

PERK

20/12/2025

SUBMITTERS (alphabetical order):

- Najwa AARAJ (Technology Innovation Institute, UAE)
- Slim BETTAIEB (Technology Innovation Institute, UAE)
- Loïc BIDOUX (Technology Innovation Institute, UAE)
- Alessandro BUDRONI (Technology Innovation Institute, UAE)
- Victor DYSERYN (XLIM, University of Limoges, FR)
- Andre ESSER (Technology Innovation Institute, UAE)
- Thibault FENEUIL (Cryptoexperts, FR)
- Philippe GABORIT (XLIM, University of Limoges, FR)
- Mukul KULKARNI (Technology Innovation Institute, UAE)
- Victor MATEU (Technology Innovation Institute, UAE)
- Marco PALUMBI (Technology Innovation Institute, UAE)
- Lucas PERIN (Technology Innovation Institute, UAE)
- Matthieu RIVAIN (Cryptoexperts, FR)
- Jean-Pierre TILLICH (INRIA, FR)
- Keita XAGAWA (Technology Innovation Institute, UAE)

CONTACT: team@pqc-perk.org

VERSION: 2.2.0

Changelog

Version 2.2.0 (20/12/2025)

- Fix several bugs in the implementation.

Version 2.1.1 (07/10/2025)

- Fix typographical errors.

Version 2.1.0 (23/09/2025)

- The design of PERK has been improved and now relies on the modeling from [BBGK24] along with the VOLEitH framework [BBD⁺23c]. As a result, PERK signature sizes have been significantly reduced.

Version 2.0.0 (05/02/2025)

- Thibault Feneuil, Matthieu Rivain and Keita Xagawa have joined the PERK team.

Version 1.1.0 (16/10/2023)

- Reduce signature sizes for short parameters set by approximately 5% using a ranking algorithm for permutation encoding ;
- Improve the implementation (reduced stack-memory usage and bug fixing).

Table of Contents

1	Introduction	4
2	Preliminaries	5
2.1	Notations	5
2.2	Standard cryptographic primitives	5
2.3	Digital signature schemes	7
2.4	Permuted Kernel Problem	8
3	Overview of PERK	10
3.1	VOLE-in-the-Head framework	10
3.2	Proof of Knowledge for PKP	12
4	Algorithmic Description	14
4.1	Object representation	14
4.2	Sampling functions	17
4.3	Hash functions and commitments	18
4.4	VOLE-in-the-Head functions	20
4.5	Proof of Knowledge functions	36
4.6	PERK	63
5	Parameter Sets and Sizes	66
5.1	PKP parameters	66
5.2	MPC and VOLE parameters	66
5.3	Signature and key sizes	68
6	Implementation and Performance Analysis	69
6.1	Reference implementation	69
6.2	Optimized implementation	70
6.3	Known Answer Test values	72
7	Security Analysis	73
7.1	Security proof	73
7.2	Known attacks against PKP	88
8	Advantages and Limitations	91
8.1	Advantages	91
8.2	Limitations	91

1 Introduction

PERK is a post-quantum digital signature scheme based on the hardness of the PERmuted Kernel Problem (PKP) [Sha90]. The scheme builds on a Zero-Knowledge Proof of Knowledge (ZK PoK) of a PKP solution computed using the modeling from [BBGK24] along with the Multi-Party Computation in the Head (MPCitH) paradigm [IKOS07]. More precisely, PERK relies on the VOLE in the Head (VOLEitH) framework [BBD⁺23c]. The ZK PoK is then converted into a signature scheme using the Fiat-Shamir transform [FS87].

Organization. We present in Section 2 some background and the notations we will use. Then, Section 3 and Section 4 respectively provide an overview and a detailed algorithmic description of PERK. Section 5 and Section 6 are dedicated to the parameters and the performances of PERK. A security analysis of the scheme is provided in Section 7. Finally, in Section 8, we summarize the main advantages and limitations of the scheme.

2 Preliminaries

2.1 Notations

Let λ denote the security parameter. For integers a, b we denote $[a, b]$ the set of integers i such that $a \leq i \leq b$. We write $[n]$ as a shorthand for $[0, n - 1]$. We denote \mathcal{S}_n the group of permutations of the set $[n]$. Let \mathbb{F}_q denote the finite field of q elements where q is the power of a prime. If S is a finite set, we denote by $x \xleftarrow{\$} S$ that x is chosen uniformly at random from S . Similarly, we write $x \xleftarrow{\$, \theta} S$, if x is sampled pseudo-randomly from the set S , based on the seed θ . We use x to denote input and denote its length by $|x|$. Vectors are denoted by bold lower-case letters and matrices by bold capital letters (e.g., $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{F}_q^n$ and $\mathbf{M} = (m_{ij})_{1 \leq i \leq k, 1 \leq j \leq n} \in \mathbb{F}_q^{k \times n}$). We denote by $\ker(\mathbf{M})$ the right kernel of the matrix \mathbf{M} .

We call a function $f : \mathbb{N} \rightarrow \mathbb{R}^+$ *negligible*, if for all $c \in \mathbb{N}$ there exists a $N_0 \in \mathbb{N}$ such that $f(n) < 1/n^c$ for all $n > N_0$. We write $\text{negl}(\lambda)$ to denote an arbitrary negligible function. We use $\text{poly}(\lambda)$ for function which is *polynomially bounded* in λ , that is there exists $c, \lambda_0 \in \mathbb{N}$ such that $\text{poly}(\lambda) \leq \lambda^c$ for all $\lambda \geq \lambda_0$. We also abbreviate *probabilistic polynomial-time* as PPT. Let X and Y be two discrete random variables defined over a finite support D . The *statistical distance* between the two distributions is defined as

$$\Delta(X, Y) := \frac{1}{2} \sum_{d \in D} |\Pr[X = d] - \Pr[Y = d]|.$$

We say two ensembles of random variables $\{X_\lambda\}_{\lambda \in \mathbb{N}}, \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are *statistically close* if there exists a negligible function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}^+$ such that $\Delta(X_\lambda, Y_\lambda) \leq \text{negl}(\lambda)$ for all $\lambda \in \mathbb{N}$. We say two ensembles of random variables $\{X_x\}_{x \in \{0,1\}^*}, \{Y_x\}_{x \in \{0,1\}^*}$ are *statistically close* if there exists a negligible function $\text{negl} : \mathbb{N} \rightarrow \mathbb{R}^+$ such that $\Delta(X_x, Y_x) \leq \text{negl}(|x|)$ for all $x \in \{0,1\}^*$.

2.2 Standard cryptographic primitives

Definition 2.1 (Salt based PRG (adapted from [BBD⁺23b])). Let $\text{PRG} : \{0,1\}^{2\lambda} \times \{0,1\}^\lambda \rightarrow \{0,1\}^\ell$ be a deterministic polynomial-time algorithm. Let \mathcal{A} be a q -query, N -batch adversary. We define an experiment $\text{Expt}_{\text{PRG}, \mathcal{A}}^{\text{mt-prg}}(q, N, \lambda)$ as in Figure 1. Let

$$\text{AdvPRG}^{\text{PRG}}[q, N, \lambda] := \left| \Pr \left[\text{out}_{\mathcal{A}} = 1 : \text{out}_{\mathcal{A}} \leftarrow \text{Expt}_{\text{PRG}, \mathcal{A}}^{\text{mt-prg}}(q, N, \lambda) \right] - \frac{1}{2} \right|.$$

Then we call PRG as $((q, N, \lambda), (t, \varepsilon))$ -secure if for every adversary \mathcal{A} running in time t , its advantage $\text{AdvPRG}^{\text{PRG}}[q, N, \lambda]$ is at most ε . If λ is obvious in the context, we will drop λ .

$$\text{Expt}_{\text{PRG}, \mathcal{A}}^{\text{mt-prg}}(q, N, \lambda)$$

```

1 :  $b \xleftarrow{\$} \{0, 1\}$ 
2 : for  $i \in [q]$  :  $\text{salt}_i \xleftarrow{\$} \{0, 1\}^{2\lambda}$ 
3 :   for  $j \in [N]$  :
4 :     if  $b = 0$  :
5 :        $\text{seed}_j \xleftarrow{\$} \{0, 1\}^\lambda$ 
6 :        $r_{i,j} := \text{PRG}(\text{salt}_i, \text{seed}_j)$ 
7 :     else :  $r_{i,j} \xleftarrow{\$} \{0, 1\}^{\hat{\ell}}$ 
8 :    $b' \leftarrow \mathcal{A}(\{\text{salt}_i\}_{i \in [q]}, \{r_{i,j}\}_{(i,j) \in [q] \times [N]})$ 
9 :   if  $b = b'$  : return 1
10 : else : return 0

```

Fig. 1: Multi-challenge security of salt-based PRG

Note that, any $((1, 1, \lambda), (t, \varepsilon_{\text{PRG}}))$ -secure salt based PRG is also $((q, N, \lambda), (t, \varepsilon'_{\text{PRG}}))$ -secure with $\varepsilon'_{\text{PRG}} \leq q \cdot N \cdot \varepsilon_{\text{PRG}}$.

Definition 2.2 (Collision Resistance). *Let $H : \{0, 1\}^{\ell_{\text{in}}} \rightarrow \{0, 1\}^{\ell_{\text{out}}}$ be a deterministic function. For any adversary making at most q queries to H , we denote its advantage in finding a collision for H by $\text{AdvColl}^H[q]$. We say that H is $(q, (t, \varepsilon))$ -secure if for every adversary \mathcal{A} running in time t , its advantage $\text{AdvColl}^H[q]$ is at most ε .*

Definition 2.3 (Multi-Target Non-Invertibility). *Let $\text{Com} : \{0, 1\}^{\ell_{\text{in}}} \rightarrow \{0, 1\}^{\ell_{\text{out}}}$ be a deterministic function. For a q -query adversary, we define its advantage as*

$$\text{AdvNI}^{\text{Com}}[q] := \Pr \left[\text{Com}(\text{inp}) = \text{com}_i : \begin{array}{l} \text{com}_0, \dots, \text{com}_{q-1} \xleftarrow{\$} \{0, 1\}^{\ell_{\text{out}}} \\ (i, \text{inp}) \leftarrow \mathcal{A}(\text{com}_0, \dots, \text{com}_{q-1}) \end{array} \right].$$

We say that Com is $(q, (t, \varepsilon))$ -secure if for every adversary \mathcal{A} running in time t , its advantage $\text{AdvNI}^{\text{Com}}[q]$ is at most ε .

Definition 2.4 (One-time PRF). *Let $F : \{0, 1\}^{2\lambda} \times \{0, 1\}^* \rightarrow \{0, 1\}^{3\lambda}$ be a deterministic polynomial-time algorithm. Let an adversary \mathcal{A} , we define its advantage as*

$$\text{AdvPRF}^F := \left| \Pr \left[b = b' : \begin{array}{l} b \xleftarrow{\$} \{0, 1\}; \text{rand} \xleftarrow{\$} \{0, 1\}^{2\lambda}; (\text{state}, \text{inp}) \leftarrow \mathcal{A}() \\ r_0 := F(\text{rand}, \text{inp}); r_1 \xleftarrow{\$} \{0, 1\}^{3\lambda}; b' \leftarrow \mathcal{A}(r_b, \text{state}) \end{array} \right] - \frac{1}{2} \right|.$$

We say that F is (t, ε) -secure if for every adversary \mathcal{A} running in time t , its advantage AdvPRF^F is at most ε .

Definition 2.5 (Joint PRF Security). Let $\text{PRG} : \{0, 1\}^{2\lambda} \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^{\hat{\ell}}$ and $\text{Com} : \{0, 1\}^* \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$ be deterministic polynomial-time algorithms. Let \mathcal{A} be a q -query, N -batch adversary in the experiment from Figure 2. Let

$$\text{AdvJPRF}^{\text{PRG}, \text{Com}}[q, N, \lambda] := \left| \Pr \left[\text{out}_{\mathcal{A}} = 1 : \text{out}_{\mathcal{A}} \leftarrow \text{Expt}_{\text{PRG}, \text{Com}, \mathcal{A}}^{\text{mt-jprf}}(q, N, \lambda) \right] - \frac{1}{2} \right|.$$

Then we call a pair of PRG and Com as $((q, N, \lambda), (t, \varepsilon))$ -secure if for every adversary \mathcal{A} running in time t , its advantage $\text{AdvJPRF}^{\text{PRG}, \text{Com}}[q, N, \lambda]$ is at most ε . If λ is obvious in the context, we will drop λ .

```

ExptPRG, Com, Amt-jprf(q, N, λ)
-----
1: b ←$ {0, 1}
2: for i ∈ [q] : salti ←$ {0, 1}2λ
3: ( {inpi,j }(i,j) ∈ [q] × [N], state) ← A ( {salti }i ∈ [q] )
4: for i ∈ [q] :
5:   for j ∈ [N] :
6:     if b = 0 :
7:       seedj ←$ {0, 1}λ
8:       ri,j := PRG(salti, seedj)
9:       comi,j := Com(inpi,j, seedj)
10:    else :
11:      ri,j ←$ {0, 1}ℓ̂
12:      comi,j ←$ {0, 1}2λ
13: b' ← A (state, { (ri,j, comi,j) }(i,j) ∈ [q] × [N])
14: if b = b' : return 1
15: else : return 0

```

Fig. 2: Multi-challenge joint PRF security of commitment and PRG

2.3 Digital signature schemes

Definition 2.6 (Signature Scheme). A signature scheme consists of three PPT algorithms $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ which work as follows:

- $\text{KeyGen}(1^\lambda)$: The key generation algorithm takes a security parameter as input and outputs a pair of keys (pk, sk) . The key sk is the private (secret) signing key and pk is the public key used for verification.

$\text{Expt}_{\text{SIG}, \mathcal{A}}^{\text{euf-cma}}(\lambda)$	$\text{Expt}_{\text{SIG}, \mathcal{A}}^{\text{suf-cma}}(\lambda)$	$\text{OSign}(\text{msg})$
1 : $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$	1 : $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda)$	1 : $\sigma \leftarrow \text{Sign}(\text{sk}, \text{msg})$
2 : $\mathcal{Q} := \emptyset$	2 : $\mathcal{Q} := \emptyset$	2 : $\mathcal{Q} := \mathcal{Q} \cup \{(\text{msg}, \sigma)\}$
3 : $(\text{msg}^*, \sigma^*) \leftarrow \mathcal{A}^{\text{OSign}(\cdot)}(\text{vk})$	3 : $(\text{msg}^*, \sigma^*) \leftarrow \mathcal{A}^{\text{OSign}(\cdot)}(\text{vk})$	3 : return σ
4 : $d_1 = \text{Verify}(\text{pk}, \text{msg}^*, \sigma^*)$	4 : $d_1 = \text{Verify}(\text{pk}, \text{msg}^*, \sigma^*)$	
5 : $d_2 = ((\text{msg}^*, \cdot) \notin \mathcal{Q})$	5 : $d_2 = ((\text{msg}^*, \sigma^*) \notin \mathcal{Q})$	
6 : return $d_1 \wedge d_2$	6 : return $d_1 \wedge d_2$	

Fig. 3: EUF-CMA and SUF-CMA games.

- $\text{Sign}(\text{sk}, \text{msg})$: The signing algorithm takes as input a secret signing key sk and a message msg from some message space (that may depend on pk). It outputs a signature $\sigma \leftarrow \text{Sign}(\text{sk}, \text{msg})$.
- $\text{Verify}(\text{pk}, \text{msg}, \sigma)$: The deterministic verification algorithm takes as input a public key pk , a message msg , and a signature σ . It outputs a bit $b := \text{Verify}(\text{pk}, \text{msg}, \sigma)$, with $b = 1$ meaning the signature-message pair is valid and $b = 0$ meaning it is invalid.

Definition 2.7 (EUF-CMA Security and EUF-NMA Security). A signature scheme $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ is existentially unforgeable under chosen-message attacks (EUF-CMA secure) if, for all PPT adversaries \mathcal{A} there is a negligible function $\text{negl}(\cdot)$ such that,

$$\Pr[\text{Expt}_{\text{SIG}, \mathcal{A}}^{\text{euf-cma}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where $\text{Expt}_{\text{SIG}, \mathcal{A}}^{\text{euf-cma}}(\lambda)$ is the security game defined in Figure 3. In addition, if we consider the game where the signing oracle OSign is removed, then the signature scheme is said to be existentially unforgeable under no-message attacks (EUF-NMA secure).

Definition 2.8 (SUF-CMA Security). A signature scheme $\text{SIG} = (\text{KeyGen}, \text{Sign}, \text{Verify})$ is strongly existentially unforgeable under chosen-message attacks (SUF-CMA secure) if, for all PPT adversaries \mathcal{A} there is a negligible function $\text{negl}(\cdot)$ such that,

$$\Pr[\text{Expt}_{\text{SIG}, \mathcal{A}}^{\text{suf-cma}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where $\text{Expt}_{\text{SIG}, \mathcal{A}}^{\text{suf-cma}}(\lambda)$ is the security game defined in Figure 3

2.4 Permuted Kernel Problem

PERK's security relies on the permuted kernel problem (PKP) introduced by Shamir in [Sha90].

Definition 2.9 (Permuted Kernel Problem (PKP)). Let (q, m, n) be positive integers such that $m < n$, $\mathbf{H} \in \mathbb{F}_q^{m \times n}$, $\mathbf{x} \in \mathbb{F}_q^n$ and $\pi \in \mathcal{S}_n$ be a permutation such that $\mathbf{H}(\pi[\mathbf{x}]) = \mathbf{0}$. Given (\mathbf{H}, \mathbf{x}) , the Permuted Kernel Problem $\text{PKP}(q, m, n)$ asks to find $\tilde{\pi} \in \mathcal{S}_n$ such that $\mathbf{H}(\tilde{\pi}[\mathbf{x}]) = \mathbf{0}$.

By convention, we call this problem the PKP problem. Hereafter, we interpret the PKP problem in matrix form namely the secret permutation π is seen as a permutation matrix $\mathbf{P} \in \mathbb{F}_q^{n \times n}$ such that $\mathbf{H}\mathbf{P}\mathbf{x} = \mathbf{0}$. In addition, our working field \mathbb{F}_q is an extension field of \mathbb{F}_2 .

The formal definition of the PKP assumption requires an instance distribution. To shorten the instance, we will employ `ExpandMatrixM`, a deterministic function $\{0, 1\}^\lambda \rightarrow \mathbb{F}_q^{m \times (n-m)}$.

Definition 2.10 (Advantage against PKP). Let (q, m, n) be positive integers such that $m < n$ and q is a power of a prime. For an adversary \mathcal{A} , we define its advantage $\text{AdvOW}_{\mathcal{A}}$ against the PKP problem as follows:

$$\text{AdvOW}_{\mathcal{A}} := \Pr \left[\begin{array}{l} \mathbf{H}_{\text{seed}} \xleftarrow{\$} \{0, 1\}^\lambda; \mathbf{M} := \text{ExpandMatrixM}(\mathbf{H}_{\text{seed}}); \\ \mathbf{H}(\tilde{\pi}[\mathbf{x}]) = \mathbf{0} : \mathbf{H} := [\mathbf{I}_m \ \mathbf{M}]; \mathbf{x}' \xleftarrow{\$} \ker(\mathbf{H}); \pi \xleftarrow{\$} \mathcal{S}_n; \\ \mathbf{x} := \pi^{-1}[\mathbf{x}']; \tilde{\pi} \leftarrow \mathcal{A}(\mathbf{H}_{\text{seed}}, \mathbf{x}) \end{array} \right].$$

We say that the PKP assumption holds if for any polynomial-time adversary \mathcal{A} , its advantage $\text{AdvOW}_{\mathcal{A}}$ is negligible in λ .

3 Overview of PERK

3.1 VOLE-in-the-Head framework

VOLE correlations. A VOLE (Vector Oblivious Linear Evaluation) correlation of length $\hat{\ell}$ over \mathbb{F}_{2^ρ} is defined by random values $(\mathbf{u}, \mathbf{v}) \in \mathbb{F}_2^{\hat{\ell}} \times \mathbb{F}_{2^\rho}^{\hat{\ell}}$, and $(\mathbf{q}, \Delta) \in \mathbb{F}_{2^\rho}^{\hat{\ell}} \times \mathbb{F}_{2^\rho}$, such that

$$q_i = u_i \Delta + v_i \quad i \in [0, 1, \dots, \hat{\ell} - 1].$$

The VOLE correlation serves as an information theoretically secure commitment to prover's random value \mathbf{u} . The mask \mathbf{v} is unknown to the verifier, this provide the hiding property, while the prover needs to guess Δ in order to open the commitment to some $\mathbf{u}' \neq \mathbf{u}$, which provides the binding property. Moreover, owing to the linearity of VOLE correlations, these commitments are linearly homomorphic. Therefore, such VOLE correlations can be used to build efficient zero-knowledge proofs of knowledge, where the prover can commit to its secret witness with help of VOLE correlations and convince the verifier by computing some public function on the witness (and other public values) which can be verified using only the verifier's VOLE correlation inputs (\mathbf{q}, Δ) .

VOLE-in-the-Head. Following the approach of [BBD+23b, BBD+23c], in order to achieve the public verifiability for our zero-knowledge proofs (and signatures) we use the *VOLE-in-the-Head* (VOLEitH) technique. In this approach the prover generates the values \mathbf{u}, \mathbf{v} and commits to these values. The prover then computes the desired public relation with the help of committed VOLE correlation inputs (\mathbf{u}, \mathbf{v}) . At this point the verifier can send Δ to the prover, and prover can send opening to the commitments to (\mathbf{u}, \mathbf{v}) , from which the verifier can compute \mathbf{q} without learning any extra information. Note that it is important that the prover learns the value of Δ (required to provide openings) only after it has committed to the VOLE correlation inputs, and to the computations of the zero-knowledge protocol (so it cannot change these after learning Δ), since after the prover knows Δ , the binding property of the linear homomorphic commitments does not hold.

In practice the VOLE correlation values are computed from uniform random strings of length $\hat{\ell}$. In order to create a single instance of VOLE correlation inputs $(\mathbf{u}, \mathbf{v}) \in \mathbb{F}_2^{\hat{\ell}} \times \mathbb{F}_{2^\rho}^{\hat{\ell}}$ the prover (signer) essentially needs to perform $O(2^\rho)$ additions and multiplications. Similarly, after receiving the opening the verifier also needs perform similar computation to acquire \mathbf{q} . The soundness error of the zero-knowledge proof of knowledge (relying on the binding property of linear homomorphic commitment based on VOLE correlations) is $2^{-\rho}$. Therefore, to achieve the desired security level we need to set $\rho \geq \lambda$, however this means that the prover and verifier will need to perform infeasible computations in order to even get started by creating the VOLE correlation inputs. This is mitigated by creating several parallel instances of VOLE correlations in a smaller field \mathbb{F}_{2^μ} and

concatenating them together to produce a single VOLE correlation instance in exponentially large field \mathbb{F}_{2^ρ} . This allows us to compute the VOLE correlations required to achieve desired security level efficiently. For each parameter set, we choose a repetition parameter, $\tau \in \mathbb{N}$ along with a VOLE field parameter, $\mu \in \mathbb{N}$ such that $\rho = \tau\mu$.

Committing to VOLE correlations. An important step in our signature scheme is to commit a vector of pseudo-random seeds, and be able to later open all-but-one of those seeds. Looking ahead these seeds will be used to generate the aforementioned VOLE correlations. The standard approach to build such an efficient commitment scheme is to derive the seeds from a tree of length-doubling PRGs. Such a construction is called an *all-but-one vector commitment scheme*, relying on a GGM tree [GGM84], as suggested in [KKW18]. Suppose a party needs to generate N seeds and then to reveal only $N - 1$ of those seeds (without knowing in advance which seed should not be revealed). The principle is to build a binary tree of depth $\lceil \log_2(N) \rceil$. The root of the tree is labeled with a master seed θ . The rest of the tree is labeled inductively by using a PRG of double extension on each parent node and splitting the output on the left and right children. To reveal all seeds except seed number i for $0 \leq i \leq N - 1$, the principle is to reveal the labels on the siblings of the paths from the root of the tree to leave i . It allows to reconstruct all seeds but seed number i at the cost of communicating $\lceil \log_2(N) \rceil$ labels, which is more effective than communicating $N - 1$ seeds. Instead of building one GGM-tree for each of the τ repetitions of the proof of knowledge, we adopted the approach of using a unified tree, as explained in [BBM⁺25], to save communication costs.

Computing VOLE correlations from seeds. The seeds obtained from the GGM tree can be used to generate VOLE correlations as follows: Let $\mathbf{r}_i = \text{PRG}(\text{seed}_i)$ be $\hat{\ell}$ -bit pseudorandom strings for an integer $i \in [0, N - 1]$ with $N := 2^\kappa$ as the number of leaves in the GGM tree. The prover computes \mathbf{u}, \mathbf{v} as

$$\mathbf{u} = \sum_{i=0}^{N-1} \mathbf{r}_i, \quad \mathbf{v} = \sum_{i=0}^{N-1} i \cdot \mathbf{r}_i.$$

with i (and thus \mathbf{v}) as an element in \mathbb{F}_{2^μ} . The verifier chooses $\Delta \in [0, N - 1] \subseteq \mathbb{F}_{2^\mu}$ uniformly at random and receives all the seeds seed_i for $i \in [0, N - 1] \setminus \Delta$ from the prover. The verifier can then compute

$$\mathbf{q} = \sum_{i=0}^{N-1} (\Delta - i) \cdot \mathbf{r}_i \in \mathbb{F}_{2^\mu}$$

We use the same approach as [BBD⁺23b] to compute the vole correlations from seeds which is detailed in Section 4.4 and algorithm 4.4. In order to achieve the desired security level, we need to repeat the above procedure τ times. However, this results in τ independent instances of VOLE correlations

$$\mathbf{q}_e = \mathbf{u}_e \Delta_e + \mathbf{v}_e \quad e \in [0, 1, \dots, \tau - 1]$$

with $(\mathbf{u}_e, \mathbf{v}_e) \in \mathbb{F}_2^{\hat{\ell}} \times \mathbb{F}_2^{\hat{\mu}}$ and $(\mathbf{q}_e, \Delta_e) \in \mathbb{F}_2^{\hat{\ell}} \times \mathbb{F}_2^{\mu}$. Let $\mathbf{u}_e \in \mathbb{F}_2^{\hat{\ell}}$ be represented as a vector of length $\hat{\ell}$ over \mathbb{F}_2 . Similarly elements in \mathbb{F}_2^{μ} such Δ_e can be represented by a vectors of length μ over \mathbb{F}_2 . Also, \mathbf{v}_e and \mathbf{q}_e are vectors with elements in \mathbb{F}_2^{μ} and therefore can be represented by matrices of dimensions $\hat{\ell} \times \mu$ over \mathbb{F}_2 . We can then write the VOLE correlation equation as

$$\mathbf{Q}_e = [\delta_{e,0}\mathbf{u}_e \ \delta_{e,1}\mathbf{u}_e \ \cdots \ \delta_{e,\mu-1}\mathbf{u}_e] + \mathbf{v}_e$$

where $(\delta_{e,0}, \delta_{e,1}, \dots, \delta_{e,\mu-1})$ is the bit decomposition of $\Delta_e \in \mathbb{F}_2^{\mu}$. If the prover can somehow modify these correlations such that all τ instances use the same \mathbf{u} value (say \mathbf{u}_0), then the prover combine (concatenate) \mathbf{v}_e and \mathbf{Q}_e matrices to build VOLE correlation values \mathbf{v} and \mathbf{q} in $\mathbb{F}_2^{\hat{\ell}}$. Similarly, prover computes $\Delta \in \mathbb{F}_2^{\rho}$ by concatenating all bits $\{\delta_{e,i}\}_{(e,i) \in [0, \dots, \tau-1] \times [0, \dots, \mu-1]}$ from individual Δ_e values. This gives us a desired VOLE correlation

$$\mathbf{Q} = \mathbf{u}_0\Delta + \mathbf{v}$$

with $(\mathbf{u}_0, \mathbf{v}) \in \mathbb{F}_2^{\hat{\ell}} \times \mathbb{F}_2^{\hat{\rho}}$ and $(\mathbf{Q}, \Delta) \in \mathbb{F}_2^{\hat{\ell}} \times \mathbb{F}_2^{\rho}$. The prover achieves this by sending the correction values $\mathbf{c}_e := \mathbf{u}_0 - \mathbf{u}_e$ for $e \in [1, \dots, \tau-1]$. These \mathbf{c}_e values can be used by the verifier to adjust its correlation inputs such that all τ VOLE correlations in \mathbb{F}_2^{μ} hold with respect to \mathbf{u}_0 .

Ensuring consistency of VOLE correlations. Note that if any of the correction values \mathbf{c}_e is inconsistent (i.e. $\mathbf{c}_e \neq \mathbf{u}_0 - \mathbf{u}_e$) then the correctness of VOLE correlations does not hold and therefore the zero-knowledge proof built using such VOLE correlations cannot guarantee its correctness either. Therefore, the verifier must check that \mathbf{c}_e values are consistent. The verifier can ensure this by asking the prover compute a random linear universal hash function of \mathbf{u}_0 and \mathbf{v} , and send the hash values (say $\tilde{\mathbf{u}}, \tilde{\mathbf{v}}$). The verifier can then compute the same function on \mathbf{Q} and then check if the VOLE correlation $\tilde{\mathbf{Q}} = \tilde{\mathbf{u}}\Delta + \tilde{\mathbf{v}}$ holds true. This consistency check was used earlier in [Roy22, BBD⁺23b]. For more details kindly refer Section 4.4 and algorithm 4.8.

3.2 Proof of Knowledge for PKP

The zero-knowledge proof of knowledge underlying PERK is based on [BBGK24]. Recall that, the prover (signer) wants to prove the knowledge of a secret permutation matrix $\mathbf{P} \in \mathbb{F}_q^{n \times n}$ such that for some given public matrix $\mathbf{H} \in \mathbb{F}_q^{m \times n}$ (for $m < n$) and a public vector $\mathbf{x} \in \mathbb{F}_q^n$, the equation $\mathbf{HP}\mathbf{x} = \mathbf{0}$ holds true. In PERK, we achieve this goal in two steps, first the prover aims to convince the verifier that it knows a permutation matrix. Then it shows that the equation $\mathbf{HP}\mathbf{x} = \mathbf{0}$ holds true for this permutation matrix. We will therefore first focus on proving that the prover knows some permutation matrix.

Elementary vectors as building blocks. An elementary row vector e_i of length n for $0 \leq i \leq (n - 1)$ is the $(n - 1 - i)^{\text{th}}$ row of an $n \times n$ identity matrix.¹ Note that, for any $n \times n$ permutation matrix its rows are also elementary vectors e_i (but in an arbitrary order). Therefore, as a first step towards proving knowledge of a specific secret permutation matrix, we can begin by trying to prove that we know a certain elementary vector. The key observation is that an elementary vector of length n , can be constructed by taking tensor product of elementary vectors of smaller size. In particular, PERK uses elementary vectors of lengths 4 and 2 to prove the knowledge of rows of secret permutation matrices of sizes 64, 92, and 118. The prover proves the elementary structure of vector $e := [e_0, e_1, e_2, e_3] \in \mathbb{F}_2^4$ by showing that the product $e_0 \cdot e_1 = 0$ and $e_2 \cdot e_3 = 0$ simultaneously. In case of elementary vector $e := [e_0, e_1] \in \mathbb{F}_2^2$ this is achieved simply by showing $e_0 \cdot e_1 = 0$. Additionally, the prover should also prove that $e_0 \oplus e_1 \oplus e_2 \oplus e_3 = 1$ (resp. $e_0 \oplus e_1 = 1$). Proving these constraints simultaneously, allows the prover to demonstrate knowledge of elementary vectors, the prover achieves this by constructing equivalent low degree polynomials which have leading coefficient equal to 0 when the constraint is satisfied.

Knowledge of permutation and PKP solution. The prover computes the proof for each row of the secret permutation vector as a tensor product of d elementary vectors (for $d \in \{3, 4\}$) by constructing equivalent degree- d polynomials with leading coefficient equal to 0 if and only if the corresponding row is an elementary vector. At this stage, the prover is able to prove to the verifier the permutation structure of the matrix. The prover then shows that these degree- d polynomials when seen as entries of $n \times n$ matrix \mathbf{P} satisfy the PKP equation $\mathbf{H}\mathbf{P}\mathbf{x} = \mathbf{0}$, where (\mathbf{H}, \mathbf{x}) corresponds to the public key of PKP. In order to prove that all these polynomials have leading coefficients equal to 0, the prover combines them all by taking a random linear combination of all the polynomials (where the coefficients of the random linear combination are provided by the verifier), adds secret masking polynomial of degree $d - 1$ to the linear combination and sends it to the verifier. The verifier checks if the received polynomial has leading coefficient equal to 0 by evaluating it at a random point.

¹ We assume the indexes start from 0, that is the elementary vector e_0 denotes the last that is $(n - 1)^{\text{th}}$ row of the identity matrix.

4 Algorithmic Description

4.1 Object representation

Finite fields. Elements of \mathbb{F}_q are stored in 16 bit unsigned integers. We also use finite field arithmetic over $\mathbb{F}_{2^{11}}, \mathbb{F}_{2^{64}}, \mathbb{F}_{2^{128}}, \mathbb{F}_{2^{132}}, \mathbb{F}_{2^{198}}, \mathbb{F}_{2^{264}}$. These fields are defined as polynomials over \mathbb{F}_2 modulo an irreducible polynomial F .

$$\begin{aligned} F_{11}(\gamma) &= 1 + \gamma^2 + \gamma^{11} \\ F_{64}(\gamma) &= 1 + \gamma^1 + \gamma^3 + \gamma^4 + \gamma^{64} \\ F_{128}(\gamma) &= 1 + \gamma^1 + \gamma^2 + \gamma^7 + \gamma^{128} \\ F_{132}(\gamma) &= 1 + \gamma^3 + \gamma^{12} \\ F_{198}(\gamma) &= 1 + \gamma^7 + \gamma^{18} \\ F_{264}(\gamma) &= 1 + \gamma^1 + \gamma^{11} + \gamma^{17} + \gamma^{24} \end{aligned}$$

We use following functions to convert bit-strings into field elements (or positive numbers) and vice versa:

- **ToField** converts a bit-string into a corresponding field element ;
- **ToBits** converts a field element into a corresponding bit-string ;
- **BitDec** converts a positive number into its binary decomposition ;
- **NumRec** takes a bit-string as the binary decomposition of a positive number and reconstructs the number.

