, [\hat{\phi}_1, \dots, \hat{\phi}_\beta], y, d)$ takes as input a list of polynomial commitments $[\hat{\phi}_1, \dots, \hat{\phi}_\beta]$, a 3-tuple $y = \langle i, (j_1 \dots j_\beta), w \rangle$, and a degree d . It outputs True if $\phi_z(i) = j_z \forall z \in \{1 \leq z \leq \beta\}$, and False otherwise.

We also use vector commitments in this work; these are also succinct commitments to a large set of data, but the data need not correspond to points on a polynomial. That is, vector commitments are cryptographic primitives that allow one to commit to an ordered sequence of values (or a vector) and later prove the value of a specific position in the vector without revealing any other information about the rest of the vector. Much like polynomial commitments, the commitment size is constant, not dependent on the size of the vector. For our implementation, we instantiate HAVEN++ with Merkle trees for vector commitments and Bulletproofs [BBB⁺18] for polynomial commitments. We note that for Bulletproofs the setup function is transparent and non-interactive meaning it does not require a trusted party or external secret information for its initialization. Instead, all the cryptographic parameters used in the commitment scheme are generated publicly and can be verified by anyone.

Defining vector commitments

This paper follows the convention of Alhaddad et al. [AVZ21] and restricts attention to commitment schemes that are *deterministic* and *homomorphic*. We leverage the deterministic property in the HAVEN++ construction so that vector and polynomial commitments do not need a separate “decommitment randomness” string that itself would need to be reliably dispersed.

This approach does introduce one security challenge: a deterministic commitment cannot be hiding when used to commit to an arbitrary secret, only for a randomly-chosen secret. This challenge is solvable through a simple blinding technique previously used by [BBB⁺18, BFS20, AVZ21, CFS17, BCC⁺16] among others. Concretely: rather than committing to an arbitrary secret directly, the committer can produce a hiding commitment to an ephemeral random secret vector (or secret polynomial), and then use the homomorphic property to construct a non-hiding proof of the correct opening of a linear combination of the desired secret and the ephemeral random secret. We omit the details here and refer interested readers to [AVZ21, §3.3] for more information.

Definition 4. A deterministic *vector commitment scheme*

$$\mathcal{V} = (\text{vSetup}, \text{vCom}, \text{vGen}, \text{vVerify})$$

comprises four algorithms that operate as follows:

- $\text{vSetup}(1^\kappa, U, L) \rightarrow \bar{\text{pp}}$ is given a security parameter κ , a set U , and a maximum vector length L . It generates public parameters $\bar{\text{pp}}$.
- $\text{vCom}(\bar{\text{pp}}, \vec{v}) \rightarrow C$ is given a vector $\vec{v} \in U^\ell$ where $\ell \leq L$. It outputs a commitment string C .
- $\text{vGen}(\bar{\text{pp}}, \vec{v}, i) \rightarrow \pi_i$ is given a vector \vec{v} and an index i . It outputs a proof string π_i .

- $\text{vVerify}(\bar{\text{pp}}, C, u_i, \pi) \rightarrow \text{True/False}$ takes as input a vector commitment C , an indexed element $u_i \in U$, and a proof string π . It outputs True if $u_i = \bar{v}[i]$ and π is a witness to this fact and False otherwise.

Security guarantees for commitments

We begin by discussing the security requirements for polynomial commitments. Just like any kind of commitment scheme, a polynomial commitment must satisfy correctness, binding, and hiding. We make two special requirements for polynomial commitments: first, we want the commitment to bind to a particular polynomial and to its (max) degree, and second, the hiding property only needs to hold for a random polynomial for the reason stated above. The specific definitions below are taken from Alhaddad et al. [AVZ21] since they also focus on the setting of polynomial commitments that are deterministic and homomorphic.

Definition 5 (Strong correctness). Let $\text{pp} \leftarrow \text{Setup}(1^\kappa, \mathbb{F}, D)$. For any polynomial $\phi(x) \in \mathbb{F}[x]$ of degree d with associated commitment $\hat{\phi} = \text{pCom}(\text{pp}, \phi, d)$:

- If $d \leq D$, then for any $i \in \mathbb{F}$ the output $y \leftarrow \text{Eval}(\text{pp}, \hat{\phi}, i)$ of evaluation is successfully verified by $\text{Verify}(\text{pp}, \hat{\phi}, y, d)$.
- If $d > D$, then no adversary can succeed with non-negligible probability at creating a commitment $\tilde{\phi}$ that is successfully verified at $d + 1$ randomly chosen indices.

Definition 6 (Evaluation binding). Let $\text{pp} \leftarrow \text{Setup}(1^\kappa, \mathbb{F}, D)$. For any PPT adversary $\mathcal{A}(\text{pp})$ that outputs a commitment $\tilde{\phi}$, a degree d , and two evaluations $y = \langle i, j, w \rangle$ and $y' = \langle i', j', w' \rangle$, there exists a negligible function $\varepsilon(\kappa)$ such that:

$$\Pr \left[(\tilde{\phi}, y, y', d) \leftarrow \mathcal{A}(\text{pp}) : i = i' \wedge j \neq j' \wedge \text{Verify}(\text{pp}, \tilde{\phi}, y, d) \wedge \text{Verify}(\text{pp}, \tilde{\phi}, y', d) \right] < \varepsilon(\kappa).$$

Definition 7 (Degree binding). Let $\text{pp} \leftarrow \text{Setup}(1^\kappa, \mathbb{F}, D)$. For any PPT adversary \mathcal{A} that outputs a polynomial ϕ of degree $\deg(\phi)$, evaluation \tilde{y} , and integer d , there exists a negligible function $\varepsilon(\kappa)$ such that:

$$\Pr \left[(\phi, \tilde{y}, d) \leftarrow \mathcal{A}(\text{pp}), \hat{\phi} = \text{pCom}(\text{pp}, \phi, \deg(\phi)) : \text{Verify}(\text{pp}, \hat{\phi}, \tilde{y}, d) \wedge \deg(\phi) > d \right] < \varepsilon(\kappa).$$

Definition 8 (Hiding for random polynomials). Let $\text{pp} \leftarrow \text{Setup}(1^\kappa, \mathbb{F}, D)$, d be an arbitrary integer less than D , and $I \subset \mathbb{F}$ be an arbitrary set of indices with $|I| \leq d$. Randomly choose a $\phi \leftarrow \mathbb{F}[x]$ of degree d and construct its commitment $\hat{\phi} = \text{pCom}(\text{pp}, \phi, d)$. For all PPT adversaries \mathcal{A} , there exists a negligible polynomial $\varepsilon(\kappa)$ such that:

$$\Pr \left[(x, y) \leftarrow \mathcal{A}(\text{pp}, \hat{\phi}, \{\text{Eval}(\text{pp}, \phi, i)\}_{i \in I}) : y = \phi(x) \wedge x \notin I \right] < \varepsilon(\kappa),$$

where the probability is taken over \mathcal{A} 's coins and the random choice of ϕ .

Finally, we discuss the security requirements for vector commitments. We omit a formal specification here because they are analogous to Definitions 5, 6, and 8 above, but with polynomial commitments and proofs replaced with vector commitments and proofs. That is: correctness requires that honestly-created opening proofs will verify, evaluation binding requires that the adversary cannot find two openings to the vector commitment at the same index i that will both verify, and hiding requires that there is a negligible probability that an adversary can produce an opening proof corresponding to a vector commitment for a random vector that the adversary was never given. There is no equivalent to degree binding for vector commitments; only evaluation binding is required.

2.4 Multi-Valued Byzantine Agreement

In this section, we provide a formal definition of multi-valued Byzantine agreement (MVBA). Looking ahead, we use this primitive as a building block in §5 to construct an asynchronous distributed key generation (ADKG) protocol.

MVBA was initially introduced by Cachin et al. [CKPS01]. It generalizes Byzantine agreement to allow for a message v that is more than one bit in length. Additionally, it introduced the ability for parties to check that the agreed-upon message satisfies some predicate Q , possibly when matched with some additional information w that is also decided during the protocol. Later, Abraham et al. [AMS19] to rule out trivial solutions in which parties always decide on some pre-defined externally valid value, they add another property to the MVBA, which they call quality. The quality property bounds the probability that the decision value was proposed by an honest party. We explicitly add the quality property and use the definition of Cachin et al. [CKPS01].

A precise definition follows.

Definition 9 (MVBA). A multi-valued Byzantine agreement scheme is an interactive protocol between n parties that satisfies the following five criteria.

- *Termination*: If all honest parties input a message that satisfies the predicate $Q(v, w)$ for some w , then all honest parties eventually output v .
- *Agreement*: If an honest party outputs v , then every honest party also terminates and outputs v .
- *External validity*: Every honest party that terminates decides v validated by w such that the predicate $Q(v, w)$ is true.
- *Integrity*: If all parties follow the protocol, and if an honest party decides v validated by w , then some party proposed v validated by w .
- *Quality*: The probability of terminating with a value v that was proposed by a correct replica is at least $1/2$.

3 Our ACSS Construction

In this section, we introduce our dual-threshold ACSS protocol, called HAVEN++ that achieves all the properties of a dual-threshold ACSS as shown in §3.2. We present the construction in two parts: first with packing of multiple secrets into a single bivariate polynomial, and then batching across multiple bivariate polynomials. For simplicity and without loss of generality we instantiate our protocol with $n = 3t + 1$ (optimal resilience) and with $p = 2t$. Note that our protocols are executed from the perspective of party i . When party i sends $\text{info}_{i,j}$ to party j , it provides the necessary information for party j to continue the protocol. Since all parties perform this operation, party i also receives information from other parties. Specifically, party i receives $\text{info}_{j,i}$ from each party j . In our notation, the first parameter of $\text{info}_{j,i}$ represents the sender (in this case, j), while the second parameter represents the receiver (in this case i).

3.1 Haven++ with Packing

HAVEN++ has two phases: a sharing phase in which the dealer distributes shares of her secret s , and a reconstruction phase in which the parties collectively reconstruct one or more secrets.

