

Non-Interactive Threshold BBS+ From Pseudorandom Correlations (Full Version)

Sebastian Faust¹ , Carmit Hazay² , David Kretzler¹ , Leandro Rometsch³ , and Benjamin Schlosser¹ 

¹ Technical University of Darmstadt, Germany
`{first.last}@tu-darmstadt.de`

² Bar-Ilan University, Israel
`carmit.hazay@biu.ac.il`

³ Hamburg University of Technology, Germany
`leandro.rometsch@tuhh.de`

Abstract. The BBS+ signature scheme is one of the most prominent solutions for realizing anonymous credentials. Its prominence is due to properties like selective disclosure and efficient protocols for creating and showing possession of credentials. Traditionally, a single credential issuer produces BBS+ signatures, which poses significant risks due to a single point of failure.

In this work, we address this threat via a novel t -out-of- n threshold BBS+ protocol. Our protocol supports an arbitrary security threshold $t \leq n$ and works in the so-called preprocessing setting. In this setting, we achieve non-interactive signing in the online phase and sublinear communication complexity in the number of signatures in the offline phase, which, as we show in this work, are important features from a practical point of view. As it stands today, none of the widely studied signature schemes, such as threshold ECDSA and threshold Schnorr, achieve both properties simultaneously. In this work, we make the observation that presignatures can be directly computed from pseudorandom correlations which allows servers to create signatures shares without additional cross-server communication. Both our offline and online protocols are actively secure in the Universal Composability model. Finally, we evaluate the concrete efficiency of our protocol, including an implementation of the online phase and the expansion algorithm of the pseudorandom correlation generator (PCG) used during the offline phase. The online protocol without network latency takes less than $14ms$ for $t \leq 30$ and credentials sizes up to 10. Further, our results indicate that the influence of t on the online signing is insignificant, $\leq 6\%$ for $t \leq 30$, and the overhead of the thresholdization occurs almost exclusively in the offline phase. Our implementation of the PCG expansion shows that even for a committee size of 10 servers, each server can expand a correlation of up to 2^{17} presignatures in less than 100 ms per presignature.

Keywords: Threshold Signature · BBS+ · Pseudorandom Correlation Functions · Pseudorandom Correlation Generators

1 Introduction

Anonymous credentials schemes, as introduced by Chaum in 1985 [32] and subsequently refined by a line of work [33, 53, 26, 27, 22, 24, 25, 10, 70], allow an issuing party to create credentials for users, which then can prove individual attributes about themselves without revealing their identities. The essential properties these schemes satisfy are *unlinkability*, ensuring that verifiers cannot link two disclosures of credentials of the same identity, and *selective disclosure*, allowing parties to decide which individual attributes of their credentials to disclose. The former makes anonymous credentials a useful privacy tool on the web, allowing clients to authenticate themselves for access to web-based services while preventing service providers from gathering information about the client’s usage patterns. The latter makes anonymous credentials an essential building block for self-sovereign identity frameworks, as it enables clients to not only take responsibility for storing their credentials but also to filter the disclosure of their credentials.

The BBS+ signature scheme [5, 23, 64] named after the group signature scheme of Boneh, Boyen, and Shacham [11] is one of the most prominent solutions for realizing anonymous credential schemes. Abstractly speaking, a BBS+ signature over a set of attributes constitutes credentials, and the holder of such credentials can prove possession of individual attributes using efficient zero-knowledge protocols. BBS+ signatures are particularly suited for anonymous credentials because of their appealing features, including the ability to sign an array of attributes while keeping the signature size constant, efficient protocols for blind signing, and efficient zero-knowledge proofs for selective disclosure of signed attributes (without having to reveal the signature). The importance of BBS+ is illustrated by the renewed attention in the research community [64, 40], several industrial implementations [1, 54, 55], ongoing standardization efforts by the W3C Verifiable Credentials Group and IETF [9, 52], and adaption in further applications [5, 34, 20, 21, 23].

In traditional credential systems, the credential issuer who is in possession of the signing key constitutes a single point of failure. A powerful and widely adapted tool mitigating such a single point of failure is to distribute the cryptographic task (e.g., [50, 43, 51, 41, 62, 30, 29, 46, 47, 3, 31, 35] and many more) via so-called *threshold cryptography*. Here, the cryptographic key is shared among a set of servers such that any subset of t servers can produce a signature, while the underlying signature scheme remains secure even if up to $t - 1$ servers are corrupted. The thresholdization of digital signature schemes comes with significant overhead in computation, communication, and round complexity. This is particularly the case for randomized signature schemes, where a random secret nonce has to be generated and used in combination with the shared secret key by the set of servers. In the signing protocol, this nonce is then used together with the shared key to produce the signature. Concretely, for BBS+ signing, we require a distributed protocol to compute the exponentiation of the inverse of the secret key added to the random nonce securely.

A straightforward approach to compute the inverse is based on the inversion protocol by Bar-Ilan and Beaver [7] and requires server interaction. In order to strengthen the protection against failure and corruption, we assess it as likely for servers to be located in different jurisdictional and geographic regions. In such a setting, any additional communication round involves a significant performance overhead. Therefore, a threshold BBS+ scheme is desired to have a non-interactive signing phase that enables servers to respond to signature requests without any cross-server interaction.

A popular approach in secure distributed computation to cope with the high complexities of protocols is to split the computation into an input-independent offline and input-dependent online phase [37, 56, 65, 66]. The offline phase provides precomputation material, which in the context of a digital signature scheme is called presignatures [42]. These presignatures are produced during idle times of the system and facilitate an efficient online phase. In recent years, Boyle et al. [14, 17, 18] put forth a novel concept to generate precomputation material. The main advantage of this concept is the generation of precomputation material with sublinear communication complexity in the amount of precomputation material generated. Recently, this technique also attracted interest for use in threshold signature protocols [3, 48]. The technique allows parties to precompute values via a pseudorandom correlation generator (PCG) or a pseudorandom correlation function (PCF). These primitives include an interactive setup phase where short keys are generated and distributed. Then, in the evaluation phase, every party locally evaluates on its key and a common input. The outputs look pseudorandom but still satisfy some correlation, e.g., vector oblivious linear evaluation (VOLE), oblivious transfer (OT), or multiplication triples.

In this work, we aim to design a threshold BBS+ signature scheme that simultaneously provides a low-latency non-interactive online signing phase and a preprocessing phase with sublinear communication complexity.

1.1 Contribution

We propose a novel t -out-of- n threshold BBS+ signature scheme in the offline-online model with an arbitrary security threshold $t \leq n$. The centerpiece of our protocol is the design of specifically tailored presignatures that can be directly instantiated from PCG or PCF evaluations and can be used by servers to create signature shares without any additional cross-server communication. This way, our scheme simultaneously provides a non-interactive online signing phase and an offline phase with sublinear communication complexity in the number of signatures. Even for the widely studied signature schemes ECDSA and Schnorr, no threshold protocol exists that achieves both features simultaneously. Moreover, we are the first to present a PCG/PCF-based protocol that supports t -out-of- n threshold, while previous protocols support only n -out-of- n . We formally analyze the active security of all our protocols in the Universal Composability framework under static corruption.

We present two instantiations of the offline phase, one based on PCGs and one based on PCFs. Conceptually, PCFs are better suited than PCGs for pre-

processing signatures as PCFs allow servers to compute presignatures only when needed. In contrast, PCGs require the generation of a large batch of presignatures at once that need to be stored on the server side. Nevertheless, existing PCG constructions provide better efficiency than PCF constructions. Therefore, we present protocols for both primitives. For the sake of presentation, we focus on PCGs in the main body and present the a new definition of PCFs and the PCF-based offline phase in the Appendix F-H.

Unlike prior work using silent preprocessing in the context of threshold signatures [3], we use the PCG and PCF primitive in a black-box way, allowing for a modular treatment. In this process, we identify several issues in using the primitives in a black-box way, extend the definitional frameworks accordingly, and prove the security of existing constructions under the adapted properties.

On a practical level, we provide an extensive evaluation of our protocol, including an implementation and experimental evaluation of the online phase and the seed expansion of the PCG-based offline phase. Since state-of-the-art PCF constructions lack concrete efficiency, we focus our evaluation on the PCG-based preprocessing. Given preprocessed presignatures, the total runtime of the online signing protocol is below 13.486 ms plus one round trip time of the slowest client-server connection for $t \leq 30$ signers and message arrays of size $k \leq 10$. Our benchmarks show that the influence of the number of signers on the runtime of the online protocol is minimal; increasing the number of signers from 2 to 30 increases the runtime by just 1.13% – 6.02% (for message array sizes between 2 and 50). Further, our results show that the cost of thresholdization occurs almost exclusively in the offline phase; a threshold signature on a single message array takes 7.286 ms in our protocol, while a non-threshold signature, including verification of the received signature, takes 6.971 ms; ignoring network delays which are the same in both settings. Our implementation of the PCG seed expansion shows that batches of up to 2^{17} presignatures can be expanded with less than 100 ms per presignature for $n \leq 10$.

1.2 Technical Overview

BBS+ Signatures. The BBS+ signature scheme was first described by Au et al. [5] as a modification of the Boneh, Boyen, and Shacham group signature scheme [11]. Recently, Tessaro and Zhu [64] showed how to reduce the size of BBS+ signatures by one group element. In this paper, we follow the improved scheme by Tessaro and Zhu [64] and discuss how to adapt our protocol to the legacy scheme of Au et al. [5] in Appendix B.

Let $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T be groups of prime order p with generators $g_1 \in \mathbb{G}_1$ and $g_2 \in \mathbb{G}_2$ and let map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ be a bilinear pairing. A BBS+ signature on a message array $\{m_\ell\}_{\ell \in [k]}$ is a tuple (A, e) with $A = (g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{\frac{1}{x+e}}$ for a random nonce $e \in_R \mathbb{Z}_p$, secret key $x \in \mathbb{Z}_p$ and a set of random elements $\{h_\ell\}_{\ell \in [k]}$ in \mathbb{G}_1 . To verify under public key g_2^x , check if $e(A, g_2^x \cdot g_2^e) = e(g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell}, g_2)$ (see Appendix A for a formal description).

Distributed Inverse Calculation. The main difficulty in thresholdizing the BBS+ signature algorithm comes from the signing operation requiring the computation of the inverse of $x + e$ without leaking x . This non-linear operation is expensive to be computed in a distributed way. Similar challenges are known from other signature schemes relying on exponentiation (or a scalar multiplication in additive notation) of the inverse of secret values, e.g., ECDSA [6, 29, 3, 69, 13]. The typical approach (cf. [7]) to compute $M^{\frac{1}{y}}$ for a group element M and a secret shared y is to separately open $B = M^a$ and $\delta = a \cdot y$ for a freshly shared random a . The desired result can be reconstructed by computing $M^{\frac{1}{y}} = B^{\frac{1}{\delta}}$.

Since δ is the product of two secret shared values, it still is a non-linear operation requiring interaction between the parties. Nevertheless, as δ is independent of the actual message, several such values can be precomputed in an *offline* phase. As explained next, a similar, yet more involved, approach can be applied to the BBS+ protocol, allowing an efficient, non-interactive online signing based on correlated precomputation material.

The Threshold BBS+ Online Protocol. We describe a simplified, n -out-of- n version of our threshold BBS+ protocol. Assume a BBS+ secret key x , elements $\{h_\ell\}_{\ell \in [k]}$ in \mathbb{G}_1 , a random blinding factor $a \in \mathbb{Z}_p$ and n servers, each having access to a preprocessed pair $(a_i, \beta_i) \in \mathbb{Z}_p^2$, in the following called *presignatures*, such that $\sum_{\ell \in [n]} \beta_\ell = a \cdot x$ for $a = \sum_{\ell \in [n]} a_\ell$.

To sign a message array $\mathbf{m} = \{m_\ell\}_{\ell \in [k]}$, each server first obtains the random nonce e via a call to a random oracle H , i.e., $e = H(\text{sid} || \text{ssid} || \mathbf{m} || [n])$, where sid and ssid are identifiers for the signing session such that the nonce is fresh for each signing request. Then, each server computes $A_i := (g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{a_i}$ and $\delta_i = \beta_i + a_i \cdot e$ and outputs a partial signature $\sigma_i := (A_i, \delta_i)$. This allows the receiver of the partial signatures to reconstruct δ and compute

$$A = \left(\prod_{i \in [n]} A_i \right)^{\frac{1}{\delta}} = \left((g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^a \right)^{\frac{1}{a(x+e)}}$$

such that the tuple (A, e) constitutes a valid BBS+ signature [64]. Each signature requires a new preprocessed pair to prevent straightforward forgeries.

The specialized layout of our presignatures allows us to realize a non-interactive signing procedure. In contrast, using plain multiplication triples, as often done in multiparty computation protocols [8, 37], would require one additional round of communication. Further, our online protocol provides active security at a low cost. This is achieved by verifying the received signatures and works since the presignatures are created securely.

The Preprocessing Protocol. The presignatures required by our online phase consist of two values. Each party P_i obtains a share a_i of a random value a and a share β_i of the product $\beta = a \cdot x$. Since a and x are both secret shared values, we can write $\beta = (a_1 + \dots + a_n) \cdot (x_1 + \dots + x_n) = a_1 x_1 + \dots + a_1 x_n + \dots + a_n x_1 + \dots + a_n x_n$, i.e. β is the sum of all cross products. While each party P_i can compute the cross product $a_i x_i$ on its own, party P_i must communicate with P_j to compute

the cross products $a_i x_j$ and $a_j x_i$. Importantly, while the share a_i must be a fresh random value for each presignature, the secret key share x_i is fixed for all presignatures. This fact allows us to obtain additive shares of the cross products via *vector oblivious linear evaluation* (VOLE) correlations. A VOLE correlation is a two-party correlation, in which party P_1 gets a tuple of two random values (a, u) and party P_2 gets a tuple (s, v) , where s is fixed for all correlation tuples and v is random. Moreover, the tuples satisfy the correlation $a \cdot s = u + v$. Using VOLE correlations, where the fixed s is set to be the secret key share x_i , parties P_i and P_j obtain additive shares $u_{j,i}, v_{j,i}$ of the cross product $a_j x_i$. The same applies for the cross product $a_i x_j$. After obtaining additive shares of all cross products, every party P_i calculates β_i by summing up the additive shares and the cross product $a_i x_i$.

An appealing choice for instantiating the preprocessing protocol is to use pseudorandom correlation generators (PCG) or functions (PCF), as they enable the efficient generation of correlated random tuples. More precisely, PCGs and PCFs allow two parties to expand short seeds to fresh correlated random tuples locally. While the distributed generation of the seeds requires more involved protocols and typically relies on general-purpose multiparty computation, the seed size and the communication complexity of the generating protocols are sublinear in the size of the expanded correlated tuples [14, 17]. Using PCG or PCF constructions, we can efficiently compute VOLE correlations [14, 17, 18, 19, 36, 57, 15]. Note that the a_i value must be the same for all cross terms, so we require the VOLE PCG/PCF to provide the reusability feature. This allows party P_i to use the same input value a_i in all VOLE correlations for the cross terms $a_i x_j$ with $j \neq i$.

Using PCGs/PCFs in a Black-Box Way. Pseudorandom correlation generators (PCGs) and pseudorandom correlation functions (PCFs) are introduced in [14] and [18]. Concrete constructions of both primitives for simple correlations, such as VOLE, are presented in a line of work including [14, 17, 16, 19, 18, 36, 57]. In our work, we aim to deal with PCGs/PCFs in a black-box way such that we can instantiate our protocols with arbitrary constructions as long as they fulfill our requirements. These requirements include supporting VOLE correlations, supporting the active security setting, and providing the opportunity to reuse inputs, as emphasized above. A first step towards black-box usage of PCGs was taken by [17]. This work defines an ideal functionality for correlated randomness, which they show can be instantiated by PCGs. However, the definition does not support reusing inputs to PCGs.

In contrast to the simulation-based approach taken by [16], [17] and [18] present game-based notions that lay the foundation of the resuability property for PCGs and PCFs. However, their definitions consider passive security only and are unsuitable for black-box usage. Therefore, we introduce new game-based notions called *reusable PCG* and *reusable PCF*, which capture the active security setting and permit black-box use.

Identical to prior definitions of PCGs and PCFs, our primitives consist of a key generation *Gen* and an expansion algorithm *Expand* or evaluation algo-

rithm Eval. The reusability feature allows both parties to specify an input to the key generation, which is used to derive the correlation tuples. Additionally, our reusable primitives must satisfy four properties. Three of these properties are stated by [17] and [18], two of which we slightly modified. Our new insight is the requirement of the *key indistinguishability* property, which we specifically introduce to cover malicious parties. The key indistinguishability property states that the adversary cannot learn information about the honest party’s input to the key generation, even if the corrupted party chooses its input arbitrarily. This property makes our notion suitable for the active security setting. We state and prove that existing PCG and PCF constructions fulfill our new definition.

The t -out-of- n Setting. So far, we discussed a setting where n -out-of- n servers must contribute to the signature creation. However, in many use cases, we need to support the more flexible t -out-of- n setting with $t \leq n$. In this setting, the secret key material is distributed to n servers, but only t must contribute to the signing protocol. A threshold $t \leq n$ improves the flexibility and robustness of the signing process, as not all servers must be online.

The typical approach in the t -out-of- n setting is to share the secret key material using Shamir’s secret sharing [61] instead of an additive sharing as done above. While additive shares are reconstructed by summation, Shamir-style shares must be aggregated using Lagrange interpolation, either on the client or server side. In this work, we reconstruct on the server side due to technical details of our precomputation protocols. Note that prior threshold signature schemes leveraging PCF/PCGs (e.g., [3, 48]) only support an n -out-of- n setting.

On a technical level, the challenge for client-side reconstruction is due to VOLE correlations providing us with two-party additive sharings of multiplications, e.g., $u_{i,j} + v_{i,j} = a_i x_j$. The additive shares $u_{i,j}$ and $v_{i,j}$ are sufficient for the n -out-of- n setting (see above). However, it is unclear how VOLE outputs can be transformed into Shamir sharings of $a \cdot x$ so that a client can reconstruct in a t -out-of- n setting.

We, therefore, incorporate a share conversion mechanism from Shamir-style shared key material into additively shared presignatures on the server side. Our mechanism consists of the servers applying the corresponding Lagrange interpolation directly to the outputs of the VOLE correlation. More precisely, as described above, each party P_i gets additive shares of the cross terms $a_i x_j$ and $a_j x_i$ for every other party P_j , where x_ℓ denotes the Shamir-style key share of P_ℓ . Let $u_{i,j}$ be the additive share of $a_i x_j$, then party P_i multiplies the required Lagrange coefficient $L_{j,\mathcal{T}}$ to this share and $L_{i,\mathcal{T}}$ to $v_{j,i}$, where \mathcal{T} is the set of t signers. The client provides the set of servers as part of the signing request to enable the servers to compute the interpolation.

While the secret key is shared using Shamir’s secret sharing, we can define the random blinding factor a as the additive sharing of all t involved signers, i.e., $a = \sum_{\ell \in \mathcal{T}} a_\ell$. Since there is at least one honest signer in \mathcal{T} , a is still random and unknown to the adversary.

1.3 Related Work

Most related to our work are the works by Gennaro et al. [44] and Doerner et al. [40], proposing threshold protocols for the BBS+ signing algorithm. While [44] focuses on a group signature scheme with threshold issuance based on the BBS signatures, their techniques can be directly applied to BBS+. [40] presents a threshold anonymous credential scheme based on BBS+. Both schemes compute the inverse using classical techniques of Bar-Ilan and Beaver [7]. Moreover, they realize the multiplication of two secret shared values by multiplying each pair of shares. While [44] uses a two-round multiplication protocol based on an additively homomorphic encryption scheme, [40] integrates a two-round OT-based multiplier. Although the OT-based multiplier requires a one-time setup, both schemes do not use precomputed values per signing request. This is in contrast to our scheme but at the cost of requiring several rounds of communication during the signing. Parts of their protocols are independent of the message that will be signed; thus, in principle, these steps can be moved to a presigning phase. In this case, the signing phase is non-interactive, but on the downside, the communication complexity of the presigning phase has linear complexity. In contrast, our protocol achieves both a non-interactive online phase and an offline phase with sublinear complexity. In addition, both works [44, 40] consider a security model tailored to the BBS+ signature scheme while we show security with respect to a more generic threshold signature ideal functionality.

Wong et al. [68] presented another threshold BBS+ signature scheme. However, their work focuses on providing strong robustness using a linearly homomorphic encryption scheme with threshold decryption. The cost of achieving this property is a four-round protocol.

Another anonymous credential scheme with threshold issuance, called Coconut, is proposed by Sonnino et al. [63] and analyzed and improved by Rial and Piotrowska [60]. Their scheme is based on the Pointcheval-Sanders (PS) signature scheme, which allows them to have a non-interactive issuance phase without coordination or precomputation. We emphasize that the PS signature scheme is less popular than BBS+ and not subject to standardization efforts. The security of PS and Coconut is based on a modified variant of the LRSW assumption introduced in [59]. This assumption is interactive in contrast to the q -Strong Diffie-Hellman assumption on which the security of BBS+ is based. While PS and Coconut also support multi-attribute credentials, the secret and public key size increases linearly in the number of attributes. In BBS+, the key size is constant. Further, PS and, therefore, the Coconut scheme relies on Type-3 pairings, while our scheme can be instantiated with any pairing type. The security of Coconut was not shown under concurrent composition while our scheme is analyzed in the Universal Composability framework.

Like our work, [3] and [48] leverage pseudorandom correlations for threshold signatures. [3] presents an ECDSA scheme, while [48] focuses on Schnorr signatures. [3] constructs a tailored PCG generating ECDSA- presignatures while our scheme uses existing VOLE PCGs/PCFs in a black-box way. Further, in contrast to our work, [3] presents an n -out-of- n protocol without a flexible threshold. [48]

introduces the new notion of a discrete log PCF and constructs a two-party protocol based on this primitive. In contrast to our work, [48] captures only the 2-out-of-2 setting. Both schemes [3, 48] require additional per-presignature communication. Depending on the phase this communication is assigned to, the schemes either have linear communication in the offline phase or require two rounds of communication in the online phase.

2 Preliminaries

Throughout this work, we denote the security parameter by $\lambda \in \mathbb{N}$, the set $\{1, \dots, k\}$ as $[k]$, the set $\{0, 1, \dots, k\}$ as $[0..k]$, the number of parties by n and a specific party by P_i . The set of indices of corrupted parties is denoted by $\mathcal{C} \subsetneq [n]$ and honest parties are denoted by $\mathcal{H} = [n] \setminus \mathcal{C}$. We denote vectors of elements via bold letters, e.g., \mathbf{a} , and the i -th element of a vector \mathbf{a} by $\mathbf{a}[i]$.

We model our protocol in the Universal Composability (UC) framework by Canetti [28]. We refer to Appendix C for a brief introduction to UC. In our constructions, we denote by \mathcal{Z} the UC environment and use `sid` and `ssid` to denote session and subsession identifier. We consider a malicious adversary statically corrupting up to $t - 1$ parties and the synchronous communication model.

We make use of a bilinear mapping following the definition of [12, 11]. A bilinear mapping is described by three cyclic groups $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T)$ of prime order p , generators $g_1 \in \mathbb{G}_1, g_2 \in \mathbb{G}_2$, and a pairing $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$. We call e a bilinear map iff it can be computed efficiently, $e(u^a, v^b) = e(u, v)^{ab}$ for all $(u, v, a, b) \in \mathbb{G}_1 \times \mathbb{G}_2 \times \mathbb{Z}_p \times \mathbb{Z}_p$, and $e(g_1, g_2) \neq 1$ for all generators g_1 and g_2 . We refer to [12] for a more formal specification.

3 Reusable Pseudorandom Correlation Generators

In this section we introduce our definition of reusable PCGs, extending the definition of programmable PCGs from [17] and [19]. We argue why existing constructions for PCGs with minor modifications satisfy our new definition in Appendix E. The extended definitional framework for PCFs, constructions, and the PCF-based instantiation of the precomputation are stated in Appendix F-H.

In a nutshell, pseudorandom correlation generators allow two parties to generate a large amount of correlated randomness from short seeds. They are useful in many two- and multiparty protocols in the offline-online-model [37, 56, 65, 66]. Examples for frequently used correlations are vector oblivious linear evaluations, oblivious transfer and multiplication triples.

Our modifications and extensions of the definitions of [16] and [19] reflect the challenges we faced when using PCGs as black-box primitives in our threshold BBS+ protocol. We present our definition and highlight these challenges and changes in the following. We note that Boyle et al. [17] presents an ideal functionality for corruptible, correlated randomness which can be instantiated by PCGs. While this notion allows to abstract from concrete PCG constructions, their ideal functionality does not cover the reusability feature required in our

setting. Therefore, we present a suitable definition extending the game-based definitions of [16] and [19].

3.1 Definition

As mentioned above, a PCF/PCG realizes a target correlation \mathcal{Y} . For some correlations, like VOLE, parts of the correlation outputs are fixed over all outputs. In the example of VOLE, where the correlation is $v = as + u$ over some ring R , the s value is fixed for all tuples.

Additionally, in a multiparty setting, we like PCG/PCF constructions that allow parties to obtain the same values for parts of the correlation output in multiple instances. Concretely, assume party P_i evaluates one VOLE PCG/PCF instance with party P_j and one with party P_k . P_i evaluates the PCG/PCF to $(a_{i,j}, u_{i,j})$ for the first instance and $(a_{i,k}, u_{i,k})$ for the second instance. Here, we want to give party P_i the opportunity to get $a_{i,j} = a_{i,k}$ when applied on the same input. This property is necessary to construct multiparty correlations from two-party PCG/PCF instances.