Algorithm 4.1: ToField(bits, k)
Public information and inputs
Public information: Maps an input bitstring $\text{bits} \in \{0, 1\}^{nk}$, for a positive integer $n \geq 1$ into a field element (or vector of n fields elements) $\mathbf{x} \in \mathbb{F}_{2^k}^n$.
<pre> 1 : let $\gamma_k \in \mathbb{F}_{2^k}$ // The γ element of \mathbb{F}_{2^k}. 2 : if $\text{bits} \in \{0, 1\}^k$: 3 : return $x := \sum_{i=0}^{k-1} \text{bits}[i] \cdot \gamma_k^i$ 4 : elseif $\text{bits} \in \{0, 1\}^{nk}$: 5 : for $i \in [n]$ 6 : $\mathbf{x}[i] := \sum_{j=0}^{k-1} \text{bits}[ni + j] \cdot \gamma_k^j$ 7 : endfor 8 : return \mathbf{x} 9 : else : 10 : return \perp 11 : endif </pre>

Algorithm 4.2: ToBits(\mathbf{x}, k, n)
Public information and inputs
Public information: Maps an input field element (or vector of n fields elements) $\mathbf{x} \in \mathbb{F}_{2^k}^n$, for a positive integer $n \geq 1$ into a bitstring $\text{bits} \in \{0, 1\}^{nk}$.
<pre> 1 : Initialize bits $\leftarrow \varepsilon$ // Empty string. 2 : Initialize bitslice $\leftarrow \varepsilon$ // Empty string. 3 : for $i \in [n]$ 4 : Parse $\mathbf{x}[i]$ as $\mathbf{x}[i] = x_0 + x_1\gamma_k + \dots + x_{k-1}\gamma_k^{k-1}$ with $x_0, x_1, \dots, x_{k-1} \in \{0, 1\}$ 5 : bitslice $:= x_0 x_1 \dots x_{k-1}$ // bitslice $\in \{0, 1\}^k$ 6 : bits $:= \text{bits} \text{bitslice}$ 7 : endfor 8 : return bits // bits $\in \{0, 1\}^{nk}$. </pre>

Integer and bits conversions									
<table border="1"> <thead> <tr> <th>Algorithm 4.3: BitDec(i, d)</th> </tr> </thead> <tbody> <tr> <td>Public information and inputs</td> </tr> <tr> <td>Public information: Decomposes an integer i into bits.</td> </tr> <tr> <td> <pre> 1 : for $j \in [d]$: 2 : $b_j := i \bmod 2$ 3 : $i := (i - b_j)/2$ 4 : endfor 5 : return $(b_0, b_1, \dots, b_{d-1})$. </pre> </td> </tr> </tbody> </table>	Algorithm 4.3: BitDec(i, d)	Public information and inputs	Public information: Decomposes an integer i into bits.	<pre> 1 : for $j \in [d]$: 2 : $b_j := i \bmod 2$ 3 : $i := (i - b_j)/2$ 4 : endfor 5 : return $(b_0, b_1, \dots, b_{d-1})$. </pre>	<table border="1"> <thead> <tr> <th>Algorithm 4.4: NumRec(d, bits)</th> </tr> </thead> <tbody> <tr> <td>Public information and inputs</td> </tr> <tr> <td>Public information: Reconstructs an integer i from a bitstring.</td> </tr> <tr> <td> <pre> 1 : Parse bits as $\text{bits} := b_0 \dots b_{d-1}$ 2 : return $\sum_{j=0}^{d-1} b_j \cdot 2^j$ </pre> </td> </tr> </tbody> </table>	Algorithm 4.4: NumRec(d, bits)	Public information and inputs	Public information: Reconstructs an integer i from a bitstring.	<pre> 1 : Parse bits as $\text{bits} := b_0 \dots b_{d-1}$ 2 : return $\sum_{j=0}^{d-1} b_j \cdot 2^j$ </pre>
Algorithm 4.3: BitDec(i, d)									
Public information and inputs									
Public information: Decomposes an integer i into bits.									
<pre> 1 : for $j \in [d]$: 2 : $b_j := i \bmod 2$ 3 : $i := (i - b_j)/2$ 4 : endfor 5 : return $(b_0, b_1, \dots, b_{d-1})$. </pre>									
Algorithm 4.4: NumRec(d, bits)									
Public information and inputs									
Public information: Reconstructs an integer i from a bitstring.									
<pre> 1 : Parse bits as $\text{bits} := b_0 \dots b_{d-1}$ 2 : return $\sum_{j=0}^{d-1} b_j \cdot 2^j$ </pre>									

Vectors and Matrices. Vectors of \mathbb{F}_q^n (respectively \mathbb{F}_q^m) are represented as arrays of length n (respectively of length m) of \mathbb{F}_q elements. Matrices $\mathbf{H} \in \mathbb{F}_q^{m \times n}$ are represented as two dimensional arrays of \mathbb{F}_q elements i.e. arrays of length m of arrays of length n .

Permutations and Witness. We use following auxiliary functions during signing process to encode positive numbers corresponding to the index of non-zero entries of the secret permutation matrix, and represent these secret indices as witness for the proof system.

- [EncodeNum-64](#) encodes a number between 0 to 63 in base-4 notation. This function is used for NIST Level-I parameter set. Given an input $\text{pos} \in [64]$,

it outputs a unique tuple of three numbers $(i, j, k) \in [4]^3$ such that $\text{pos} = 16k + 4j + i$;

- **EncodeNum-128** encodes a number between 0 to 127 using hybrid base-4 and base-2 notation. This function is used for NIST Level-III and NIST Level-V parameter sets. On input $\text{pos} \in [127]$, it outputs a unique tuple of four numbers $(i, j, k, b) \in [4]^3 \times \{0, 1\}$ such that $\text{pos} = 64b + 16k + 4j + i$;
- **EncodeNum** is a wrapper function that internally calls either **EncodeNum-64** or **EncodeNum-128** depending on the security level ;
- **LeftShift** shifts the bits of an input bit-string to left by a specified amount given as input.

Algorithm 4.5: EncodeNum-64(pos)
Public information and inputs
Public information: Encodes input number $\text{pos} \in [64]$ in base-4. order.
Output
Array encPosArray of length 3, encoding the input position pos .

<pre> 1 : Initialize $\text{encPosArray} \leftarrow [\text{null}, \text{null}, \text{null}]$ 2 : $k \leftarrow \lfloor \frac{\text{pos}}{16} \rfloor$ 3 : $j \leftarrow \lfloor \frac{\text{pos} - (k * 16)}{4} \rfloor$ 4 : $i \leftarrow \text{pos} - (k * 16) - (j * 4)$ 5 : $\text{encPosArray} \leftarrow [i, j, k]$ 6 : return encPosArray </pre>

Algorithm 4.6: EncodeNum-128(pos)
Public information and inputs
Public information: Encodes input number $\text{pos} \in [128]$ in hybrid base-4/base-2.
Output
Array encPosArray of length 4, encoding the input position pos .

<pre> 1 : Initialize $\text{encPosArray} \leftarrow [\text{null}, \text{null}, \text{null}, \text{null}]$ 2 : $b \leftarrow \lfloor \frac{\text{pos}}{64} \rfloor$ // $b \in \{0, 1\}$ 3 : $\text{temp} \leftarrow \text{EncodeNum-64}(\text{pos} - (b * 64))$ // Parse temp as $\text{temp} := [i, j, k]$ where $i, j, k \in [0, 3]$ 4 : $\text{encPosArray} \leftarrow [i, j, k, b]$ 5 : return encPosArray </pre>

Algorithm 4.7: EncodeNum(pos)
Public information and inputs
Public information: Wrapper function for selecting encoding function for input number pos based on desired security level.
Output
Array encPosArray encoding the input position pos.

<p style="text-align: center;"><u>Security Level 1</u></p> 1: if $\lambda = 128$: 2: return EncodeNum-64 (pos)
<p style="text-align: center;"><u>Security Levels 3 and 5</u></p> 3: elseif $\lambda \in \{192, 256\}$: 4: return EncodeNum-128 (pos) 5: else : 6: return \perp 7: endif

Algorithm 4.8: LeftShift(bits, shift)
Public information and inputs
Public information: Shifts input bitstring bits $\in \{0, 1\}^*$ to left by shift positions if shift is smaller than length of bits.

1: if shift \in [len(bits)]: 2: return bits \ll shift 3: else : 4: return \perp 5: endif

4.2 Sampling functions

The `randombytes` function provided by the NIST is used to sample uniformly at random the salt and various seeds (e.g., \mathbf{H}_{seed} , ker_{seed} , $\text{perm}_{\text{seed}}$). The PRG function is instantiated using SHAKE-128 for $\lambda = 128$ and SHAKE-256 otherwise, along with domain separators.

Random elements of \mathbb{F}_q are obtained by sampling $\log_2(q) = 11$ random bits from the PRG. Random vectors in \mathbb{F}_q^n (respectively matrices in $\mathbb{F}_q^{m \times n}$) are sampled uniformly by sampling in order n (respectively $m \times n$) elements in \mathbb{F}_q .

$\text{ExpandMatrixM}(\mathbf{H}_{\text{seed}}) \rightarrow \mathbb{F}_q^{m \times (n-m)}$: Samples a matrix $\mathbf{M} \in \mathbb{F}_q^{m \times (n-m)}$ uniformly at random using the PRG with seed \mathbf{H}_{seed} and domain separator $\text{dom} = 0 \times 00$.

$\text{ExpandKernelVector}(\text{ker}_{\text{seed}}, \mathbf{H}) \rightarrow \ker(\mathbf{H})$: Samples a vector $\mathbf{x}' \in \mathbb{F}_q^n$ in $\ker(\mathbf{H})$ uniformly at random. Specifically, we derive a basis $\mathbf{k}_1, \dots, \mathbf{k}_{n-m} \in \mathbb{F}_q^n$ for $\ker(\mathbf{H})$ by taking the rows of the matrix $(\mathbf{M}^\top | \mathbf{I}_{n-m})$. Then we sample $n - m$ random scalars $c_1, \dots, c_m \in \mathbb{F}_q$ using the PRG with seed ker_{seed} and domain separator $\text{dom} = 0 \times 10$, and compute the resulting vector in $\ker(\mathbf{H})$ as $\mathbf{x}' = \sum_{i=0}^{n-m} c_i \cdot \mathbf{k}_i$. For an adversary, its advantage in distinguishing the output of $\text{ExpandKernelVector}$ from a random kernel vector is denoted by $\text{AdvPR}^{\text{ExpandKernelVector}}$.

$\text{ExpandPermutation}(\text{perm}_{\text{seed}}) \rightarrow \mathcal{S}_n$: Samples permutation π of length n uniformly at random. Specifically, we first construct a vector $\mathbf{v} = (v_0, \dots, v_{n-1}) = (0, 1, \dots, n-1)$. Then, we use the PRG with seed $\text{perm}_{\text{seed}}$ and domain separator $\text{dom} = 0 \times 20$ for sampling a random vector $\mathbf{e} = (e_0, \dots, e_{n-1}) \in (\mathbb{F}_2^{16})^n$. We construct the vector $\mathbf{p} = (p_0, \dots, p_{n-1})$, where the high-order and low-order bits of p_i corresponds e_i and v_i , respectively. Finally, we sort this integer sequence in constant time using `djbsort` [Ber19], and extract the permutation π from the lower-order bits \mathbf{p} . If there are any duplicate values in the vector \mathbf{e} , we discard it and restart the procedure. For an adversary, its advantage in distinguishing the output of ExpandPermutation from a random permutation is denoted by $\text{AdvPR}^{\text{ExpandPermutation}}$.

4.3 Hash functions and commitments

In the following, we instantiate functions from the **SHA3** family, choosing the output length according to the security parameter λ : for $\lambda = 128$ we use **SHA3-256**, for $\lambda = 192$ we use **SHA3-384**, and for $\lambda = 256$ we use **SHA3-512**. Throughout, we denote a binary string as $x \in \{0, 1\}^*$.

Pseudorandom Generators. We make use of two distinct pseudorandom generators, denoted by PRG_1 and PRG_2 . Specifically, PRG_1 is instantiated from either the **SHA3** family or **Rijndael**, while PRG_2 is instantiated from either **SHAKE** or **Rijndael**.

- $\text{PRG}_1(\text{salt} \parallel \text{index} \parallel \text{seed})$ expands a seed into a binary tree according to the GGM construction, producing outputs of length 2λ . The length of `index` depends on the instantiation: 2 bytes for **SHA3** and 4 bytes for **AES/Rijndael**. Let $\text{salt} = (\text{salt}_0 \parallel \text{salt}_1)$.

$$\text{PRG}_1(\text{salt} \parallel \text{index} \parallel \text{seed}) := \text{SHA3-}\lambda(\text{salt} \parallel \text{index} \parallel \text{seed} \parallel \text{dom})$$

$$\text{PRG}_1(\text{salt} \parallel \text{index} \parallel \text{seed}) := (\text{high} \parallel \text{low})$$

where,

$$\begin{aligned} \text{high} &= \begin{cases} \text{AES-128}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0 \times 00 \parallel \text{index} \parallel \text{dom})) & \text{if } \lambda = 128 \\ \text{Rijndael-256}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0 \times 00 \parallel \text{index} \parallel \text{dom})) & \text{otherwise.} \end{cases} \\ \text{low} &= \begin{cases} \text{AES-128}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0 \times 01 \parallel \text{index} \parallel \text{dom})) & \text{if } \lambda = 128 \\ \text{Rijndael-256}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0 \times 01 \parallel \text{index} \parallel \text{dom})) & \text{otherwise.} \end{cases} \end{aligned}$$

- $\text{PRG}_2(\text{salt} \parallel \text{seed})$ converts seeds into instances for the Vole protocol, with output length $\hat{\ell}$.

$$\text{PRG}_2(\text{salt} \parallel \text{seed}) := \begin{cases} \text{SHAKE-128}(\text{salt} \parallel \text{seed} \parallel \text{dom}) & \text{if } \lambda = 128, \\ \text{SHAKE-256}(\text{salt} \parallel \text{seed} \parallel \text{dom}) & \text{if } \lambda \in \{192, 256\}. \end{cases}$$

$$\text{PRG}_2(\text{salt} \parallel \text{seed}) := \begin{cases} \text{AES-128}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (\text{ctr} \parallel \text{z} \parallel \text{dom})) & \text{if } \lambda = 128, \\ \text{Rijndael-256}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (\text{ctr} \parallel \text{z} \parallel \text{dom})) & \text{otherwise,} \end{cases}$$

where ctr is a counter byte that starts from 0×00 and gets updated for every block produced, and $\text{z} = (0 \times 00 \parallel 0 \times 00 \parallel 0 \times 00 \parallel 0 \times 00)$.

In both cases, domain separation is enforced via an explicit tag $\text{dom} \in 0 \times 05, 0 \times 06$, ensuring independence between PRG_1 and PRG_2 . For the security level parameter $\lambda = 192$, we extend salt_0 to 256 bits by appending 0×00 s, and we truncate the output of Rijndael-256 to 192 bits.

Hash functions.

- $H_1 : \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$ defined as

$$H_1(x) := \text{SHA3-}\lambda(x \parallel \text{dom}), \text{ where } \text{dom} = 0 \times 01.$$

- $H_2^1 : \{0, 1\}^* \rightarrow \{0, 1\}^{5\lambda+8}$ defined as

$$H_2^1(x) := \begin{cases} \text{SHAKE-128}(x \parallel \text{dom}) & \text{if } \lambda = 128, \text{ dom} = 0 \times 21, \\ \text{SHAKE-256}(x \parallel \text{dom}) & \text{if } \lambda \in \{192, 256\}, \text{ dom} = 0 \times 21. \end{cases}$$

- $H_2^2 : \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$ defined as

$$H_2^2(x) := \text{SHA3-}\lambda(x \parallel \text{dom}), \text{ where } \text{dom} = 0 \times 22.$$

- $H_2^3 : \{0, 1\}^* \rightarrow \{0, 1\}^{\tau_0 \kappa_0 + \tau_1 \kappa_1 + w}$ defined as

$$H_2^3(x) := \begin{cases} \text{SHAKE-128}(x \parallel \text{dom}) & \text{if } \lambda = 128, \text{ dom} = 0 \times 23, \\ \text{SHAKE-256}(x \parallel \text{dom}) & \text{if } \lambda \in \{192, 256\}, \text{ dom} = 0 \times 23. \end{cases}$$

- $H_3 : \{0, 1\}^* \rightarrow \{0, 1\}^{3\lambda}$ defined as

$$H_3(x) := \begin{cases} \text{SHAKE-128}(x \parallel \text{dom}) & \text{if } \lambda = 128, \text{ dom} = 0\text{x}03, \\ \text{SHAKE-256}(x \parallel \text{dom}) & \text{if } \lambda \in \{192, 256\}, \text{ dom} = 0\text{x}03. \end{cases}$$

- $H_4 : \{0, 1\}^* \rightarrow \{0, 1\}^{\rho(cn+n+m)}$ defined as

$$H_4(x) := \begin{cases} \text{SHAKE-128}(x \parallel \text{dom}) & \text{if } \lambda = 128, \text{ dom} = 0\text{x}04, \\ \text{SHAKE-256}(x \parallel \text{dom}) & \text{if } \lambda \in \{192, 256\}, \text{ dom} = 0\text{x}04. \end{cases}$$

Commitments. For the commitments, we consider the two following approaches.

Commitment Com_1 . This scheme is instantiated from either the **SHA3** family or **Rijndael**. Let τ denote the subtree and n the index of the leaf in the GGM tree array.

- **SHA3** instantiation: we absorb τ as a single byte and n as two bytes.
- **Rijndael** instantiation: we concatenate τ and n into a 32-bit string (4 bytes).

We denote this combined value as index . Let $\text{salt} = (\text{salt}_0 \parallel \text{salt}_1)$.

Option 1: **SHA3** based commitment.

$\text{Com}_1(\text{salt} \parallel \text{index} \parallel \text{seed}) := \text{SHA3-}\lambda(\text{salt} \parallel \text{index} \parallel \text{seed} \parallel \text{dom})$, where $\text{dom} = 0\text{x}07$.

Option 2: **Rijndael**-based commitment.

$\text{Com}_1(\text{salt} \parallel \text{index} \parallel \text{seed}) := (\text{high} \parallel \text{low})$, where

$$\begin{aligned} \text{high} &= \begin{cases} \text{AES-128}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0\text{x}00 \parallel \text{index} \parallel \text{dom})) & \text{if } \lambda = 128 \\ \text{Rijndael-256}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0\text{x}00 \parallel \text{index} \parallel \text{dom})) & \text{otherwise.} \end{cases} \\ \text{low} &= \begin{cases} \text{AES-128}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0\text{x}01 \parallel \text{index} \parallel \text{dom})) & \text{if } \lambda = 128 \\ \text{Rijndael-256}(k = \text{seed}, \text{msg} = \text{salt}_0 \oplus (0\text{x}01 \parallel \text{index} \parallel \text{dom})) & \text{otherwise.} \end{cases} \end{aligned}$$

For the security level parameter $\lambda = 192$, we extend salt_0 to 256 bits by appending $0\text{x}00\text{s}$, and we truncate the output of **Rijndael-256** to 192 bits. In all three cases, we take $\text{dom} = 0\text{x}07$.

Commitment Com_2 . This scheme is derived from the **SHA3** family

$$\text{Com}_2(x) := \text{SHA3-}\lambda(x \parallel \text{dom}), \text{ where } \text{dom} = 0\text{x}08.$$

4.4 VOLE-in-the-Head functions

The VOLE correlations form one of the foundational building block of our scheme. In this section, we describe how to perform basic operations on VOLE correlations and how to construct them from the batch all-but-one vector commitments construction. We also explain how, within the scheme, the prover commits to these VOLE correlations and checks their consistency.

VOLE correlations. Let $\mathbf{q} = u\Delta \oplus \mathbf{v}$ be a VOLE correlation with $(u, \mathbf{v}) \in \mathbb{F}_2 \times \mathbb{F}_{2^\mu}$, and $(\mathbf{q}, \Delta) \in \mathbb{F}_{2^\mu} \times \mathbb{F}_{2^\mu}$.² Such a generic VOLE correlation corresponds to a linear (degree-1) commitment to u denoted as $f_u(X) = uX + v$, with the evaluation $q = f_u(\Delta)$ given to the verifier. From here onward in this document, we denote degree-1 commitment to a bit u (or a bit-string $\mathbf{u} \in \mathbb{F}_2^{\hat{\ell}}$) by $\llbracket u \rrbracket$ (or $\llbracket \mathbf{u} \rrbracket$ respectively). This notion can be extended to polynomial-based commitments with higher degree polynomials. We write $\llbracket s \rrbracket^{(d)}$ to denote a degree- d commitment to a secret value $s \in \mathbb{F}_2$ where the prover holds $f_s(X) = \sum_{i=0}^d a_i X^i$ with coefficients $a_i \in \mathbb{F}_{2^\mu}$ and a_d equal to s lifted to \mathbb{F}_{2^μ} while the verifier holds $q_s = f_s(\Delta) \in \mathbb{F}_{2^\mu}$.

It is possible to compute arbitrary linear combinations of a given set of input VOLE correlations, due to their linear homomorphic property. [Algorithm 4.9 LinearCombination](#) from [\[BBD⁺23b\]](#) given below shows how VOLE correlations for linear functions of secret values $u_1, \dots, u_n \in \mathbb{F}_2$ can be computed. In fact, it is also possible to combine k VOLE correlations $\mathbf{q}_i = u_i\Delta \oplus \mathbf{v}_i, i \in [k]$ with $(u_i, \mathbf{v}_i) \in \mathbb{F}_2 \times \mathbb{F}_{2^\mu}$ and $(\mathbf{q}_i, \Delta) \in \mathbb{F}_{2^\mu} \times \mathbb{F}_{2^\mu}$ for $i \in [k]$, to obtain a single VOLE correlation $\mathbf{q} = \mathbf{u}\Delta \oplus \mathbf{v}$ with $(\mathbf{u}, \mathbf{v}) \in \mathbb{F}_{2^k} \times \mathbb{F}_{2^\mu}$ and $(\mathbf{q}, \Delta) \in \mathbb{F}_{2^\mu} \times \mathbb{F}_{2^\mu}$ for an arbitrary $\mathbb{F}_{2^k} \subseteq \mathbb{F}_{2^\mu}$. To achieve this, the prover computes $\mathbf{u} := \sum_{i=0}^{k-1} u_i X^i$ where $\{1, X, \dots, X^{k-1}\}$ is the power-basis of \mathbb{F}_{2^k} over \mathbb{F}_2 . Prover also computes \mathbf{v} as $\mathbf{v} := \sum_{i=0}^{k-1} \mathbf{v}_i X^i$. While the verifier can compute \mathbf{q} as $\mathbf{q} := \sum_{i=0}^{k-1} \mathbf{q}_i \Delta^i$ (and Δ remains unchanged). We can perform homomorphic operations on the degree- d -commitments locally by the prover and the verifier with the help of the [Algorithm 4.10 Add](#) and [Algorithm 4.11 Multiply](#) given below. In the following, let $d_1 \geq d_2$ without loss of generality, also let $d = d_1 + d_2$.

Given a VOLE correlation $q = u\Delta + v$ for a random $u \in \mathbb{F}_2$, it is possible to embed arbitrary value $w \in \mathbb{F}_2$. To do so, the prover computes $t = w \oplus u$ and sends t to the verifier. Since u is uniform random and unknown to the verifier, t does not leak any information about w . The verifier then computes $q_w = q + t\Delta$. Note that the prover and verifier now possess their respective parts for the VOLE correlation $q_w = w\Delta + v$. After embedding the witness w , one can use the VOLE correlations to establish PoK of the witness for relations which can be modeled as polynomial functions of the commitment to the witness.

² Here we consider $\hat{\ell} = 1$ for readability and easy to read notation. Therefore, u and Δ can be seen as scalars represented by 1 and μ -bits respectively. Whereas, \mathbf{v}, \mathbf{q} are vectors of length μ over \mathbb{F}_2 . All of the discussion in [Section 4.4](#) naturally extends to VOLE correlations with $\hat{\ell} > 1$, where we view Δ as a scalar represented by μ -bits, \mathbf{u} is a vector of length $\hat{\ell}$ over \mathbb{F}_2 , and \mathbf{V}, \mathbf{Q} are seen as $\hat{\ell} \times \mu$ matrices over \mathbb{F}_2 .

Algorithm 4.9: $\text{LinearCombination}(c_0, c_1, \dots, c_n, (\llbracket u_i \rrbracket)_{i \in [1, n]})$
Prover's computation: $\text{P.LinearCombination}(c_0, c_1, \dots, c_n, (\llbracket u_i \rrbracket)_{i \in [1, n]})$
<p>Prover's input: Coefficients of linear combination $c_0, c_1, \dots, c_n \in \mathbb{F}_2$, VOLE correlation inputs $(u_1, v_1), \dots, (u_n, v_n) \in (\mathbb{F}_2 \times \mathbb{F}_{2^\rho})^n$.</p> <p>Prover's output: VOLE correlations (u, v) for the linear combination of secret inputs.</p>

Computes $u := c_0 + \sum_{i=1}^n c_i u_i$ and $v := \sum_{i=1}^n c_i v_i$.
Verifier's computation: $\text{V.LinearCombination}(c_0, c_1, \dots, c_n, \Delta, q_1, \dots, q_n)$
<p>Verifier's input: Coefficients of linear combination $c_0, c_1, \dots, c_n \in \mathbb{F}_2$, VOLE correlation inputs $\Delta, q_1, \dots, q_n \in \mathbb{F}_{2^\rho}^{n+1}$</p> <p>Verifier's output: VOLE correlation q for linear combination of secret inputs.</p>

Computes $q := c_0 \Delta + \sum_{i=1}^n c_i q_i$.

Algorithm 4.10: $\text{Add}(\llbracket s_1 \rrbracket^{(d_1)}, \llbracket s_2 \rrbracket^{(d_2)})$
Public information:
Degrees of input VOLE correlations d_1, d_2 .
Prover's computation: $\text{P.Add}(\llbracket s_1 \rrbracket^{(d_1)}, \llbracket s_2 \rrbracket^{(d_2)})$
<p>Prover's input: VOLE correlations represented as polynomials $f_{s_1}(X)$ and $f_{s_2}(X)$.</p> <p>Prover's output: VOLE correlation $\llbracket s \rrbracket^{(d_1)}$ for addition of secret inputs.</p>

Computes $\llbracket s \rrbracket^{(d_1)} = f_s(X) := f_{s_1}(X) + f_{s_2}(X)X^{d_1-d_2}$ where $s = s_1 + s_2$.
Verifier's computation: $\text{V.Add}(\Delta, q_{s_1}, q_{s_2})$
<p>Verifier's input: $\Delta, q_{s_1} = f_{s_1}(\Delta)$, and $q_{s_2} = f_{s_2}(\Delta)$</p> <p>Verifier's output: VOLE correlation q_s for addition of secret inputs.</p>

Computes $q_s := q_{s_1} + q_{s_2} \Delta^{d_1-d_2}$.

Algorithm 4.11: $\text{Multiply}(\llbracket s_1 \rrbracket^{(d_1)}, \llbracket s_2 \rrbracket^{(d_2)})$
Public information:
Degrees of input VOLE correlations d_1, d_2 .
Prover's computation: $\text{P.Multiply}(\llbracket s_1 \rrbracket^{(d_1)}, \llbracket s_2 \rrbracket^{(d_2)})$
Prover's input: VOLE correlations represented as polynomials $f_{s_1}(X)$ and $f_{s_2}(X)$. Prover's output: VOLE correlation $\llbracket s \rrbracket^{(d_1)}$ for multiplication of secret inputs.

Computes $\llbracket s \rrbracket^{(d)} = f_s(X) := f_{s_1}(X)f_{s_2}(X)$ where $s = s_1s_2$.
Verifier's computation: $\text{V.Multiply}(q_{s_1}, q_{s_2})$
Verifier's input: $q_{s_1} = f_{s_1}(\Delta)$ and $q_{s_2} = f_{s_2}(\Delta)$ Verifier's output: VOLE correlation q_s for multiplication of secret inputs.

Computes $q_s := q_{s_1}q_{s_2}$.

Committing to VOLE correlations. Recall that in order to achieve the desired security level (soundness), the atomic zero-knowledge protocol based on VOLE correlations should be repeated τ times. This means that the prover needs to generate τ GGM trees instances (or a forest of τ GGM trees). Let N_e be number of leaves (or seeds to be committed to) in each of such τ trees for $e \in [\tau]$ and let $N := \sum_{e=0}^{\tau-1} N_e$. Recently authors of [BBM⁺25] showed how the communication cost of such commitments can be further reduced by using *batch all-but-one vector commitment (BAVC) schemes*. The idea is to generate a single big GGM tree with N leaves instead of τ individual trees per instance with N_e leaves each. Note that in both cases (either with a forest of τ smaller trees or with single unified tree) the prover needs to hide a total of τ leaves. Intuitively, opening all-but- τ leaves of the unified tree is more efficient than opening all-but-one leaves of τ smaller trees, if the leaves to be opened in the big tree are relatively close to each other (or share some ancestor node in the tree). To achieve this, following the authors of [BBM⁺25] we interleave the leaves of the τ instances. That is, the first τ leaves of the big tree correspond to the first entry of the individual τ vector commitments, the next τ leaves correspond to the second entries, and so on. We also use a fixed threshold value T_{open} to ensure that the revealed path is not too long (thus avoiding long signature sizes). The opening algorithm aborts if the number of nodes exceeds T_{open} . This results in rejection sampling during the opening, which reduces the entropy of the challenge space. Fortunately, in [BBM⁺25] the authors showed that security is actually unaffected: since each rejection sampling step results in the prover computing a hash function, which can be considered as a proof of work done during each signing operation. We refer the interested readers to [BBM⁺25] for further details.

The batch BAVC consists of the following algorithms.

- **Algorithm 4.1** called as **VC.Commit** generates a vector commitment using a master seed \mathbf{mseed} and a salt value \mathbf{salt} as inputs with the help of length-doubling salt based PRG PRG_1 . It outputs the leaves of the tree as seeds $\{\mathbf{seed}_{e,i}\}_{\tau,N}$, a commitment h_{com} to all the seeds and decommitment information \mathbf{decom} which consists of commitments $\{\mathbf{com}_{e,i}\}_{\tau,N}$ to individual seeds and all intermediate nodes required to construct the unified GGM tree.
- **Algorithm 4.2** **VC.Open** takes the decommitment information \mathbf{decom} and the index set \mathbf{i}^* of size τ corresponding to the indexes of the hidden leaves. It outputs a partial decommitment \mathbf{pdecom} which can be used to recompute all-but- τ seeds in the GGM tree, along with τ commitments $\{\mathbf{com}_{e,i^*}\}_{\tau,\tau}$ to the hidden seeds.
- **Algorithm 4.3** **VC.Reconstruct** on inputs the index set \mathbf{i}^* of size τ (same as in **VC.Open**), \mathbf{pdecom} , and salt value \mathbf{salt} outputs all-but- τ seeds (hidden by the indexes in \mathbf{i}^*) by reconstructing the GGM tree, and also outputs the commitment h_{com} to all the seeds.

In the following, we assume that $N_0 = \dots = N_{\tau_0-1} \geq N_{\tau_0} = \dots = N_{\tau-1}$ for some τ_0 and define N as $N := \sum_{e=0}^{\tau-1} N_e$. We define ψ as

$$\psi(e, i) = \begin{cases} i \cdot \tau + e & \text{if } i < N_{\tau_0} \\ N_{\tau_0} \cdot \tau + (i - N_{\tau_0}) \cdot \tau_0 + e & \text{otherwise} \end{cases} \quad (1)$$

We also use $\{\mathbf{seed}_{e,i}\}_{\tau,N} := (\mathbf{seed}_{0,0}, \mathbf{seed}_{0,1}, \dots, \mathbf{seed}_{(\tau-1), (N_{\tau-1}-1)})$ to denote an ordered tuple of N seeds where $e \in [\tau]$ and $i \in [N_e]$. Similarly, we use $\{\mathbf{com}_{e,i}\}_{\tau,N} := (\mathbf{com}_{0,0}, \mathbf{com}_{0,1}, \dots, \mathbf{com}_{(\tau-1), (N_{\tau-1}-1)})$ to denote an ordered tuple of N commitments where $e \in [\tau]$ and $i \in [N_e]$. We also use $\{\mathbf{com}_{e,i^*}\}_{\tau,\tau} := (\mathbf{com}_{0,i^*[0]}, \mathbf{com}_{1,i^*[1]}, \dots, \mathbf{com}_{(\tau-1), i^*[\tau-1]})$ to denote an ordered tuple of τ commitments where $e \in [\tau]$, and $\mathbf{i}^* \in [N_0] \times [N_1] \times \dots \times [N_{\tau-1}]$ is an ordered list of τ indexes corresponding to hidden leaves.

Algorithm 4.1: VC.Commit(mseed, salt)
Public information and inputs
Public information: A number of iterations τ , a number of parties $N = \sum_{e=0}^{\tau-1} N_e$. Prover's input: A master seed $mseed \in \{0, 1\}^\lambda$ and a salt $\in \{0, 1\}^{2\lambda}$
Output
N seeds $\{\text{seed}_{e,i}\}_{\tau,N} \in (\{0, 1\}^\lambda)^N$, a commitment $h_{\text{com}} \in \{0, 1\}^{2\lambda}$ and a decommitment auxiliary variable decom .
<pre> 1 : nodes[0] := mseed 2 : for i ∈ [N − 1]: 3 : (nodes[2i + 1], nodes[2i + 2]) := PRG₁(salt, nodes[i] :: 2λ) 4 : endfor 5 : for e ∈ [τ]: 6 : for i ∈ [N_e]: 7 : seed_{e,i} := nodes[N − 1 + ψ(e, i)] 8 : com_{e,i} := Com₁(salt, seed_{e,i} :: 2λ) 9 : endfor 10 : endfor 11 : h_{com} := Com₂(salt, (cmt_{0,0}, . . . , cmt_{(τ−1),(N_{τ−1}−1)}) :: 2λ) 12 : return ({seed_{e,i}}_{τ,N}, h_{com}, decom := (nodes, {com_{e,i}}_{τ,N})) </pre>

Algorithm 4.2: VC.Open(decom, i^*)
Public information and inputs
<p>Public information: A number of iterations τ, a number of parties $N = \sum_{e=0}^{\tau-1} N_e$, a rejection parameter T_{open}</p> <p>Prover's input: A decommitment auxiliary variable $\text{decom} := (\text{nodes}, \{\text{com}_{e,i}\}_{\tau,N})$ and $i^* \in [N_0] \times [N_1] \times \dots \times [N_{\tau-1}]$ is an ordered list of τ indexes corresponding to hidden leaves.</p>
Output
<p>A sibling path pdecom and unopened commitments $\{\text{com}_{e,i^*}\}_{\tau,\tau}$</p> <hr/> <pre> // Selecting hidden leaves from interleaved tree 1: hidden := {N - 1 + $\psi(e, i^*[e])$: $e \in [\tau]$} 2: revealed := {N - 1, ..., 2N - 2} \setminus hidden 3: for i from N - 2 downto 0 : // Adding parent node to revealed if both children nodes are in revealed. 4: if (2i + 1) \in revealed and (2i + 2) \in revealed : 5: revealed := (revealed \setminus {2i + 1, 2i + 2}) \cup {i} 6: endfor 7: if len(revealed) > T_{open} : 8: return \perp 9: pdecom := \emptyset 10: for i \in [2N - 1] : 11: if i \in revealed : 12: pdecom := (pdecom nodes[i]) 13: endfor 14: return (pdecom, {com_{e,i*}}_{τ,τ}) </pre>

Algorithm 4.3: VC.Reconstruct(\mathbf{i}^* , pdecom, $\{\text{com}_{e,\mathbf{i}^*}\}_{\tau,\tau}$, salt)

Public information and inputs

Public information: A number of iterations τ , a number of parties $N = \sum_{e=0}^{\tau-1} N_e$.