Algorithm 1 Sharing phase of HAVEN++, for server P_i and tag $ID.d$.

- 1: UPON RECEIVING ($ID.d, in, share, s_1 \dots s_b$): ▷ only if party is the dealer P_d
 - 2: Uniformly sample $\phi(X, Y)$ with degree $2t$ in X and t in Y such that $\phi(-k + 1, 0) = s_k, \forall k \in [1, b]$.
 - 3: **for** $i \in [1, n]$ **do**
 - 4: $\phi_i(Y) = \phi(i, Y)$
 - 5: $\hat{\phi}_i = \text{pCom}(\text{pp}, \phi_i(Y), t)$ ▷ Commit to every column polynomial
 - 6: **for** $i \in [1, n]$ **do**
 - 7: compute $\vec{y}_i = [\text{Eval}(\text{pp}, \phi_j(Y), i) \text{ for } j \in [1, n]]$ ▷ evaluate and create witnesses for every point on every column polynomial
 - 8: send “ $ID.d, send, \text{set}_i$ ” to party P_i , where $\text{set}_i = \{[\hat{\phi}_1 \dots \hat{\phi}_n], \vec{y}_i\}$ ▷ send every party i all column polynomial commitments and the i^{th} evaluation of every column polynomial with the proper opening proof

 - 9: UPON RECEIVING ($ID.d, send, \text{set}_i$) from P_d for the first time: ▷ echo stage
 - 10: **if** $\forall j, \text{Verify}(\text{pp}, \phi_j, \vec{y}_i[j], t)$ and all points $(j, \vec{y}_i[j])$ form a degree $2t$ polynomial **then**
 - 11: $C = \text{vCom}(\text{pp}, [\hat{\phi}_1 \dots \hat{\phi}_n])$ ▷ commit to all polynomial commitments
 - 12: **for** $j \in [1, n]$ **do**
 - 13: $\pi_j = \text{vGen}(\text{pp}, [\hat{\phi}_1 \dots \hat{\phi}_n], j)$ ▷ send message to each party P_j
 - 14: send “ $ID.d, echo, \text{info}_{i,j}$ ” to P_j , where $\text{info}_{i,j} = \{C, \hat{\phi}_j, \pi_j, \vec{y}_i[j]\}$

 - 15: UPON RECEIVING ($ID.d, echo, \text{info}_{m,i}$) from P_m for the first time: ▷ ready stage
 - 16: **if** $\text{vVerify}(\text{pp}, C, \hat{\phi}_i, i, \pi_i)$ and $\text{Verify}(\text{pp}, \hat{\phi}_i, \vec{y}_i[m], t) = \text{True}$ **then**
 - 17: **if** not yet sent **ready** and received $2t + 1$ valid **echo** with the same C **then**
 - 18: interpolate $\phi_i = \phi(i, Y)$ from any $t + 1$ valid $\vec{y}_i[m]$ in the received **echo** ▷ interpolate column i with the help of other honest parties
 - 19: send “ $ID.d, ready, \text{info}_{i,j}$ ” to P_j where $\text{info}_{i,j} = (y_i = \text{Eval}(\text{pp}, \phi_i, j), \hat{\phi}_i, \pi_i, C)$ ▷ completes Bracha consensus on C , and also sends to party j a point on their row

 - 20: UPON RECEIVING ($ID.d, ready, \text{info}_{m,i}$) from P_m for the first time:
 - 21: **if** $\text{vVerify}(\text{pp}, C, \hat{\phi}_m, m, \pi_m)$ and $\text{Verify}(\text{pp}, \hat{\phi}_m, \vec{y}_m[i], t) = \text{True}$ **then**
 - 22: **if** not yet sent **ready** and received $t + 1$ valid **ready** with this C **then**
 - 23: wait to receive $t + 1$ valid **echo** with this C ▷ must happen eventually
 - 24: interpolate $\phi_i = \phi(i, Y)$ from any $t + 1$ valid $\vec{y}_i[m]$ ▷ interpolate column i with the help of other honest parties
 - 25: send “ $ID.d, ready, \text{info}_{i,j}$ ” to P_j where $\text{info}_{i,j} = (y_i = \text{Eval}(\text{pp}, \phi_i, j), \hat{\phi}_i, \pi_i, C)$ ▷ Bracha consensus on C , while also sending to party j , a point on their row

 - 26: **if** received $2t + 1$ valid **ready** messages **then**
 - 27: interpolate $\phi(X, i)$ from the $2t + 1$ valid **ready** messages ▷ construct the row polynomial from the column points of other parties
 - 28: output ($ID.d, out, shared$) ▷ locally halt
-

Algorithm 2 Reconstruction phase of HAVEN++ for all packed secrets, for server P_i and tag $ID.d$

- 1: UPON RECEIVING ($ID.d, in, reconstruct$) from P_i for the first time:
 - 2: send ($ID.d, reconstruct-share, \hat{\phi}_m, y_m^* = \langle 0, \phi(m, 0), w \rangle$) to all parties \triangleright from Party P_m
 - 3: UPON RECEIVING ($ID.d, reconstruct-share, \hat{\phi}_m, y_m^*$): \triangleright from Party P_m
 - 4: **if** $\hat{\phi}_m$ in C and $Verify(pp, \hat{\phi}_m, y_m^*, t)$ **then**
 - 5: **if** received $2t + 1$ valid **reconstruct-share** messages **then**
 - 6: interpolate $\phi(X, 0)$ from the $2t + 1$ valid points
 - 7: output ($ID.d, out, reconstructed, \phi(-k + 1, 0) = s_k, \forall k \in [1, b]$)
-

Algorithm 3 Reconstruction phase of HAVEN++ for a packed secret share j , for server P_i and tag $ID.d$

- 1: UPON RECEIVING ($ID.d, in, reconstruct, j$) from P_i for the first time:
 - 2: send ($ID.d, reconstruct-share, \phi(-j, m)$) to all parties \triangleright from Party P_m
 - 3: UPON RECEIVING ($ID.d, reconstruct-share, \phi(-j, m)$): \triangleright from Party P_m
 - 4: **if** received at least $t + 1$ **reconstruct-share** **then** \triangleright Run Online Error correcting code
 - 5: $\hat{\phi}_{-j} = ECC(points, t, e)$ \triangleright with e initialized to 1, attempt to interpolate the column polynomial at $-j, \phi(-j, Y)$
 - 6: **if** $\hat{\phi}_{-j} \neq \perp$ **then**
 - 7: output ($ID.d, out, reconstructed, \phi_{-j}(0)$)
 - 8: $e = e + 1$ \triangleright increase the number of errors by one with each failed decoding
-

Sharing phase. The construction of HAVEN++ is heavily influenced by HAVEN [AVZ21]. It operates in three rounds of communication, and follows the same communication pattern as Bracha’s asynchronous reliable broadcast [Bra87]. However, unlike Haven, our construction uses a bivariate polynomial when producing the shares and uses a distributed check to make sure that recovery polynomials are consistent with the shares.

Below, we describe the protocol for the optimal resilience case of $n = 3t + 1$. Conceptually, the protocol contains three distinct phases.

1. The broadcast phase (lines 1-8): The dealer samples a random bivariate polynomial $\phi(X, Y)$ such that each row polynomial is of degree $2t$ and each column polynomial is of degree t . Also, the row polynomial at index 0, encodes b secrets. Each packed secret $s_k \in s_1 \dots s_b$, is packed at $\phi(-k + 1, 0)$. Using pCom, the dealer commits to the first n column polynomial of degree t (lines 3-5). The dealer evaluates each column polynomial $\phi_j(Y)$ at indices $1 \dots n$ and at the same time produces n proofs for every point on every column by calling Eval (lines 6-7). Remember that Eval returns both the (x, y) coordinate as well as a proof that this coordinate is on the committed polynomial. The dealer then sends each row (not column) of evaluation proofs with all n polynomial commitments to every party (lines 6-8).
2. The echo phase (lines 9-14): when a party i receives the first broadcast message from the dealer: The party verifies that the row evaluation proofs are consistent with the column polynomial polynomial commitments and checks that the evaluation points are on a polynomial of degree $2t$ (line 10). If both check pass, then the party commits using a vector commitment to all the polynomial commitments (line 11) in the same order it got from the dealer and produces proofs of inclusion (lines 12-13). It then sends an echo message to every party p_j containing the vector commitment C , the

polynomial commitment $\hat{\phi}_j$, the corresponding inclusion proof π_j and the evaluation proof at j , $\vec{y}_i[j]$ (lines 12-14). Note that when party p_i sends an echo message, she doesn't yet know whether her polynomial commitments will become the consensus ones, because the Bracha broadcast protocol on C might not be complete

3. The ready phase (lines 15-28): When a party p_i hears $2t + 1$ echo messages from different parties with the same vector commitment C , the proper column polynomial commitment and its inclusion proof at position i and with $2t + 1$ valid evaluation proofs on the column polynomial. p_i interpolates its own column polynomial and generates n evaluation proofs (line 18-19). p_i then sends each party j the j^{th} evaluation on its own column polynomial (every point on the column of this party, is also a point on the row of another party) with its own polynomial commitment, inclusion proof and C (line 19).

To guarantee that every honest party sends a ready, just like Bracha broadcast protocol, we also have an amplification step. If a party has not yet sent a ready message because it hasn't received $2t + 1$ valid echo messages, but then receives $t + 1$ ready messages from other parties, it adjusts its waiting condition. Instead of waiting for $2t + 1$ valid echo messages, the party now waits until it has received $t + 1$ valid echo messages on top of the $t + 1$ ready messages it has already received, and then sends its own ready message. This condition must be met eventually because at least one honest party heard $2t + 1$ echo messages where $t + 1$ must have come from honest parties. Once $t + 1$ valid echo messages are heard the protocol continues as before. i.e. p_i interpolates its own column polynomial and generates n evaluation proofs (line 24). p_i then sends each party j the j^{th} evaluation on its own column polynomial with its own polynomial commitment, inclusion proof and C (line 25).

