Elementary-Vector Modeling for Shorter VOLE-Based Signatures from SD and PKP

Boyuan Gao[1], Zongyang Zhang[1], Weihan Li[2], Jianwei Liu[1], and Jianting Ning[3,4]

[1] School of Cyber Science and Technology, Beihang University
{boyuangao,zongyangzhang,liujianwei}@buaa.edu.cn
[2] Ant Group
liweihan.lwh@antgroup.com
[3] Zhejiang Key Laboratory of Digital Fashion and Data Governance,
College of Information Science and Engineering, Zhejiang Sci-Tech University
[4] Faculty of Data Science, City University of Macau
jtning88@gmail.com

Abstract. Zero-knowledge proofs based on VOLE have recently become a promising approach for designing post-quantum signature schemes. Such schemes are constructed by having a signer prove knowledge of a secret input for a prescribed one-way function. In this work, we optimize VOLE-based signature schemes from syndrome decoding (SD) and the permuted kernel problem (PKP). Despite relying on different one-way functions, these schemes share a common modeling method: elementary vectors are compressed into shorter witnesses and then reconstructed through low-degree polynomials. We propose a unified framework for elementary-vector encoding and reconstruction. By choosing different parameters, this framework captures previous modeling approaches for SD and PKP as special cases. Based on this framework, we propose an optimized encoding and reconstruction approach that reduces the extended witness size and hence the final signature size. At NIST Level-I security, our SD- and PKP-based schemes achieve "signature + public key" sizes of 3 496 and 3 341 bytes, respectively, improving over the corresponding 3 775 bytes of SDitH v2.0 and 3 577 bytes of PERK v2.2. Furthermore, we construct a 3 673-byte post-quantum ring signature for rings of 2^{20} members, which is 27% and 70% smaller than those of Chiang *et al.* (CCS '25) and BLOOM (ASIACRYPT '22), respectively.

Keywords: post-quantum signature · VOLE-in-the-Head · syndrome decoding · permuted kernel problem · ring signature

1 Introduction

Zero-knowledge arguments of knowledge (ZKAoKs) [GMR85, BCC88] enable a prover to convince a verifier of the knowledge of a witness for an NP relation, while revealing nothing beyond the validity of the statement. Identification protocols [FFS87], naturally derived from public-coin ZKAoKs, can be compiled into a standard non-interactive signature scheme using the Fiat–Shamir transformation [FS86].

Multi-party Computation in the Head (MPCitH) [IKOS07, FR23, FR25] and Vector Oblivious Linear Evaluation in the Head (VOLEitH) [BBdSG+23b] are two powerful techniques for constructing ZKAoKs. Despite being motivated by different conceptual starting points, MPCitH and VOLEitH follow similar high-level frameworks and rely on common building blocks, making them naturally suited to analogous cryptographic applications such as post-quantum signatures. In the third round of NIST's additional post-quantum digital signature standardization process [NIS22, NIS26], 3 of the 9 candidates belong to the MPC/VOLEitH family.

In MPC/VOLEitH-based signature schemes, a ZKAoK typically proves knowledge of a preimage for a prescribed one-way function (OWF). The OWFs used in these schemes may rely on diverse assumptions, including well-studied NP-hard problems such as syndrome decoding [FJR22, MGH+23, CCJ23, CLY+24, OTX24, MBB+25], the permuted kernel problem [BBGK24, ABB+25a], and the multivariate quadratic problem [BFR24, BKPV24, BBFR25], as well as block ciphers such as AES [BBdSG+23b, BBM+24, BBB+25a]. The chosen OWF is then represented as a system of polynomial constraints evaluated on the witness, and the signature size usually scales linearly with the witness length and the maximum degree of polynomials. Therefore, compressing the witness and

reducing the polynomial degrees are important for the size reduction of MPC/VOLEitH-based signatures. We refer to *modeling* as the approach to representing a problem by polynomial constraints.

This paper focuses on signatures from syndrome decoding (SD) and the permuted kernel problem (PKP). SD for random linear codes is one of the central hardness assumptions in code-based cryptography. Its relevance to post-quantum signatures is further reflected by SDitH [MBB+25], a VOLEitH-based signature from SD, being selected as a third-round candidate in NIST's additional post-quantum signature process due to its conservative hardness assumption [NIS26].

Despite relying on different OWFs, SD- and PKP-based signatures use similar modeling approaches. They first represent the secret inputs of OWFs as elementary vectors. Then, they encode each elementary vector into a compressed vector and build the corresponding reconstruction polynomial constraint system. This system recovers the original elementary vector from the compressed one while proving that it has weight one. Differences lie in the modeling approaches for elementary vectors, which lead to different trade-offs in performance. Specifically, CLY+24 [CLY+24] explicitly include the entire elementary vectors in the witness and prove that each vector has a Hamming weight of one, which incurs a relatively simple constraint but a large witness. In contrast, BBG+24 [BBGK24] and OTX24 [OTX24] use the position indices of the non-zero coordinates as the stored witness and reconstruct the elementary vectors through Lagrange interpolation, yielding an optimal-size witness at the price of high-degree constraints. SDitH [MBB+25] generalizes the approaches in BBG+24 and OTX24 by allowing each elementary vector to be split into multiple shorter elementary vectors, exploring a better trade-off in signature size via a larger witness but reduced polynomial degrees. Chou25 [Cho25] further allows the shorter vectors to have higher prescribed Hamming weights. Taken together, these works show that the choice of representation directly affects signature size, while leaving open the characterization of the most compact elementary-vector representation achievable under a given polynomial-degree bound.

In this paper, we ask: *How much can we minimize the sizes of SD/PKP-based signature schemes by optimizing the modeling approaches of elementary vectors?*

1.1 Our Contribution

The contributions of this paper are as follows.

A framework for elementary-vector modeling approaches. We propose a general framework for modeling elementary vectors, which describes how to encode an elementary vector into a compressed binary string and reconstruct it using low-degree polynomials. This framework uniformly explains the constructions in existing SD/PKP-based works [CLY+24, OTX24, BBGK24, MBB+25, Cho25], guides our modeling approach, and may be helpful for future work using elementary vectors.

A more compact VOLEitH-friendly modeling approach. Based on our framework, we give a more compact modeling approach for elementary vectors in VOLEitH-based signatures. We further prove the upper bound of elementary-vector length given a fixed witness length and degree. Our concrete constructions attain the bound for practical parameters.

Table 1. Comparison of elementary-vector modeling approaches in SD/PKP-based signature schemes under NIST's L1 security ($\lambda = 128$)

Problem	Modeling	Witness length $\lvert\boldsymbol{w}\rvert/8$ (bytes)	Constraint degree d	Extended witness size $\big(\lvert\boldsymbol{w}\rvert + \lambda \cdot (d-1)\big)/8$ (bytes)
SD	BBG+24 [BBGK24]	72	6	152
	SDitH [MBB+25]	77	4	125
	Chou25 [Cho25]	63	4	111
	This work	55	4	103
PKP	BBG+24 [BBGK24]	48	6	128
	PERK [ABB+25a]	72	3	104
	Chou25 [Cho25]	64	3	96
	This work	55	3	87

Table 1 compares the resulting witness and extended-witness sizes at NIST Level I. For SD, our extended witness is 32.2%, 17.6%, and 7.2% smaller than those of BBG+24 [BBGK24], SDitH [MBB+25], and Chou25 [Cho25], respectively. For PKP, the corresponding reductions over BBG+24 [BBGK24], PERK [ABB+25a], and Chou25 [Cho25] are 32.0%, 16.3%, and 9.4%. At NIST Levels III and V, the SD extended witness is 16.5% and 17.2% smaller than that of SDitH [MBB+25], while the PKP extended witness is 26.2% and 19.3% smaller than that of PERK [ABB+25a].

Shorter and efficient signatures from SD and PKP. Our optimized modeling approach leads to shorter VOLEitH-based signature schemes without sacrificing efficiency. At NIST's Level I, our SD- and PKP-based signatures are 3 431 and 3 239 bytes, respectively, which are 7.4% and 6.7% smaller than the NIST submissions of SDitH [MBB+25] and PERK [ABB+25a]. Under the standard "signature + public key" metric, the corresponding sizes are 3 496 and 3 341 bytes; both are smaller than ML-DSA and more than 50% smaller than SLH-DSA. At higher security levels, our SD/PKP-based signature schemes inherit the compactness. At NIST Levels III and V, the SD-based scheme is 6.7% and 6.7% shorter than SDitH [MBB+25], while the PKP-based scheme is 10.7% and 8.0% shorter than PERK [ABB+25a].

By representing longer elementary vectors more compactly, our modeling further allows the joint search to select larger SD parameters for shorter signatures. Our evaluation optimization offsets the resulting computational cost, making our SD-based scheme faster than SDitH for all parameter sets.

Applications to t-out-of-n proofs and ring signatures. We further extend the framework from elementary-vector targets to arbitrary target vectors, yielding compact t-out-of-n proofs. Combining the resulting one-out-of-n selector with the generic ring-signature construction of Feneuil and Rivain [FR25], we obtain an SD-based ring signature whose size increases from 3 563 bytes for 2^{10} members to only 3 673 bytes for 2^{20} members. Thus, increasing the ring size by a factor of 1 024 adds only 110 bytes to the signature. At $N = 2^{20}$, our signature is 27% and 49% smaller than the AES- and SD-based signatures in CDD+25 [CDD+25] and BBG+24 [BBGK24], respectively. The comparison with the lattice-based BLOOM [LN22] shows a signature-size reduction of more than 70%.

1.2 Technical Overview

A framework for elementary-vector modeling approaches. Our general framework separates the modeling problem into two parts. The first part is to choose a compact representative set $S = \{\boldsymbol{s}_i\}_{i\in[\mu]}$, where $\boldsymbol{s}_i \in \mathbb{F}_2^{\eta}$, for the set $\mathcal{B}_\mu = \{\boldsymbol{e}_i\}_{i\in[\mu]}$ of length-μ elementary vectors. The second part is to reconstruct each elementary vector $\boldsymbol{e}_i$ from its representative $\boldsymbol{s}_i$ using low-degree polynomials. We denote this reconstruction map by $\phi : S \to \mathcal{B}_\mu$, where $\phi(\boldsymbol{s}_i) = \boldsymbol{e}_i$.

In this view, previous constructions can be seen as special cases through different choices of representatives. CLY+24 [CLY+24] uses the elementary vector itself. The index-based construction of OTX24 [OTX24] and BBG+24 [BBGK24] uses the binary representation of the non-zero position. The tensor-product construction in SDitH [MBB+25] represents an elementary vector using several smaller weight-one vectors, while Chou25 [Cho25] uses several high-weight vectors for the same purpose.

However, an arbitrary choice of S does not automatically give low-degree reconstruction polynomials. Fixing a S, it may be infeasible to realize ϕ using polynomials within the desired degree bound d. Thus, we also need a criterion that tells when the low-degree reconstruction is possible.

We give a clean criterion. We observe that once S is fixed, the reconstruction polynomials can be obtained by interpolation. If $\boldsymbol{A}$ denotes their coefficient matrix, the interpolation conditions are exactly $\boldsymbol{G}_{S,d}\boldsymbol{A} = \boldsymbol{I}_\mu$, where $\boldsymbol{I}_\mu$ is a matrix composed of μ elementary vectors. The interpolation succeeds exactly when the public evaluation matrix $\boldsymbol{G}_{S,d}$, obtained by evaluating all degree-$\leq d$ multilinear monomials on S, has full row rank. Thus, the problem becomes choosing compact representatives while preserving this full-rank property, which turns the modeling problem into a specific linear-algebraic design problem.

Generalizing block-wise representations of elementary vectors. Based on our framework, we explore how to reduce the extended witness size and the resulting signature size. Our starting point is the block-wise representation used in previous works: SDitH [MBB+25] represents an elementary vector by several smaller weight-one vectors, whereas Chou25 [Cho25] allows these sub-vectors to have higher prescribed weights. Writing their concatenation as $\boldsymbol{s} = (\boldsymbol{s}^{(1)}, \ldots, \boldsymbol{s}^{(\ell)})$, where $\boldsymbol{s}^{(j)}$ has length η_j, a weight-one sub-vector provides η_j possible values, while a weight-r_j sub-vector provides $\binom{\eta_j}{r_j}$. Higher weights can therefore represent more elementary vectors with the same number of bits, or equivalently reduce the witness and signature sizes for a prescribed elementary-vector length.

Our framework generalizes these representations by allowing the number and lengths of the subvectors and their allowed weight tuples $(\mathsf{wt}(\boldsymbol{s}^{(1)}), \ldots, \mathsf{wt}(\boldsymbol{s}^{(\ell)}))$ to vary, while previous approaches are recovered through particular parameter choices. The framework also explains why this family admits low-degree reconstruction. If every allowed weight tuple has a total weight of at most d, then each representative has a support monomial of degree at most d. Ordering the representatives by Hamming weight makes these monomials form a triangular evaluation submatrix with unit diagonal. Hence, $\boldsymbol{G}_{S,d}$ has full row rank, and ϕ can be realized by polynomials of degree at most d.

Because these reconstruction polynomials are only guaranteed to behave correctly on S, we must additionally prove that the stored representative belongs to S, which we call the membership proof. The prescribed weight conditions make this convenient: a weight upper bound r_j on $\boldsymbol{s}^{(j)}$ is enforced by requiring every product of $r_j + 1$ distinct coordinates to vanish, while exact weight r_j additionally requires the sum of all r_j-fold products to be one, ruling out smaller weights. When $r_j + 1 \leq d$, these constraints do not increase the reconstruction degree. Thus, the generalization gives a broad structured family supporting both low-degree reconstruction and membership, while the remaining restrictions on S motivate the more general constructions considered next.

Optimal elementary-vector representations beyond Hamming weights. The block-wise structure above makes both reconstruction and membership feasible to establish, but its Hamming-weight restrictions are sufficient conditions rather than intrinsic requirements. We therefore study the maximum size of an admissible representative set $S \subseteq \mathbb{F}_2^\eta$ for a fixed representation length η and degree bound d. This maximum determines the longest elementary vector that can be represented by the same η bits without increasing the polynomial degree.

Our key observation is that the size of S is governed by how much of the available polynomial space vanishes on S. The space $\mathcal{P}_{\eta,d}$ of degree-$\leq d$ multilinear polynomials has fixed dimension $k_d(\eta) = \sum_{j=0}^{d} \binom{\eta}{j}$. Since $\boldsymbol{G}_{S,d}$ has full row rank, evaluation on S has rank $|S|$, and its kernel consists exactly of the polynomials that vanish on S. Thus, $|S| + \dim\ker(\boldsymbol{G}_{S,d}) = k_d(\eta)$. The membership equations must come from this kernel. We prove that every admissible S requires $\dim\ker(\boldsymbol{G}_{S,d}) \geq \lceil \eta/d \rceil - 1$, which yields the universal upper bound

$$|S| \leq k_d(\eta) - \left\lceil \frac{\eta}{d} \right\rceil + 1$$

for our modeling framework. The Hamming-weight constructions above have larger kernels, leaving fewer dimensions to distinguish representatives. In contrast, our explicit constructions attain the minimum kernel dimension for practical parameters, using all remaining dimensions to enlarge S. For $d = 4$, our constructions improve over the best block-wise Hamming-ball constructions by 32.3%, 37.6%, 43.6%, and 50.0% for $\eta = 9, 10, 11, 12$, respectively. The resulting sets contain dozens to hundreds of representatives, covering the sizes required by our concrete SD- and PKP-based signatures.

Instantiating SD and PKP signatures. We instantiate our framework with SD and PKP as OWFs to obtain two smaller signature schemes, respectively.

For parameter selection, we choose the OWF and modeling parameters jointly. Notably, in the SD instantiation, we do not simply reuse the parameters of SDitH, but perform a search over candidate parameters. For each target security level, we first keep only OWF parameters whose estimated attack cost meets the required security, and then enumerate admissible modeling parameters and check whether they can represent the required elementary-vector length while keeping the reconstruction and membership constraints within the desired degree. Among all valid candidates, we take the minimal extended witness size.