Verifier's input: An ordered list $\mathbf{i}^* \in [N_0] \times [N_1] \times \dots \times [N_{\tau-1}]$ of τ indexes corresponding to hidden leaves, a sibling path pdecom and unopened commitments $\{\text{com}_{e,\mathbf{i}^*}\}_{\tau,\tau}$

Output

A commitment $h_{\text{com}} \in \{0,1\}^{2\lambda}$, and $\{\text{seed}_{e,i}\}_{\tau,N} := (\text{seed}_{0,0}, \text{seed}_{0,1}, \dots, \text{seed}_{(\tau-1),(N-1)})$ to denote an ordered tuple of N seeds where $e \in [\tau]$ and $i \in [N_e]$. Note that the τ seeds corresponding to the hidden leaves denoted by an ordered tuple $\{\text{seed}_{e,\mathbf{i}^*}\}_{\tau,\tau} := (\text{seed}_{0,\mathbf{i}^*[0]}, \text{seed}_{1,\mathbf{i}^*[1]}, \dots, \text{seed}_{(\tau-1),\mathbf{i}^*[\tau-1]})$ are set to \perp .

```

// Selecting hidden leaves from interleaved tree
1 : hidden := {N - 1 + ψ(e, i*[e]) : e ∈ [τ]}
2 : revealed := {N - 1, ..., 2N - 2} \ hidden
3 : for i from N - 2 downto 0 :
    // Adding parent node to revealed if both children nodes are in revealed.
4 :   if (2i + 1) ∈ revealed and (2i + 2) ∈ revealed :
5 :     revealed := (revealed \ {2i + 1, 2i + 2}) ∪ {i}
6 :   endif
    // Check if pdecom is well formed by checking that number of nodes
    // in revealed matches with pdecom and is ≤ Topen.
7 :   if len(revealed) ≠ len(pdecom) or len(revealed) > Topen :
8 :     return ⊥
9 :   endif
10 : nodes[0], ..., nodes[2N - 2] := ∅, ..., ∅
11 : for i ∈ [N - 1] :
12 :   if i ∈ revealed :
13 :     (nodes[i], pdecom) := pdecom
14 :   if nodes[i] ≠ ∅ :
15 :     nodes[2i + 1], nodes[2i + 2] := PRG1(salt, nodes[i] :: 2λ)
16 :   endif
17 : for e ∈ [τ] :
18 :   for i ∈ [Ne] :
19 :     if i ≠ i*[e] :
20 :       seede,i := nodes[N - 1 + ψ(e, i)]
21 :       come,i := Com1(salt, seede,i :: 2λ)
22 :     else :
23 :       seede,i := ⊥
24 :       come,i := come,i*[e]
25 :     endif
26 :   endif
27 : hcom := Com2(salt, (cmt0,0, ..., cmt(τ-1),i(Nτ-1-1)) :: 2λ)
28 : return (hcom, {seede,i}_{τ,N})

```

Computing VOLE from seed. The N seeds committed via the vector commitments are then converted to VOLE correlations using the algorithms specified below. Note that we use the exact same algorithm [Algorithm 4.4 ConvertToVOLE](#) as [\[Roy22, BBD⁺23b\]](#) for this conversion which uses divide-and-conquer approach to compute the VOLE correlations iteratively. The only difference is sometimes we lift the elements from \mathbb{F}_{2^κ} to \mathbb{F}_{2^μ} to ensure that the finite field arithmetic between objects is compatible.

The signing algorithm [PERK.Sign](#) uses the outputs of vector commitment and [ConvertToVOLE](#) algorithms to commit to τ instances of VOLE correlations. This is achieved by [Algorithm 4.5 VOLECommit](#) which creates τ instances of VOLE correlations by running [ConvertToVOLE](#) τ times. It then also computes the correction values \mathbf{c}_e for $e \in [1, \dots, \tau - 1]$ and outputs the VOLE signers correlation inputs \mathbf{u}, \mathbf{V} , correction vales \mathbf{c}_e for $e \in [1, \dots, \tau - 1]$, and commitment and decommitment information from vector commitment.

The challenge decoding algorithm [Algorithm 4.6 ChallDec](#) takes an input challenge string ch of length $\log(N)$ where N is the number of leaves in the GGM tree, and outputs an index set \mathbf{i}^* indicating the indexes of the hidden leaves. \mathbf{i}^* is computed by parsing τ chunks of input ch and converting them to an integer $\in [N_e]$ for $e \in [\tau]$. [ChallDec](#) is used by both signer to create opening information for the verifier and by the verifier to reconstruct the VOLE correlation inputs from the GGM tree.

[Algorithm 4.7 VOLEReconstruct](#) is used by the verifier to compute its VOLE correlation inputs \mathbf{Q} and Δ . The values of Δ is computed from the indexes of \mathbf{i}^* obtained by running the [ChallDec](#) procedure. And \mathbf{Q} is computed by reconstructing the committed seeds from the GGM tree with the help of [VC.Reconstruct](#).

Algorithm 4.4: ConvertToVOLE $\left(N_{e^*}, (\text{seed}_{e^*,i})_{i \in [N_{e^*}]}, \text{salt}, \mu :: \hat{\ell}\right)$
Public information and inputs
<p>Inputs: A number of parties $N_{e^*} = 2^\kappa$, a tuple of N_{e^*} seeds $(\text{seed}_{e^*,i})_{i \in [N_{e^*}]} \in (\{0,1\}^\lambda)^{N_{e^*}}$ for some fixed $e^* \in [\tau]$, a $\text{salt} \in \{0,1\}^{2\lambda}$, and $\hat{\ell} \in \mathbb{N}$ denoting the number of VOLE correlations output by the function. Recall that $\hat{\ell} := \ell_{\text{VOLEHashMask}} + \ell + \ell_{\text{CZMask}}$.</p>
Output
<p>Outputs $\hat{\ell}$ VOLE correlations $(\mathbf{u}_k, \mathbf{v}_k) \in \mathbb{F}_2 \times \mathbb{F}_2^\mu$ for $k \in [\hat{\ell}]$. All these VOLE correlations can be seen as linear (degree-1) polynomials, $[[u_k]] = f_{u_k}(X) := u_k X + v_k \in \mathbb{F}_2^\mu[X]$.</p> <hr style="border-top: 1px dashed black;"/> <pre> 1: $\kappa := \log_2(N_{e^*})$ 2: if $\text{seed}_0 = \perp$: 3: $\mathbf{r}_{0,0} := 0^{\hat{\ell}}$ 4: else : 5: $\mathbf{r}_{0,0} := \text{PRG}_2(\text{salt}, \text{seed}_{e^*,0} :: \hat{\ell})$ 6: for $i \in [1, N - 1]$: 7: $\mathbf{r}_{0,i} := \text{PRG}_2(\text{salt}, \text{seed}_{e^*,i} :: \hat{\ell})$ 8: endfor 9: $\mathbf{v}_0 = \dots = \mathbf{v}_{\kappa-1} := 0^{\hat{\ell}}$ 10: for $j \in [\kappa]$: 11: for $i \in [\frac{N_{e^*}}{2^j+1}]$: 12: $\mathbf{v}_j := \mathbf{v}_j \oplus \mathbf{r}_{j,2i+1}$ 13: $\mathbf{r}_{j+1,i} := \mathbf{r}_{j,2i} \oplus \mathbf{r}_{j,2i+1}$ 14: endfor 15: endfor 16: $\mathbf{u} := \mathbf{r}_{\kappa,0}$ 17: if $\mu > \kappa$: 18: for $i \in [\mu - \kappa]$: 19: $\mathbf{v}_{\kappa+i} = 0^{\hat{\ell}}$ 20: endfor 21: return $(\mathbf{u}, \mathbf{v}_0, \dots, \mathbf{v}_{\kappa-1}, \mathbf{v}_\kappa, \dots, \mathbf{v}_{\mu-1})$ 22: else : 23: return $(\mathbf{u}, \mathbf{v}_0, \dots, \mathbf{v}_{\kappa-1})$ // This will happen only if $\mu = \kappa$. </pre>

Algorithm 4.5: VOLECommit($\text{mseed}, \text{salt} :: \hat{\ell}$)
Public information and inputs
Public information: A number of iterations τ , a number of parties $N := \sum_{e=0}^{\tau-1} N_e$, $k_e := \log_2(N_e) \in \{\kappa_0, \kappa_1\}$, $\rho := \mu_0 \tau'_0 + \mu_1 \tau'_1$ Prover's input: A master seed $\text{mseed} \in \{0, 1\}^\lambda$ and a salt $\in \{0, 1\}^{2\lambda}$, a length $\hat{\ell} \in \mathbb{N}$
Output
Prover's output: A commitment $h_{\text{com}} \in \{0, 1\}^{2\lambda}$, a decommitment auxiliary variable decom , VOLE corrections $(\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1})$, VOLE correlation secrets $\mathbf{u} \in \mathbb{F}_2^{\hat{\ell}}$, VOLE correlation \mathbf{v} -vectors $\mathbf{V} \in \mathbb{F}_2^{\hat{\ell} \times \rho}$
<pre> 1: $(\{\text{seed}_{e,i}\}_{\tau,N}, h_{\text{com}}, \text{decom} := (\text{nodes}, \{\text{com}_{e,i}\}_{\tau,N})) := \text{VC.Commit}(\text{mseed}, \text{salt})$ 2: for $e \in [\tau]$: 3: $(\mathbf{u}_e, \mathbf{v}_{e,0}, \dots, \mathbf{v}_{e,\mu_e-1}) := \text{ConvertToVOLE}(N_e, (\text{seed}_{e,i})_{i \in [N_e]}, \text{salt}, \mu_e :: \hat{\ell})$ 4: $\mathbf{V}_e := [\mathbf{v}_{e,0} \cdots \mathbf{v}_{e,\mu_e-1}] \in \mathbb{F}_2^{\hat{\ell} \times \mu_e}$ 5: endfor 6: $\mathbf{V} := [\mathbf{V}_0 \cdots \mathbf{V}_{\tau-1}] \in \mathbb{F}_2^{\hat{\ell} \times \rho}$ 7: $\mathbf{u} := \mathbf{u}_0$ 8: for $e \in [1, \tau-1]$: 9: $\mathbf{c}_e := \mathbf{u} \oplus \mathbf{u}_e$ 10: endfor 11: return $(h_{\text{com}}, \text{decom}, \mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}, \mathbf{u}, \mathbf{V})$.</pre>

Algorithm 4.6: ChallDec(ch)
Public information and inputs
Public information: A number of iterations τ , a number of parties $N = \sum_{e=1}^{\tau} N_e$, $\kappa_e = \log_2(N_e)$ Inputs: A challenge $\text{ch} \in \{0, 1\}^{\tau_0 \kappa_0 + \tau_1 \kappa_1}$
Output
An ordered index set $\mathbf{i}^* := (i_0^*, \dots, i_{\tau-1}^*) \in [N_0] \times [N_1] \times \cdots \times [N_{\tau-1}]$ of size τ .
<pre> 1: $\text{lo} := 0$ 2: for $e \in [\tau]$: 3: $i_e^* := \text{NumRec}(\kappa_e, \text{ch}[\text{lo} : \text{lo} + \kappa_e - 1])$ 4: $\text{lo} := \text{lo} + \kappa_e$ 5: endfor 6: return $\mathbf{i}^* := (i_0^*, \dots, i_{\tau-1}^*)$</pre>

Algorithm 4.7: VOLEReconstruct $\left(\mathbf{i}^*, \text{pdecom}, \{\text{com}_{e, \mathbf{i}^*}\}_{\tau, \tau}, \text{salt}\right)$	
Public information and inputs	
<p>Public information: A number of iterations τ, a number of parties $N = \sum_{e=0}^{\tau-1} N_e$.</p> <p>Verifier's input: An ordered list $\mathbf{i}^* \in [N_0] \times [N_1] \times \dots \times [N_{\tau-1}]$ of τ indexes corresponding to hidden leaves, a sibling path pdecom, unopened commitments $\{\text{com}_{e, \mathbf{i}^*}\}_{\tau, \tau}$, and a salt $\in \{0, 1\}^{2\lambda}$.</p>	
Output	
<p>A commitment $h_{\text{com}} \in \{0, 1\}^{2\lambda}$, and verifier's VOLE correlation inputs $\mathbf{Q}' := [\mathbf{Q}'_0 \dots \mathbf{Q}'_{\tau-1}] \in \mathbb{F}_2^{\hat{\ell} \times \rho}$</p> <hr/> <pre> 1: out := VC.Reconstruct $\left(\mathbf{i}^*, \text{pdecom}, \{\text{com}_{e, \mathbf{i}^*}\}_{\tau, \tau}, \text{salt}\right)$ 2: if out = \perp : 3: return \perp 4: else : 5: Parse out as out := $(h_{\text{com}}, \{\text{seed}_{e, \mathbf{i}^*}\}_{\tau, N})$. 6: for $e \in [\tau]$ 7: $\Delta_e := \mathbf{i}^*[e]$ 8: for $i \in [N_e]$: $\text{seed}'_{e, i} := \text{seed}_{e, i \oplus \Delta_e}$ // permute $\text{seed}_{e, i}$ by using Δ_e 9: $(\mathbf{u}'_e, \mathbf{q}_{e, 0}, \dots, \mathbf{q}_{e, \mu_e - 1}) := \text{ConvertToVOLE}(N_e, (\text{seed}'_{e, i})_{i \in [N_e]}, \text{salt}, \mu_e :: \hat{\ell})$ 10: $\mathbf{Q}'_e := [\mathbf{q}_{e, 0} \dots \mathbf{q}_{e, \mu_e - 1}] \in \mathbb{F}_2^{\hat{\ell} \times \mu_e}$ 11: endfor 12: return $(h_{\text{com}}, \mathbf{Q}'_0, \dots, \mathbf{Q}'_{\tau-1})$. </pre>	

Ensuring VOLE consistency. It is crucial for the verifier to ensure that the correction values \mathbf{c}_e sent by the prover are consistent with the committed VOLE correlation input \mathbf{u}_0 . The verifier ensures this by asking the prover to compute a random linear universal hash. In this section, we explain this process in detail.³ We begin by defining family of linear universal hash functions, since it will be used to conduct the consistency checks.

Definition 4.1 (Linear universal hash functions). *A family of linear hash functions is a family of matrices $\mathcal{H} \subseteq \mathbb{F}_q^{r \times n}$. The family is ε -almost universal, if*

³ This technique is independent of the underlying PoK scheme or signature scheme since this generically helps the verifier to check the consistency of the VOLE correlations committed by the prover. Therefore, in PERK we use the exact same techniques and algorithm for these checks as those implemented in FAEST [BBD⁺23b]. Due to this reason we simply (re)state important definitions, propositions, and lemmas which are essentially same as those presented in [BBD⁺23b].

for any $\mathbf{x} \in \mathbb{F}_q^n \setminus \{\mathbf{0}\}$,

$$\Pr[\mathbf{H}\mathbf{x} = \mathbf{0} : \mathbf{H} \xleftarrow{\$} \mathcal{H}] \leq \varepsilon.$$

The family is ε -almost uniform, if for any $\mathbf{x} \in \mathbb{F}_q^n \setminus \{\mathbf{0}\}$ and for any $\mathbf{v} \in \mathbb{F}_q^r$,

$$\Pr[\mathbf{H}\mathbf{x} = \mathbf{v} : \mathbf{H} \xleftarrow{\$} \mathcal{H}] \leq \varepsilon.$$

In order to preserve the zero-knowledge property when the hash outputs are shared with the verifier, the hashes used in PERK must also satisfy the hiding property given below.

Definition 4.2. A matrix $\mathbf{H} \in \mathbb{F}_q^{r \times (h+n)}$ is \mathbb{F}_q^n -hiding if for $\mathbf{v}[0, h-1] \xleftarrow{\$} \mathbb{F}_q^h$ the distribution of $\mathbf{H}\mathbf{v}$ is independent from $\mathbf{v}[h, n+h-1]$. A hash family $\mathcal{H} \subseteq \mathbb{F}_q^{r \times (h+n)}$ is \mathbb{F}_q^n -hiding if every $\mathbf{H} \in \mathcal{H}$ is \mathbb{F}_q^n -hiding.

Proposition 4.1. Let $\mathcal{H} \subseteq \mathbb{F}_q^{r \times n}$ be ε -almost uniform hash family. Let $\mathcal{H}' \subseteq \mathbb{F}_q^{r \times (r+n)}$ be the family $\{[\mathbf{I}_r \ \mathbf{H}] : \mathbf{H} \in \mathcal{H}\}$, where \mathbf{I}_r is the $r \times r$ identity matrix. Then, (a) \mathcal{H}' is ε -almost universal, and (b) \mathcal{H}' is \mathbb{F}_q^n -hiding.

Proof. Let $\mathbf{x} := \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \end{bmatrix}$ be non-zero, with $\mathbf{x}_0 \in \mathbb{F}_q^r$ and $\mathbf{x}_1 \in \mathbb{F}_q^n$. If $\mathbf{H}' \in \mathcal{H}'$ then, $\mathbf{H}'\mathbf{x} = \mathbf{0}$ implies that $-\mathbf{x}_0 = \mathbf{H}\mathbf{x}_1$. Since \mathbf{x}_0 and \mathbf{x}_1 cannot be equal to zero simultaneously (because \mathbf{x} is non-zero), this implies that $\mathbf{x}_1 \neq \mathbf{0}$. Therefore from ε -almost uniform property of \mathcal{H} , we can conclude that \mathcal{H}' is ε -almost universal. The hiding property of \mathcal{H}' holds because when the first r elements (\mathbf{x}_0) of the input are chosen uniformly at random then they perfectly mask the rest of the output component $\mathbf{H}\mathbf{x}_1$. \square

Standard constructions of linear universal hash families. Following [BBD⁺23b] we also use the matrix hash family and polynomial-based hash family as building blocks of our VOLE consistency checks [CW79, BJKS94]. The matrix hash family $\mathcal{H} = \mathbb{F}_q^{r \times n}$ is q^{-r} -almost uniform. In polynomial-based hash, the input $\mathbf{x} \in \mathbb{F}_q^n$ is seen as the coefficients of a polynomial with degree $\leq n-1$. Sampling a hash function is implemented by evaluating such a polynomial at a uniform random point in \mathbb{F}_q . Since the polynomial has at most $n-1$ roots, the polynomial hash family is $\frac{n-1}{q}$ -almost universal. If the random point is chosen from a set S with cardinality $|S|$, then the polynomial hash family is $\frac{n-1}{|S|}$ -almost universal.

Composition and truncation of hashes. We also recall the properties of composition and truncation of hashes originally proved by the authors of [Sti92, Roy22, BBD⁺23b]. These properties will be useful in proving that the prover can successfully bypass the consistency check for VOLE correlations with extremely low probability.

Proposition 4.2. *Let \mathcal{H} and \mathcal{H}' be two independent ε and ε' -almost universal hash families respectively. Then the concatenation $\left\{ \begin{bmatrix} \mathbf{H} \\ \mathbf{H}' \end{bmatrix} : \mathbf{H} \in \mathcal{H}, \mathbf{H}' \in \mathcal{H}' \right\}$ is $\varepsilon\varepsilon'$ -almost universal.*

Proof. This holds true because of the independence of \mathcal{H} and \mathcal{H}' . \square

Proposition 4.3. *Let $\mathcal{H} \subseteq \mathbb{F}_q^{r' \times n}$ be ε -almost universal and $\mathcal{H}' \subseteq \mathbb{F}_q^{r \times r'}$ be ε' -almost uniform. Then the product $\{\mathbf{H}'\mathbf{H} : \mathbf{H} \in \mathcal{H}, \mathbf{H}' \in \mathcal{H}'\}$ is $(\varepsilon + \varepsilon')$ -almost uniform.*

Proof. Let $\mathbf{x} \in \mathbb{F}_q^n \setminus \{\mathbf{0}\}$. Then $\Pr[\mathbf{x}' \neq \mathbf{0} : \mathbf{x}' = \mathbf{H}\mathbf{x}] \geq 1 - \varepsilon$ since \mathcal{H} is ε -almost universal. For $\mathbf{x}' \neq \mathbf{0}$, $\Pr[\mathbf{H}'\mathbf{x}' \neq \mathbf{v} : \mathbf{v} \in \mathbb{F}_q^r] \geq 1 - \varepsilon'$ as \mathcal{H}' is ε' -almost uniform. Therefore $\Pr[\mathbf{H}'\mathbf{H}\mathbf{x} \neq \mathbf{v}] \geq (1 - \varepsilon)(1 - \varepsilon') \geq 1 - \varepsilon - \varepsilon'$, which implies that the product is $(\varepsilon + \varepsilon')$ -almost uniform. \square

Proposition 4.4. *Let $\delta \in \mathbb{N}$ and $\mathcal{H} \subseteq \mathbb{F}_q^{r \times n}$ be ε -almost uniform hash family. Then, the truncated family $\{\mathbf{H}_{0,r-\delta-1} : \mathbf{H} \in \mathcal{H}\}$, where $\mathbf{H}_{0,r-\delta-1}$ denotes the first $(r - \delta)$ rows of \mathbf{H} , is εq^δ -almost uniform.*

Proof. For each $\mathbf{H} \in \mathcal{H}$, let $\mathbf{H} := \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix}$ where $\mathbf{H}_0 \in \mathbb{F}_q^{(r-\delta) \times n}$ and $\mathbf{H}_1 \in \mathbb{F}_q^{\delta \times n}$.

Let $\mathbf{x} \in \mathbb{F}_q^n$ be a non-zero vector, and $\mathbf{y} := \begin{bmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \end{bmatrix} \in \mathbb{F}_q^r$. If $\mathbf{H} \xleftarrow{\$} \mathcal{H}$, then $\Pr[\mathbf{H}\mathbf{x} = \mathbf{y}] \leq \varepsilon$. We can then apply conditional probability to obtain,

$$\begin{aligned} \Pr[\mathbf{H}\mathbf{x} = \mathbf{y}] &\leq \varepsilon \\ \Pr[\mathbf{H}_0\mathbf{x}_0 = \mathbf{y}_0 \wedge \mathbf{H}_1\mathbf{x}_1 = \mathbf{y}_1] &\leq \varepsilon \\ \Pr[\mathbf{H}_0\mathbf{x}_0 = \mathbf{y}_0] \cdot \Pr[\mathbf{H}_1\mathbf{x}_1 = \mathbf{y}_1 \mid \mathbf{H}_0\mathbf{x}_0 = \mathbf{y}_0] &\leq \varepsilon \\ \Pr[\mathbf{H}_0\mathbf{x}_0 = \mathbf{y}_0] &\leq \varepsilon \cdot \Pr[\mathbf{H}_1\mathbf{x}_1 = \mathbf{y}_1 \mid \mathbf{H}_0\mathbf{x}_0 = \mathbf{y}_0]^{-1} \\ &\leq \varepsilon q^\delta \end{aligned}$$

where the final inequality comes from fixing a $\mathbf{y}_1 \in \mathbb{F}_q^\delta$, that maximizes $p := \Pr[\mathbf{H}_1\mathbf{x}_1 = \mathbf{y}_1 \mid \mathbf{H}_0\mathbf{x}_0 = \mathbf{y}_0]$, which implies p is at least $q^{-\delta}$. \square

VOLE universal hash. In order to verify the consistency of VOLE correlation inputs in $\mathbb{F}_2^{\hat{\ell}}$, we need a hash family that is linear over \mathbb{F}_2 . Also recall that, $\hat{\ell} := \ell_{\text{VOLEHashMask}} + \ell + \ell_{\text{CZMask}}$, where $\ell_{\text{VOLEHashMask}} := \lambda + B$, and $B = 16$ is a parameter chosen for security.⁴ To compute the hash, we start by mapping the seed seed into $(r_0, r_1, r_2, r_3, s, t) \in \mathbb{F}_{2^\lambda}^5 \times \mathbb{F}_{2^{64}}$. The input $\mathbf{x} \in \mathbb{F}_2^{\hat{\ell}}$ is split into

⁴ Similar to [BBD⁺23b] PERK actually calls **VOLEHash** on inputs represented as $\hat{\ell} \times \rho$ matrix, which is translated into computing the hash on each column separately, with the same seed.

$(\mathbf{x}_0, \mathbf{x}_1)$, where $\mathbf{x}_1 \in \mathbb{F}_2^{\ell + \ell_{\text{CZMask}}}$, and then \mathbf{x}_1 is parsed twice, first as a vector $\hat{\mathbf{y}}$ of \mathbb{F}_{2^λ} elements, and then as a vector $\bar{\mathbf{y}}$ of $\mathbb{F}_{2^{64}}$ elements.⁵ Then compute,

$$h_0 := \hat{y}_0 s^{\frac{\hat{\ell}}{\lambda} - 1} + \hat{y}_1 s^{\frac{\hat{\ell}}{\lambda} - 2} + \cdots + \hat{y}_{\frac{\hat{\ell}}{\lambda} - 2} s + \hat{y}_{\frac{\hat{\ell}}{\lambda} - 1} \in \mathbb{F}_{2^\lambda},$$

$$h_1 := \bar{y}_0 t^{\frac{\hat{\ell}}{64} - 1} + \bar{y}_1 t^{\frac{\hat{\ell}}{64} - 2} + \cdots + \bar{y}_{\frac{\hat{\ell}}{64} - 2} t + \bar{y}_{\frac{\hat{\ell}}{64} - 1} \in \mathbb{F}_{2^{64}}$$

Viewing h_1 as an element of \mathbb{F}_{2^λ} (by padding zeros), the hash is then computed in \mathbb{F}_{2^λ} as,

$$\begin{bmatrix} h_2 \\ h_3 \end{bmatrix} := \begin{bmatrix} r_0 & r_1 \\ r_2 & r_3 \end{bmatrix} \begin{bmatrix} h_0 \\ h_1 \end{bmatrix}$$

Finally we take the first $\ell_{\text{VOLEHashMask}}$ (i.e. $\lambda + B$) bits of the concatenation of the field elements h_2 and h_3 , and XOR it with \mathbf{x}_0 . We argue security of this construction below (same as [BBD⁺23b]). Like [BBD⁺23b], we also aim for $\varepsilon := 2^{-\ell_{\text{VOLEHashMask}}} = 2^{-\lambda - B}$, with $B = 16$, in order to compensate for $\binom{\tau}{2}$ security loss shown by [BBD⁺23c].

Lemma 4.1. *VOLEHash is an ε_v -almost universal hash family in $\mathbb{F}_2^{\ell_{\text{VOLEHashMask}} \times \hat{\ell}}$, for $\varepsilon_v = 2^{-\ell_{\text{VOLEHashMask}}} (1 + 2^{B-50})$, if $\hat{\ell} \leq 2^{13}$. Furthermore, VOLEHash is $\mathbb{F}_2^{\ell + \ell_{\text{CZMask}}}$ -hiding.*

Proof. We show ε_v -almost uniform property of the hash that outputs the first $\ell_{\text{VOLEHashMask}}$ (i.e. $\lambda + B$) bits of (h_2, h_3) , that is without adding \mathbf{x}_0 . By Proposition 4.1, this implies the hiding and ε_v -almost universal property of the final hash. The first part of the hash which computes h_0, h_1 , is a concatenation of two polynomial hashes, over either \mathbb{F}_{2^λ} or $\mathbb{F}_{2^{64}}$. These are ε -almost universal with $\varepsilon = \frac{d}{|\mathbb{F}|}$, where d is the degree of the polynomial and \mathbb{F} is the field, and we have $d \leq \frac{\hat{\ell}}{64}$. Both these hashes are \mathbb{F}_2 -linear, as the binary field multiplication is bilinear over \mathbb{F}_2 . Thus, applying Proposition 4.2 we get that the concatenation of the two polynomial hashes is then ε_0 -almost universal with $\varepsilon_0 \leq \frac{\hat{\ell}^2}{2^{\lambda+76}}$. Therefore, for $\hat{\ell} \leq 2^{13}$, we have $\varepsilon_0 \leq 2^{-\lambda-50}$.

The second part of the hash starts with a 2×2 matrix hash, which is $2^{-2\lambda}$ -almost uniform. After the truncation, the result is ε_1 -almost uniform for $\varepsilon_1 = 2^{-\ell_{\text{VOLEHashMask}}}$ due to the Proposition 4.4. The final combined hash is a product of these two parts, so from Proposition 4.3 and summing the probabilities, we get that for all $\hat{\ell} \leq 2^{13}$, the hash is ε_v -almost uniform for $\varepsilon_v = \varepsilon_0 + \varepsilon_1 \leq 2^{-\ell_{\text{VOLEHashMask}}} (1 + 2^{B-50})$. \square

Following algorithms help prove consistency of the VOLE correlations.

⁵ In order to allow for better parsing in PERK we swap the order of \mathbf{x}_0 and \mathbf{x}_1 from [BBD⁺23b]. That is \mathbf{x}_0 serves as a mask in PERK, where as [BBD⁺23b] uses \mathbf{x}_1 as a mask.

Algorithm 4.8: $\text{VOLEHash}(\text{seed}, \mathbf{x} :: \ell_{\text{VOLEHashMask}})$
Public information and inputs
Inputs: A seed, represented as a tuple $(r_0, r_1, r_2, r_3, s, t) \in \{0, 1\}^{5\lambda+64}$, and a vector to be hashed $\mathbf{x} \in \{0, 1\}^\ell$, represented as a pair $(\mathbf{x}_0, \mathbf{x}_1) \in \{0, 1\}^{\ell_{\text{VOLEHashMask}}} \times \{0, 1\}^{\ell + \ell_{\text{CZMask}}}$
Output
A VOLE hash $\mathbf{h} \in \{0, 1\}^{\ell_{\text{VOLEHashMask}}}$
<hr/> <pre> 1 : $r_i := \text{ToField}(r_i, \lambda)$ for $i \in [4]$ 2 : $s := \text{ToField}(s, \lambda)$ 3 : $t := \text{ToField}(t, 64)$ 4 : $\ell' := \lceil (\ell + \ell_{\text{CZMask}}) / \lambda \rceil$ 5 : $\mathbf{x}_1 := \mathbf{x}_1 \ 0^{\ell' - (\ell + \ell_{\text{CZMask}})}$ (pad to multiple of λ) 6 : $\hat{\mathbf{y}} := \text{ToField}(\mathbf{x}_1, \lambda)$ 7 : $\bar{\mathbf{y}} := \text{ToField}(\mathbf{x}_1, 64)$ 8 : $h_0 := \sum_{i=0}^{\ell'/\lambda-1} s^{\ell'/\lambda-1-i} \hat{\mathbf{y}}[i]$ 9 : $h_1 := \sum_{i=0}^{\ell'/64-1} s^{\ell'/64-1-i} \bar{\mathbf{y}}[i]$ 10 : $h'_1 := \text{ToField}(\text{ToBits}(h_1) \ 0^{\lambda-64}, \lambda)$ 11 : $(h_2, h_3) := (r_0 h_0 + r_1 h'_1, r_2 h_0 + r_3 h'_1)$ 12 : $\mathbf{h} := (\text{ToBits}(h_2) \ \text{ToBits}(h_3)[0..\ell_{\text{VOLEHashMask}} - 1]) \oplus \mathbf{x}_0$ 13 : return \mathbf{h} </pre>

4.5 Proof of Knowledge functions

In this section, we present the various algorithms used by the prover and by the verifier to prove (respectively verify) the knowledge of the secret permutation which serves as solution to the PKP instance defined by the public key. This is achieved with help of VOLE correlations created as described in [Section 4.4](#), and then using such VOLE correlations to compute linear functions over polynomials to ascertain that different constraints related to the elementary vector structure, and the satisfiability of the PKP solution are fulfilled by these polynomials.

The prover begins by proving the knowledge of elementary vectors of length 4 or 2, and then generates polynomials corresponding to the rows of the secret permutation matrix by computing the tensor products of these smaller elementary vectors. In order to prove the elementary structure of the length 4 (or length 2) vector, the prover needs to show that following constraints are obeyed by the underlying vector:

1. For a vector $e := [e_0, e_1, e_2, e_3] \in \mathbb{F}_2^4$ the product $e_0 \cdot e_1 = 0$ and $e_2 \cdot e_3 = 0$ simultaneously. In case of elementary vector $e := [e_0, e_1] \in \mathbb{F}_2^2$ this is achieved simply by showing $e_0 \cdot e_1 = 0$.
2. Additionally, the prover should also prove that $e_0 \oplus e_1 \oplus e_2 \oplus e_3 = 1$ (respectively $e_0 \oplus e_1 = 1$).

Note that proving constraint 2, actually allows for optimization leading to saving of a single bit per elementary vector. The prover simply sends 3 bits (or 1 bit) and the remaining bit of the elementary is computed by the verifier as: $e_3 := 1 \oplus e_0 \oplus e_1 \oplus e_2$ (respectively $e_1 := 1 \oplus e_0$). The algorithms [algorithm 4.9 CompWit](#) and [algorithm 4.10 ExpWit](#) describe this process, these algorithms are used by the prover and the verifier respectively.

Witness Compression and Expansion

Algorithm 4.9: CompWit(w')

Public information and inputs

Public information: Compress input witness w' into shorter witness w by dropping bits.

```

1 : Initialize  $w := \varepsilon$ 
2 : if  $\text{len}(w') = 12$ :
    // Parse  $w'$  as  $w' := w'_0 || w'_1 || w'_2$ 
    // where,
    //  $w'_i := w'_{i,0} || w'_{i,1} || w'_{i,2} || w'_{i,3}$ .
3 :   for  $i \in [3]$ :
4 :     for  $j \in [3]$ :
5 :        $w_{i,j} := w'_{i,j}$ 
6 :     endfor
7 :   endif
8 :   assert ( $\text{len}(w) = 9$ )
9 :   return  $w$ 
10 : elseif  $\text{len}(w') = 14$ :
    // Parse  $w'$  as
    //  $w' := w'_0 || w'_1 || w'_2 || w'_3$ 
    // with  $w'_i$  defined as in line 2
    // for  $i \in [3]$  and
    //  $w'_3 := w'_{3,0} || w'_{3,1}$ .
11 :    $w := \text{CompWit}(w'_0 || w'_1 || w'_2)$ 
12 :    $w := w || w'_{3,0}$ 
13 :   assert ( $\text{len}(w) = 10$ )
14 :   return  $w$ 
15 : else :
16 :   return  $\perp$ .
17 : endif

```

Algorithm 4.10: ExpWit(w)

Public information and inputs

Public information: Expand compressed input witness w into witness w' .

```

1 : Initialize  $w' := \varepsilon$ 
2 : if  $\text{len}(w) = 9$ :
    // Parse  $w$  as
    //  $w := w_{0,0} || \dots || w_{i,j} || \dots || w_{2,2}$ .
    // For  $i \in [3]$  and  $j \in [3]$ .
3 :   for  $i \in [3]$ :
4 :      $w'_{i,3} := 1$ 
5 :     for  $j \in [3]$ :
6 :        $w'_{i,j} := w_{i,j}$ 
7 :        $w'_{i,3} := w'_{i,3} \oplus w_{i,j}$ 
8 :      $w' := w' || w'_{i,j}$ 
9 :   endifor
10 :    $w' := w' || w'_{i,3}$ 
11 : endifor
12 : assert ( $\text{len}(w') = 12$ )
13 : return  $w'$ 
14 : elseif  $\text{len}(w) = 10$ :
    // Parse  $w$  as
    //  $w := \bar{w} || w_{3,0}$ , and
    //  $\bar{w} := w_{0,0} || \dots || w_{i,j} || \dots || w_{2,2}$ .
    // For  $i \in [3]$  and  $j \in [3]$ .
15 :    $w' := \text{ExpWit}(\bar{w})$ 
16 :    $w' := w' || w_{3,0} || (1 \oplus w_{3,0})$ 
17 :   assert ( $\text{len}(w') = 14$ )
18 :   return  $w'$ 
19 : else :
20 :   return  $\perp$ .
21 : endif

```

Prover. In this section we present all the algorithms that will be used by the prover (signer) to produce the proof of knowledge. The prover begins by encoding each row of the secret permutation matrix into smaller elementary vectors. Algorithms [EncPosArrayToWitness](#) and [PosToWitness](#) together take an input position corresponding to the non-zero element in a row of permutation matrix,

and output its corresponding unique witness encoded as blocks of elementary vectors. Let e_i be an elementary vector of length 4, generated by shifting the bit-string ‘0001’ to left by i positions. Therefore, $e_0 := ‘0001’$, $e_1 := ‘0010’$, $e_2 := ‘0100’$, and $e_3 := ‘1000’$. Also, let e'_b be an elementary vector of length 2, generated by shifting the bit-string ‘01’ to left by either b positions for $b \in \{0, 1\}$. Therefore, $e'_0 := ‘01’$ and $e'_1 := ‘10’$.