To formally model the abovementioned properties, we define a *target correlation* as a tuple of probabilistic algorithms $(\text{Setup}, \mathcal{Y})$, where **Setup** takes two inputs and creates a master key mk . These inputs enable fixing parts of the correlation, e.g., the fixed value s in VOLE correlations, and obtaining the same values over multiple instances, e.g., by reusing the values s in multiple VOLE correlations. Algorithm \mathcal{Y} uses the master key to sample correlation outputs.

Finally, we follow [16, 19] and require a target correlation to be reverse-sampleable to facilitate a suitable definition of PCGs. In contrast to [17, 19], our definition of a target correlation explicitly considers the reusability of values over multiple invocations.

Definition 1 (Reverse-sampleable target correlation with setup). *Let $\ell_0(\lambda), \ell_1(\lambda) \leq \text{poly}(\lambda)$ be output length functions. Let $(\text{Setup}, \mathcal{Y})$ be a tuple of probabilistic algorithms, such that **Setup** on input 1^λ and two parameters ρ_0, ρ_1 returns a master key mk ; algorithm \mathcal{Y} on input 1^λ and mk returns a pair of outputs $(y_0^{(i)}, y_1^{(i)}) \in \{0, 1\}^{\ell_0(\lambda)} \times \{0, 1\}^{\ell_1(\lambda)}$.*

*We say that the tuple $(\text{Setup}, \mathcal{Y})$ defines a reverse-sampleable target correlation with setup if there exists a probabilistic polynomial time algorithm **RSample** that takes as input $1^\lambda, \text{mk}, \sigma \in \{0, 1\}, y_\sigma^{(i)} \in \{0, 1\}^{\ell_\sigma(\lambda)}$ and outputs $y_{1-\sigma}^{(i)} \in \{0, 1\}^{\ell_{1-\sigma}(\lambda)}$, such that for all $\sigma \in \{0, 1\}$, for all mk, mk' in the range of **Setup** for arbitrary but fixed input ρ_σ the following distributions are statistically close:*

$$\begin{aligned} & \{(y_0, y_1) \mid (y_0, y_1) \xleftarrow{\$} \mathcal{Y}(1^\lambda, \text{mk})\} \\ & \{(y_0, y_1) \mid (y'_0, y'_1) \xleftarrow{\$} \mathcal{Y}(1^\lambda, \text{mk}'), \\ & \quad y_\sigma \leftarrow y'_\sigma, y_{1-\sigma} \leftarrow \text{RSample}(1^\lambda, \text{mk}, \sigma, y_\sigma)\}. \end{aligned}$$

Given the definition of a reverse-sampleable correlation with setup, we define our primitive called *reusable PCG* (*rPCG*).

The security properties in the original notion of programmable PCGs assumes randomly selected seeds that are inserted into the key generation. This reflects a passive or semi-honest setting in which the adversary cannot deviate from the protocol description such that the seeds are indeed random. We are interested in the active security setting, where an adversary can insert arbitrary seeds into the key generation. Therefore, we adapted the pseudorandom \mathcal{Y} -correlated outputs and the security property. Additionally, we introduced a key indistinguishability property. This property is helpful in security proofs of protocols where we want to use PCGs in a black-box way. In particular, it allows us to switch from fake to real PCG seeds/keys and show indistinguishability between two hybrids via a reduction to the property. Next, we formally propose the notion of *rPCG*.

Definition 2 (Reusable pseudorandom correlation generator (rPCG)).

Let $(\text{Setup}, \mathcal{Y})$ be a reverse-sampleable correlation with setup which has output length functions $\ell_0(\lambda), \ell_1(\lambda)$, let $\lambda \leq \eta(\lambda) \leq \text{poly}(\lambda)$ be the sample size function. Let $(\text{PCG.Gen}_p, \text{PCG.Expand})$ be a pair of algorithms with the following syntax:

- $\text{PCG.Gen}_p(1^\lambda, \rho_0, \rho_1)$ is a probabilistic polynomial-time algorithm that on input the security parameter 1^λ and reusable inputs ρ_0, ρ_1 outputs a pair of keys (k_0, k_1) .
- $\text{PCG.Expand}(\sigma, k_\sigma)$ is a deterministic polynomial-time algorithm that on input $\sigma \in \{0, 1\}$ and key k_σ outputs $\mathbf{y}_\sigma \in \{0, 1\}^{\ell_\sigma(\lambda) \times \eta(\lambda)}$, i.e. an array of size $\eta(\lambda)$ with elements being bit-strings of length $\ell_\sigma(\lambda)$.

We say $(\text{PCG.Gen}_p, \text{PCG.Expand})$ is a reusable pseudorandom correlation generator (rPCG) for $(\text{Setup}, \mathcal{Y})$, if the following conditions hold:

- **Programmability.** There exist public efficiently computable functions ϕ_0, ϕ_1 , such that for all $\rho_0, \rho_1 \in \{0, 1\}^*$

$$\Pr \left[\begin{array}{l} (k_0, k_1) \xleftarrow{\$} \text{PCG.Gen}_p(1^\lambda, \rho_0, \rho_1) \\ (\mathbf{x}_0, \mathbf{z}_0) \leftarrow \text{PCG.Expand}(0, k_0), \\ (\mathbf{x}_1, \mathbf{z}_1) \leftarrow \text{PCG.Expand}(1, k_1) \end{array} : \begin{array}{l} \mathbf{x}_0 = \phi_0(\rho_0) \\ \mathbf{x}_1 = \phi_1(\rho_1) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

- **Pseudorandom \mathcal{Y} -correlated outputs.** For every non-uniform adversary \mathcal{A} of size $\text{poly}(\lambda)$ it holds that

$$\left| \Pr[\text{Exp}_{\mathcal{A}}^{\text{pr-g}}(\lambda) = 1] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

for all sufficiently large λ , where $\text{Exp}_{\mathcal{A}}^{\text{pr-g}}(\lambda)$ is as defined in Figure 1.

- **Security.** For each $\sigma \in \{0, 1\}$ and non-uniform adversary \mathcal{A} of size $\text{poly}(\lambda)$, it holds that

$$\left| \Pr[\text{Exp}_{\mathcal{A}, \sigma}^{\text{sec-g}}(\lambda) = 1] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

for all sufficiently large λ , where $\text{Exp}_{\mathcal{A}, \sigma}^{\text{sec-g}}(\lambda)$ is as defined in Figure 1.

- **Key indistinguishability.** For any $\sigma \in \{0, 1\}$ and non-uniform adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, it holds

$$\Pr[\text{Exp}_{\mathcal{A}, \sigma}^{\text{key-ind-g}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda)$$

for all sufficiently large λ , where $\text{Exp}_{\mathcal{A}, \sigma}^{\text{key-ind-g}}$ is as defined in Figure 1.

$\text{Exp}_{\mathcal{A}}^{\text{pr-g}}(\lambda) :$ $b \xleftarrow{\$} \{0, 1\}, N \leftarrow \eta(\lambda), (\rho_0, \rho_1) \leftarrow \mathcal{A}_0(1^\lambda)$ $\text{mk} \xleftarrow{\$} \text{Setup}(1^\lambda, \rho_0, \rho_1)$ $(k_0, k_1) \xleftarrow{\$} \text{PCG.Gen}_p(1^\lambda, \rho_0, \rho_1)$ $\text{if } b = 0 \text{ then } (y_0, y_1) \xleftarrow{\$} \mathcal{Y}(1^\lambda, \text{mk})$ $\text{else } y_\sigma \leftarrow \text{PCG.Expand}(\sigma, k_\sigma) \text{ for } \sigma \in \{0, 1\}$ $b' \leftarrow \mathcal{A}_1(1^\lambda, y_0, y_1), \text{ return } b' = b$	$\text{Exp}_{\mathcal{A}, \sigma}^{\text{key-ind-g}}(\lambda) :$ $b \xleftarrow{\$} \{0, 1\}, \rho_\sigma \leftarrow \mathcal{A}_0(1^\lambda)$ $\rho_{1-\sigma}^{(0)}, \rho_{1-\sigma}^{(1)} \xleftarrow{\$} \{0, 1\}^*, \rho_{1-\sigma} \leftarrow \rho_{1-\sigma}^{(b)}$ $(k_0, k_1) \leftarrow \text{PCG.Gen}_p(1^\lambda, \rho_0, \rho_1)$ $b' \leftarrow \mathcal{A}_1(1^\lambda, k_\sigma, \rho_{1-\sigma}^{(0)}) \text{ return } b' = b$
$\text{Exp}_{\mathcal{A}, \sigma}^{\text{sec-g}}(\lambda) :$ $b \xleftarrow{\$} \{0, 1\}, N \leftarrow \eta(\lambda), (\rho_0, \rho_1) \leftarrow \mathcal{A}_0(1^\lambda)$ $\text{mk} \xleftarrow{\$} \text{Setup}(1^\lambda, \rho_0, \rho_1), (k_0, k_1) \xleftarrow{\$} \text{PCG.Gen}_p(1^\lambda, \rho_0, \rho_1)$ $y_\sigma \xleftarrow{\$} \text{PCG.Expand}(\sigma, k_\sigma)$ $\text{if } b = 0 \text{ then } (y_{1-\sigma}) \leftarrow \text{PCG.Expand}(1-\sigma, k_{1-\sigma})$ $\text{else } y_{1-\sigma} \leftarrow \text{RSample}(1^\lambda, \text{mk}, \sigma, y_\sigma)$ $b' \leftarrow \mathcal{A}_1(1^\lambda, y_0, y_1), \text{ return } b' = b$	

Fig. 1: Security games for reusable PCGs.

3.2 Correlations

Our VOLE correlation of size N over a finite field \mathbb{F} is given by $\mathbf{z}_1 = \mathbf{x}_0 \cdot x_1 + \mathbf{z}_0$, where $\mathbf{z}_1, \mathbf{x}_0, \mathbf{z}_0 \in \mathbf{F}_p^N$ and x_1 is a fixed scalar in \mathbf{F}_p . Moreover, we require \mathbf{x}_0 being computed by a pseudorandom generator (PRG). We formally define the reverse-sampleable target correlation with setup $(\text{Setup}_{\text{VOLE}}, \mathcal{Y}_{\text{VOLE}})$ of size N over field \mathbb{F}_p as

$$\begin{aligned} \text{mk} &= (\rho, x_1) \leftarrow \text{Setup}_{\text{VOLE}}(1^\lambda, \rho, x_1), \\ ((F(\rho), \mathbf{z}_0), (x_1, \mathbf{z}_1)) &\leftarrow \mathcal{Y}_{\text{VOLE}}(1^\lambda, \text{mk}) \quad \text{such that} \\ \mathbf{z}_0 &\xleftarrow{\$} \mathbb{F}_p^N \text{ and } \mathbf{z}_1 = F_0(\rho_0) \cdot x_1 + \mathbf{z}_0, \end{aligned}$$

where F being a PRG. The reverse-sampling algorithm is defined such that $(x_1, F(\rho) \cdot \mathbf{x}_1 + \mathbf{z}_0) \leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, \text{mk}, 0, (F(\rho), \mathbf{z}_0))$ and $(F(\rho), \mathbf{z}_1 - F(\rho) \cdot x_1) \leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, \text{mk}, 1, (x_1, \mathbf{z}_1))$.

We state a PCG construction realizing this definition of a VOLE correlation in Appendix E.

4 Threshold Online Protocol

In this section, we present our threshold BBS+ protocol. This protocol yields a signing phase without interaction between the signers and a flexible threshold parameter t .

We show the security of our protocol against a malicious adversary statically corrupting up to $t - 1$ parties in the UC framework. We show that our scheme implements a modification of the generic ideal functionality for threshold signature schemes introduced by Canetti et al. [29]. We deliberately chose the generic threshold signature functionality by Canetti et al. [29] over a specific BBS+ functionality such as the one used in [40]. Proving security under a generic threshold functionality enables our threshold BBS+ protocol to be used whenever a threshold signature scheme is required (e.g., for the construction of a more complex protocol such as an anonymous credential system). We present the ideal functionality and discuss the changes with respect to the original version in Appendix I.

Our protocol uses precomputation to accelerate online signing and requires a random oracle. An intuitive description of the random oracle usage and the precomputation used is given in Section 1.2. We formally model the random oracle and the precomputation by describing our protocol in a hybrid model where parties can access a global random oracle functionality \mathcal{H} and a hybrid preprocessing functionality $\mathcal{F}_{\text{Prep}}$. Using a hybrid model allows us to abstract from the concrete instantiation of the preprocessing functionality. We present concrete instantiations of $\mathcal{F}_{\text{Prep}}$ in Section 5 and Appendix H.

4.1 Global Random Oracle Functionality

We follow the formalism of Canetti et al. [29] to model a global random oracle in the UC framework. This model constitutes an abstraction of a hash function that is used within and outside of the analyzed system. We refer the reader to [29] and references within for further discussions.

Functionality \mathcal{H}

Parameter: Output domain \mathbb{Z}_p .

Upon receiving (query, x) from party P , do:

- If a tuple (x, a) is stored, then output (answer, a) to P .
- Else, sample $a \leftarrow \mathbb{Z}_p$ and store (x, a) .

Output (answer, a) to P .

4.2 Ideal Preprocessing Functionality

The preprocessing functionality consists of two phases. First, the *Initialization* phase samples a private/public key pair. Second, the *Tuple* phase provides correlated tuples upon request. In the second phase, the output values of the honest parties are reverse sampled, given the corrupted parties' outputs. To explicitly model the Tuple phase as non-interactive, we require the simulator to specify a function **Tuple** during the Initialization. This function defines the corrupted parties' output values in the Tuple phase and is computed first to reverse sample the honest parties' outputs.

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

The functionality $\mathcal{F}_{\text{Prep}}$ interacts with parties P_1, \dots, P_n and ideal-world adversary \mathcal{S} . The functionality is parameterized by a threshold parameter t . During the initialization, \mathcal{S} provides a tuple function $\text{Tuple}(\cdot, \cdot, \cdot) \rightarrow \mathbb{Z}_p^2$.

Initialization. Upon receiving **(init, sid)** from all parties,

1. Sample the secret key $\text{sk} \xleftarrow{\$} \mathbb{Z}_p$.
2. Send $\text{pk} = (g_2^{\text{sk}})$ to \mathcal{S} . Upon receiving **(ok, Tuple(\cdot, \cdot, \cdot))** from \mathcal{S} , send pk to every honest party.

Tuple. On input **(tuple, sid, ssid, \mathcal{T})** from party P_i where $i \in \mathcal{T}$, $\mathcal{T} \subseteq [n]$ of size t do:

- If $(\text{ssid}, \mathcal{T}, \{(a_\ell, \beta_\ell)\}_{\ell \in \mathcal{T}})$ is stored, send (a_i, β_i) to P_i .
Else, compute $(a_j, \beta_j) \leftarrow \text{Tuple}(\text{ssid}, \mathcal{T}, j)$ for every corrupted party P_j where $j \in \mathcal{C} \cap \mathcal{T}$. Next, sample $a \xleftarrow{\$} \mathbb{Z}_p$ and tuples (a_j, β_j) over \mathbb{Z}_p for $j \in \mathcal{H} \cap \mathcal{T}$ such that

$$\sum_{\ell \in \mathcal{T}} a_\ell = a \quad \sum_{\ell \in \mathcal{T}} \beta_\ell = a \cdot \text{sk} \quad (1)$$

Store $(\text{sid}, \text{ssid}, \mathcal{T}, \{(a_\ell, \beta_\ell)\}_{\ell \in \mathcal{T}})$ and send $(\text{sid}, \text{ssid}, a_i, \beta_i)$ to honest P_i .

Abort. On input **(abort, sid)** from \mathcal{S} , send **abort** to all honest parties and halt.

4.3 Online Signing Protocol

Next, we formally state our threshold BBS+ protocol. We refer the reader to the technical overview in Section 1.2 for a high-level description of our protocol. Further, we discuss extensions for anonymous credentials systems and blind signing in Appendix D.

Construction 1: $\pi_{\text{TBS}+}$

We describe the protocol from the perspective of an honest party P_i .

Public Parameters. Number of parties n , maximal number of signatures N , size of message arrays k , security threshold t , a bilinear mapping tu-

ple $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, g_1, g_2, \mathbf{e})$ and randomly sampled \mathbb{G}_1 elements $\{h_\ell\}_{\ell \in [k]}$. Let $\text{Verify}_{\text{pk}}(\cdot, \cdot)$ be the BBS+ verification algorithm as defined in Appendix A.

KeyGen.

- Upon receiving $(\text{keygen}, \text{sid})$ from \mathcal{Z} , send $(\text{init}, \text{sid})$ to $\mathcal{F}_{\text{Prep}}$ and receive pk in return.
- Upon receiving $(\text{pubkey}, \text{sid})$ from \mathcal{Z} output $(\text{pubkey}, \text{sid}, \text{Verify}_{\text{pk}}(\cdot, \cdot))$.

Sign. Upon receiving $(\text{sign}, \text{sid}, \text{ssid} \in [N], \mathcal{T}, \mathbf{m} = \{m_\ell\}_{\ell \in [k]})$ from \mathcal{Z} with $P_i \in \mathcal{T}$ and no tuple $(\text{sid}, \text{ssid}, \cdot)$ is stored, perform the following steps:

1. Send $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ to $\mathcal{F}_{\text{Prep}}$ and receive tuple $(\text{sid}, \text{ssid}, a_i, \beta_i)$.
2. Store $(\text{sid}, \text{ssid}, \mathbf{m})$, send $(\text{query}, \text{sid} \parallel \text{ssid} \parallel \mathbf{m} \parallel \mathcal{T})$ to \mathcal{H} and receive (answer, e) . Then, send $(\text{sid}, \text{ssid}, \mathcal{T}, A_i := (g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{a_i}, \delta_i := \beta_i + a_i \cdot e)$ to each party $P_j \in \mathcal{T}$.
3. Once $(\text{sid}, \text{ssid}, \mathcal{T}, A_j, \delta_j)$ is received from every party $P_j \in \mathcal{T} \setminus \{P_i\}$,
 - (a) Send $(\text{query}, \text{sid} \parallel \text{ssid} \parallel \mathbf{m} \parallel \mathcal{T})$ to \mathcal{H} and receive (answer, e) . Next, compute $\epsilon = (\sum_{\ell \in \mathcal{T}} \delta_\ell)^{-1}$ and $A = (\prod_{\ell \in \mathcal{T}} A_\ell)^\epsilon$.
 - (b) If $\text{Verify}_{\text{pk}}(\mathbf{m}, (A, e)) = 1$, set $\text{out} = \sigma = (A, e)$. Otherwise, set $\text{out} = \text{abort}$. Then, output $(\text{sig}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, \text{out})$.

Verify. Upon receiving $(\text{verify}, \text{sid}, \mathbf{m} = \{m_\ell\}_{\ell \in [k]}, \sigma, \text{Verify}_{\text{pk}'}(\cdot, \cdot))$ from \mathcal{Z} output $(\text{verified}, \text{sid}, \mathbf{m}, \sigma, \text{Verify}_{\text{pk}'}(\mathbf{m}, \sigma))$.

Remark. While we simplified our UC model to capture the scenario where every signer obtains the final signature, we expect real-world scenarios to have a dedicated client which is the only party to obtain the signature. In the latter case, the signers send the partial signature in Step 2 only to the client and Steps 3a and 3b are performed by the client. We stress that in both cases the communication follows a request-response pattern which is the minimum for MPC protocols. Moreover, note that the $(\text{tuple}, \cdot, \cdot, \cdot)$ -call to $\mathcal{F}_{\text{Prep}}$ does not involve additional communication when being instantiated based on PCGs or PCFs as done in this work. Using such an instantiation, the $(\text{tuple}, \cdot, \cdot, \cdot)$ -call is realized by local evaluation of the PCF or local expansion of the PCG so that no interaction between the parties is needed.

Theorem 1. *Assuming the strong unforgeability of BBS+, protocol $\pi_{\text{TBSB}+}$ UC-realizes $\mathcal{F}_{\text{tsig}}$ in the $(\mathcal{F}_{\text{Prep}}, \mathcal{H})$ -hybrid model in the presence of malicious adversaries controlling up to $t - 1$ parties.*

The proof is given in Appendix J.

5 PCG-based Threshold Preprocessing Protocol

We state our threshold BBS+ signing protocol in Section 4 in a $\mathcal{F}_{\text{Prep}}$ -hybrid model. Now, we present an instantiation of the $\mathcal{F}_{\text{Prep}}$ functionality using pseudorandom correlation generators (PCGs). Our $\pi_{\text{Prep}}^{\text{PCG}}$ protocol builds solely on PCGs for VOLE correlations. The resulting protocol consists of an interactive *Initialization* and a non-interactive *Tuple* phase, consisting only of the retrieval of stored PCG tuples and additional local computation.

Our preprocessing protocol consists of four steps: the first three are part of the Initialization phase, and the fourth one builds the Tuple phase. First, the parties set up a secret and corresponding public key. For the BBS+ signature scheme, the public key is $\mathbf{pk} = g_2^x$, while the secret key is $\mathbf{sk} = x$, which is secret-shared using Shamir's secret sharing. This procedure constitutes a standard distributed key generation protocol for a DLOG-based cryptosystem. Therefore, we abstract from the concrete instantiation of this protocol and model the key generation as a hybrid functionality \mathcal{F}_{KG} . Second, the parties set up the keys for the PCG instances. The protocol uses two-party PCGs, meaning each pair of parties sets up required instances. We model the PCG key generation as a hybrid functionality $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$. Third, every party expands the local seeds to the required VOLE correlations and store them in their storage. The fourth step constitutes the Tuple phase and is executed by every party in the signer set \mathcal{T} of a signing request. In this phase party P_i generates (a_i, β_i) , where the values fulfill correlation (1). For a signing request with ssid , P_i uses the ssid -th component of the previously expanded correlation vector \mathbf{a} as a_i . Note that the a_i values constitute an additive secret sharing of a value a (cf. (1)). Then, $\sum_{\ell \in \mathcal{T}} \beta_\ell = a \cdot \mathbf{sk}$ can be rewritten as $a \cdot \mathbf{sk} = \sum_{\ell \in \mathcal{T}} a_\ell \cdot \sum_{j \in \mathcal{T}} L_{j, \mathcal{T}} \cdot \mathbf{sk}_j = \sum_{\ell \in \mathcal{T}} \sum_{j \in \mathcal{T}} a_\ell \cdot \mathbf{sk}_j \cdot L_{j, \mathcal{T}}$. Recall that the \mathbf{sk}_j values constitute a Shamir sharing of \mathbf{sk} so that we need to integrate LaGrange coefficients for interpolation. The multiplication $a_\ell \mathbf{sk}_j$ is equal to the additive shares of a VOLE correlation, i.e., $c_1 - c_0 = a_\ell \mathbf{sk}_j$. For $\ell \neq j$, the parties obtain these additive shares via the stored VOLE correlations that were expanded in the third step. Note that the parties use again the ssid -th component of the vectors to get consistent values. Finally, party P_i applies the appropriate LaGrange coefficients to the output of its PCG expansion and adds $a_i \cdot \mathbf{sk}_i \cdot L_{i, \mathcal{T}}$ to get an additive sharing of $a \cdot \mathbf{sk}$. Note that while the value a is a fresh random values for each signing request, \mathbf{sk} is fixed. Therefore, the parties can use VOLE correlations to compute $a \cdot \mathbf{sk}$.

Note that party P_i uses PCG instances for computing additive shares of $a_i \mathbf{sk}_j$ and $a_i \mathbf{sk}_\ell$ for two different parties P_j and P_ℓ . Since a_i must be the same for both products, we use reusable PCGs so parties can fix a_i over multiple PCG instances. Based on these two requirements, our protocol relies on reusable PCGs defined in Section 3.

Key Generation Functionality. We abstract from the concrete instantiation of the key generation. Therefore, we state a very simple key generation functionality for discrete logarithm-based cryptosystems similar to the functionality of [67]. The functionality describes a standard distributed key generation for discrete logarithm-based cryptosystems and can be realized by [45, 67] or the key generation phase of [29] or [40].

Functionality \mathcal{F}_{KG}

The functionality is parameterized by the order of the group from which the secret key is sampled p , a generator for the group of the public key g_2 , and a

threshold parameter t . The key generation functionality interacts with parties P_1, \dots, P_n and ideal-world adversary \mathcal{S} .

Key Generation. Upon receiving $(\text{keygen}, \text{sid})$ from every party P_i and $(\text{corruptedShares}, \text{sid}, \{\text{sk}_j\}_{j \in \mathcal{C}})$ from \mathcal{S} :

1. Sample random polynomial $F \in \mathbb{Z}_p[X]$ of degree $t-1$ such that $F(j) = \text{sk}_j$ for every $j \in \mathcal{C}$.
2. Set $\text{sk} = F(0)$, $\text{pk} = g_2^{\text{sk}}$, $\text{sk}_\ell = F(\ell)$ and $\text{pk}_\ell = g_2^{\text{sk}_\ell}$ for $\ell \in [n]$.
3. Send $(\text{sid}, \text{sk}_i, \text{pk}, \{\text{pk}_\ell\}_{\ell \in [n]})$ to every party P_i .

Setup Functionality. The setup functionality gets random values, secret key shares, and partial public keys as input from every party. Then, it first checks if the secret key shares and the partial public keys match and next generates the PCG keys using the random values. Finally, it returns the generated PCG keys to the parties.