Although our parameter choices are larger than those of SDitH, which would normally lead to more multiplications and slower proving, our optimization reduces the number of required multiplications, so the shorter signatures do not come at a prohibitive cost in running time.

Extending the framework beyond elementary vectors. Our framework is not restricted to reconstructing elementary vectors. The identity matrix in $\boldsymbol{G}_{S,d}\boldsymbol{A} = \boldsymbol{I}_\mu$ arises only because the reconstruction targets above are the elementary vectors. More generally, for any target set $\mathcal{Y} = \{\boldsymbol{y}_i\}_{i\in[\mu]} \subseteq \mathbb{F}_2^n$, let $\boldsymbol{Y}$ be the matrix whose i-th row is $\boldsymbol{y}_i^T$ and replace the interpolation system by $\boldsymbol{G}_{S,d}\boldsymbol{A} = \boldsymbol{Y}$. If $\boldsymbol{G}_{S,d}$ has full row rank, this system has a solution for every $\boldsymbol{Y}$. Hence, the same representative set can reconstruct arbitrary target vectors, including vectors of higher Hamming weight.

Choosing $\mathcal{Y}$ as the set of length-n binary vectors of Hamming weight t makes the reconstructed output $\boldsymbol{y} = (y_1, \ldots, y_n)$ a t-out-of-n selector. The prover can then aggregate the selected relations as

$\sum_{i=1}^{n} r_i y_i \mathcal{C}_i(\boldsymbol{w}_i) = 0$, where each r_i is random challenge, $y_i \in \{0,1\}$ identifies whether the statement is selected, and $\mathcal{C}_i(\boldsymbol{w}_i) = 0$ validates its corresponding witness $\boldsymbol{w}_i$. This gives a compact proof of knowledge of witnesses for t among n statements.

When $t = 1$, this becomes a one-out-of-n selector and can be used in the generic ring-signature construction of Feneuil and Rivain [FR25, Section 6.6]. In their construction, the selector chooses one public key pk_i from the ring, after which the prover proves knowledge of a corresponding secret key $\boldsymbol{w}$ satisfying $F(\boldsymbol{w}) = \mathsf{pk}_i$; we use our framework to represent this selector compactly. For a ring of 2^{20} members, we take a size-2^{20} subset of the product of two length-45 Hamming balls, both of radius 2. The product contains $(1 + 45 + \binom{45}{2})^2 = 1\,073\,296 > 2^{20}$ representatives, so the selector requires only 90 witness bits.

2 Preliminaries

2.1 Notations

$[n]$ denotes the set $\{1, \ldots, n\}$ and $[m, n]$ denotes $\{m, m+1, \ldots, n\}$ for non-negative integers $m < n$. Bold small letters denote vectors. Given a vector $\boldsymbol{a}$, use $\boldsymbol{a}[i..j]$ to represent a sub-vector of $\boldsymbol{a}$ from its i-th entry to the j-th entry, and use $(\boldsymbol{a})_i$ to denote its i-th entry. Bold capital letters denote matrices. Let $f \in \mathbb{F}[X_1, \ldots, X_\eta]_{\leq d}$ be a polynomial over the finite field $\mathbb{F}$ with η variables and a degree bound d for each variable.

2.2 Syndrome Decoding Problem

This section recalls the definitions of the SD problem and its variant named as regular syndrome decoding (RSD) problem, which we directly use for OWFs and has been proven to be tightly reduced into the SD problem.

Definition 1 (Syndrome Decoding Problem). *Let q, n, k and w be positive integers such that $n > k$ and $n > w$. The syndrome decoding problem with parameters (q, n, k, w) is the following problem:*

Let $\boldsymbol{H}$, $\boldsymbol{x}$ and $\boldsymbol{y}$ be such that:

1. *$\boldsymbol{H}$ is uniformly sampled from $\mathbb{F}_q^{(n-k)\times n}$,*
2. *$\boldsymbol{x}$ is uniformly sampled from $\{\boldsymbol{x} \in \mathbb{F}_q^n : \mathsf{wt}(\boldsymbol{x}) = w\}$,*
3. *$\boldsymbol{y}$ is defined as $\boldsymbol{y} := \boldsymbol{H}\boldsymbol{x}$.*

From $(\boldsymbol{H}, \boldsymbol{y})$, find $\boldsymbol{x}$.

Definition 2 (Regular Syndrome Decoding Problem). *Let q, n, k and w be positive integers such that $n > k$, $n > w$, and $w \mid n$. The regular syndrome decoding problem with parameters (q, n, k, w) is the following problem:*

Let $\boldsymbol{H}$, $\boldsymbol{x}$ and $\boldsymbol{y}$ be such that:

1. *$\boldsymbol{H}$ is uniformly sampled from $\mathbb{F}_q^{(n-k)\times n}$,*
2. *$\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_w) \in \mathbb{F}_q^n$ is uniformly sampled such that $\boldsymbol{x}_i \in \mathbb{F}_q^{n/w}$ and $\mathsf{wt}(\boldsymbol{x}_i) = 1$ for all $i \in [w]$,*
3. *$\boldsymbol{y}$ is defined as $\boldsymbol{y} := \boldsymbol{H}\boldsymbol{x}$.*

From $(\boldsymbol{H}, \boldsymbol{y})$, find $\boldsymbol{x}$.

Following prior work [FJR22, BBGK24, MBB+25], we use regular syndrome decoding as a modeling tool rather than as an underlying hardness assumption. Concretely, we instantiate the ZKPoK for a regular syndrome decoding instance for which there exists a provable reduction to a secure unstructured syndrome decoding instance. This ensures that our construction relies only on the standard syndrome decoding problem while benefiting from the efficiency of modeling solutions with regular structure. We formally describe the reduction below.

Theorem 3 ([MBB+25]). *Let $\mathbb{F}_q$ be a finite field. Let $\theta \in [1, \infty)$. Let n, k_{SD}, k_{RSD}, w be positive integers such that $k_{\mathsf{SD}} \leq k_{\mathsf{RSD}} \leq n$, $w < n$, $w \mid n$, and*

$$k_{\mathsf{RSD}} \leq n - w \cdot \log_2\left(\frac{n}{w}\right) - \log_2 \theta.$$

Let $\mathcal{A}_{\mathsf{RSD}}$ be an algorithm solving a random $(q, n, k_{\mathsf{RSD}}, w)$-instance of the regular syndrome decoding problem in time t with success probability $\varepsilon_{\mathsf{RSD}}$. Then there exists an algorithm $\mathcal{A}_{\mathsf{SD}}$ solving a random $(q, n, k_{\mathsf{SD}}, w)$-instance of the standard syndrome decoding problem in time t with probability $\varepsilon_{\mathsf{SD}}$, where

$$\varepsilon_{\mathsf{SD}} \geq e^{-1/\theta} \cdot \frac{\left(\frac{n}{w}\right)^w}{\binom{n}{w}} \cdot \varepsilon_{\mathsf{RSD}}.$$

2.3 Permuted Kernel Problem

Definition 4 (Permuted Kernel Problem). *Let q, n and m be positive integers such that $m < n$. The permuted kernel problem (PKP) with parameters (q, n, m) is the following problem:*

Let $\boldsymbol{H}$, $\boldsymbol{x}$ and π be such that:
1. *$\boldsymbol{H} \in \mathbb{F}_q^{m \times n}$,*
2. *$\boldsymbol{x} \in \mathbb{F}_q^n$,*
3. *$\pi \in S_n$ is a permutation such that $\boldsymbol{H}\pi[\boldsymbol{x}] = \boldsymbol{0}$.*

From $(\boldsymbol{H}, \boldsymbol{x})$, find a $\tilde{\pi} \in S_n$ such that $\boldsymbol{H}\tilde{\pi}[\boldsymbol{x}] = \boldsymbol{0}$.

As in BBG+24 [BBGK24], we can write a permutation π as the corresponding permutation matrix $\boldsymbol{P}$, so that $\pi[\boldsymbol{x}] = \boldsymbol{P}\boldsymbol{x}$ and $\boldsymbol{H}\boldsymbol{P}\boldsymbol{x} = \boldsymbol{0}$.

3 A Framework for Elementary-Vector Encoding

This section develops a unified framework for encoding and reconstructing elementary vectors using low-degree polynomials. We first review the Lagrange-basis encoding [OTX24, BBGK24], tensor-product encoding [MBB+25], and support-monomial encoding [Cho25] under a common viewpoint. We then present our general framework for degree-bounded polynomial reconstruction and admissible representative sets, which recovers existing approaches as special cases through different parameter choices and structural restrictions. Finally, we derive a universal upper bound on the representative-set size and give constructions attaining it for practical parameters. Table 2 summarizes the notation used throughout this section.

Table 2. Main notations used in Section 3

Notation	Meaning
$\mathcal{B}_\mu$	Standard basis of $\mathbb{F}_2^\mu$
$\boldsymbol{e}_i$	The i-th elementary vector in $\mathcal{B}_\mu$
μ	Length of the reconstructed elementary vector
S	Representative set for elementary vectors
η	Length of the embedded binary representation
η'	Stored witness length after possible compression
ϕ	Reconstruction bijection from S to $\mathcal{B}_\mu$
d	Degree bound for reconstruction and membership polynomials
$\mathcal{H}_{\eta,r}$	Set of binary vectors in $\mathbb{F}_2^\eta$ of weight at most r
$\mathcal{E}_{\eta,r}$	Set of binary vectors in $\mathbb{F}_2^\eta$ of weight exactly r
ℓ	Number of blocks
η_j	Length of the j-th block
r_j	Weight parameter of the j-th block
$\rho_{j,T}$	Polynomial enforcing the weight upper bound of the j-th block
σ_j	Polynomial enforcing the weight lower bound of the j-th block

3.1 Warm-up: Existing Elementary-Vector Encodings

As recalled in Section 2, solutions to both RSD and PKP can be represented as elementary vectors. To reduce the signature size, these vectors are encoded into shorter representations and later reconstructed during signature generation. Prior works instantiate this step by representing an elementary vector as the binary encoding of its nonzero position [OTX24, BBGK24], as a tuple of smaller elementary vectors [MBB+25], or as a structured binary vector of fixed Hamming weight [Cho25]. This subsection reviews these approaches through their encoding and reconstruction mechanisms before presenting our framework in Section 3.2.

A common viewpoint. Let η, μ be positive integers. Let $\mathcal{B}_\mu = \{\boldsymbol{e}_i\}_{i\in[\mu]}$ be the standard basis of $\mathbb{F}_2^\mu$, where $\boldsymbol{e}_i$ is the i-th elementary vector. Let $S = \{\boldsymbol{s}_i\}_{i\in[\mu]} \subseteq \mathbb{F}_2^\eta$ be a representative set of size μ, where each $\boldsymbol{s}_i$ represents the elementary vector $\boldsymbol{e}_i$. We formalize this correspondence by a public bijection $\phi : S \to \mathcal{B}_\mu$ satisfying $\phi(\boldsymbol{s}_i) = \boldsymbol{e}_i$ for every $i \in [\mu]$. We call ϕ the reconstruction map and its inverse $\phi^{-1} : \mathcal{B}_\mu \to S$ the encoding map.

Since S contains μ distinct vectors, we necessarily have $\eta \geq \lceil \log_2 \mu \rceil$. On the other hand, using the elementary vector itself corresponds to $\eta = \mu$. We therefore restrict attention to $\lceil \log_2 \mu \rceil \leq \eta \leq \mu$. Under this viewpoint, different encoding or reconstruction approaches correspond to different choices of S and ϕ.

Lagrange-basis encoding. In OTX24 [OTX24] and BBG+24 [BBGK24], an elementary vector is represented by the binary index of its non-zero entry position. Assume that μ is a power of two. For every $j \in [\mu]$, let $(j_1, \ldots, j_{\log_2 \mu})$ denote the bit decomposition of $j-1$. Over $\mathbb{F}_2$, the reconstruction polynomial for the j-th coordinate is $f_j^{\mathsf{Lag}}(X_1, \ldots, X_{\log_2 \mu}) = \prod_{k=1}^{\log_2 \mu}(1 + j_k + X_k)$.

For every $i, j \in [\mu]$, we have $f_j^{\mathsf{Lag}}(i_1, \ldots, i_{\log_2 \mu}) = (\boldsymbol{e}_i)_j$. This construction fits the common viewpoint above as follows:

- The representative set is $S = \mathbb{F}_2^{\log_2 \mu}$, consisting of the binary representations of the integers in $\{0, \ldots, \mu - 1\}$.
- The bijection $\phi\colon \mathbb{F}_2^{\log_2 \mu} \to \mathcal{B}_\mu$ is defined as

$$\phi : \quad (i_1, \ldots, i_{\log_2 \mu}) \quad \longmapsto \quad (f_1^{\mathsf{Lag}}(i_1, \ldots, i_{\log_2 \mu}), \ldots, f_\mu^{\mathsf{Lag}}(i_1, \ldots, i_{\log_2 \mu})).$$

- The encoding of $\boldsymbol{e}_i$ is $(i_1, \ldots, i_{\log_2 \mu})$, of length $\eta = \log_2 \mu$.
- Every reconstruction polynomial has degree $\log_2 \mu$.

The representation length is optimal, but the reconstruction degree is not necessarily well-suited to MPC/VOLEitH. In typical parameter regimes [BBGK24], it ranges from 6 to 12, substantially larger than the degree-2/3 constraints in the modeling of many other OWFs [BBFR25,ABB+25b,BBB+25a].

Tensor-product encoding. SDitH [MBB+25] uses tensor-product decomposition to obtain lower-degree reconstruction at the cost of a longer representation. Let $t_1, \ldots, t_d$ be positive integers satisfying $\prod_{r=1}^d t_r = \mu$. Every elementary vector $\boldsymbol{e}_i \in \mathcal{B}_\mu$ has a unique decomposition $\boldsymbol{e}_i = \boldsymbol{b}_1 \otimes \cdots \otimes \boldsymbol{b}_d$, where $\boldsymbol{b}_r \in \mathcal{B}_{t_r}$. Since every $\boldsymbol{b}_r$ has Hamming weight one, its last coordinate can be recovered as $(\boldsymbol{b}_r)_{t_r} = 1 - \sum_{a=1}^{t_r-1}(\boldsymbol{b}_r)_a$. Omitting it from each block reduces the representation length from $\eta = \sum_{r=1}^d t_r$ to $\eta' = \sum_{r=1}^d (t_r - 1) = \eta - d$.

We now express this reconstruction using polynomials, which will serve as a basis for our generalization. For every $j \in [\mu]$, let $(j_1, \ldots, j_d)$ denote the mixed-radix representation of $j-1$ with radices $(t_1, \ldots, t_d)$; namely, it is the unique tuple satisfying $j_r \in [t_r]$ for every $r \in [d]$ and $j - 1 = \sum_{r=1}^d (j_r - 1)\prod_{\ell=r+1}^d t_\ell$, where $\prod_{\ell=d+1}^d t_\ell = 1$. The resulting reconstruction polynomial is $f_j^{\mathsf{Ten}}(\boldsymbol{X}_1, \ldots, \boldsymbol{X}_d) = \prod_{r=1}^d (\boldsymbol{X}_r)_{j_r}$. For notational convenience, each $\boldsymbol{X}_r$ denotes a vector of indeterminates, although f_j^{Ten} uses only its j_r-th coordinate. Evaluating f_j^{Ten} at $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d)$ gives 1 exactly when j_r is the position of the nonzero entry of $\boldsymbol{b}_r$ for every $r \in [d]$. Hence, evaluating all the reconstruction polynomials recovers $\boldsymbol{b}_1 \otimes \cdots \otimes \boldsymbol{b}_d$.