When a party hears ready messages from $2t + 1$ different parties p_m that contain the same vector commitment C , the associated column polynomial commitment ϕ_m and its inclusion proof π_m at position m and with $2t + 1$ valid evaluation proofs for $\phi_m(i)$ (line 26): then, the party interpolates its row polynomial and finishes the dispersal.

Reconstruction phase. We provide two algorithms for reconstruction; the first one (Algorithm 2) enables reconstruction of all secrets at the same time (or one secret that requires a high threshold to recover), and the second one (Algorithm 3) allows selective opening of a specific secret. Both protocols are simple: each party sends one or more points on the bivariate polynomial to the recipient along with corresponding proofs. The most important difference between the two protocols is *which* points are included. As shown in Fig. 2, reconstructing all secrets requires interpolation on a polynomial of degree p , whereas reconstructing a single secret requires interpolation on a different polynomial of degree t . Both types of reconstruction are possible because each party holds one row and one column of data after the sharing phase.

Theorem 1. *Assuming that the underlying polynomial and vector commitment schemes satisfy Definitions 5-8, then the HAVEN++ protocol is a dual-threshold ACSS with $O(\kappa n^2 \log n)$ communication complexity, where κ is the security parameter.*

3.2 Security Analysis

In this section, we provide the proofs for each of the properties in Definition 1.

3.2.1 Proof of Liveness

If the dealer P_d is honest, then P_d will send everyone the same univariate polynomial commitments. Also, each party i receives one evaluation at index i on every univariate

polynomial (a row polynomial). All of the checks on line 10 pass, so the honest parties can vector commit to the same list of polynomial commitments and produce the same vector commitment C along with the same proofs of inclusion for the polynomial commitments. This will enable every party to echo the right evaluations and corresponding proofs to everyone. As a result, all honest parties will pass the checks on lines 17 and interpolate their own univariate column polynomial and will be able to send ready messages with the right evaluations and proofs. This will enable every honest party to reconstruct their row polynomial from the correct ready messages linked to the vector commitment C send by at least $n - t$ honest parties. If any dishonest party tries to send a malformed commitment or evaluation in their echo message or ready, then the evaluation binding property of the underlying commitment (cf. Definition 6) ensures that it will not link back to the same root commitment, so honest parties will eventually disregard this message. Finally, if an honest party completes dispersal and invokes $\text{Reconstruct}(j)$, then every honest party ($n - t$) will evaluate their row polynomial at j and return it. This enables the honest party to run the online error correcting and reconstruct the univariate polynomial of degree $2t$ since there are only t possible errors that can happen.

3.2.2 Proof of Secrecy

Without loss of generality let's assume dispersal has been done with t packed secrets. We will first analyze what the attacker learns during dispersal and then look at the reconstruction step with an inductive approach.

During dispersal, each party's view consists of n polynomial commitments, one row and one column polynomial together with evaluation proofs. As a result an attacker that can corrupt t parties has access to t full rows and t full columns. The hiding property of the polynomial commitments (Def. 8) guarantees that this is insufficient to distinguish any other point on the column polynomial from random with non-negligible probability. It is left to show that the evaluations themselves don't leak information.

Let us first consider the information available for the attacker after dispersal has finished. The attacker has t points on every column polynomial (including the ones holding the secret) and knows t full columns. This is because the dealer sends a row polynomial to every party and during echo every party helps every other party reconstruct its column. Although the attacker knows t points on each column polynomial of degree t . Column-wise, Shamir secret sharing guarantees that the Adv learns nothing about any particular secret, unless they will learn one more point on that column. As a result column-wise from that information alone, a packed secret i is safe. Row-wise the attacker has t points on every row polynomial including the row 0 polynomial of degree $2t$ holding the $t + 1$ packed secrets. For row 0, Shamir secret sharing also guarantees that the Adv learns nothing about any of the $t + 1$ packed secrets. This is because information theoretically, it is easy to see that even if the attacker holds the t shares of a degree $2t$ polynomial, the other $t + 1$ points are still indistinguishable from random. Ergo, the packed secrets could be anything. For example, let the packed secrets be the vector of all 0, interpolate a new polynomial made of degree $2t$ where $t + 1$ points have the value 0 (the packed points), and the other t points are the attacker's points.

Let us consider the new information learned by the attacker when the reconstruction algorithm for the first packed secret i is called. The attacker learns the column polynomial i of degree t . i.e the attacker learns a new point on every row polynomial of degree $2t$, including the row 0 polynomial of degree $2t$ holding the $t + 1$ packed secrets. Still, information theoretically the other t packed secrets are still indistinguishable from random. Even if the attacker holds $t + 1$ shares of a degree $2t$ polynomial the other t points are free. Ergo, the other t packed secrets could still be anything.

Let us consider the new information learned by the attacker when the reconstruction algorithm is called the t^{th} time. The attacker has learned t column polynomials each

of degree t on top of the ones they learned from dispersal. i.e the attacker has learned $2t$ points on every row polynomial of degree $2t$, including the the row 0 polynomial of degree $2t$ holding the 1 secret left. Still, information theoretically the last packed secrets is still indistinguishable from random. Even if the attacker holds $2t$ shares of a degree $2t$ polynomial, there is one point that is free. Ergo, the last packed secret could still be anything.

Lemma 1. *Let $\phi_1 \dots \phi_n$ be a list of column polynomials of degree t . Suppose there exists a set $S \subset \{1, \dots, n\}$ of size $t + 1$ such that for all $i \in S$, the row polynomial formed by interpolating $\phi_1(i), \phi_2(i), \dots, \phi_n(i)$ is of degree p . Then, there exists a unique bivariate polynomial $f(x, y)$ of degree p in one dimension and t in the other dimension where $\phi_i(\cdot) = f(\cdot, i)$.*

Proof. Let $S = \{x_1, \dots, x_{t+1}\}$. Let $\psi_{x_1} \dots \psi_{x_{t+1}}$ be the row polynomials of degree p given by the statement of the lemma. Define $f(x, y) = \sum_{i=1}^{t+1} \psi_{x_i}(y)L_i(x)$, where $L_i(x)$ the appropriate Lagrange coefficient (namely, the unique degree- t univariate polynomial that vanishes at x_j for $i \neq j$ and is 1 at x_i).

Observe that f is of degree p in one variable and t in the other. Now we need to prove $f(x, y) = \phi_y(x)$. If $x \in S$, then this is true by construction, because $f(x, y) = \psi_x(y)$ (because there is only Lagrange coefficient that doesn't vanish at x), which is equal to $\phi_y(x)$ by definition of ψ . Since $f(\cdot, y)$ and $\phi_y(\cdot)$ are degree- t polynomials that agree on $t + 1$ points (namely, all points in S), they must be equal as polynomials, and thus the statement is true for all x , not just $x \in S$.

We now need to prove uniqueness. Observe that for every i , $f(x, i)$, as a polynomial of degree less than n in x , is unique if it agrees with $\phi_i(x)$ in n points (because two different univariate polynomials of degree less than n cannot agree on n points). Thus, viewing $f(x, y)$ as a univariate polynomial in y that evaluates to polynomials in x , we know that its n evaluations are unique. Since it has degree less than n in y , it must also be unique. \square

3.2.3 Proof of Correctness

Correctness states that all nonfaulty parties who complete reconstruction of the k^{th} secret should agree on the same secret, which in turn should be the same as the one used by the dealer if it was honest.

We reason about correctness in the following steps. First, our use of Bracha's broadcast ensures that all honest parties have agreement over the root commitment by the end of the sharing phase. Second, for the broadcast to succeed, at least $t + 1$ honest parties must have received the actual vector of polynomial commitments in the dealer's sharing phase (or else they would not have echoed the root commitment, or anything else for that matter) and have checked that it forms a bivariate polynomial (lemma 1). These parties collectively hold enough data to reconstruct the secrets. Moreover, in Algorithm 1 they will provide every honest party with their column polynomial and its commitment. Finally, this implies that the honest parties have enough information at reconstruction for the online error correction to terminate and produce the correct secret.

3.3 Haven++ with Batching and Aggregate Proofs

In this section, we explain how we can batch multiple invocations of HAVEN++ to save on amortized communication complexity and also enhance run time. The core idea is that instead of the dealer generating one packed bivariate polynomial, the dealer has to generate a batch of them. Batching enables the dealer to always generate n^2 proofs regardless of how many secrets are being batched.

3.3.1 Batching for multiple polynomial evaluation

Consider a set of β polynomial commitments $\{\hat{\phi}_1 \dots \hat{\phi}_\beta\}$ and a single evaluation point j . If a prover wants to prove that $\phi_i(j) = y_i$ for all $i \in \{1 \dots \beta\}$, then the prover has to send β witnesses, one for each $(y_i, \hat{\phi}_i)$. However, one can build one succinct proof for all β polynomials at index j .

To achieve this, hbACSS [YLF⁺22] built their own “Batch Inner-Product” that generalizes the inner product argument of Bulletproofs [BBB⁺18] and used it to empower their ACSS construction. Similarly Alhaddad et al. [ADVZ21], built a generic construction that works for any additive homomorphic polynomial commitments. For this work, we instantiate the generic construction of Alhaddad et al. [ADVZ21] using both Bulletproofs [BBB⁺18] and AMT [TCZ⁺20] with some modifications that we describe below.

To recall, the scheme of Alhaddad et al. [ADVZ21] is a three-round sigma protocol that is made non-interactive using the Fiat-Shamir transform. It has three steps: (1) a commitment, (2) a public coin challenge, and (3) a response.

1. Commitment: prover commits to β different polynomials $\{\phi_1 \dots \phi_\beta\}$ with the same degree d . For each $i \in \{1, \dots, \beta\}$, the prover runs $\hat{\phi}_i = \text{pCom}(\text{pp}, \phi_i, d)$ and sends the pair $(\phi_i(j), \hat{\phi}_i)$.
2. Challenge: verifier generates a random point $c \in \mathbb{F}$ and sends it to the prover.
3. Response: prover interpolates the polynomial ϕ_c from $\{\phi_1, \dots, \phi_\beta\}$ using Lagrange interpolation: $\phi_c(X) = \sum_{i=1}^{\beta} \phi_i(X) \cdot \ell_i(c)$ where $\ell_i(c) = \prod_{\substack{k=1 \\ k \neq i}}^{\beta} \frac{c-k}{i-k}$. After computing ϕ_c , the prover runs $\text{Eval}(\text{pp}, \hat{\phi}_c, j)$ to obtain the witness w , and then sends w to the verifier.