In order to provide modularity, we abstract from concrete instantiation by specifying this functionality. Nevertheless, $\mathcal{F}_{\text{Setup}}$ can be instantiated using general-purpose MPC or tailored protocols similar to distributed seed generation protocols from prior work [19, 3]. We leave a formal specification of a tailored protocol as future work.

Functionality $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$

Let $\text{PCG}_{\text{VOLE}}.\text{Gen}_p$ be an rPCG for VOLE correlations. The setup functionality interacts with parties P_1, \dots, P_n .

Setup. Upon receiving $(\text{setup}, \text{sid}, \rho_a^{(i)}, \text{sk}_i, \{\text{pk}_\ell^{(i)}\}_{\ell \in [n]})$ from every party P_i :

1. Check if $g_2^{\text{sk}_i} = \text{pk}_i^{(i)}$ for every $\ell, i \in [n]$. If the check fails, send **abort** to all parties.
 Else, compute $(k_{i,j,0}^{\text{VOLE}}, k_{i,j,1}^{\text{VOLE}}) \leftarrow \text{PCG}_{\text{VOLE}}.\text{Gen}_p(1^\lambda, \rho_a^{(i)}, \text{sk}_j)$ for each tuple (P_k, P_j) with $k, j \in [n], k \neq j$.
2. Send keys $(\text{sid}, \{k_{i,j,0}^{\text{VOLE}}, k_{j,i,1}^{\text{VOLE}}\}_{j \in [n] \setminus i})$ to party P_i for $i \in [n]$.

PCG-based Preprocessing Protocol. In this section, we formally present our PCG-based preprocessing protocol in the $(\mathcal{F}_{\text{KG}}, \mathcal{F}_{\text{Setup}}^{\text{PCG}})$ -hybrid model.

Construction 2: $\pi_{\text{Prep}}^{\text{PCG}}$

Let $(\text{PCG}_{\text{VOLE}}.\text{Gen}_p, \text{PCG}_{\text{VOLE}}.\text{Expand})$ be an rPCG for VOLE correlations.

We describe the protocol from the perspective of P_i .

Initialization. Upon receiving input $(\text{init}, \text{sid})$, do:

1. Send $(\text{keygen}, \text{sid})$ to \mathcal{F}_{KG} .
2. Upon receiving $(\text{sid}, \text{sk}_i, \text{pk}, \{\text{pk}_\ell^{(i)}\}_{\ell \in [n]})$ from \mathcal{F}_{KG} , sample $\rho_a^{(i)} \in \{0, 1\}^\lambda$ and send $(\text{setup}, \text{sid}, \rho_a^{(i)}, \text{sk}_i, \{\text{pk}_\ell^{(i)}\}_{\ell \in [n]})$ to $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$.
3. Upon receiving $(\text{sid}, \{k_{i,j,0}^{\text{VOLE}}, k_{j,i,1}^{\text{VOLE}}\}_{j \neq i})$ from $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$, compute and store for every $j \in [N] \setminus \{i\}$:

- (a) $(\mathbf{a}_i, \mathbf{c}_{i,j,0}^{\text{VOLE}}) = \text{PCG}_{\text{VOLE}}.\text{Expand}(0, \mathbf{k}_{i,j,0}^{\text{VOLE}})$,
 - (b) $(\mathbf{sk}_i, \mathbf{c}_{j,i,1}^{\text{VOLE}}) = \text{PCG}_{\text{VOLE}}.\text{Expand}(1, \mathbf{k}_{j,i,0}^{\text{VOLE}})$,
 - 4. Output \mathbf{pk} .
 - Tuple.** Upon receiving input $(\mathbf{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$, compute:
 - 5. Let $a_i = \mathbf{a}_i[\text{ssid}]$, $c_{(i,j,0)}^{\text{VOLE}} = \mathbf{c}_{(i,j,0)}^{\text{VOLE}}[\text{ssid}]$, $c_{(j,i,1)}^{\text{VOLE}} = \mathbf{c}_{(j,i,1)}^{\text{VOLE}}[\text{ssid}]$ for $j \in \mathcal{T} \setminus \{i\}$.
 - 6. Compute $\beta_i = a_i \cdot L_{i,\mathcal{T}} \cdot \mathbf{sk}_i + \sum_{j \in \mathcal{T} \setminus \{i\}} (L_{i,\mathcal{T}} \cdot c_{j,i,1}^{\text{VOLE}} - L_{j,\mathcal{T}} \cdot c_{i,j,0}^{\text{VOLE}})$
 - 7. Output $(\text{sid}, \text{ssid}, a_i, \beta_i)$.

Theorem 2. Let PCG_{VOLE} be an $r\text{PCG}$ for VOLE correlations. Then, protocol $\pi_{\text{Prep}}^{\text{PCG}}$ UC-realizes $\mathcal{F}_{\text{Prep}}$ in the $(\mathcal{F}_{\text{KG}}, \mathcal{F}_{\text{Setup}}^{\text{PCG}})$ -hybrid model in the presence of malicious adversaries controlling up to $t - 1$ parties.

We state our simulator, a proof sketch and the full indistinguishability proof in Appendix K, L, and M.

6 Evaluation

This section presents the evaluation of the online and the offline phase of our protocol and compares it to related work. We refer the reader to Appendix N for an evaluation of our threshold protocol based on the legacy BBS+ scheme (cf. Appendix B).

6.1 Online, Signing Request-Dependent Phase

We evaluate the online, signing request-dependent phase by implementing the protocol, running benchmarks, and reporting the runtime and the communication complexity. For comparison, we also implement and benchmark the non-threshold BBS+ signing algorithm. We open-source our prototype implementation to foster future research in this area.¹

Implementation and Experimental Setup. Our implementation and benchmarks of the online phase are written in Rust (compiled using rustc 1.68.2 (9eb3afe9e)) and based on the BLS12_381 curve.² Note, since the BLS12_381 curve defines an elliptic curve, we use the additive group notation in the following. This differs from the multiplicative group notation used in the protocol description. Our code, including the benchmarks and rudimentary tests, comprises 1,400 lines.

For our benchmarks, we split the protocol into four phases: *Adapt* (Steps 6 of protocol $\pi_{\text{Prep}}^{\text{PCG}}$), *Sign* (Step 2 of $\pi_{\text{BBS+}}$), *Reconstruct* (Step 3a of $\pi_{\text{BBS+}}$) and *Verify* (Step 3b of $\pi_{\text{BBS+}}$). Adapt and Sign are executed by the servers. Reconstruct and Verify are executed by the client. Together, these phases cover the whole online signing protocol. The runtime of our protocol is influenced by the security threshold t and the message array size k . We perform benchmarks

¹ <https://github.com/AppliedCryptoGroup/NI-Threshold-BBS-Plus-Code>

² We have used [4] for all curve operations.

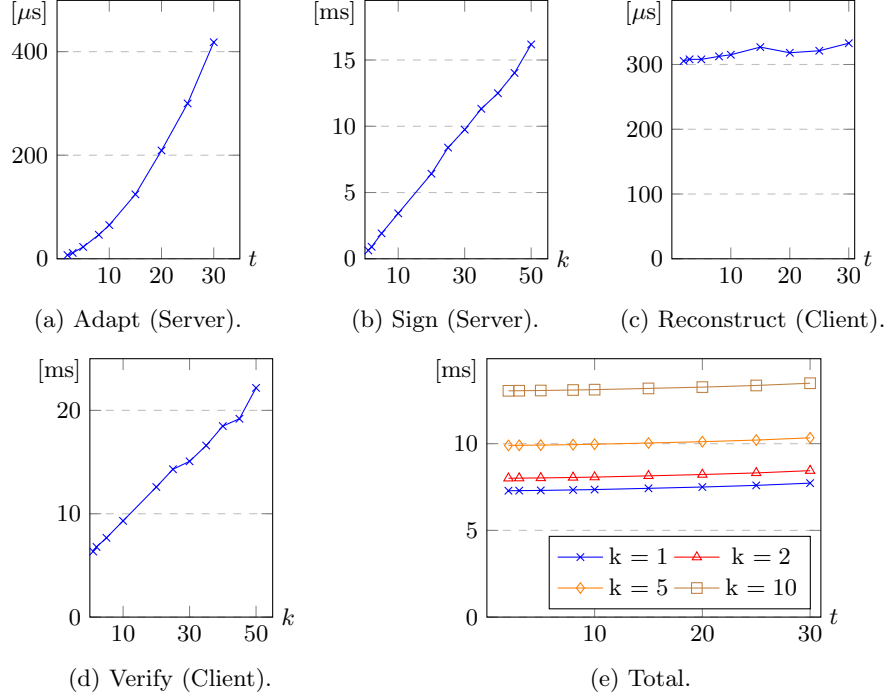


Fig. 2: The runtime of individual protocol phases (a)-(d) and the total online protocol (e). The *Adapt* phase, describing Steps 6 of protocol $\pi_{\text{Prep}}^{\text{PCG}}$, and the *Reconstruct* phase, describing Step 3a of π_{TBS^+} , depend on security threshold t . The *Sign* phase, describing Step 2 of π_{TBS^+} , and the signature verification, describing Step 3b of π_{TBS^+} , depend on the message array size k .

for $2 \leq t \leq 30$ and $1 \leq k \leq 50$. The range for parameter t is chosen to provide comparability with [40] and we deem $k \leq 50$ a realistic setting for the use-cases of credential certificates. Moreover, both ranges illustrate the trend for increasing parameters. The influence of the total number of servers n is insignificant to non-existent. Our benchmarks do not account for network latency, which heavily depends on the location of clients and servers. Network latency, in our protocol, incurs the same overhead as in the non-threshold setting. It can be incorporated by adding the round-trip time of messages up to 2KB over the client's (slowest) server connection to the total runtime. As the online phase of our protocol is non-interactive, we benchmark servers and clients individually. We execute all benchmarks on a single machine with a 14-core Intel Xeon Gold 5120 CPU @ 2.20GHz processor and 64GB of RAM. We repeat each benchmark 100 times to account for statistical deviations and report the average. For comparability, we report the runtime of basic arithmetic operations in Table 1 in Appendix O.

Experimental Results. We report the results of our benchmarks in Figure 2. These results reflect our expectations as outlined in the following. The *Adapt* phase transforming PCF/PCG outputs to signing request-dependent presignatures involves only field operations and is much faster than the other phases for small t . The runtime increases for larger t due to the computation of LaGrange coefficient which scales quadratically with the number of signers. The *Sign* phase requires the servers to compute $k + 1$ scalar multiplications in \mathbb{G}_1 , each taking 100 times more time than the slowest field operation (cf. Appendix O). The *Reconstruct* phase involves a single \mathbb{G}_1 scalar multiplication and less expensive hashing and field operations. The *Verify* phase requires the client to compute two pairing operations, a single scalar multiplication in \mathbb{G}_2 , k scalar multiplications \mathbb{G}_1 , and multiple additions in \mathbb{G}_1 and \mathbb{G}_2 . The pairing operations and the scalar multiplication in \mathbb{G}_2 are responsible for the constant costs visible in the graph. The scalar multiplications in \mathbb{G}_1 cause the linear increase. The influence of \mathbb{G}_1 and \mathbb{G}_2 additions is insignificant because they take at most 1.4% of scalar multiplication in \mathbb{G}_1 . The *Total* runtime mainly depends on the size of the signed message array, k , due to the scalar multiplications in the signing and verification step. The number of signers, t , has only a minor influence on the online runtime; increasing the number of signers from 2 to 30 increases the runtime by 1.13% – 6.02%. Consequently, the online protocol can essentially support any number of servers as long as the preprocessing, which is expected to scale worse, can be instantiated efficiently for the number of servers and the storage complexity of the generated preprocessing material does not exceed the servers’ capacities (cf. Section 6.2).

To measure the *overhead of thresholdization*, we compare the runtime of our online protocol to the runtime of signature creation in the non-threshold setting in Figure 3. The overhead of our online protocol mainly consists of a single scalar multiplication in \mathbb{G}_1 , assuming that clients also verify received signatures in the non-threshold setting. This observation reflects our protocol pushing all the overhead of the distributed signing to the offline phase.




Fig. 3: The total runtime of our online protocol compared to plain, non-threshold signing with and without signature verification in dependence of k . The number of signers t is insignificant (cf. Figure 2e).

Communication Complexity. The client has to send one signing request of size $(k \cdot \lceil \log p \rceil) + (t \cdot \lceil \log n \rceil)$ bits to each of the t selected servers. By deriving the signer set via a random oracle, we can reduce the size of the request to $(k \cdot \lceil \log p \rceil)$. Each selected server has to send a partial signature of size $(\lceil \log p \rceil + |\mathbb{G}_1|)$. In case of the BLS12_381 curve, $\lceil \log p \rceil$ equals 381 bits whereas $|\mathbb{G}_1|$ equals 762 bits. Parties can also encode \mathbb{G}_1 elements with 381 bits by only sending the x -coordinate of the curve point and requiring the sender to compute the y -coordinate itself.

Note that our UC functionality models a scenario where every signer obtains the final signature. Therefore, the partial signatures are sent to all other signers. However, by incorporating a dedicated client into the model, the signers can send the partial signatures only to the client. While we expect this to be sufficient for real-life settings, it makes the model messier. We emphasize that this request-response behavior is the minimum interaction for distributed protocols. As there is no interaction between the servers, this setting is referred to as non-interactive in the literature [29, 3].

6.2 Offline, Signing Request-Independent Phase

In this section, we provide the evaluation of the offline, signing request-independent phase. We focus on the PCG-based precomputation as PCFs lack efficient instantiations. We implement the seed expansion of the PCGs (Step 3 of protocol $\pi_{\text{Prep}}^{\text{PCG}}$), run benchmarks and report the runtime. We further compute the storage complexity of the computed correlations.

Remark. While we focus the following evaluation on the protocol presented in this paper, we have also conducted evaluations of the original version of the BBS+ protocol, i.e., one without the optimization introduced by [64] that includes an additional s value into the signature. To this end, we instantiate the PCG in the offline phase based on the Ring-LPN assumption, with which we cannot only generate VOLE correlations but also oblivious linear evaluation (OLE) correlations required for the adapted version of our protocol.

Parameters. In the following, we denote the security parameter by λ , the number of servers by n , the security threshold by t , the size of the signed message arrays by k , the number of generated precomputation tuples by N , the order of the elliptic curve's groups \mathbb{G}_1 and \mathbb{G}_2 by p and assume PCGs based on the Ring LPN problem with static leakage and security parameters c and τ , i.e., the $R^c\text{-LPN}_{p,\tau}$ assumption.³ This assumption is common to state-of-the-art PCG instantiations for OLE correlations [19].

Experimental setup. Our implementation⁴ and benchmarks are written in Go. Our code, including the benchmarks and rudimentary tests, comprises 5 450 lines. We compiled our code using go 1.21.3. Again, we execute all benchmarks

³ For 128-bit security and $N = 2^{20}$, [19] reports $(c, \tau) \in \{(2, 76), (4, 16), (8, 5)\}$.

⁴ <https://github.com/leandro-ro/Threshold-BBS-Plus-PCG>

on machines with a 14-core Intel Xeon Gold 5120 CPU @ 2.20GHz processor and 64GB of RAM. Due to the complexity of the benchmarks and the high amount of repetition within a single protocol run, we execute the benchmarks for each choice of parameters just once.

The runtime of the seed expansion is influenced by the number of parties n , the number of generated precomputation tuples N and the Module Ring LPN security parameters (c, τ) . For security parameters we fix $c = 4$ and $\tau = 16$ which corresponds to 128-bit security [19]. We compute over a cyclotomic ring as proposed by [19] and fix the prime p to be the order of the BLS12_381 curve. Our tests have shown that this choice of parameters yields the best performance of the possible choices for the same security level. For the number of parties, we consider $2 \leq n \leq 10$.⁵ Further, we consider both the t -out-of- n setting and the n -out-of- n setting as the latter has tremendous potential for optimization as discussed below. For the number of generated triples, we consider $N \in [2^{11}, \dots, 2^{17}]$.

Our benchmarks cover the seed expansion phase of the required PCGs (Step 3 of $\pi_{\text{Prep}}^{\text{PCG}}$). As our PCG instantiations compute over a ring, they also return ring elements each representing an array of N field elements. For example, for a batch of N VOLE correlations $\mathbf{a} \cdot b = \mathbf{c} + \mathbf{d}$ ($\mathbf{a}, \mathbf{c}, \mathbf{d} \in \mathbb{Z}_p^N, b \in \mathbb{Z}_p$), the PCG returns one field element b and three degree- N polynomials A, C, D such that $A \cdot b = C + D$. By choosing the ring appropriately (cf. [19]), each polynomial can be split into N independent VOLE correlations over \mathbb{Z}_p . This step does not need to happen in a batch but can be done individually. We report the computation time of the PCG seed expansion, yielding the ring elements, and the time to split a single VOLE correlation over \mathbb{Z}_p from a ring element, separately.

n -out-of- n vs. t -out-of- n . The runtime of the seed expansion strongly depends on whether we consider a t -out-of- n or a n -out-of- n setting. To understand this dependency, recall the basic concepts of PCGs (cf. PCG constructions in Appendix E). Parties first compute the desired correlation with sparse polynomials as input values. Then, they expand these preliminary sparse correlations to dense randomized correlations by applying an LPN-based randomization. In our protocol specification, parties do this for each individual VOLE-correlation and then combine the randomized correlations to get the final precomputation tuples. However, in a real implementation parties can first combine the sparse correlations and then apply the LPN-based randomization, effectively reducing the amount of randomization operations. In the t -out-of- n setting, the signer set is only known during the online phase, i.e., after the randomization step. As the combination itself largely depends on the signer set, parties can only perform the combination steps that are independent of the set. In the n -out-of- n setting, the signer set is already known during the offline phase, i.e., every party has to sign. Parties can therefore perform combinations before randomization. More precisely, in the n -out-of- n setting, each party has to perform two randomizations and split two polynomials, while in the t -out-of- n setting each party performs $1 + 2 \cdot (n - 1)$ randomizations and splits just as many polynomials.

⁵ The only prior work implementing the seed expansion [3] is restricted to $n \in \{2, 3\}$.

Experimental results. In Figure 4 and Figure 5, we display the computation time per signature of the PCG expansion in the n -out-of- n setting and the t -out-of- n setting. The computation time per signature increases superlinear with the number of signatures (note that the x-axis has a logarithmic scale) and linear with the number of parties n . The superlinear increase is due to the fact that the seed expansion requires multiplication of degree- N polynomials. We perform the multiplication via the Fast Fourier Transformation which scales superlinear with the degree of the polynomial. Both graphs show that the runtime increases with the number N . Nevertheless, as the correlations are expanded from small keys, a large batch size N still benefits from the sublinear communication complexity in N of a distributed seed generation.

We further benchmarked the computation time to extract one of N field elements from a degree- N polynomial. The results range from 0.1ms for $N = 2^{11}$ to 1.6ms for $N = 2^{17}$. This step essentially represents a polynomial evaluation executed via the Horner’s method which explains the linear increase in the computation time.




Fig. 4: Computation time in ms per signature of the seed expansion of all required PCGs in the n -out-of- n setting for different committee sizes ($n \in \{2, \dots, 10\}$) dependent on the number of generated precomputation tuples N .

Storage complexity. When instantiating the precomputation with PCGs, servers must evaluate all of the PCGs’ outputs at once. The resulting precomputation material occupies

$$\log p \cdot N \cdot (1 + 2 \cdot (n - 1))$$

bits of storage. In [3], the authors report $N = 94\,019$ as a reasonable parameter for a PCG-based setup protocol. In [19], the authors base their analysis on $N = 2^{20} = 1\,048\,576$. To efficiently apply Fast Fourier Transformation algorithms during the seed expansion, it is necessary to choose N such that it divides $p - 1$.

Figure 6 reports the storage complexity depending on the number of servers n for different N . Note that the dependency on the number of servers n stems from the fact that we support any threshold $t \leq n$. In a n -out-of- n settings, servers execute Step 6 of $\pi_{\text{Prep}}^{\text{PCG}}$ during the preprocessing, and hence, only store $\log p \cdot 2N$ bits of preprocessing material.




Fig. 5: Computation time in ms per signature of the seed expansion of all required PCGs in the t -out-of- n setting for different committee sizes ($n \in \{2, \dots, 10\}$) dependent on the number of generated precomputation tuples N .




Fig. 6: Storage complexity of the precomputation material required for $N \in \{94\,019, 1\,048\,576\}$ signatures depending on the number of servers n .

6.3 Comparison to [40]

Independently of our work, [40] presented the first t -out-of- n threshold protocol for the legacy BBS+ scheme. This protocol incorporates a lightweight setup independent of the number of generated signatures but requires an interactive sign-




Fig. 7: Runtime of the signing protocol of [40] compared to the network adjusted runtime of our signing protocol in the LAN and WAN setting.

ing protocol. In contrast, our scheme offers a trade-off that provides an efficient, non-interactive online phase at the cost of a more complex offline phase. This trade-off aims to minimize the time it takes to answer a signing request. To show that our online phase’s efficiency indeed benefits from the costly preprocessing, we compare our online signing phase to their interactive signing protocol.⁶ We stress that the advantage of our online phase comes at the cost of a significantly more complex offline phase. However, our online phase is independent of the concrete instantiation of the offline phase. In particular, less memory-consuming but more communication-intensive instantiations, e.g., based on oblivious transfer or somewhat homomorphic encryption, are also possible.

As our implementation, their implementation is in Rust and based on the BLS12_381 curve. When comparing the benchmarking machines, \mathbb{G}_1 and \mathbb{G}_2 scalar multiplications are 15 – 25% faster on our machine, while signature verifications are 20% faster on their machine. Although not explicitly stated, the numbers strongly indicate the choice $k = 1$ in [40]; the reported runtime of non-threshold BBS+ signing is slightly larger than three \mathbb{G}_1 scalar multiplications. Due to the interactivity of their protocol, their benchmarks incorporate network delays for different settings (LAN, WAN). We add network delays to our results to compare our benchmarks to theirs. All machines used in their evaluation are *Google Cloud c2d-standard-4* instances. In the LAN setting, all instances are located at the us-east1-c zone. [38] reports a LAN latency of 0.146 ms for this zone. We add a delay of 0.3 ms to our results. In the WAN setting, the first 12 instances in their benchmarks are located in the US, while other machines are in Europe or the US. According to [49], we add 100 ms to our results for $t < 13$ and 150 ms for $t \geq 13$.

In Figure 7, we compare the runtime, including latency, of our online signing protocol to the runtimes reported in [40] for the LAN and the WAN setting. The graphs show that our protocol outperforms the one of [40] in both settings for every number of servers. The only exception is the runtime for $t = 2$ in

⁶ We thank the authors of [40] for sharing concrete numbers of their evaluation.

the WAN setting. This exception seems caused by an unusually low connection latency between the first two servers and the client in [40]. The overhead of [40] is mainly caused by the two additional rounds of cross-server interaction. This overhead rises with the number of servers as each server has to communicate with each other servers and is especially severe in the WAN setting.

Due to the high efficiency and non-interactivity of our online phase, our protocol is more suited for settings where servers have a sufficiently long setup interval and storage capacities to deal with the complexity of the preprocessing phase. On the other hand, the protocol of [40] is more suited for use cases with more lightweight servers, especially in a LAN environment where the network delay of the additional communication is less significant.

Acknowledgments. The first, third, and fifth authors were supported by the German Research Foundation (DFG) CRC 1119 CROSSING (project S7), by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE, by the European Research Council (ERC) under the European Union’s Horizon 2020 and Horizon Europe research and innovation programs (grant CRYPTOLAYER-101044770), and by the Hessen Agentur, LOEWE-Förderrichtlinie 3, 1376/22-81. The second author was partially supported by the Algorand Centres of Excellence programme managed by Algorand Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Algorand Foundation and the United States-Israel Binational Science Foundation (BSF) through Grant No. 2020277.

References

1. Github - trinsic-id/bbs. <https://github.com/trinsic-id/bbs>. (Accessed on 10/14/2024).
2. Masayuki Abe and Tatsuaki Okamoto. Provably secure partially blind signatures. In *CRYPTO*, 2000.
3. Damiano Abram, Ariel Nof, Claudio Orlandi, Peter Scholl, and Omer Shlomovits. Low-bandwidth threshold ECDSA via pseudorandom correlation generators. In *IEEE SP*, 2022.
4. Algorand. BLS12-381 Rust crate. <https://github.com/algorand/pairing-plus>. (Accessed on 10/14/2024).
5. Man Ho Au, Willy Susilo, Yi Mu, and Sherman S. M. Chow. Constant-size dynamic k -times anonymous authentication. *IEEE Syst. J.*, 2013.
6. Jean-Philippe Aumasson, Adrian Hamelink, and Omer Shlomovits. A survey of ECDSA threshold signing. *IACR Cryptol. ePrint Arch.*, 2020.
7. Judit Bar-Ilan and Donald Beaver. Non-cryptographic fault-tolerant computing in constant number of rounds of interaction. In *PODC*, 1989.
8. Donald Beaver. Efficient multiparty protocols using circuit randomization. In *CRYPTO*, 1991.