This construction fits the common viewpoint above as follows:

- The representative set is $S = \mathcal{B}_{t_1} \times \cdots \times \mathcal{B}_{t_d}$, and hence $|S| = \prod_{r=1}^d t_r = \mu$.
- The bijection $\phi : S \to \mathcal{B}_\mu$ is defined by

$$\phi : \quad (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d) \quad \longmapsto \quad (f_1^{\mathsf{Ten}}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d), \ldots, f_\mu^{\mathsf{Ten}}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d)).$$

- The encoding of $\boldsymbol{e}_i$ is the unique tuple $\phi^{-1}(\boldsymbol{e}_i)$. Omitting the last coordinate of every block reduces the stored length from $\eta = \sum_{r=1}^{d} t_r$ to $\eta' = \sum_{r=1}^{d}(t_r - 1) = \eta - d$.
- Every reconstruction polynomial has degree d, and substituting the affine recovery expressions does not increase this degree. In SDitH, d is a small constant, typically at most 4 [MBB+25].

Membership proof. After recovering the omitted coordinates, it remains to verify that every $\boldsymbol{b}_r$ belongs to $\mathcal{B}_{t_r}$. For $r \in [d]$ and $1 \le u < u' \le t_r$, define $\rho_{r,u,u'}(\boldsymbol{X}_r) = (\boldsymbol{X}_r)_u(\boldsymbol{X}_r)_{u'}$. The equations $\rho_{r,u,u'}(\boldsymbol{b}_r) = 0$ ensure that $\boldsymbol{b}_r$ has at most one nonzero coordinate, while the recovery equation $(\boldsymbol{b}_r)_{t_r} = 1 - \sum_{u=1}^{t_r-1}(\boldsymbol{b}_r)_u$ rules out the all-zero vector. Together, these equations characterize $\boldsymbol{b}_r \in \mathcal{B}_{t_r}$. Each polynomial has degree at most 2, so the membership proof does not increase the overall degree bound when $d \ge 2$.

Support-monomial encoding. Chou25 [Cho25] represents an elementary vector by a binary vector $\boldsymbol{v} \in \mathbb{F}_2^{\eta}$ with $\mathsf{wt}(\boldsymbol{v}) = d$. For every $\boldsymbol{v} \in \mathbb{F}_2^{\eta}$, define $\mathrm{prod}_d(\boldsymbol{v}) = \left(\prod_{j\in T} v_j\right)_{T\in\binom{[\eta]}{d}}$. For every $T \in \binom{[\eta]}{d}$, the corresponding coordinate equals 1 if and only if $v_j = 1$ for all $j \in T$. Since $\mathsf{wt}(\boldsymbol{v}) = d$, there is a unique such T, namely the set of nonzero positions of $\boldsymbol{v}$. Multiplication by the public matrix $\boldsymbol{M}$ therefore selects the corresponding column. Chou25 assigns selected columns of $\boldsymbol{M}$ to target elementary vectors and sets the remaining columns to zero, so a representation of $\boldsymbol{e}_i$ satisfies $\boldsymbol{e}_i = \boldsymbol{M}\,\mathrm{prod}_d(\boldsymbol{v})$.

For the RSD application, $d = 4$ and $\boldsymbol{v} = (\boldsymbol{v}^{(1)}, \boldsymbol{v}^{(2)}) \in \mathbb{F}_2^{\eta_1} \times \mathbb{F}_2^{\eta_2}$, where $\eta_1 + \eta_2 = \eta$. For a set $\mathcal{R}$ of block-weight pairs satisfying $0 \le r_i \le \eta_i$ where $i \in [2]$ and $r_1 + r_2 = 4$, define

$$S_{\mathcal{R}} := \left\{ \boldsymbol{v} = (\boldsymbol{v}^{(1)}, \boldsymbol{v}^{(2)}) \in \mathbb{F}_2^{\eta_1} \times \mathbb{F}_2^{\eta_2} : \left(\mathsf{wt}(\boldsymbol{v}^{(1)}), \mathsf{wt}(\boldsymbol{v}^{(2)})\right) \in \mathcal{R} \right\}.$$

Chou25 takes $\mathcal{R}_A := \{(2,2)\}$ in [Cho25, Sec. 5.1] and $\mathcal{R}_B := \{(1,3),(2,2),(3,1)\}$ in [Cho25, Sec. 5.2]. For either choice, set $S = S_{\mathcal{R}} = \{\boldsymbol{s}_i\}_{i\in[\mu]}$, where $\mu = |S|$. For every $i \in [\mu]$, the corresponding index set and reconstruction polynomial are

$$T_i := \{j \in [\eta] : (\boldsymbol{s}_i)_j = 1\}, \qquad f_i^{\mathsf{Supp}}(X_1, \ldots, X_\eta) := \prod_{j\in T_i} X_j.$$

For the distinct weight-4 representatives, $f_i^{\mathsf{Supp}}(\boldsymbol{s}_k) = 1$ if and only if $i = k$.

This construction fits the common viewpoint above as follows:

- The representative set is $S = S_{\mathcal{R}}$, with $\mathcal{R} \in \{\mathcal{R}_A, \mathcal{R}_B\}$.
- The bijection $\phi : S \to \mathcal{B}_\mu$ is defined as

$$\phi : \quad \boldsymbol{v} \quad \longmapsto \quad \left(f_1^{\mathsf{Supp}}(\boldsymbol{v}), \ldots, f_\mu^{\mathsf{Supp}}(\boldsymbol{v})\right).$$

- The encoding of $\boldsymbol{e}_i$ is $\boldsymbol{s}_i$. For $\mathcal{R}_A$, one coordinate can be recovered from each fixed-weight block, giving $\eta' = \eta - 2$; for $\mathcal{R}_B$, the fixed total weight permits one omitted coordinate overall, giving $\eta' = \eta - 1$.
- Every reconstruction polynomial has degree 4, and substituting the affine recovery expressions does not increase this degree.

Chou25's original reconstruction $\boldsymbol{M}\,\mathrm{prod}_4 : S \to \mathcal{B}_N$, where $N \le \mu$, is surjective rather than bijective. In our common viewpoint, the same construction is expressed using the bijection $\phi : S \to \mathcal{B}_\mu$ above, followed by the public linear map $\boldsymbol{D} : \mathbb{F}_2^{\mu} \to \mathbb{F}_2^{N}$ induced by Chou25's column assignment. Thus, $\boldsymbol{M}\,\mathrm{prod}_4(\boldsymbol{v}) = \boldsymbol{D}\phi(\boldsymbol{v})$ for every $\boldsymbol{v} \in S$.

Membership proof. For $\boldsymbol{z} = (z_1, \ldots, z_m)$, write $e_t(\boldsymbol{z}) := \sum_{T\in\binom{[m]}{t}} \prod_{j\in T} z_j$ for the t-th elementary symmetric polynomial, and abbreviate $e_t^{(i)} := e_t(\boldsymbol{v}^{(i)})$. For $S_{\mathcal{R}_A}$, all cubic monomials in each block vanish and $e_2^{(i)} = 1$ for $i \in \{1,2\}$, giving block weights $(2,2)$. For $S_{\mathcal{R}_B}$, all quartic monomials vanish, and Chou25 imposes $e_1^{(i)} + e_2^{(i)} + e_3^{(i)} = 1$ for $i \in \{1,2\}$, $e_1^{(1)} + e_1^{(2)} = 0$, and $e_2^{(1)}e_1^{(2)} + e_1^{(1)}e_2^{(2)} = e_1^{(1)}$. These relations leave exactly $(1,3)$, $(2,2)$, and $(3,1)$. Thus, the two systems characterize $S_{\mathcal{R}_A}$ and $S_{\mathcal{R}_B}$, respectively, using polynomials of degree at most 4.

3.2 Generalizing Elementary-Vector Encoding and Reconstruction

The preceding constructions illustrate a trade-off between the representation length and the reconstruction degree. We now consider an arbitrary representative set S and characterize when its bijection $\phi : S \to \mathcal{B}_\mu$ can be realized by degree-$\leq d$ polynomials. In addition, it must be verified that the supplied binary vector lies in S using polynomial constraints of degree at most d.

Constructing the reconstruction polynomials. Recall that $S = \{\boldsymbol{s}_i\}_{i\in[\mu]} \subseteq \mathbb{F}_2^\eta$ and that the reconstruction bijection $\phi : S \to \mathcal{B}_\mu$ satisfies $\phi(\boldsymbol{s}_i) = \boldsymbol{e}_i$ for every $i \in [\mu]$. We first give a general polynomial realization of ϕ.

Let $\mathcal{P}_{\eta,d}$ denote the $\mathbb{F}_2$-vector space of all multilinear polynomials of degree at most d in η variables. The reconstruction requirement for $f_1, \ldots, f_\mu \in \mathcal{P}_{\eta,d}$ is

$$\phi(\boldsymbol{s}_i) = \big(f_1(\boldsymbol{s}_i), \ldots, f_\mu(\boldsymbol{s}_i)\big) = \boldsymbol{e}_i \qquad \text{for every } i \in [\mu].$$

We express these reconstruction polynomials in the degree-$\leq d$ monomial basis. Let $g^{\langle 1\rangle}, \ldots, g^{\langle k\rangle}$ denote all multilinear monomials of total degree at most d in η variables. These monomials form a basis of $\mathcal{P}_{\eta,d}$, where $k_d(\eta) = \dim(\mathcal{P}_{\eta,d}) = \sum_{t=0}^{d} \binom{\eta}{t}$, and we write $k = k_d(\eta)$. For each $j \in [\mu]$, we can therefore write

$$f_j(X_1, \ldots, X_\eta) = \sum_{h=1}^{k} \alpha_j^{\langle h\rangle} g^{\langle h\rangle}(X_1, \ldots, X_\eta), \tag{1}$$

where $\alpha_j^{\langle h\rangle} \in \mathbb{F}_2$ is the coefficient of $g^{\langle h\rangle}$ in f_j. The required coefficients need not exist for an arbitrary representative set S; we next characterize when they do.

Define $\boldsymbol{G}_{S,d} \in \mathbb{F}_2^{\mu\times k}$ by $(\boldsymbol{G}_{S,d})_{i,h} = g^{\langle h\rangle}(\boldsymbol{s}_i)$, and let $\boldsymbol{A} \in \mathbb{F}_2^{k\times\mu}$ be the coefficient matrix whose j-th column is $(\alpha_j^{\langle 1\rangle}, \ldots, \alpha_j^{\langle k\rangle})^\top$. Then the conditions $\phi(\boldsymbol{s}_i) = \boldsymbol{e}_i$ for all $i \in [\mu]$ are equivalent to

$$\boldsymbol{G}_{S,d}\boldsymbol{A} = \boldsymbol{I}_\mu. \tag{2}$$

Hence, the linear system has a solution if and only if $\boldsymbol{G}_{S,d}$ has full row rank, namely, $\mathrm{rank}(\boldsymbol{G}_{S,d}) = \mu$. Therefore, the full-row-rank condition exactly characterizes whether ϕ can be realized by reconstruction polynomials of degree at most d. Since the monomials $g^{\langle 1\rangle}, \ldots, g^{\langle k\rangle}$ are fixed and public, satisfying this condition depends only on the choice of the representative set S.

Arbitrary reconstruction targets. The identity matrix on the right-hand side of Equation (2) arises only because the reconstruction targets are the elementary vectors. More generally, for any family of target vectors $(\boldsymbol{y}_i)_{i\in[\mu]}$ in $\mathbb{F}_2^n$, let $\boldsymbol{Y} \in \mathbb{F}_2^{\mu\times n}$ be the matrix whose i-th row is $\boldsymbol{y}_i^T$. If $\boldsymbol{G}_{S,d}$ has full row rank, then

$$\boldsymbol{G}_{S,d}\boldsymbol{A}_{\mathcal{Y}} = \boldsymbol{Y}$$

has a solution $\boldsymbol{A}_{\mathcal{Y}} \in \mathbb{F}_2^{k\times n}$ for any $\boldsymbol{Y}$. Then, arbitrary targets can be reconstructed. This only changes $\boldsymbol{A}_{\mathcal{Y}}$, while S and its membership constraints remain.

Admissible representative sets. The reconstruction polynomials are only required to behave correctly on S, so it is additionally necessary to verify that a given binary vector belongs to S. For this purpose, we introduce a membership map $\varphi : \mathbb{F}_2^\eta \to \mathbb{F}_2^t$ for some positive integer t. We express φ in terms of coordinate polynomials $q_1, \ldots, q_t \in \mathcal{P}_{\eta,d}$ as

$$\varphi(\boldsymbol{x}) = \big(q_1(\boldsymbol{x}), \ldots, q_t(\boldsymbol{x})\big), \qquad \boldsymbol{x} \in \mathbb{F}_2^\eta.$$

We say that S admits membership constraints of degree at most d if φ vanishes exactly on S, namely, for every $\boldsymbol{x} \in \mathbb{F}_2^\eta$,

$$\boldsymbol{x} \in S \quad \Longleftrightarrow \quad \varphi(\boldsymbol{x}) = \boldsymbol{0}.$$

Let $\boldsymbol{B} \in \mathbb{F}_2^{k\times t}$ be the coefficient matrix of $q_1, \ldots, q_t$ with respect to the same monomial basis $g^{\langle 1\rangle}, \ldots, g^{\langle k\rangle}$ used for the reconstruction polynomials. Since $\phi(\boldsymbol{s}_i) = \boldsymbol{e}_i$ and $\varphi(\boldsymbol{s}_i) = \boldsymbol{0}$ for every $i \in [\mu]$, their polynomial realizations satisfy

$$\boldsymbol{G}_{S,d}\begin{bmatrix}\boldsymbol{A} & \boldsymbol{B}\end{bmatrix} = \begin{bmatrix}\boldsymbol{I}_\mu & \boldsymbol{0}_{\mu\times t}\end{bmatrix}.$$

Thus, the reconstruction and membership polynomials use the same monomial basis and the same evaluation matrix $\boldsymbol{G}_{S,d}$, but satisfy different evaluation conditions on S. Consequently, no polynomial can serve as both a reconstruction polynomial and a membership polynomial.

We call a representative set $S \subseteq \mathbb{F}_2^\eta$ admissible under degree bound d if $\boldsymbol{G}_{S,d}$ has full row rank and S admits membership constraints of degree at most d. For fixed η and d, our goal is to construct an admissible representative set S that is as large as possible, since $|S| = \mu$ determines the length of the elementary vectors represented by an η-bit vector. The next subsection gives structured constructions of admissible representative sets using high-weight sub-vectors.

3.3 Structured Representative-Set Constructions

Recall that the tensor-product-based construction in Section 3.1 splits each representative into several weight-one blocks. A weight-one block of length η_j provides only $\binom{\eta_j}{1} = \eta_j$ choices. Allowing weight r_j instead provides $\binom{\eta_j}{r_j}$ choices and therefore gives a natural heuristic for choosing a larger representative set S. This leads to two issues that we address below: whether $\boldsymbol{G}_{S,d}$ still has full row rank when the sub-vectors have larger Hamming weights, and whether membership in S can still be expressed by polynomial equations of degree at most d.

At first glance, the full Hamming ball $\mathcal{H}_{\eta,d}$ appears ideal: it has $k_d(\eta)$ representatives, and Lemma 5 guarantees full reconstruction rank. However, characterizing membership in $\mathcal{H}_{\eta,d}$ requires the degree-$(d+1)$ monomials to vanish, which exceeds the degree budget d. Replacing it by $\mathcal{H}_{\eta,d-1}$ keeps membership within degree d, but reduces the number of representatives from $\Theta(\eta^d)$ to $\Theta(\eta^{d-1})$ for fixed d. This motivates the block-wise construction below, which retains $\Theta(\eta^d)$ representatives while keeping both reconstruction and membership constraints within degree d.