The verifier computes $\phi_c(j) = \sum_{i=1}^{\beta} \phi_i(j) \cdot \ell_i(c)$ by interpolating all $\phi_i(j)$, where $i \in \{1, \dots, \beta\}$ and $\ell_i(c)$ is the i -th Lagrange basis polynomial evaluated at c , defined by $\ell_i(c) = \prod_{\substack{k=1 \\ k \neq i}}^{\beta} \frac{c-k}{i-k}$. Subsequently, the verifier calculates $\hat{\phi}_c = \sum_{i=1}^{\beta} \hat{\phi}_i \cdot \ell_i(c)$. The verifier can compute $\hat{\phi}_c$ thanks to the additive homomorphic property of the employed polynomial commitment scheme. The verifier accepts the proof if $\text{Verify}(\text{pp}, \hat{\phi}_c, (j, \phi_c(j), w), d)$ returns True, and rejects it otherwise.

Consider a prover who wishes to perform n simultaneous evaluation proofs for a verifier. If implemented naively with Fiat-Shamir, the prover has to generate n different challenges and interpolate n different polynomials. Instead, we adopt the same approach used in hbACSS, where we employ a common challenge for all proofs across all verifiers.

If a prover intends to create multiple proofs across various indices, it is not possible to send every set of evaluations to each verifier individually. To address this, we construct a Merkle tree in which each leaf node contains all of the transcripts associated with a specific verifier, and the root hash of the Merkle tree serves as the shared challenge. Subsequently, we provide each verifier with a Merkle branch, allowing them to reconstruct the root hash and verify that it fully encompasses all of the verifier’s transcripts.

We formally define our batch proof in Algorithm 4. Note that our algorithm is named Aggregate Batch Proof rather than Double Batch Proof, as referred to in Yurek et al. [YLF⁺22], despite identical functionality and parameters. This terminology is chosen to emphasize the dual enhancements of aggregating multiple proofs and improving prover efficiency through the batching of multiple evaluations.

Construction. To support efficient batched calls to HAVEN++, our batched HAVEN++ protocol makes some changes to the sharing phase in Algorithm 1. The main changes can be summarized with the following:

Algorithm 4 Aggregate Batch Proof Algorithm

Require: A polynomial commitment scheme \mathcal{P} that is additively homomorphic. Public parameters \mathbf{pp} of the polynomial commitment scheme. A batch of β independent polynomials $[\phi_1(X), \dots, \phi_\beta(X)]$, each of degree at most d , defined over a finite field \mathbb{F} . The number of evaluations n (e.g., the total number of parties). A collision resistant hash function $H : \{0, 1\}^* \rightarrow \mathbb{F}$ mapping to field elements of \mathbb{F} .

- 1: $\text{AGGBATCHPROOF}(\mathbf{pp}, [\phi_1(X), \dots, \phi_\beta(X)], n, d)$
- 2: Initialize an empty list W
- 3: **for** $i = 1$ to n **do**
- 4: Compute the vector of evaluations $\mathbf{y}_i = [\phi_1(i), \dots, \phi_\beta(i)]$
- 5: Compute the leaf node $\text{leaf}_i = H(\mathbf{pp}, i, \mathbf{y}_i)$
- 6: Construct $\text{root} = \text{MerkleTree.Create}(\text{leaf}_1, \dots, \text{leaf}_n)$ \triangleright Build a Merkle tree over the leaves $\{\text{leaf}_1, \dots, \text{leaf}_n\}$
- 7: For each i , compute the Merkle proof branch $_i$ (path from leaf $_i$ to root)
- 8: Compute $\phi_{\text{root}}(X) = \sum_{k=1}^{\beta} \phi_k(X) \cdot \ell_k(\text{root})$ where $\ell_k(\text{root}) = \prod_{\substack{j=1 \\ j \neq k}}^{\beta} \frac{\text{root} - j}{k - j}$ \triangleright
- Interpolate the polynomial $\phi_{\text{root}}(X)$ from $[(1, \phi_1(X)), \dots, (\beta, \phi_\beta(X))]$ using Lagrange interpolation
- 9: $y_1 \dots y_n = \text{BatchProof}(\mathbf{pp}, \phi_{\text{root}}, n, d)$
- 10: **for** $i = 1$ to n **do**
- 11: Parse y as $\langle i, \phi_{\text{root}}(i), w_i \rangle$
- 12: Construct the tuple $w'_i = \langle i, \mathbf{y}_i, \text{branch}_i || w_i \rangle$
- 13: Append w'_i to the list W
- 14: **return** W
- 15: $\text{AGGBATCHVERIFY}(\mathbf{pp}, [\hat{\phi}_1, \dots, \hat{\phi}_\beta], y, d)$
- 16: Parse y as $\langle i, \mathbf{y}_i, \text{branch}_i || w_i \rangle$
- 17: Compute $\text{leaf}_i = H(\mathbf{pp}, i, \mathbf{y}_i)$
- 18: **if** not $\text{MerkleTree.Verify}(\text{leaf}_i, \text{branch}_i)$ **then**
- 19: **return** **False**
- 20: Let $\text{root} = \text{branch}_i[0]$ \triangleright get the root of the Merkle tree from the Merkle proof
- 21: Compute $\hat{\phi}_{\text{root}} = \sum_{k=1}^{\beta} \hat{\phi}_k \cdot \ell_k(\text{root})$
- 22: Compute $\phi_{\text{root}}(i) = \sum_{k=1}^{\beta} \phi_k(i) \cdot \ell_k(\text{root})$ using \mathbf{y}_i
- 23: **return** $\text{Verify}(\mathbf{pp}, \hat{\phi}_{\text{root}}, w_i, d)$

1. The dealer creates β batches of bivariate polynomials each packing b secrets.
2. The dealer calls `AggBatchProof` to generate n proofs for β univariate column polynomials at a time.
3. Instead of every party vector committing to all the univariate polynomial commitment directly. They commit to lists of column polynomial commitments (of size β) where each column polynomial commitment is from a different bivariate polynomial.
4. Every party can verify β points at time (each point is on a column that belong to a different bivariate polynomial) using `AggBatchVerify`.

We now offer more details into the full protocol:

- The broadcast phase: Instead of generating one bivariate polynomial, the dealer generates a batch β of bivariate polynomials. Each bivariate polynomial packs b secrets $s_1 \dots s_b$ (line 2). To cater for a list of bivariate polynomials, the dealer commits using `pCom` to n univariate polynomials of every bivariate polynomial (lines 3-5). Instead of producing proofs for every evaluation point, we now batch across all β bivariate polynomials using the algorithm `AggBatchProof` 4. In more details, for a specific column j , $\phi_{i,j}$ of every bivariate polynomial $i \in [1, \beta]$ is used as input to the new algorithm `AggBatchProof`, which will produce n proofs i.e `AggBatchProof(pp, [\phi_{1,j} \dots \phi_{\beta,j}], n, t)`. Party p_i will receive the i^{th} proof. Instead of sending one row to every party, the dealer sends β rows (i^{th} row of every bivariate polynomial) at a time with n proofs, regardless of how big β can be. Still, the dealer has to send all $(n * \beta)$ polynomial commitments and polynomial evaluations to every party.
- The echo phase (lines 9-14): Each party p_i verifies that all row evaluations and proofs are consistent with the column polynomials. One batch of β univariate column polynomial at a time. Let C_j be the j^{th} univariate polynomial commitment of every bivariate added together in a list and let $r_1 \dots r_\beta$ be the row evaluation of every bivariate polynomial. The verifier calls `AggBatchVerify(C_j, \langle j, [r_1[j] \dots r_\beta[j]], \vec{y}_i[j], t \rangle)`. This allows the verifier check β points each from every row at the same time, for one proof. Also, just like before, each row of evaluations must be on a degree $2t$ polynomial (line 10).

Instead of committing to all univariate polynomial commitments directly (line 11), party p_i vector commits to every C_i (after hashing it) instead, producing one vector commitment C . Remember that every C_i is a list of column polynomial commitments, where each polynomial commitment is at index i of the list of batched bivariate polynomials. Just as before the party also produces proofs of inclusion with the exception that that every element inside of the vector commitment is a hash to a list of polynomial commitments instead of hash of a single polynomial commitment. (lines 12-13) The party then sends an echo message to every party p_j containing the vector commitment C , the polynomial commitments C_j , the corresponding inclusion proof π_j and the evaluation proof at j , $\vec{y}_i[j]$ (lines 12-14).

- The ready phase (lines 15-28): When a party p_i hears $2t + 1$ echo messages from different parties with the same vector commitment C , the proper column polynomial commitments and its inclusion proof at position i and with $2t + 1$ valid evaluation proofs for all column polynomials. p_i interpolates its own β column polynomials and generates n evaluation proofs by calling `AggBatchProof` on the list of column polynomials it just interpolated. p_i then sends each party j the j^{th} evaluation on its own column polynomial for every bivariate polynomial with its own polynomial commitments C_i , inclusion proof and C (line 19).

Amortized costs per secret per party as a function of the number of batches




Figure 3: The computational cost per party in HAVEN++ is reduced as the number of batches increases. Each batch packs $t+1$ secrets. (Each curve has a fixed number of parties n , and lower is better.)

If a party didn't send a `ready` and hears $t+1$ `ready`. The party waits until it hears $t+1$ valid echo messages instead of $2t+1$. Once $t+1$ valid echo messages are heard the protocol continues as before. i.e. p_i interpolates its own β column polynomials and generates n evaluation proofs by calling `AggBatchProof` (line 24). p_i then sends each party j the j^{th} evaluation on its own column polynomial for every bivariate polynomial with its own polynomial commitments C_i , inclusion proof and C (line 25).