9. Greg Bernstein and Manu Sporny. Data integrity bbs cryptosuites v1.0. <https://w3c.github.io/vc-di-bbs/>, April 2024. (Accessed on 10/14/2024).
10. Johannes Blömer, Jan Bobolz, Denis Diemert, and Fabian Eidens. Updatable anonymous credentials and applications to incentive systems. In *CCS*, 2019.
11. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *CRYPTO*, 2004.
12. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the weil pairing. In *CRYPTO*, 2001.
13. Alexandre Bouez and Kalpana Singh. One round threshold ECDSA without roll call. In *CT-RSA*, 2023.
14. Elette Boyle, Geoffroy Couteau, Niv Gilboa, and Yuval Ishai. Compressing vector OLE. In *CCS*, 2018.
15. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Nicolas Resch, and Peter Scholl. Correlated pseudorandomness from expand-accumulate codes. In *CRYPTO*, 2022.
16. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, Peter Rindal, and Peter Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In *CCS*, 2019.
17. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In *CRYPTO*, 2019.
18. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Correlated pseudorandom functions from variable-density LPN. In *FOCS*, 2020.
19. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators from ring-lpn. In *CRYPTO*, 2020.
20. Ernie Brickell and Jiangtao Li. A pairing-based DAA scheme further reducing TPM resources. In *TRUST*, 2010.
21. Ernie Brickell and Jiangtao Li. Enhanced privacy ID from bilinear pairing for hardware authentication and attestation. *Int. J. Inf. Priv. Secur. Integr.*, 2011.
22. Jan Camenisch. Anonymous credentials: Opportunities and challenges. In *SEC*, 2006.
23. Jan Camenisch, Manu Drijvers, and Anja Lehmann. Anonymous attestation using the strong diffie hellman assumption revisited. In *TRUST*, 2016.
24. Jan Camenisch, Maria Dubovitskaya, Kristiyan Haralambiev, and Markulf Kohlweiss. Composable and modular anonymous credentials: Definitions and practical constructions. In *ASIACRYPT*, 2015.
25. Jan Camenisch, Stephan Krenn, Anja Lehmann, Gert Læssøe Mikkelsen, Gregory Neven, and Michael Østergaard Pedersen. Formal treatment of privacy-enhancing credential systems. In *SAC*, 2015.
26. Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT*, 2001.
27. Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In *CRYPTO*, 2004.
28. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
29. Ran Canetti, Rosario Gennaro, Steven Goldfeder, Nikolaos Makriyannis, and Udi Peled. UC non-interactive, proactive, threshold ECDSA with identifiable aborts. In *CCS*, 2020.

30. Guilhem Castagnos, Dario Catalano, Fabien Laguillaumie, Federico Savasta, and Ida Tucker. Bandwidth-efficient threshold EC-DNA. In *PKC*, 2020.
31. Guilhem Castagnos, Fabien Laguillaumie, and Ida Tucker. Threshold linearly homomorphic encryption on $\mathbf{Z}/2^k\mathbf{Z}$. In *ASIACRYPT*, 2022.
32. David Chaum. Security without identification: Transaction systems to make big brother obsolete. *Commun. ACM*, 1985.
33. Lidong Chen. Access with pseudonyms. In *Cryptography: Policy and Algorithms*, 1995.
34. Liqun Chen. A DAA scheme requiring less TPM resources. In *Information Security and Cryptology*, 2009.
35. Hien Chu, Paul Gerhart, Tim Ruffing, and Dominique Schröder. Practical schnorr threshold signatures without the algebraic group model. In *CRYPTO*, 2023.
36. Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In *CRYPTO*, 2021.
37. Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, 2012.
38. Rick Jones Derek Phanekham. How much is google cloud latency (gcp) between regions? <https://cloud.google.com/blog/products/networking/using-netperf-and-ping-to-measure-network-latency>, June 2020. (Accessed on 10/14/2024).
39. Samuel Dittmer, Yuval Ishai, Steve Lu, and Rafail Ostrovsky. Authenticated garbling from simple correlations. In *CRYPTO*, 2022.
40. Jack Doerner, Yash Kondi, Eysa Lee, abhi shelat, and LakYah Tyner. Threshold bbs+ signatures for distributed anonymous credential issuance. In *IEEE SP*, 2023.
41. Jack Doerner, Yashvanth Kondi, Eysa Lee, and Abhi Shelat. Threshold ECDSA from ECDSA assumptions: The multiparty case. In *SP*, 2019.
42. Shimon Even, Oded Goldreich, and Silvio Micali. On-line/off-line digital signatures. *J. Cryptol.*, 1996.
43. Rosario Gennaro and Steven Goldfeder. Fast multiparty threshold ECDSA with fast trustless setup. In *CCS*, 2018.
44. Rosario Gennaro, Steven Goldfeder, and Bertrand Ithurburn. Fully distributed group signatures, 2019.
45. Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, and Tal Rabin. Secure distributed key generation for discrete-log based cryptosystems. In *EUROCRYPT*, 1999.
46. Chelsea Komlo and Ian Goldberg. FROST: flexible round-optimized schnorr threshold signatures. In *SAC*, 2020.
47. Yashvanth Kondi, Bernardo Magri, Claudio Orlandi, and Omer Shlomovits. Refresh when you wake up: Proactive threshold wallets with offline devices. In *SP*, 2021.
48. Yashvanth Kondi, Claudio Orlandi, and Lawrence Roy. Two-round stateless deterministic two-party schnorr signatures from pseudorandom correlation functions. In *CRYPTO*, 2023.
49. Chandan Kumar. How much is google cloud latency (gcp) between regions? <https://geekflare.com/google-cloud-latency/>, March 2022. (Accessed on 10/14/2024).
50. Yehuda Lindell. Fast secure two-party ECDSA signing. In *CRYPTO*, 2017.
51. Yehuda Lindell and Ariel Nof. Fast secure multiparty ECDSA with practical distributed key generation and applications to cryptocurrency custody. In *CCS*, 2018.

52. Tobias Looker, Vasilis Kalos, Andrew Whitehead, and Mike Lodder. The BBS Signature Scheme. Internet-Draft draft-irtf-cfrg-bbs-signatures-02, Internet Engineering Task Force, March 2023. (Work in Progress).
53. Anna Lysyanskaya, Ronald L. Rivest, Amit Sahai, and Stefan Wolf. Pseudonym systems. In *SAC*, 1999.
54. MATTR. mattrglobal/bbs-signatures: An implementation of bbs+ signatures for node and browser environments. <https://github.com/mattrglobal/bbs-signatures>. (Accessed on 10/14/2024).
55. Microsoft. microsoft/bbs-node-reference: Typescript/node reference implementation of bbs signature. <https://github.com/microsoft/bbs-node-reference>. (Accessed on 10/14/2024).
56. Jesper Buus Nielsen, Peter Sebastian Nordholt, Claudio Orlandi, and Sai Sheshank Burra. A new approach to practical active-secure two-party computation. In *CRYPTO*, 2012.
57. Claudio Orlandi, Peter Scholl, and Sophia Yakubov. The rise of paillier: Homomorphic secret sharing and public-key silent OT. In *EUROCRYPT*, 2021.
58. Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *CRYPTO*, 1991.
59. David Pointcheval and Olivier Sanders. Short randomizable signatures. In *CT-RSA*, 2016.
60. Alfredo Rial and Ania M. Piotrowska. Security analysis of coconut, an attribute-based credential scheme with threshold issuance. *IACR Cryptol. ePrint Arch.*, 2022.
61. Adi Shamir. How to share a secret. *Commun. ACM*, 1979.
62. Nigel P. Smart and Younes Talibi Alaoui. Distributing any elliptic curve based protocol. In *IMA*, 2019.
63. Alberto Sonnino, Mustafa Al-Bassam, Shehar Bano, Sarah Meiklejohn, and George Danezis. Coconut: Threshold issuance selective disclosure credentials with applications to distributed ledgers. In *NDSS*, 2019.
64. Stefano Tessaro and Chenzhi Zhu. Revisiting BBS signatures. In *EUROCRYPT*, 2023.
65. Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and efficient maliciously secure two-party computation. In *CCS*, 2017.
66. Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-scale secure multiparty computation. In *CCS*, 2017.
67. Douglas Wikström. Universally composable DKG with linear number of exponentiations. In *SCN*, 2004.
68. Harry W. H. Wong, Jack P. K. Ma, and Sherman S. M. Chow. Secure multiparty computation of threshold signatures made more efficient. In *NDSS*, 2024.
69. Harry W. H. Wong, Jack P. K. Ma, Hoover H. F. Yin, and Sherman S. M. Chow. Real threshold ECDSA. In *NDSS*, 2023.
70. Zuoxia Yu, Man Ho Au, and Rupeng Yang. Accountable anonymous credentials. In *Advances in Cyber Security: Principles, Techniques, and Applications*. 2019.

Appendix:

Non-Interactive Threshold BBS+ From Pseudorandom Correlations

A The BBS+ Signature Scheme

We present the BBS+ signature scheme as proposed by Tessaro and Zhu [64]. Their version is an improved variant of the original BBS+ signature scheme as described by Au et al. [5]. For the sake of completeness, we present the version of Au et al. in Appendix B

Let k be the size of the message arrays, $\mathcal{G} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, p, g_1, g_2, \mathbf{e})$ be a bilinear mapping tuple and $\{h_\ell\}_{\ell \in [0..k]}$ be random elements of \mathbb{G}_1 . The BBS+ signature scheme is defined as follows:

- **KeyGen**(λ): Sample $x \xleftarrow{\$} \mathbb{Z}_p^*$, compute $y = g_2^x$, and output $(\mathbf{pk}, \mathbf{sk}) = (y, x)$.
- **Sign**_{sk}($\{m_\ell\}_{\ell \in [k]} \in \mathbb{Z}_p^k$): Sample $e \xleftarrow{\$} \mathbb{Z}_p$, compute $A := (g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{\frac{1}{x+e}}$ and output $\sigma = (A, e)$.
- **Verify**_{pk}($\{m_\ell\}_{\ell \in [k]} \in \mathbb{Z}_p^k, \sigma$): Output 1 iff $\mathbf{e}(A, y \cdot g_2^e) = \mathbf{e}(g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell}, g_2)$

The BBS+ signature scheme is proven strong unforgeable under the q -strong Diffie Hellman (SDH) assumption for pairings of type 1, 2, and 3 [64]. Intuitively, strong unforgeability states that the attacker is not possible to come up with a forgery even for messages that have been signed before. We refer to [64] for further details.

B The BBS+ Signature Scheme by Au et al. [5]

In this section, we describe the BBS+ signature scheme as originally defined by Au et al. [5]. Moreover, we describe the implications on our online and offline protocol.

Let $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T be groups of prime order p with generators $g_1 \in \mathbb{G}_1$ and $g_2 \in \mathbb{G}_2$ and let map $\mathbf{e} : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ be a bilinear paring. A BBS+ signature according to the definition of [5], which we call a legacy BBS+ signature in the following, on a message array $\{m_\ell\}_{\ell \in [k]}$ is a tuple (A, e, s) with $A = (g_1 \cdot h_0^s \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{\frac{1}{x+e}}$ for random nonces $e, s \in_R \mathbb{Z}_p$, secret key $x \in \mathbb{Z}_p$ and a set of random elements $\{h_\ell\}_{\ell \in [0..k]}$ in \mathbb{G}_1 . To verify under public key g_2^x , check if $\mathbf{e}(A, g_2^x \cdot g_2^e) = \mathbf{e}(g_1 \cdot h_0^s \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell}, g_2)$. In contrast to the optimized BBS+ scheme of [64], the legacy BBS+ signature has the additional s value. While originally used to prove the security of the scheme, [64] shows how to prove security without the additional nonce s , resulting in an optimized version.

Based on the legacy version of the BBS+ signature scheme, our online and offline protocols must be adapted. In particular, we adapt the definitions of presignatures, which are generated by the offline phase and consumed by the

online phase. We describe the idea on the example of an n -out-of- n threshold protocol.

Assume a BBS+ secret key x , elements $\{h_\ell\}_{\ell \in [0..k]}$ in \mathbb{G}_1 , a random blinding factor $a \in \mathbb{Z}_p$ and n servers, each having access to a preprocessed tuple $(a_i, s_i, \beta_i, \alpha_i) \in \mathbb{Z}_p^4$, in the following called presignatures, such that

$$\begin{aligned} \sum_{i \in [n]} \beta_i &= a \cdot x, & \sum_{i \in [n]} \alpha_i &= a \cdot s \\ \text{for } a &= \sum_{i \in [n]} a_i, & s &= \sum_{i \in [n]} s_i. \end{aligned} \tag{2}$$

To sign a message array $\{m_\ell\}_{\ell \in [k]}$, each server first gets e from the random oracle, then it computes $A_i := (g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{a_i} \cdot h_0^{\alpha_i}$ and outputs a partial signature $\sigma_i := (A_i, \delta_i := \beta_i + a_i \cdot e, s_i)$. This allows the receiver of the partial signatures to reconstruct δ and s and compute

$$A = \left(\prod_{i \in [n]} A_i \right)^{\frac{1}{\delta}} = \left((g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^a \cdot h_0^{as} \right)^{\frac{1}{a(x+e)}}$$

so that the tuple (A, e, s) constitutes a valid BBS+ signature.

During the offline phase for the legacy BBS+ scheme, the parties must additionally compute shares of α , which is the product of two secret shared values a and s . Recall that we use VOLE correlations to compute shares of the product of $a \cdot x$. We use VOLE correlations, since the secret key x is constant over all presignatures. In contrast, the nonce s must be random for each presignature. Therefore, we use OLE correlations instead VOLE correlations. Using OLE correlations follows the same approach as using VOLE correlations.

We describe reusable PCG and PCF constructions for OLE correlations in Appendix E and G and evaluate our threshold protocol for the legacy BBS+ scheme in Appendix N.

C Universal Composability Framework ([28])

We formally model and prove the security of our protocols in the Universal Composability framework (UC). The framework was introduced by Canetti in 2001 [28] to analyze the security of protocols formally. The universal composability property guarantees the security of a protocol holds even under concurrent composition. We give a brief intuition and defer the reader to [28] for all details.

Like simulation-based proofs, the framework differentiates between real-world and ideal-world execution. The real-world execution consists of n parties P_1, \dots, P_n executing protocol π , an adversary \mathcal{A} , and an environment \mathcal{Z} . All parties are initialized with security parameter λ and a random tape, and \mathcal{Z} runs on some advice string z . In this work, we consider only static corruption, where the adversary corrupts parties at the onset of the execution. After corruption, the adversary

may instruct the corrupted parties to deviate arbitrarily from the protocol specification. The environment provides inputs to the parties, instructs them to continue the execution of π , and receives outputs from the parties. Additionally, \mathcal{Z} can interact with the adversary.

The real-world execution finishes when \mathcal{Z} stops activating parties and outputs a decision bit. We denote the output of the real-world execution by $\text{REAL}_{\pi, \mathcal{A}, \mathcal{Z}}(\lambda, z)$.

The ideal-world execution consists of n dummy parties, an ideal functionality \mathcal{F} , an ideal adversary \mathcal{S} , and an environment \mathcal{Z} . The dummy parties forward messages between \mathcal{Z} and \mathcal{F} , and \mathcal{S} may corrupt dummy parties and act on their behalf in the following execution. \mathcal{S} can also interact with \mathcal{F} directly according to the specification of \mathcal{F} . Additionally, \mathcal{Z} and \mathcal{S} may interact. The goal of \mathcal{S} is to simulate a real-world execution such that the environment cannot tell apart if it is running in the real or ideal world. Therefore, \mathcal{S} is also called the simulator.

Again, the ideal-world execution ends when \mathcal{Z} outputs a decision bit. We denote the output of the ideal-world execution by $\text{IDEAL}_{\mathcal{F}, \mathcal{S}, \mathcal{Z}}(\lambda, z)$.

Intuitively, a protocol is secure in the UC framework if the environment cannot distinguish between real-world and ideal-world execution. Formally, protocol π UC-realizes \mathcal{F} if for every probabilistic polynomial-time (PPT) adversary \mathcal{A} there exists a PPT simulator \mathcal{S} such that for every PPT environment \mathcal{Z}

$$\{\text{REAL}_{\pi, \mathcal{A}, \mathcal{Z}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} = \{\text{IDEAL}_{\mathcal{F}, \mathcal{S}, \mathcal{Z}}(1^\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}.$$

D Anonymous Credentials and Blind Signing

Our online protocol defined in Section 4.3 describes a threshold variant of the BBS+ signature scheme. Since anonymous credentials are one prominent application of BBS+ signatures, we elaborate on this application in the following.

BBS+ signatures can be used to design anonymous credential schemes as follows. To receive a credential, a client sends a signing request to the servers in the form of a message array, which contains its public and private credential information. Public parts of the credentials are sent in clear, while private information is blinded. The client can add zero-knowledge proofs that blinded messages satisfy some predicate. These proofs enable the issuing servers to enforce a signing policy even though they blindly sign parts of the messages. Given a credential, clients can prove in zero-knowledge that their credential fulfills certain predicates without leaking their signature.

Our scheme must be extended by a blind-signing property to realize the described blueprint. Precisely, we require a property called *partially blind* signatures [2]. This property prevents the issuer from learning private information about the message to be signed.

To transform our scheme into a partially blind signature scheme, we follow the approach of [5]. Let $\{m_\ell\}_{\ell \in [k]}$ be the set of messages representing the client's credential information. Without loss of generality, we assume that m_k is the public part. In order to blind its messages, the client computes a Pedersen

commitment [58] on the private messages: $C = h_0^{s'} \cdot \prod_{\ell \in [k-1]} h_\ell^{m_\ell}$ for a random s' and a zero-knowledge proof π that C is well-formed, i.e., that the client knows $(s', \{m_\ell\}_{\ell \in [k-1]})$. The client sends $(\mathcal{T}, C, \pi, m_k)$ and potential zero-knowledge proofs for signing policy enforcement to the servers. Each server P_i for $i \in \mathcal{T}$ replies with $(A_i = (g_1 \cdot C \cdot h_k^{m_k})^{a_i}, \delta_i)$. While creating the responses, the servers compute the nonce $e = \text{H}(\text{ssid}||C||\mathcal{T})$, i.e., in contrast to our online protocol presented in Section 4.3, the servers take the commitment instead of the set of messages as input to the random oracle. This approach in combination with a random value s' in the commitment ensures that the server does not learn any information about the messages. Note that the client can perform the same random oracle evaluation and upon receiving the server responses, the client computes A as before and outputs a valid signature (A, e) . Due to the random nonce s' of the Pedersen commitment, the final signature represents a signature on the message array (s', m_1, \dots, m_k) .

As the blinding mechanism and the resulting signatures are equivalent in the non-threshold BBS+ setting, we can use existing zero-knowledge proofs for policy enforcement and credential usage from the non-threshold setting [5, 23, 64].

E Reusable PCG Constructions

In this section, we present a construction of reusable PCGs for VOLE correlations according to the definitions provided in Section 3 together with the required building blocks and security assumptions. Since our threshold protocol for the legacy BBS+ signature scheme additionally requires OLE correlations (cf. Appendix B), we also present a construction for this type of correlations. The constructions are derived from the one of [19].

Notation Let R be a ring. For two column vectors $\mathbf{u} = (u_1, \dots, u_t) \in R^t$ and $\mathbf{v} = (v_1, \dots, v_t) \in R^t$, we define the *outer sum* $\mathbf{u} \boxplus \mathbf{v}$ be the vector $(u_i + v_j)_{i,j \in [t]} \in R^{t^2}$. Similar, we define the *outer product* (or tensor product) $\mathbf{u} \otimes \mathbf{v}$ to be $(u_i \cdot v_j)_{i,j \in [t]} \in R^{t^2}$. The *inner product* of two t -size vectors $\langle \mathbf{u}, \mathbf{v} \rangle$ is defined as $(\sum_{i \in [t]} u_i \cdot v_i) \in R$.

Ring Module LPN Assumption The following definition of the Module Ring LPN assumption introduced by [19] is taken almost verbatim from the original [19, Definition 3.2] but adapted to our notation.

Definition 3. Module-LPN

Let $c \geq 2$ be an integer, let $R = \mathbb{Z}_p/F(X)$ for a prime p and degree- N polynomial $F(X) \in \mathbb{Z}_p[X]$ and let $\tau \in \mathbb{N}$ be an integer. Further, let $\mathcal{HW}_{R,\tau}$ denote the distribution of “sparse polynomials” over R obtained by sampling τ noise positions $\alpha \leftarrow [N]^\tau$ and τ payloads $\beta \leftarrow (\mathbb{Z}_p^*)^\tau$ uniformly at random and outputting $e(X) := \sum_{i \in [\tau]} \beta[i] \cdot X^{\alpha[i]-1}$. Then, for $R = R(\lambda), m = m(\lambda), \tau = \tau(\lambda)$, we

say the $R^c\text{-LPN}_{R,m,\tau}$ problem is hard if for every nonuniform polynomial-time distinguisher \mathcal{A} , it holds that

$$\begin{aligned} & |\Pr[\mathcal{A}(\{\langle \mathbf{a}^{(i)}, \langle \mathbf{a}^{(i)}, \mathbf{e} \rangle + f^{(i)} \rangle\}_{i \in [m]} = 1] \\ & - \Pr[\mathcal{A}(\{\langle \mathbf{a}^{(i)}, u^{(i)} \rangle\}_{i \in [m]} = 1)]| \leq \text{negl}(\lambda) \end{aligned}$$

where the probabilities are taken over $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(m)} \leftarrow R^{c-1}$, $u^{(1)}, \dots, u^{(m)} \leftarrow R$, $\mathbf{e} \leftarrow \mathcal{HW}_{R,\tau}^{c-1}$, $f^{(1)}, \dots, f^{(m)} \leftarrow \mathcal{HW}_{R,\tau}$.

Distributed Sum of Point Functions (DSPF) We use distributed sum of functions. The definition is taken partially verbatim from [19, 3] but adapted to our notation.

Definition 4 (Distributed Sum of Point Functions). Let \mathbb{G} be an Abelian group, N, τ be positive integers, $f_{\alpha,\beta} : [N] \rightarrow \mathbb{G}$ be a sum of τ point functions, parametrized for $\alpha \in [N]^\tau$ and $\beta \in \mathbb{G}^\tau$, such that $f_{\alpha,\beta}(x) = 0 + \sum_{(i \in [\tau] \text{ s.t. } \alpha[i]=x)} \beta[i]$. A 2-party distributed sum of point functions (DSPF) with domain $[N]$, codomain \mathbb{G} , and weight τ is a pair of PPT algorithms (DSPF.Gen, DSPF.Eval) with the following syntax.

- DSPF.Gen takes as input the security parameter 1^λ and a description of the sum of point functions $f_{\alpha,\beta}$, specifically, the special positions $\alpha \in [N]^\tau$ and the non-zero elements $\beta \in \mathbb{G}^\tau$. The output is two keys (K_0, K_1) .
- DSPF.Eval takes as input a DPF key K_σ , index $\sigma \in \{0, 1\}$ and a value $x \in [N]$, outputting an additive share v_σ of $f_{\alpha,\beta}(x)$.

A DSPF should satisfy the following properties:

- **Correctness.** For every set of special positions $\alpha \in [N]^\tau$, set of non-zero elements $\beta \in \mathbb{G}^\tau$ and element $x \in [N]$, we have that

$$\begin{aligned} & \Pr[v_0 + v_1 = f_{\alpha,\beta}(x) | (K_0, K_1) \leftarrow \text{DSPF.Gen}(1^\lambda, \alpha, \beta), \\ & v_\sigma \leftarrow \text{DSPF.Eval}(K_\sigma, \sigma, x) \text{ for } \sigma \in \{0, 1\}] = 1 \end{aligned}$$

- **Security.** There exists a PPT simulator \mathcal{S} such that, for every corrupted party $\sigma \in \{0, 1\}$, set of special positions $\alpha \in [N]^\tau$ and set of non-zero elements $\beta \in \mathbb{G}^\tau$, the output of $\mathcal{S}(1^\lambda, \sigma)$ is computationally indistinguishable from

$$\{K_\sigma | (K_0, K_1) \leftarrow \text{DSPF.Gen}(1^\lambda, \alpha, \beta)\}$$

We denote the execution of $\text{DSPF.Eval}(K_\sigma, \sigma, x)$ for every $x \in [N]$, i.e. the evaluation over the whole domain $[N]$, by $\text{DSPF.FullEval}(K_\sigma, \sigma)$.

OLE correlation. Analogously to the formal definition of a VOLE correlation given in Section 3.2, we formally define an OLE correlation. Our OLE correlation of size N over a finite field \mathbb{F}_p is given by $\mathbf{z}_1 = \mathbf{x}_0 \cdot \mathbf{x}_1 + \mathbf{z}_0$, where $\mathbf{x}_0, \mathbf{x}_1, \mathbf{z}_0, \mathbf{z}_1 \in \mathbb{F}_p^N$. Moreover, we require \mathbf{x}_0 and \mathbf{x}_1 being computed by a pseudorandom generator (PRG). Formally, we define the reverse-sampleable target correlation with setup $(\text{Setup}_{\text{OLE}}, \mathcal{Y}_{\text{OLE}})$ of size N over a field \mathbb{F}_p as

$$\begin{aligned} \text{mk} &= (\rho_0, \rho_1) \leftarrow \text{Setup}_{\text{OLE}}(1^\lambda, \rho_0, \rho_1), \\ ((F_0(\rho_0), \mathbf{z}_0), (F_1(\rho_1), \mathbf{z}_1)) &\leftarrow \mathcal{Y}_{\text{OLE}}(1^\lambda, \text{mk}) \quad \text{such that} \\ \mathbf{z}_0 &\stackrel{\$}{\leftarrow} \mathbb{F}_p^N \text{ and } \mathbf{z}_1 = F_0(\rho_0) \cdot F_1(\rho_1) + \mathbf{z}_0, \end{aligned} \quad (3)$$

where F_0, F_1 being pseudorandom generators (PRG). Note that while the Setup algorithm for our OLE and VOLE correlation essentially is the identity function, the algorithm might be more complex for other correlations. The reverse-sampling algorithm is defined such that

$$\begin{aligned} (F_1(\rho_1), F_0(\rho_0) \cdot F_1(\rho_1) + \mathbf{z}_0) &\leftarrow \text{RSample}_{\text{OLE}}(1^\lambda, \text{mk}, 0, (F_0(\rho_0), \mathbf{z}_0)) \text{ and} \\ (F_0(\rho_0), \mathbf{z}_1 - F_0(\rho_0) \cdot F_1(\rho_1)) &\leftarrow \text{RSample}_{\text{OLE}}(1^\lambda, \text{mk}, 1, (F_1(\rho_1), \mathbf{z}_1)). \end{aligned}$$

PCG constructions. The OLE construction is derived from [19, Fig. 1]. However, we extend it by the reusability feature by deriving the sparse polynomials normally sampled in PCG.Gen by applying a random oracle on seeds provided as input to the programmable key generation PCG.Gen_p .