Products of Hamming balls. We allow each sub-vector to have Hamming weight at most r_j. This increases the representational capacity of S while keeping the sub-vector length unchanged. Formally, let $\mathcal{H}_{\eta,r} = \{\boldsymbol{x} \in \mathbb{F}_2^\eta \mid \mathsf{wt}(\boldsymbol{x}) \le r\}$ be the Hamming ball of radius r in $\mathbb{F}_2^\eta$. Write each representative as $\boldsymbol{s} = (\boldsymbol{s}^{(1)}, \ldots, \boldsymbol{s}^{(\ell)})$, where $\boldsymbol{s}^{(j)} \in \mathbb{F}_2^{\eta_j}$ and $\sum_{j=1}^{\ell} \eta_j = \eta$, choose $0 \le r_j \le \eta_j$, and define $S = \mathcal{H}_{\eta_1,r_1} \times \cdots \times \mathcal{H}_{\eta_\ell,r_\ell}$. We show that this high-weight block-wise choice of S satisfies the full-rank requirements for polynomial reconstruction, as long as the blocks of S have a total Hamming weight of at most d (*i.e.*, $\sum_{j\in[\ell]} r_j \le d$), which is the degree bound of $f_j(X_1, \ldots, X_\eta)$ as shown in Equation (1).

Lemma 5. *Let $S \subseteq \mathbb{F}_2^\eta$ satisfy $\mathsf{wt}(\boldsymbol{s}) \le d$ for every $\boldsymbol{s} \in S$. Then the matrix $\boldsymbol{G}_{S,d}$ in Equation (2) has full row rank.*

Appendix A.1 proves Lemma 5 and derives its support-monomial refinement.

The proof and support-monomial refinement above show that there exist reconstruction polynomials $f_1, \ldots, f_\mu \in \mathcal{P}_{\eta,d}$ such that, for every $i \in [\mu]$,

$$\phi(\boldsymbol{s}_i) = \big(f_1(\boldsymbol{s}_i), \ldots, f_\mu(\boldsymbol{s}_i)\big) = \boldsymbol{e}_i.$$

This yields a degree-d reconstruction from an $\eta = \sum_j \eta_j$-bit representative to an elementary vector of length

$$\mu = \prod_{j=1}^{\ell} \sum_{t=0}^{r_j} \binom{\eta_j}{t}. \tag{3}$$

Note that by definition of Hamming balls, we have $\eta_j \ge r_j$ for each j. Since $\sum_{t=0}^{r_j} \binom{\eta_j}{t}$ is increasing in r_j, using the entire degree budget maximizes the number of representatives μ. Hence we take $\sum_{j\in[\ell]} r_j = d$ by default.

Similar to the tensor-product-based approach in Section 3.1, it remains to verify $\boldsymbol{s}_i^{(j)} \in \mathcal{H}_{\eta_j,r_j}$ for every block $j \in [\ell]$. Let $\binom{[n]}{m}$ denote the collection of all m-element subsets of $[n]$. For each $j \in [\ell]$ and $T \in \binom{[\eta_j]}{r_j+1}$, define

$$\rho_{j,T}(X_1, \ldots, X_{\eta_j}) = \prod_{t\in T} X_t. \tag{4}$$

For every $\boldsymbol{x} \in \mathbb{F}_2^{\eta_j}$, we have $\boldsymbol{x} \in \mathcal{H}_{\eta_j,r_j}$ if and only if $\rho_{j,T}(\boldsymbol{x}) = 0$ for all $T \in \binom{[\eta_j]}{r_j+1}$: a nonzero evaluation identifies $r_j + 1$ nonzero coordinates, while any vector of weight greater than r_j contains such a set of coordinates. Since $\deg(\rho_{j,T}) = r_j + 1$, these polynomials have degree at most d whenever $r_j + 1 \le d$ for every $j \in [\ell]$.

Products of Hamming spheres. We also consider the Hamming-sphere variant, where each subvector has a fixed Hamming weight rather than weight at most r_j. Let $\mathcal{E}_{\eta,r} = \{\boldsymbol{x} \in \mathbb{F}_2^{\eta} : \mathsf{wt}(\boldsymbol{x}) = r\}$ be the Hamming sphere of weight r. We set the representative set to $S = \mathcal{E}_{\eta_1,r_1} \times \cdots \times \mathcal{E}_{\eta_\ell,r_\ell}$.

For each $\boldsymbol{s}_j = (\boldsymbol{s}_j^{(1)}, \ldots, \boldsymbol{s}_j^{(\ell)}) \in S$, set $U_{j,k} := \mathrm{supp}(\boldsymbol{s}_j^{(k)}) \in \binom{[\eta_k]}{r_k}$ and define

$$f_j(X_1, \ldots, X_\eta) := \prod_{k=1}^{\ell} \prod_{t \in U_{j,k}} X_{k,t},$$

where $X_{k,t}$ is the t-th variable in block k. For every $\boldsymbol{s} \in S$,

$$f_j(\boldsymbol{s}) = 1 \quad \Longleftrightarrow \quad U_{j,k} \subseteq \mathrm{supp}(\boldsymbol{s}^{(k)}) \text{ for all } k \in [\ell] \quad \Longleftrightarrow \quad \boldsymbol{s} = \boldsymbol{s}_j,$$

because both sets in each inclusion have cardinality r_k. Thus f_j evaluates to $\boldsymbol{e}_j$ on S, and $\deg(f_j) = \sum_k r_k$. For $\sum_k r_k = d$, it is a degree-d monomial. For $\mathcal{H}_{\eta,r}$, the inclusion $U \subseteq \mathrm{supp}(\boldsymbol{x})$ may be strict, so reconstruction may require linear combinations of support monomials.

For each $j \in [\ell]$, let U range over the r_j-element subsets of $[\eta_j]$. To characterize $\mathcal{E}_{\eta_j,r_j}$, use the polynomials $\rho_{j,T}$ in Equation (4) together with

$$\sigma_j(X_1, \ldots, X_{\eta_j}) = 1 + \sum_{U} \prod_{u \in U} X_u. \tag{5}$$

If $\rho_{j,T}(\boldsymbol{x}) = 0$ for every $T \in \binom{[\eta_j]}{r_j+1}$, then $\mathsf{wt}(\boldsymbol{x}) \le r_j$. Under this restriction, the sum in Equation (5) equals $\binom{\mathsf{wt}(\boldsymbol{x})}{r_j}$, so $\sigma_j(\boldsymbol{x}) = 0$ if and only if $\mathsf{wt}(\boldsymbol{x}) = r_j$. Consequently, these equations characterize $\mathcal{E}_{\eta_j,r_j}$ with degrees $r_j + 1$ and r_j, both at most d whenever $r_j + 1 \le d$.

As in SDitH, a fixed-weight block of length η_j can be stored using its first $\eta_j - 1$ bits, with the last coordinate recovered from $(x_1, \ldots, x_{\eta_j-1})$ as $x_{\eta_j} \equiv r_j - \sum_{t=1}^{\eta_j-1} x_t \pmod 2$. Thus, elementary vectors of length $\mu = \prod_{j=1}^{\ell} \binom{\eta_j}{r_j}$ can be encoded using $\eta' = \sum_j \eta_j - \ell$ bits.

Finite unions of products of Hamming spheres. Finite unions of products of Hamming spheres strictly generalize both products of Hamming spheres and products of Hamming balls. Let $\mathcal{R}$ be a finite set of block-weight patterns $\mathbf{r} = (r_1, \ldots, r_\ell)$, where $0 \le r_j \le \eta_j$ for every $j \in [\ell]$ and $\sum_{j=1}^{\ell} r_j \le d$. The corresponding representative set is

$$S_{\mathcal{R}} := \bigcup_{\mathbf{r} \in \mathcal{R}} \left(\mathcal{E}_{\eta_1,r_1} \times \cdots \times \mathcal{E}_{\eta_\ell,r_\ell}\right).$$

The union is disjoint, and hence $|S_{\mathcal{R}}| = \sum_{\mathbf{r} \in \mathcal{R}} \prod_{j \in [\ell]} \binom{\eta_j}{r_j}$.

For every $\boldsymbol{s} \in S_{\mathcal{R}}$, there is some $\mathbf{r} \in \mathcal{R}$ such that $\mathsf{wt}(\boldsymbol{s}) = \sum_{j=1}^{\ell} r_j \le d$. Lemma 5 therefore implies that $\boldsymbol{G}_{S_{\mathcal{R}},d}$ has full row rank, and hence $S_{\mathcal{R}}$ admits reconstruction polynomials of degree at most d.

The two preceding constructions are recovered by restricting $\mathcal{R}$. Taking $\mathcal{R}$ to be a singleton gives a product of Hamming spheres. Taking $\mathcal{R} = \{0, \ldots, r_1\} \times \cdots \times \{0, \ldots, r_\ell\}$ gives $S_{\mathcal{R}} = \mathcal{H}_{\eta_1,r_1} \times \cdots \times \mathcal{H}_{\eta_\ell,r_\ell}$, namely a product of Hamming balls. A general $\mathcal{R}$ need be neither a singleton nor a Cartesian product of weight ranges, and may instead contain any selected collection of block-weight patterns.

Accordingly, the finite-union description guarantees degree-$\le d$ reconstruction, while admissibility additionally depends on a degree-$\le d$ membership proof tailored to the chosen $\mathcal{R}$. Such a proof must use polynomials whose common zero set is exactly $S_{\mathcal{R}}$; in particular, the polynomials used for a single product of Hamming spheres do not directly characterize an arbitrary union.

3.4 Previous Constructions as Special Cases

Our framework provides a unified way to view existing encoding approaches. In our framework, an encoding is specified by the number of blocks, the size of each block, and the allowed block-weight patterns. Different choices of these parameters recover existing approaches as special cases of our framework. Table 3 summarizes these choices. We detail the instantiations summarized in Table 3.

Table 3. Summary of existing modeling approaches under the unified framework

Modeling	# sub-vectors ℓ	Sub-vector sizes η or $(\eta_1, \ldots, \eta_\ell)$	Sub-vector weights r or $(r_1, \ldots, r_\ell)^\dagger$
OTX24 [OTX24], BBG+24 [BBGK24]	η	$(1, 1, \ldots, 1)$	$\leq (1, 1, \ldots, 1)$
SDitH [MBB+25]	d	$(\eta_1, \ldots, \eta_d)$	$= (1, 1, \ldots, 1)$
CLY+24 [CLY+24]	1	η	$= 1$
Chou25 [Cho25, Sec. 5.1]	2	(η_1, η_2)	$= (2, 2)$
Chou25 [Cho25, Sec. 5.2]	2	(η_1, η_2)	$\in \{(1, 3), (2, 2), (3, 1)\}$

† $\leq$ denotes the Hamming ball case, while $=$ denotes the Hamming sphere case.
$\in$ denotes the union of the listed block-weight patterns.

Lagrange-basis encoding. Set $\ell = \eta$ and $\eta_j = r_j = 1$ for every $j \in [\ell]$. Then $S = \mathcal{H}_{1,1} \times \cdots \times \mathcal{H}_{1,1} \cong \mathbb{F}_2^\eta$, $\mu = 2^\eta$, and the reconstruction degree is $d = \eta$. The resulting reconstruction polynomials are precisely the point indicators used by the Lagrange-basis encoding of [OTX24, BBGK24].

Tensor-product-based encoding. Set $\ell = d$ and $r_t = 1$ for every $t \in [d]$. Then $S = \mathcal{E}_{\eta_1,1} \times \cdots \times \mathcal{E}_{\eta_d,1}$ and $\mu = \prod_{t=1}^d \eta_t$. For every $\boldsymbol{s} \in S$, the reconstruction coordinate indexed by $(j_1, \ldots, j_d)$ is $\prod_{t=1}^d (\boldsymbol{s}^{(t)})_{j_t}$, exactly as in the tensor-product-based encoding of SDitH [MBB+25]. Omitting one coordinate from each block gives stored witness length $\sum_{t=1}^d (\eta_t - 1)$.

Identity encoding. Set $\ell = 1$, $\eta = \mu$, and $S = \mathcal{B}_\mu = \mathcal{E}_{\mu,1}$, with ϕ the identity. This recovers the CLY+24 construction [CLY+24], whose membership proof uses the sketching technique of [BGI16]; our Hamming-sphere construction instead uses explicit polynomial constraints.

Support-monomial encoding. For $d = 4$ and $\ell = 2$, Chou25 takes $\mathcal{R}_A := \{(2, 2)\}$ in [Cho25, Sec. 5.1] and $\mathcal{R}_B := \{(1, 3), (2, 2), (3, 1)\}$ in [Cho25, Sec. 5.2]. Hence,

$$S_{\mathcal{R}_A} = \mathcal{E}_{\eta_1,2} \times \mathcal{E}_{\eta_2,2},$$
$$S_{\mathcal{R}_B} = (\mathcal{E}_{\eta_1,1} \times \mathcal{E}_{\eta_2,3}) \cup (\mathcal{E}_{\eta_1,2} \times \mathcal{E}_{\eta_2,2}) \cup (\mathcal{E}_{\eta_1,3} \times \mathcal{E}_{\eta_2,1}).$$

Thus, the first is a product of Hamming spheres, while the second is a finite union of three such products.

3.5 Optimal Representative Sets

All structured representative sets considered above are defined through block-wise Hamming-weight restrictions, including exact weights and finite collections of allowed weight patterns. These restrictions provide a tractable way to design degree-bounded reconstruction and membership proofs, but confine S to a particular combinatorial form. This section maximizes $|S|$ over arbitrary admissible subsets of $\mathbb{F}_2^\eta$, separating the loss caused by the bounded-weight constructions from the unavoidable loss imposed by the degree bound.

An upper bound on the representative-set size. Recall that $\mathcal{P}_{\eta,d}$ has dimension $k_d(\eta)$, and let $\mu_d(\eta)$ be the maximum cardinality of an admissible subset of $\mathbb{F}_2^\eta$. For an admissible set S, define $c = k_d(\eta) - |S| = \dim \ker(\boldsymbol{G}_{S,d})$.

Theorem 6 (Universal upper bound). *For all η and d, every admissible $S \subseteq \mathbb{F}_2^\eta$ satisfies the following bound:*

$$c \geq \left\lceil \frac{\eta}{d} \right\rceil - 1, \qquad \textit{or equivalently} \qquad \mu_d(\eta) \leq k_d(\eta) - \left\lceil \frac{\eta}{d} \right\rceil + 1 =: U_d(\eta).$$

The proof is given in Appendix A.2. Briefly, a basis $q_1, \ldots, q_c$ of $\ker(\boldsymbol{G}_{S,d})$ determines an indicator polynomial $\mathbf{1}_S = \prod_{i=1}^c (1 + q_i)$ of degree at most cd. If $d(c+1) < \eta$, the rows of $\boldsymbol{G}_{S,d}$ sum to zero, contradicting full row rank.

Although the bound is universal, it need not be tight for every (η, d). Appendix A.2 gives $(\eta, d) = (5, 3)$ as such an example. We therefore claim equality only for parameter pairs for which an explicit construction attains the bound.

Constructions attaining the bound. Although we do not give constructions attaining the bound for arbitrary (η, d), Appendix B gives optimal constructions S_η for $d = 4$ with $9 \le \eta \le 12$, proving $\mu_4(\eta) = k_4(\eta) - 2$, together with useful optimal constructions for $d = 3$. Their sizes range from a few dozen to several hundred and cover the SD- and PKP-based parameters considered later.

To compare with the structured constructions in Section 3.3, let $B_4(\eta)$ be the largest cardinality of a product $S = \mathcal{H}_{\eta_1,r_1} \times \cdots \times \mathcal{H}_{\eta_\ell,r_\ell}$ for which that section provides reconstruction polynomials and a membership proof of degree at most 4. For $9 \le \eta \le 12$, the maximum among these products of Hamming balls is attained by $(\eta_1, \eta_2) = (2, \eta - 2)$ and $(r_1, r_2) = (1, 3)$, giving $B_4(\eta) = 3\sum_{t=0}^{3} \binom{\eta-2}{t}$, with respective values 192, 279, 390, and 528. For $9 \le \eta \le 12$, $\mu_4(\eta)$ exceeds $B_4(\eta)$ by 32.3%–50.0%. In particular, $\mu = 384$ requires $\eta \ge 11$ with Hamming balls but only $\eta = 10$ with S_{10}. Sections 4.3 and 4.5 show that these larger representative sets yield smaller extended-witness sizes for the resulting SD- and PKP-based signatures.