When a party hears $2t+1$ `ready` messages from every different party p_m with the same vector commitment C , the proper column polynomial commitments C_m and its inclusion proof π_m at position m and with $2t+1$ valid evaluation proofs (line 26). The party interpolates β row polynomials and finishes the dispersal.

We describe the asymptotic communication complexity of HAVEN++ in §5.2.3.

4 Experimental Results

We have implemented the HAVEN++ protocol in Python, and our open-source code is available at <https://github.com/nicolas3355/AMPC>. In this section, we describe some features of our code, and we provide detailed experimental results (with a few different polynomial commitment schemes) and comparisons to the hbACSS family of protocols.

Implementation details. The implemented version of the HAVEN++ protocol generally performs the same computation as shown in Algorithms 1-3 and described in §3. However, there are a few noteworthy differences that we describe below.

First, in the dealing step, the dealer sends both a row polynomial and column polynomial to every party, rather than just a row polynomial. This has no impact on confidentiality or asymptotic communication, and provides a small efficiency boost in the honest dealer setting. This is because if a party p_i has received a column and row polynomial from the dealer for a specific root commitment C , and other parties are sending echos and ready messages with evaluations with the root commitment C , then p_i can disregard verifying and using those evaluations because p_i already has that information from the dealer.

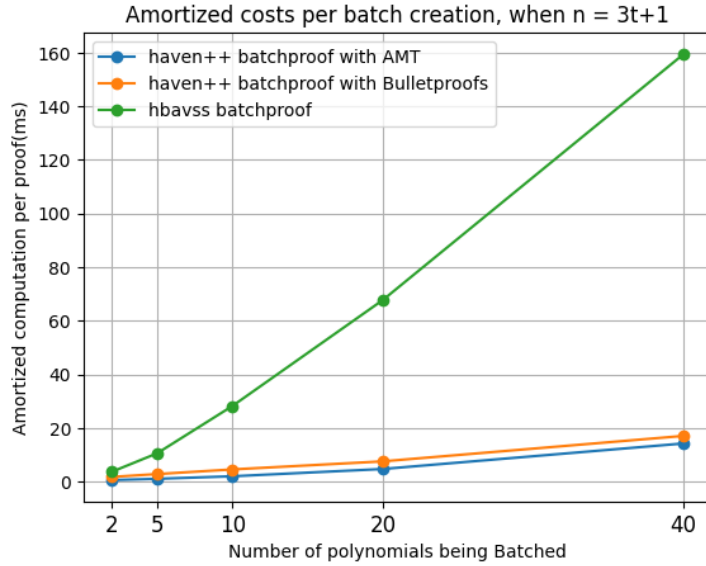


Figure 4: HAVEN++ batching instantiated with AMT [TCZ⁺20] and Bulletproofs [BBB⁺18] vs hbACSS batch proof [YLF⁺22]. Our batch proof substantially beats hbACSS in batch creation for any number of polynomials batched.

Second, the implemented version uses fast Fourier transforms over the roots of unity in order to evaluate and interpolate the polynomials efficiently. The implemented version also batches multiple bivariate polynomials as described in §3.3, and it is also a multi-core implementation that delegates each bivariate polynomial to a different core.

Experimental setup. We implemented the HAVEN++ protocol using Python 3.7.3. on top of the existing open source implementation of hbACSS [YLF⁺22]. For easy comparison, we use the same Python wrapper that implements the elliptic curve (BLS12-381) in Rust (field size is 381 bits), the same setting of $n = 3t + 1$, and also run in the same docker container as hbACSS for easy reproducibility of both our work and theirs. The containers run on our machine which has an AMD 3700X CPU (released in 2019) with 8 main cores and 80GB RAM. Note that HAVEN++ comfortably runs with less than 8GB of RAM.

We use `asyncio` for managing concurrent communication. All parties are simulated using a single core and run one after the other in a queue. It is important to note that our dealer makes use of multiple cores when generating the proofs to simulate a real use case scenario and exploit the way we batch proofs. Also, parties use multiple cores to verify the dealer message. However, when verifying messages from each other no multi-core is used.

The experimental results shown in Yurek et al. [YLF⁺22] are extrapolated: they run their protocol with dummy polynomial commitments, and use this to estimate the total runtime. By contrast, the figures shown in this work are based on actual executions of their protocol using the primitives that they have developed.

Experimental results. We show the results of our experiments in Figures 1-5. The overall results are that our construction scales better to a large message with a large number of parties, using less CPU and RAM resources and benefiting from batching.

Figure 1 compares the performance of HAVEN++ and the two fastest variations of hbACSS, namely variants 0 and 2. Variant 0, is the fastest in the optimistic case where there are no faults, while Variant 2 is better in the pessimistic case where there are t faults.




Figure 5: HAVEN++ batching instantiated with AMT [TCZ+20] and Bulletproofs [BBB+18] vs hbACSS’s batch proof verification [YLF+22]. Our batch proof verification substantially beats hbACSS in batch creation for any number of polynomials batched.

We show results for a constant batch size, and for a batch size that scales linearly in the number of parties. The figure shows that as n scales, our amortized computational cost per party remains low whereas hbACSS grows rapidly.

The remaining figures show the impact of batching. Figure 3 shows how HAVEN++ becomes more efficient as the batch size grows, even while holding fixed the number of parties and the number of secrets packed into each polynomial. Figures 4 and 5 show the costs of creating and verifying a proof, respectively. They show a 2-8 \times speedup for HAVEN++ relative to hbACSS.

5 Applications of ACSS

ACSS is a building block for a wide array of applications, including Byzantine Agreement, Weak Leader Election (and Weak Common Coin), ADKG and Asynchronous Multi Party Computation in general. Our ACSS serves as a drop in replacement, for many of the AVSS and ACSS primitives in those applications.

We will provide a brief overview of each application before focusing primarily on our ADKG application. This focus is twofold: (1) to introduce a novel method of combining randomness, which could be of independent interest, and (2) to highlight new asymptotic improvements in amortized word complexity. Note that for simplicity we will demonstrate the application of HAVEN++ without batching and use word complexity instead of communication complexity. However, for real world applications, batching would enhance the communication complexity by an extra log factor when batching at least $\log(n)$ invocation of HAVEN++.

1. **Byzantine Agreement and Weak Leader Election:** HAVEN++ shares the same interface with Bingo: (1) it allows a dealer to share up to k secrets where $k \leq t + 1$, (2) it allows reconstruction of any single secret if $t + 1$ parties are present, and (3) it allows to reconstruct the sum of those secrets. As such, HAVEN++ can also be used to create Validated Byzantine Agreement and weak common coin in the same

way that Bingo did, i.e., by using the techniques of Abraham et al. [AJM⁺23b] of running gather and of Canetti et al. [CR93b] to generate weak common coin from verifiable secret sharing (note that this implies weak leader election).

2. **Asynchronous MPC:** Most modern MPC constructions rely on pre-processing of shared correlated randomness in the offline phase, which it can then use later to speed up the online phase of the computation. A prominent example of such a technique is the generation of Beaver triples [Bea91], where parties maintain secret shares of three variables a , b , and c with the condition that $c = ab$. While we conjecture that HAVEN++ could enhance the amortized communication complexity for generating these triples, in this work we focus on another type of correlated randomness: HAVEN++ enables the generation of dual-threshold shares, which enable multiplication in the online phase with just one round of communication [BTH07, DN07]. We provide details in §5.3. When the distributed randomness technique of §5.1 is applied, HAVEN++ enables the generation of the pre-processing material with a word complexity of just $O(\kappa n^2)$ without the need for a trusted dealer with $b = 1$, $n = 3t + 1$, and $p = 2t$.
3. **ADKG:** HAVEN++ can be used to produce both low and high threshold ADKG; we provide constructions of both in §5.2. For low threshold ADKG our amortized word complexity is $O(\kappa n)$. While our High threshold ADKG is $O(\kappa n^2)$. To do this, we introduce a new way of generating randomness from bivariate polynomials (discussed in §5.1) and use it as a stepping stone for our ADKG construction described in §5.2. For simplicity, we assume that we have a black box access to an Multi-Valued Byzantine Agreement (MVBA) protocol, rather than building one from HAVEN++.

5.1 Generating Distributed Secret Shared Randomness

A common technique to generate distributed secret shared common randomness among n parties is to have every party secret-share a random secret using ACSS. An MVBA protocol (as defined in §2.4) is run by the parties to agree on a set of $n - t$ (sometimes $t + 1$) parties that finished the sharing phase. Afterward, every party sums up the shares from the set of parties that finished dispersal and reveals the sum of the shares. The common random secret is computed as the sum of all those secrets. The idea is that as long as one party is honest and chooses its own secret uniformly at random then the output is random. Notice, that this technique requires to sacrifice n calls to the ACSS disperse to get one common random number.

In this section, we show how to use n calls to HAVEN++ (with $b = t + 1$) in order to produce $O(n^2)$ distributed random secrets that are secret shared using low threshold or $O(n)$ distributed random secrets that are secret shared using high threshold ($b = 1$). We summarize our Distributed Randomness protocol in Algorithm 5 and describe it below.

Sharing Phase. Each party p_i samples b uniformly random numbers from a finite field $s_1 \dots s_b$. Each party p_i then calls the dispersal phase of HAVEN++ and passes the secrets as input.