Construction 3: Reusable PCG for $\mathcal{Y}_{\text{OLE}}^R$

Let λ be the security parameter, $\tau = \tau(\lambda)$ be the noise weight, $c \geq 2$ the compression factor, $p = p(\lambda)$ a modulus, $N = N(\lambda)$ a degree, and $R_p = \mathbb{Z}_p[X]/F(X)$ be a ring for a degree- N $F(X) \in \mathbb{Z}_p[X]$. Further, let $(\text{DSPF.Gen}, \text{DSPF.Eval})$ be a FSS scheme for sums of τ^2 -point functions with domain $[2N]$ and range \mathbb{Z}_p . Finally, let $\mathcal{H} : \{0, 1\}^\lambda \rightarrow ([N]^\tau \times (\mathbb{Z}_p^*)^\tau)^c$ be a random oracle.

Correlation: The target correlation $\mathcal{Y}_{\text{OLE}}^R$ over ring R_p is defined as

$$\begin{aligned} \text{mk} &= (\rho_0, \rho_1) \leftarrow \text{Setup}_{\text{OLE}}^R(1^\lambda, \rho_0, \rho_1) \\ ((x_0, z_0), (x_1, z_1)) &\leftarrow \mathcal{Y}_{\text{OLE}}^R(1^\lambda, \text{mk}) \text{ such that} \\ x_0 &= F_0(\rho_0), x_1 = F_1(\rho_1), z_0 \stackrel{\$}{\leftarrow} R_p, z_1 = x_0 \cdot x_1 - z_0 \\ (x_\sigma, x_0 \cdot x_1 - z_\sigma) &\leftarrow \text{RSample}_{\text{OLE}}^R(1^\lambda, \text{mk}, \sigma, (x_\sigma, z_\sigma)) \text{ where} \\ x_0 &= F_0(\rho_0), x_1 = F_1(\rho_1) \end{aligned}$$

with F_0 and F_1 being PRGs. As proposed by [19], R_p can be constructed to be isomorphic to N copies of \mathbb{Z}_p . This allows the direct transformation of one OLE over R_p into N independent OLEs over \mathbb{Z}_p .

Public Input: Random R^c - LPN polynomials $a_2, \dots, a_c \in R_p$, defining the vector $\mathbf{a} = (1, a_2, \dots, a_c)$.

$\text{PCG.Gen}_p(1^\lambda, \rho_0, \rho_1)$:

1. Compute $\{(\alpha_\sigma^i, \beta_\sigma^i)\}_{i \in [c]} \leftarrow \mathbf{H}(\rho_\sigma)$ for $\sigma \in \{0, 1\}$ where each $\alpha_\sigma^i \in [N]^\tau$ and each $\beta_\sigma^i \in (\mathbb{Z}_p^*)^\tau$.
2. For $i, j \in [c]$, sample FSS keys $(K_0^{(i,j)}, K_1^{(i,j)}) \xleftarrow{\$} \text{DSPF.Gen}(1^\lambda, \alpha_0^i \boxplus \alpha_1^j, \beta_0^i \boxtimes \beta_1^j)$.
3. For $\sigma \in \{0, 1\}$, define $\mathbf{k}_\sigma = (\{(\alpha_\sigma^i, \beta_\sigma^i)\}_{i \in [c]}, \{K_\sigma^{(i,j)}\}_{i,j \in [c]})$.
4. Output $(\mathbf{k}_0, \mathbf{k}_1)$.

PCG.Expand $(\sigma, \mathbf{k}_\sigma)$:

5. Parse \mathbf{k}_σ as $(\{(\alpha_\sigma^i, \beta_\sigma^i)\}_{i \in [c]}, \{K_\sigma^{(i,j)}\}_{i,j \in [c]})$.
6. For $i \in [c]$, define (over \mathbb{Z}_p) the degree $< N$ polynomial:

$$e_\sigma^i(X) = \sum_{k \in [\tau]} \beta_\sigma^i[k] \cdot X^{\alpha_\sigma^i[k]}$$

and compose all e_σ^i (for $i \in [c]$) to a length- c vector \mathbf{e}_σ .

7. For $i, j \in [c]$, compute $u_\sigma^{i+c(j-1)} \leftarrow \text{DSPF.FullEval}(\sigma, K_\sigma^{(i,j)})$ and view this as a degree $< 2N$ polynomial. Compose all u_σ^i (for $i \in [c^2]$) to a length- c^2 vector $\mathbf{v}_\sigma \bmod F(X)$.
8. Compute $x_\sigma = \langle \mathbf{a}, \mathbf{e}_\sigma \rangle \bmod F(X)$ and $z_\sigma = \langle \mathbf{a} \otimes \mathbf{a}, \mathbf{v}_\sigma \rangle \bmod F(X)$.
9. Output (x_σ, z_σ) .

From the previous construction, we derive a VOLE construction in a straightforward way.

Construction 4: Reusable PCG for $\mathcal{Y}_{\text{VOLE}}^R$

Let λ be the security parameter, $\tau = \tau(\lambda)$ be the noise weight, $c \geq 2$ the compression factor, $p = p(\lambda)$ a modulus, $N = N(\lambda)$ a degree, and $R_p = \mathbb{Z}_p[X]/F(X)$ be a ring for a degree- N $F(X) \in \mathbb{Z}_p[X]$. Further, let $(\text{DSPF.Gen}, \text{DSPF.Eval})$ be a FSS scheme for sums of τ point functions with domain $[N]$ and range \mathbb{Z}_p . Finally, let $\mathcal{H} : \{0, 1\}^\lambda \rightarrow ([N]^\tau \times (\mathbb{Z}_p^*)^\tau)^c$ be a random oracle.

Correlation: The target correlation $\mathcal{Y}_{\text{VOLE}}^R$ over ring R_p is defined as

$$\begin{aligned} \mathbf{mk} &= (\rho, x) \leftarrow \text{Setup}_{\text{OLE}}^R(1^\lambda, \rho, x) \\ ((y, z_0), (x, z_1)) &\leftarrow \mathcal{Y}_{\text{OLE}}^R(1^\lambda, \mathbf{mk}) \text{ such that} \\ y &= F(\rho), z_0 \xleftarrow{\$} R_p, z_1 = x \cdot y - z_0 \\ (x, x \cdot F(\rho) - z_0) &\leftarrow \text{RSample}_{\text{VOLE}}^R(1^\lambda, \mathbf{mk}, 0, (F(\rho), z_0)) \\ (F(\rho), x \cdot y - z_1) &\leftarrow \text{RSample}_{\text{VOLE}}^R(1^\lambda, \mathbf{mk}, 1, (x, z_1)) \end{aligned}$$

with F being a PRG. As proposed by [19], R_p can be constructed to be isomorphic to N copies of \mathbb{Z}_p . This allows the direct transformation of one VOLE over R_p into N independent VOLEs over \mathbb{Z}_p .

Public Input: Random R^c -LPN polynomials $a_2, \dots, a_c \in R_p$, defining the vector $\mathbf{a} = (1, a_2, \dots, a_c)$.

PCG.Gen_p $(1^\lambda, \rho_0, \rho_1)$:

1. Parse ρ_1 as x and compute $\{(\alpha^i, \beta^i)\}_{i \in [c]} \leftarrow \mathbf{H}(\rho_0)$ where $x \in \mathbb{Z}_p^*$, each $\alpha^i \in [N]^\tau$ and each $\beta^i \in (\mathbb{Z}_p^*)^\tau$.

2. For $i \in [c]$, sample FSS keys $(K_0^i, K_1^i) \xleftarrow{\$} \text{DSPF.Gen}(1^\lambda, \alpha^i, x \cdot \beta^i)$.
3. For $\sigma \in \{0, 1\}$, define $\mathbf{k}_\sigma = (\rho_\sigma, \{K_\sigma^i\}_{i \in [c]})$.
4. Output $(\mathbf{k}_0, \mathbf{k}_1)$.

PCG.Expand $(\sigma, \mathbf{k}_\sigma)$:

5. If $\sigma = 0$, parse \mathbf{k}_0 as $(\rho_0, \{K_0^i\}_{i \in [c]})$ and compute $\{(\alpha^i, \beta^i)\}_{i \in [c]} \leftarrow \mathcal{H}(\rho_0)$ where each $\alpha^i \in [N]^\tau$ and each $\beta^i \in (\mathbb{Z}_p^*)^\tau$. Then, for $i \in [c]$, define (over \mathbb{Z}_p) the degree $< N$ polynomial:

$$e^i(X) = \sum_{k \in [\tau]} \beta^i[k] \cdot X^{\alpha^i[k]}$$

and compose all e^i (for $i \in [c]$) to a length- c vector \mathbf{e} .

6. If $\sigma = 1$, parse \mathbf{k}_1 as $(x, \{K_1^i\}_{i \in [c]})$.
7. For $i \in [c]$, compute $u_\sigma^i \leftarrow \text{DSPF.FullEval}(\sigma, K_\sigma^i)$ and view the result as a degree $< N$ polynomial. Compose all u_σ^i (for $i \in [c]$) to a length- c vector $\mathbf{v}_\sigma \bmod F(X)$.
8. Compute $z_\sigma = \langle \mathbf{a}, \mathbf{v}_\sigma \rangle \bmod F(X)$.
9. If
 - $\sigma = 0$, compute $y = \langle \mathbf{a}, \mathbf{e} \rangle \bmod F(X)$ and output (y, z_0)
 - $\sigma = 1$, output (x, z_1) .

Security. We state the following Theorems:

Theorem 3. *Assume the $R^c\text{-LPN}_{R_p, 1, \tau}$ assumption holds and that DSPF is a secure instantiation of a distributed sum of point functions. Then, Construction 3 is a secure reusable PCG for OLE correlations over R_p in the random oracle model.*

Theorem 4. *Assume the $R^c\text{-LPN}_{R_p, 1, \tau}$ assumption holds and that DSPF is a secure instantiation of a distributed sum of point functions. Then, Construction 4 is a secure reusable PCG for VOLE correlations over R_p in the random oracle model.*

In the following, we provide a proof sketch for Theorem 3. A proof sketch for Theorem 4 follows in a straight-forward way.

Proof. To show that Construction 3 is a secure reusable PCG, we need to show programmability, pseudorandom \mathcal{Y} -correlated outputs, security and key indistinguishability.

Programmability can be shown, by defining ϕ_σ as a function, that first computes $\{(\alpha_\sigma^i, \beta_\sigma^i)\}_{i \in [c]} \leftarrow \mathcal{H}(\rho_\sigma)$, expands these to $\mathbf{e}_\sigma \in R_p^c$ as done in the PCG.Expand algorithm, and then outputs $\langle \mathbf{a}, \mathbf{e}_\sigma \rangle$.

Pseudorandom \mathcal{Y} -correlated outputs can be shown via a sequence of games. First, we replace the PRG F_σ in \mathcal{Y} by ϕ_σ . As the random oracle ensures that the secrets c_σ^* are sampled uniformly at random, indistinguishability can be shown via a reduction to the $R^c\text{-LPN}_{R_p, 1, \tau}$ assumption. Next, we skip the DSPF key

generation and full evaluation during the expansion. Instead, we directly sample $z_0 \in_R R_p$ and define $z_1 = x_0 \cdot x_1 - z_0$. Here, indistinguishability can be shown analogously to the correctness proof in [19]. Note that in the previous game for every $i, j \in [c]$, it holds that

$$e_0^i(X) \cdot e_1^j(X) = \sum_{k, l \in [\tau]} \beta_0^i[k] \cdot \beta_1^j[l] \cdot X^{\alpha_0^i[k] \cdot \alpha_1^j[l]}.$$

Therefore, parties can obtain an additive sharing of this product by fully evaluating the (i, j) -th DSPF instance. It follows that $u_0^{i+c(j-1)} + u_1^{i+c(j-1)} = e_0^i(X) \cdot e_1^j(X)$, and hence, $\mathbf{v} = \mathbf{e}_0 \otimes \mathbf{e}_1$. This observation yields the following relation of the outputs:

$$\begin{aligned} z_0 + z_1 &= \langle \mathbf{a} \otimes \mathbf{a}, \mathbf{v}_0 + \mathbf{v}_1 \rangle = \langle \mathbf{a} \otimes \mathbf{a}, \mathbf{e}_0 + \mathbf{e}_1 \rangle \\ &= \langle \mathbf{a}, \mathbf{e}_0 \rangle \cdot \langle \mathbf{a}, \mathbf{e}_1 \rangle = x_0 \cdot x_1 \end{aligned}$$

As the correlation of (x_0, x_1, z_0, z_1) is the same in both games, the computation of x_0 and x_1 remains untouched, and the DSPF implies that each z_σ is individually pseudorandom, both games are computationally indistinguishable. In the resulting game, the challenger executes the exact same steps independent of the coin b . Therefore, it follows that any adversary wins the final game with probability exactly $\frac{1}{2}$ which implies that any adversary wins the original security game with probability at most $\frac{1}{2} + \text{negl}$.

As $\text{RSample}_{\text{OLE}}^R$ executes the same steps as the forward sampling $\mathcal{Y}_{\text{OLE}}^R$, security can be shown analogously to the pseudorandom \mathcal{Y} -correlated outputs property.

Key indistinguishability follows from the security property of the DSPF scheme via a sequence of game hops. We replace one by one the DSPF-keys in \mathbf{k}_σ with ones produced by the DSPF-simulator. Indistinguishability between games can be proven via reductions to the security property of the DSPF scheme. Finally, we remove in one more game the PCG key generation and the assignment of $\rho_{1-\sigma}$ as both steps become redundant. The final game is completely independent of the choice of b such that the success probability of \mathcal{A} is exactly $\frac{1}{2}$ which shows that the success probability of \mathcal{A} in the initial game is at most $\frac{1}{2} + \text{negl}(\lambda)$.

F Reusable Pseudorandom Correlation Function

On a high level, a pseudorandom correlation function (PCF) allows two parties to generate a large amount of correlated randomness from short seeds. PCF extends the notion of a pseudorandom correlation generator (PCG) in a similar way as a pseudorandom function extends a pseudorandom generator. While a PCG generates a large batch of correlated randomness during one-time expansion, a PCF allows the creation of correlation samples on the fly.

A PCF consists of two algorithms, **Gen** and **Eval**. The **Gen** algorithm computes a pair of short keys distributed to two parties. Then, each party can locally

evaluate the **Eval** algorithm using its key and public input to generate an output of the target correlation. One example of such a correlation is the oblivious linear evaluation (OLE) correlation, defined by a pair of random values (y_0, y_1) where $y_0 = (a, u)$ and $y_1 = (s, v)$ such that $v = as + u$. Other meaningful correlations are oblivious transfer (OT) and multiplication triples.

PCFs are helpful in two- and multi-party protocols, where parties first set up correlated randomness and then use this data to speed up the computation [39, 3, 48].

This section presents our definition of reusable PCFs, extending the definition of programmable PCFs from [18]. Furthermore, we state constructions of reusable PCFs and argue why they satisfy our new definition in Appendix G.

Our modifications and extensions of the definition [18] reflect the challenges we faced when using PCFs as black-box primitives in our threshold BBS+ protocol. We present our definition and highlight these challenges and changes in the following.

F.1 Definition

Similar to PCGs, PCFs realize a target correlation \mathcal{Y} . While PCFs output single correlation outputs instead of a bunch of correlation as PCGs, we need to slightly adapt the definition of a target correlation. We emphasize the modification in the following.

We formally define a *target correlation* as a tuple of probabilistic algorithms $(\text{Setup}, \mathcal{Y})$, where **Setup** takes two inputs and creates a master key **mk**. These inputs enable fixing parts of the correlation, e.g., the fixed value s . Algorithm \mathcal{Y} uses the master key and an index i to sample correlation outputs. The index i helps to sample the same value if one of the **Setup** inputs is identical for multiple invocations. The input i is not necessary for correlations for PCGs since the output of PCG expansion is a bunch of correlation. For PCFs, the output of the evaluation is a single correlation tuple. Thus, we need the index i to sample the same value if one of the **Setup** inputs is identical for multiple PCF invocations.

Definition 5 (Reverse-sampleable and indexable target correlation with setup). Let $\ell_0(\lambda), \ell_1(\lambda) \leq \text{poly}(\lambda)$ be output length functions. Let $(\text{Setup}, \mathcal{Y})$ be a tuple of probabilistic algorithms, such that **Setup** on input 1^λ and two parameters ρ_0, ρ_1 returns a master key **mk**; algorithm \mathcal{Y} on input 1^λ , **mk**, and index i returns a pair of outputs $(y_0^{(i)}, y_1^{(i)}) \in \{0, 1\}^{\ell_0(\lambda)} \times \{0, 1\}^{\ell_1(\lambda)}$.

We say that the tuple $(\text{Setup}, \mathcal{Y})$ defines a reverse-sampleable and indexable target correlation with setup if there exists a probabilistic polynomial time algorithm **RSample** that takes as input 1^λ , **mk**, $\sigma \in \{0, 1\}$, $y_\sigma^{(i)} \in \{0, 1\}^{\ell_\sigma(\lambda)}$ and i , and outputs $y_{1-\sigma}^{(i)} \in \{0, 1\}^{\ell_{1-\sigma}(\lambda)}$, such that for all $\sigma \in \{0, 1\}$, for all **mk**, **mk'** in the range of **Setup** for arbitrary but fixed input ρ_σ , and all $i \in \{0, 1\}^*$ the following

distributions are statistically close:

$$\begin{aligned} & \{(y_0^{(i)}, y_1^{(i)}) | (y_0^{(i)}, y_1^{(i)}) \xleftarrow{\$} \mathcal{Y}(1^\lambda, \text{mk}, i)\} \\ & \{(y_0^{(i)}, y_1^{(i)}) | (y_0'^{(i)}, y_1'^{(i)}) \xleftarrow{\$} \mathcal{Y}(1^\lambda, \text{mk}', i), \\ & y_\sigma^{(i)} \leftarrow y_\sigma'^{(i)}, y_{1-\sigma}^{(i)} \leftarrow \text{RSample}(1^\lambda, \text{mk}, \sigma, y_\sigma, i)\}. \end{aligned}$$

Given the definition of a reverse-sampleable and indexable correlation with setup, we define our primitive called *strong reusable PCF* (srPCF). Our definition builds on the definition of a strong PCF of Boyle et al. [18] and extends it by a reusability feature. Note that [18] presents a separate definition of this reusability feature for PCFs, but this property also affects the other properties of a PCF. Therefore, we merge these definitions. Additionally, the reusability definition of Boyle et al. works only for the semi-honest setting, while our definition covers malicious adversaries. The crucial point to cover malicious adversaries is to allow the corrupted party to choose an arbitrary value as its input to the key generation. Our definitions give this power to the adversary, while the definitions of Boyle et al. use randomly chosen inputs.

A PCF must fulfill two properties. First, the pseudorandomness property intuitively states that the joint outputs of the **Eval** algorithm are computationally indistinguishable from outputs of the correlation \mathcal{Y} . Second, the security property intuitively guarantees that the PCF output of party $P_{1-\sigma}$ is indistinguishable from a reverse-sampled value. Indistinguishability holds even if the adversary corrupts party P_σ and learns its key. Hence, this property provides security against an insider.

Similarly to the notions of weak and strong PRFs, there exist the notions of *weak* and *strong* PCFs. For a weak PCF, we consider the **Eval** algorithm to be executed on randomly chosen inputs, while for a strong PCF, we consider arbitrarily chosen inputs. Boyle et al. [18] showed a generic transformation from a weak to a strong PCF using a hash function modeled as a programmable random oracle. In Appendix G, we present constructions for weak srPCFs, which then yield strong srPCFs based on the transformation of Boyle et al.

A PCF needs to meet two additional requirements to satisfy the reusability features. First, an adversary cannot learn any information about the other party's input used for the key generation from its own key. This is modeled by the key indistinguishability property and the corresponding game in Figure 10. In the game, the challenger samples two random values and uses one for the key generation. Then, given the corrupted party's key and the random values, the adversary has to identify which of the two random value was used. Second, two efficiently computable functions must exist to compute the reusable parts of the correlation from the setup input and the public evaluation input. Formally, we state the definition of a strong reusable PCF next.

Definition 6 (Strong reusable pseudorandom correlation function (srPCF)). Let $(\text{Setup}, \mathcal{Y})$ be a reverse-sampleable and indexable correlation with setup which has output length functions $\ell_0(\lambda), \ell_1(\lambda)$, and let $\lambda \leq \eta(\lambda) \leq \text{poly}(\lambda)$

be an input length function. Let $(\text{PCF.Gen}, \text{PCF.Eval})$ be a pair of algorithms with the following syntax:

- $\text{PCF.Gen}(1^\lambda, \rho_0, \rho_1)$ is a probabilistic polynomial-time algorithm that on input the security parameter 1^λ and reusable inputs ρ_0, ρ_1 outputs a pair of keys (k_0, k_1) .
- $\text{PCF.Eval}(\sigma, k_\sigma, x)$ is a deterministic polynomial-time algorithm that on input $\sigma \in \{0, 1\}$, key k_σ and input value $x \in \{0, 1\}^{\eta(\lambda)}$ outputs a value $y_\sigma \in \{0, 1\}^{\ell_\sigma(\lambda)}$.

We say $(\text{PCF.Gen}, \text{PCF.Eval})$ is a strong reusable pseudorandom correlation function (srPCF) for $(\text{Setup}, \mathcal{Y})$, if the following conditions hold:

- **Strong pseudorandom \mathcal{Y} -correlated outputs.** For every non-uniform adversary \mathcal{A} of size $\text{poly}(\lambda)$ asking at most $\text{poly}(\lambda)$ queries to the oracle $\mathcal{O}_b(\cdot)$, it holds

$$\left| \Pr[\text{Exp}_{\mathcal{A}}^{\text{s-pr}}(\lambda) = 1] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

for all sufficiently large λ , where $\text{Exp}_{\mathcal{A}}^{\text{s-pr}}(\lambda)$ is as defined in Figure 8.

- **Strong security.** For each $\sigma \in \{0, 1\}$ and non-uniform adversary \mathcal{A} of size $\text{poly}(\lambda)$ asking at most $\text{poly}(\lambda)$ queries to oracle $\mathcal{O}_b(\cdot)$, it holds

$$\left| \Pr[\text{Exp}_{\mathcal{A}, \sigma}^{\text{s-sec}}(\lambda) = 1] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

for all sufficiently large λ , where $\text{Exp}_{\mathcal{A}, \sigma}^{\text{s-sec}}(\lambda)$ is as defined in Figure 9.

- **Programmability.** There exist public efficiently computable functions f_0, f_1 , such that for all $x \in \{0, 1\}^{\eta(\lambda)}$ and all $\rho_0, \rho_1 \in \{0, 1\}^*$

$$\Pr \left[\begin{array}{l} (k_0, k_1) \leftarrow \text{PCF.Gen}(1^\lambda, \rho_0, \rho_1) \\ (a, c) \leftarrow \text{PCF.Eval}(0, k_0, x), \\ (b, d) \leftarrow \text{PCF.Eval}(1, k_1, x) \end{array} : \begin{array}{l} a = f_0(\rho_0, x) \\ b = f_1(\rho_1, x) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

- **Key indistinguishability.** For any $\sigma \in \{0, 1\}$ and non-uniform adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, it holds

$$\Pr[\text{Exp}_{\mathcal{A}, \sigma}^{\text{key-ind}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda)$$

for all sufficiently large λ , where $\text{Exp}_{\mathcal{A}, \sigma}^{\text{key-ind}}$ is as defined in Figure 10.

F.2 Correlations

Here, we state the correlations required for our PCF-based precomputation protocol (cf. Appendix H). As these correlations differ slightly from the correlations required by our PCG-based offline phase (cf. Section 5), we state them in the

Exp_{\mathcal{A}}^{s-pr}(λ) :	$\mathcal{O}_0(x)$:
$(\rho_0, \rho_1) \leftarrow \mathcal{A}_0(1^\lambda)$	if $(x, y_0, y_1) \in \mathcal{Q}$:
mk $\leftarrow \text{Setup}(1^\lambda, \rho_0, \rho_1)$	return (y_0, y_1)
$(k_0, k_1) \leftarrow \text{PCF.Gen}(1^\lambda, \rho_0, \rho_1)$	else :
$\mathcal{Q} = \emptyset$	$(y_0, y_1) \leftarrow \mathcal{Y}(1^\lambda, \text{mk}, x)$
$b \xleftarrow{\$} \{0, 1\}$	$\mathcal{Q} = \mathcal{Q} \cup \{(x, y_0, y_1)\}$
$b' \leftarrow \mathcal{A}_1^{\mathcal{O}_b(\cdot)}(1^\lambda)$	return (y_0, y_1)
if $b = b'$ return 1	$\mathcal{O}_1(x)$:
else return 0	for $\sigma \in \{0, 1\}$:
	$y_\sigma \leftarrow \text{PCF.Eval}(\sigma, k_\sigma, x)$
	return (y_0, y_1)

Fig. 8: Strong pseudorandom \mathcal{Y} -correlated outputs of a PCF.