The full Hamming ball $\mathcal{H}_{\eta,4}$ has $k_4(\eta)$ elements and full reconstruction rank, but the universal upper bound rules out a degree-4 membership proof for this set. Our construction loses only two representatives while providing both degree-4 membership polynomials and full reconstruction rank. Hence, it attains the maximum cardinality among all admissible subsets of $\mathbb{F}_2^\eta$, whereas products of Hamming balls are smaller because of their additional block-wise restrictions.

4 Shorter Signature Schemes from Our Framework

This section instantiates the encoding framework of Section 3 in the VOLEitH paradigm, yielding shorter signature schemes from RSD and PKP. We first recall the VOLEitH protocol, and then describe two instantiations based on RSD and PKP. For each instantiation, we specify the polynomial constraints and compare the resulting extended witness size with previous modeling approaches.

4.1 The VOLEitH technique

Before giving our signature schemes, we recall the VOLEitH framework from [BBdSG+23b]. We refer to [BBdSG+23b] for a detailed description, and here we only recall those aspects needed in the constructions below.

VOLE correlations. Let λ be the security parameter. A VOLE correlation of length ℓ consists of random bits $\boldsymbol{u} = (u_1, \ldots, u_\ell) \in \mathbb{F}_2^\ell$ and random VOLE tags $\boldsymbol{v} = (v_1, \ldots, v_\ell) \in \mathbb{F}_{2^\lambda}^\ell$, together with a random global key $\Delta \in \mathbb{F}_{2^\lambda}$ and VOLE keys $\boldsymbol{q} = (q_1, \ldots, q_\ell) \in \mathbb{F}_{2^\lambda}^\ell$ such that

$$q_i = u_i\Delta + v_i, \qquad \forall i \in [\ell].$$

The vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ are known only to the prover, while Δ and $\boldsymbol{q}$ are given to the verifier. To derive authenticated shares for a witness $\boldsymbol{w} \in \mathbb{F}_2^\ell$, the prover sends $\boldsymbol{d} = \boldsymbol{w} - \boldsymbol{u}$. The verifier then updates its VOLE keys by computing $q_i' \coloneqq q_i + d_i\Delta$.

Polynomial constraints verification. Let $F_1, \ldots, F_t \in \mathbb{F}_2[X_1, \ldots, X_\ell]_{\le d}$ be a set of polynomials. Given the prover's authenticated witness representation $(\boldsymbol{w}, \boldsymbol{v})$ and the verifier's VOLE share $(\Delta, \boldsymbol{q}')$, the protocol allows the prover to prove knowledge of a witness $\boldsymbol{w} \in \mathbb{F}_2^\ell$ satisfying

$$F_j(\boldsymbol{w}) = 0, \qquad \forall j \in [t].$$

For each $j \in [t]$, let $\mathcal{M}_j$ be the set of exponents $\alpha = (\alpha_1, \ldots, \alpha_\ell)$ whose monomial $\prod_{i=1}^{\ell} X_i^{\alpha_i}$ appears in F_j with coefficient $c_j^{\langle\alpha\rangle}$. We write

$$F_j(X_1, \ldots, X_\ell) = \sum_{\alpha \in \mathcal{M}_j} c_j^{\langle\alpha\rangle} \prod_{i=1}^{\ell} X_i^{\alpha_i}.$$

The protocol handles these polynomial constraints by viewing each pair (w_i, v_i) as the coefficients of the degree-1 polynomial

$$p_i(X) = w_i X + v_i.$$

In this way, the condition $F_j(\boldsymbol{w}) = 0$ is equivalent to requiring the following polynomial to have a zero leading coefficient:

$$P_j(X) = \sum_{\alpha \in \mathcal{M}_j} c_j^{\langle \alpha \rangle} X^{d-(\alpha_1+\cdots+\alpha_\ell)} \prod_{i=1}^{\ell} p_i(X)^{\alpha_i}$$

To verify this, the verifier samples random scalars $\boldsymbol{\gamma} \in \mathbb{F}_{2^\lambda}^t$, and the prover sends a degree-$\leq d-1$ polynomial $\tilde{P}(X)$ claimed to equal $\sum_{j=1}^t \gamma_j P_j(X)$. Using its VOLE share $(\Delta, \boldsymbol{q}')$, the verifier checks $\sum_{j=1}^t \gamma_j P_j(\Delta) = \tilde{P}(\Delta)$. If the degree-$d$ coefficient of the batch is nonzero, the difference is a nonzero degree-$\leq d$ polynomial, so acceptance probability is at most $d \cdot 2^{-\lambda}$. For zero knowledge, the protocol additionally generates $\lambda(d-1)$ random VOLEs and uses them to mask the coefficients of $\tilde{P}(X)$ before revealing them.

VOLE-in-the-Head. The polynomial-checking procedure described above is naturally a designated-verifier protocol, since verification relies on the verifier's VOLE share $(\Delta, \boldsymbol{q}')$. The VOLEitH paradigm makes this check publicly verifiable by committing to locally generated VOLE correlations and opening all but one of them via an all-but-one vector commitment.

Signature sizes from VOLE-in-the-Head. The signature size of a VOLEitH-based scheme mainly comes from two parts. The first part is the VOLEitH overhead, including the seeds, commitments, and openings used to generate and verify the VOLE correlations, typically organized through a GGM tree and an all-but-one vector commitment [BBM+24]. The second part is the cost of proving the OWF-specific polynomial relation, which is governed by the authenticated witness length and by the degree of the polynomial constraints. We focus on optimizing the second part. For a witness $\boldsymbol{w}$ and degree bound d, define $L_{\text{ext}} := |\boldsymbol{w}| + \lambda(d-1)$. We call L_{ext} the extended witness size. For a fair comparison, we use the same VOLEitH parameters as the corresponding baseline whenever possible, so the GGM-tree and commitment overheads are kept the same.

4.2 Signatures from SD

In our SD-based signature, VOLEitH proves knowledge of an RSD solution. We compactly encode its elementary-vector blocks with representative sets and express both block membership and the syndrome equation as low-degree polynomial constraints. RSD is used only for this modeling: Theorem 3 converts an algorithm solving the selected RSD instances into one solving the associated SD instances, so the resulting signature is based on standard SD.

Witness representation. Let $w \mid n$ and set $\mu = n/w$. Given $\boldsymbol{H} \in \mathbb{F}_2^{(n-k)\times n}$ and $\boldsymbol{y} \in \mathbb{F}_2^{n-k}$, the RSD problem asks for

$$\boldsymbol{x} = \big(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_w\big), \qquad \boldsymbol{x}_i \in \mathcal{B}_\mu \text{ for every } i \in [w], \qquad \boldsymbol{H}\boldsymbol{x} = \boldsymbol{y}.$$

Let $S \subseteq \mathbb{F}_2^\eta$ be admissible under degree bound d, with $|S| = \mu$, and use the reconstruction and membership maps $\phi : S \to \mathcal{B}_\mu$ and $\varphi : \mathbb{F}_2^\eta \to \mathbb{F}_2^t$ defined in Section 3.2. Writing $\boldsymbol{x} = \big(\phi(\boldsymbol{s}_1), \ldots, \phi(\boldsymbol{s}_w)\big)$ represents the RSD witness by $(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_w) \in S^w$, containing $w\eta$ bits.

Polynomial constraints. We impose the following constraints on $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_w$.

- *Membership in S.* For every $i \in [w]$, impose $\varphi(\boldsymbol{s}_i) = \boldsymbol{0}$; equivalently, all $q_r(\boldsymbol{s}_i)$ vanish for $r \in [t]$. These equations ensure that every reconstructed block $\phi(\boldsymbol{s}_i)$ belongs to $\mathcal{B}_\mu$.
- *Syndrome equation.* The reconstructed vector must satisfy

$$\boldsymbol{H}\big(\phi(\boldsymbol{s}_1), \ldots, \phi(\boldsymbol{s}_w)\big) = \boldsymbol{y}.$$

Together, these degree-$\leq d$ constraints characterize the RSD relation exactly.

Signature scheme and security. The resulting scheme uses $\boldsymbol{x}$ as the secret key and $(\boldsymbol{H}, \boldsymbol{y})$ as the public key. The public key has size $\lambda + n - k$ bits when $\boldsymbol{H}$ is derived from a λ-bit seed and only $\boldsymbol{y}$ is stored explicitly. Concrete RSD parameters $(n, n-k, w)$ and modeling parameters (η, d) are presented in Section 4.3.

EUF-CMA security in the random oracle model follows VOLEitH and FAEST [BBdSG+23b, BBdSG+23a, BBB+25a]. The zero-knowledge simulator answers signing queries by programming the Fiat–Shamir random oracles; extraction from a valid forgery yields an RSD solution, which Theorem 3 converts into an algorithm for random SD instances with the stated loss in success probability. We defer the detailed security analysis to the full version of this paper.

4.3 Concrete Parameters for Signatures from SD

Modeling approaches. Table 4 gives the complete RSD and modeling parameters for our constructions and the previous approaches used in the comparison. For BBG+24 and SDitH, both the parameters and modeling are taken from the corresponding works; Chou25's modeling is evaluated using the SDitH parameters. OTX24 [OTX24] is omitted because its modeling matches that of BBG+24, while its concrete parameters rely directly on RSD. The block-related columns do not apply to Section 3.5 because its representative sets are not defined through a block decomposition.

Table 4. Complete RSD and modeling parameters used in the comparison

Category	Modeling	RSD Parameters		Block Structure			(η, d)
		$(n, n-k, w)$	μ	ℓ	Block lengths	Block weights	
I	BBG+24 [BBGK24]	$(6\,080, 701, 95)$	64	6	$(1, \ldots, 1)$	$\leq (1, \ldots, 1)$	$(6, 6)$
	SDitH [MBB+25]	$(10\,360, 432, 56)$	185	4	$(4, 4, 4, 3)$	$= (1, 1, 1, 1)$	$(11, 4)$
	Chou25 [Cho25]†	$(10\,360, 432, 56)$	185	2	$(5, 5)$	$\in \mathcal{R}$	$(9, 4)$
	Section 3.3	$(16\,940, 384, 44)$	385	2	$(2, 9)$	$\leq (1, 3)$	$(11, 4)$
	Section 3.5	$(16\,896, 392, 44)$	384	–	–	–	$(10, 4)$
III	BBG+24 [BBGK24]	–	–	–	–	–	–
	SDitH [MBB+25]	$(18\,396, 592, 73)$	252	4	$(4, 4, 4, 4)$	$= (1, 1, 1, 1)$	$(12, 4)$
	Chou25 [Cho25]†	$(18\,396, 592, 73)$	252	2	$(5, 6)$	$\in \mathcal{R}$	$(10, 4)$
	Section 3.3	$(24\,245, 576, 65)$	373	2	$(2, 9)$	$\leq (1, 3)$	$(11, 4)$
	Section 3.5	$(31\,900, 536, 58)$	550	–	–	–	$(11, 4)$
V	BBG+24 [BBGK24]	$(12\,160, 1\,405, 190)$	64	6	$(1, \ldots, 1)$	$\leq (1, \ldots, 1)$	$(6, 6)$
	SDitH [MBB+25]	$(19\,864, 800, 104)$	191	4	$(4, 4, 4, 3)$	$= (1, 1, 1, 1)$	$(11, 4)$
	Chou25 [Cho25]†	$(19\,864, 800, 104)$	191	2	$(5, 5)$	$\in \mathcal{R}$	$(9, 4)$
	Section 3.3	$(25\,392, 752, 92)$	276	2	$(2, 8)$	$\leq (1, 3)$	$(10, 4)$
	Section 3.5	$(26\,255, 736, 89)$	295	–	–	–	$(12, 3)$

† Parameters from Chou25 [Cho25, Sec. 5.2].
$\mathcal{R} := \{(1, 3), (2, 2), (3, 1)\}$ denotes the possible pairs of block weights.

Parameter selection. We consider NIST security Categories I, III, and V, corresponding to $\kappa \in \{143, 207, 272\}$. Following the methodology of SDitH, we select RSD parameters admitting a provable reduction to unstructured binary SD and estimate the security of the associated SD instances under the same cost model. The parameters are required to satisfy the following conditions.

- We consider unstructured binary SD instances with parameters (n, k_{SD}, w), where k_{SD} is chosen so that the instance lies close to the Gilbert–Varshamov bound. Its estimated security must be at least $\kappa' = \kappa + \log_2\left(\binom{n}{w}/(n/w)^w\right) + 1/(\theta \ln 2)$.
- For each SD instance, we choose an RSD instance with parameters (n, k, w), where k is the largest value satisfying $k_{\mathrm{SD}} \leq k \leq n - w\log_2(n/w) - \log_2 \theta$. The reduction in Theorem 3 therefore bases the security of the resulting RSD instance on binary SD.
- The concrete SD parameters are evaluated using the Cryptographic Estimators[5] [EVZB24] under the logarithmic memory-access cost model. As in SDitH, the memory limit is 2^{143} for Category I and 2^{160} for Categories III and V.

For every selected RSD instance, let $\mu = n/w$. We use the upper bound to identify the best representation length that could be achieved within our encoding framework.

1. Enumerate valid RSD parameters $(n, n-k, w)$ and set $\mu = n/w$.
2. For every candidate degree d, find the smallest η satisfying $\mu \leq U_d(\eta)$.

[5] https://github.com/Crypto-TII/CryptographicEstimators

3. For every feasible candidate, compute $L_{\text{ext}} = w\eta + \lambda(d-1)$.
4. Minimize this quantity over the enumerated RSD parameters and degrees. The resulting values of (η, d) identify the representative-set constructions required for the selected parameters.

Throughout this subsection, η denotes the number of bits actually stored for each representative, including any coordinate-elimination optimization provided by the corresponding construction.

Selection using the upper bound. We determine (η, d) by minimizing $L_{\text{ext}} = w\eta + \lambda(d-1)$, subject to $\mu \le U_d(\eta)$. Once the cost formula of another proof system is known, the same procedure can guide parameter selection by replacing L_{ext} with the corresponding objective. For example, a smaller factor multiplying $d-1$ may favor a larger d, which can allow a smaller η and potentially a shorter signature. The complete search procedure is given above.

Because $U_d(\eta)$ is only an upper bound, this calculation initially gives a lower bound on the achievable L_{ext}. It does not by itself prove that the required S exists. For the parameter pairs selected below, however, Section 3.5 provides explicit admissible sets of size at least μ.

If a construction yields $|S| > \mu$, we may use any subset $S' \subseteq S$ of size μ. Deleting rows preserves full row rank, and adding one degree-$\le d$ polynomial to the existing membership proof excludes $S \setminus S'$, so S' remains admissible.

The upper bound shows that representative sets of sizes 384, 550, and 295 require $\eta \ge 10$, 11, and 12 at degrees $d = 4$, 4, and 3, respectively. The constructions in Section 3.5 attain these sizes.

Comparison of extended witness sizes. Table 5 reports $|\boldsymbol{w}|$ and L_{ext} for the selected constructions and previous approaches. The values are given in bytes and rounded upward. The optimal representative sets reduce the extended witness size relative to Chou25 [Cho25, Sec. 5.2] by 8, 12, and 15 bytes and relative to the Hamming-ball construction in Section 3.3 by 6, 10, and 13 bytes for Categories I, III, and V, respectively. Relative to SDitH [MBB+25], the reductions are 17.6%, 16.5%, and 17.2%; relative to BBG+24 [BBGK24], they are 32.2% and 34.7% for NIST-I and NIST-V, respectively.