Algorithm 4.11 EncPosArrayToWitness takes the unique encoded position array $[i, j, k]$ corresponding to some position $\text{pos} \in [64]$ generated from **EncodeNum-64**, and outputs its corresponding unique witness $\mathbf{w}' := e_k || e_j || e_i \in \mathbb{F}_2^{12}$.

Algorithm 4.12 PosToWitness on an input position pos , first computes the unique encoded position array encPosArray corresponding to pos by calling **EncodeNum**. If encPosArray contains exactly 3 elements then, **PosToWitness** outputs corresponding unique witness for pos as $\mathbf{w}' := e_k || e_j || e_i \in \mathbb{F}_2^{12}$ by subsequently calling **EncPosArrayToWitness**. Otherwise if $\text{encPosArray} := [i, j, k, b]$ contains exactly 4 elements then, it computes $e_k || e_j || e_i$ as described above. It then outputs corresponding unique witness for pos as $\mathbf{w}' := e'_b || e_k || e_j || e_i \in \mathbb{F}_2^{14}$.

Next, the prover creates the VOLE correlations that will be used in proof generation, by embedding the witness (now represented as elementary vectors). This is achieved by **Algorithm 4.13 EmbedWitnessBlock** which takes the witness generated by **PosToWitness** as input and embeds it block-by-block inside random VOLE correlation also given as inputs. The output of this algorithm are the VOLE correlations $[[\beta'_i]]$ corresponding to the elementary blocks of witness. **Algorithm 4.14 EmbedWitness** aggregates the VOLE correlations embedding the witness from individual elementary vectors of lengths 4 and 2, and outputs VOLE correlations $[[\beta']]$ corresponding to the aggregation of 3 elementary vectors of lengths 4 (and in case of L3 and L5 parameters, another elementary vector of length 2).

Once the prover possesses the VOLE correlations (linear polynomials) corresponding to the elementary vector entries, it then computes the degree- d VOLE correlations (degree- d polynomials) corresponding to each individual row of the secret permutation matrix with help of **Algorithm 4.16 TensorProductToElementaryVector**, which internally calls **Algorithm 4.15 TensorProduct** to compute the tensor product between two blocks. At this stage the prover has computed n polynomials per row each of degree d , $[[\mathbf{z}]]^{(d)} := ([[z_0]]^{(d)}, \dots, [[z_{n-1}]]^{(d)})$ which are the VOLE correlations corresponding to the rows of the secret permutation matrix, viewed as an elementary vector of length n . The **Algorithm 4.17 VOLE-ElementaryVector** produces the VOLE correlations $[[\mathbf{z}]]^{(d)}$ along with intermediate values $[[\beta']]$, and masked witness \mathbf{t} which will be sent to the verifier.

After obtaining the VOLE correlations for each of the n rows by running **Algorithm 4.17 VOLE-ElementaryVector** n times, the prover further computes extra n degree- d polynomials, $[[\text{ColCheck}_j]]^{(d)}$ which ensure that each column of secret matrix adds upto exactly 1. This is described in **Algorithm 4.18 VOLE-**

Permutation. The check that columns add upto exactly 1 along with the elementary structure of the individual rows is sufficient to prove the permutation structure of the secret matrix.

The prover proves the elementary structure by computing the polynomials which have leading coefficient equal to 0 if and only if constraint 1 mentioned above ($e_0 \cdot e_1 = 0$ and $e_2 \cdot e_3 = 0$) is satisfied using [Algorithm 4.19 Check-ElementaryBlock](#), and [Algorithm 4.20 Check-ElementaryVector](#).

So far we have seen that, the prover has generated many degree- d polynomials, which should all have leading coefficients equal to 0. In order to verify, this the prover needs to send these polynomials to the verifier which can then evaluate them at a random point of its choice and check if the all the leading coefficients are equal to zero as expected. However, sending so many degree- d polynomials is inefficient, the prover can instead send a single degree- d polynomial by taking a random linear combination of all the polynomials, where the coefficients of the random linear combination are chosen by the verifier.

Note that, in order to convince the verifier the prover still needs to send a degree- d polynomial to the verifier whose leading coefficient is supposed to be zero. However, since all these polynomials are computed by taking tensor products and other linear functions of secret witness (embedded in the VOLE correlations), each coefficient of this polynomial obtained by the prover contains information about the secret. As the prover only needs to reveal that leading coefficient is zero, it should add a masking polynomial to blind the remaining coefficients. The [Algorithm 4.21 CheckZero](#) achieves this and outputs a masked polynomial with leading coefficient equal to zero which can be used by the verifier.

Finally, the [Algorithm 4.22 Check-PKP](#) puts all of the checks for checking the elementary structure of rows (blocks), column sums equaling to 1, and satisfiability of PKP equation together by computing degree- d polynomials with zero as leading coefficients if and only if these constraints are satisfied. These polynomials are then merged together into a single polynomial by computing (verifier dictated) random linear combination, which is then masked with the help of [Algorithm 4.21 CheckZero](#). As a result, the prover should send the masked witness \mathbf{t} obtained from [Algorithm 4.18 VOLE-Permutation](#) along with the masked polynomial proof output by [Algorithm 4.22 Check-PKP](#) to the verifier.

Algorithm 4.11: EncPosArrayToWitness(encPosArray)

Public information and inputs

Public information: Given encoded position array encPosArray of length 3 with all input elements in $[4]$, outputs its corresponding unique witness $\mathbf{w}' \in \{0,1\}^{12}$ with hamming weight 3.

```
1: Initialize  $\mathbf{w}' := \varepsilon$  // Empty string
2: if  $\text{len}(\text{encPosArray}) = 3$  :
    // Parse  $\text{encPosArray}$  as  $\text{encPosArray} := [i, j, k]$  with  $i, j, k \in [4]$ , else output error.
3:   for  $\text{index} \in [3]$ :
4:      $\mathbf{w}'_{\text{index}} := \text{LeftShift}('0001', \text{encPosArray}[\text{index}])$ 
5:      $\mathbf{w}' := \mathbf{w}' || \mathbf{w}'_{\text{index}}$ 
6:   endfor
7:   assert ( $\text{len}(\mathbf{w}') = 12$ )
8:   return  $\mathbf{w}'$ 
9: else :
10:  return  $\perp$ 
11: endif
```


Algorithm 4.12: PosToWitness(pos)

Public information and inputs

Public information: Given secret input pos outputs its corresponding unique witness w' .

```
1 : encPosArray := EncodeNum (pos)

   Security Level 1
2 : if len(encPosArray) = 3 :
3 :    $w' := \text{EncPosArrayToWitness}(\text{encPosArray})$ 
4 :   return  $w'$ 

   Security Levels 3 and 5
5 : elseif len(encPosArray) = 4 :
6 :    $w' := \text{EncPosArrayToWitness}(\text{encPosArray}[0 : 2])$ 
7 :    $w'_3 := \text{LeftShift}('01', \text{encPosArray}[3])$ 
8 :    $w' := w' || w'_3$ 
9 :   assert (len( $w'$ ) = 14)
10 :  return  $w'$ 
11 : else :
12 :  return  $\perp$ 
13 : endif
```

Algorithm 4.13: P.EmbedWitnessBlock($\mathbf{w}'_i, (\llbracket u_k \rrbracket)_{k \in [3]}$)
Public information and inputs
<p>Public information: Length of the witness block = 4.</p> <p>Prover's input: i^{th} secret witness block $\mathbf{w}'_i := w'_{i,0} w'_{i,1} w'_{i,2} w'_{i,3} \in \mathbb{F}_2^4$, 3 VOLE correlations $\llbracket u_k \rrbracket$ for random $(u_k, v_k) \in \mathbb{F}_2 \times \mathbb{F}_{2^\rho}$ represented as polynomials $f_{u_k}(X) = u_k X + v_k$ for $k \in [3]$.</p>
Output
<p>Prover's output: VOLE correlations $\llbracket \beta'_i \rrbracket := (\llbracket \beta'_{i,0} \rrbracket, \llbracket \beta'_{i,1} \rrbracket, \llbracket \beta'_{i,2} \rrbracket, \llbracket \beta'_{i,3} \rrbracket)$, where $\llbracket \beta'_{i,j} \rrbracket \in \mathbb{F}_2 \times \mathbb{F}_{2^\rho}$ for $j \in [4]$.</p> <hr/> <p><u>Construct VOLE correlations with witness</u></p> <ol style="list-style-type: none"> 1 : // Parse \mathbf{w}'_i as $\mathbf{w}'_i := w'_{i,0} w'_{i,1} w'_{i,2} w'_{i,3}$. 2 : for $j \in [3]$: 3 : $\beta'_{i,j}(X) := w'_{i,j}X + v_j$ // $\beta'_{i,j}$ are polynomials with coefficients in \mathbb{F}_{2^ρ}. 4 : endfor 5 : $\beta'_{i,3}(X) := w'_{i,3}X + v_0 + v_1 + v_2$ 6 : $\llbracket \beta'_i \rrbracket := (\llbracket \beta'_{i,0} \rrbracket, \llbracket \beta'_{i,1} \rrbracket, \llbracket \beta'_{i,2} \rrbracket, \llbracket \beta'_{i,3} \rrbracket)$ 7 : return $\llbracket \beta'_i \rrbracket$

Algorithm 4.14: $\text{P.EmbedWitness}(\mathbf{w}', (\llbracket u_k \rrbracket)_{k \in [\ell_{\text{row}]})$	
Public information and inputs	
<p>Public information: Length of the expanded secret witness $\mathbf{w}' := 2\ell_{\text{row}} - 6$.</p> <p>Prover's input: Secret witness $\mathbf{w}' \in \mathbb{F}_2^{(2\ell_{\text{row}}-6)}$, ℓ_{row} VOLE correlations $\llbracket u_k \rrbracket$ for random $(u_k, v_k) \in \mathbb{F}_2 \times \mathbb{F}_{2^\rho}$ represented as polynomials $f_{u_k}(X) = u_k X + v_k$ for $k \in [\ell_{\text{row}}]$.</p>	
Output	
<p>Prover's output: $(2\ell_{\text{row}} - 6)$ VOLE correlations $\llbracket \beta' \rrbracket$ with elements in $\mathbb{F}_2 \times \mathbb{F}_{2^\rho}$.</p> <hr/> <p style="text-align: center;"><u>Construct VOLE correlations with witness</u></p> <pre> 1 : if $\text{len}(\mathbf{w}') = 12$: // Security level 1 // Parse \mathbf{w}' as $\mathbf{w}' := \mathbf{w}'_0 \mathbf{w}'_1 \mathbf{w}'_2$, where, $\mathbf{w}'_i := w'_{i,0} w'_{i,1} w'_{i,2} w'_{i,3}$. 2 : for $i \in [3]$: 3 : $\llbracket \beta'_i \rrbracket := \text{P.EmbedWitnessBlock}(\mathbf{w}'_i, (\llbracket u_k \rrbracket)_{k \in [3i, 3i+2]})$ 4 : endfor 5 : $\llbracket \beta' \rrbracket := (\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket)$ 6 : return $\llbracket \beta' \rrbracket$ 7 : elseif $\text{len}(\mathbf{w}') = 14$: // Security levels 3 and 5 // Parse \mathbf{w}' as $\mathbf{w}' := \mathbf{w}'_0 \mathbf{w}'_1 \mathbf{w}'_2 w'_{3,0} w'_{3,1}$, where, \mathbf{w}'_i are as above (line 1) 8 : $(\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket) := \text{P.EmbedWitness}(\mathbf{w}'_0 \mathbf{w}'_1 \mathbf{w}'_2, (\llbracket u_k \rrbracket)_{k \in [9]})$ 9 : $\beta'_{3,0}(X) := w'_{3,0}X + v_9$ 10 : $\beta'_{3,1}(X) := w'_{3,1}X + v_9$ 11 : $\llbracket \beta'_3 \rrbracket := (\llbracket \beta'_{3,0} \rrbracket, \llbracket \beta'_{3,1} \rrbracket)$ 12 : $\llbracket \beta' \rrbracket := (\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket, \llbracket \beta'_3 \rrbracket)$ 13 : return $\llbracket \beta' \rrbracket$ 14 : else : 15 : return \perp 16 : endif </pre>	

Algorithm 4.15: P.TensorProduct($\llbracket \beta'_i \rrbracket^{(d_i)}, \llbracket \beta'_j \rrbracket^{(d_j)}, \text{num}$)
Public information and inputs
<p>Public information: Sizes n_i and n_j of two blocks of VOLE correlations $\llbracket \beta'_i \rrbracket^{(d_i)}$ and $\llbracket \beta'_j \rrbracket^{(d_j)}$ respectively. Each element of $\llbracket \beta'_i \rrbracket^{(d_i)}$ (resp. $\llbracket \beta'_j \rrbracket^{(d_j)}$) is represented as degree d_i (resp. d_j) polynomial with coefficients in \mathbb{F}_{2^ρ}.</p> <p>Prover's input: $(n_i + n_j)$ secret polynomials $\llbracket \beta'_{i,0} \rrbracket^{(d_i)}, \dots, \llbracket \beta'_{i,(n_i-1)} \rrbracket^{(d_i)}, \llbracket \beta'_{j,0} \rrbracket^{(d_j)}, \dots, \llbracket \beta'_{j,(n_j-1)} \rrbracket^{(d_j)}$.</p>
Output
<p>Prover's output: num^* polynomials, $\llbracket \zeta \rrbracket^{(d_i+d_j)} := (\llbracket \zeta_0 \rrbracket^{(d_i+d_j)}, \dots, \llbracket \zeta_{\text{num}^*-1} \rrbracket^{(d_i+d_j)})$ each of degree $d_i + d_j$ which can be seen as VOLE correlations, where $\text{num}^* := \min(n_i n_j, \text{num})$.</p> <hr/> <p style="text-align: center;"><u>Computing tensor product</u></p> <pre> 1 : counter := 0 2 : for index_j ∈ [n_j]: 3 : for index_i ∈ [n_i]: 4 : if counter ∈ [num]: 5 : $\llbracket \zeta_{\text{counter}} \rrbracket^{(d_i+d_j)} := \text{P.Multiply}(\llbracket \beta'_{j,\text{index}_j} \rrbracket^{(d_j)}, \llbracket \beta'_{i,\text{index}_i} \rrbracket^{(d_i)})$ 6 : counter := counter + 1 7 : else : 8 : break 9 : endfor 10 : endfor 11 : $\llbracket \zeta \rrbracket^{(d_i+d_j)} := (\llbracket \zeta_0 \rrbracket^{(d_i+d_j)}, \dots, \llbracket \zeta_{\text{counter}-1} \rrbracket^{(d_i+d_j)})$ 12 : return $\llbracket \zeta \rrbracket^{(d_i+d_j)}$ </pre>

Algorithm 4.16: P.TensorProductToElementaryVector($\llbracket \beta' \rrbracket, n$)
Public information and inputs
Public information: Length of the secret elementary vector n , degree $d := \lceil \log_4(n) \rceil$. Prover's input: $(2\ell_{\text{row}} - 6)$ secret degree-1 (linear) polynomials represented as $\llbracket \beta' \rrbracket$.
Output
Prover's output: n polynomials, $\llbracket \mathbf{z} \rrbracket^{(d)} := (\llbracket z_0 \rrbracket^{(d)}, \dots, \llbracket z_{n-1} \rrbracket^{(d)})$ each of degree d which can be seen as VOLE correlations corresponding to the secret elementary vector.

<p style="text-align: center;"><u>Compute elementary vector using tensor product</u></p> <pre> 1 : if len($\llbracket \beta' \rrbracket$) = 12: // Security level 1 // Parse $\llbracket \beta' \rrbracket$ as $\llbracket \beta' \rrbracket := (\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket)$. // Generating 16 degree-2 VOLE correlations in $\mathbb{F}_2 \times \mathbb{F}_{2^\rho}$. 2 : $\llbracket \zeta_{0,1} \rrbracket^{(2)} := \text{P.TensorProduct}(\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, n)$ // Generating 64 degree-3 VOLE correlations in $\mathbb{F}_2 \times \mathbb{F}_{2^\rho}$. 3 : $\llbracket \mathbf{z} \rrbracket^{(3)} := \text{P.TensorProduct}(\llbracket \zeta_{0,1} \rrbracket^{(2)}, \llbracket \beta'_2 \rrbracket, n)$ 4 : return $\llbracket \mathbf{z} \rrbracket^{(3)}$ 5 : elseif len($\llbracket \beta' \rrbracket$) = 14: // Security levels 3 and 5 // Parse $\llbracket \beta' \rrbracket$ as $\llbracket \beta' \rrbracket := (\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket, \llbracket \beta'_3 \rrbracket)$. // Generating 16 degree-2 VOLE correlations in $\mathbb{F}_2 \times \mathbb{F}_{2^\rho}$. 6 : $\llbracket \zeta_{0,1} \rrbracket^{(2)} := \text{P.TensorProduct}(\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, n)$ // Generating 64 degree-3 VOLE correlations in $\mathbb{F}_2 \times \mathbb{F}_{2^\rho}$. 7 : $\llbracket \zeta_{0,1,2} \rrbracket^{(3)} := \text{P.TensorProduct}(\llbracket \zeta_{0,1} \rrbracket^{(2)}, \llbracket \beta'_2 \rrbracket, n)$ 8 : // Generating n degree-4 VOLE correlations in $\mathbb{F}_2 \times \mathbb{F}_{2^\rho}$. 9 : $\llbracket \mathbf{z} \rrbracket^{(4)} := \text{P.TensorProduct}(\llbracket \zeta_{0,1,2} \rrbracket^{(3)}, \llbracket \beta'_3 \rrbracket, n)$ 10 : return $\llbracket \mathbf{z} \rrbracket^{(4)}$ 11 : else : 12 : return \perp 13 : endif </pre>

Algorithm 4.17: P.VOLE-ElementaryVector($\text{pos}, (\llbracket u_k \rrbracket)_{k \in [\ell_{\text{row}}]}$)
Public information and inputs
<p>Public information: Length of the secret elementary vector n, degree $d := \lceil \log_4 n \rceil$, size of the compressed secret witness $\ell_{\text{row}} := d + 6$.</p> <p>Prover's input: Secret position $\text{pos} \in [n]$, ℓ_{row} VOLE correlations $\llbracket u_k \rrbracket$ for random $(u_k, v_k) \in \mathbb{F}_2 \times \mathbb{F}_{2^d}$ represented as polynomials $f_{u_k}(X) = u_k X + v_k$ for $k \in [\ell_{\text{row}}]$.</p>
Output
<p>Prover's output: Masked compressed secret $\mathbf{t} \in \mathbb{F}_2^{\ell_{\text{row}}}$, n polynomials $\llbracket \mathbf{z} \rrbracket^{(d)} := (\llbracket z_0 \rrbracket^{(d)}, \dots, \llbracket z_{n-1} \rrbracket^{(d)})$ each of degree d which can be seen as VOLE correlations corresponding to the secret elementary vector, along with $(2\ell_{\text{row}} - 6)$ secret degree-1 (linear) polynomials represented as $\llbracket \beta' \rrbracket$.</p>
<hr style="border-top: 1px dashed black;"/> <p style="text-align: center;"><u>Compute masked secret</u></p> <p>1 : $\mathbf{w}' := \text{PosToWitness}(\text{pos})$ 2 : $\mathbf{w} := \text{CompWit}(\mathbf{w}')$ 3 : $\mathbf{t} := \mathbf{w} \oplus \mathbf{u}$</p> <p style="text-align: center;"><u>Construct VOLE correlations with witness</u></p> <p>4 : $\llbracket \beta' \rrbracket := \text{P.EmbedWitness}(\mathbf{w}', (\llbracket u_k \rrbracket)_{k \in [\ell_{\text{row}}]})$</p> <p style="text-align: center;"><u>Compute elementary vector using tensor product</u></p> <p>5 : $\llbracket \mathbf{z} \rrbracket^{(d)} := \text{P.TensorProductToElementaryVector}(\llbracket \beta' \rrbracket, n)$ 6 : // Note that $\llbracket \mathbf{z} \rrbracket^{(d)} := (\llbracket z_0 \rrbracket^{(d)}, \dots, \llbracket z_{n-1} \rrbracket^{(d)})$. 7 : return $(\mathbf{t}, \llbracket \beta' \rrbracket, \llbracket \mathbf{z} \rrbracket^{(d)})$</p>

Algorithm 4.18: P.VOLE-Permutation($\mathbf{P}, (\llbracket u_{i,k} \rrbracket)_{i \in [n], k \in [\ell_{\text{row}]})$)
Public information and inputs
<p>Public information: Matrix dimension n of the secret permutation matrix, degree $d := \lceil \log_4 n \rceil$, size of the compressed secret witness $\ell_{\text{row}} := d + 6$.</p> <p>Prover's input: Secret permutation matrix \mathbf{P} represented as n positions ($\text{pos}_0, \dots, \text{pos}_{n-1}$), $n \cdot \ell_{\text{row}}$ VOLE correlations $\llbracket u_{i,k} \rrbracket$ for random $(u_{i,k}, v_{i,k}) \in \mathbb{F}_2 \times \mathbb{F}_{2^d}$ represented as polynomials $f_{u_{i,k}}(X) = u_{i,k}X + v_{i,k}$ for $(i, k) \in [n] \times [\ell_{\text{row}}]$.</p>
Output
<p>Prover's output: Masked compressed secret $\mathbf{t} \in \mathbb{F}_2^\ell$, where $\ell := n \cdot \ell_{\text{row}}$. Also, n^2 polynomials $\llbracket \mathbf{z} \rrbracket^{(d)} := (\llbracket z_{0,0} \rrbracket^{(d)}, \dots, \llbracket z_{(n-1), (n-1)} \rrbracket^{(d)})$ each of degree d which can be seen as VOLE correlations corresponding to the individual entries of the secret permutation matrix \mathbf{P}. Along with $n(2\ell_{\text{row}} - 6)$ secret degree-1 (linear) polynomials represented as $\llbracket \boldsymbol{\beta} \rrbracket$, and n polynomials $\llbracket \text{ColCheck} \rrbracket^{(d)} := (\llbracket \text{ColCheck}_0 \rrbracket^{(d)}, \dots, \llbracket \text{ColCheck}_{n-1} \rrbracket^{(d)})$ each of degree d which can be seen as VOLE correlations corresponding to the sums of the individual columns of the secret permutation matrix \mathbf{P}.</p> <hr style="border-top: 1px dashed black;"/> <p style="text-align: center;"><u>Compute \mathbf{P} row-wise as n elementary vectors</u></p> <p>1 : for $i \in [n]$:</p> <p>2 : $(\mathbf{t}_i, \llbracket \boldsymbol{\beta}_i \rrbracket, \llbracket \mathbf{z}_i \rrbracket^{(d)}) := \text{P.VOLE-ElementaryVector}(\text{pos}_i, (\llbracket u_{i,k} \rrbracket)_{k \in [\ell_{\text{row}]})$</p> <p>3 : endfor</p> <p style="text-align: center;"><u>Compute \mathbf{P} columns check</u></p> <p>4 : for $j \in [n]$:</p> <p>5 : $\llbracket \text{ColSum}_j \rrbracket^{(d)} := \sum_{i=0}^{n-1} \llbracket z_{i,j} \rrbracket^{(d)}$</p> <p>6 : $\llbracket \text{ColCheck}_j \rrbracket^{(d)} := \llbracket \text{ColSum}_j \rrbracket^{(d)} - X^d$</p> <p>7 : endfor</p> <p>8 : $\mathbf{t} := (\mathbf{t}_0, \dots, \mathbf{t}_{n-1})$</p> <p>9 : $\llbracket \boldsymbol{\beta} \rrbracket := (\llbracket \boldsymbol{\beta}_0 \rrbracket, \dots, \llbracket \boldsymbol{\beta}_{n-1} \rrbracket)$</p> <p>10 : $\llbracket \mathbf{z} \rrbracket^{(d)} := (\llbracket \mathbf{z}_0 \rrbracket^{(d)}, \dots, \llbracket \mathbf{z}_{n-1} \rrbracket^{(d)})$</p> <p>11 : $\llbracket \text{ColCheck} \rrbracket^{(d)} := (\llbracket \text{ColCheck}_0 \rrbracket^{(d)}, \dots, \llbracket \text{ColCheck}_{n-1} \rrbracket^{(d)})$</p> <p>12 : return $(\mathbf{t}, \llbracket \boldsymbol{\beta} \rrbracket, \llbracket \mathbf{z} \rrbracket^{(d)}, \llbracket \text{ColCheck} \rrbracket^{(d)})$</p>

Algorithm 4.19: P.Check-ElementaryBlock($\llbracket \beta'_i \rrbracket$)
Public information and inputs
<p>Public information: Size $n' \in \{2, 4\}$ of a secret block of VOLE correlations $\llbracket \beta'_i \rrbracket$.</p> <p>Prover's input: A block of secret VOLE correlations $\llbracket \beta'_i \rrbracket$, where each element of $\llbracket \beta'_i \rrbracket$ is represented as degree-1 (linear) polynomial with coefficients in \mathbb{F}_{2^ρ}.</p>
Output
<p>Prover's output: $\frac{n'}{2}$ quadratic polynomials, $\llbracket e'_i \rrbracket^{(2)} := \left(\llbracket e'_{i,0,1} \rrbracket^{(2)}, \llbracket e'_{i,2,3} \rrbracket^{(2)} \right)$ (or $\llbracket e'_{i,0,1} \rrbracket^{(2)}$) which can be seen as VOLE correlations.</p> <hr style="border-top: 1px dashed black;"/> <pre> 1: if $\text{len}(\llbracket \beta'_i \rrbracket) = 4$: // Parse $\llbracket \beta'_i \rrbracket$ as $\llbracket \beta'_i \rrbracket := (\llbracket \beta'_{i,0} \rrbracket, \llbracket \beta'_{i,1} \rrbracket, \llbracket \beta'_{i,2} \rrbracket, \llbracket \beta'_{i,3} \rrbracket)$. 2: $\llbracket e'_{i,0,1} \rrbracket^{(2)} := \text{P.Multiply}(\llbracket \beta'_{i,0} \rrbracket, \llbracket \beta'_{i,1} \rrbracket)$ 3: $\llbracket e'_{i,2,3} \rrbracket^{(2)} := \text{P.Multiply}(\llbracket \beta'_{i,2} \rrbracket, \llbracket \beta'_{i,3} \rrbracket)$ 4: $\llbracket e'_i \rrbracket^{(2)} := (\llbracket e'_{i,0,1} \rrbracket^{(2)}, \llbracket e'_{i,2,3} \rrbracket^{(2)})$ 5: return $\llbracket e'_i \rrbracket^{(2)}$ 6: elseif $\text{len}(\llbracket \beta'_i \rrbracket) = 2$: // Parse $\llbracket \beta'_i \rrbracket$ as $\llbracket \beta'_i \rrbracket := (\llbracket \beta'_{i,0} \rrbracket, \llbracket \beta'_{i,1} \rrbracket)$. 7: $\llbracket e'_{i,0,1} \rrbracket^{(2)} := \text{P.Multiply}(\llbracket \beta'_{i,0} \rrbracket, \llbracket \beta'_{i,1} \rrbracket)$ 8: return $\llbracket e'_{i,0,1} \rrbracket^{(2)}$ 9: else : 10: return \perp 11: endif </pre>

Algorithm 4.20: P.Check-ElementaryVector($\llbracket \beta' \rrbracket$)
Public information and inputs
Public information: Length of the secret elementary vector n , degree $d := \lceil \log_4(n) \rceil$, $\ell_{\text{row}} := d + 6$. Prover's input: $(2\ell_{\text{row}} - 6)$ secret degree-1 (linear) polynomials represented as $\llbracket \beta' \rrbracket$.
Output
Prover's output: $(d + 3)$ degree- d polynomials, $\llbracket e' \rrbracket^{(d)}$ which can be seen as VOLE correlations.
<pre> 1: if len($\llbracket \beta' \rrbracket$) = 12: // Security level 1 // Parse $\llbracket \beta' \rrbracket$ as $\llbracket \beta' \rrbracket := (\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket)$. 2: for $i \in [3]$: 3: $\llbracket e'_i \rrbracket^{(2)} := \text{P.Check-ElementaryBlock}(\llbracket \beta'_i \rrbracket)$ 4: $\llbracket e'_i \rrbracket^{(3)} := X \cdot \llbracket e'_i \rrbracket^{(2)}$ 5: endfor 6: $\llbracket e' \rrbracket^{(3)} := (\llbracket e'_0 \rrbracket^{(3)}, \llbracket e'_1 \rrbracket^{(3)}, \llbracket e'_2 \rrbracket^{(3)})$ 7: return $\llbracket e' \rrbracket^{(3)}$ 8: elseif len($\llbracket \beta' \rrbracket$) = 14: // Security levels 3 and 5 // Parse $\llbracket \beta' \rrbracket$ as $\llbracket \beta' \rrbracket := (\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket, \llbracket \beta'_3 \rrbracket)$. 9: $(\llbracket e'_0 \rrbracket^{(3)}, \llbracket e'_1 \rrbracket^{(3)}, \llbracket e'_2 \rrbracket^{(3)}) := \text{P.Check-ElementaryVector}(\llbracket \beta'_0 \rrbracket, \llbracket \beta'_1 \rrbracket, \llbracket \beta'_2 \rrbracket)$ 10: for $i \in [3]$: 11: $\llbracket e'_i \rrbracket^{(4)} := X \cdot \llbracket e'_i \rrbracket^{(3)}$ 12: endfor 13: $\llbracket e'_{3,0,1} \rrbracket^{(2)} := \text{P.Check-ElementaryBlock}(\llbracket \beta'_3 \rrbracket)$ 14: $\llbracket e'_{3,0,1} \rrbracket^{(4)} := X^2 \cdot \llbracket e'_{3,0,1} \rrbracket^{(2)}$ 15: $\llbracket e' \rrbracket^{(4)} := (\llbracket e'_0 \rrbracket^{(4)}, \llbracket e'_1 \rrbracket^{(4)}, \llbracket e'_2 \rrbracket^{(4)}, \llbracket e'_{3,0,1} \rrbracket^{(4)})$ 16: return $\llbracket e' \rrbracket^{(4)}$ 17: else : 18: return \perp 19: endif </pre>

Algorithm 4.21: $\text{P.CheckZero}(\llbracket w \rrbracket^{(d)}, (\llbracket u_{i,k} \rrbracket)_{(i,k) \in [d-1] \times [\rho]})$
Public information and inputs
<p>Public information: Degree of input VOLE correlation $\llbracket w \rrbracket^{(d)}$ (seen as polynomial) d.</p> <p>Prover's input: Degree-d VOLE correlation $\llbracket w \rrbracket^{(d)}$, $(d-1)\rho$ random VOLE correlation represented as $f_{u_{i,k}}(X) = u_{i,k}X + v_{i,k}$ where $(u_{i,k}, v_{i,k}) \in \mathbb{F}_2 \times \mathbb{F}_{2^\rho}$ for $(i, k) \in [d-1] \times [\rho]$.</p>
Output
<p>Prover's output: Polynomial $\llbracket a \rrbracket$. Note that for an honest prover, the leading coefficient (coefficient of X^d term will be equal to 0) and therefore $\llbracket a \rrbracket$ will consists of only d coefficients for terms X^i for $i \in [d]$.</p> <hr style="border-top: 1px dashed black;"/> <pre> // Generating VOLE correlations in $\mathbb{F}_{2^\rho} \times \mathbb{F}_{2^\rho}$ 1 : for $i \in [d-1]$: 2 : $u'_i := \sum_{k=0}^{\rho-1} u_{i,k} \gamma_\rho^k$ // $\{\gamma_\rho^k\}_{k=0}^{\rho-1}$ is the power basis of \mathbb{F}_{2^ρ} with coefficients in \mathbb{F}_2. 3 : $v'_i := \sum_{k=0}^{\rho-1} v_{i,k} \gamma_\rho^k$ 4 : $f_{u'_i}(X) := u'_i X + v'_i$ // (u'_i, v'_i) in $\mathbb{F}_{2^\rho} \times \mathbb{F}_{2^\rho}$. 5 : endfor 6 : $f_{\text{mask}}(X) := \sum_{i=0}^{d-2} f_{u'_i}(X) \cdot X^i$ 7 : $\llbracket a \rrbracket := f_w(X) + f_{\text{mask}}(X)$ 8 : return $\llbracket a \rrbracket$ </pre>

Algorithm 4.22:

P.Check-PKP(\mathbf{P} , pk , $(\llbracket u_{i,k} \rrbracket)_{(i,k) \in [n] \times [\ell_{\text{row}}]}$, $(\llbracket u_{i',k'} \rrbracket)_{(i',k') \in [d-1] \times [\rho]}$, seed)

Public information

Public information: Matrix dimension n of the secret permutation matrix, degree $d := \lceil \log_4 n \rceil$, size of the compressed secret witness $\ell_{\text{row}} := d + 6$.

Prover's input: Secret permutation matrix \mathbf{P} represented as n positions $(\text{pos}_0, \dots, \text{pos}_{n-1})$, public key $\text{pk} = (\mathbf{H}, \mathbf{x})$, $n \cdot \ell_{\text{row}}$ VOLE correlations $\llbracket u_{i,k} \rrbracket$ for random $(u_{i,k}, v_{i,k}) \in \mathbb{F}_2 \times \mathbb{F}_{2^\rho}$ represented as polynomials $f_{u_{i,k}}(X) = u_{i,k}X + v_{i,k}$ for $(i,k) \in [n] \times [\ell_{\text{row}}]$, $(d-1)\rho$ VOLE correlations $\llbracket u_{i',k'} \rrbracket$ for random $u_{i',k'}, v_{i',k'} \in \mathbb{F}_2 \times \mathbb{F}_{2^\rho}$ represented as $f_{u_{i',k'}}(X) = u_{i',k'}X + v_{i',k'}$ for $(i',k') \in [d-1] \times [\rho]$, $\text{seed} \in \{0,1\}^{2^\lambda}$.

Output

Prover's output: Degree- d polynomial $\llbracket a \rrbracket$ as a proof to show that \mathbf{P} is a solution to the PKP instance defined by pk .