Exp_{\mathcal{A}, σ}^{s-sec}(λ) :	$\mathcal{O}_0(x)$:
$(\rho_0, \rho_1) \leftarrow \mathcal{A}_0(1^\lambda)$	$y_{1-\sigma} \leftarrow \text{PCF.Eval}(1-\sigma, k_{1-\sigma}, x)$
mk $\leftarrow \text{Setup}(1^\lambda, \rho_0, \rho_1)$	return $y_{1-\sigma}$
$(k_0, k_1) \leftarrow \text{PCF.Gen}(1^\lambda, \rho_0, \rho_1)$	$\mathcal{O}_1(x)$:
$b \xleftarrow{\$} \{0, 1\}$	$y_\sigma \leftarrow \text{PCF.Eval}(\sigma, k_\sigma, x)$
$b' \leftarrow \mathcal{A}_1^{\mathcal{O}_b(\cdot)}(1^\lambda, \sigma, k_\sigma)$	$y_{1-\sigma} \leftarrow \text{RSample}(1^\lambda, \text{mk}, \sigma, y_\sigma, x)$
if $b = b'$ return 1	return $y_{1-\sigma}$
else return 0	

Fig. 9: Strong security of a PCF.

following for completeness. Note that our protocol for the optimized BBS+ signature scheme requires only VOLE correlations. In contrast, our protocol for the legacy BBS+ signature scheme (cf. Appendix B) additionally requires OLE correlations. Therefore, we present definitions for both of them in the following. Our OLE correlation over ring R is given by $c_1 = ab + c_0$, where $a, b, c_0, c_1 \in R$. Moreover, we require a and b being computed by a weak pseudorandom function (PRF). Formally, we define the reverse-sampleable and indexable target correlation with setup $(\text{Setup}_{\text{OLE}}, \mathcal{V}_{\text{OLE}})$ over ring R as

$$\begin{aligned}
 (k, k') &\leftarrow \text{Setup}_{\text{OLE}}(1^\lambda, k, k'), \\
 ((F_k(i), u), (F_{k'}(i), v)) &\leftarrow \mathcal{V}_{\text{OLE}}(1^\lambda, (k, k'), i) \quad \text{such that} \\
 v &= F_k(i) \cdot F_{k'}(i) + u,
 \end{aligned}$$

$$\text{Exp}_{\mathcal{A}, \sigma}^{\text{key-ind}}(\lambda) :$$

$$b \xleftarrow{\$} \{0, 1\}$$

$$\rho_{1-\sigma}^{(0)}, \rho_{1-\sigma}^{(1)} \xleftarrow{\$} \{0, 1\}^*$$

$$\rho_{1-\sigma} \leftarrow \rho_{1-\sigma}^{(b)}$$

$$\rho_{\sigma} \leftarrow \mathcal{A}_0(1^\lambda)$$

$$(k_0, k_1) \leftarrow \text{PCF.Gen}(1^\lambda, \rho_0, \rho_1)$$

$$b' \leftarrow \mathcal{A}_1(1^\lambda, k_\sigma, \rho_{1-\sigma}^{(0)}, \rho_{1-\sigma}^{(1)})$$

$$\text{if } b' = b \text{ return } 1$$

$$\text{else return } 0$$

Fig. 10: Key Indistinguishability of a reusable PCF.

where $u \xleftarrow{\$} R, u \in R$ and F being a (PRF) with key k, k' . Note that while the **Setup** algorithm for our OLE and VOLE correlation essentially is the identity function, the algorithm might be more complex for other correlations. The reverse-sampling algorithm is defined such that $(F_{k'}(i), F_k(i) \cdot F_{k'}(i) + u) \leftarrow \text{RSample}_{\text{OLE}}(1^\lambda, (k, k'), 0, (F_k(i), u), i)$ and $(F_k(i), v - F_k(i) \cdot F_{k'}(i)) \leftarrow \text{RSample}_{\text{OLE}}(1^\lambda, (k, k'), 1, (F_{k'}(i), v), i)$.

Our VOLE correlation is the same as OLE but the value b is fixed over multiple correlation samples, i.e., $\vec{c}_1 = \vec{a}b + \vec{c}_0$, where each correlation sample contains one component of the vectors. We formally define the reverse-sampleable and indexable target correlation with setup $(\text{Setup}_{\text{VOLE}}, \mathcal{Y}_{\text{VOLE}})$ over ring R as

$$\begin{aligned} (k, b) &\leftarrow \text{Setup}_{\text{VOLE}}(1^\lambda, k, b), \\ ((F_k(i), u), (b, v)) &\leftarrow \mathcal{Y}_{\text{VOLE}}(1^\lambda, (k, b), i) \quad \text{such that} \\ v &= F_k(i) \cdot b + u, \end{aligned}$$

where $u \xleftarrow{\$} R, b, v \in R$ and F being a weak pseudorandom function (PRF) with key k . Note that b is fixed over all correlation samples, while u and v are not. The reverse-sampling algorithm is defined such that

$$\begin{aligned} (b, F_k(i) \cdot b + u) &\leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, (k, b), 0, (F_k(i), u), i) \text{ and} \\ (F_k(i), v - F_k(i) \cdot b) &\leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, (k, b), 1, (b, v), i). \end{aligned}$$

We state PCF constructions realizing these definitions of OLE and VOLE correlations in Appendix G. The VOLE PCF construction is taken from [18], and the OLE PCF follows a straightforward adaptation of the VOLE PCF.

G Reusable PCF Constructions

This section presents construction of reusable PCFs for VOLE and OLE correlations as defined in Appendix F.2. We first present the reusable PCF for VOLE and then for OLE.

The VOLE construction heavily builds on the constructions of [18], which provides only weak PCF. However, Boyle et al. presented a generic transformation from weak to strong PCF using a programmable random oracle. This transformation is also straightforwardly applicable to reusable PCFs. Therefore, we state a weak reusable PCF in the following and emphasize that this construction can be extended to a strong reusable PCF in the programmable random oracle model.

The following construction is taken from [18, Fig. 22]. It builds on a weak PRF F and a function secret sharing for the multiplication of F with a scalar.

Construction 5: Reusable PCF for $\mathcal{V}_{\text{VOLE}}$

Let $\mathcal{F} = \{F_k : \{0, 1\}^\eta \rightarrow R\}_{k \in \{0, 1\}^\lambda}$ be a weak PRF and $\text{FFS} = (\text{FFS.Gen}, \text{FFS.Eval})$ an FSS scheme for $\{c \cdot F_k\}_{c \in R, k \in \{0, 1\}^\lambda}$ with weak pseudorandom outputs. Let further $\rho_0 \in \{0, 1\}^\lambda, \rho_1 \in R$.
 $\text{PCF.Gen}_p(1^\lambda, \rho_0, \rho_1)$:

1. Set the weak PRF key $k \leftarrow \rho_0$ and $b \leftarrow \rho_1$.
2. Sample a pair of FSS keys $(K_0^{\text{FFS}}, K_1^{\text{FFS}}) \leftarrow \text{FFS.Gen}(1^\lambda, b \cdot F_k)$.
3. Output the keys $k_0 = (K_0^{\text{FFS}}, k)$ and $k_1 = (K_1^{\text{FFS}}, b)$.

$\text{PCF.Eval}(\sigma, k_\sigma, x)$: On input a random x :

- If $\sigma = 0$:
 1. Let $c_0 = -\text{FFS.Eval}(0, K_0^{\text{FFS}}, x)$.
 2. Let $a = F_k(x)$.
 3. Output (a, c_0) .
- If $\sigma = 1$:
 1. Let $c_1 = \text{FFS.Eval}(1, K_1^{\text{FFS}}, x)$.
 2. Output (b, c_1) .

Theorem 5. *Let $R = R(\lambda)$ be a finite commutative ring. Suppose there exists an FSS scheme for scalar multiples of a family of weak pseudorandom functions $\mathcal{F} = \{F_k : \{0, 1\}^\eta \rightarrow R\}_{k \in \{0, 1\}^\lambda}$. Then, there is a reusable PCF for the VOLE correlation over R as defined in Appendix F.2, given by Construction 5.*

Proof. Boyle et al. showed in their proof of [18, Theorem 5.3] that Construction 5 satisfies pseudorandom $\mathcal{V}_{\text{VOLE}}$ -correlated outputs and security. Although we slightly adapted our definition to consider reusable inputs, their argument still holds. Further, it is easy to see that programmability holds for functions $f_0(\rho_0, x) = F_{\rho_0}(x)$ and $f_1(\rho_1, x) = \rho_1$. Finally, key indistinguishability follows from the secrecy property of the FSS scheme. The secrecy property states that for every function f of the function family, there exists a simulator $\mathcal{S}(1^\lambda)$ such

that the output of \mathcal{S} is indistinguishable from the FSS keys generated correctly using the FFS.Gen-algorithm.

To briefly sketch the proof of key indistinguishability, we define a hybrid experiment, where inside the PCF key generation, we use \mathcal{S} to simulate FSS keys. These simulated FSS keys are used inside the PCF key, which is given to \mathcal{A}_1 . We can show via a reduction to the FSS secrecy that the original $\text{Exp}^{\text{key-ind}}$ game is indistinguishable from the hybrid experiment. For the hybrid experiment, it is easy to see that the adversary can only guess bit b' since the simulated PCF key is independent of $\rho_{1-\sigma}^{(0)}, \rho_{1-\sigma}^{(1)}$ and hence also independent of b . It follows that $\Pr[\text{Exp}_{\mathcal{A},\sigma}^{\text{key-ind}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda)$.

The construction of the reusable PCF for OLE correlations follows the same blueprint as our PCF construction for VOLE.

The following construction is generically based on a weak PRF and function secret sharing (FSS) for products of two weak PRFs.

Construction 6: Reusable PCF for \mathcal{Y}_{OLE}

Let $\mathcal{F} = \{F_k : \{0, 1\}^\eta \rightarrow R\}_{k \in \{0,1\}^\lambda}$ be a weak PRF and $\text{FFS} = (\text{FFS.Gen}, \text{FFS.Eval})$ an FSS scheme for $\{F_{k_0} \cdot F_{k_1}\}_{k_0, k_1 \in \{0,1\}^\lambda}$ with weak pseudorandom outputs. Let further $\rho_0, \rho_1 \in \{0, 1\}^\lambda$.

PCF.Gen_p($1^\lambda, \rho_0, \rho_1$):

1. Set the weak PRF keys $k \leftarrow \rho_0$ and $k' \leftarrow \rho_1$.
2. Sample a pair of FSS keys $(K_0^{\text{FFS}}, K_1^{\text{FFS}}) \leftarrow \text{FFS.Gen}(1^\lambda, F_k F_{k'})$.
3. Output the keys $k_0 = (K_0^{\text{FFS}}, k)$ and $k_1 = (K_1^{\text{FFS}}, k')$.

PCF.Eval(σ, k_σ, x): On input a random x :

- If $\sigma = 0$:
 1. Let $c_0 = -\text{FFS.Eval}(0, K_0^{\text{FFS}}, x)$.
 2. Let $a = F_k(x)$.
 3. Output (a, c_0) .
- If $\sigma = 1$:
 1. Let $c_1 = \text{FFS.Eval}(1, K_1^{\text{FFS}}, x)$.
 2. Let $b = F_{k'}(x)$.
 3. Output (b, c_1) .

Theorem 6. *Let $R = R(\lambda)$ be a finite commutative ring. Suppose there exists an FSS scheme for multiplications of two elements of a family of weak pseudorandom functions $\mathcal{F} = \{F_k : \{0, 1\}^\eta \rightarrow R\}_{k \in \{0,1\}^\lambda}$. Then, there is a reusable PCF for the OLE correlation over R , given by Construction 6.*

We omit the proof as it follows the same arguments as the proof of Theorem 5.

H PCF-based Threshold Preprocessing Protocol

In this section, we state the PCF-based instantiation of $\mathcal{F}_{\text{Prep}}$. As it is conceptually very similar to the PCG-based instantiation in Section 5, we omit a detailed

description and intuition here. We refer the reader to Section 5 for an intuition and a detailed description.

Our protocol $\pi_{\text{Prep}}^{\text{PCF}}$ builds on reusable PCFs for VOLE correlations. As ssid , which is used to evaluate the PCFs, is provided by the environment, we require strong reusable PCFs.

H.1 Setup Functionality

The setup functionality is identical to $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ just that the functionality generates PCF keys instead of PCG keys. For the sake of completeness, we formally state $\mathcal{F}_{\text{Setup}}^{\text{PCF}}$ next.

Functionality $\mathcal{F}_{\text{Setup}}^{\text{PCF}}$

Let $(\text{PCF}_{\text{VOLE}}.\text{Gen}, \text{PCF}_{\text{VOLE}}.\text{Eval})$ be an srPCF for VOLE correlations. The setup functionality interacts with parties P_1, \dots, P_n and ideal-world adversary \mathcal{S} .

Setup:

Upon receiving $(\text{setup}, \text{sid}, \rho_a^{(i)}, \text{sk}_i, \{\text{pk}_\ell^{(i)}\}_{\ell \in [n]})$ from every party P_i send (setup) to \mathcal{S} and do:

1. Check if $g_2^{\text{sk}_\ell} = \text{pk}_\ell^{(i)}$ for every $\ell, i \in [n]$. If the check fails, send **abort** to all parties and \mathcal{S} .

Else, compute for every pair of parties (P_i, P_j) :

$$(k_{i,j,0}^{\text{VOLE}}, k_{i,j,1}^{\text{VOLE}}) \leftarrow \text{PCF}_{\text{VOLE}}.\text{Gen}(1^\lambda, \rho_a^{(i)}, \text{sk}_j).$$

2. Send keys $(\text{sid}, \{k_{i,j,0}^{\text{VOLE}}, k_{j,i,1}^{\text{VOLE}}\}_{j \neq i})$ to every party P_i .

H.2 PCF-based Preprocessing Protocol

In this section, we formally present our PCF-based preprocessing protocol in the $(\mathcal{F}_{\text{KG}}, \mathcal{F}_{\text{Setup}})$ -hybrid model.

Construction 7: $\pi_{\text{Prep}}^{\text{PCF}}$

Let $(\text{PCF}_{\text{VOLE}}.\text{Gen}, \text{PCF}_{\text{VOLE}}.\text{Eval})$ be an srPCF for VOLE correlations.

We describe the protocol from the perspective of P_i .

Initialization. Upon receiving input $(\text{init}, \text{sid})$, do:

1. Send $(\text{keygen}, \text{sid})$ to \mathcal{F}_{KG} .
2. Upon receiving $(\text{sid}, \text{sk}_i, \text{pk}, \{\text{pk}_\ell^{(i)}\}_{\ell \in [n]})$ from \mathcal{F}_{KG} , sample $\rho_a^{(i)} \in \{0, 1\}^\lambda$ and send $(\text{setup}, \text{sid}, \rho_a^{(i)}, \text{sk}_i, \{\text{pk}_\ell^{(i)}\}_{\ell \in [n]})$ to $\mathcal{F}_{\text{Setup}}$.
3. Upon receiving $(\text{sid}, \{k_{i,j,0}^{\text{VOLE}}, k_{j,i,1}^{\text{VOLE}}\}_{j \neq i})$ from $\mathcal{F}_{\text{Setup}}$, output pk .

Tuple. Upon receiving input $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$, compute:

4. for $j \in \mathcal{T} \setminus \{i\}$:

- (a) $(a_i, c_{i,j,0}^{\text{VOLE}}) = \text{PCF}_{\text{VOLE}}.\text{Eval}(0, k_{i,j,0}^{\text{VOLE}}, \text{ssid})$,
 (b) $(\text{sk}_i, c_{j,i,1}^{\text{VOLE}}) = \text{PCF}_{\text{VOLE}}.\text{Eval}(1, k_{j,i,1}^{\text{VOLE}}, \text{ssid})$,
 5. $\beta_i = a_i \cdot L_{i,\tau} \text{sk}_i + \sum_{j \in \mathcal{T} \setminus \{i\}} (L_{i,\tau} c_{j,i,1}^{\text{VOLE}} - L_{j,\tau} c_{i,j,0}^{\text{VOLE}})$

Finally, output $(\text{sid}, \text{ssid}, a_i, \beta_i)$.

Theorem 7. *Let PCF_{VOLE} be an srPCF for VOLE correlations as defined in Appendix F.2. Then, protocol $\pi_{\text{Prep}}^{\text{PCF}}$ UC-realizes $\mathcal{F}_{\text{Prep}}$ in the $(\mathcal{F}_{\text{KG}}, \mathcal{F}_{\text{Setup}})$ -hybrid model in the presence of malicious adversaries controlling up to $t - 1$ parties.*

The proof works analogously to the proof of Theorem 2, which is presented in Appendix M. Therefore, we omit the proof of Theorem 7 for the sake of conciseness.

I Ideal Threshold Signature Functionality

In this section, we state our ideal threshold functionality $\mathcal{F}_{\text{tsig}}$ on which we base our security analysis of the online protocol (cf. Theorem 1). The functionality is presented in the universal composability (UC) framework and we refer the reader to Appendix C for a brief introduction into the UC framework and its notation. $\mathcal{F}_{\text{tsig}}$ is a modification of the functionality proposed by Canetti et al. [29]. First, we allow the parties to specify a set of signers \mathcal{T} during the signing request. This allows us to account for a flexible threshold of signers instead of requiring all n parties to sign. Second, we model the signed message as an array of messages. This change accounts for signature schemes allowing signing k messages simultaneously, such as BBS+. Third, we remove the identifiability property, the key-refresh, and the corruption/decorruption interface. The key-refresh and the corruption/decorruption interface are not required in our scenario as we consider a static adversary in contrast to the mobile adversary in [29]. Fourth, we allow every party to sign only one message per ssid . Finally, at the end of the signing phase, honest parties might output **abort** instead of a valid signature. This modification is due to our protocol not providing robustness or identifiable abort.

Next, we state the full formal description of our threshold signature functionality $\mathcal{F}_{\text{tsig}}$.

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

The functionality is parameterized by a threshold parameter t . We denote a set of t parties by \mathcal{T} . For a specific session id sid , the sub-procedures *Signing* and *Verification* can only be executed once a tuple $(\text{sid}, \mathcal{V})$ is recorded.

Key-generation.

1. Upon receiving $(\text{keygen}, \text{sid})$ from some party P_i , interpret $\text{sid} = (\dots, \mathbf{P})$, where $\mathbf{P} = (P_1, \dots, P_n)$.
 - If $P_i \in \mathbf{P}$, send to \mathcal{S} and record $(\text{keygen}, \text{sid}, i)$.

- Otherwise ignore the message.
- 2. Once $(\text{keygen}, \text{sid}, i)$ is recorded for all $P_i \in \mathbf{P}$, send $(\text{pubkey}, \text{sid})$ to the adversary \mathcal{S} and do:
 - (a) Upon receiving $(\text{pubkey}, \text{sid}, \mathcal{V})$ from \mathcal{S} , record $(\text{sid}, \mathcal{V})$.
 - (b) Upon receiving $(\text{pubkey}, \text{sid})$ from $P_i \in \mathbf{P}$, output $(\text{pubkey}, \text{sid}, \mathcal{V})$ if it is recorded. Else ignore the message.

Signing.

1. Upon receiving $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m} = (m_1, \dots, m_k))$ with $\mathcal{T} \subseteq \mathbf{P}$, from $P_i \in \mathcal{T}$ and no tuple $(\text{sign}, \text{sid}, \text{ssid}, \cdot, \cdot, i)$ is stored, send to \mathcal{S} and record $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, i)$.
2. Upon receiving $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m} = (m_1, \dots, m_k), i)$ from \mathcal{S} , record $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, i)$ if $P_i \in \mathcal{C}$. Else ignore the message.
3. Once $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, i)$ is recorded for all $P_i \in \mathcal{T}$, send $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m})$ to the adversary \mathcal{S} .
4. Upon receiving $(\text{sig}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, \sigma, \mathcal{I})$ from \mathcal{S} , where $\mathcal{I} \subseteq \mathcal{T} \setminus \mathcal{C}$, do:
 - If there exists a record $(\text{sid}, \mathbf{m}, \sigma, 0)$, output an error.
 - Else, record $(\text{sid}, \mathbf{m}, \sigma, \mathcal{V}(\mathbf{m}, \sigma))$, send $(\text{sig}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, \sigma)$ to all $P_i \in \mathcal{T} \setminus (\mathcal{C} \cup \mathcal{I})$ and send $(\text{sig}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, \text{abort})$ to all $P_i \in \mathcal{T} \cap \mathcal{I}$.

Verification.

Upon receiving $(\text{verify}, \text{sid}, \mathbf{m} = (m_1, \dots, m_k), \sigma, \mathcal{V}')$ from a party Q , send the tuple $(\text{verify}, \text{sid}, \mathbf{m}, \sigma, \mathcal{V}')$ to \mathcal{S} and do:

- If $\mathcal{V}' = \mathcal{V}$ and a tuple $(\text{sid}, \mathbf{m}, \sigma, \beta')$ is recorded, then set $\beta = \beta'$.
- Else, if $\mathcal{V}' = \mathcal{V}$ and less than t parties in \mathbf{P} are corrupted, set $\beta = 0$ and record $(\text{sid}, \mathbf{m}, \sigma, 0)$.
- Else, set $\beta = \mathcal{V}'(\mathbf{m}, \sigma)$.

Output $(\text{verified}, \text{sid}, \mathbf{m}, \sigma, \beta)$ to Q .

J Proof of Theorem 1

This section presents the proof of our online protocol, i.e., Theorem 1.

Proof. We construct a simulator \mathcal{S} that interacts with the environment and the ideal functionality $\mathcal{F}_{\text{tsig}}$. Since the security statement for UC requires that for every real-world adversary \mathcal{A} , there is a simulator \mathcal{S} , we allow \mathcal{S} to execute \mathcal{A} internally. In the internal execution of \mathcal{A} , \mathcal{S} acts as the environment and the honest parties. In particular, \mathcal{S} forwards all messages between its environment and \mathcal{A} . The adversary \mathcal{A} creates messages for the corrupted parties. These messages are sent to \mathcal{S} in the internal execution. Note that this scenario also covers dummy adversaries, which just forward messages received from the environment. An output of \mathcal{S} indistinguishable from the output of \mathcal{A} in the real-world execution is created by simulating a protocol transcript towards \mathcal{A} that is indistinguishable

from the real-world execution and outputting whatever \mathcal{A} outputs in the simulated execution. Since the protocol π_{TBSB^+} is executed in the $(\mathcal{F}_{\text{Prep}}, \mathbf{H})$ -hybrid model, \mathcal{S} impersonates the hybrid functionality $\mathcal{F}_{\text{Prep}}$ and the global random oracle functionality \mathbf{H} in the internal execution.

We start with presenting our simulator \mathcal{S} .

Simulator \mathcal{S}

Oracle Queries.

Upon receiving (query, x) from P , do:

- If a tuple (x, a) is stored, then output (answer, a) to P .
- Else, sample $a \leftarrow \mathbb{Z}_p$ and store (x, a) .
Output (answer, a) to P .

KeyGen.

1. Upon receiving $(\text{init}, \text{sid})$ from corrupted party P_j , send $(\text{keygen}, \text{sid})$ on behalf of P_j to $\mathcal{F}_{\text{tsig}}$.
2. Upon receiving $(\text{pubkey}, \text{sid})$ from $\mathcal{F}_{\text{tsig}}$ simulate the initialization phase of $\mathcal{F}_{\text{Prep}}$ to get pk . In particular, sample $\text{sk} \xleftarrow{\$} \mathbb{Z}_p$ and send $\text{pk} = g_2^{\text{sk}}$ to \mathcal{A} .
3. Upon receiving $(\text{ok}, \text{Tuple}(\cdot, \cdot, \cdot))$ from \mathcal{A} , send $(\text{pubkey}, \text{sid}, \text{Verify}_{\text{pk}}(\cdot, \cdot))$ to $\mathcal{F}_{\text{tsig}}$.

Sign.