Table 5. Constraint degree, witness size, and extended witness size for RSD

	NIST-I			NIST-III			NIST-V		
Modeling	d	$\lvert\boldsymbol{w}\rvert$	L_{ext}	d	$\lvert\boldsymbol{w}\rvert$	L_{ext}	d	$\lvert\boldsymbol{w}\rvert$	L_{ext}
BBG+24 [BBGK24]	6	72	152	6	–	–	6	143	303
SDitH [MBB+25]	4	77	125	4	110	182	4	143	239
Chou25 [Cho25, Sec. 5.2]	4	63	111	4	92	164	4	117	213
Section 3.3	4	61	109	4	90	162	4	115	211
Section 3.5	4	55	**103**	4	80	**152**	3	134	**198**

Remark 7 (How the modeling yields shorter signatures). At first glance, the improvements shown above may appear to arise mainly from choosing different RSD parameters. In fact, the signature-size advantage of our modeling is twofold:

- **Under the same RSD parameters.** Using the same RSD parameters as SDitH, we use S_9 with $|S_9| = 254$, so $\eta = 9$ suffices at all three security levels, compared with $\eta = 11$, 12, and 11, respectively, in SDitH. This gives extended-witness sizes of 111, 155, and 213 bytes, corresponding to reductions of 11.2%, 14.8%, and 10.9%, respectively. At Category I, for example, this change alone gives a 3,551-byte short signature, a reduction of 4.2%. Under the same parameters, our η matches Chou25 at Categories I and V and is smaller by one bit at Category III.
- **Through joint parameter selection.** For the parameter pairs used here, our constructions attain the upper bound on $|S|$ for the given representation length and constraint degree. This greater capacity enlarges the feasible RSD parameter space, enabling the joint search to select, at the same security level, parameters with a larger n and a smaller w, thereby further reducing the extended witness.

Thus, our signature-size gains arise both directly from shorter representations under the same RSD parameters and additionally from the broader parameter choices enabled by those representations.

4.4 Signatures from PKP

Our PKP-based construction uses VOLEitH to prove knowledge of a permutation represented by its permutation matrix, while compactly encoding its elementary-vector columns and expressing the PKP relation through low-degree polynomial constraints.

Witness representation. Let $q = 2^p$. Given $\boldsymbol{H} \in \mathbb{F}_q^{m \times n}$ and $\boldsymbol{x} \in \mathbb{F}_q^n$, the PKP problem asks for a permutation matrix $\boldsymbol{P} \in \mathbb{F}_2^{n \times n}$ satisfying

$$\boldsymbol{HPx} = \boldsymbol{0}.$$

Let $S \subseteq \mathbb{F}_2^\eta$ be admissible under degree bound d, with $|S| = n$, and use the reconstruction and membership maps $\phi : S \to \mathcal{B}_n$ and $\varphi : \mathbb{F}_2^\eta \to \mathbb{F}_2^t$ defined in Section 3.2. Writing $\boldsymbol{P} = \big(\phi(\boldsymbol{s}_1) \mid \cdots \mid \phi(\boldsymbol{s}_n)\big)$ represents the PKP witness by $(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n) \in S^n$, containing $n\eta$ bits.

Polynomial constraints. We impose the following constraints on $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n$.

- *Membership in S.* For every $i \in [n]$, impose $\varphi(\boldsymbol{s}_i) = \boldsymbol{0}$; equivalently, all $q_r(\boldsymbol{s}_i)$ vanish for $r \in [t]$. Consequently, every reconstructed column $\phi(\boldsymbol{s}_i)$ belongs to $\mathcal{B}_n$, as required.
- *PKP equation.* The reconstructed matrix must satisfy

  $$\boldsymbol{H}\big(\phi(\boldsymbol{s}_1) \mid \cdots \mid \phi(\boldsymbol{s}_n)\big)\boldsymbol{x} = \boldsymbol{0}.$$

  Since $q = 2^p$, every equation over $\mathbb{F}_q$ is expanded into p binary polynomial constraints over a fixed $\mathbb{F}_2$-basis of $\mathbb{F}_q$. This binary expansion preserves the degree of the original equation.
- *Permutation condition.* For every $j \in [n]$, impose

  $$\sum_{i=1}^{n} \big(\phi(\boldsymbol{s}_i)\big)_j = 1.$$

  Membership in S gives $\sum_{j=1}^{n} \mathsf{wt}(\boldsymbol{P}_{j,*}) = n$. The row-sum equations give $\mathsf{wt}(\boldsymbol{P}_{j,*}) \equiv 1 \pmod 2$, and hence $\mathsf{wt}(\boldsymbol{P}_{j,*}) \geq 1$, for each of the n rows. These lower bounds already account for all n ones, so every row has weight exactly one and $\boldsymbol{P}$ is a permutation matrix.

Together, these degree-$\leq d$ constraints characterize the PKP relation exactly.

Signature scheme and security. The resulting scheme uses $\boldsymbol{P}$ as the secret key and $(\boldsymbol{H}, \boldsymbol{x})$ as the public key. The public key has size $\lambda + n \log_2 q$ bits when $\boldsymbol{H}$ is derived from a λ-bit seed and only $\boldsymbol{x}$ is stored explicitly. Concrete PKP parameters (q, n, m) and modeling parameters (η, d) are presented in Section 4.5. EUF-CMA security in the random oracle model follows VOLEitH and FAEST [BBdSG+23b, BBdSG+23a, BBB+25a]. The zero-knowledge simulator answers signing queries by programming the Fiat–Shamir random oracles; extraction yields representatives satisfying the constraints, which form a permutation matrix $\boldsymbol{P}$ satisfying $\boldsymbol{HPx} = \boldsymbol{0}$. We defer the detailed security analysis to the full version of this paper.

4.5 Concrete Parameters for Signatures from PKP

Parameter selection. We consider NIST security Categories I, III, and V, with target values $\kappa \in \{143, 207, 272\}$. We focus on the classical mono-dimensional PKP defined in Definition 4. Our PKP parameters are obtained by adapting those of PERK v2.2 [ABB+25a] to the modeling parameters required by our representative sets, while preserving the target security levels. For each candidate PKP instance (q, n, m), we choose m near the balanced regime $n!/q^m \approx 1$ and use the Cryptographic Estimators [EVZB24] to evaluate the concrete costs of known PKP attacks, including KMP [KMRP19] and SBC [SBC24]. Table 6 lists the resulting parameters.

Modeling approaches. Table 6 gives the complete PKP and modeling parameters for our constructions and the previous approaches used in the comparison; if $\mu > n$, only n representatives are used. For the previous constructions, the parameters and modeling are taken from the corresponding works; Chou25 reports parameters only for Categories I and V.

Table 6. Complete PKP and modeling parameters used in the comparison

Category	Modeling	PKP Parameters		Block Structure			(η, d)
		(q, n, m)	μ	ℓ	Block lengths	Block weights	
I	BBG+24 [BBGK24]	$(2048, 64, 27)$	64	6	$(1, \ldots, 1)$	$\leq (1, \ldots, 1)$	$(6, 6)$
	PERK v2.2 [ABB+25a]	$(2048, 64, 27)$	64	3	$(4, 4, 4)$	$= (1, 1, 1)$	$(9, 3)$
	Chou25 [Cho25]†	$(2048, 64, 27)$	70	2	$(4, 5)$	$\in \{(1, 2), (2, 1)\}$	$(8, 3)$
	Section 3.3	$(2048, 62, 26)$	64	2	$(3, 5)$	$\leq (1, 2)$	$(8, 3)$
	Section 3.5	$(2048, 62, 26)$	62	–	–	–	$(7, 3)$
III	BBG+24 [BBGK24]	–	–	–	–	–	–
	PERK v2.2 [ABB+25a]	$(2048, 92, 43)$	128	4	$(4, 4, 4, 2)$	$= (1, 1, 1, 1)$	$(10, 4)$
	Chou25 [Cho25]†	–	–	–	–	–	–
	Section 3.3	$(4096, 84, 35)$	88	2	$(3, 6)$	$\leq (1, 2)$	$(9, 3)$
	Section 3.5	$(2048, 90, 42)$	90	–	–	–	$(8, 3)$
V	BBG+24 [BBGK24]	$(4096, 109, 49)$	128	7	$(1, \ldots, 1)$	$\leq (1, \ldots, 1)$	$(7, 7)$
	PERK v2.2 [ABB+25a]	$(2048, 118, 59)$	128	4	$(4, 4, 4, 2)$	$= (1, 1, 1, 1)$	$(10, 4)$
	Chou25 [Cho25]†	$(4096, 109, 49)$	135	2	$(5, 6)$	$\in \{(1, 2), (2, 1)\}$	$(10, 3)$
	Section 3.3	$(4096, 109, 49)$	110	2	$(4, 6)$	$\leq (1, 2)$	$(10, 3)$
	Section 3.5	$(2048, 118, 59)$	128	–	–	–	$(9, 3)$

† Parameters from Chou25 [Cho25, Sec. 5.3].

Comparison of extended witness sizes. Table 7 compares complete concrete instantiations because the modeling approaches use different PKP instances. Both $|\boldsymbol{w}|$ and L_{ext} are given in bytes and rounded upward. The optimal representative sets reduce the extended witness size relative to Chou25 [Cho25, Sec. 5.3] by 9 and 4 bytes for Categories I and V and relative to the Hamming-ball construction in Section 3.3 by 7, 5, and 4 bytes for Categories I, III, and V. Relative to PERK v2.2 [ABB+25a], the reductions are 16.3%, 26.2%, and 19.3%; relative to BBG+24 [BBGK24], they are 32.0% and 31.6% for NIST-I and NIST-V, respectively.

Table 7. Constraint degree, witness size, and extended witness size for PKP

	NIST-I			NIST-III			NIST-V		
Modeling	d	$\lvert\boldsymbol{w}\rvert$	L_{ext}	d	$\lvert\boldsymbol{w}\rvert$	L_{ext}	d	$\lvert\boldsymbol{w}\rvert$	L_{ext}
BBG+24 [BBGK24]	6	48	128	–	–	–	7	96	288
PERK v2.2 [ABB+25a]	3	72	104	4	115	187	4	148	244
Chou25 [Cho25, Sec. 5.3]	3	64	96	–	–	–	3	137	201
Section 3.3	3	62	94	3	95	143	3	137	201
Section 3.5	3	55	**87**	3	90	**138**	3	133	**197**

4.6 Optimizations

Recall that in VOLEitH, each witness bit w_i and its VOLE tag $v_i \in \mathbb{F}_{2^\lambda}$ define the degree-one polynomial $p_i(X) = w_i X + v_i$. Evaluating the constraint polynomials $F_1, \ldots, F_t$ requires products of these degree-one polynomials. Binary coefficients $w_i \in \mathbb{F}_2$ can be handled inexpensively, whereas products of the tags $v_i \in \mathbb{F}_{2^\lambda}$ require multiplications in $\mathbb{F}_{2^\lambda}$ and dominate the evaluation cost. We therefore use the number of multiplications in $\mathbb{F}_{2^\lambda}$ as our cost metric. We describe the optimizations for the RSD constraints and report concrete multiplication counts for Category I. Similar optimizations can also be applied to PKP.

Membership constraints. We optimize the block-wise Hamming-weight conditions $\rho_{j,T}$ and σ_j defined in Equations (4) and (5), respectively. Since $\deg(\sigma_j) \leq r_j$, the dominant cost comes from the degree-$(r_j + 1)$ polynomials $\rho_{j,T}$. Instead of evaluating every $\rho_{j,T}$ independently, we compute each tag product once within its block and reuse it across all constraints containing that product.

For the Category-I Hamming-ball construction, $(\eta_1, r_1) = (2, 1)$ and $(\eta_2, r_2) = (9, 3)$, reducing the membership cost from 1387 to 247 large-field multiplications per representative vector, a factor of approximately 5.62. The exact operation count is derived below. This optimization does not apply to the general representative sets in Section 3.5 because their membership polynomials do not have this block-wise Hamming-weight form.

Syndrome-equation constraints. The syndrome equation is

$$\boldsymbol{H}\big(\phi(\boldsymbol{s}_1), \ldots, \phi(\boldsymbol{s}_w)\big) = \boldsymbol{y},$$

where $\boldsymbol{H} = [\boldsymbol{h}_1 \mid \cdots \mid \boldsymbol{h}_n] \in \mathbb{F}_2^{(n-k)\times n}$. Using the row-packing technique introduced in SDitH [MBB+25], we pack every λ rows over $\mathbb{F}_2$ into one row over $\mathbb{F}_{2^\lambda}$ and take a random linear combination of the resulting $\nu = \lceil (n-k)/\lambda \rceil$ rows. Writing the batched public coefficients and syndrome as $\tilde{h}_1, \ldots, \tilde{h}_n$ and $\tilde{y}$, respectively, gives

$$\sum_{i=1}^{w} \sum_{q=1}^{\mu} \tilde{h}_{(i-1)\mu+q} f_q(\boldsymbol{s}_i) = \tilde{y}, \tag{6}$$

where $f_1, \ldots, f_\mu$ are the reconstruction polynomials. Evaluating all μ coordinate polynomials independently and then incorporating them into Equation (6) requires $\mu d(d+1)/2$ large-field multiplications per representative vector.

This computation can be fused with the public linear relation for every admissible representative set. Let $K = k_d(\eta)$, and write each reconstruction polynomial in the degree-$\leq d$ monomial basis as $f_q = \sum_{h=1}^{K} \alpha_q^{\langle h \rangle} g^{\langle h \rangle}$. Define $\widehat{h}_{i,h} := \sum_{q=1}^{\mu} \tilde{h}_{(i-1)\mu+q} \alpha_q^{\langle h \rangle} \in \mathbb{F}_{2^\lambda}$. Substituting the monomial expansion into Equation (6) gives

$$\sum_{i=1}^{w} \sum_{q=1}^{\mu} \tilde{h}_{(i-1)\mu+q} f_q(\boldsymbol{s}_i) = \sum_{i=1}^{w} \sum_{h=1}^{K} \widehat{h}_{i,h} g^{\langle h \rangle}(\boldsymbol{s}_i) = \tilde{y}, \tag{7}$$

The prover therefore evaluates the monomial basis only once and combines these evaluations using the transformed public coefficients $\widehat{h}_{i,h}$. This transformation only uses the linearity of the public relation and therefore applies to both product-structured and general representative sets.

Product-structured representative sets. For products of Hamming balls or spheres, the fused relation can be evaluated recursively using the product structure and only the supporting monomials. The complete recursion and its operation count are derived below. For the Category-I Hamming-ball construction, this reduces the syndrome-equation cost from 3900 to 395 multiplications. Including the membership constraints, the total cost decreases from 5287 to 642, a factor of approximately 8.23.

General representative sets. The general constructions in Section 3.5 do not admit the preceding block recursion, but Equation (7) still applies. For the Category-I construction with $(\eta, d) = (10, 4)$ and $|S| = 384$, the optimized evaluation requires 760 large-field multiplications per representative vector, compared with 3840 for evaluating the reconstruction polynomials separately, a reduction factor of approximately 5.05. The detailed evaluation is given below. The membership map φ reuses the same cached monomials and requires no additional large-field multiplications beyond this total.

Detailed Constraint-Evaluation Costs Membership-constraint evaluation. For the j-th block, direct evaluation computes every constraint $\rho_{j,T}$ independently. Let $V_B := \prod_{i \in B} v_i$ for $B \subseteq T$ and $|B| \geq 2$. Direct evaluation recomputes V_B in every constraint containing B; instead, we compute it once within its block and reuse it. For ℓ blocks, the direct and optimized costs are

$$C_{\text{mem}}^{\text{dir}} = \sum_{j=1}^{\ell} \binom{\eta_j}{r_j + 1} \big(2^{r_j+1} - r_j - 2\big), \quad C_{\text{mem}}^{\text{opt}} = \sum_{j=1}^{\ell} \sum_{t=2}^{r_j+1} \binom{\eta_j}{t}.$$

For the Category-I Hamming-ball construction, $(\eta_1, r_1) = (2, 1)$ and $(\eta_2, r_2) = (9, 3)$. Hence $C_{\text{mem}}^{\text{dir}} = 1387$, whereas $C_{\text{mem}}^{\text{opt}} = 247$.