Agreement Phase. During the agreement phase, parties run a multi-valued validated Byzantine agreement (MVBA) protocol to reach consensus on a subset of terminated dispersals of HAVEN++. Specifically, each party i waits for a set S_i of $2t + 1$ instances to finish. Party i subsequently inputs S_i into the MVBA protocol. Additionally, party p_i maintains a comprehensive set S representing all instances that have terminated so far that is incrementally updated. For a given value S_j provided to the MVBA by party p_j , party

Algorithm 5 Generating Distributed Randomness for party i

SHARING PHASE:

- 1: $S \leftarrow \{\}$
- 2: Sample b random secrets $s_1, \dots, s_b \leftarrow \mathbb{Z}_q$
- 3: $\phi_1, \dots, \phi_b = \text{HAVEN++}(s_1, \dots, s_b)$ \triangleright Let ϕ_0, \dots, ϕ_b be the column univariate polynomial used in HAVEN++ dispersal
- 4: $S \leftarrow S \cup \{j\}$ when j -th HAVEN++ dispersal terminates at party p_i

AGREEMENT PHASE:

- 5: **if** $|S| = 2t + 1$ **then**
- 6: Let $S_i \leftarrow S$, invoke $MVBA(S_i)$ with predicate $P(S_j, S)$ \triangleright S_j is the input value of some party p_j , S is party p_i 's local variable defined in the Sharing Phase. $P(S_j, S)$ only returns 1 once $S_j \subseteq S$.

RANDOMNESS EXTRACTION PHASE:

- 7: Let T be the output of the $MVBA$ protocol after picking exactly the first $2t + 1$
- 8: Let $B_{row} = []$ \triangleright used to store row i of every bivariate polynomial, there is going to be b of them
- 9: Let $B_{col} = []$ \triangleright used to store column i of every bivariate polynomial, there is going to be b of them
- 10: Let $B_{com} = []$ \triangleright used to store the column polynomial commitments of the new bivariate polynomials, there is going to be b of them. It is enough to store $2t + 1$ column polynomial commitments
- 11: Let $O_{col,1}, O_{col,n}$ be the columns that p_i has dispersed during the HAVEN++ dispersal and let $O_{col,-b}, O_{col,0}$ be the columns holding the packed secrets.
- 12: **for** each $j \in T$ **do**
- 13: Let $O_{row,i}$ be row i that p_j has dispersed during the HAVEN++ dispersal.
- 14: Let $\hat{\phi}_1 \dots \hat{\phi}_n$, be all n column polynomial commitments that has been dispersed by p_j during the HAVEN++ dispersal.
- 15: Compute $\hat{\phi}_{-b} \dots \hat{\phi}_0$ homomorphically from $\hat{\phi}_1 \dots \hat{\phi}_{2t+1}$
- 16: **for** each $k \in [-b, 0]$ **do**
- 17: $B_{com}[-k].append(\hat{\phi}_{-k})$
- 18: $B_{col}[-k].append(O_{col,-k})$
- 19: $B_{row}[-k].append(O_{row,j}(k))$
- 20: **if** $|B_{row}[-k]| = 2t + 1$ **then**
- 21: $B_{row}[-k] = interpolate(B_{row}[-k])$
- 22: **for** each $j \setminus T$ **do** \triangleright only needed if the higher threshold of every bivariate need to be opened
- 23: **for** each $k \in [-b, 0]$ **do**
- 24: send $(-k, row_i(j))$ to P_j \triangleright ensure that all parties have columns

DISTRIBUTION PHASE:

- 25: **if** $i \setminus T$ **then** \triangleright only needed if the higher threshold of every bivariate need to be opened
- 26: UPON RECEIVING((k, val) from party P_m for the first time with k :
- 27: $B_{col}[k].append(m, val)$
- 28: **if** $|B_{col}[k]| \geq t + 1$ **then** \triangleright Run Online Error correcting code
- 29: $\phi_{col} = ECC(B_{col}[k], t, e_k)$ \triangleright with e_k initialized to 1, attempt to interpolate the column polynomial of degree t
- 30: **if** $\phi_{col} \neq \perp$ **then**
- 31: $B_{col}[k] = \phi_{col}$
- 32: $e_k = e_k + 1$ \triangleright increase the number of errors by one with each failed decoding
- 33: **if** for all $k \in [-b, 0]$ $B_{col}[k]$ is a polynomial of degree t **then**
- 34: output($B_{com}, B_{col}, B_{row}$)
- 35: **else**
- 36: output($B_{com}, B_{col}, B_{row}$)

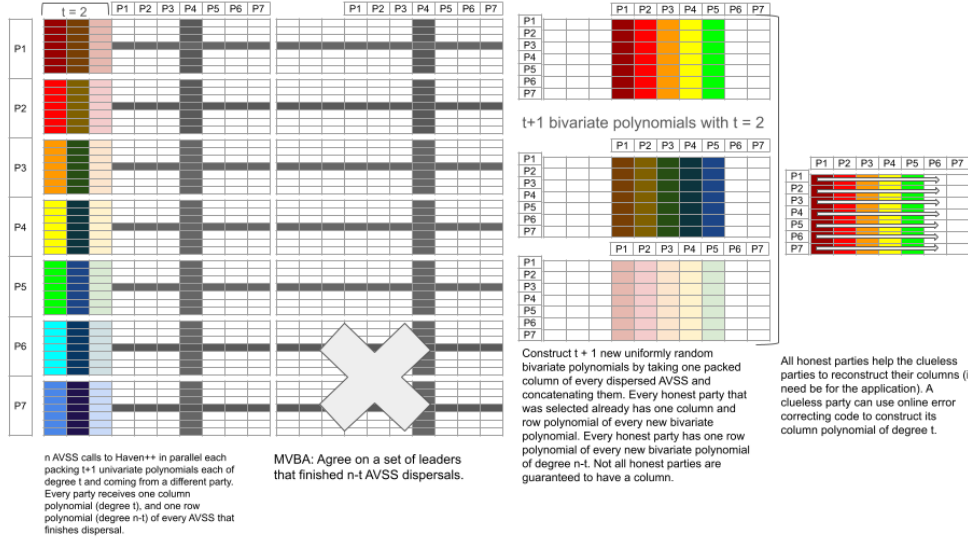


Figure 6: Producing $t + 1$ uniform random bivariate polynomials from HAVEN++. In this figure, $n = 3t + 1$ where $t = 2$.

p_i leverages the predicate $P(S_j, S)$ to verify that $|S_j| \geq 2t + 1$ and $S_j \subseteq S$, confirming that all dispersal instances within S_j have indeed terminated at party p_i . Upon completion, the MVBA protocol yields a set T , where $|T| \geq 2t + 1$. Following the determination of set T by the MVBA protocol, every party p_i pick the first $2t + 1$ parties of the set T and remove the rest.

Randomness Extraction Phase. For every party p_j that was in the set T , party p_i saves the row polynomial that it got from the dispersal phase of that party and all the column polynomial commitments. The first part is to set the column polynomial commitments for each new bivariate polynomial. To this end p_i has to first compute homomorphically the column polynomials that has been shared by every p_j . Let $\hat{\phi}_{-b} \dots \hat{\phi}_0$ be the polynomial commitments for party p_j that has been computed by p_i . Party p_i considers each $\hat{\phi}_k$ as the j^{th} polynomial commitment of the $-k$ bivariate polynomial where, $-b \leq k \leq 0$. Now that p_i has polynomial commitments for every new bivariate polynomial, p_i has to compute its row polynomials. Party p_i first evaluates the row polynomial row_j (acquired from HAVEN++'s dispersal of p_j) at b points from $[-b, 0]$. Recall that at those locations, each evaluation, would be a share of the packed secret with threshold t (as shown in Fig 6). Party p_i considers each $(1, row_1(-k)), (2, row_2(-k)) \dots (2t + 1, row_{2t+1}(-k))$ as row_i of the $-k$ bivariate polynomial where, $-b \leq k \leq 0$.

Distribution Phase. To guarantee that every party has columns and not just rows, party p_i sends every party p_j , $r_1(j) \dots r_b(j)$ evaluations where $r_1 \dots r_b$ are the row polynomials for the newly distributed random biviates (as shown in Figure 6). If a party doesn't have some columns (because it wasn't picked as part of the MVBA set), then that party uses online error correcting code to reconstruct its columns from other parties.

5.2 Low and High Threshold ADKG

In this section, we show how the protocol from §5.1 can be extended, with the help of private polynomial commitments, to build either a low, or high threshold asynchronous distributed key generation protocol. We begin by defining the properties of a private polynomial commitment, followed by detailing how it can be instantiated on top of any additively homomorphic polynomial commitment. Finally, we present the construction of a low- and high-threshold asynchronous distributed key generation protocol.

5.2.1 Defining Private Polynomial Commitments

For our ADKG construction, we require an additional property of our polynomial commitment. A prover must demonstrate that a value is a correct polynomial evaluation in the exponent, while maintaining the privacy of the polynomial evaluation itself. The verifier, upon receiving the polynomial commitment, the index, the evaluation in the exponent, and the proof, can return true if the claim is correct or false otherwise.

1. $PrivateEval(pp, \phi, i) \rightarrow \langle i, g^{\phi(i)}, w \rangle$ is given a polynomial ϕ as well as an index $i \in \mathbb{F}$. It outputs a 3-tuple containing i , the evaluation in the exponent $g^{\phi(i)}$, and witness string w_i . Here g is a random generator of a group G where discrete log is hard.
2. $PrivateVerify(pp, \hat{\phi}, y, d) \rightarrow \text{True/False}$ takes as input a commitment $\hat{\phi}$, a 3-tuple $y = \langle i, j, w \rangle$, and a degree d . It outputs True if the verification succeeds, and False otherwise.

We instantiate private polynomial commitments on top of any additively homomorphic polynomial commitment in the next section.

5.2.2 Constructing Private Polynomial Commitments

Private polynomial commitments (as defined in §5.2.1) can be instantiated using zk-SNARKs. However, we make the observation that Bulletproofs [BBB⁺18] already supports this primitive, with a small modification. In fact, we can construct private polynomial commitments from any additively homomorphic polynomial commitment such that the same field is used to specify the polynomial and the exponents of a Diffie-Hellman group.