1. Upon receiving $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m} = \{m_\ell\}_{\ell \in [k]}, i)$ from $\mathcal{F}_{\text{tsig}}$ for honest party P_i , simulate the tuple phase of $\mathcal{F}_{\text{Prep}}$ to get (a_i, β_i) for P_i and simulate the random oracle \mathbf{H} call to get (answer, e) on input $(\text{query}, \text{sid} || \text{ssid} || \mathbf{m} || \mathcal{T})$. Then, compute $(A_i := (g_1 \cdot \prod_{\ell \in [k]} h_\ell^{m_\ell})^{a_i} \cdot h_0^{\alpha_i}, \delta_i := \beta_i + a_i \cdot e)$ and send it to the corrupted parties in \mathcal{T} in the internal execution.
2. Upon receiving $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m})$ from \mathcal{Z} to corrupted party P_j , send message to P_j in the internal execution and do:
 - (a) Upon receiving $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ on behalf of $\mathcal{F}_{\text{Prep}}$ from corrupted party P_j with $j \in \mathcal{T}$ return $(a_j, \beta_j) \leftarrow \text{Tuple}(\text{ssid}, \mathcal{T}, j)$ to P_j .
 - (b) Forward $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, j)$ to $\mathcal{F}_{\text{tsig}}$ and define an empty set $\widehat{\mathcal{L}}_j = \emptyset$ of honest parties that received signature shares from corrupted party P_j .
 - (c) Upon receiving $(\text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, A'_{j,i}, \beta'_{j,i})$ from P_j to honest party P_i in the internal execution, add P_i to $\widehat{\mathcal{L}}_j$.
3. Upon receiving $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m})$ from $\mathcal{F}_{\text{tsig}}$, do:
 - Use tuple (a_j, β_j) and random oracle output (answer, e) on input $(\text{query}, \text{sid} || \text{ssid} || \mathbf{m} || \mathcal{T})$ to compute honestly generated (A_j, β_j) for $P_j \in \mathcal{T} \cap \mathcal{C}$. Compute honestly generated signature $\sigma = (A, e)$ as honest parties do using (A_ℓ, β_ℓ) for $P_\ell \in \mathcal{T}$.
 - For each honest party P_i recompute signature σ_i obtained by P_i as honest parties do by using $A'_{j,i}, \beta'_{j,i}$ for $P_j \in \mathcal{T} \cap \mathcal{C}$.

- We define set \mathcal{I} of honest parties that obtained no or an invalid signature. First set, $\mathcal{I} = (\mathcal{T} \setminus \mathcal{C}) \setminus (\bigcap_{j \in \mathcal{T} \cap \mathcal{C}} \widehat{\mathcal{I}}_j)$, i.e., add all honest parties to \mathcal{I} that did not receive signature shares from all corrupted parties in \mathcal{T} . Next, compute $\mathcal{I} = \mathcal{I} \cup \{i : \sigma_i \neq \sigma\}$, i.e., add all honest parties that obtained a signature different to the honestly generated signature. If there exists $\sigma_i \neq \sigma$ such that $\text{Verify}_{\text{pk}}(\mathbf{m}, \sigma_i) = 1$ and $(\text{sig}, \text{sid}, \text{ssid}, \cdot, \mathbf{m}, \sigma_i, \cdot)$ was not sent to $\mathcal{F}_{\text{tsig}}$ before, output **fail** and stop the execution.
- Finally, send $(\text{sig}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, \sigma, \mathcal{I})$ to $\mathcal{F}_{\text{tsig}}$.

Verify. Upon receiving $(\text{verify}, \text{sid}, \mathbf{m}, \sigma, \text{Verify}_{\text{pk}'}(\cdot, \cdot))$ from $\mathcal{F}_{\text{tsig}}$ check if

- $\text{Verify}_{\text{pk}'}(\cdot, \cdot) = \text{Verify}_{\text{pk}}(\cdot, \cdot)$,
- $(\text{sig}, \text{sid}, \text{ssid}, \cdot, \mathbf{m}, \sigma, \cdot)$ was not sent to $\mathcal{F}_{\text{tsig}}$ before
- $\text{Verify}_{\text{pk}}(\mathbf{m}, \sigma) = 1$.

If the checks hold, output **fail** and stop the execution.

Lemma 1. *If simulator \mathcal{S} does not outputs **fail**, protocol π_{TBB^+} UC-realizes $\mathcal{F}_{\text{tsig}}$ in the $(\mathcal{F}_{\text{Prep}}, \text{H})$ -hybrid model in the presence of malicious adversaries controlling up to $t - 1$ parties.*

Proof. If the simulator \mathcal{S} does not outputs **fail**, it behaves precisely as the honest parties in real-world execution. Therefore, the simulation is perfect, and no environment can distinguish between the real and ideal worlds.

Lemma 2. *Assuming the strong unforgeability of BBS+, the probability that \mathcal{S} outputs **fail** is negligible.*

Proof. We show Lemma 2 via contradiction. Given a real-world adversary \mathcal{A} such that simulator \mathcal{S} outputs **fail** with non-negligible probability, we construct an attacker \mathcal{B} against the strong unforgeability (SUF) of BBS+ with non-negligible success probability. \mathcal{B} simulates the protocol execution towards \mathcal{A} like \mathcal{S} except the following aspects:

1. During the simulation of the initialization phase of $\mathcal{F}_{\text{Prep}}$, instead of sampling $\text{sk} \xleftarrow{\$} \mathbb{Z}_p$ and computing $\text{pk} = g_2^{\text{sk}}$, \mathcal{B} returns pk^* obtained from the SUF-challenger. Since the SUF-challenger samples the key exactly as the simulator \mathcal{S} , this step of the simulations is indistinguishable towards \mathcal{A} .
2. During the simulation of *oracle queries*, upon receiving $(\text{query}, x := \text{sid} || \text{ssid} || \mathbf{m} || \mathcal{T})$ from P and no tuple (x, a) is stored, do the following:
 - Request the signing oracle of the SUF-game on message \mathbf{m} to obtain signature $\sigma = (A, e)$. Store the tuple $(\text{sig} - \text{req}, \text{sid}, \text{ssid}, \mathbf{m}, \mathcal{T}, \sigma)$ and return (answer, e) to P .
3. During the *Sign* phase, upon receiving $(\text{sign}, \text{sid}, \text{ssid}, \mathcal{T}, \mathbf{m}, i)$ from $\mathcal{F}_{\text{tsig}}$ for honest party P_i , the computation of signature shares of the honest parties is modified as follows:

- If $(\mathbf{sig} - \mathbf{req}, \mathbf{sid}, \mathbf{ssid}, \mathbf{m}, \mathcal{T}, \sigma)$ is stored, forward σ to $\mathcal{F}_{\text{tsig}}$ on receiving $(\mathbf{sig}, \mathbf{sid}, \mathbf{ssid}, \mathcal{T}, \mathbf{m})$ from $\mathcal{F}_{\text{tsig}}$.
- If not tuple $(\mathbf{sig} - \mathbf{req}, \mathbf{sid}, \mathbf{ssid}, \mathbf{m}, \mathcal{T}, \cdot)$ is stored, request the signing oracle of the SUF-game and program the random oracle as in step 2. Forward the obtained signature to $\mathcal{F}_{\text{tsig}}$ on receiving $(\mathbf{sig}, \mathbf{sid}, \mathbf{ssid}, \mathcal{T}, \mathbf{m})$ from $\mathcal{F}_{\text{tsig}}$.
- Compute $(a_j, \beta_j) \leftarrow \text{Tuple}(\mathbf{ssid}, \mathcal{T}, j)$ and (A_j, δ_j) according to the protocol specification for every corrupted party $P_j \in \mathcal{T} \cap \mathcal{C}$.
- Sample random index $k \xleftarrow{\$} \mathcal{T} \setminus \mathcal{C}$.
- For all honest parties except P_k sample random signature share, i.e., $\forall P_i \in (\mathcal{T} \setminus \mathcal{C}) \setminus \{P_k\} : (A_i, \delta_i) \xleftarrow{\$} (\mathbb{G}_1, \mathbb{Z}_p)$.
- For P_k sample random $\delta_k \xleftarrow{\$} \mathbb{Z}_p$ and compute

$$A_k = \frac{A^{\sum_{\ell \in \mathcal{T}} \delta_\ell}}{\prod_{\ell \in \mathcal{T} \setminus \{k\}} A_\ell}.$$

Now, we must show two aspects. First, we show that the view of the adversary in the reduction is indistinguishable to the adversary's view in the ideal world. Second, we show that a successful adversary yields a valid attacker against the SUF game.

In order to show the first aspect, we argue that the changes in the reduction do not modify the distribution of the adversary's view. The first change is the sampling of the keys. Note that the public keys are both random elements in \mathbb{G}_2 and are identically distributed. In the reduction, \mathcal{B} does not know the secret key but the secret key is also never required later. Next, note that the value e is sampled at random by both, \mathcal{S} , and \mathcal{B} . This is due to the fact, that \mathcal{S} samples it random during the simulation of the random oracle queries, and \mathcal{B} obtains the value from the signing oracle of the SUF-game. In this game, the signing oracle samples the e value at random as well. During the signing phase, δ_i is the sum of β_i , which is a share of $a \cdot \text{sk}$, and $a_i \cdot e$. Although e is public, the a_i share, which is also part of a , works as a random mask. Therefore, δ_i is distributed randomly in the simulation by \mathcal{S} and by \mathcal{B} . Finally, the A_i values yield a valid signature in \mathcal{B} . Note that by programming the output of the oracle to the value e obtained by the signing oracle, we ensure that the output of the random oracle query on input $\mathbf{sid}||\mathbf{ssid}||\mathbf{m}||\mathcal{T}$ is consistent with the signature obtained for the signing request $(\mathbf{sig}, \mathbf{sid}, \mathbf{ssid}, \mathbf{m}, \mathcal{T})$. Therefore, the simulation of the *Sign* phase of \mathcal{B} and \mathcal{S} are indistinguishable to \mathcal{A} .

Next, we show that a successful adversary yields a valid attacker against the SUF game. Here, we use the fact that \mathcal{S} outputs **fail** with non-negligible probability either in the *Sign* or the *Verify* phase. As the interaction of \mathcal{B} with \mathcal{A} is indistinguishable, \mathcal{B} outputs **fail** with non-negligible probability as well. Whenever \mathcal{B} outputs **fail**, it forwards the pair (\mathbf{m}^*, σ^*) obtained in the *Sign* or *Verify* phase to the SUF-challenger.

It remains to show that \mathcal{B} successfully wins the SUF-game. In order to be a valid forger, it must hold that (1) $\text{Verify}_{\text{pk}^*}(\mathbf{m}^*, \sigma^*) = 1$ and (2) (\mathbf{m}^*, σ^*)

was not returned by the signing oracle before. (1) is trivially true, since \mathcal{B} only outputs **fail** if this condition holds. For (2), we note that \mathcal{A} has never seen σ^* as output from $\mathcal{F}_{\text{tsig}}$, since \mathcal{B} checks that $(\text{sig}, \text{sid}, \text{ssid}, \cdot, \mathbf{m}^*, \sigma^*, \cdot)$ was not sent to $\mathcal{F}_{\text{tsig}}$ before. However, it might happen that \mathcal{B} obtained σ^* as response to a signing request for message \mathbf{m}^* without forwarding it to $\mathcal{F}_{\text{tsig}}$ (this happens if the environment does not instruct all parties in \mathcal{T} to sign). Since the signing oracle samples e at random from \mathbb{Z}_p , the probability that σ^* was returned by the signing oracle is $\leq \frac{q}{p}$, where q is the number of oracle requests and p is the size of the field. While q is a polynomial, p is exponential in the security parameter. Thus, the probability that σ^* hits an unseen response from the signing oracle is negligible in the security parameter. It follows that (\mathbf{m}^*, σ^*) is a valid forgery and \mathcal{B} wins the SUF-game.

Since this contradicts the strong unforgeability of BBS+, it follows that the probability that \mathcal{S} outputs **fail** is negligible.

Combining Lemma 1 and Lemma 2 concludes the proof of Theorem 1.

K Simulator for PCG-based Preprocessing

Here, we state our simulator for proving security of our PCG-based preprocessing. Formally, the security is stated in Theorem 2. We provide a proof sketch of our indistinguishability argument in Appendix L and state the full proof in Appendix M.

Simulator for Preprocessing \mathcal{S}

Without loss of generality, we assume the adversary corrupts parties P_1, \dots, P_{t-1} and parties P_t, \dots, P_n are honest. \mathcal{S} internally uses adversary \mathcal{A} .

Initialization.

- 1: • Upon receiving $(\text{keygen}, \text{sid})$ on behalf of \mathcal{F}_{KG} from corrupted party P_j , send $(\text{init}, \text{sid})$ on behalf of corrupted P_j to $\mathcal{F}_{\text{Prep}}$. Then, wait to receive $(\text{corruptedShares}, \text{sid}, \{\text{sk}_j\}_{j \in \mathcal{C}})$ from \mathcal{A} .
- 2: • Upon receiving pk from $\mathcal{F}_{\text{Prep}}$, set $\text{pk}_j = g_2^{\text{sk}_j}$ for $j \in \mathcal{C}$ and compute $\text{pk}_i = \left(\text{pk} / (\text{pk}_1^{L_{1,\mathcal{T}}} \cdot \dots \cdot \text{pk}_{t-1}^{L_{t-1,\mathcal{T}}}) \right)^{1/L_{i,\mathcal{T}}}$, where $\mathcal{T} := \mathcal{C} \cup \{i\}$, for every honest party P_i . Then, send $(\text{sid}, \text{sk}_j, \text{pk}, \{\text{pk}_\ell\}_{\ell \in [n]})$ to every corrupted party P_j .
 - Upon receiving $(\text{setup}, \text{sid}, \rho_a^{(j)}, \text{sk}'_j, \{\text{pk}_\ell^{(j)}\}_{\ell \in [n]})$ on behalf of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ from every corrupted party P_j , check that $\text{pk}_\ell^{(j)} = \text{pk}_\ell$ and $g_2^{\text{sk}'_j} = \text{pk}_j$ for $j \in \mathcal{C}$ and $\ell \in [n]$. If any check fails, send $(\text{abort}, \text{sid})$ to $\mathcal{F}_{\text{Prep}}$. Otherwise sample $\rho_a^{(i)}$ and a dummy secret key share $\hat{\text{sk}}_i$ for every honest party P_i and simulate the computation of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ (i.e., compute all the PCG keys using the values received from the corrupted parties and the values sampled for the honest parties).

- 3: • Send keys $(\text{sid}, \{k_{j,\ell,0}^{\text{VOLE}}, k_{\ell,j,1}^{\text{VOLE}}\}_{\ell \neq j})$ to every corrupted party P_j .
 • Send $(\text{ok}, \text{Tuple}(\cdot, \cdot, \cdot))$ to $\mathcal{F}_{\text{Prep}}$, where $\text{Tuple}(\text{ssid}, \mathcal{T}, j)$ computes (a_j, β_j) for corrupted party P_j exactly as P_j computes its tuple in the protocol description.
 First, expand for every $\ell \in \mathcal{T} \setminus \{j\}$:

$$\begin{aligned} (\mathbf{a}_j, \mathbf{c}_{j,\ell,0}^{\text{VOLE}}) &= \text{PCG}_{\text{VOLE}}.\text{Expand}(0, k_{j,\ell,0}^{\text{VOLE}}), \\ (\text{sk}_j, \mathbf{c}_{\ell,j,1}^{\text{VOLE}}) &= \text{PCG}_{\text{VOLE}}.\text{Expand}(1, k_{\ell,j,1}^{\text{VOLE}}). \end{aligned}$$

Next, set $a_j = \mathbf{a}_j[\text{ssid}]$, $c_{(j,\ell,0)}^{\text{VOLE}} = \mathbf{c}_{(j,\ell,0)}^{\text{VOLE}}[\text{ssid}]$, and $c_{(\ell,j,1)}^{\text{VOLE}} = \mathbf{c}_{(\ell,j,1)}^{\text{VOLE}}[\text{ssid}]$ for $\ell \in \mathcal{T} \setminus \{j\}$ and compute

$$\beta_j = a_j \cdot L_{j,\mathcal{T}} \text{sk}_j + \sum_{\ell \in \mathcal{T} \setminus \{j\}} (L_{j,\mathcal{T}} c_{\ell,j,1}^{\text{VOLE}} - L_{\ell,\mathcal{T}} c_{j,\ell,0}^{\text{VOLE}}).$$

Tuple. Upon receiving $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ from \mathcal{Z} on behalf of corrupted party P_j , forward message $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ to \mathcal{A} and output whatever \mathcal{A} outputs.

L Indistinguishability Proof Sketch of Theorem 2

We prove indistinguishability between the ideal-world execution and the real-world execution via a sequence of hybrid experiments. We start with **Hybrid₀** which is the ideal-world execution and end up in **Hybrid₇** being identical to the real-world execution. By showing indistinguishability between each subsequent pair of hybrids, it follows that the ideal and real-world execution are indistinguishable. In particular, we show indistinguishability between the joint distribution of the adversary's view and the outputs of the honest parties in **Hybrid_i** and **Hybrid_{i+1}** for $i = 0, \dots, 6$. In the following we sketch the proof outline and defer the full proof to Appendix M.

Hybrid₁: In this hybrid experiment, we inline the description of the simulator \mathcal{S} , the ideal functionality $\mathcal{F}_{\text{Prep}}$ and the outputs of the honest parties. Since this is only a syntactical change, the distribution is identical to the one of **Hybrid₀**.

Hybrid₂: Instead of sampling the secret key sk at random from \mathbb{Z}_p , we sample a random polynomial $F(x) \in \mathbb{Z}_p[X]$ of degree $t-1$ such that $F(j) = \text{sk}_j$ for every $j \in \mathcal{C}$. The secret key is then defined as $\text{sk} = F(0)$.

Note that the adversary knows only $t-1$ shares of the polynomial which give no information about sk . This is due to the information-theoretically secrecy of Shamir's secret sharing. It follows that **Hybrid₁** and **Hybrid₂** are perfectly indistinguishable.

Hybrid₃: In this hybrid, we change the way honest parties' secret key shares are defined. Instead of sampling random dummy key shares, we derive the key shares from the polynomial introduced in the last hybrid. In more detail, the key share

of honest party P_i is computed as $\text{sk}_i = F(i)$. This change effects the PCG key generation as the dummy key share is replaced by sk_i for honest party P_i .

To show indistinguishability between Hybrid_2 and Hybrid_3 , we reduce to the key indistinguishability property of the PCG_{VOLE} primitive. More specifically, we introduce a sequence of intermediate hybrids where we only change the secret key of a single honest party in each step.

Hybrid₄: In this hybrid, we change the computation of the honest party P_i 's public key share pk_i . Instead of interpolating pk_i it is defined as $\text{pk}_i = g_2^{\text{sk}_i}$. As both ways are equivalent, Hybrid_4 is perfectly indistinguishable from Hybrid_3 .

Hybrid₅: In this hybrid, we make the sampling of the honest parties' outputs of the tuple phase explicit. To this end, we compute the tuple values in two steps. First, we sample values for a_i , then we compute β_i . For sampling, we distinguish between two cases. (1) For every pair of two honest parties (P_i, P_ℓ) the values are sampled from $\mathcal{Y}_{\text{VOLE}}$. (2) For every pair of one honest party P_i and one corrupted party P_j , we use the reverse-sampling algorithm of the VOLE correlation to compute the correlation outputs of the honest party. We illustrate the idea for a_i and β_i in the following.

After simulating the PCG key generation of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$, the experiment computes once and stores for every $i, \ell \in ([N] \setminus \mathcal{C})$ with $i \neq \ell$:

$$\begin{aligned} ((\mathbf{a}_i, \mathbf{c}_{i,\ell,0}^{\text{VOLE}}), \cdot) &\in \mathcal{Y}_{\text{VOLE}}(1^\lambda, (\rho_a^{(i)}, \text{sk}_\ell), [N]), \\ (\cdot, (\text{sk}_i, \mathbf{c}_{\ell,i,1}^{\text{VOLE}})) &\in \mathcal{Y}_{\text{VOLE}}(1^\lambda, (\rho_a^{(\ell)}, \text{sk}_i), [N]), \end{aligned}$$

and for every $i \in ([N] \setminus \mathcal{C}), j \in ([N] \cap \mathcal{C})$:

$$\begin{aligned} (\mathbf{a}_i, \mathbf{c}_{i,j,0}^{\text{VOLE}}) &\leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, (\rho_a^{(i)}, \text{sk}_j), 0, (\text{sk}_j, \mathbf{c}_{i,j,1}^{\text{VOLE}}), [N]) \\ (\text{sk}_i, \mathbf{c}_{j,i,1}^{\text{VOLE}}) &\leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, (\rho_a^{(j)}, \text{sk}_i), 1, (\mathbf{a}_j, \mathbf{c}_{j,i,0}^{\text{VOLE}}), [N]), \end{aligned}$$

where $(\mathbf{a}_j, \mathbf{c}_{j,i,0}^{\text{VOLE}}) = \text{PCG}_{\text{VOLE}}.\text{Expand}(0, \mathbf{k}_{j,i,0}^{\text{VOLE}})$ and $(\text{sk}_j, \mathbf{c}_{i,j,1}^{\text{VOLE}}) = \text{PCG}_{\text{VOLE}}.\text{Expand}(1, \mathbf{k}_{i,j,1}^{\text{VOLE}})$.

Then, during the tuple phase, for every $j \in \mathcal{T} \setminus \{i\}$ let $a_i = \mathbf{a}_i[\text{ssid}]$, $c_{i,j,0}^{\text{VOLE}} = \mathbf{c}_{i,j,0}^{\text{VOLE}}[\text{ssid}]$, and $c_{j,i,1}^{\text{VOLE}} = \mathbf{c}_{j,i,1}^{\text{VOLE}}[\text{ssid}]$ and compute

$$\beta_i = a_i \cdot L_{i,\mathcal{T}} \text{sk}_i + \sum_{\ell \in \mathcal{T} \setminus \{i\}} (L_{i,\mathcal{T}} c_{\ell,i,1}^{\text{VOLE}} - L_{\ell,\mathcal{T}} c_{i,\ell,0}^{\text{VOLE}}).$$

A straightforward calculation shows that resulting tuple values satisfy correlation (1). Note that the reverse-sampling and the correlation sampling outputs uniform correlation outputs and hence the correlation is identically distributed as in Hybrid_4 . It follows that the view of the environment is indistinguishable in Hybrid_4 and Hybrid_5 .

Hybrid₆: Now, we replace the sampling of correlation outputs for calculating honest parties' tuples (cf. case (1) of previous hybrid) with the expansion of the PCG keys, i.e., instead of using outputs of the $\mathcal{Y}_{\text{VOLE}}$ correlation, we run the PCG_{VOLE} expansions. For running the PCG expansions, we use the keys obtained during the simulation of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ in step (2).

Indistinguishability between Hybrid_5 and Hybrid_6 can be shown via reductions to the pseudorandom $\mathcal{V}_{\text{VOLE}}$ -correlated output property of the PCG_{VOLE} primitive. More precisely, a series of intermediate hybrids can be introduced, where in each hop only a single correlation output is replaced by the output of PCG expansions.

Hybrid₇: Finally, we replace the reverse-sampling in case (2) of Hybrid_5 with the PCG expansion. The indistinguishability between Hybrid_6 and Hybrid_7 can be shown via a reduction to the security property of the rPCGs.

Hybrid_7 is the real-world execution, which concludes the proof.

M Full Indistinguishability Proof of Theorem 2

In this section, we provide the full indistinguishability proof of Theorem 2. The simulator is given in Appendix K.

Hybrid₀: The initial experiment Hybrid_0 denotes the ideal-world execution where simulator \mathcal{S} is interacting with the corrupted parties, ideal functionality $\mathcal{F}_{\text{Prep}}$ and internally runs real-world adversary \mathcal{A} .

Hybrid₁: In this hybrid, we inline the description of the simulator \mathcal{S} , the ideal functionality $\mathcal{F}_{\text{Prep}}$ and the outputs of the honest parties. Since this is only a syntactical change, the joint distribution of the adversary's view and the output of the honest parties is identical to the one of Hybrid_0 . We state Hybrid_1 as the starting point, and emphasize only on the changes in the following hybrids.

Hybrid₁

Without loss of generality, we assume the adversary corrupts parties P_1, \dots, P_{t-1} and parties P_t, \dots, P_n are honest. \mathcal{S} internally uses adversary \mathcal{A} .

Initialization.

- 1: • Upon receiving $(\text{keygen}, \text{sid})$ on behalf of \mathcal{F}_{KG} from corrupted party P_j , store $(\text{init}, \text{sid}, P_j)$. Then, wait to receive $(\text{corruptedShares}, \text{sid}, \{\text{sk}_j\}_{j \in \mathcal{C}})$ from \mathcal{A} .
 - Upon receiving $(\text{init}, \text{sid})$ from every honest party, sample the secret key $\text{sk} \xleftarrow{\$} \mathbb{Z}_p$ and set $\text{pk} = g_2^{\text{sk}}$. Further, set $\text{pk}_j = g_2^{\text{sk}_j}$ for $j \in \mathcal{C}$ and compute $\text{pk}_i = \left(\text{pk} / (\text{pk}_1^{L_1, \tau} \cdot \dots \cdot \text{pk}_{t-1}^{L_1, \tau}) \right)^{1/L_i, \tau}$, where $\mathcal{T} := \mathcal{C} \cup \{i\}$, for every honest party P_i .
- 2: • Send $(\text{sid}, \text{sk}_j, \text{pk}, \{\text{pk}_\ell\}_{\ell \in [n]})$ to every corrupted party P_j .
 - Upon receiving $(\text{setup}, \text{sid}, \rho_a^{(j)}, \text{sk}'_j, \{\text{pk}_\ell^{(j)}\}_{\ell \in [n]})$ on behalf of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ from every corrupted party P_j , check that $\text{pk}_\ell^{(j)} = \text{pk}_\ell$ and $g_2^{\text{sk}'_j} = \text{pk}_j$ for $j \in \mathcal{C}$ and $\ell \in [n]$. If any check fails, honest parties output **abort**. Otherwise sample $\rho_a^{(i)}$ and a dummy secret key share $\widehat{\text{sk}}_i$ for every honest party P_i and simulate the computation of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ (i.e., compute all the PCG keys using the values received from the corrupted parties and the values sampled for the honest parties).