Product-structured syndrome evaluation. For $S = S_1 \times \cdots \times S_\ell$ constructed from Hamming balls or spheres, Equation (8) permits us to restrict to the supporting monomials and exploit their product

structure recursively. Define $D_j = |S_j|$ and $L_j = \sum_{t=1}^{r_j} \binom{\eta_j}{t}$. For an evaluation order $(j_1, \ldots, j_\ell)$, let $R_m = \sum_{q<m} r_{j_q}$. Successively contracting the public coefficient tensor along these blocks requires

$$C_{\mathrm{syn}}(j_1, \ldots, j_\ell) = \sum_{m=1}^{\ell} \left(\prod_{q>m} D_{j_q} \right) (R_m + 1) L_{j_m}$$

large-field multiplications per representative vector; we choose the block order minimizing this expression. For the Category-I Hamming-ball construction, $(D_1, L_1) = (3, 2)$ and $(D_2, L_2) = (130, 129)$. Processing the weight-three block first gives $C_{\mathrm{syn}}^{\mathrm{opt}} = 3 \cdot 129 + 4 \cdot 2 = 395$, whereas direct evaluation requires 3900 multiplications. Including membership, the total cost decreases from $1387 + 3900 = 5287$ to $247 + 395 = 642$.

General representative sets. For general representative sets, we first compute and cache all nonlinear monomials in the authenticated variables. We then transform the public reconstruction coefficients and evaluate the fused relation using the Taylor/zeta transform and Horner evaluation. For the Category-I construction with $(\eta, d) = (10, 4)$, $|S| = 384$, and $K = k_4(10) = 386$, the two parts of the optimized computation cost

$$C_{\mathrm{mon}} = \sum_{t=2}^{4} \binom{10}{t} = 375, \qquad C_{\mathrm{eval}} = K - 1 = 385.$$

The total cost is therefore 760 large-field multiplications per representative vector, compared with $|S|d(d+1)/2 = 3840$ for evaluating all reconstruction polynomials separately. The two membership polynomials reuse the cached monomials and require no additional large-field multiplications.

5 Implementation and Benchmarks

This section presents our implementation and benchmarks. We compare our signature schemes with the post-quantum signatures selected for standardization by NIST and the third-round candidates based on MPC/VOLEitH techniques in NIST's Additional Digital Signature Standardization Process[6]. Ring-signature performance is considered separately in Section 6.

5.1 Implementation Methodology

We implement our proposed signature schemes based on the AVX2-optimized implementation of FAEST v2 [BBB+25a][7]. The implementation is available online.[8] The improved all-but-one vector commitment from [BBB+25b] is not adopted. Although the latter improves computational efficiency, its available security analysis relies on a pseudorandomness assumption for the OWF underlying the signature scheme. Our SD and PKP instantiations assume only one-wayness and do not provide such pseudorandomness; consequently, we retain the commitment of Baum *et al.* [BBM+24].

The SD-based scheme adopts the proof-system parameters of SDitH [MBB+25], whereas the PKP-based scheme adopts those of PERK [ABB+25a]. The two complete sets of proof-system parameters are reported in Appendix C.1. These choices are intended to provide direct comparisons with SDitH and PERK, respectively, while reducing the influence of proof-system parameter choices as much as possible; dedicated parameter optimization may yield better size–time trade-offs for our schemes than those reported here.

For a degree-d relation, the VOLEitH signature size is

$$|\mathsf{sig}| = \tau\big(|\mathsf{wit}| + (d-1)\lambda + \lambda + 16\big) + T_{\mathsf{open}}\lambda + 2\tau\lambda + 32 + \lambda + 128.$$

All terms in this expression are measured in bits. By incorporating the optimizations from FAEST v2 [BBB+25a], the above formula saves $3\lambda - 128$ bits compared with the corresponding signature-size formula of SDitH [MBB+25]. This corresponds to 32, 56, and 80 bytes at NIST Levels I, III, and V,

[6] https://csrc.nist.gov/Projects/pqc-dig-sig/round-3-additional-signatures
[7] https://github.com/faest-sign/faest-arch-opt
[8] https://github.com/gby2023/Shorter-VOLEitH-Signatures

Table 8. Comparisons of sizes and runtimes for our schemes and existing post-quantum signatures at NIST security level I

Scheme	Sizes (bytes)			Runtimes (ms)		Assumption
	|pk|	|sig|	|sig| + |pk|	sign	verify	
ML-DSA-44 [LDK+22]	1312	2420	3732	0.071	0.026	Lattice
FN-DSA-512 [PFH+22]	897	666	1563	0.225	0.049	Lattice
SLH-DSA-SHAKE-128s [HBD+22]	32	7856	7888	320.440	0.400	Hash
SLH-DSA-SHAKE-128f [HBD+22]	32	17088	17120	15.517	1.058	Hash
SLH-DSA-SHA2-128s [HBD+22]	32	7856	7888	148.477	0.224	Hash
SLH-DSA-SHA2-128f [HBD+22]	32	17088	17120	7.249	0.591	Hash
FAEST [BBB+25a] (short)	32	4506	4538	4.277	3.298	AES
FAEST [BBB+25a] (fast)	32	5924	5956	0.568	0.463	AES
MQOM-gf2 [BBFR25] (short)	52	2820	2872	3.789	3.655	MQ over $\mathbb{F}_2$
MQOM-gf2 [BBFR25] (fast)	52	3144	3196	1.685	1.579	MQ over $\mathbb{F}_2$
MQOM-gf16 [BBFR25] (short)	60	2916	2976	2.063	1.926	MQ over $\mathbb{F}_{16}$
MQOM-gf16 [BBFR25] (fast)	60	3280	3340	0.875	0.694	MQ over $\mathbb{F}_{16}$
PERK [ABB+25a] (short)	104	3473	3577	9.469	8.240	PKP
PERK [ABB+25a] (fast)	104	4312	4416	2.352	1.695	PKP
This work, PKP (short)	**102**	**3239**	**3341**	**6.161**	**5.804**	PKP
This work, PKP (fast)	**102**	**3988**	**4090**	**1.493**	**1.254**	PKP
SDitH-gf2 [MBB+25] (short)	70	3705	3775	9.809	7.950	SD over $\mathbb{F}_2$
SDitH-gf2 [MBB+25] (fast)	70	4484	4554	2.973	2.590	SD over $\mathbb{F}_2$
This work, SD (short)	**65**	**3431**	**3496**	**5.949**	**5.699**	SD over $\mathbb{F}_2$
This work, SD (fast)	**65**	**4100**	**4165**	**1.480**	**1.325**	SD over $\mathbb{F}_2$

respectively. We refer to [BBB+25a] for the definitions of the proof-system parameters and the details of these optimizations.

We benchmark all schemes locally on a single core of a 3.8 GHz AMD Ryzen 9 3900X machine running Ubuntu 22.04 with 64 GB of memory. All timings are arithmetic means over 1000 runs per parameter set and are reported in milliseconds.

5.2 Comparisons with Other Post-Quantum Signatures

Table 8 compares the sizes and runtimes of our SD- and PKP-based schemes with representative post-quantum signatures at NIST security level I.

As shown in Table 8, the combined size of a public key and a signature for our SD-based short variant is 3,496 bytes. This is smaller than the 3,732 bytes of ML-DSA and less than half of the 7,888 bytes of the short variants of SLH-DSA. Signing and verification take approximately 6 ms in the short variant and approximately 1.5 ms in the fast variant. Compared with SLH-DSA-SHA2, our signing is approximately 25.0× faster for the short variants and 4.9× faster for the fast variants. SLH-DSA-SHA2 retains faster verification, by approximately 25.4× and 2.2×, respectively.

For PKP, the combined public-key and signature sizes are 3,341 bytes for the short variant and 4,090 bytes for the fast variant. Compared with our SD-based scheme, these represent reductions of approximately 4.4% and 1.8%, respectively.

Table 9 compares our SD-based signature in more detail with three NIST third-round candidates, namely FAEST [BBB+25a], SDitH [MBB+25], and MQOM [BBFR25]. All three are based on MPCitH or VOLEitH techniques. Comparisons with PERK [ABB+25a] and Chou25 [Cho25] are summarized separately below.

Comparison with SDitH. SDitH also uses VOLEitH and relies on syndrome decoding, so this gives the closest comparison for our SD-based scheme. Our signatures are smaller at all three security levels and for both variants, with reductions ranging from 6.7% to 8.6%. Our public keys are 6.1%–7.1% smaller. The concrete reductions in Table 9 additionally include the shorter witness obtained

Table 9. Comparisons of our SD-based scheme with parameter sets from FAEST v2 [BBB+25a], SDitH v2 [MBB+25], and MQOM v2 [BBFR25] at NIST security levels I, III, and V

Scheme	Sizes (bytes)			Runtimes (ms)		Assumption
	|pk|	|sig|	|sig| + |pk|	sign	verify	
Ours-L1-short	**65**	**3431**	**3496**	**5.949**	**5.699**	SD over $\mathbb{F}_2$
Ours-L1-fast	**65**	**4100**	**4165**	**1.480**	**1.325**	SD over $\mathbb{F}_2$
Ours-L3-short	**91**	**7428**	**7519**	**22.395**	**21.296**	SD over $\mathbb{F}_2$
Ours-L3-fast	**91**	**9140**	**9231**	**3.654**	**3.205**	SD over $\mathbb{F}_2$
Ours-L5-short	**124**	**13180**	**13304**	**34.932**	**34.816**	SD over $\mathbb{F}_2$
Ours-L5-fast	**124**	**16148**	**16272**	**7.226**	**7.178**	SD over $\mathbb{F}_2$
SDitH-L1-gf2-short	70	3705	3775	9.809	7.950	SD over $\mathbb{F}_2$
SDitH-L1-gf2-fast	70	4484	4554	2.973	2.590	SD over $\mathbb{F}_2$
SDitH-L3-gf2-short	98	7964	8062	44.183	39.679	SD over $\mathbb{F}_2$
SDitH-L3-gf2-fast	98	9916	10014	7.719	6.803	SD over $\mathbb{F}_2$
SDitH-L5-gf2-short	132	14121	14253	62.886	57.097	SD over $\mathbb{F}_2$
SDitH-L5-gf2-fast	132	17540	17672	11.210	10.065	SD over $\mathbb{F}_2$
FAEST-128-short	32	4506	4538	4.277	3.298	AES
FAEST-128-fast	32	5924	5956	0.568	0.463	AES
FAEST-192-short	48	11260	11308	18.194	14.349	AES
FAEST-192-fast	48	14948	14996	2.610	2.352	AES
FAEST-256-short	48	20696	20744	23.252	23.070	AES
FAEST-256-fast	48	26548	26596	3.329	3.154	AES
FAEST-EM-128-short	32	3906	3938	3.117	2.480	EM-AES
FAEST-EM-128-fast	32	5060	5092	0.459	0.364	EM-AES
FAEST-EM-192-short	48	9340	9388	12.066	11.357	EM-AES
FAEST-EM-192-fast	48	12380	12428	1.841	1.642	EM-AES
FAEST-EM-256-short	64	17984	18048	18.749	17.995	EM-AES
FAEST-EM-256-fast	64	23476	23540	2.751	2.602	EM-AES
MQOM-L1-gf2-short	52	2820	2872	3.789	3.655	MQ over $\mathbb{F}_2$
MQOM-L1-gf2-fast	52	3144	3196	1.685	1.579	MQ over $\mathbb{F}_2$
MQOM-L3-gf2-short	78	6280	6358	16.484	15.690	MQ over $\mathbb{F}_2$
MQOM-L3-gf2-fast	78	7414	7492	8.085	8.029	MQ over $\mathbb{F}_2$
MQOM-L5-gf2-short	104	11564	11668	35.533	35.112	MQ over $\mathbb{F}_2$
MQOM-L5-gf2-fast	104	13124	13228	19.856	19.998	MQ over $\mathbb{F}_2$
MQOM-L1-gf16-short	60	2916	2976	2.063	1.926	MQ over $\mathbb{F}_{16}$
MQOM-L1-gf16-fast	60	3280	3340	0.875	0.694	MQ over $\mathbb{F}_{16}$
MQOM-L3-gf16-short	90	6496	6586	9.873	8.718	MQ over $\mathbb{F}_{16}$
MQOM-L3-gf16-fast	90	7738	7828	3.119	2.974	MQ over $\mathbb{F}_{16}$
MQOM-L5-gf16-short	122	12014	12136	15.606	14.965	MQ over $\mathbb{F}_{16}$
MQOM-L5-gf16-fast	122	13772	13894	6.941	6.563	MQ over $\mathbb{F}_{16}$

from our elementary-vector modeling. For all parameter sets, our signing is 1.55×–2.11× faster, while verification is 1.40×–2.12× faster.

As discussed in Section 4.3, our modeling provides a direct size reduction under fixed RSD parameters and also enables the joint parameter choices used here. At Categories I, III, and V, the selected code lengths increase from 10,360, 18,396, and 19,864 in SDitH to 16,896, 31,900, and 26,255, respectively, corresponding to factors of 1.63×, 1.73×, and 1.32×. Such larger instances would normally increase the evaluation cost. However, the optimizations described in Section 4 combine the reconstruction polynomials before evaluation and reuse common intermediate results, substantially reducing the cost of evaluating the RSD constraints. This optimization substantially offsets the additional cost and helps explain the faster performance of our implementation relative to SDitH.

Comparison with FAEST. The FAEST v2 parameter sets instantiate their underlying OWFs using either AES or EM-AES [BBB+25a]. Across the three security levels and both variants, our signatures are 23.9%–39.2% smaller than those of FAEST and 12.2%–31.2% smaller than those of FAEST-EM. FAEST and FAEST-EM are faster in every case: FAEST signing and verification are approximately 1.23×–2.61× and 1.36×–2.86× faster than ours, respectively; for FAEST-EM, the corresponding ranges are 1.86×–3.23× and 1.88×–3.64×.

Comparison with MQOM. MQOM uses TCitH [FR25], a variant of MPCitH, and its security relies on the multivariate quadratic problem. Our signatures are larger than those of MQOM. Compared with MQOM over $\mathbb{F}_2$, MQOM is faster for the L1 and L3 short variants, whereas our fast variants are faster at all three security levels. MQOM over $\mathbb{F}_{16}$ is faster for all six parameter sets.

Comparison with PERK and Chou25. Detailed comparisons with PERK and Chou25 are reported in Appendix C. Compared with PERK v2.2 [ABB+25a], our PKP signatures are 6.7%–12.5% smaller, with signing and verification speedups of 1.54×–2.28× and 1.35×–2.37×, respectively, for all six parameter sets.

Chou25 [Cho25] reports signature sizes but does not report implementation results. For SD, our signatures are 2.9%–3.8% smaller across the reported security levels and variants. For PKP, our short-variant signature is essentially the same size as Chou25's at Level I and approximately 1.4% smaller at Level V.

6 Application to Ring Signatures

Ring signatures allow one member of a set of public keys to sign without revealing which secret key was used. This section presents our SD-based ring-signature instantiation, its security analysis, and its concrete performance.

6.1 Compact Selectors for Ring Signatures

As shown in Section 3.2, our framework can reconstruct arbitrary target vectors. Assigning representatives to all weight-t binary vectors of length N yields a compact t-out-of-N proof; the case $t = 1$ gives a one-out-of-N selector for the generic ring-signature construction of Feneuil and Rivain [FR25, Section 6.6]. Let $\boldsymbol{H} \leftarrow \mathbb{F}_2^{(n-k)\times n}$ be sampled uniformly once as a global public parameter, and let $R = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N)$ be an ordered ring of distinct public keys generated with this common matrix. We use the RSD representative set S and the maps ϕ and φ defined in Section 4.2. For the selector, let S_N be a representative set with reconstruction map $\phi_N : S_N \to \mathcal{B}_N$ and membership map φ_N. Writing $\phi_N(\boldsymbol{s}_0) = (\phi_{N,1}(\boldsymbol{s}_0), \ldots, \phi_{N,N}(\boldsymbol{s}_0))$, the witness consists of a selector representative $\boldsymbol{s}_0$ and RSD representatives $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_w$, subject to

$$\varphi_N(\boldsymbol{s}_0) = \boldsymbol{0}, \qquad \varphi(\boldsymbol{s}_j) = \boldsymbol{0} \quad \text{for every } j \in [w],$$
$$\boldsymbol{H}\big(\phi(\boldsymbol{s}_1), \ldots, \phi(\boldsymbol{s}_w)\big) = \sum_{i=1}^{N} \phi_{N,i}(\boldsymbol{s}_0)\boldsymbol{y}_i.$$

If $\phi_N(\boldsymbol{s}_0) = \boldsymbol{e}_\ell$, the right-hand side is $\boldsymbol{y}_\ell$; hence the relation proves knowledge of an RSD solution for one ring member while keeping ℓ hidden. The Fiat–Shamir statement includes the message, a domain-separation tag, and the injectively encoded ordered ring. We otherwise follow the ring-signature template of Feneuil and Rivain [FR25, Section 6.6].