Compute \mathbf{P} in matrix form

1 : $(\mathbf{t}, \llbracket \beta \rrbracket, \llbracket \mathbf{z} \rrbracket^{(d)}, \llbracket \text{ColCheck} \rrbracket^{(d)}) := \text{P.VOLE-Permutation}(\mathbf{P}, (\llbracket u_{i,k} \rrbracket)_{(i,k) \in [n] \times [\ell_{\text{row}}]})$
 // Parse $\llbracket \text{ColCheck} \rrbracket^{(d)}$ as $(f_0(X), f_1(X), \dots, f_{n-1}(X))$

Check elementary vectors

2 : **for** $i \in [n]$
 3 : $\llbracket \text{ElemVecCheck}_i \rrbracket^{(d)} := \text{P.Check-ElementaryVector}(\llbracket \beta_i \rrbracket)$
 // Each $\llbracket \text{ElemVecCheck}_i \rrbracket^{(d)}$ contains 6 degree- d polynomials if $\lambda = 128$,
 // and 7 degree- d polynomials if $\lambda \in \{192, 256\}$.
 4 : **endfor**
 5 : $\llbracket \text{ElemVecCheck} \rrbracket^{(d)} := (\llbracket \text{ElemVecCheck}_0 \rrbracket^{(d)}, \dots, \llbracket \text{ElemVecCheck}_{n-1} \rrbracket^{(d)})$
 // Parse $\llbracket \text{ElemVecCheck} \rrbracket^{(d)}$ as $(f_n(X), f_{n+1}(X), \dots, f_{cn+n-1}(X))$,
 // where, $c = 6$ if $\lambda = 128$, and $c = 7$ if $\lambda \in \{192, 256\}$.

Compute $\mathbf{x}' = \mathbf{P}\mathbf{x}$

6 : **for** $i \in [n]$
 7 : $\llbracket \mathbf{x}'_i \rrbracket^{(d)} := \sum_{j=0}^{n-1} \llbracket z_{i,j} \rrbracket^{(d)} \cdot \mathbf{x}_j$
 8 : **endfor**

Compute $\mathbf{y} = \mathbf{H}\mathbf{x}'$

9 : **for** $i \in [m]$
 10 : $\llbracket \mathbf{y}_i \rrbracket^{(d)} = f_{i+cn+n}(X) := \sum_{j=0}^{n-1} h_{i,j} \cdot \llbracket \mathbf{x}'_j \rrbracket^{(d)}$
 11 : **endfor**

Merge polynomials and run CheckZero

12 : $\alpha := \text{H}_4(\text{seed} :: \rho \cdot (cn + n + m))$ // α should be parsed as $\in \mathbb{F}_{2^\rho}^{cn+n+m}$
 13 : $f(X) := \sum_{j=0}^{cn+n+m-1} \alpha_j \cdot f_j(X)$
 14 : $\llbracket a \rrbracket := \text{P.CheckZero}(f(X), (\llbracket u_{i',k'} \rrbracket)_{(i',k') \in [d-1] \times [\rho]})$

15 : $\text{proof} := \llbracket a \rrbracket$
 16 : **return** proof

Verifier. In this section we present all the algorithms that will be used by the verifier to verify that the prover has knowledge of the secret permutation which serves as a solution to the PKP instance corresponding to the public key. As expected in VOLE-in-the-Head (or MPC-in-the-Head) type PoK, the verifier’s algorithm bear a close resemblance to those used by the prover. In the case of PERK, the main difference is that while prover’s algorithms explained in [Section 4.5](#) take polynomials as inputs and manipulate them, the verifier’s algorithms described in this section perform analogous manipulations on evaluations of corresponding polynomials. The verifier possesses the VOLE correlation inputs \mathbf{q} and Δ , and it also receives the masked (compressed) witness \mathbf{t} and masked polynomial $\llbracket a \rrbracket$ from the prover.

The first step the verifier performs is to update the VOLE correlation inputs \mathbf{q} with the help of the masked witness \mathbf{t} , to ensure that they satisfy the VOLE correlation with respect to witness \mathbf{w} (instead of \mathbf{u}). This is achieved by [Algorithm 4.23 EmbedMaskedWitnessBlock](#). The output of this algorithm are the VOLE correlations \mathbf{q}'_{β_i} corresponding to the elementary blocks of witness. Similar to the prover’s case, the [Algorithm 4.24 EmbedMaskedWitness](#) aggregates the VOLE correlations corresponding to the individual elementary vectors of lengths 4 and 2, and outputs VOLE correlations $\mathbf{q}'_{\beta'}$ corresponding to the aggregation of 3 elementary vectors of lengths 4 (and in case of L3 and L5 parameters, another elementary vector of length 2). Once the verifier possesses the VOLE correlation inputs ($\mathbf{q}'_{\beta'}$) corresponding to the elementary vector entries, it then computes VOLE correlations \mathbf{q}_z corresponding to each individual row of the secret permutation matrix with help of [Algorithm 4.26 TensorProductToElementaryVector](#), which internally calls [Algorithm 4.25 TensorProduct](#) to compute the tensor product between two blocks. The [Algorithm 4.27 VOLE-ElementaryVector](#) produces the VOLE correlations $\mathbf{q}'_{\beta'}$ along with \mathbf{q}_z .

After obtaining the VOLE correlations for each of the n rows by running [Algorithm 4.27 VOLE-ElementaryVector](#) n times, the verifier proceeds to compute the extra n VOLE correlation values, $\mathbf{q}_{\text{ColCheck}}$ which ensure that each column of secret matrix adds upto exactly 1. This is described in [Algorithm 4.28 VOLE-Permutation](#). The verifier checks the elementary structure by computing the values \mathbf{q}'_e , using [Algorithm 4.29 Check-ElementaryBlock](#), and [Algorithm 4.30 Check-ElementaryVector](#).

In order to check that the leading coefficient of the masked polynomial $\llbracket a \rrbracket$ is zero, the verifier first checks the degree of the polynomial is equal to $d-1$, it then evaluates the masking polynomial using its inputs \mathbf{q}, Δ and then finally checks if the evaluation of the polynomial $\llbracket a \rrbracket$ sent by the prover at point Δ matches the addition of the \mathbf{q} value obtained from the computations checking constraints related to the PKP problem and the evaluation of masking polynomial. The [Algorithm 4.31 CheckZero](#) achieves this and outputs 1 when all the values match and outputs 0 otherwise indicating the failure to verify the prover’s claim.

Finally, as in the prover's case, the [Algorithm 4.32 Check-PKP](#) puts all of the checks for checking the elementary structure of rows (blocks), column sums equaling to 1, and satisfiability of PKP equation together by evaluating degree- d polynomials at Δ . These evaluations are then merged together into a single value by computing (verifier dictated) random linear combination, which is then checked with the help of [Algorithm 4.31 CheckZero](#).

Algorithm 4.23: V.EmbedMaskedWitnessBlock($\Delta, \mathbf{t}'_i, (q_k)_{k \in [3]}$)
Public information and inputs
<p>Public information: Length of the masked witness block = 4.</p> <p>Verifier's input: VOLE correlation challenge $\Delta \in \mathbb{F}_{2^\rho}$, i^{th} block $\mathbf{t}'_i := t'_{i,0} t'_{i,1} t'_{i,2} t'_{i,3} \in \mathbb{F}_2^4$ of the masked witness \mathbf{t}', 3 VOLE correlation inputs (q_0, q_1, q_2) with each of them in \mathbb{F}_{2^ρ}.</p>
Output
<p>Verifier's output: VOLE correlation values $\mathbf{q}'_{\beta'_i} := (q'_{\beta'_i,0}, q'_{\beta'_i,1}, q'_{\beta'_i,2}, q'_{\beta'_i,3}) \in \mathbb{F}_{2^\rho}^4$</p> <hr/> <p style="text-align: center;"><u>Construct VOLE correlations with witness</u></p> <p>1: // Parse \mathbf{t}'_i as $\mathbf{t}'_i := t'_{i,0} t'_{i,1} t'_{i,2} t'_{i,3}$.</p> <p>2: for $j \in [3]$:</p> <p>3: $q'_{\beta'_i,j} := t'_{i,j} \cdot \Delta + q_j$</p> <p>4: endfor</p> <p>5: $q'_{\beta'_i,3} := t'_{i,3} \cdot \Delta + q_0 + q_1 + q_2$</p> <p>6: $\mathbf{q}'_{\beta'_i} := (q'_{\beta'_i,0}, q'_{\beta'_i,1}, q'_{\beta'_i,2}, q'_{\beta'_i,3})$</p> <p>7: return $\mathbf{q}'_{\beta'_i}$</p>

Algorithm 4.24: V.EmbedMaskedWitness($\Delta, \mathbf{t}', \mathbf{q}$)
Public information and inputs
<p>Public information: Length of the compressed masked witness ℓ_{row}, and length of the masked witness $\mathbf{t}' := 2 \cdot \ell_{\text{row}} - 6$.</p> <p>Verifier's input: VOLE correlation challenge $\Delta \in \mathbb{F}_{2^\rho}$, the masked witness \mathbf{t}', ℓ_{row} VOLE correlation inputs $\mathbf{q} := (q_0, \dots, q_{\ell_{\text{row}}-1})$ with each of them in \mathbb{F}_{2^ρ}.</p>
Output
<p>Verifier's output: $\mathbf{q}'_{\beta'}$ with elements in \mathbb{F}_{2^ρ}.</p> <hr/> <p style="text-align: center;"><u>Construct VOLE correlations with witness</u></p> <pre> 1: if len(\mathbf{t}') = 12: // Security level 1 // Parse \mathbf{t}' as $\mathbf{t}' := \mathbf{t}'_0 \mathbf{t}'_1 \mathbf{t}'_2$, where, $\mathbf{t}'_i := t'_{i,0} t'_{i,1} t'_{i,2} t'_{i,3}$. 2: for $i \in [3]$: 3: $\mathbf{q}'_{\beta'_i} := \text{V.EmbedMaskedWitnessBlock}(\Delta, \mathbf{t}'_i, \mathbf{q}[3i : 3i + 2])$ 4: endfor 5: $\mathbf{q}'_{\beta'} := (\mathbf{q}'_{\beta'_0}, \mathbf{q}'_{\beta'_1}, \mathbf{q}'_{\beta'_2})$ 6: return $\mathbf{q}'_{\beta'}$ 7: elseif len(\mathbf{t}') = 14: // Security levels 3 and 5 // Parse \mathbf{t}' as $\mathbf{t}' := \mathbf{t}'_0 \mathbf{t}'_1 \mathbf{t}'_2 \mathbf{t}'_{3,0} \mathbf{t}'_{3,1}$, where, \mathbf{t}'_i are as above (line 1) 8: $(\mathbf{q}'_{\beta'_0}, \mathbf{q}'_{\beta'_1}, \mathbf{q}'_{\beta'_2}) := \text{V.EmbedMaskedWitness}(\Delta, \mathbf{t}'_0 \mathbf{t}'_1 \mathbf{t}'_2, \mathbf{q}[0 : 8])$ 9: $\mathbf{q}'_{\beta'_{3,0}} := \mathbf{t}'_{3,0} \cdot \Delta + q_9$ 10: $\mathbf{q}'_{\beta'_{3,1}} := \mathbf{t}'_{3,1} \cdot \Delta + q_9$ 11: $\mathbf{q}'_{\beta'_3} := (\mathbf{q}'_{\beta'_{3,0}}, \mathbf{q}'_{\beta'_{3,1}})$ 12: $\mathbf{q}'_{\beta'} := (\mathbf{q}'_{\beta'_0}, \mathbf{q}'_{\beta'_1}, \mathbf{q}'_{\beta'_2}, \mathbf{q}'_{\beta'_3})$ 13: return $\mathbf{q}'_{\beta'}$ 14: else : 15: return \perp 16: endif </pre>

Algorithm 4.25: V.TensorProduct($\mathbf{q}'_{\beta'_i}, \mathbf{q}'_{\beta'_j}, \text{num}$)
Public information and inputs
<p>Public information: Sizes n_i and n_j of the two blocks of verifier's VOLE correlation inputs $\mathbf{q}'_{\beta'_i}$ and $\mathbf{q}'_{\beta'_j}$ respectively. Each element of $\mathbf{q}'_{\beta'_i}$ (resp. $\mathbf{q}'_{\beta'_j}$) is in \mathbb{F}_{2^ρ}.</p> <p>Verifier's input: $(n_i + n_j)$ values $\mathbf{q}'_{\beta'_i} := (q'_{\beta'_i,0}, \dots, q'_{\beta'_i,(n_i-1)})$ and $\mathbf{q}'_{\beta'_j} := (q'_{\beta'_j,0}, \dots, q'_{\beta'_j,(n_j-1)})$.</p>
Output
<p>Verifier's output: $\mathbf{q}_\zeta := (q_{\zeta_0}, \dots, q_{\zeta_{\text{num}^* - 1}})$ is a block of num^* values in \mathbb{F}_{2^ρ}, where $\text{num}^* := \min(n_i n_j, \text{num})$.</p>
<p style="text-align: center;"><u>Computing tensor product</u></p> <pre> 1 : counter := 0 2 : for index_j ∈ [n_j]: 3 : for index_i ∈ [n_i]: 4 : if counter ∈ [num]: 5 : q_{ζ_{counter}} := q'_{β'_j, index_j} · q'_{β'_i, index_i} 6 : counter := counter + 1 7 : else : 8 : break 9 : endfor 10 : endfor 11 : q_ζ := (q_{ζ₀}, ..., q_{ζ_{num* - 1}}) 12 : return q_ζ </pre>

Algorithm 4.26: $V.TensorProductToElementaryVector(q'_{\beta'}, num)$	
Public information and inputs	
<p>Public information: Length of the input vector $num \in [128]$. Lengths of the blocks $q'_{\beta'_i} := 4$ for $i \in [3]$, and $q'_{\beta'_3} := 2$</p> <p>Verifier's input: VOLE correlation input $q'_{\beta'}$ with elements in \mathbb{F}_{2^ρ}.</p>	
Output	
<p>Verifier's output: q_z a block of values in \mathbb{F}_{2^ρ}.</p> <hr/> <p style="text-align: center;"><u>Compute elementary vector using tensor product</u></p> <pre> 1: if $len(q'_{\beta'}) = 12$: // Security level 1 // Parse $q'_{\beta'}$ as $q'_{\beta'} := (q'_{\beta'_0}, q'_{\beta'_1}, q'_{\beta'_2})$. 2: $q_{\zeta_{0,1}} := V.TensorProduct(q'_{\beta'_0}, q'_{\beta'_1}, num)$ 3: $q_z := V.TensorProduct(q_{\zeta_{0,1}}, q'_{\beta'_2}, num)$ 4: return q_z 5: elseif $len(q'_{\beta'}) = 14$: // Security levels 3 and 5 // Parse $q'_{\beta'}$ as $q'_{\beta'} := (q'_{\beta'_0}, q'_{\beta'_1}, q'_{\beta'_2}, q'_{\beta'_3})$. 6: $q_{\zeta_{0,1}} := V.TensorProduct(q'_{\beta'_0}, q'_{\beta'_1}, num)$ 7: $q_{\zeta_{0,1,2}} := V.TensorProduct(q_{\zeta_{0,1}}, q'_{\beta'_2}, num)$ 8: $q_z := V.TensorProduct(q_{\zeta_{0,1,2}}, q'_{\beta'_3}, num)$ 9: return q_z 10: else : 11: return \perp 12: endif </pre>	

Algorithm 4.27: V.VOLE-ElementaryVector($\Delta, \mathbf{t}, \mathbf{q}$)
Public information and inputs
<p>Public information: Length of the elementary vector n, length of the compressed masked witness $\mathbf{t} := \ell_{\text{row}}$.</p> <p>Verifier's input: VOLE correlation challenge $\Delta \in \mathbb{F}_{2^\rho}$, the compressed masked witness \mathbf{t}, ℓ_{row} VOLE correlation inputs $\mathbf{q} := (q_0, \dots, q_{\ell_{\text{row}}-1})$ with each of them in \mathbb{F}_{2^ρ}.</p>
Output
<p>Verifier's output: Tuple of $(2 \cdot \ell_{\text{row}} - 6)$ VOLE correlation values corresponding to the shares β' held by the prover, along with another tuple of n VOLE correlation values $\mathbf{q}_z := (q_{z_0}, q_{z_1}, \dots, q_{z_{n-1}})$ corresponding to the secret elementary vector of length n held by the prover. All VOLE correlation values output are in \mathbb{F}_{2^ρ}.</p> <hr/> <p><u>Compute masked secret</u></p> <p>1 : $\mathbf{t}' := \text{ExpWit}(\mathbf{t})$</p> <p>2 : $\mathbf{q}'_{\beta'} := \text{V.EmbedMaskedWitness}(\Delta, \mathbf{t}', \mathbf{q})$</p> <p><u>Compute elementary vector using tensor product</u></p> <p>3 : $\mathbf{q}_z := \text{V.TensorProductToElementaryVector}(\mathbf{q}'_{\beta'}, n)$</p> <p>4 : return $(\mathbf{q}'_{\beta'}, \mathbf{q}_z)$</p>

Algorithm 4.28: V.VOLE-Permutation($\Delta, \mathbf{t}, \mathbf{q}$)
Public information and inputs
<p>Public information: Matrix dimension n of the secret permutation matrix, length of the compressed masked witness $\mathbf{t} := \ell$. Note that $\ell = n\ell_{\text{row}}$.</p> <p>Verifier's input: VOLE correlation challenge $\Delta \in \mathbb{F}_{2^{\rho}}$, the compressed masked witness $\mathbf{t} := (\mathbf{t}_0, \mathbf{t}_1, \dots, \mathbf{t}_{n-1})$, where each $\mathbf{t}_i \in \mathbb{F}_2^{\ell_{\text{row}}}$, ℓ VOLE correlation inputs $\mathbf{q} := (\mathbf{q}_0, \dots, \mathbf{q}_{n-1})$ with each \mathbf{q}_i consists of ℓ_{row} values in $\mathbb{F}_{2^{\rho}}$.</p>
Output
<p>Verifier's output: Tuple of VOLE correlation values corresponding to the shares β held by the prover, tuple of n VOLE correlation values $\mathbf{q}_z := (q_{z_0}, q_{z_1}, \dots, q_{z_{n-1}})$ corresponding to the secret elementary vector of length n held by the prover, along with tuple of n VOLE correlation values $\mathbf{q}_{\text{ColCheck}} := (q_{\text{ColCheck}_0}, q_{\text{ColCheck}_1}, \dots, q_{\text{ColCheck}_{n-1}})$ corresponding to (sum of column entries - 1) for each column of the secret permutation matrix held by the prover. All VOLE correlation values output are in $\mathbb{F}_{2^{\rho}}$.</p> <hr/> <p style="text-align: center;"><u>Compute \mathbf{P} row-wise as n elementary vectors</u></p> <p>1: for $i \in [n]$:</p> <p>2: $(\mathbf{q}_{\beta_i}, \mathbf{q}_{z_i}) := \text{V.VOLE-ElementaryVector}(\Delta, \mathbf{t}_i, \mathbf{q}_i)$</p> <p>3: endfor</p> <p style="text-align: center;"><u>Compute \mathbf{P} columns check</u></p> <p>4: for $j \in [n]$:</p> <p>5: $q_{\text{ColSum}_j} = \sum_{i=0}^{n-1} q_{z_i, j}$</p> <p>6: $q_{\text{ColCheck}_j} = q_{\text{ColSum}_j} - \Delta^d$</p> <p>7: endfor</p> <p>8: $\mathbf{q}_{\beta} := (\mathbf{q}_{\beta_0}, \mathbf{q}_{\beta_1}, \dots, \mathbf{q}_{\beta_{n-1}})$</p> <p>9: $\mathbf{q}_z := (\mathbf{q}_{z_0}, \mathbf{q}_{z_1}, \dots, \mathbf{q}_{z_{n-1}})$</p> <p>10: $\mathbf{q}_{\text{ColCheck}} := (q_{\text{ColCheck}_0}, q_{\text{ColCheck}_1}, \dots, q_{\text{ColCheck}_{n-1}})$</p> <p>11: return $(\mathbf{q}_{\beta}, \mathbf{q}_z, \mathbf{q}_{\text{ColCheck}})$</p>

Algorithm 4.29: V.Check-ElementaryBlock($\mathbf{q}'_{\beta'_i}$)	
Public information and inputs	
<p>Public information: Length of the elementary vector block to be checked $\in \{2, 4\}$.</p> <p>Verifier's input: A tuple of VOLE correlation inputs $\mathbf{q}'_{\beta'_i}$ with same size as the elementary vector block to be checked, with each element of $\mathbf{q}'_{\beta'_i}$ in \mathbb{F}_{2^ρ}.</p>	
Output	
<p>Verifier's output: VOLE correlation values $\mathbf{q}'_{e'_i}$ (which can be a single element or a tuple) to help verify the elementary vector structure the secret vector held by the prover.</p> <hr style="border-top: 1px dashed black;"/> <pre> 1 : if $\text{len}(\mathbf{q}'_{\beta'_i}) = 4$: // Parse $\mathbf{q}'_{\beta'_i}$ as $\mathbf{q}'_{\beta'_i} := (q'_{\beta'_i,0}, q'_{\beta'_i,1}, q'_{\beta'_i,2}, q'_{\beta'_i,3})$. 2 : $q'_{e'_i,0,1} := q'_{\beta'_i,0} \cdot q'_{\beta'_i,1}$ 3 : $q'_{e'_i,0,2} := q'_{\beta'_i,2}, q'_{\beta'_i,3}$ 4 : $\mathbf{q}'_{e'_i} := (q'_{e'_i,0,1}, q'_{e'_i,0,2})$ 5 : return $\mathbf{q}'_{e'_i}$ 6 : elseif $\text{len}(\mathbf{q}'_{\beta'_i}) = 2$: // Parse $\mathbf{q}'_{\beta'_i}$ as $\mathbf{q}'_{\beta'_i} := (q'_{\beta'_i,0}, q'_{\beta'_i,1})$. 7 : $q'_{e'_i,0,1} := q'_{\beta'_i,0} \cdot q'_{\beta'_i,1}$ 8 : return $q'_{e'_i,0,1}$ 9 : else : 10 : return \perp 11 : endif </pre>	

Algorithm 4.30: V.Check-ElementaryVector($\Delta, \mathbf{q}'_{\beta'}$)	
Public information and inputs	
<p>Public information: Length of the tuple of VOLE correlation inputs corresponding to the masked witness $\mathbf{q}'_{\beta'} := 2 \cdot \ell_{\text{row}} - 6$. Lengths of the blocks $\mathbf{q}'_{\beta'_i} := 4$ for $i \in [3]$, and $\mathbf{q}'_{\beta'_3} := 2$</p> <p>Verifier's input: VOLE correlation challenge $\Delta \in \mathbb{F}_{2^\rho}$, a tuple of VOLE correlation inputs $\mathbf{q}'_{\beta'}$ with each of element in \mathbb{F}_{2^ρ}.</p>	
Output	
<p>Verifier's output: A tuple of VOLE correlation values $\mathbf{q}'_{e'}$ to help verify the elementary vector structure the secret vector held by the prover.</p> <hr/> <pre> 1 : if len($\mathbf{q}'_{\beta'}$) = 12: // Security level 1 // Parse $\mathbf{q}'_{\beta'}$ as $\mathbf{q}'_{\beta'} := (\mathbf{q}'_{\beta'_0}, \mathbf{q}'_{\beta'_1}, \mathbf{q}'_{\beta'_2})$. 2 : for $i \in [3]$: 3 : $\mathbf{q}'_{e'_i} := \text{V.Check-ElementaryBlock}(\mathbf{q}'_{\beta'_i})$ 4 : $\mathbf{q}'_{e'_i} := \Delta \cdot \mathbf{q}'_{e'_i}$ // Here each element of the tuple $\mathbf{q}'_{e'_i}$ should be multiplied by Δ. 5 : endfor 6 : $\mathbf{q}'_{e'} := (\mathbf{q}'_{e'_0}, \mathbf{q}'_{e'_1}, \mathbf{q}'_{e'_2})$ 7 : return $\mathbf{q}'_{e'}$ 8 : elseif len($\mathbf{q}'_{\beta'}$) = 14: // Security levels 3 and 5 // Parse $\mathbf{q}'_{\beta'}$ as $\mathbf{q}'_{\beta'} := (\mathbf{q}'_{\beta'_0}, \mathbf{q}'_{\beta'_1}, \mathbf{q}'_{\beta'_2}, \mathbf{q}'_{\beta'_3})$. 9 : $(\mathbf{q}'_{e'_0}, \mathbf{q}'_{e'_1}, \mathbf{q}'_{e'_2}) := \text{V.Check-ElementaryVector}(\mathbf{q}'_{\beta'_0}, \mathbf{q}'_{\beta'_1}, \mathbf{q}'_{\beta'_2})$ 10 : for $i \in [3]$: 11 : $\mathbf{q}'_{e'_i} := \Delta \cdot \mathbf{q}'_{e'_i}$ // Here each element of the tuple $\mathbf{q}'_{e'_i}$ should be multiplied by Δ. 12 : endfor 13 : $q'_{e'_{3,0,1}} := \text{V.Check-ElementaryBlock}(\mathbf{q}'_{\beta'_3})$ 14 : $q'_{e'_{3,0,1}} := \Delta^2 \cdot q'_{e'_{3,0,1}}$ 15 : $\mathbf{q}'_{e'} := (\mathbf{q}'_{e'_0}, \mathbf{q}'_{e'_1}, \mathbf{q}'_{e'_2}, q'_{e'_{3,0,1}})$ 16 : return $\mathbf{q}'_{e'}$ 17 : else : 18 : return \perp 19 : endif </pre>	

Algorithm 4.31: $V.\text{CheckZero}(\Delta, q_f, (q_{u_{i,k}})_{(i,k) \in [d-1] \times [\rho]}, \llbracket a \rrbracket)$
Public information and inputs
Public information: Degree of input polynomial $\llbracket a \rrbracket := d - 1$. Verifier's input: Polynomial $\llbracket a \rrbracket$, $q_f, q_{u_{i,k}} = f_{u_{i,k}}(\Delta)$ for $(i, k) \in [d - 1] \times [\rho]$, $\Delta \in \mathbb{F}_{2^\rho}$.
Output
Verifier's output: Boolean indicating if leading coefficient (coefficient of X^d) of some degree- d polynomial $f(X)$ (for which verifier already holds $q_f \in \mathbb{F}_{2^\rho}$) is equal to zero or not.
<hr/> <pre> 1: if degree of $\llbracket a \rrbracket \neq d - 1$: 2: return \perp 3: else : 4: // Generating $q'_{u_i} \in \mathbb{F}_{2^\rho}$ for $i \in [d - 1]$ 5: for $i \in [d - 1]$: 6: $q'_{u_i} := \sum_{k=0}^{\rho-1} q_{u_{i,k}} \cdot \gamma_\rho^k$ // $\{\gamma_\rho^k\}_{k=0}^{\rho-1}$ is the power basis of \mathbb{F}_{2^ρ}. 7: endfor 8: // Parse $\llbracket a \rrbracket$ as $\llbracket a \rrbracket := a_{d-1}X^{d-1} + \dots + a_1X + a_0$ 9: Compute $\tilde{q} := \sum_{i=0}^{d-1} a_i \cdot \Delta^i$ 10: $q = q_f + \sum_{i=0}^{d-2} q'_{u_i} \cdot \Delta^i$ 11: $b := (q \stackrel{?}{=} \tilde{q})$ 12: return b </pre>

Algorithm 4.32: V.Check-PKP(seed, pk, \mathbf{t} , Δ , \mathbf{q} , \mathbf{q}_{cz} , $\llbracket a \rrbracket$)
Public information
<p>Public information: Matrix dimension n of the secret permutation matrix, $d = \lceil \log_4 n \rceil$.</p> <p>Verifier's input: VOLE correlation challenge $\Delta \in \mathbb{F}_{2^\rho}$, the compressed masked witness $\mathbf{t} := (\mathbf{t}_0, \mathbf{t}_1, \dots, \mathbf{t}_{n-1})$, where each $\mathbf{t}_i \in \mathbb{F}_2^{\ell_{\text{row}}}$, ℓ VOLE correlation inputs $\mathbf{q} := (\mathbf{q}_0, \dots, \mathbf{q}_{n-1})$ with each \mathbf{q}_i consists of ℓ_{row} values in \mathbb{F}_{2^ρ}, $(d-1)$ VOLE correlations \mathbf{q}_{cz} in \mathbb{F}_{2^ρ}, polynomial $\llbracket a \rrbracket$, and seed $\in \{0, 1\}^{2^\lambda}$.</p>
Output
<p>Verifier's output: Boolean value b indicating if the proof is accepted or not.</p> <hr/> <p style="text-align: center;"><u>Compute VOLE correlations corresponding to shares of \mathbf{P} in matrix form</u></p> <p>1 : $(\mathbf{q}_\beta, \mathbf{q}_z, \mathbf{q}_{\text{ColCheck}}) := \mathbf{V.VOLE-Permutation}(\Delta, \mathbf{t}, \mathbf{q})$ // Parse $\mathbf{q}_{\text{ColCheck}}$ as $(q_{f_0}, q_{f_1}, \dots, q_{f_{n-1}})$</p> <p style="text-align: center;"><u>Check elementary vectors</u></p> <p>2 : for $i \in [0, n-1]$</p> <p>3 : $\mathbf{q}_{\text{ElemVecCheck}_i} := \mathbf{V.Check-ElementaryVector}(\Delta, \mathbf{q}_{\beta_i})$ // Each $\mathbf{q}_{\text{ElemVecCheck}_i}$ contains 6 elements if $\lambda = 128$, // and 7 elements if $\lambda \in \{192, 256\}$.</p> <p>4 : endfor</p> <p>5 : $\mathbf{q}_{\text{ElemVecCheck}} := (\mathbf{q}_{\text{ElemVecCheck}_0}, \dots, \mathbf{q}_{\text{ElemVecCheck}_{n-1}})$ // Parse $\mathbf{q}_{\text{ElemVecCheck}}$ as $(q_{f_n}, q_{f_{n+1}}, \dots, q_{f_{cn+n-1}})$, // where, $c = 6$ if $\lambda = 128$, and $c = 7$ if $\lambda \in \{192, 256\}$.</p> <p style="text-align: center;"><u>Compute $\mathbf{x}' = \mathbf{P}\mathbf{x}$</u></p> <p>6 : for $i \in [0, n-1]$</p> <p>7 : $q_{\mathbf{x}'_i} = \sum_{j=0}^{n-1} q_{z_{i,j}} \cdot \mathbf{x}_j$</p> <p>8 : endfor</p> <p style="text-align: center;"><u>Compute $\mathbf{y} = \mathbf{H}\mathbf{x}'$</u></p> <p>9 : for $i \in [0, m-1]$</p> <p>10 : $q_{\mathbf{y}_i} = q_{f_{i+cn+n}} = \sum_{j=0}^{n-1} h_{i,j} \cdot q_{\mathbf{x}'_j}$</p> <p>11 : endfor</p> <p style="text-align: center;"><u>Merge polynomials and run CheckZero</u></p> <p>12 : $\alpha := \mathbf{H}_4(\text{seed} :: \rho \cdot (cn + n + m))$ // α should be parsed as $\in \mathbb{F}_{2^\rho}^{cn+n+m}$</p> <p>13 : $q_f = \sum_{j=0}^{cn+n+m-1} \alpha_j \cdot q_{f_j}$</p> <p>14 : $b := \mathbf{V.CheckZero}(\Delta, q_f, \mathbf{q}_{\text{cz}}, \llbracket a \rrbracket)$</p> <p>15 : return b</p>

4.6 PERK

In this section, we describe the `PERK.KeyGen`, `PERK.Sign` and `PERK.Verify` algorithms. The PERK key generation algorithm `PERK.KeyGen` takes the public parameters for the PKP $\text{param}_{\text{PKP}} := (q, m, n)$ as input along with the security parameter λ . The key generation algorithm outputs a public key represented by a seed of length λ which generates the matrix \mathbf{H} (in deterministic manner), and a vector $\mathbf{x} \in \mathbb{F}_q^n$; along with the secret key also represented by a separate seed of length λ which generates the secret permutation in a deterministic way.

The key generation algorithm, first samples 3 distinct seeds \mathbf{H}_{seed} , $\text{perm}_{\text{seed}}$, and ker_{seed} of length λ independently and at uniform random. It expands the seed \mathbf{H}_{seed} to compute a pseudorandom matrix $\mathbf{M} \in \mathbb{F}_q^{m \times (n-m)}$ and sets the public matrix in its canonical form as $\mathbf{H} := [\mathbf{I}_m \ \mathbf{M}]$. Then it computes the basis of kernel of \mathbf{H} and samples a random vector \mathbf{x}' by taking a random linear combination of the basis vectors, this random linear combination is derived using ker_{seed} . The key generation algorithm also samples a permutation π using $\text{perm}_{\text{seed}}$ and sets $\mathbf{x} := \pi^{-1}(\mathbf{x}')$. The algorithm outputs $\text{pk} := (\mathbf{H}_{\text{seed}}, \mathbf{x})$, and $\text{sk} := \text{perm}_{\text{seed}}$.

Algorithm 4.33: PERK.KeyGen()
Public information and inputs
Public information: Public parameters $\text{param}_{\text{PKP}} := (q, m, n)$, and security parameter λ .
Output
The public key as a seed that generates the public matrix concatenated with the public vector, and the private key as the seed that generates the secret permutation.

<u>Sampling randomness</u>
1 : $\mathbf{H}_{\text{seed}} \xleftarrow{\mathbb{S}} \{0, 1\}^\lambda$
2 : $\text{ker}_{\text{seed}} \xleftarrow{\mathbb{S}} \{0, 1\}^\lambda$
3 : $\text{perm}_{\text{seed}} \xleftarrow{\mathbb{S}} \{0, 1\}^\lambda$
<u>Construct PKP instance</u>
4 : $\mathbf{M} \leftarrow \text{ExpandMatrixM}(\mathbf{H}_{\text{seed}})$ in $\mathbb{F}_q^{m \times (n-m)}$
5 : $\mathbf{H} = [\mathbf{I}_m \ \mathbf{M}]$ in $\mathbb{F}_q^{m \times n}$
6 : $\mathbf{x}' \leftarrow \text{ExpandKernelVector}(\text{ker}_{\text{seed}}, \mathbf{H})$ in $\ker(\mathbf{H})$
7 : $\pi \leftarrow \text{ExpandPermutation}(\text{perm}_{\text{seed}})$ in \mathcal{S}_n
8 : $\mathbf{x} = \pi^{-1}(\mathbf{x}')$
9 : return ($\text{pk} = (\mathbf{H}_{\text{seed}}, \mathbf{x})$, $\text{sk} = (\text{perm}_{\text{seed}})$).