We describe the main ideas here, and show the full construction in Algorithm 6. Bulletproofs demonstrate the application of inner product arguments to construct polynomial commitments. To verify a polynomial f 's evaluation at a point i , the prover discloses an inner product involving two vectors, v_1 and v_2 . Here, v_1 represents the coefficients of the polynomial f . The evaluation point i is exponentiated across a range from 0 to d , with d being the polynomial's degree, to form v_2 : $v_2 = \langle 1, i, i^2, \dots, i^d \rangle$. Employing this technique reveals details about the coefficients within the vector v_1 . To make it confidential, a standard technique [BBB⁺18] would be to ask the verifier for a challenge c , the prover sends both $f(i)$ and $f'(i)$ and proves the evaluation for $(f + cf')(i)$ instead of $f(i)$ where f' is picked uniformly at random from the field. Hence, v_1 contains information about the coefficients of $f + cf'$ instead of f . The same methods can be used to build private polynomial commitments on top of discrete log systems. Instead of the prover sending $f(i)$ and $f'(i)$, the prover sends $g^{f(i)}$ and $g^{f'(i)}$ (standard Feldman Commitments) and proves the evaluation for $(f + cf')(i)$ as before. The verifier can check the proof for $(f + cf')(i)$ using the inner product argument and then check in the exponent that indeed $g^{f(i)}(g^{f'(i)})^c = g^{(f+cf')(i)}$. In this method, in the same way as the standard technique, it is acceptable to reveal information about $(f + cf')(i)$ because the polynomial f' serves as a one-time pad that hides f from the verifier. Soundness follows from the fully binding property of the Feldman commitments ($g^{f(i)}$, $g^{f'(i)}$ and $g^{(f+cf')(i)}$) and the correctness of the Polynomial Commitment (in this case Bulletproofs).

Algorithm 6 Private Polynomial Commitments from Additively Homomorphic Polynomial Commitments

Require: ϕ is a polynomial of degree d where every coefficient is uniformly sampled from \mathbb{Z}_p , $i \in \mathbb{Z}_p$ the index that the polynomial need to be opened at, $g \in G$ of order p , and pp denotes the public parameters used by the polynomial commitment.

- 1: **P's input:** (\mathbf{g}, ϕ, i)
- 2: **V's input:** $(\mathbf{g}, \hat{\phi}, i)$
- 3: $P : \phi' \leftarrow \text{random} \in \mathbb{Z}_p^d$
- 4: $P \rightarrow V : \hat{\phi}' = \text{pCom}(\text{pp}, \phi', d)$
- 5: $V : c \leftarrow \text{random} \in \mathbb{Z}_p$
- 6: $V \rightarrow P : c$
- 7: $P \rightarrow V : g^{\phi'(i)}, g^{\phi(i)}$ and $\langle i, (\phi + c\phi')(i), w_i \rangle = \text{Eval}(\text{pp}, (\phi + c\phi'), i)$
- 8: V computes:
- 9: $\widehat{\phi + c\phi'} = \text{Hom}(\text{pp}, \hat{\phi}, \hat{\phi}', c)$
- 10: **if** $\text{Verify}(\text{pp}, \widehat{\phi + c\phi'}, \langle i, (\phi + c\phi')(i), w_i \rangle)$ and $g^{(\phi + c\phi')(i)} = (g^{\phi(i)}) * (g^{\phi'(i)})^c$ **then**
- 11: **return** V_{accepts}
- 12: **else**
- 13: **return** V_{rejects}

We show in Algorithm 6 a formal construction for the case of Bulletproofs, where the main modification occurs in lines 7 and 10. We believe the technique generalizes to work for any additively homomorphic polynomial commitment.

5.2.3 Constructing Low and High Threshold ADKG

To recall, at the end of our distributed randomness protocol in Algorithm 5, every party has a row and column of every bivariate polynomial. In other words, for every bivariate polynomial, every party p_i has $t + 1$ secret shares of $t + 1$ different secrets (each of degree t), or one secret share of one secret that is of degree $2t$. Moreover, every party has a univariate polynomial commitment for every column of every bivariate polynomial. What's left is to have public keys that correspond to packed univariate polynomial commitment(s). To this end, we use private polynomial commitments as defined in §5.2.1.

Let g be a random generator of a group G where discrete log is hard. We show how to get public keys for one distributed bivariate polynomial but without lost of generality the same can be done with all bivariate. Let ϕ_i be the column polynomial that party p_i has. p_i runs $\text{PrivateEval}(\text{pp}, \phi_i, 0)$ to generate the tuple $y_i = \langle i, g^{\phi_i(0)}, w \rangle$ and sends it to all parties. Every party j verifies that the tuple (and the public key is correct) by checking that $\text{PrivateVerify}(\text{pp}, \hat{\phi}_i, y_i, d)$ returns True. Recall that Party j already has consensus over all polynomial commitments (party i doesn't have to send it). Once party p_i hears $2t + 1$ valid private polynomial commitments from $2t + 1$ different parties. For every packed secret k , where $-b \leq k \leq 0$, party p_i computes $g^k = \prod g^{\phi_j(0)L_j(k)}$ where j is the index of the party that sent a valid private polynomial commitment and L is the Lagrange polynomial.

In case of high threshold, $b = 0$, there will only be one public key per bivariate polynomial. While in case of low threshold, $b = t + 1$, every bivariate would be packing $t + 1$ public keys that can be evaluated from $z = -b$ to $z = 0$.

Word Complexity Analysis. Our ADKG construction can be decomposed into the sum of the costs for our ACSS, MVBA, and the transmission of private polynomial commitment evaluations. The total cost is as follows:

- HAVEN++ incurs a word complexity of $O(\kappa n^2)$ when instantiated with Bulletproofs. In our ADKG, we have n calls to our HAVEN++, which results in a word complexity of $O(\kappa n^3)$.
- The MVBA incurs a word complexity of at most $O(\kappa n^3)$, if instantiated with the latest MVBA [DWZ23, AJM⁺23a, AJM⁺23b].
- Broadcasting the private polynomial commitments incurs a word complexity of $O(\kappa n)$ per party per bivariate polynomial. We have a batch of $O(n)$ $(t + 1)$ polynomials and we have n parties. Thus a total word complexity of $O(\kappa n^3)$.

For high threshold ADKG, it is possible to construct $O(n)$ $(t + 1)$ ADKGs (when $b = 1$), with the total amortized word complexity being $O(\kappa n^2)$. For low threshold ADKG, it is possible to construct $O(n^2)$ $((t + 1) \times (t + 1))$ ADKGs with $b = t + 1$. The total amortized word complexity is thus $O(\kappa n)$.

Communication Complexity and Runtime Analysis. Each proof is of size $O(\log(\beta n)\kappa)$ where β is the number of bivariate polynomials being batched rather than $O(\beta \log(n)\kappa)$. Notice that β is inside the log. So, the (non-amortized) concrete communication complexity when batching across β bivariate polynomials is: $6\beta n^2\kappa + 3n^2(\log(\beta n)) + n^2\kappa + n^2\kappa \log(n)$. This complexity is independent of packing: it's the same for one secret as for packing $t + 1$ secrets per polynomial. In the latter case, the amortized asymptotic communication complexity for each one of the $\beta(t + 1)$ secrets is $O(\beta n\kappa + n \log(\beta n) + n \log(n)\kappa)$. Note that if $\beta = \log(n)$, then we get $O(n\kappa)$; there is no extra log factor.

As for the concrete runtime performance boost for both the verifier and dealer: intuitively, batching achieves this because the proof size is $\log(\beta n)$ rather than $\beta \log(n)$. The runtime to produce all the proofs is: $O(\beta n\kappa) + O(n^2 \log(\beta n)\kappa)$ when batching across β bivariate polynomials. Compare this with $O(\beta n^2 \log(n)\kappa)$ if no batching is used. For multicore, you can divide by the number of cores.

5.3 Asynchronous MPC

Just like with Beaver triples [Bea91], *dual secret sharing* (or *random double sharing*) [BTH07, DN07] is a pre-processing building block for performing multiplication gates in MPC. This technique requires having two secret different sharing of the same random secret under two different thresholds t , and $2t$ respectively. For each multiplication, the pre-processing material is consumed and never used again. Our construction from Algorithm 5 instantiated with $b = 1$, can be used to generate distributed dual secret-shared randomness with no trusted dealer. To open the secret under a $2t$ threshold, one can use Algorithm 2, while to open the same secret under a t threshold, one can use Algorithm 3. We remind the reader how to use $\{t, 2t\}$ dual sharing to do multiplication in the honest but curious case. Note that support for multiplication with malicious security can be achieved, if additively homomorphic polynomial commitments are used.

Multiplication in the Honest But Curious Setting. Assume there are two secrets s_1 and s_2 secret shared using Shamir secret sharing, such that any $t + 1$ parties can reconstruct either secrets. If each party i multiplies the two shares it has, this will result in a polynomial S of degree $2t$ where $S(0) = s_1 s_2$ and each party i has the share $S(i)$. If the parties need to do more multiplications, the parties need to lower the threshold from $2t$ to t while still maintaining the same constant coefficient. The trick involves a pre-processing step, given a random number r , secret shared among n parties using two random polynomials HR and R of degree $2t$ and t respectively, such that $HR(0) = R(0) = r$ and where each party i has the pair of shares $HR(i), R(i)$. Each party opens $HR(i) - S(i)$. It is okay to open this value because $HR(i)$ is random and will act as a blinding factor

to the value $S(i)$. Given $2t$ shares, the polynomial $HR - S$ is opened at 0 in the clear. Each party i locally adds $R(i)$ to the constant $(HR - S)(0)$. It is clear to see that this will result with every party having shares of s_1s_2 under a degree t polynomial.

6 Conclusion

This work presents a dual-threshold ACSS construction called HAVEN++, which doesn't require trusted setup and is free from any complaint phase. HAVEN++ accommodates both low and high thresholds to reconstruct the entire message, and it boasts an optimal number of rounds and communication complexity while also demonstrating practical efficiency. Furthermore, we introduce an innovative method to construct an ADKG system by utilizing secret shared bivariate polynomials coupled with private polynomial commitments.

Acknowledgments

The authors thank the anonymous reviewers and shepherd for their helpful feedback. This material is based on work supported by DARPA under Agreement No. HR00112020021 and by the National Science Foundation under Grants No. 1801564, 1915763, 2209194, 2217770, and 2228610.