- 3: • Send keys $(\text{sid}, \{k_{j,\ell,0}^{\text{VOLE}}, k_{\ell,j,1}^{\text{VOLE}}\}_{\ell \neq j})$ to every corrupted party P_j .
 • Store $(\text{ok}, \text{Tuple}(\cdot, \cdot, \cdot))$, where $\text{Tuple}(\text{ssid}, \mathcal{T}, j)$ computes (a_j, β_j) for corrupted party P_j exactly as P_j computes its tuple in the protocol description.

First, expand for every $\ell \in \mathcal{T} \setminus \{j\}$:

$$\begin{aligned} (a_j, c_{j,\ell,0}^{\text{VOLE}}) &= \text{PCG}_{\text{VOLE}}.\text{Expand}(0, k_{j,\ell,0}^{\text{VOLE}}), \\ (\text{sk}_j, c_{\ell,j,1}^{\text{VOLE}}) &= \text{PCG}_{\text{VOLE}}.\text{Expand}(1, k_{\ell,j,1}^{\text{VOLE}}). \end{aligned}$$

Next, set $a_j = a_j[\text{ssid}]$, $c_{(j,\ell,0)}^{\text{VOLE}} = c_{(j,\ell,0)}^{\text{VOLE}}[\text{ssid}]$, and $c_{(\ell,j,1)}^{\text{VOLE}} = c_{(\ell,j,1)}^{\text{VOLE}}[\text{ssid}]$ for $\ell \in \mathcal{T} \setminus \{j\}$ and compute

$$\beta_j = a_j \cdot L_{j,\mathcal{T}} \text{sk}_j + \sum_{\ell \in \mathcal{T} \setminus \{j\}} (L_{j,\mathcal{T}} c_{\ell,j,1}^{\text{VOLE}} - L_{\ell,\mathcal{T}} c_{j,\ell,0}^{\text{VOLE}}).$$

- The honest parties P_t, \dots, P_n output pk .

Tuple.

- Upon receiving $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ from \mathcal{Z} on behalf of corrupted party P_j , forward message $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ to \mathcal{A} and output whatever \mathcal{A} outputs.
- Upon receiving $(\text{tuple}, \text{sid}, \text{ssid}, \mathcal{T})$ from \mathcal{Z} on behalf of honest party P_i , if $(\text{sid}, \text{ssid}, \mathcal{T}, \{(a_\ell, \beta_\ell)\}_{\ell \in \mathcal{T}})$ is stored, output $(\text{sid}, \text{ssid}, a_i, \beta_i)$. Otherwise, compute $(a_j, \beta_j) \leftarrow \text{Tuple}(\text{ssid}, \mathcal{T}, j)$ for every corrupted party P_j where $j \in \mathcal{C} \cap \mathcal{T}$ and sample $a \xleftarrow{\$} \mathbb{Z}_p$ and tuples (a_i, β_i) over \mathbb{Z}_p for $i \in \mathcal{H} \cap \mathcal{T}$ such that

$$\sum_{\ell \in \mathcal{T}} a_\ell = a \quad \sum_{\ell \in \mathcal{T}} \beta_\ell = a \cdot \text{sk}$$

Store $(\text{sid}, \text{ssid}, \mathcal{T}, \{(a_\ell, \beta_\ell)\}_{\ell \in \mathcal{T}})$ and honest party P_i outputs $(\text{sid}, \text{ssid}, a_i, \beta_i)$.

Hybrid₂: In this hybrid, we change the sampling of the secret key sk . Instead of sampling sk in step 1 from \mathbb{Z}_p , we sample a random polynomial $F \in \mathbb{Z}_p[X]$ of degree $t - 1$ such that $F(j) = \text{sk}_j$ for every $j \in \mathcal{C}$. Further, we define $\text{sk} = F(0)$. Since the polynomial is of degree $t - 1$, t evaluation points are required to fully determine $F(x)$. As the adversary knows only $t - 1$ shares, it cannot learn anything about sk . In detail, for every $\text{sk}' \in \mathbb{Z}_p$ there exists a t -th share that defined the polynomial $F(x)$ such that $F(x) = \text{sk}'$. It follows that the views of the adversary are distributed identically and hence **Hybrid₂** and **Hybrid₃** are perfectly indistinguishable.

Hybrid₃: Next, we use the polynomial $F(x)$ sampled in step 1 to determine the honest parties' secret key shares. In particular, for every honest party P_i the

experiment samples $\text{sk}_i = F(i)$. The secret key shares $\{\text{sk}_i\}_{i \in \mathcal{H}}$ are then used for the simulation of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$ instead of the dummy key shares. In particular, the correctly sampled key shares of the honest parties are used as input to $\text{PCG}_{\text{Vole}}.\text{Gen}$ whenever a secret key share of the honest party is used. Since the experiment does not use the dummy key shares at all after these changes, we remove them completely. Note that the sampling of the honest parties' key shares and the generation of the PCG keys are exactly as in the real-world execution.

Indistinguishability between Hybrid_2 and Hybrid_3 can be shown via a series of reductions to the key indistinguishability property of the VOLE PCG. We briefly sketch the proof outline in the following. We define intermediate hybrids $\text{Hybrid}_{2,\ell,k}$ for $\ell \in \{0, \dots, n - (t - 1)\}$ and $k \in [n]$, which only differ in the honest parties' key shares that are used in the generation of the VOLE PCG keys. Recall that for every party P_ℓ we generate a VOLE PCG for every other party P_k , where P_ℓ uses its secret key shares as input. We define $\text{Hybrid}_{2,\ell,k}$ such that the key shares derived from polynomial $F(x)$ are used for the first ℓ honest parties in all VOLE PCG instances and for the $(\ell + 1)$ -th honest party in the VOLE PCG instances with the first k other parties. For all other VOLE PCG instances, the dummy key shares are used for the honest parties' key shares.

Note that $\text{Hybrid}_{2,0,0} = \text{Hybrid}_2$ and $\text{Hybrid}_{2,n-(t-1),n} = \text{Hybrid}_3$. To show indistinguishability between $\text{Hybrid}_{2,\ell,k}$ and $\text{Hybrid}_{2,\ell,k+1}$ for every $\ell \in \{0, \dots, n - (t - 1)\}$, we make a reduction to the key indistinguishability property of the VOLE PCG. In particular, we construct an adversary $\mathcal{A}^{\text{key-ind}}$ from a distinguisher \mathcal{D}_ℓ which distinguishes between $\text{Hybrid}_{2,\ell,k}$ and $\text{Hybrid}_{2,\ell,k+1}$. Upon receiving the shares of the corrupted parties in the hybrid execution, $\mathcal{A}^{\text{key-ind}}$ forwards the key share of the $k + 1$ -th corrupted party to the security game. Then, the security game samples two possible key shares for the ℓ -th honest party $\rho_1^{(0)}, \rho_1^{(1)}$, uses one of them in the VOLE PCG key generation and sends the key k_1 for the corrupted party and $\rho_1^{(0)}$ to $\mathcal{A}^{\text{key-ind}}$. Next, $\mathcal{A}^{\text{key-ind}}$ continues the simulation of hybrid $\text{Hybrid}_{2,\ell,k}$ or $\text{Hybrid}_{2,\ell,k+1}$ by sampling the polynomial $F(x)$ using the corrupted key shares and $\rho_1^{(0)}$. Since $\rho_1^{(0)}$ is a random value in \mathbb{Z}_p , $F(x)$ is also a random polynomial. Finally, $\mathcal{A}^{\text{key-ind}}$ uses k_1 as the output of the simulation of $\mathcal{F}_{\text{Setup}}$.

If k_1 was sampled using $\rho_1^{(0)}$, then the simulated experiment is identical to $\text{Hybrid}_{2,\ell,k+1}$ and otherwise it is identical to $\text{Hybrid}_{2,\ell,k}$. It is easy to see that a successful distinguisher between these two hybrids allows to easily win the key indistinguishability game. Since we assume the VOLE PCG to satisfy the key indistinguishability property, this leads to a contradiction. Thus, the two hybrids are indistinguishable.

Hybrid₄: In this hybrid, we derive the honest parties public key shares pk_i from the secret key shares sk_i instead of interpolating them from pk and the corrupted shares. More precisely, in Hybrid_3 the public key share of honest party P_i was computed as

$$\text{pk}_i = \left(\text{pk} / (\text{pk}_1^{L_{1,\tau}} \cdot \dots \cdot \text{pk}_{t-1}^{L_{1,\tau}}) \right)^{1/L_{i,\tau}},$$

where $\mathcal{T} := \mathcal{C} \cup \{i\}$. In Hybrid_4 the public key share is instead computed as $\text{pk}_i = g_2^{\text{sk}_i}$. We show that both definitions are equivalent.

To this end, note that $\text{sk} = \sum_{\ell \in \mathcal{T}} L_{\ell, \mathcal{T}} \text{sk}_{\ell}$ for every set \mathcal{T} of size t , $\text{pk} = g_2^{\text{sk}}$ and $\text{pk}_j = g_2^{\text{sk}_j}$ for $j \in \mathcal{C}$. Using this equation we get for $\mathcal{T} = \mathcal{C} \cup \{i\}$

$$\begin{aligned}
\text{pk}_i &= \left(\frac{\text{pk}}{\text{pk}_1^{L_{1, \mathcal{T}}} \cdots \text{pk}_{t-1}^{L_{t-1, \mathcal{T}}}} \right)^{1/L_{i, \mathcal{T}}} \\
&\Leftrightarrow \text{pk}_i = \left(\frac{g_2^{\text{sk}}}{g_2^{L_{1, \mathcal{T}} \text{sk}_1} \cdots g_2^{L_{t-1, \mathcal{T}} \text{sk}_{t-1}}} \right)^{1/L_{i, \mathcal{T}}} \\
&\Leftrightarrow \text{pk}_i = \left(\frac{g_2^{\sum_{\ell \in \mathcal{T}} L_{\ell, \mathcal{T}} \text{sk}_{\ell}}}{g_2^{L_{1, \mathcal{T}} \text{sk}_1} \cdots g_2^{L_{t-1, \mathcal{T}} \text{sk}_{t-1}}} \right)^{1/L_{i, \mathcal{T}}} \\
&\Leftrightarrow \text{pk}_i = \left(g_2^{L_{i, \mathcal{T}} \text{sk}_i} \right)^{1/L_{i, \mathcal{T}}} \\
&\Leftrightarrow \text{pk}_i = g_2^{\text{sk}_i}
\end{aligned}$$

As public key shares are equivalent in both hybrids, the view of the adversary is identical distributed. Hence, Hybrid_3 and Hybrid_4 are perfectly indistinguishable.

Hybrid₅: In this hybrid, we derive the sampling the honest parties' outputs of the tuple phase from correlation samples and reverse sampling. To this end, we distinguish two cases. (1) For every pair of honest parties (P_i, P_{ℓ}) , the values are sampled from $\mathcal{Y}_{\text{VOLE}}$. (2) For every pair of one honest party P_i and one corrupted party P_j , we take the output of P_j 's PCG expansion and reverse-sample the output of the honest party. More specifically, after simulating the PCG key generation of $\mathcal{F}_{\text{Setup}}^{\text{PCG}}$, the experiment computes once and stores for every $i, \ell \in ([N] \setminus \mathcal{C})$ with $i \neq \ell$:

$$\begin{aligned}
((\mathbf{a}_i, \mathbf{c}_{i, \ell, 0}^{\text{VOLE}}), \cdot) &\in \mathcal{Y}_{\text{VOLE}}(1^{\lambda}, (\rho_a^{(i)}, \text{sk}_{\ell})), [N]), \\
(\cdot, (\text{sk}_i, \mathbf{c}_{\ell, i, 1}^{\text{VOLE}})) &\in \mathcal{Y}_{\text{VOLE}}(1^{\lambda}, (\rho_a^{(\ell)}, \text{sk}_i), [N]),
\end{aligned}$$

and for every $i \in ([N] \setminus \mathcal{C}), j \in ([N] \cap \mathcal{C})$:

$$\begin{aligned}
(\mathbf{a}_i, \mathbf{c}_{i, j, 0}^{\text{VOLE}}) &\leftarrow \text{RSample}_{\text{VOLE}}(1^{\lambda}, (\rho_a^{(i)}, \text{sk}_j), 0, (\text{sk}_j, \mathbf{c}_{j, i, 1}^{\text{VOLE}}), [N]) \\
(\text{sk}_i, \mathbf{c}_{j, i, 1}^{\text{VOLE}}) &\leftarrow \text{RSample}_{\text{VOLE}}(1^{\lambda}, (\rho_a^{(j)}, \text{sk}_i), 1, (\mathbf{a}_j, \mathbf{c}_{j, i, 0}^{\text{VOLE}}), [N]),
\end{aligned}$$

where

$$\begin{aligned}
(\mathbf{a}_j, \mathbf{c}_{j, i, 0}^{\text{VOLE}}) &= \text{PCG}_{\text{VOLE}}.\text{Expand}(0, \mathbf{k}_{j, i, 0}^{\text{VOLE}}), \\
(\text{sk}_j, \mathbf{c}_{i, j, 1}^{\text{VOLE}}) &= \text{PCG}_{\text{VOLE}}.\text{Expand}(1, \mathbf{k}_{i, j, 1}^{\text{VOLE}}),
\end{aligned}$$

Then, during the tuple phase, for every $j \in \mathcal{T} \setminus \{i\}$ let $a_i = \mathbf{a}_i[\text{ssid}]$, $c_{i,j,0}^{\text{VOLE}} = \mathbf{c}_{i,j,0}^{\text{VOLE}}[\text{ssid}]$, and $c_{j,i,1}^{\text{VOLE}} = \mathbf{c}_{j,i,1}^{\text{VOLE}}[\text{ssid}]$, and compute according to the protocol specification

$$\beta_i = a_i \cdot L_{i,\mathcal{T}} \mathbf{sk}_i + \sum_{\ell \in \mathcal{T} \setminus \{i\}} (L_{i,\mathcal{T}} c_{\ell,i,1}^{\text{VOLE}} - L_{\ell,\mathcal{T}} c_{i,\ell,0}^{\text{VOLE}}).$$

We show that the resulting tuple outputs satisfy the same correlation as before. In particular, we show $\sum_{\ell \in \mathcal{T}} \delta_\ell = a \cdot \mathbf{sk}$, where $a = \sum_{\ell \in \mathcal{T}} a_\ell$.

$$\begin{aligned} \sum_{\ell \in \mathcal{T}} \beta_\ell &= \sum_{\ell \in \mathcal{T}} \left(a_\ell \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell + \sum_{k \in \mathcal{T} \setminus \{\ell\}} L_{\ell,\mathcal{T}} c_{k,\ell,1}^{\text{VOLE}} - L_{k,\mathcal{T}} c_{\ell,k,0}^{\text{VOLE}} \right) \\ &= \sum_{\ell \in \mathcal{T}} a_\ell \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell + \sum_{\ell \in \mathcal{T}} \sum_{k \in \mathcal{T} \setminus \{\ell\}} L_{\ell,\mathcal{T}} c_{k,\ell,1}^{\text{VOLE}} - L_{\ell,\mathcal{T}} c_{k,\ell,0}^{\text{VOLE}} \\ &= \sum_{\ell \in \mathcal{T}} a_\ell \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell + \sum_{\ell \in \mathcal{T}} \sum_{k \in \mathcal{T} \setminus \{\ell\}} L_{\ell,\mathcal{T}} a_k \mathbf{sk}_\ell \\ &= \sum_{\ell \in \mathcal{T}} a_\ell \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell + \sum_{\ell \in \mathcal{T}} \sum_{k \in \mathcal{T} \setminus \{\ell\}} a_k \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell \\ &= \sum_{\ell \in \mathcal{T}} \sum_{k \in \mathcal{T}} a_k \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell \\ &= \sum_{k \in \mathcal{T}} \sum_{\ell \in \mathcal{T}} a_k \cdot L_{\ell,\mathcal{T}} \mathbf{sk}_\ell \\ &= \sum_{k \in \mathcal{T}} a_k \sum_{\ell \in \mathcal{T}} L_{\ell,\mathcal{T}} \mathbf{sk}_\ell \\ &= a \cdot \mathbf{sk} \end{aligned}$$

As the tuple values of the honest parties still satisfy the same correlation as in Hybrid_4 , Hybrid_4 and Hybrid_5 are indistinguishable. Note that the reverse-sampling and the correlation sampling outputs uniform correlation outputs and hence the correlation is identically distributed as in Hybrid_4 .

Hybrid₆: In this hybrid, we replace the correlation sampling of values of a pair of honest parties with PCG expansions (cf. case (1) of Hybrid_5). For example, instead of sampling $((\mathbf{a}_i, \mathbf{c}_{i,\ell,0}^{\text{VOLE}}), \cdot) \in \mathcal{Y}_{\text{VOLE}}(1^\lambda, (\rho_a^{(i)}, \mathbf{sk}_\ell), [N])$, party P_i computes $(\mathbf{a}_i, \mathbf{c}_{i,\ell,0}^{\text{VOLE}}) = \text{PCG}_{\text{VOLE}}.\text{Expand}(0, \mathbf{k}_{i,\ell,0}^{\text{VOLE}})$. The same change is applied to all VOLE correlations.

Indistinguishability can be shown via a series of reductions to the pseudorandom $\mathcal{Y}_{\text{VOLE}}$ -correlated output property of the PCGs. In more detail, we construct a sequence of hybrid experiments where only a single correlation sampling is replaced by a PCG expansion. Then, in the reduction to the pseudorandom correlated output property, in case the challenge bit is 0, the reduction simulates the hybrid where the output is sampled from the correlation, and in case the challenge bit is 1, the output is the PCG expansion. A distinguisher between

any pair of hybrid experiments in the sequence helps to construct a successful adversary against the pseudorandom correlated output property. We conclude that Hybrid_5 and Hybrid_6 are indistinguishable under the assumption of reusable PCGs.

Hybrid₇: Finally, we replace the reverse-sampling in case (2) of Hybrid_5 with the corresponding PCG expansion. For instance, instead of computing

$$(\mathbf{a}_i, \mathbf{c}_{i,j,0}^{\text{VOLE}}) \leftarrow \text{RSample}_{\text{VOLE}}(1^\lambda, (\rho_a^{(i)}, \text{sk}_j), 0, (\text{sk}_j, \mathbf{c}_{i,j,1}^{\text{VOLE}}), [N])$$

the honest party computes

$$(\mathbf{a}_i, \mathbf{c}_{i,j,0}^{\text{VOLE}}) = \text{PCG}_{\text{VOLE}}.\text{Expand}(0, \mathbf{k}_{i,j,0}^{\text{VOLE}}).$$

The same change is applied for all other reverse-sampling algorithms.

Analog to the indistinguishability between Hybrid_5 and Hybrid_6 , we can show indistinguishability between Hybrid_6 and Hybrid_7 via a sequence of hybrid experiments. In each hybrid one reverse sampling is replaced by the PCG expansion. Indistinguishability between adjacent hybrids is reduced to the security property of the PCG. Since the only change between two adjacent hybrids is the fact whether the correlation output of an honest party is reverse-sampled given the output of a corrupted party or taken as the PCG expansion, it is easy to see that a distinguisher between these hybrids can be used to construct a successful adversary against the security property.

We end up in Hybrid_7 where all correlation outputs and reverse-sampling outputs are replaced by PCG expansions. As this hybrid does not use any reverse-sampling anymore, we can get rid of the tuple function `Tuple`.

Now, Hybrid_7 is identical to the real-world execution which concludes the proof.

N Evaluation of the Legacy BBS+ Scheme

While we based our threshold BBS+ scheme on an optimized version of the BBS+ protocol due to Tessaro and Zhu [64], which removes the s variable from the signature, our scheme can be extended in a straight-forward way to thresholdize the legacy BBS+ signature scheme (cf. Appendix B). The major difference is that the legacy scheme requires OLE-correlations from the preprocessing. This chapter repeats our evaluation presented in Section 6 for the thresholdized legacy BBS+ signature scheme.

Online, signing request-dependent phase. For the online phase, we have implemented the thresholdized legacy BBS+ signature scheme and executed benchmarks. The scope of the implementation and the setup of our benchmarks remains unchanged. The results of our benchmarks are reported in Figure 11. The comparison to the legacy non-threshold BBS+ signature scheme is displayed in Figure 12. The size of signing requests does not change. The size of partial signatures sent by the servers increases to $(2\lceil \log p \rceil + |\mathbb{G}_1|)$.

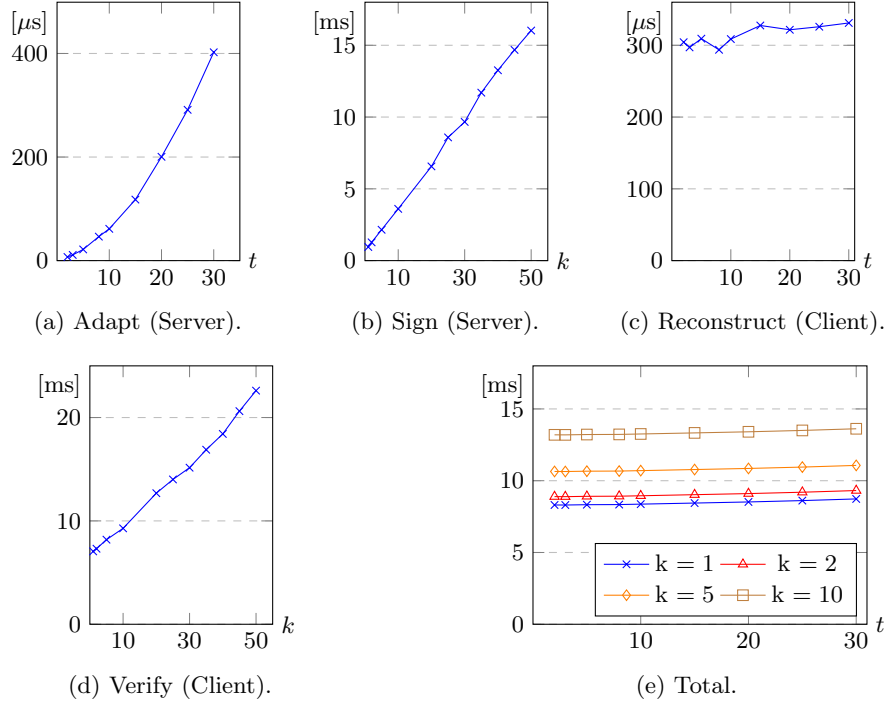


Fig. 11: The runtime of individual phases (a)-(d) and the total online protocol (e) in the thresholdized legacy BBS+ protocol. The *Adapt* phase and the *Reconstruct* phase depend on security threshold t . The *Sign* phase and the signature verification depend on the message array size k .




Fig. 12: The total runtime of the online phase of the PCG-based thresholdized legacy BBS+ protocol in comparison to the non-threshold legacy BBS+ signing algorithm with and without signature verification in dependence of the size of the message array k . As depicted in Figure 11e, the influence of the number of signers t is insignificant. We choose $t = 10$.

Offline, signing request-independent phase. For the offline phase, we extend our implementation and benchmarks by a PCG for OLE correlations. The results are displayed in Figure 13 and Figure 14. In the n -out-of- n setting, each party performs four randomization and splits four polynomials. In the t -out-of- n setting, each party performs $2 + 3 \cdot (n - 1)$ randomizations and splits just as many polynomials. The time to extract one of the N field elements from a degree- N polynomial remains unchanged.




Fig. 13: Computation time of the seed expansion of all required PCGs in the PCG-based thresholdized legacy BBS+ protocol in the n -out-of- n setting for different committee sizes ($n \in \{2, \dots, 10\}$) dependent on the number of generated precomputation tuples N .




Fig. 14: Computation time of the seed expansion of all required PCGs in the PCG-based thresholdized legacy BBS+ protocol in the t -out-of- n setting for different committee sizes ($n \in \{2, \dots, 10\}$) dependent on the number of generated precomputation tuples N .

Storage complexity When instantiating the precomputation with PCGs, the expanded precomputation material occupies

$$\log p \cdot (1 + N \cdot (2 + 4 \cdot (n - 1)))$$

bits of storage. We illustrate the complexity in Figure 15. In the n -out-of- n setting, the storage complexity reduces to $\log p \cdot 4N$ bits of storage.




Fig. 15: Storage complexity of the preprocessing material in the PCG-based thresholdized legacy BBS+ protocol required for $N \in \{98\,304, 1\,048\,576\}$ signatures depending on the number of servers n .

O Benchmarks of Basic Arithmetic Performance

We report the runtime of basic arithmetic operations in Table 1. The presented numbers might help the reader to assess the performance of system used for benchmarking and provides details for comparisons.

Table 1: Runtime of basic arithmetic operations in the BLS12_381 curve on our evaluation machine. The bit-size of the curve’s group order p is 255. The error terms report standard deviation.

Operation	Time
\mathbb{Z}_p addition	5.315 ns \pm 1.105 ns
\mathbb{Z}_p multiplication	34.311 ns \pm 1.752 ns
\mathbb{Z}_p inverse	2.907 μ s \pm 90.82 ns
\mathbb{G}_1 addition	1.18 μ s \pm 33.597 ns
\mathbb{G}_2 addition	3.971 μ s \pm 113.992 ns
\mathbb{G}_1 scalar multiplication	297.86 μ s \pm 13.59 μ s
\mathbb{G}_2 scalar multiplication	1.024 ms \pm 0.054 μ s
Pairing	2.638 ms \pm 57.258 μ s