Algorithm 4.34: PERK.Sign(msg, sk, pk)
Public information and inputs
<p>Public information: Security parameter λ, public parameters $\text{param}_{\text{PKP}} := (q, m, n)$, $\text{param}_{\text{VOLE}} := (\tau, \mu, \rho)$, $\text{param}_{\text{TreePRG}} := (\tau, \kappa, N, T_{\text{open}}, w)$ where, $N := \tau 2^\kappa$ is the number of leaves of the VOLE commitment. Public key (expanded) $\text{pk} := (\mathbf{H}, \mathbf{x})$.</p> <p>Prover's input: Secret key sk represented as an array of positions of non-zero entries of the secret permutation matrix (row-wise), message msg to be signed.</p>
Output
Signature σ for message msg generated using prover's secret key sk .
<hr/> <p style="text-align: center;"><u>Initialization</u></p> <p>1 : $\tilde{\mu} := H_1(\text{pk} \parallel \text{msg} :: 2\lambda)$ 2 : $\text{rand} \xleftarrow{\\$} \{0, 1\}^{2\lambda}$ 3 : $(\text{mseed}, \text{salt}) := H_3(\text{sk} \parallel \tilde{\mu} \parallel \text{rand} :: 3\lambda) \in \{0, 1\}^\lambda \times \{0, 1\}^{2\lambda}$</p> <p style="text-align: center;"><u>VOLE construction, commitments and consistency checks</u></p> <p>4 : $(\mathbf{h}_{\text{com}}, \text{decom}, \mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}, \mathbf{u}, \mathbf{V}) := \text{VOLECommit}(\text{mseed}, \text{salt} :: \hat{\ell})$ 5 : $\text{ch}_1 := H_2^1(\tilde{\mu} \parallel \mathbf{h}_{\text{com}} \parallel \mathbf{c}_1 \parallel \dots \parallel \mathbf{c}_{\tau-1} \parallel \text{salt} :: 5\lambda + 64)$ 6 : $\tilde{\mathbf{u}} := \text{VOLEHash}(\text{ch}_1, \mathbf{u} :: \ell_{\text{VOLEHashMask}}) \in \{0, 1\}^{\ell_{\text{VOLEHashMask}}}$ 7 : $\tilde{\mathbf{V}} := \text{VOLEHash}(\text{ch}_1, \mathbf{V} :: \ell_{\text{VOLEHashMask}}) \in \{0, 1\}^{\ell_{\text{VOLEHashMask}} \times \rho}$ // hash column-wise 8 : $h_V := H_1(\tilde{\mathbf{V}} :: 2\lambda)$ // hash in column major order</p> <p style="text-align: center;"><u>Committing to witness and PKP proof</u></p> <p>9 : $(\mathbf{t}, \llbracket \beta \rrbracket, \llbracket \mathbf{z} \rrbracket^{(d)}, \llbracket \text{ColCheck} \rrbracket^{(d)}) := \text{P.VOLE-Permutation}(\text{sk}, (\llbracket u_{i,k} \rrbracket)_{i \in [n], k \in [\ell]})$ 10 : $\text{ch}_2 := H_2^2(\text{ch}_1 \parallel \tilde{\mathbf{u}} \parallel h_V \parallel \mathbf{t} :: 2\lambda)$ 11 : $\llbracket a \rrbracket := \text{P.Check-PKP}(\text{sk}, \text{pk}, (\llbracket u_{i,k} \rrbracket)_{i \in [n], k \in [\ell]}, (\llbracket u_{i',k'} \rrbracket)_{(i',k') \in [d-1] \times [\rho]}, \text{ch}_2)$</p> <p style="text-align: center;"><u>VOLE decommitments and opening VC</u></p> <p>12 : $\text{ctr} := 0$ 13 : while True: 14 : $\text{ch}_3 := H_2^3(\text{ch}_2 \parallel \llbracket a \rrbracket \parallel \text{ctr} :: \tau_0 \kappa_0 + \tau_1 \kappa_1 + w)$ // $\tau_0 \kappa_0 + \tau_1 \kappa_1 + w$ bits 15 : if $\text{ch}_3[\tau_0 \kappa_0 + \tau_1 \kappa_1 : \tau_0 \kappa_0 + \tau_1 \kappa_1 + w - 1] = 0^w$: 16 : $\mathbf{i}^* := \text{ChallDec}(\text{ch}_3[0 : \tau_0 \kappa_0 + \tau_1 \kappa_1 - 1])$ 17 : $(\text{pdecom}, (\text{com}_{e, \mathbf{i}^*[e]})_{e \in [\tau]}) := \text{VC.Open}(\text{decom}, \mathbf{i}^*)$ 18 : if the above output is not \perp break 19 : $\text{ctr} = \text{ctr} + 1$</p> <p>20 : return $\sigma := (\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}, \tilde{\mathbf{u}}, \mathbf{t}, \llbracket a \rrbracket, \text{pdecom}, (\text{com}_{e, \mathbf{i}^*[e]})_{e \in [\tau]}, \text{ch}_3, \text{ctr}, \text{salt})$</p>

Algorithm 4.35: PERK.Verify(pk, msg, σ)

Public information and inputs

Public information: Security parameter λ , public parameters $\text{param}_{\text{PKP}} := (q, m, n)$, $\text{param}_{\text{VOLE}} := (\tau, \mu, \rho)$, $\text{param}_{\text{TreePRG}} := (\tau, \kappa, N, T_{\text{open}}, w)$ where, $N := \tau 2^\kappa$ is the number of leaves of the VOLE commitment. Public key (expanded) $\text{pk} := (\mathbf{H}, \mathbf{x})$.
 Verifier's input: Public key (expanded) $\text{pk} := (\mathbf{H}, \mathbf{x})$, message msg , and a signature σ .

Output

Bit b indicating if the signature σ verifies as a valid signature for message msg with pk as the public key. If $b = 1$ the signature is accepted as valid, if $b = 0$ it is rejected.

Initialization

- 1: Parse $\sigma := (\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}, \tilde{\mathbf{u}}, \mathbf{t}, \llbracket a \rrbracket, \text{pdecom}, (\text{com}_{e, i^*[e]})_{e \in [0, \tau]}, \text{ch}_3, \text{ctr}, \text{salt})$
- 2: **if** $\text{ch}_3[\tau_0 \kappa_0 + \tau_1 \kappa_1 : \tau_0 \kappa_0 + \tau_1 \kappa_1 + w - 1] \neq 0^w$ **then return** 0
- 3: $\mu := \text{H}_1(\text{pk} \parallel \text{msg} :: 2\lambda)$
- 4: $i^* := \text{ChallDec}(\text{ch}_3[0 : \tau_0 \kappa_0 + \tau_1 \kappa_1 - 1])$

5: Reconstruct VOLEs and check commitments

- 6: $\text{out} := \text{VOLEReconstruct}(i^*, \text{pdecom}, \{\text{com}_{e, i^*}\}_{\tau, \tau}, \text{salt})$

- 7: **if** $\text{out} = \perp$:

- 8: **return** 0

- 9: **else** :

- 10: Parse out as $\text{out} := (h_{\text{com}}, \mathbf{Q}'_0, \dots, \mathbf{Q}'_{\tau-1})$.

- 11: $\bar{\text{ch}}_1 := \text{H}_2^1(\tilde{\mu} \parallel h_{\text{com}} \parallel \mathbf{c}_1 \parallel \dots \parallel \mathbf{c}_{\tau-1} \parallel \text{salt} :: 5\lambda + 64)$

12: Check VOLE's consistency

- 13: **for** $e \in [\tau]$

- 14: $(\delta_0, \dots, \delta_{\kappa_e-1}) := \text{BitDec}(i^*[e])$

- 15: **if** $\mu_e > \kappa_e$:

- 16: $\text{zeropad} := \mathbf{0}^{\ell_{\text{VOLEHashMask}} \times (\mu_e - \kappa_e)}$

- 17: **else** :

- 18: $\text{zeropad} := \varepsilon$

- 19: $\tilde{\mathbf{D}}_e := [\delta_0 \cdot \tilde{\mathbf{u}} \cdots \delta_{\kappa_e-1} \cdot \tilde{\mathbf{u}} \parallel \text{zeropad}] \in \{0, 1\}^{\ell_{\text{VOLEHashMask}} \times \mu_e}$

- 20: **if** $e = 0$ **then** $\mathbf{Q}_0 := \mathbf{Q}'_0$

- 21: **if** $e > 0$ **then** $\mathbf{Q}_e := \mathbf{Q}'_e \oplus [\delta_0 \cdot \mathbf{c}_e \cdots \delta_{\kappa_e-1} \cdot \mathbf{c}_e \parallel \text{zeropad}]$

- 22: **endfor**

- 23: $\mathbf{Q} := [\mathbf{Q}_0 \cdots \mathbf{Q}_{\tau-1}]$

- 24: $\tilde{\mathbf{Q}} := \text{VOLEHash}(\bar{\text{ch}}_1, \mathbf{Q} :: \ell_{\text{VOLEHashMask}}) \in \{0, 1\}^{\ell_{\text{VOLEHashMask}} \times \rho}$

- 25: $h_V := \text{H}_1(\tilde{\mathbf{Q}} \oplus [\tilde{\mathbf{D}}_0 \cdots \tilde{\mathbf{D}}_{\tau-1}] :: 2\lambda)$

- 26: $\bar{\text{ch}}_2 := \text{H}_2^2(\bar{\text{ch}}_1 \parallel \tilde{\mathbf{u}} \parallel h_V \parallel \mathbf{t} :: 2\lambda)$

27: Check PKP's consistency

- 28: $\bar{\text{ch}}_3 := \text{H}_2^3(\bar{\text{ch}}_2 \parallel \llbracket a \rrbracket \parallel \text{ctr} :: \tau_0 \kappa_0 + \tau_1 \kappa_1 + w)$

- 29: $b_V := \text{V.Check-PKP}(\text{seed}, \text{pk}, \mathbf{t}, \Delta, \mathbf{q}, \mathbf{g}, \bar{\text{ch}}_3 \parallel a)$

- 30: **if** $b_V = 1$ and $\text{ch}_3 = \bar{\text{ch}}_3$ **then**

- 31: **return** 1

- 32: **else return** 0

5 Parameter Sets and Sizes

We provide several parameter sets using the nomenclature PERK-X-Y where $X \in \{1, 3, 5\}$ denotes the security level and $Y \in \{\text{Short}, \text{Fast}\}$ refers to size / performance trade-off considered for the parameter set.

5.1 PKP parameters

The PKP parameters $\text{param}_{\text{PKP}} := (q, m, n)$ used in PERK are given in Table 1. Parameters were chosen to minimize the signature size while offering concrete bit-security of PKP above the NIST specified thresholds for category 1, 3 and 5. We give full details on estimating the security of PKP in Section [Section 7.2.1](#).

Instance	q	n	m
PERK-1	2048	64	27
PERK-3	2048	92	43
PERK-5	2048	118	59

Table 1: PKP parameters used in PERK

5.2 MPC and VOLE parameters

The tree parameters $\text{param}_{\text{TreePRG}} := (\tau, \kappa, N, T_{\text{open}}, w)$ are given in Table 2. The number of parties and iterations is governed by the knowledge soundness of the protocol. The MPC parameters are also chosen to guarantee a soundness probability of $2^{-\lambda}$ for $\lambda \in \{128, 192, 256\}$ for category 1, 3 and 5 respectively. Following common practice we propose two different parameter sets, a *short* variant using $N' \in \{2048, 4096\}$ and a *fast* variant using $N' \in \{128, 256\}$, where N' denotes the number of leaf nodes in each individual tree. The table below lists the total number of leaves $N := \tau 2^\kappa$ in all the trees together.

The VOLE related parameters $\text{param}_{\text{VOLE}} := (\tau, \mu, \rho)$ are given in Table 3. The parameter ρ denotes the dimension of finite field (\mathbb{F}_{2^ρ}) in which the VOLE correlations (seen as polynomials) reside. We also need to compute the PKP relation $\mathbf{HP}\mathbf{x}$ using the VOLE correlations therefore the fields should maintain the following tower relationship: $\mathbb{F}_2 \subset \mathbb{F}_q \subset \mathbb{F}_{2^\rho}$ and $\mathbb{F}_2 \subset \mathbb{F}_{2^\kappa} \subseteq \mathbb{F}_{2^\mu} \subset \mathbb{F}_{2^\rho}$.

The witness parameters $\text{param}_{\text{witness}} := (n, d, \ell_{\text{row}}, \ell, \ell_{\text{CZMask}}, \ell_{\text{VOLEHashMask}}, \hat{\ell})$ are given in Table 4. To prove the knowledge of the secret witness $\mathbf{P} \in \mathbb{F}_q^{n \times n}$, we are using $\hat{\ell}$ VOLE correlations that are computed using $\text{param}_{\text{witness}}$ where:

- n is the number of rows and columns in the permutation matrix ;

- $d := \lceil \log_4(n) \rceil$ is the degree of polynomials, which will help prove the knowledge of the witness using VOLE correlations ;
- $\ell_{\text{row}} := d + 6$ is the length of (compressed) witness (in bits) for each row of the secret permutation matrix ;
- $\ell := n \cdot \ell_{\text{row}}$ is the length of the (compressed) witness (in bits) for full secret permutation matrix ; This is essentially the witness size that affects the signature sizes ;
- $\ell_{\text{CZMask}} := (d - 1) \cdot \rho$ is the number of bits required to construct masking VOLE correlations in $\mathbb{F}_{2^\rho} \times \mathbb{F}_{2^\rho}$ required in [P.CheckZero](#) ;
- $\ell_{\text{VOLEHashMask}} := \lambda + B$ is the number of bits required for masking [VOLEHash](#). We always set $B := 16$ for all our parameter sets and instances ;
- $\hat{\ell} := \ell_{\text{VOLEHashMask}} + \ell + \ell_{\text{CZMask}}$ is the number of bits that should be communicated in each round. Therefore, the signature size is affected by the value $(\tau - 1) \cdot \hat{\ell}$.

Note that in [Tables 2 and 3](#), τ is the number of repetitions required to desired security level $\rho \geq \lambda$ (in our case $\rho > \lambda$ for all cases). Therefore, following criteria must always be satisfied:

- $\tau := \tau_0 + \tau_1 = \tau'_0 + \tau'_1$;
- $\rho := \tau'_0 \mu_0 + \tau'_1 \mu_1$;
- $\rho \geq \lambda$;
- $\tau_0 \kappa_0 + \tau_1 \kappa_1 + w - \log_2(d) \geq \lambda$.

Instance	τ	τ_0	τ_1	κ_0	κ_1	N	T_{open}	w
PERK-1-Short	11	11	0	11	0	3200	106	9
PERK-1-Fast	16	9	7	8	7	22528	110	
PERK-3-Short	16	8	8	12	11	5120	166	10
PERK-3-Fast	24	16	8	8	7	49152		
PERK-5-Short	22	8	14	12	11	7424	222	8
PERK-5-Fast	32	26	6	8	7	61440	220	

Table 2: BAVC parameters used in PERK

Instance	τ	τ'_0	τ'_1	μ_0	μ_1	ρ
PERK-1-Short	11	11	0	12	0	132
PERK-1-Fast	16	4	12	9	8	
PERK-3-Short	16	6	10	13	12	198
PERK-3-Fast	24	6	18	9	8	
PERK-5-Short	22	22	0	12	0	264
PERK-5-Fast	32	8	24	9	8	

Table 3: VOLE fields parameters used in PERK

Instance	B	d	ℓ_{row}	ℓ	ℓ_{CZMask}	$\hat{\ell}$
PERK-1		3	9	576	264	984
PERK-3	16	4	10	920	594	1722
PERK-5		4	10	1180	792	2244

Table 4: VOLE correlations parameters used in PERK

5.3 Signature and key sizes

Table 5 presents the public key, secret key, and signature sizes of PERK. The size of the public key pk is $\lambda + n \lceil \log_2(q) \rceil$ bits while the size of the secret key sk is λ bits. In practice, our implementations concatenate the public key within the secret key in order to respect the API provided by the NIST.

A PERK signature consists of:

- a salt, a hash value ch_3 and a counter ctr making a subtotal of $3\lambda + 66$ bits ;
- $(\tau - 1)$ VOLE correlation $\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}$ each of length $\hat{\ell}$ bits ;
- τ commitments $(\text{com}_{e, i^*[e]})_{e \in [\tau]}$ each of size 2λ bits ;
- a masked witness \mathbf{t} of size $n \cdot \ell_{\text{row}}$ bits ;
- a VOLEHash value $\tilde{\mathbf{u}}$ of size $\lambda + 16$ bits ;
- some opening information pdecom of size $\lambda \cdot T_{\text{open}}$ bits ;
- a polynomial $\llbracket a \rrbracket$ represented as coefficients with $\lceil \log_4(n) \rceil \cdot \rho$ bits.

Overall, for a security level λ , the signature size is given by:

$$\begin{aligned}
|\sigma| = & \underbrace{4\lambda + 82}_{\text{salt, ctr, ch}_3, \tilde{\mathbf{u}}} + \underbrace{\tau \cdot 2\lambda}_{\text{commitments}} + \underbrace{(\tau - 1) \cdot \hat{\ell}}_{\text{VOLE correlations}} \\
& + \underbrace{\lceil \log_4(n) \rceil \cdot \rho}_{\llbracket a \rrbracket} + \underbrace{\lambda \cdot T_{\text{open}}}_{\text{pdecom}} + \underbrace{n \cdot (\lceil \log_4(n) \rceil + 6)}_{\mathbf{t}}.
\end{aligned}$$

Instance	sk	pk	σ
PERK-1-Short	16 B	0.10 kB	3.48 kB
PERK-1-Fast	16 B	0.10 kB	4.32 kB
PERK-3-Short	24 B	0.15 kB	8.32 kB
PERK-3-Fast	24 B	0.15 kB	10.43 kB
PERK-5-Short	32 B	0.19 kB	14.83 kB
PERK-5-Fast	32 B	0.19 kB	18.22 kB

Table 5: Keys and signature sizes of PERK

6 Implementation and Performance Analysis

This section provides performance measurement of our PERK implementations.

Benchmark platform. The benchmarks have been done on a machine running Ubuntu 22.04.2 LTS, that has 64 GB of memory and an Intel[®] Core[™] i9-13900K @ 3.00 GHz for which the Hyper-Threading and Turbo Boost features were disabled. For each parameter set, the results have been obtained by computing the average from 10 and 500 random instances for the reference and optimized implementation, respectively. The scheme has been compiled with `gcc` (version 11.4.0) and the following third party libraries have been used: `XKCP` (commit 7fa59c0ec4) and `djbsort` (version 20190516).

Remark on the instantiation of PERK. The overall efficiency of PERK is determined to a large extent by the symmetric primitives employed in its construction. The PERK’s pseudorandom generator (PRG) may be instantiated using `AES/Rijndael` or `SHA3`, while hash functions are realized through `SHA3`. For benchmarking purposes, we provide results for two instantiations of the commitment scheme: one using `AES/Rijndael` and another using `SHA3`. It is important to emphasize that the choice of instantiation can lead to substantial variations in performance.

6.1 Reference implementation

The reference implementation is written in C and have been compiled using the compilation flags `-O3 -funroll-loops -flto`. The performances of our reference implementation on the aforementioned benchmark platform are described in Table 6-8.

Instance	Keygen	Sign	Verify
PERK-1-Short	50 K	13615 M	13608 M
PERK-1-Fast	50 K	1947 M	1943 M
PERK-3-Short	93 K	89057 M	89184 M
PERK-3-Fast	93 K	9360 M	9323 M
PERK-5-Short	146 K	111810 M	111622 M
PERK-5-Fast	130 K	13678 M	13591 M

Table 6: Performances of the reference implementation (in CPU cycles) - AES instantiation for both PRG and commitments

Instance	Keygen	Sign	Verify
PERK-1-Short	47 K	11242 M	11235 M
PERK-1-Fast	45 K	1610 M	1606 M
PERK-3-Short	93 K	74494 M	74613 M
PERK-3-Fast	93 K	7842 M	7805 M
PERK-5-Short	136 K	93676 M	93507 M
PERK-5-Fast	127 K	11486 M	11402 M

Table 7: Performances of the reference implementation (in CPU cycles) - AES for PRG and SHA3 for commitments

Instance	Keygen	Sign	Verify
PERK-1-Short	40 K	125 M	115 M
PERK-1-Fast	38 K	30.9 M	26.3 M
PERK-3-Short	82 K	401 M	354 M
PERK-3-Fast	83 K	121 M	82.3 M
PERK-5-Short	124 K	822 M	759 M
PERK-5-Fast	120 K	254 M	193 M

Table 8: Performances of the reference implementation (in CPU cycles) - SHA3 instantiation for both PRG and commitments

6.2 Optimized implementation

A constant-time optimized implementation leveraging **AVX2** instructions have been provided. Its performances on the aforementioned benchmark platform are described in Tables 9-11. The following optimization flags have been used during

compilation: -O3 -funroll-loops -march=native -mavx2 -mpclmul -msse4.2 -maes.

Instance	Keygen	Sign	Verify
PERK-1-Short	34 K	16.3 M	12.6 M
PERK-1-Fast	33 K	5.0 M	3.4 M
PERK-3-Short	66 K	143 M	130 M
PERK-3-Fast	65 K	33.5 M	23.2 M
PERK-5-Short	102 K	191 M	172 M
PERK-5-Fast	100 K	53.2 M	37.4 M

Table 9: Performances of the reference implementation (in CPU cycles) - AES instantiation for both PRG and commitments

Instance	Keygen	Sign	Verify
PERK-1-Short	33 K	27.8 M	24.3 M
PERK-1-Fast	33 K	6.6 M	5.0 M
PERK-3-Short	66 K	155 M	142 M
PERK-3-Fast	65 K	34.9 M	24.5 M
PERK-5-Short	101 K	236 M	217 M
PERK-5-Fast	99 K	58.4 M	42.7 M

Table 10: Performances of the reference implementation (in CPU cycles) - AES for PRG and SHA3 for commitments

Instance	Keygen	Sign	Verify
PERK-1-Short	35 K	47.4 M	43.5 M
PERK-1-Fast	34 K	9.4 M	7.9 M
PERK-3-Short	66 K	142 M	129 M
PERK-3-Fast	65 K	33.4 M	23.2 M
PERK-5-Short	103 K	274 M	255 M
PERK-5-Fast	102 K	63.1 M	47.2 M

Table 11: Performances of the reference implementation (in CPU cycles) - SHA3 instantiation for both PRG and commitments

6.3 Known Answer Test values

Known Answer Test (KAT) values have been generated using the script provided by the NIST. They are available in the folder `KATs` and files are the same for both reference and optimized implementation. In addition, examples with intermediate values have also been provided in these folders. Notice that one can generate the aforementioned test files using respectively the `kat` and `verbose` modes of our implementation. The procedure to follow in order to do so is detailed in the technical documentation.

7 Security Analysis

7.1 Security proof

Our signature scheme is strongly existentially unforgeable under chosen-message attacks (SUF-CMA-secure) in the random oracle model (ROM)⁶ under the assumption of hardness of PKP. The proof of SUF-CMA security, written below, happens in two stages:

- We first show that the slightly modified signature, which we call PERK', is existentially unforgeable under no-message attacks (EUF-NMA-secure) in the ROM assuming the hardness of the PKP problem.
- We then show that the signature scheme PERK is SUF-CMA-secure in the ROM by assuming that PERK' is EUF-NMA-secure in the ROM and some computational hardness of the functions.

We also discuss the beyond unforgeability features (BUFF) securities.

Notations. To simplify notations, we define the following variables:

- $\tilde{\mu} := H_1(\text{pk} \parallel \text{msg})$;
- $a_1 := (h_{\text{com}}, c_1, \dots, c_{\tau-1})$;
- $\text{ch}_1 := H_2^1(\tilde{\mu} \parallel a_1 \parallel \text{salt})$;
- $a_2 := (\tilde{\mathbf{u}}, h_V, \mathbf{t})$;
- $\text{ch}_2 := H_2^2(\text{ch}_1 \parallel a_2)$;
- $a_3 := \llbracket a \rrbracket$;
- $\text{ch}_3 := H_2^3(\text{ch}_2 \parallel a_3)$;
- $a_4 := (c_1, \dots, c_{\tau-1}, \tilde{\mathbf{u}}, \mathbf{t}, \llbracket a \rrbracket, \text{pdecom}, (\text{com}_{e, i[e]}), \text{ctr}, \text{salt})$.

By these notations, we have $\sigma = (\text{ch}_3, a_4)$. In what follows, we denote the internally reproduced values in the verifications `PERK.Verify(pk, msg, σ)` and `PERK.Verify(pk, msg*, σ)` by $\bar{\cdot}$ and $\bar{\cdot}^*$, respectively.

Preliminaries. Hereafter, we show the extractable-binding and multi-hiding properties of VOLE commitments. The following lemma helps us estimate the bound for one-wayness and collision-resistant property of random oracles.

Lemma 7.1 (Random oracle graph). *Let $H: \mathcal{X} \rightarrow \mathcal{Y}$ be the random oracle. We consider the following random oracle graph game between a challenger and an adversary:*

1. *The challenger initializes two sets \mathcal{V} and \mathcal{E} by \emptyset and runs the adversary.*
2. *The adversary can query the random oracle H in the following two ways:*
 - *The adversary queries $y \in \mathcal{Y}$ to the random oracle H ;*
 - *If $y \notin \mathcal{V}$, then the challenger adds node y to \mathcal{V} .*

⁶ If the Rijndael-based commitment is employed for Com_1 , then we also require the ideal-cipher model (ICM).

- The adversary queries $x \in \mathcal{X}$ to the random oracle H ; Let $H(x) = y$.
 - If $y \in \mathcal{V}$ but there is no edge $(x', y) \in \mathcal{E}$, then the adversary wins (because it breaks onewayness).
 - If an edge exists $(x', y) \in \mathcal{E}$ with $x' \neq x$, then the adversary wins (because it finds a collision).
 - Else, the challenger adds nodes x and y to \mathcal{V} and an edge $e = (x, y)$ from x to y to \mathcal{E} .

If the adversary makes at most Q queries to H , then the adversary's advantage is at most $Q^2/|\mathcal{Y}|$.

Extractable-binding property. Hereafter, we start by showing the extractable binding property of **VOLECommit**.

Lemma 7.2 (Extractable Binding). *Let \mathcal{A} be an adversary that makes $q_{\text{com},1}$ and $q_{\text{com},2}$ queries to Com_1 and Com_2 , respectively, where Com_1 and Com_2 are modeled as random oracles.⁷ We consider the following security game, which uses an extractor Ext defined later:*

1. $(h_{\text{com}}, \text{salt}) \leftarrow \mathcal{A}^{\text{Com}_1, \text{Com}_2}(1^\lambda, \text{commit})$.
2. $(\mathbf{u}_e^*, \mathbf{V}_e^*)_{e \in [\tau]} \leftarrow \text{Ext}(\mathcal{E}_{\text{Com}_1}^1, \mathcal{E}_{\text{Com}_2}^2, h_{\text{com}}, \text{salt})$, where $\mathcal{E}_{\text{Com}_1}^1$ and $\mathcal{E}_{\text{Com}_2}^2$ are the lists for the random oracles Com_1 and Com_2 .
3. $(\text{ch}_3, \text{pdecom}, \text{com}) \leftarrow \mathcal{A}^{\text{Com}_1, \text{Com}_2}(1^\lambda, \text{open})$.
4. $\mathbf{i} \leftarrow \text{ChallDec}(\text{ch})$.
5. For $e \in [\tau]$, $(\delta_{e,0}, \dots, \delta_{e,\mu_e-1}) := \text{BitDec}(\mathbf{i}^*[e])$.
6. $(\bar{h}_{\text{com}}, \mathbf{Q}_0, \dots, \mathbf{Q}_{\tau-1}) \leftarrow \text{VOLEReconstruct}^{\text{Com}_1, \text{Com}_2}(\mathbf{i}, \text{pdecom}, \text{com}, \text{salt})$.
7. Output False if
 - (a) $\bar{h}_{\text{com}} = \perp$ or $\bar{h}_{\text{com}} \neq h_{\text{com}}$; or
 - (b) $\mathbf{Q}_e = \mathbf{V}_e^* \oplus [\delta_{e,0}\mathbf{u}_e^* \cdots \delta_{e,\mu_e}\mathbf{u}_e^*]$ for all $i \in [\tau]$.
8. Output True otherwise.

Let $\text{AdvExt}^{\text{VOLE}} = \Pr[\mathcal{A} \text{ wins}]$ be \mathcal{A} 's advantage. We have

$$\text{AdvExt}^{\text{VOLE}} \leq (q_{\text{com},1} + N)^2/2^{2\lambda} + (q_{\text{com},2} + 1)^2/2^{2\lambda}.$$

Proof. The proof is essentially the same as that in FAEST's specification. We define the straight-line extractor Ext as follows:

1. Given h_{com} , find a preimage $\{\text{com}_{e,i}\}$ under Com_2 from the list $\mathcal{E}_{\text{Com}_2}^2$. If there is no preimage or multiple ones, then output \perp and abort.
2. For each $e \in [\tau]$ and $i \in [N_e]$: find preimages $\text{seed}_{e,i}$ of $\text{com}_{e,i}$ from $\mathcal{E}_{\text{Com}_1}^1$. If there is no such preimage, then set $\text{seed}_{e,i} = \perp$. If there are multiple preimages, then output \perp and abort.
3. For each $e \in [\tau]$, compute $(\mathbf{u}_e, \mathbf{V}_e)$ as follows:
 - Case 1: If Ext finds all preimages $\text{seed}_{e,i}$ for e , then it computes \mathbf{V}_e and \mathbf{u}_e honestly via **ConvertToVOLE**.

⁷ If the Rijndael-based commitment is employed for Com_1 , then we also require the ideal-cipher model (ICM).

- Case 2: If a single preimage is missing, then set $\Delta_e = j^*$. It then computes $(\mathbf{u}_e, \mathbf{q}_{e,0}, \dots, \mathbf{q}_{e,\mu_e-1})$ via [ConvertToVOLE](#) with permuted seed with Δ_e and sets $\mathbf{Q}_e = [\mathbf{q}_{e,0} \cdots \mathbf{q}_{e,\mu_e-1}]$. It sets $\mathbf{V}_e := \mathbf{Q}_e \oplus [\delta_{e,0}\mathbf{u}_e \cdots \delta_{e,\mu_e-1}\mathbf{u}_e]$.
- Case 3: If multiple preimages are missing, then output \perp and abort.

4. Output $(\mathbf{u}_e, \mathbf{V}_e)_{e \in [\tau]}$.

5. If it fails to extraction, then it adds the image that missed the preimage to $\mathcal{V}_{\text{Com}}^1$ or $\mathcal{V}_{\text{Com}}^2$.

Let **Fail** be the event that the extractor fails to output $(\mathbf{u}_e, \mathbf{V}_e)_e$. Because of the random oracle graph game, we have

$$\Pr[\text{Fail}] \leq (q_{\text{com},1} + N)^2 / 2^{2\lambda} + (q_{\text{com},2} + 1)^2 / 2^{2\lambda},$$

where we add N for $q_{\text{com},1}$ since we have at most N commitments missing preimages.

If $\neg\text{Fail}$, then h_{com} uniquely determines $(\text{com}_{e,i})_{e,i}$ and each $\text{com}_{e,i}$ uniquely determines $\text{seed}_{e,i}$ for $e \in [\tau]$ and $i \in [N_e]$. (NOTE: For each $e \in [\tau]$, at most one of $\text{seed}_{e,i}$ can be \perp .) In both cases (case 1 or case 2), the extraction is perfect, as explained in FAEST's specification. Thus, the extractor's failing probability is at most $((q_{\text{com},1} + N)^2 + (q_{\text{com},2} + 1)^2) \cdot 2^{-2\lambda}$.

Multi-hiding property. We next show the multi-hiding property of [VOLECommit](#). To do so, we first define the simulation algorithm, [SimVOLECommit](#).

Algorithm 7.1: SimVOLECommit(\mathbf{i} , salt, $\hat{\ell}$)
Public information and inputs
Public information: A number of iterations τ , a number of parties $N = \sum_{e=0}^{\tau-1} N_e$, $\hat{k}_e = \log_2(N_e)$, $\rho = k_0\tau_0 + k_1\tau_1$ Prover's input: \mathbf{i} and salt and $\hat{\ell}$. We assume that \mathbf{i} is accepted.
Output
A commitment $\mathbf{h}_{\text{com}} \in \{0, 1\}^{2\lambda}$, a sibling path pdecom , unopened commitments $(\text{com}_{e, \mathbf{i}[e]})_e$, VOLE corrections $(\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1})$, and VOLE correlation secrets \mathbf{u}
<pre> 1 : hidden := {N - 1 + $\psi(e, \mathbf{i}[e])$: $e \in [0, \tau)$} 2 : opened := {N - 1, \dots, 2N - 2} \setminus hidden 3 : for i from N - 2 downto 0 do 4 : if $2i + 1 \in \text{opened}$ and $2i + 2 \in \text{opened}$ then 5 : opened = opened \cup {i} 6 : endfor 7 : nodes[0], \dots, nodes[2N - 2] = $\emptyset, \dots, \emptyset$ 8 : for $i \in [0, N - 1)$ do 9 : if $2i + 1 \in \text{opened}$ and $2i + 2 \in \text{opened}$ then 10 : (nodes[2i + 1], nodes[2i + 2]) \leftarrow PRG(nodes[i], salt) 11 : else 12 : (nodes^(j)[2i + 1], nodes^(j)[2i + 2]) \leftarrow {0, 1}^{2λ} 13 : for $e \in [0, \tau)$ do 14 : for $i \in [0, N_e)$ do 15 : seed_{e, i} = nodes[N - i + $\psi(e, i)$] 16 : if $i = \mathbf{i}[e]$ then 17 : com_{e, i} \leftarrow $\{0, 1\}^{2\lambda}$ 18 : else 19 : com_{e, i} \leftarrow Com₁(salt, e, i, seed_{e, i}) 20 : h_{com} \leftarrow Com₂(salt, {com_{e, i}}) 21 : decom = (nodes, (com_{e, i})) 22 : (pdecom, (com_{$e, \mathbf{i}[e]$})_{e}) \leftarrow VC.Open(decom, \mathbf{i}) 23 : ($\mathbf{u}, \mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}$) \leftarrow $(\mathbb{F}_2^{\hat{\ell}})^\tau$ 24 : return (h_{com}, pdecom, (com_{$e, \mathbf{i}[e]$})_{e}, $\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}, \mathbf{u}$). </pre>

Lemma 7.3 (Multi Hiding). *Let us consider the following Q -multi-hiding game for VOLECommit between an adversary \mathcal{A} and a challenger. Let $b^* \in \{0, 1\}$.*

1. For $j \in [Q]$:
 - (a) Take $(r^{(j)}, \text{salt}^{(j)})$ uniformly at random.
 - (b) Take a random challenge $\text{ch}_3^{(j)} \leftarrow \{0, 1\}^{\tau_0\kappa_0 + \tau_1\kappa_1 + w}$ until it is accepted:

- i. If $\text{ch}_3^{(j)}[\tau_0\kappa_0 + \tau_1\kappa_1 : \tau_0\kappa_0 + \tau_1\kappa_1 + w - 1] \neq 0^w$, then re-sample.
 - ii. Compute $\mathbf{i}^{(j)}$ by $\text{ChallDec}(\text{ch}_3^{(j)}[0 : \tau_0\kappa_0 + \tau_1\kappa_1 - 1])$. If the indices lead to $\text{len}(\text{revealed}) > T_{\text{open}}$,⁸ then re-sample.
- (c) If $b^* = 0$:
- i. Compute $(h_{\text{com}}^{(j)}, \text{decom}^{(j)}, \mathbf{c}_1^{(j)}, \dots, \mathbf{c}_{\tau-1}^{(j)}, \mathbf{u}^{(j)}, \mathbf{V}^{(j)}) \leftarrow \text{VOLECommit}(r^{(j)}, \text{salt}^{(j)}, \hat{\ell})$, where $\text{decom}^{(j)} = (\text{nodes}^{(j)}, (\text{com}_{e, \mathbf{i}^{(j)}}^{(j)}))$.
 - ii. Compute $(\text{pdecom}^{(j)}, (\text{com}_{e, \mathbf{i}^{(j)}[e]}^{(j)})) \leftarrow \text{VC.Open}(\text{decom}^{(j)}, \mathbf{i}^{(j)})$.
- (d) If $b^* = 1$:
- i. Compute $(h_{\text{com}}^{(j)}, \text{pdecom}^{(j)}, (\text{com}_{e, \mathbf{i}^{(j)}[e]}^{(j)})_e, \mathbf{c}_1^{(j)}, \dots, \mathbf{c}_{\tau-1}^{(j)}, \mathbf{u}^{(j)}) \leftarrow \text{SimVOLECommit}(\mathbf{i}^{(j)}, \text{salt}^{(j)})$.
2. $b \leftarrow \mathcal{A}\left((h_{\text{com}}^{(j)}, \text{pdecom}^{(j)}, (\text{com}_{e, \mathbf{i}^{(j)}[e]}^{(j)})_e, \mathbf{c}_1^{(j)}, \dots, \mathbf{c}_{\tau-1}^{(j)}, \mathbf{u}^{(j)}, \text{ch}_3^{(j)})_{j \in [Q]}\right)$.
 3. Output True if $b = b^*$; otherwise, output False.