References

- [AAPP24] Ittai Abraham, Gilad Asharov, Shravani Patil, and Arpita Patra. Perfect asynchronous mpc with linear communication overhead. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024*, pages 280–309, Cham, 2024. Springer Nature Switzerland. doi:10.1007/978-3-031-58740-5_10.
- [ADD⁺22a] Nicolas Alhaddad, Sourav Das, Sisi Duan, Ling Ren, Mayank Varia, Zhuolun Xiang, and Haibin Zhang. Balanced byzantine reliable broadcast with near-optimal communication and improved computation. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC'22, page 399–417, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519270.3538475.
- [ADD⁺22b] Nicolas Alhaddad, Sourav Das, Sisi Duan, Ling Ren, Mayank Varia, Zhuolun Xiang, and Haibin Zhang. Brief announcement: Asynchronous verifiable information dispersal with near-optimal communication. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC'22, page 418–420, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519270.3538476.
- [ADVZ21] Nicolas Alhaddad, Sisi Duan, Mayank Varia, and Haibin Zhang. Succinct erasure coding proof systems. Cryptology ePrint Archive, Paper 2021/1500, 2021. <https://eprint.iacr.org/2021/1500>. URL: <https://eprint.iacr.org/2021/1500>.
- [AJM⁺23a] Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, and Gilad Stern. Bingo: Adaptivity and asynchrony in verifiable secret sharing and distributed key generation. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 39–70, Cham, 2023. Springer Nature Switzerland. doi:10.1007/978-3-031-38557-5.

- [AJM⁺23b] Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Reaching consensus for asynchronous distributed key generation. *Distributed Computing*, 36(3):219–252, Sep 2023. doi:[10.1007/s00446-022-00436-8](https://doi.org/10.1007/s00446-022-00436-8).
- [AMS19] Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. Asymptotically optimal validated asynchronous byzantine agreement. In Peter Robinson and Faith Ellen, editors, *38th ACM PODC*, pages 337–346. ACM, July / August 2019. doi:[10.1145/3293611.3331612](https://doi.org/10.1145/3293611.3331612).
- [AVZ21] Nicolas Alhaddad, Mayank Varia, and Haibin Zhang. High-threshold AVSS with optimal communication complexity. In Nikita Borisov and Claudia Díaz, editors, *FC 2021, Part II*, volume 12675 of *LNCS*, pages 479–498. Springer, Berlin, Heidelberg, March 2021. doi:[10.1007/978-3-662-64331-0_25](https://doi.org/10.1007/978-3-662-64331-0_25).
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018. doi:[10.1109/SP.2018.00020](https://doi.org/10.1109/SP.2018.00020).
- [BCC⁺16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357. Springer, Berlin, Heidelberg, May 2016. doi:[10.1007/978-3-662-49896-5_12](https://doi.org/10.1007/978-3-662-49896-5_12).
- [BCG93] Michael Ben-Or, Ran Canetti, and Oded Goldreich. Asynchronous secure computation. In *25th ACM STOC*, pages 52–61. ACM Press, May 1993. doi:[10.1145/167088.167109](https://doi.org/10.1145/167088.167109).
- [BDK13] Michael Backes, Amit Datta, and Aniket Kate. Asynchronous computational VSS with reduced communication complexity. In Ed Dawson, editor, *CT-RSA 2013*, volume 7779 of *LNCS*, pages 259–276. Springer, Berlin, Heidelberg, February / March 2013. doi:[10.1007/978-3-642-36095-4_17](https://doi.org/10.1007/978-3-642-36095-4_17).
- [Bea91] Donald Beaver. Secure multiparty protocols and zero-knowledge proof systems tolerating a faulty minority. *Journal of Cryptology*, 4:75–122, 1991. doi:[10.1007/BF00196771](https://doi.org/10.1007/BF00196771).
- [BFS20] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Cham, May 2020. doi:[10.1007/978-3-030-45721-1_24](https://doi.org/10.1007/978-3-030-45721-1_24).
- [BG18] Jonathan Bootle and Jens Groth. Efficient batch zero-knowledge arguments for low degree polynomials. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 561–588. Springer, Cham, March 2018. doi:[10.1007/978-3-319-76581-5_19](https://doi.org/10.1007/978-3-319-76581-5_19).
- [BGdMM05] Lucas Ballard, Matthew Green, Breno de Medeiros, and Fabian Monrose. Correlation-resistant storage via keyword-searchable encryption. Cryptology ePrint Archive, Paper 2005/417, 2005. <https://eprint.iacr.org/2005/417>. URL: <https://eprint.iacr.org/2005/417>.
- [Bol03] Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 31–46. Springer, Berlin, Heidelberg, January 2003. doi:[10.1007/3-540-36288-6_3](https://doi.org/10.1007/3-540-36288-6_3).

- [Bra87] Gabriel Bracha. Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75(2):130–143, 1987. doi:10.1016/0890-5401(87)90054-X.
- [BTH07] Zuzana Beerliová-Trubíniová and Martin Hirt. Simple and efficient perfectly-secure asynchronous mpc. In Kaoru Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007*, pages 376–392, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. doi:10.1007/978-3-540-76900-2_23.
- [Can96] Ran Canetti. *Studies in secure multiparty computation and applications*. PhD thesis, Citeseer, 1996.
- [CF13] Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 55–72. Springer, Berlin, Heidelberg, February / March 2013. doi:10.1007/978-3-642-36362-7_5.
- [CFS17] Alessandro Chiesa, Michael A. Forbes, and Nicholas Spooner. A zero knowledge sumcheck and its applications. *Electron. Colloquium Comput. Complex.*, 24:57, 2017.
- [CKLS02] Christian Cachin, Klaus Kursawe, Anna Lysyanskaya, and Reto Strohli. Asynchronous verifiable secret sharing and proactive cryptosystems. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 88–97. ACM Press, November 2002. doi:10.1145/586110.586124.
- [CKPS01] Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 524–541. Springer, Berlin, Heidelberg, August 2001. doi:10.1007/3-540-44647-8_31.
- [CKS05] Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in Constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, July 2005. doi:10.1007/s00145-005-0318-0.
- [CR93a] Ran Canetti and Tal Rabin. Fast asynchronous byzantine agreement with optimal resilience. In *25th ACM STOC*, pages 42–51. ACM Press, May 1993. doi:10.1145/167088.167105.
- [CR93b] Ran Canetti and Tal Rabin. Fast asynchronous byzantine agreement with optimal resilience. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '93, page 42–51, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/167088.167105.
- [CT05] Christian Cachin and Stefano Tessaro. Asynchronous verifiable information dispersal. In *SRDS*, pages 191–201. IEEE, 2005. doi:10.1109/RELDIS.2005.9.
- [DN07] Ivan Damgård and Jesper Buus Nielsen. Scalable and unconditionally secure multiparty computation. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 572–590. Springer, Berlin, Heidelberg, August 2007. doi:10.1007/978-3-540-74143-5_32.
- [DWZ23] Sisi Duan, Xin Wang, and Haibin Zhang. Fin: Practical signature-free asynchronous common subset in constant time. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, CCS '23, page 815–829, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3576915.3616633.

- [DXKKR23] Sourav Das, Zhuolun Xiang, Lefteris Kokoris-Kogias, and Ling Ren. Practical asynchronous high-threshold distributed key generation and distributed polynomial sampling. In Joseph A. Calandrino and Carmela Troncoso, editors, *USENIX Security 2023*, pages 5359–5376. USENIX Association, August 2023.
- [DXR21] Sourav Das, Zhuolun Xiang, and Ling Ren. Asynchronous data dissemination and its applications. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 2705–2721. ACM Press, November 2021. doi:10.1145/3460120.3484808.
- [GJKR07] Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. *Journal of Cryptology*, 20(1):51–83, January 2007. doi:10.1007/s00145-006-0347-3.
- [GS24] Jens Groth and Victor Shoup. Fast batched asynchronous distributed key generation. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024*, pages 370–400, Cham, 2024. Springer Nature Switzerland. doi:10.1007/978-3-031-58740-5_13.
- [HMW18] Timo Hanke, Mahnush Movahedi, and Dominic Williams. DFINITY technology overview series, consensus system. *CoRR*, abs/1805.04548, 2018.
- [JLS24] Xiaoyu Ji, Junru Li, and Yifan Song. Linear-communication asynchronous complete secret sharing with optimal resilience. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 418–453, Cham, 2024. Springer Nature Switzerland. doi:10.1007/978-3-031-68397-8_13.
- [KHG12] Aniket Kate, Yizhou Huang, and Ian Goldberg. Distributed key generation in the wild. *Cryptology ePrint Archive*, Report 2012/377, 2012. URL: <https://eprint.iacr.org/2012/377>.
- [KMS20] Eleftherios Kokoris-Kogias, Dahlia Malkhi, and Alexander Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1751–1767. ACM Press, November 2020. doi:10.1145/3372297.3423364.
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Berlin, Heidelberg, December 2010. doi:10.1007/978-3-642-17373-8_11.
- [LY10] Benoît Libert and Moti Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 499–517. Springer, Berlin, Heidelberg, February 2010. doi:10.1007/978-3-642-11799-2_30.
- [MS77] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The theory of error-correcting codes*, volume 16. Elsevier, 1977.
- [SS24] Victor Shoup and Nigel P. Smart. Lightweight asynchronous verifiable secret sharing with optimal resilience. *Journal of Cryptology*, 37(3):27, 2024. doi:10.1007/s00145-024-09505-6.

-
- [TCZ⁺20] Alin Tomescu, Robert Chen, Yiming Zheng, Ittai Abraham, Benny Pinkas, Guy Golan Gueta, and Srinivas Devadas. Towards scalable threshold cryptosystems. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 877–893, 2020. doi:[10.1109/SP40000.2020.00059](https://doi.org/10.1109/SP40000.2020.00059).
- [YLF⁺22] Thomas Yurek, Licheng Luo, Jaiden Fairuze, Aniket Kate, and Andrew Miller. hbACSS: How to robustly share many secrets. In *NDSS 2022*, San Diego, CA, USA, April 24–28 2022. The Internet Society. doi:[10.14722/ndss.2022.23120](https://doi.org/10.14722/ndss.2022.23120).