Then, the adversary's advantage

$$\text{AdvHide}^{\text{VOLE}}[Q] := |\Pr[\mathcal{A} \text{ wins} \mid b^* = 0] - \Pr[\mathcal{A} \text{ wins} \mid b^* = 1]|$$

is at most $\hat{k} \cdot \text{AdvPRG}^{\text{PRG}_1}[Q, \tau] + \text{AdvPRF}^{\text{PRG}_2, \text{Com}_1}[Q, \tau]$, where $\hat{k} = \lceil \log_2(N) \rceil + 1$.

Proof. We use the standard GGM tree argument.

G₀: This is the original game with $b^* = 0$. Thus, we have $\Pr[W_0] = \Pr[\mathcal{A} \text{ wins} \mid b^* = 0]$.

G₁: We gradually replace GGM trees constructed in **VC.Commit** by following $\mathbf{i}^{(j)}$. We define $G_{1,k}$ for $k = 0, \dots, \hat{k}$ as follows:

G_{1,k}: We modify **VC.Commit** invoked from $\text{VOLECommit}(r^{(j)}, \text{salt}^{(j)}, \hat{\ell})$ as follows:

1. Define $\text{opened}^{(j)} := \{N - 1, \dots, 2N - 2\} \setminus \{N - 1 + \psi(e, \mathbf{i}^{(j)}[e]) : e \in [\tau]\}$.
2. For i from $N - 2$ downto 0: if both $2i + 1, 2i + 2 \in \text{opened}^{(j)}$ then $\text{opened}^{(j)} := \text{opened}^{(j)} \cup \{i\}$.
3. For $i \in [N - 1]$:
 - (a) If $i \in [2^k]$ and $i \notin \text{opened}^{(j)}$, then $(\text{nodes}^{(j)}[2i + 1], \text{nodes}^{(j)}[2i + 1]) \leftarrow \{0, 1\}^{2\lambda}$;
 - (b) else, $(\text{nodes}^{(j)}[2i + 1], \text{nodes}^{(j)}[2i + 1]) \leftarrow \text{PRG}_1(\text{salt}, \text{nodes}^{(j)}[i] :: 2\lambda)$.

We note that the width of co-path is at most τ . Hence, this introduces the difference at most $\text{AdvPRG}^{\text{PRG}_1}[Q, \tau]$.

We note that, at the final game $G_{1,\hat{k}}$, the hidden $\text{seed}_{e, \mathbf{i}^{(j)}[e]}^{(j)} = \text{nodes}^{(j)}[N - 1 + \psi(e, \mathbf{i}^{(j)}[e])]$ are chosen uniformly at random.

G₂: We next replace $(\text{com}_{e, \mathbf{i}^{(j)}[e]}^{(j)})^{(j)}$ with random string and $\mathbf{r}_{0, \mathbf{i}^{(j)}[e]}^{(j)}$ with random vector for all $e \in [0, \tau]$ and $j \in [Q]$ in the computation of **ConvertToVOLE** $(N_e, (\text{seed}_{e, \mathbf{i}^{(j)}[e]}^{(j)})_{i \in [0, N_e]}, \text{salt}^{(j)})$ invoked by $\text{VOLECommit}(r^{(j)}, \text{salt}^{(j)}, \hat{\ell})$. This is justified by the joint PRF security of PRG_2 in **ConvertToVOLE** and Com_1 since $\text{seed}_{e, \mathbf{i}^{(j)}[e]}^{(j)}$ are hidden from the adversary. Thus, the difference is at most $\text{AdvPRF}^{\text{PRG}_2, \text{Com}_1}[Q, \tau]$.

⁸ See **VC.Open** $(\text{decom}, \mathbf{i}^*)$. Note that decom is independent of the computation of revealed.

G₃: Next, we replace $\mathbf{u}_e^{(j)}$ for all e and j with random vectors in the computation of $\text{ConvertToVOLE}(N_e, (\text{seed}_{e,i}^{(j)})_{i \in [0, N_e]}, \text{salt}^{(j)}, \hat{\ell})$ invoked by $\text{VOLECommit}(r^{(j)}, \text{salt}^{(j)}, \hat{\ell})$. Since $\mathbf{u}_e^{(j)} = \sum_{i \in [0, N_e]: i \neq i^{(j)}[e]} \text{PRG}_2(\text{salt}^{(j)}, \text{seed}_{e,i}^{(j)} :: \hat{\ell}) \oplus \mathbf{r}_{0, i^{(j)}[e]}^{(j)}$, this does not change the distribution from the previous game and we have $\Pr[W_2] = \Pr[W_3]$. Now, in this game, every $\mathbf{u}_e^{(j)}$ for $e \in [\tau]$ is random. Thus, $\mathbf{c}_i^{(j)} = \mathbf{u}_0^{(j)} \oplus \mathbf{u}_i^{(j)}$ is also random and $(\mathbf{u}^{(j)}, \mathbf{c}_1^{(j)}, \dots, \mathbf{c}_{\tau-1}^{(j)})$ is random. Therefore, this game is equivalent to the game for $b^* = 1$, and we have $\Pr[W_3] = \Pr[\mathcal{A} \text{ wins} \mid b^* = 1]$.

Wrapping up, we have $\hat{k} \cdot \text{AdvPRG}^{\text{PRG}_1}[Q, \tau] + \text{AdvPRF}^{\text{PRG}_2, \text{Com}_1}[Q, \tau]$ as the upper bound.

Security proofs for the EUF-NMA security. We here consider the security of the slightly modified signature scheme denoted by PERK' . Concretely speaking, we replace “(mseed, salt) \leftarrow H₃(sk||μ||rand)” in line 3 of $\text{PERK}.\text{Sign}$ with “(mseed, salt) \leftarrow $\{0, 1\}^{\lambda+2\lambda}$ ”. Our proof mainly follows that in FAEST’s specification, but we modify several points to adopt their proof in our setting, e.g., showing the formal proof for grinding and optimizations.

Theorem 7.1 (EUF-NMA security in the ROM). *Let \mathcal{B} be an adversary against the EUF-NMA security of PERK' . Let $q_{\text{com},1}, q_{\text{com},2}, q_1, q_{2,1}, q_{2,2}, q_{2,3}$, and q_4 be the number of queries \mathcal{B} made to $\text{Com}_1, \text{Com}_2, \text{H}_1, \text{H}_2^1, \text{H}_2^2, \text{H}_2^3$, and H_4 , respectively. Suppose that VOLEHash is an ϵ_v -almost universal hash family (Lemma 4.1). We have another adversary \mathcal{C} against the PKP assumption (Definition 2.10) such that*

$$\begin{aligned} \text{Adv}_{\mathcal{B}}^{\text{euf-nma}} &\leq \text{AdvOW}_{\mathcal{C}} + \text{AdvPR}^{\text{ExpandKernelVector}} + \text{AdvPR}^{\text{ExpandPermutation}} \\ &\quad + (q_{\text{com},1} + q_{\text{com},2} + q_1 + (N+2)q_{2,1} + 2q_{2,2} + q_{2,3})^2 \cdot 2^{-2\lambda} \\ &\quad + q_{2,1} \binom{\tau}{2} \cdot \epsilon_v + q_{2,2} \cdot 2^{-\rho} + (q_4 + q_{2,2})^2 \cdot 2^{-(3\lambda+64)} \\ &\quad + q_{2,3}d \cdot 2^{-(\tau_0\kappa_0 + \tau_1\kappa_1 + w)} + 2^{-\lambda+1}. \end{aligned}$$

The running time of \mathcal{C} is about that of \mathcal{B} .

We follow the games defined in FAEST’s specification and show the bounds of the differences between the games.

G₁: The original EUF-NMA game. We have

$$\Pr[W_1] = \text{Adv}_{\mathcal{B}}^{\text{euf-nma}}.$$

We note that the random oracles and commitment are implemented by lazy sampling and the lists, $\mathcal{E}_{\text{Com}}^1, \mathcal{E}_{\text{Com}}^2, \mathcal{E}_1, \mathcal{E}_2^1, \mathcal{E}_2^2, \mathcal{E}_2^3$, and \mathcal{E}_4 . Here, we use \mathcal{E} because they will be edge sets of the random oracle graphs.

G₂: In this game, we consider the random oracle graph games. To do so, we additionally consider the vertex sets $\mathcal{V}_{\text{Com}}^1, \mathcal{V}_{\text{Com}}^2, \mathcal{V}_1, \mathcal{V}_2^1, \mathcal{V}_2^2$, and \mathcal{V}_2^3 . If the adversary wins the random oracle graph game, then we abort.

Let Fail_i be the event that the adversary wins the random oracle graph game. For any G_i , we have $\Pr[W_i] = \Pr[\text{Fail}_i] \cdot \Pr[W_i \mid \text{Fail}_i] + \Pr[\neg \text{Fail}_i] \cdot \Pr[W_i \mid \neg \text{Fail}_i] \leq \Pr[\text{Fail}_i] + \Pr[W_i \mid \neg \text{Fail}_i]$. We note that if Fail_2 does not occur, then G_2 is equivalent to G_1 . Thus, we have $\Pr[W_1] = \Pr[W_2 \mid \neg \text{Fail}_2]$.

G₃: We next use *the extractor* Ext in the proof of [Lemma 7.2](#) on every query $(\tilde{\mu}, h_{\text{com}}, \dots, \text{salt})$ to H_2^1 to obtain $(\mathbf{u}_e, \mathbf{V}_e)_{e \in [\tau]}$. If the adversary's forgery is valid but one of the extracting test fails, then the adversary *loses*. Because of the modification introduced G_2 , we have $\Pr[W_2 \mid \neg \text{Fail}_2] = \Pr[W_3 \mid \neg \text{Fail}_3]$. But, the number of queries blows up because of the queries the extractor made. We have

$$\bar{q}_{\text{com},1} = q_{\text{com},1} + Nq_{2,1} \text{ and } \bar{q}_{\text{com},2} = q_{\text{com},2} + q_{2,1},$$

since Ext is invoked at most $q_{2,1}$ times and each invocation adds at most N queries or one query to Com_1 or Com_2 , respectively.

G₄: We next modify H_2^1 ; if the query is extractable, then it runs the VOLE consistency checks with extracted $(\mathbf{u}_e, \mathbf{V}_e)_{e \in [\tau]}$. While we omit the details of the check, FAEST's specification shows that the probability that the consistency check fails $\epsilon_v \binom{\tau}{2}$, where ϵ_v is universality of [VOLEHash](#), by using [[BBD⁺23c](#), Thm.2]. Since we take [VOLEHash](#) from FAEST, we have the same probability. Thus, we have

$$\Pr[W_3 \mid \neg \text{Fail}_3] \leq \Pr[W_4 \mid \neg \text{Fail}_4] + q_{2,1} \binom{\tau}{2} \epsilon_v,$$

where ϵ_v is universality of [VOLEHash](#).

G₅: We next modify H_2^2 to force the chain of hash values. On query $(\text{ch}_1, h_V, \dots)$ to H_2^2 , we do as follows:

- If ch_1 has no preimage, then query ch_1 under H_2^1 ; that is, if $\text{ch}_1 \notin \mathcal{V}_2^1$, then add ch_1 to \mathcal{V}_2^1 .
- If h_V has no preimage, then query h_V under H_1 ; that is, if $h_V \notin \mathcal{V}_1$, then add h_V to \mathcal{V}_1 .

We have $\Pr[W_4 \mid \neg \text{Fail}_4] = \Pr[W_5 \mid \neg \text{Fail}_5]$. By this modification, the bound of the random oracle queries to the random oracle graph game are

$$\bar{q}_1 = q_1 + q_{2,2} \text{ and } \bar{q}_{2,1} = q_{2,1} + q_{2,2}.$$

G₆: We modify H_2^2 as follows: On each new query $(\text{ch}_1, h_V, \dots)$ to H_2^2 , if the query related to ch_1 is extractable and VOLE-consistent, then do as follows:

1. Take a random $\text{ch}_2 \leftarrow_{\$} \{0, 1\}^{3\lambda+64}$. If ch_2 is already queried to H_4 , then it aborts.
2. If the ZK soundness check below fails, then abort.
 - (a) Let $\alpha \xleftarrow{\$, \text{ch}_2} \mathbb{F}_{2^\rho}^{cn+m}$. That is, we take a random sample $\alpha \leftarrow_{\$} \mathbb{F}_{2^\rho}^{cn+m}$ and put (ch_2, α) to \mathcal{E}_4 for H_4 .

- (b) Compute $\mathbf{e} \in \mathbb{F}_{2^\rho}^{cn+m}$, which is a vector consisting of the degree d coefficients of $f_j(X)$ for $j = 0, \dots, cn + m - 1$ constructed from the extracted witness.
 - (c) If $\mathbf{e} \neq \mathbf{0}$ but $\sum_{j=0}^{cn+m-1} \alpha_j e_j = 0$, then the output “fail”. Otherwise, output “success”.
3. Otherwise, look up a preimage $\tilde{\mathbf{V}}$ of h_V . If a preimage does not exist, then abort.
 4. Otherwise, return ch_2 .

On the collision test for step 1, we have a bound $(q_4 + q_{2,2})^2 / 2^{3\lambda+64}$. If α is uniformly at random, then the probability that the check in step 2 fails is at most $\epsilon_{zk} = 1/2^\rho$.

Taking a union bound, we have

$$\Pr[W_5 \mid \neg\text{Fail}_5] \leq \Pr[W_6 \mid \neg\text{Fail}_6] + q_{2,2} \cdot 2^{-\rho} + (q_4 + q_{2,2})^2 \cdot 2^{-(3\lambda+64)}.$$

G₇: We next modify H_2^3 as follows: On query $(\text{ch}_2, \llbracket a \rrbracket, \text{ctr})$ to H_2^3 , we do as follows:

- If ch_2 has no preimage, then query ch_2 under H_2^2 ; that is, if $\text{ch}_2 \notin \mathcal{V}_2^2$, then add ch_2 to \mathcal{V}_2^2 .

We have $\Pr[W_6 \mid \neg\text{Fail}_6] = \Pr[W_7 \mid \neg\text{Fail}_7]$. The number of random oracle queries is

$$\bar{q}_{2,2} = q_{2,2} + q_{2,3}.$$

G₈: We modify the handling H_2^3 on a query $(\text{ch}_2, \llbracket a \rrbracket, \text{ctr})$ as follows:

1. If ch_2 has no preimage, then abort. Otherwise, let $(\text{ch}_1, h_V, \dots)$ be the preimage of ch_2 .
2. If h_V has no preimage, then abort.
3. Otherwise, extract \mathbf{u} and \mathbf{V} , and define the witness \mathbf{w} .
4. Sample $\text{ch}_3 \leftarrow \{0, 1\}^{\tau_0 \kappa_0 + \tau_1 \kappa_1 + w}$.
5. If \mathbf{w} doesn't satisfy the constraints and the bad event occurs, then abort the game. The bad event is the event that, for uniformly random $\text{ch}_3 \in \{0, 1\}^{\tau_0 \kappa_0 + \tau_1 \kappa_1 + w}$ in step 3, 1) ch_3 passes the verification test (T_{open} and 0^w check) and 2) for uniformly random $\text{ch}_3 \in \{0, 1\}^{\tau_0 \kappa_0 + \tau_1 \kappa_1 + w}$, we have $\Delta = \Delta'$, where Δ is computed from ch_3 and Δ' is computed from the query $\llbracket a \rrbracket$ by the adversary.

The probability 1) is $|\#\text{accepted challenges}| / 2^{\tau_0 \kappa_0 + \tau_1 \kappa_1 + w}$ and the probability 2) is at most $d / |\#\text{accepted challenges}|$ since $\llbracket a \rrbracket$ is treated as $(d-1)$ -degree polynomial $a_{d-1}X^{d-1} + \dots + a_1X + a_0$ and the verification algorithm checks if $q_f + \sum_{i=0}^{d-2} q_{s_i} \cdot \Delta^i = \sum_{i=0}^{d-1} a_i \Delta^i$ in V.CheckZero invoked by V.CheckPKP . Thus, we have the bound

$$\begin{aligned} \Pr[W_7 \mid \neg\text{Fail}_7] &\leq \Pr[W_8 \mid \neg\text{Fail}_8] + q_{2,3} \cdot \frac{d}{|\#\text{accepted challenges}|} \cdot \frac{|\#\text{accepted challenges}|}{2^{\tau_0 \kappa_0 + \tau_1 \kappa_1 + w}} \\ &= \Pr[W_8 \mid \neg\text{Fail}_8] + q_{2,3} \cdot d \cdot 2^{-(\tau_0 \kappa_0 + \tau_1 \kappa_1 + w)}. \end{aligned}$$

G₉: We next modify the key-generation algorithm. In the experiment, the key-generation algorithm chooses $\mathbf{x}' \xleftarrow{\$} \ker(\mathbf{H})$ and $\pi \xleftarrow{\$} \mathcal{S}_n$ instead of computing $\mathbf{x}' \leftarrow \text{ExpandKernelVector}(\ker_{\text{seed}}, \mathbf{H})$ and $\pi \leftarrow \text{ExpandPermutation}(\text{perm}_{\text{seed}})$.

$$\Pr[W_8 \mid \neg\text{Fail}_8] \leq \Pr[W_9 \mid \neg\text{Fail}_9] + \text{AdvPR}^{\text{ExpandKernelVector}} + \text{AdvPR}^{\text{ExpandPermutation}}.$$

We then discuss the evaluation of $\Pr[\text{Fail}_9]$ and the reduction to the PKP problem.

Evaluation of $\Pr[\text{Fail}_9]$. The numbers of queries to the random oracles are now

$$\begin{aligned} \bar{q}_{\text{com},1} &= q_{\text{com},1} + q_{2,1}N, \bar{q}_{\text{com},2} = q_{\text{com},2} + q_{2,1}, \\ \bar{q}_1 &= q_1 + q_{2,2}, \bar{q}_{2,1} = q_{2,1} + q_{2,2}, \text{ and } \bar{q}_{2,2} = q_{2,2} + q_{2,3}. \end{aligned}$$

Thus, the advantage of the random oracle games is at most

$$\begin{aligned} \Pr[\text{Fail}_9] &\leq \bar{q}_{\text{com},1}^2/2^{2\lambda} + \bar{q}_{\text{com},2}^2/2^{2\lambda} + \bar{q}_1^2/2^{2\lambda} + \bar{q}_{2,1}^2/2^{5\lambda+64} + \bar{q}_{2,2}^2/2^{2\lambda} \\ &\leq (q_{\text{com}} + q_{\text{com},2} + q_1 + (N+2)q_{2,1} + 2q_{2,2} + q_{2,3})^2/2^{2\lambda}. \end{aligned}$$

Reduction to PKP.

We construct a reduction algorithm \mathcal{C} using \mathcal{B} in the final game G_9 conditioned on that Fail_9 does not occur. The reduction algorithm \mathcal{C} against our PKP assumption (Definition 2.10) is defined as follows:

1. It is given a random seed $\mathbf{H}_{\text{seed}} \xleftarrow{\$} \{0,1\}^\lambda$ and $\mathbf{x}' \in \mathbb{F}_q^n$, where $\mathbf{M} \leftarrow \text{ExpandMatrixM}(\mathbf{H}_{\text{seed}})$, $\mathbf{H} := [\mathbf{I}_m \ \mathbf{M}] \in \mathbb{F}_q^{m \times n}$, $\mathbf{x}' \xleftarrow{\$} \ker(\mathbf{H})$, $\pi \xleftarrow{\$} \mathcal{S}_n$, and $\mathbf{x} := \pi^{-1}(\mathbf{x}')$. It wants to output $\tilde{\pi}$ such that $\mathbf{H}(\tilde{\pi}(\mathbf{x})) = \mathbf{0}$.
2. It sets $\text{pk} = (\mathbf{H}_{\text{seed}}, \mathbf{x})$ and run \mathcal{B} in G_9 .
3. Finally, \mathcal{B} outputs (m^*, σ^*) and stops. If \mathcal{B} wins and Fail_9 does not occur, then \mathcal{C} extracts the witness \mathbf{P} from the random oracle queries, and outputs $\tilde{\pi}$ corresponding to \mathbf{P} .

Since the simulation of \mathcal{C} is perfect, if \mathcal{B} wins the game, then \mathcal{C} can extract the witness \mathbf{P} from \mathcal{B} 's queries to the random oracles. Thus, we have

$$\Pr[W_9 \mid \neg\text{Fail}_9] \leq \text{AdvOW}_{\mathcal{C}}.$$

Security proofs for the SUF-CMA security. We follow the proof in FAEST's specification while we consider an optimized version and we will repair some gaps in the original proof. Thus, we need to consider *a special form of Fiat-Shamir with aborts*, which only samples ch_3 instead of the whole. Hence, we enhanced the repaired proofs for 3-round Fiat-Shamir with aborts in Devevey *et al.* [DFPS23] and Barbosa *et al.* [BBD⁺23a]. To consider SUF-CMA security, we also employ the techniques in [KX24b].

Theorem 7.2 (SUF-CMA security in the ROM). *Let \mathcal{A} be an adversary against the EUF-CMA security of PERK. Let $Q_{\text{prg},1}$, $Q_{\text{com},1}$, $Q_{\text{com},2}$, Q_1 , $Q_{2,1}$, $Q_{2,2}$, and $Q_{2,3}$ be the number of queries \mathcal{A} made to PRG_1 , Com_1 , Com_2 , H_1 , H_2^1 ,*

H_2^2 , and H_2^3 , respectively. Let Q_{sig} be the number of queries \mathcal{A} made to the signing oracle. Let \tilde{Q}_{sig} be the number of queries the signing oracle made to H_2^3 . We then have an adversary \mathcal{B} against the EUF-NMA security of PERK' satisfying

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{suf-cma}} &\leq \text{Adv}_{\mathcal{B}}^{\text{euf-nma}} + Q_{\text{sig}} \cdot \text{AdvPRF}^{\text{H}_3} \\ &\quad + \text{AdvColl}^{\text{PRG}_1} [Q_{\text{prg},1} + (N-1)Q_{\text{sig}}] \\ &\quad + \text{AdvColl}^{\text{Com}_1} [Q_{\text{com}} + NQ_{\text{sig}}] + \text{AdvColl}^{\text{Com}_2} [Q_{\text{com},2} + Q_{\text{sig}}] \\ &\quad + \text{AdvColl}^{\text{H}_1} [Q_1 + 2Q_{\text{sig}}] + \text{AdvColl}^{\text{H}_2^1} [Q_{2,1} + Q_{\text{sig}}] \\ &\quad + \text{AdvColl}^{\text{H}_2^2} [Q_{2,2} + Q_{\text{sig}}] + \text{AdvColl}^{\text{H}_2^3} [Q_{2,3} + \tilde{Q}_{\text{sig}}] \\ &\quad + Q_{\text{sig}}(Q_{2,1} + Q_{\text{sig}})2^{-2\lambda} + Q_{\text{sig}}(Q_{2,2} + Q_{\text{sig}})2^{-(5\lambda+64)} + \tilde{Q}_{\text{sig}}(Q_{2,3} + \tilde{Q}_{\text{sig}})2^{-2\lambda} \\ &\quad + \hat{k} \cdot \text{AdvPRG}^{\text{PRG}_1} [Q_{\text{sig}}, \tau] + \text{AdvJPRF}^{\text{PRG}_2, \text{Com}_1} [Q_{\text{sig}}, \tau] \\ &\quad + \text{AdvNI}^{\text{Com}_1} [\tau Q_{\text{sig}}]. \end{aligned}$$

Remark 7.1. Let us discuss the order of \tilde{Q}_{sig} . Let $r := |\# \text{ accepted challenges}|/2^{70\kappa_0 + \tau_1\kappa_1 + w}$. We let $\tilde{Q}_{\text{sig}} := (2/r) \cdot Q_{\text{sig}}$. This \tilde{Q}_{sig} gives us $\Pr[\# \text{ success is less than } Q_{\text{sig}}] \leq \exp(-Q_{\text{sig}}/4)$ and the right-hand side is negligible if $Q_{\text{sig}} = \omega(\log(\lambda))$. (See e.g., [KX24a, Pf. of Thm.1].)

Corollary 7.1 (EUF-CMA security in the ROM). *Let the parameters be the same as Theorem 7.2. We then have an adversary \mathcal{B} against the EUF-NMA security of PERK' satisfying*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{euf-cma}} &\leq \text{Adv}_{\mathcal{B}}^{\text{euf-nma}} + Q_{\text{sig}} \cdot \text{AdvPRF}^{\text{H}_3} \\ &\quad + \text{AdvColl}^{\text{H}_1} [Q_1 + 2Q_{\text{sig}}] + \text{AdvColl}^{\text{H}_2^1} [Q_{2,1} + Q_{\text{sig}}] + \text{AdvColl}^{\text{H}_2^2} [Q_{2,2} + Q_{\text{sig}}] \\ &\quad + Q_{\text{sig}}(Q_{2,1} + Q_{\text{sig}})2^{-2\lambda} + Q_{\text{sig}}(Q_{2,2} + Q_{\text{sig}})2^{-(5\lambda+64)} + \tilde{Q}_{\text{sig}}(Q_{2,3} + \tilde{Q}_{\text{sig}})2^{-2\lambda} \\ &\quad + \hat{k} \cdot \text{AdvPRG}^{\text{PRG}_1} [Q_{\text{sig}}, \tau] + \text{AdvJPRF}^{\text{PRG}_2, \text{Com}_1} [Q_{\text{sig}}, \tau]. \end{aligned}$$

In what follows, we denote W_i the event that the adversary wins in G_i .

G_1 : The original SUF-CMA game. We have

$$\Pr[W_1] = \text{Adv}_{\mathcal{A}}^{\text{suf-cma}}.$$

G_2 : The signing oracle chooses $(\text{mseed}, \text{salt}) \leftarrow_{\$} \{0, 1\}^{\lambda+2\lambda}$ instead of $(\text{mseed}, \text{salt}) := H_3(\text{sk} || \tilde{\mu} || \text{rand})$.

Since mseed is kept secret in the whole procedure, we can treat mseed as a one-time secret key of PRF. Thus, the difference between G_1 and G_2 is

$$|\Pr[W_1] - \Pr[W_2]| \leq Q_{\text{sig}} \cdot \text{AdvPRF}^{\text{H}_3}.$$

Remark 7.2. If the signature scheme is deterministic, we use the argument for the ROM in [BPS16, Thm.4], which leads to the inequality

$$\Pr[W_1] \leq 2 \cdot \Pr[W_2].$$

G₃: Next, we introduce a collision check for PRG for the GGM tree, Com_1 , Com_2 , H_1 , H_2^1 , H_2^2 , and H_2^3 . If there is a collision among the queries to those oracles in the whole security game, then the challenger aborts the game. We easily have

$$\begin{aligned} |\Pr[W_2] - \Pr[W_3]| &\leq \text{AdvColl}^{\text{PRG}_1}[Q_{\text{prg},1} + (N-1)Q_{\text{sig}}] \\ &\quad + \text{AdvColl}^{\text{Com}_1}[Q_{\text{com}} + NQ_{\text{sig}}] + \text{AdvColl}^{\text{Com}_2}[Q_{\text{com},2} + Q_{\text{sig}}] \\ &\quad + \text{AdvColl}^{H_1}[Q_1 + 2Q_{\text{sig}}] + \text{AdvColl}^{H_2^1}[Q_{2,1} + Q_{\text{sig}}] \\ &\quad + \text{AdvColl}^{H_2^2}[Q_{2,2} + Q_{\text{sig}}] + \text{AdvColl}^{H_2^3}[Q_{2,3} + \tilde{Q}_{\text{sig}}]. \end{aligned}$$

Remark 7.3. If we only consider the EUF-CMA security, then we do not need collision checks for PRG_1 for the GGM tree, Com_1 , Com_2 , and H_2^3 .

G₄: The signing oracle programs H_2^1 by choosing $\text{ch}_1 \leftarrow_{\$} \{0, 1\}^{5\lambda+64}$. Since salt are chosen uniformly at random over $\{0, 1\}^\lambda$, a_1 's min-entropy is at least 2λ . The adaptive reprogramming technique shows that

$$|\Pr[W_3] - \Pr[W_4]| \leq Q_{\text{sig}}(Q_{2,1} + Q_{\text{sig}}) \cdot 2^{-2\lambda}.$$

NOTE: Since we use 2λ -bit salt to compute ch_1 , we do not need additional games that make h_{com} random.

G₅: The signing oracle programs H_2^2 by choosing $\text{ch}_2 \leftarrow_{\$} \{0, 1\}^{2\lambda}$.

We write a reduction as follows: The reduction algorithm first computes a_1 , chooses ch_1 , and a_2 . It then queries (ch_1, a_2) to its oracle and obtains ch_2 .

We have

$$|\Pr[W_4] - \Pr[W_5]| \leq Q_{\text{sig}}(Q_{2,2} + Q_{\text{sig}}) \cdot 2^{-(5\lambda+64)}.$$

NOTE: We do not need to program H_4 .

G₆: The signing oracle programs H_2^3 by choosing $\text{ch}_3 \leftarrow_{\$} \{0, 1\}^{\tau_0\kappa_0 + \tau_1\kappa_1 + w}$.

We write a reduction as follows: The reduction algorithm first computes a_1 , chooses ch_1 , computes a_2 , and chooses ch_2 , and computes a_3 . It then queries (ch_2, a_3) to its oracle and obtains ch_3 . Since ch_2 has a min-entropy at least 2λ , we have

$$|\Pr[W_5] - \Pr[W_6]| \leq \tilde{Q}_{\text{sig}}(Q_{2,3} + \tilde{Q}_{\text{sig}}) \cdot 2^{-2\lambda},$$

G₇: We then modify the order of sampling;

1. Sample ch_1 and ch_2 uniformly at random.
2. Sample $\text{ch}_{3,0}, \dots, \text{ch}_{3,B-1} \leftarrow \{0, 1\}^{\tau_0\kappa_0 + \tau_1\kappa_1 + w}$ until it is accepted; define $\text{ch}_3 := \text{ch}_{3,B-1}$.
3. Run the prover algorithm.
4. Reprogram $H_2^1(\mu || h_{\text{com}} || \mathbf{c}_1 || \dots || \mathbf{c}_{\tau-1} || \text{salt}) := \text{ch}_1$.
5. Reprogram $H_2^2(\text{ch}_1 || \tilde{\mathbf{u}} || h_V || \mathbf{t}) := \text{ch}_2$.
6. For $\text{ctr} \in [B]$, reprogram $H_2^3(\text{ch}_2 || [a] || \text{ctr}) := \text{ch}_{3,\text{ctr}}$.

This is just a conceptual change of the order, and we have

$$\Pr[W_6] = \Pr[W_7].$$

G_8 : Next, we make the signature uniformly at random as possible as FAEST's proof. Intuitively speaking, this corresponds to the replacement of the prover algorithms in the signing oracle with the simulator algorithms.

- $G_{8,1}$: On each signing query, we use $(h_{\text{com}}, \text{pdecom}, (\text{com}_{e,i[e]}), \mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}, \mathbf{u}) \leftarrow \text{SimVOLECommit}(i, \text{salt}, \hat{\ell})$ and adjust \mathbf{V} to be consistent with Δ and \mathbf{q} induced by ch_3 . We stress that SimVOLECommit samples \mathbf{u} and $\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}$ uniformly at random. Following FAEST's proof, this modification is justified by the multi-hiding property of VOLECommit . Concretely speaking, the upper bound is

$$|\Pr[W_7] - \Pr[W_{8,1}]| \leq \text{AdvHide}^{\text{VOLE}}[Q_{\text{sig}}].$$

Now, the distribution of (\mathbf{u}, \mathbf{V}) is independent of $\mathbf{c}_1, \dots, \mathbf{c}_{\tau-1}$.

- $G_{8,2}$: On each signing query, we sample $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{V}}$ at random instead of computing VOLEHash ; after that, we adjust the last $\ell_{\text{VOLEHashMask}} = \lambda + B$ rows of \mathbf{u} and \mathbf{V} to make them be consistent with ch_1 . Since VOLEHash is $\mathbb{F}_2^{\ell+\lambda}$ -hiding (Lemma 1 in FAEST's spec), the distributions are the same. We have

$$\Pr[W_{8,1}] = \Pr[W_{8,2}].$$

- $G_{8,3}$: Next, on each signing query, we choose $\llbracket a \rrbracket$ uniformly at random instead of computing $f_{\text{mask}}(X) + \sum_i \alpha_i f_i(X)$ and adjust the middle ℓ_{CZMask} rows of \mathbf{u} and \mathbf{V} . Correctly speaking, we compute $f_i(X)$, choose $\llbracket a \rrbracket$ uniformly at random while it is consistent with \mathbf{q} and Δ , then compute $f_{\text{mask}}(X) = \llbracket a \rrbracket - \sum_i \alpha_i f_i(X)$; we then adjust $\llbracket u_{i',k'} \rrbracket_{i' \in [d-1], k' \in [\rho]}$, which decide the coefficients of $(d-1)$ -degree polynomial $f_{\text{mask}}(X)$. Since $\llbracket u_{i',k'} \rrbracket_{i' \in [d-1], k' \in [\rho]}$ are hidden from the adversary, the modification in this game does not change anything. Thus, we have

$$\Pr[W_{8,2}] = \Pr[W_{8,3}].$$

- $G_{8,4}$: Finally, on each signing query, we replace $\mathbf{t} = \mathbf{w} + \mathbf{u}[\ell]$ with random one since $\llbracket a \rrbracket$ is independent of \mathbf{w} . This modification doesn't change the distribution because $\mathbf{u}[\ell]$ is chosen uniformly at random and is not used elsewhere. Thus, we have

$$\Pr[W_{8,3}] = \Pr[W_{8,4}].$$

G_9 : In this game, the adversary *loses* if there exists $(\text{msg}, \sigma = (\text{ch}_3, a_4)) \in \mathcal{Q}$ such that

- $(\bar{a}_1, \bar{\text{ch}}_1) = (\bar{a}_1^*, \bar{\text{ch}}_1^*)$ and $\text{ch}_3 \neq \text{ch}_3^*$; or
- $(\bar{a}_1, \bar{\text{ch}}_1, \dots, \bar{a}_3, \text{ch}_3) = (\bar{a}_1^*, \bar{\text{ch}}_1^*, \dots, \bar{a}_3^*, \text{ch}_3^*)$ and $a_4 \neq a_4^*$.

If there is a difference between G_8 and G_9 , then the adversary submits new $(\text{msg}^*, (\text{ch}_3^*, a_4^*))$ such that there exists $(\text{msg}, (\text{ch}_3, a_4)) \in \mathcal{Q}$ satisfying either one of the conditions. Let \bar{h}_{com} and \bar{h}_{com}^* be the results of VC.Reconstruct from ch_3 and ch_3^* respectively.

Case 1: Suppose that $(\bar{a}_1, \bar{ch}_1) = (\bar{a}_1^*, \bar{ch}_1^*)$ and $ch_3 \neq ch_3^*$. Since both challenges are accepted, we have $\mathbf{i} \neq \mathbf{i}^*$, where $\mathbf{i} := \text{ChallDec}(ch_3[0 : \tau_0\kappa_0 + \tau_1\kappa_1 - 1])$ and $\mathbf{i}^* := \text{ChallDec}(ch_3^*[0 : \tau_0\kappa_0 + \tau_1\kappa_1 - 1])$. On the other hand, we have $\bar{h}_{\text{com}} = \bar{h}_{\text{com}}^*$ since $\bar{a}_1 = \bar{a}_1^*$. Since we checked the collision for Com_2 , this implies $\{\bar{\text{com}}_{e,i}\} = \{\bar{\text{com}}_{e,i}^*\}$. This breaks the multi-target non-invertibility of Com_1 since the adversary outputs the preimage of $\text{com}_{e,i[e]}$ for some e such that $\mathbf{i}[e] \neq \mathbf{i}^*[e]$, where $\text{com}_{e,i[e]}$ is chosen uniformly at random in SimVOLECommit . Thus, the difference is at most $\text{AdvNI}^{\text{Com}_1}[\tau Q_{\text{sig}}]$, in which the adversary receives random τQ_{sig} strings and outputs one of preimages of the strings.

Case 2: Next, suppose that $(\bar{a}_1, \bar{ch}_1, \dots, \bar{a}_3, ch_3) = (\bar{a}_1^*, \bar{ch}_1^*, \dots, \bar{a}_3^*, ch_3^*)$ and $a_4 \neq a_4^*$. Note that the condition $ch_3 = ch_3^*$ leads to $\mathbf{i} = \mathbf{i}^*$. Since $(\bar{a}_1, \bar{a}_2, \bar{a}_3) = (\bar{a}_1^*, \bar{a}_2^*, \bar{a}_3^*)$, we have $(\text{pdecom}, \{\text{com}_{e,i[e]}\}_{e \in [\tau]}, \text{ctr}, \text{salt}) \neq (\text{pdecom}^*, \{\text{com}_{e,i[e]}^*\}_{e \in [\tau]}, \text{ctr}^*, \text{salt}^*)$.

We have four sub-cases:

- Suppose that $\text{pdecom} \neq \text{pdecom}^*$. In the computation of VC.Reconstruct , we have $\text{nodes}[i^+] \neq \text{nodes}^*[i^+]$ for some $i^+ \in \text{revealed}$ calculated from $\mathbf{i} = \mathbf{i}^*$. We then consider the sub-tree of the GGM tree whose root is i^+ . Let I_{all} be the indices of the nodes of the sub-tree and I_{leaves} be the indices of the leaf nodes of the sub-tree.

We have two cases:

- If $\{\text{nodes}[i]\}_{i \in I_{\text{leaves}}} = \{\text{nodes}^*[i]\}_{i \in I_{\text{leaves}}}$, we then have a collision in the sub-tree. That is, there is an index $j \in I_{\text{all}}$ such that $\text{nodes}[j] \neq \text{nodes}^*[j]$ but $\text{PRG}_1(\text{nodes}[j], \text{salt} :: 2\lambda) = \text{PRG}_1(\text{nodes}^*[j], \text{salt}^* :: 2\lambda)$. However, this case is already excluded by the collision check for PRG_1 introduced in G_3 .
- If $\{\text{nodes}[i]\}_{i \in I_{\text{leaves}}} \neq \{\text{nodes}^*[i]\}_{i \in I_{\text{leaves}}}$, we have at least one index $j^+ \in I_{\text{leaves}}$ satisfying $\text{nodes}[j^+] \neq \text{nodes}^*[j^+]$. Let (e, i) be a pair induced by j^+ (that is, $j^+ = N - 1 + \psi(e, i)$). If $\text{com}_{e,i} = \text{com}_{e,i}^*$, then we find a collision for Com_1 . However, this case is already excluded by the collision check for Com_1 introduced in G_3 . Otherwise, if $\text{com}_{e,i} \neq \text{com}_{e,i}^*$, then we find a collision for Com_2 since $\bar{a}_1 = \bar{a}_1^*$ implies $\bar{h}_{\text{com}} = \bar{h}_{\text{com}}^*$. However, this cannot happen since we already exclude this event by the collision check for Com_2 introduced in G_3 .
- If $\{\text{com}_{e,i[e]}\}_e \neq \{\text{com}_{e,i^*[e]}^*\}$, then we have $\{\bar{\text{com}}_{e,i}\}_{e,i} \neq \{\bar{\text{com}}_{e,i}^*\}_{e,i}$. However, $\bar{a}_1 = \bar{a}_1^*$ implies $\bar{h}_{\text{com}} = \bar{h}_{\text{com}}^*$ and we have a collision for Com_2 . However, this case is already excluded by the collision check for Com_2 introduced in G_3 .
- If $\text{ctr} \neq \text{ctr}^*$, then we find a collision for H_2^3 since $(\bar{ch}_2 || \bar{a}_3 || \text{ctr}) \neq (\bar{ch}_2^* || \bar{a}_3^* || \text{ctr}^*)$ but $ch_3 = H_2^3(\bar{ch}_2 || \bar{a}_3 || \text{ctr}) = H_2^3(\bar{ch}_2^* || \bar{a}_3^* || \text{ctr}^*) = ch_3^*$. However, this case is already excluded by the collision check for H_2^3 introduced in G_3 .
- If $\text{salt} \neq \text{salt}^*$, then we find a collision for H_2^1 since $\bar{ch}_1 = \bar{ch}_1^*$ but its corresponding inputs differs. However, this case is already excluded by the collision check for H_2^1 introduced in G_3 .

Summing up, the adversary cannot output a valid forgery satisfying $(\bar{a}_1, \bar{ch}_1, \dots, \bar{a}_3, ch_3) = (\bar{a}_1^*, \bar{ch}_1^*, \dots, \bar{a}_3^*, ch_3^*)$ and $a_4 \neq a_4^*$ with $(\text{msg}, (ch_3, a_4)) \in \mathcal{Q}$.

The bound for both cases, we obtain that

$$|\Pr[W_{8,4}] - \Pr[W_9]| \leq \text{AdvNI}^{\text{Com}_1}[\tau Q_{\text{sig}}].$$

Reduction to EUF-NMA. We prove that the adversary's forgery in G_9 never involves the reprogrammed points by contradiction. Suppose that the adversary's forgery involves some of the reprogrammed points when the signing oracle computes $\sigma = (\text{ch}_3, a_4)$ on a query msg . Those points are related to $\text{ch}_1 = H_2^1(\tilde{\mu}||a_1||\text{salt})$, $\text{ch}_2 = H_2^2(\text{ch}_1||a_2)$, and $\text{ch}_{3,i} = H_2^3(\text{ch}_2||a_3||i)$, where $i \in [B]$. In any case, because of the collision checks, the first point $\text{ch}_1 = H_2^1(\tilde{\mu}||a_1||\text{salt})$ should be reprogrammed.⁹ Hence, the adversary's forgery satisfies $\text{ch}_1 = \overline{\text{ch}_1}^*$, $\tilde{\mu} = \tilde{\mu}^*$, $\bar{a}_1 = \bar{a}_1^*$, and $\text{salt} = \text{salt}^*$. Especially, $\tilde{\mu} = \tilde{\mu}^*$ implies $\text{msg} = \text{msg}^*$ because of the collision check of H_1 introduced in G_3 . Furthermore, $\text{msg} = \text{msg}^*$ implies $(\text{ch}_3, a_4) \neq (\text{ch}_3^*, a_4^*)$. Due to the first check in G_9 , if $(\bar{a}_1, \overline{\text{ch}_1}) = (\bar{a}_1^*, \overline{\text{ch}_1}^*)$ then ch_3 should be equal to ch_3^* . Thus, the condition is boiled down to $(\text{msg}, \bar{a}_1, \overline{\text{ch}_1}, \text{ch}_3) = (\text{msg}^*, \bar{a}_1^*, \overline{\text{ch}_1}^*, \text{ch}_3^*)$ and $a_4 \neq a_4^*$. By the way, since $\text{ch}_3 = \text{ch}_3^*$ holds and there must not be collisions for H_2^3 , we have $(\overline{\text{ch}_2}, \bar{a}_3) = (\overline{\text{ch}_2}^*, \bar{a}_3^*)$. Furthermore, $\overline{\text{ch}_2} = \overline{\text{ch}_2}^*$ and the collision check for H_2^2 implies $(\overline{\text{ch}_1}, \bar{a}_2) = (\overline{\text{ch}_1}^*, \bar{a}_2^*)$. Thus, we have $(\bar{a}_1, \overline{\text{ch}_1}, \dots, \bar{a}_3, \text{ch}_3) = (\bar{a}_1^*, \overline{\text{ch}_1}^*, \dots, \bar{a}_3^*, \text{ch}_3^*)$ and $a_4 \neq a_4^*$, but if this holds the adversary loses in G_9 . This is a contradiction, and now, we can conclude that the adversary's forgery does not involve the points reprogrammed by the signing oracle in G_9 .

Thus, we can easily construct an adversary \mathcal{B} against the EUF-NMA security of the signature scheme satisfying

$$\Pr[W_9] \leq \text{Adv}_{\mathcal{B}}^{\text{euf-nma}}.$$

Remark 7.4. : If we only consider the EUF-CMA security, we can skip G_9 and we have $\Pr[W_{8,4}] \leq \text{Adv}_{\mathcal{B}}^{\text{euf-nma}}$. The argument follows:

Since we consider the EUF-CMA security, the adversary's output should msg^* such that $(\text{msg}^*, \sigma) \notin Q$ for any σ . Suppose that the adversary's forgery involves some of the reprogrammed points when the signing oracle computes $\sigma = (\text{ch}_3, a_4)$ on a query $\text{msg} \neq \text{msg}^*$. We note that, due to the collision check for H_1 introduced in G_3 , we have $\mu = H_1(\text{pk}||\text{msg}) \neq \mu^* = H_1(\text{pk}||\text{msg}^*)$.

1. If the first reprogrammed point $\text{ch}_1 = H_2^1(\mu||a_1||\text{salt})$ is involved, then we have $(\mu, a_1, \text{salt}) = (\mu^*, \bar{a}_1^*, \text{salt}^*)$ and $\mu = H_1(\text{pk}||\text{msg}) = H_1(\text{pk}||\text{msg}^*) = \mu^*$. But, this contradicts with $\mu \neq \mu^*$. Thus, the first reprogrammed point should not be involved in the forgery.
2. If the second reprogrammed point $\text{ch}_2 = H_2^2(\text{ch}_1||a_2)$ is involved, then we have $(\text{ch}_1, a_2) = (\overline{\text{ch}_1}^*, \bar{a}_2)$. But, $\mu \neq \mu^*$ and the collision check for H_2^1 implies that this cannot happen. Thus, the first and second reprogrammed points should not be involved in the forgery. In addition, we have $\text{ch}_1 \neq \overline{\text{ch}_1}^*$.

⁹ Otherwise, we can find the collision for H_2^1 or H_2^2 introduced in G_3 , but such event is eliminated by the collision check.

3. If the third reprogrammed point $\text{ch}_3 = \text{H}_2^3(\text{ch}_2 \| a_3 \| \text{ctr})$ for some $\text{ctr} \in [B]$ is involved, then we have $(\text{ch}_2, a_3) = (\overline{\text{ch}_2}, \overline{a_3})$. But, $\text{ch}_1 \neq \overline{\text{ch}_1}$ and the collision check for H_2^1 implies that this cannot happen. Thus, all three reprogrammed points should not be involved in the forgery.

Thus, the forgery does not involve the reprogrammed point, and the reduction is easily obtained. ¹⁰

Security proofs for the BUFF securities. Section 4.B.4 of the call for proposal lists additional desirable security properties beyond standard unforgeability. In this section, we evaluate the so-called BUFF securities (message-bound signatures, exclusive ownership, and non re-signability) of our proposal. For the definitions of BUFF securities, see [CDF⁺21]. The proofs in [KX24b] showed that several MPCitH signatures achieve some BUFF securities. We adopt their proofs in the context of the VOLEitH signature and contain the concrete proofs below for completeness.

Message-bounding signatures (MBS). The MBS security shows that any efficient adversary cannot output pk and σ with two different messages msg and msg' such that $(\text{pk}, \text{msg}, \sigma)$ and $(\text{pk}, \text{msg}', \sigma)$ are both valid. Let \mathcal{A} be an adversary against the MBS security of PERK: \mathcal{A} takes 1^λ as input and outputs pk , msg , msg' , and σ with $\text{msg} \neq \text{msg}'$ satisfying $\text{PERK.Verify}(\text{pk}, \text{msg}, \sigma) = \text{PERK.Verify}(\text{pk}, \text{msg}', \sigma) = 1$. Let us denote the internally reproduced values in the verifications $\text{PERK.Verify}(\text{pk}, \text{msg}, \sigma)$ and $\text{PERK.Verify}(\text{pk}, \text{msg}', \sigma)$ by $\bar{\cdot}$ and $\bar{\cdot}'$, respectively.

- Due to the definition of PERK.Verify , we have $\text{ch}_3 = \bar{\text{ch}}_3 = \bar{\text{ch}}_3'$, where $\bar{\text{ch}}_3 = \text{H}_2^3(\bar{\text{ch}}_2 \| \bar{a}_3 \| \text{ctr})$ and $\bar{\text{ch}}_3' = \text{H}_2^3(\bar{\text{ch}}_2' \| \bar{a}_3' \| \text{ctr})$. If $(\bar{\text{ch}}_2, \bar{a}_3) \neq (\bar{\text{ch}}_2', \bar{a}_3')$, then we find the collision for H_2^3 .
- Otherwise, we have $\bar{\text{ch}}_2 = \bar{\text{ch}}_2'$, where $\bar{\text{ch}}_2 = \text{H}_2^2(\bar{\text{ch}}_1 \| \bar{a}_2)$ and $\bar{\text{ch}}_2' = \text{H}_2^2(\bar{\text{ch}}_1' \| \bar{a}_2')$. If $(\bar{\text{ch}}_1, \bar{a}_2) \neq (\bar{\text{ch}}_1', \bar{a}_2')$, then we find the collision for H_2^2 .
- Otherwise, we have $\bar{\text{ch}}_1 = \bar{\text{ch}}_1'$, where $\bar{\text{ch}}_1 = \text{H}_2^1(\bar{\mu} \| \bar{a}_1 \| \text{salt})$ and $\bar{\text{ch}}_1' = \text{H}_2^1(\bar{\mu}' \| \bar{a}_1' \| \text{salt})$. If $(\bar{\mu}, \bar{a}_1, \text{salt}) \neq (\bar{\mu}', \bar{a}_1', \text{salt})$, then we find the collision for H_2^1 .
- Otherwise, we have $\bar{\mu} = \bar{\mu}'$, where $\bar{\mu} = \text{H}_1(\text{pk} \| \text{msg})$ and $\bar{\mu}' = \text{H}_1(\text{pk} \| \text{msg}')$. Since $\text{msg} \neq \text{msg}'$, thus, we find the collision for H_1 .

In any case, we can find a collision for either H_1 , H_2^1 , H_2^2 , or H_2^3 . Thus, if they are collision-resistant, then PERK is MBS-secure.

Malicious strong universal exclusive ownership (M-S-UEO). We consider M-S-UEO, which is the strongest form of exclusive ownership.¹¹ The M-S-UEO security shows that any efficient adversary cannot output two different public key

¹⁰ We notice that FAEST's proof did not consider the collision-resistance property of H_1 , H_2^1 , and H_2^2 .

¹¹ If the scheme is M-S-UEO-secure, then it also Strong destructive exclusive ownership (S-DEO-secure) and Strong conservative exclusive ownership (S-CEO-secure).

pk and pk' , two (possibly different) messages msg and msg' , and a signature σ such that $(\text{pk}, \text{msg}, \sigma)$ and $(\text{pk}', \text{msg}', \sigma)$ are both valid. Let \mathcal{A} be an adversary against the M-S-UEO security of PERK: \mathcal{A} takes 1^λ as input and outputs pk , pk' , msg , msg' , and σ with $\text{pk} \neq \text{pk}'$ such that $\text{PERK.Verify}(\text{pk}, \text{msg}, \sigma) = \text{PERK.Verify}(\text{pk}', \text{msg}', \sigma) = 1$. By a similar argument to the MBS security above, we have a collision for H_3^2 , H_2^2 , or H_2^1 by ch_3 , ch_2 , or ch_1 , respectively. If $\bar{\mu} = \bar{\mu}'$, then we have $H_1(\text{pk}||\text{msg}) = H_1(\text{pk}'||\text{msg}')$ and obtain a collision $(\text{pk}, \text{msg}) \neq (\text{pk}', \text{msg}')$ for H_1 since $\text{pk} \neq \text{pk}'$. Thus, if the hash functions H_1 , H_2^1 , H_2^2 , and H_3^2 are collision resistant, then PERK is M-S-UEO-secure.

Weak non-resignability (WNR). Roughly speaking, (weak) non-resignability shows that, given pk and σ on a hidden message msg (with some leakage)¹², any efficient adversary cannot output pk' and σ' such that $(\text{pk}', \text{msg}, \sigma')$ is valid. There are some generic conversion for EUF-CMA-secure signature scheme in the (Q)ROM.

Since the signature is produced on $\bar{\mu} := H_1(\text{pk}||\text{msg})$ instead of msg itself, our signature scheme inherently implements the BUFF transform with a random oracle H_1 . Don, Fehr, Huang, Liao, and Struck [DFH⁺24] showed the BUFF transform allows us to achieve $\text{sNR}^{H_1, \perp}$. Thus, PERK satisfies $\text{sNR}^{H_1, \perp}$.

There is another weakened NR, $\text{NR}^{H_1, \perp}$, in [DFHS24], but this requires the \mathcal{S} -BUFF transform that computes $y = H_1(\text{pk}||\text{msg}||s)$ with salt s .

7.2 Known attacks against PKP

7.2.1 Overview of known attacks. The Permuted Kernel Problem (PKP) problem was introduced by Shamir in 1990 [Sha90]. Despite its long standing history in cryptographic applications [Sha90, BFK⁺19, Beu20, BG23] and consequently many cryptanalytic efforts [Geo92, BCCG93, PC94, JJ01, LP11, KMP19, SBC23], algorithms to solve the PKP are still rather simple adaptations of combinatorial enumeration and meet-in-the-middle techniques. Indeed, the best attack on standard PKP is a meet-in-the-middle adaptation known as the KMP algorithm by Koussa, Macario-Rat and Patarin [KMP19]. Even though there has been some recent progress on attacks [SBC23], those do not improve over the KMP algorithm in the case of standard PKP on which PERK is based.

7.2.2 KMP algorithm on PKP. In this section we briefly sketch the KMP algorithm to solve the PKP. Fully fledged descriptions, analysis and estimation scripts are given for example in [KMP19, SBC23, EVZB24]. The algorithm by Koussa, Macario-Rat and Patarin [KMP19] is a slight variant of previously known combinatorial techniques [Geo92, BCCG93, PC94, JJ01]. The algorithm was first proposed for the inhomogeneous version of PKP, where $\mathbf{H}\pi(\mathbf{x}) = \mathbf{y}$ for a given vector $\mathbf{y} \in \mathbb{F}_q^m$ [KMP19]. The algorithm was then recently extended to the multi-dimensional case [SBC23], i.e. the case where multiple \mathbf{x}_i and \mathbf{y}_i are provided and the the solution is a permutation π , with $\mathbf{H}\pi(\mathbf{x}_i) = \mathbf{y}_i$ for all i .

¹² The definitions vary depending on how the information of msg is leaked to the adversary.

Santini, Baldi and Chiaraluce also introduced further improvements to this generalized KMP algorithm. However, as it only improves for $i > 1$ pairs $(\mathbf{x}_i, \mathbf{y}_i)$, we do not consider it for the security analysis of PERK, which uses $i = 1$.

Initially, the matrix \mathbf{H} is transformed into semi-systematic form by applying a change of basis (modelled by the invertible matrix \mathbf{Q})

$$\mathbf{QH} = \begin{pmatrix} \mathbf{I}_{m-u} & \mathbf{H}_1 \\ \mathbf{0} & \mathbf{H}_2 \end{pmatrix},$$

where $\mathbf{H}_1 \in \mathbb{F}_q^{(m-u) \times (n-m+u)}$, $\mathbf{H}_2 \in \mathbb{F}_q^{u \times (n-m+u)}$ and u is an optimization parameter of the algorithm. For the inhomogeneous variant, where $\mathbf{y} \neq \mathbf{0}$, one maintains the validity of the PKP identity by multiplying the syndrome \mathbf{y} by the same matrix \mathbf{Q}

$$\begin{aligned} \mathbf{QH}\pi(\mathbf{x}) &= \begin{pmatrix} \mathbf{I}_{m-u} & \mathbf{H}_1 \\ \mathbf{0} & \mathbf{H}_2 \end{pmatrix} \pi(\mathbf{x}) = \begin{pmatrix} \mathbf{I}_{m-u} & \mathbf{H}_1 \\ \mathbf{0} & \mathbf{H}_2 \end{pmatrix} \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} = (\mathbf{x}_1 + \mathbf{H}_1\mathbf{x}_2, \mathbf{H}_2\mathbf{x}_2)^\top \\ &= (\mathbf{y}_1, \mathbf{y}_2)^\top = \mathbf{Q}\mathbf{y}, \end{aligned}$$

where $\mathbf{Q}\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2) \in \mathbb{F}_q^{m-u} \times \mathbb{F}_q^u$ and $\pi(\mathbf{x}) = (\mathbf{x}_1, \mathbf{x}_2) \in \mathbb{F}_q^{m-u} \times \mathbb{F}_q^u$. The algorithm now focuses on solving the identity $\mathbf{H}_2\mathbf{x}_2 = \mathbf{y}_2$. For any found \mathbf{x}_2 satisfying the identity it is then checked if $\mathbf{x}_1 = \mathbf{y}_1 - \mathbf{H}_1\mathbf{x}_2$ and \mathbf{x}_2 together form a permutation of \mathbf{x} .

Candidates for \mathbf{x}_2 are obtained by a meet-in-the-middle enumeration strategy. Therefore \mathbf{x}_2 is further split as $\mathbf{x}_2 = (\mathbf{x}_{21}, \mathbf{x}_{22})$, with $\mathbf{x}_{21}, \mathbf{x}_{22} \in \mathbb{F}_q^{u \times ((n-m+u)/2)}$ to obtain the meet-in-the-middle identity

$$\mathbf{H}_2(\mathbf{x}_{21}, \mathbf{0}) = \mathbf{y}_2 - (\mathbf{0}, \mathbf{x}_{22}). \quad (2)$$

Then the algorithm enumerates all candidates for \mathbf{x}_{21} and \mathbf{x}_{22} , that is all permutations of any selection of $(n-m+u)/2$ entries of \mathbf{x} . For each such vector the left (resp. right) side of Equation (2) is stored in a list L_1 (resp. L_2). In a final step the algorithm searches for matches between the lists L_1 and L_2 yielding the candidates for \mathbf{x}_2 . From there \mathbf{x}_1 can be computed as $\mathbf{x}_1 = \mathbf{y}_1 - \mathbf{H}_1\mathbf{x}_2$. If $(\mathbf{x}_1, \mathbf{x}_2)$ forms a permutation of \mathbf{x} this yields the solution π .

The complexity of the algorithm is (up to polynomial factors) linear in the sizes of the lists L_1 , L_2 and L , where L is the list of matches. The expected sizes are

$$|L_1| = |L_2| = \binom{n}{(n-m+u)/2} ((n-m+u)/2)! \quad \text{and} \quad |L| = \frac{|L_1| \times |L_2|}{q^u}$$

7.2.3 Relation between PKP and the Code Equivalence Problem. In their recent work Santini et al. [SBC23] formalized the equivalence between PKP and the subcode equivalence problem. Namely, PKP asks to find a permutation that sends the one dimensional code \mathbf{x} into the code with parity check matrix

H, which defines the problem. For variants of PKP using higher dimensional codes this can have some implications regarding security, depending on the concrete choice of code dimension. However, again, for standard PKP using a 1-dimensional code \mathbf{x} this has no effect on security.

7.2.4 PKP over extension fields. In the context of PERK, we consider PKP over \mathbb{F}_q with $q = p^r$ being a prime power. To the best of our knowledge, no algorithms are known that exploit the structure of the extension field \mathbb{F}_q . Especially, there are no known adaptations or enhancements of the KMP algorithm that leverage this characteristic.

Note that, this observation aligns with the evidence from the closely related syndrome decoding problem over \mathbb{F}_q . The most effective algorithm for this problem, when $q > 2$, relies on an enumeration routine that is conceptually similar to the KMP algorithm. Yet, even in the case of syndrome decoding over extension fields, no algorithmic improvements have been identified that exploit the extension field's structure. This was recently confirmed in [EW25].

7.2.5 Parameter selection. For parameter selection we fix $q = 2048$. We then, for any choice of n rely on a standard choice of m . That is, for any choice of n we choose m minimal such that the expected amount of solutions to a random instance of the PKP(q, m, n) is smaller than one. Subsequently we use the *CryptographicEstimators* library¹³ [EVZB24], for the concrete complexity estimation of the KMP algorithm for any set of parameters (q, n, m) . Eventually, for any security level we choose n such that the complexity estimation yields a comfortable margin to the NIST specified security levels.

Note that in addition to this margin, the estimation of the KMP algorithm via the *CryptographicEstimators* is already a lower bound on the algorithm's complexity by neglecting some factors. Furthermore the KMP algorithm suffers from a memory complexity that is equal to its time complexity. Therefore, realistic attacks have to resort to time-memory trade-offs further adding to this margin. The proposed parameters can therefore be seen as conservative choices with respect to security.

Overall the detailed procedure leads to the choices of parameters given in Section 5 whose estimated bit complexity is given in Table 12.

¹³ https://github.com/Crypto-TII/cryptographic_estimators

Instance	q	n	m	Bit Security
PERK-1	2048	64	27	150
PERK-3	2048	92	43	220
PERK-5	2048	118	59	286

Table 12: Bit security estimates of PERK parameters

8 Advantages and Limitations

We now discuss some advantages and limitations of PERK.

8.1 Advantages

Some advantages of our design are:

- + PERK features very small public key and secret key sizes along with moderate signature sizes. Therefore, on the combined metric of `pk + signature` size, PERK produces sizes of approximately 3.5 kB for NIST security level 1 which compares well with other signature schemes.
- + Contrarily to many post-quantum schemes, the security of PERK is not based on a problem relying on cyclic structure or ring structure.
- + Resilience against PKP attacks: A large part of the signature size scales with the security parameter λ (due to the seed trees and commitments) and not directly with the PKP parameters. As a consequence, increasing the PKP parameters has a limited impact on the total size of the signature.
- + PERK performances are constrained by numerous calls to symmetric cryptographic primitives. Any speedup to the implementation of these primitives directly benefit PERK. In particular, hardware acceleration support for such primitives improves the performance of the scheme.

8.2 Limitations

In the following, we point out the limitations of PERK.

- While PKP was initially defined over prime fields, PERK relies on PKP defined over \mathbb{F}_q where \mathbb{F}_q is an extension field of \mathbb{F}_2 . While, no algorithms exploiting the structure of the extension field are known, this variant has been less studied than the original long time standing PKP problem.
- While PERK’s performance profile is comparable to other MPCitH-based constructions, those can not compete with the fastest post-quantum secure schemes, usually based on structured lattices.

References

- BBD⁺23a. Manuel Barbosa, Gilles Barthe, Christian Doczkal, Jelle Don, Serge Fehr, Benjamin Grégoire, Yu-Hsuan Huang, Andreas Hülsing, Yi Lee, and Xiaodi Wu. Fixing and mechanizing the security proof of Fiat-Shamir with aborts and Dilithium. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 358–389. Springer, Cham, August 2023.
- BBD⁺23b. Carsten Baum, Lennart Braun, Cyprien Delpéch de Saint Guilhem, Michael Kloof, Christian Majenz, Shibam Mukherjee, Emmanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl. FAEST. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- BBD⁺23c. Carsten Baum, Lennart Braun, Cyprien Delpéch de Saint Guilhem, Michael Kloof, Emmanuela Orsini, Lawrence Roy, and Peter Scholl. Publicly verifiable zero-knowledge and post-quantum signatures from VOLE-in-the-head. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 581–615. Springer, Cham, August 2023.
- BBGK24. Slim Bettaiieb, Loïc Bidoux, Philippe Gaborit, and Mukul Kulkarni. Modelings for generic PoK and applications: Shorter SD and PKP based signatures. Cryptology ePrint Archive, Report 2024/1668, 2024.
- BBM⁺25. Carsten Baum, Ward Beullens, Shibam Mukherjee, Emmanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl. One tree to rule them all: Optimizing ggm trees and owfs for post-quantum signatures. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 463–493. Springer, 2025.
- BCCG93. Thierry Baritaud, Mireille Campana, Pascal Chauvaud, and Henri Gilbert. On the security of the permuted kernel identification scheme. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 305–311. Springer, Berlin, Heidelberg, August 1993.
- Ber19. Daniel J. Bernstein. [djbsort](https://sorting.cr.yp.to/). <https://sorting.cr.yp.to/>, 2019. [Online; accessed 20-June-2023].
- Beu20. Ward Beullens. Sigma protocols for MQ, PKP and SIS, and Fishy signature schemes. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 183–211. Springer, Cham, May 2020.
- BFK⁺19. Ward Beullens, Jean-Charles Faugère, Eliane Koussa, Gilles Macario-Rat, Jacques Patarin, and Ludovic Perret. PKP-based signature scheme. In Feng Hao, Sushmita Ruj, and Sourav Sen Gupta, editors, *INDOCRYPT 2019*, volume 11898 of *LNCS*, pages 3–22. Springer, Cham, December 2019.
- BG23. Loïc Bidoux and Philippe Gaborit. Compact post-quantum signatures from proofs of knowledge leveraging structure for the PKP, SD and RSD problems. In *Codes, Cryptology and Information Security (C2SI)*, pages 10–42. Springer, 2023.
- BJKS94. Jürgen Bierbrauer, Thomas Johansson, Gregory Kabatianskii, and Ben Smeets. On families of hash functions via geometric codes and concatenation. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 331–342. Springer, Berlin, Heidelberg, August 1994.

- BPS16. Mihir Bellare, Bertram Poettering, and Douglas Stebila. From identification to signatures, tightly: A framework and generic transforms. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 435–464. Springer, Berlin, Heidelberg, December 2016.
- CDF⁺21. Cas Cremers, Samed Düzlülü, Rune Fiedler, Marc Fischlin, and Christian Janson. BUFFing signature schemes beyond unforgeability and the case of post-quantum signatures. In *2021 IEEE Symposium on Security and Privacy*, pages 1696–1714. IEEE Computer Society Press, May 2021.
- CW79. J. Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979.
- DFH⁺24. Jelle Don, Serge Fehr, Yu-Hsuan Huang, Jyun-Jie Liao, and Patrick Struck. Hide-and-seek and the non-resignability of the BUFF transform. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part III*, volume 15366 of *LNCS*, pages 347–370. Springer, Cham, December 2024.
- DFHS24. Jelle Don, Serge Fehr, Yu-Hsuan Huang, and Patrick Struck. On the (in)security of the BUFF transform. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 246–275. Springer, Cham, August 2024.
- DFPS23. Julien Devevey, Pouria Fallahpour, Alain Passelègue, and Damien Stehlé. A detailed analysis of Fiat-Shamir with aborts. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 327–357. Springer, Cham, August 2023.
- EVZB24. Andre Esser, Javier A. Verbel, Floyd Zweydinger, and Emanuele Bellini. SoK: CryptographicEstimators - a software library for cryptographic hardness estimation. In Jianying Zhou, Tony Q. S. Quek, Debin Gao, and Alvaro A. Cárdenas, editors, *ASIACCS 24*. ACM Press, July 2024.
- EW25. Freja Elbro and Violetta Weger. Can we speed up information set decoding by using extension field structure? *Cryptology ePrint Archive*, Paper 2025/1402, 2025.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Berlin, Heidelberg, August 1987.
- Geo92. Jean Geogiades. Some remarks on the security of the identification scheme based on permuted kernels. *Journal of Cryptology*, 5(2):133–137, January 1992.
- GGM84. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the cryptographic applications of random functions. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 276–288. Springer, Berlin, Heidelberg, August 1984.
- IKOS07. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 21–30. ACM Press, June 2007.
- JJ01. Éliane Jaulmes and Antoine Joux. Cryptanalysis of PKP: A new approach. In Kwangjo Kim, editor, *PKC 2001*, volume 1992 of *LNCS*, pages 165–172. Springer, Berlin, Heidelberg, February 2001.
- KKW18. Jonathan Katz, Vladimir Kolesnikov, and Xiao Wang. Improved non-interactive zero knowledge with applications to post-quantum signatures. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 525–537. ACM Press, October 2018.

- KMP19. Eliane Koussa, Gilles Macario-Rat, and Jacques Patarin. On the complexity of the permuted kernel problem. *Cryptology ePrint Archive*, Report 2019/412, 2019.
- KX24a. Haruhisa Kosuge and Keita Xagawa. Probabilistic hash-and-sign with retry in the quantum random oracle model. In Qiang Tang and Vanessa Teague, editors, *PKC 2024, Part I*, volume 14601 of *LNCS*, pages 259–288. Springer, Cham, April 2024.
- KX24b. Mukul Kulkarni and Keita Xagawa. Strong existential unforgeability and more of MPC-in-the-head signatures. *Cryptology ePrint Archive*, Report 2024/1069, 2024.
- LP11. Rodolphe Lampe and Jacques Patarin. Analysis of some natural variants of the pkp algorithm. *Cryptology ePrint Archive*, 2011.
- PC94. Jacques Patarin and Pascal Chauvaud. Improved algorithms for the permuted kernel problem. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 391–402. Springer, Berlin, Heidelberg, August 1994.
- Roy22. Lawrence Roy. SoftSpokenOT: Quieter OT extension from small-field silent VOLE in the minicrypt model. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 657–687. Springer, Cham, August 2022.
- SBC23. Paolo Santini, Marco Baldi, and Franco Chiaraluce. Computational hardness of the permuted kernel and subcode equivalence problems. *IEEE Transactions on Information Theory*, 2023.
- Sha90. Adi Shamir. An efficient identification scheme based on permuted kernels (extended abstract) (rump session). In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 606–609. Springer, New York, August 1990.
- Sti92. Douglas R. Stinson. Universal hashing and authentication codes. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 74–85. Springer, Berlin, Heidelberg, August 1992.