For every reported ring size, let $\widetilde{S}_N = \mathcal{H}_{\eta_1,2} \times \mathcal{H}_{\eta_2,2}$ be the candidate set. It contains

$$C(\eta_1, \eta_2) = \left(1 + \eta_1 + \binom{\eta_1}{2}\right)\left(1 + \eta_2 + \binom{\eta_2}{2}\right)$$

representatives. Subject to $C(\eta_1, \eta_2) \geq N$, we minimize $\eta_{\mathsf{sel}} = \eta_1 + \eta_2$, take the first N elements of $\widetilde{S}_N$ in lexicographic order as S_N, and map them to $(\boldsymbol{e}_i)_{i\in[N]}$. Every element of $\widetilde{S}_N$ has weight at most 4, so Lemma 5 gives full row rank of $\boldsymbol{G}_{\widetilde{S}_N,4}$ and hence of $\boldsymbol{G}_{S_N,4}$. By the arbitrary-target property in Section 3.2, there is a degree-at-most-4 polynomial p_N that is 0 on S_N and 1 on $\widetilde{S}_N \setminus S_N$. The Hamming-ball constraints together with $p_N = 0$ therefore characterize exactly the N valid selectors.

6.2 Security Analysis

We provide a security analysis in the classical random-oracle model. The proof follows the security framework and strategy of CDD+25 [CDD+25] for the $t = 1$ setting. At most u honestly generated users share $\boldsymbol{H}$. A forgery (m^*, R^*, σ^*) is counted only if every key in R^* is honest and uncorrupted and no member has signed the same pair (m^*, R^*). Building on the analysis of Feneuil and Rivain [FR25], the reduction further accounts for the multi-user loss caused by reusing $\boldsymbol{H}$ across these users. An informal theorem follows, and its complete proof is deferred to a future version.

Theorem 8. *Assume a complete and adaptively zero-knowledge non-interactive ZKAoK. Assume further that it is weakly simulation-extractable. For every classical PPT forger $\mathcal{A}$ in the above setting, the construction is correct, has anonymity advantage at most $2\varepsilon_{\mathsf{zk}}$, and yields an algorithm $\mathcal{B}$ against standard SD such that*

$$\mathrm{Adv}^{\mathsf{uf}}_{\mathsf{RS}}(\mathcal{A}) \le \varepsilon_{\mathsf{coll}} + \varepsilon_{\mathsf{zk}} + \varepsilon_{\mathsf{wse}} + \frac{u}{c_{\mathsf{reg}}} \mathrm{Adv}_{\mathsf{SD}}(\mathcal{B}).$$

Here, $c_{\mathsf{reg}} := e^{-1/\theta}(n/w)^w / \binom{n}{w}$ is the success factor in Theorem 3, while $\varepsilon_{\mathsf{coll}}$, $\varepsilon_{\mathsf{zk}}$, and $\varepsilon_{\mathsf{wse}}$ denote the honest-key collision, zero-knowledge, and extraction advantages, respectively.

Proof (Proof sketch). A valid proof establishes knowledge of $\boldsymbol{s}_0, \ldots, \boldsymbol{s}_w$ with $\phi_N(\boldsymbol{s}_0) = \boldsymbol{e}_j$ and regular weight-w vector $\boldsymbol{x} = (\phi(\boldsymbol{s}_1), \ldots, \phi(\boldsymbol{s}_w))$ satisfying $\boldsymbol{H}\boldsymbol{x} = \boldsymbol{y}_j$. Completeness and zero knowledge give correctness and hide j. For unforgeability, simulated signing and weak simulation-extractability recover $(j, \boldsymbol{x})$ from a fresh forgery; freshness makes this a common-matrix multi-user RSD solution for an uncorrupted honest user. Embedding a single-target RSD challenge at a random honest user succeeds with probability $1/u$, and applying the RSD-to-SD reduction gives the stated bound. Because the user is identified only after extraction, the factor u affects only the final SD term.

For the fixed-ring instantiations below, $u = N$. Since the c_{reg} loss is already included in our RSD parameter selection, the random-index step contributes the additional $\log_2 N$-bit loss—at most 20 bits for $N = 2^{20}$.

6.3 Concrete Parameters and Performance

The reduction in Section 6.2 incurs a multi-user security loss. We therefore raise the target security of the underlying RSD instance while retaining the NIST Level-I short VOLEitH parameters. With $\mu = 540$ and $(\eta_{\mathrm{RSD}}, d)=(11, 4)$, the resulting instantiation uses $(n, n-k, w)=(25\,380, 432, 47)$. This RSD instance provides 164 bits of security, so it still meets NIST Level I after the maximum 20-bit loss for rings of up to 2^{20} members.

With these parameters, the selector-free component of the signature is 3 541 bytes. Since the proof system uses $\tau = 11$ repetitions, an η_{sel}-bit selector gives

$$|\mathsf{sig}_{\mathsf{ring}}| = 3\,541 + 11 \left\lceil \frac{\eta_{\mathrm{sel}}}{8} \right\rceil \text{ bytes.}$$

The signature therefore grows from 3 563 bytes for 2^{10} members to 3 673 bytes for 2^{20} members. Increasing the ring size by a factor of 1 024 adds only 110 bytes.

Table 10. Post-quantum ring-signature sizes (bytes) at approximately 128-bit security

Scheme	2^6	2^{10}	2^{12}	2^{18}	2^{20}
CDD+25 [CDD+25] (AES-EM)	4 418	4 594	4 682	–	5 034
CDD+25 [CDD+25] (AES)	4 858	5 034	5 122	–	5 474
TCitH-AES-EM [FR25]	≈ 6 840	≈ 6 960	≈ 7 120	–	≈ 8 270
TCitH-AES [FR25]	≈ 7 900	≈ 8 020	≈ 8 180	–	≈ 9 390
TCitH-MQ [FR25]	≈ 4 330	≈ 4 450	≈ 4 600	–	≈ 5 620
TCitH-SD [FR25]	≈ 7 510	≈ 8 240	≈ 8 400	–	≈ 10 090
BBG+24 [BBGK24] (SD)	4 842	5 513	5 849	–	7 191
SMILE [LNS21]	–	≈ 17 270	–	–	≈ 20 150
BLOOM [LN22]	≈ 13 100	–	≈ 14 100	≈ 15 100	–
This work	**3 552**	**3 563**	**3 574**	**3 629**	**3 673**

Signature size. The comparison includes only ordinary, non-linkable ring signatures. Approximate byte values are converted from sizes reported in KB by the original papers. At $N = 2^{20}$, our signature is the smallest among the schemes compared in the table. Among the SD-based schemes, it is 49% smaller than BBG+24 and 64% smaller than TCitH-SD. It is also 27% smaller than the AES-EM construction in CDD+25 and 35% smaller than TCitH-MQ. SMILE and BLOOM are lattice-based ring signatures. Our signature is 82% smaller than SMILE at $N = 2^{20}$ and more than 70% smaller than BLOOM at every common ring size.

Running time. Following the online/offline decomposition of Xue *et al.* [XLAZ24], public-ring preprocessing depends only on the fixed ordered ring and is reused for subsequent messages. Table 11 reports complete times including this preprocessing and per-message times after its reuse. Our reported values are single-core medians over 1000 local runs on an AMD Ryzen 9 3900X. For CDD+25, we run the 128-bit AVX2 AES variant of their public implementation[9] with $(\tau, w, T) = (11, 7, 102)$. The hot-vector dimensions are 1, 2, 2, and 4 for $N = 2^6$, 2^{10}, 2^{12}, and 2^{20}, respectively. The CDD+25 values are medians over 50 runs on the same processor and logical core as ours.

Table 11. Median ring-signature running times in milliseconds

	Complete operation		After preprocessing	
Ring size N	Signing	Verification	Signing	Verification
2^6	6.599	6.032	6.588	6.021
2^{10}	6.285	6.213	6.112	6.041
2^{12}	6.791	6.763	6.098	6.071
2^{14}	9.257	8.933	6.489	6.165
2^{16}	18.600	17.603	7.530	6.533
2^{18}	55.012	52.493	10.649	8.130
2^{20}	200.665	191.332	23.362	14.029

At $N = 2^6$, 2^{10}, and 2^{12}, our complete signing is 1.16×, 3.42×, and 9.32× faster than CDD+25, while complete verification is 2.36×, 3.80×, and 7.53× faster. At $N = 2^{20}$, the corresponding speedups are 93.64× and 76.03×. After preprocessing, our signing and verification take 23.362 and 14.029 ms, giving amortized ratios of 804× and 1 037× relative to the complete CDD+25 timings. CDD+25 does not separate its execution into online and offline phases, so these ratios compare our post-preprocessing times with its complete signing and verification times. Our speedup comes from reusing $\boldsymbol{H}$ across the RSD public keys, so processing the ring requires only combining public syndromes, whereas the AES evaluations in CDD+25 grow with the ring size.

7 Conclusion

We introduce a general framework for representing elementary vectors by short binary strings and reconstructing them with low-degree polynomials. The framework gives an exact rank criterion for reconstruction and incorporates low-degree constraints for checking valid representatives. We establish a general upper bound on the number of vectors that can be represented and give optimal constructions for the parameter ranges used in our schemes. These constructions reduce the length of extended witness in VOLE-in-the-Head proofs for SD and PKP. They lead to shorter SD- and PKP-based signature schemes, with Level-I signature-plus-public-key sizes of 3 496 and 3 341 bytes, respectively. The framework can also reconstruct arbitrary target vectors. As one application, it yields a 3 673-byte ring signature for 2^{20} members.

Two main questions remain open. First, for general degree bounds d and representation lengths η, how can one explicitly construct an admissible representative set S of maximum size? Second, which cryptographic relations beyond the t-out-of-N selectors considered here can benefit from arbitrary-target reconstruction? For structured target families, their shared structure may enable more efficient constructions than treating each target independently.

[9] https://github.com/jachiang/PQ-Threshold-Ring-Sigs-from-VOLEitH.

References

ABB+25a. N. Aaraj, S. Bettaieb, L. Bidoux, A. Budroni, V. Dyseryn, A. Esser, T. Feneuil, P. Gaborit, M. Kulkarni, V. Mateu, M. Palumbi, L. Perin, M. Rivain, J.-P. Tillich, and K. Xagawa. PERK. NIST's Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes Project (Round 2), https://pqc-perk.org/, 2025.

ABB+25b. N. Aragon, M. Bardet, L. Bidoux, J.-J. Chi-Domínguez, V. Dyseryn, T. Feneuil, P. Gaborit, A. Joux, R. Neveu, M. Rivain, J.-P. Tillich, and A. Vinçotte. RYDE. NIST's Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes Project (Round 2), https://pqc-ryde.org/, 2025.

BBB+25a. C. Baum, W. Beullens, L. Braun, C. D. de Saint Guilhem, M. Klooß, C. Majenz, S. Mukherjee, E. Orsini, S. Ramacher, C. Rechberger, L. Roy, and P. Scholl. FAEST v2: Algorithm specifications. NIST's Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes Project, Round 2, Version 2.0, 2025.

BBB+25b. C. Baum, W. Beullens, L. Braun, C. D. de Saint Guilhem, M. Klooß, C. Majenz, S. Mukherjee, E. Orsini, S. Ramacher, C. Rechberger, L. Roy, and P. Scholl. Shorter, tighter, FAESTer: Optimizations and improved (QROM) analysis for VOLE-in-the-head signatures. In *CRYPTO 2025, Part VI*, *Lecture Notes in Computer Science* 16005, pages 124–156. Springer, Cham, 2025.

BBdSG+23a. C. Baum, L. Braun, C. D. de Saint Guilhem, M. Klooß, C. Majenz, S. Mukherjee, S. Ramacher, C. Rechberger, E. Orsini, L. Roy, and P. Scholl. FAEST: Algorithm specifications. NIST's Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes Project, Round 1, Version 1.0, 2023.

BBdSG+23b. C. Baum, L. Braun, C. D. de Saint Guilhem, M. Klooß, E. Orsini, L. Roy, and P. Scholl. Publicly verifiable zero-knowledge and post-quantum signatures from VOLE-in-the-head. In *CRYPTO 2023, Part V*, *LNCS* 14085, pages 581–615. Springer, 2023.

BBFR25. R. Benadjila, C. Bouillaguet, T. Feneuil, and M. Rivain. MQOM. NIST's Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes Project (Round 2), https://mqom.org/, 2025.

BBGK24. S. Bettaieb, L. Bidoux, P. Gaborit, and M. Kulkarni. Modelings for generic PoK and applications: Shorter SD and PKP based signatures. Cryptology ePrint Archive, Paper 2024/1668, 2024.

BBM+24. C. Baum, W. Beullens, S. Mukherjee, E. Orsini, S. Ramacher, C. Rechberger, L. Roy, and P. Scholl. One tree to rule them all: Optimizing GGM trees and OWFs for post-quantum signatures. In *ASIACRYPT 2024, Part I*, *LNCS* 15484, pages 463–493. Springer, 2024.

BCC88. G. Brassard, D. Chaum, and C. Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156–189, 1988.

BFR24. R. Benadjila, T. Feneuil, and M. Rivain. MQ on my mind: Post-quantum signatures from the non-structured multivariate quadratic problem. In *2024 IEEE 9th European Symposium on Security and Privacy (EuroS&P)*, pages 468–485. IEEE, 2024.

BGI16. E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing: Improvements and extensions. In *ACM CCS 2016*, pages 1292–1303. ACM, 2016.

BKPV24. L. Bettale, D. Kahrobaei, L. Perret, and J. A. Verbel. Biscuit: New MPCitH signature scheme from structured multivariate polynomials. In *Applied Cryptography and Network Security – ACNS 2024, Part I*, *LNCS* 14583, pages 457–486. Springer, 2024.

CCJ23. E. Carozza, G. Couteau, and A. Joux. Short signatures from regular syndrome decoding in the head. In *EUROCRYPT 2023, Part V*, *LNCS* 14008, pages 532–563. Springer, 2023.

CDD+25. J. H.-Y. Chiang, I. Damgård, W. R. Duro, S. Engan, S. Kolby, and P. Scholl. Post-quantum threshold ring signature applications from VOLE-in-the-head. In *ACM CCS 2025*, pages 4664–4678. ACM, 2025.

Cho25. T. Chou. Representations of elementary vectors in VOLE-in-the-head-based schemes. Cryptology ePrint Archive, Paper 2025/1336, 2025.

CLY+24. H. Cui, H. Liu, D. Yan, K. Yang, Y. Yu, and K. Zhang. ReSolveD: Shorter signatures from regular syndrome decoding and VOLE-in-the-head. In *PKC 2024, Part I*, *LNCS* 14601, pages 229–258. Springer, 2024.

EVZB24. A. Esser, J. A. Verbel, F. Zweydinger, and E. Bellini. SoK: `CryptographicEstimators` – a software library for cryptographic hardness estimation. In *ACM ASIACCS 2024*, pages 560–574. ACM, 2024.

FFS87. U. Feige, A. Fiat, and A. Shamir. Zero knowledge proofs of identity. In *STOC*, pages 210–217. ACM, 1987.

FJR22. T. Feneuil, A. Joux, and M. Rivain. Syndrome decoding in the head: Shorter signatures from zero-knowledge proofs. In *CRYPTO 2022, Part II*, *LNCS* 13508, pages 541–572. Springer, 2022.

FR23. T. Feneuil and M. Rivain. Threshold linear secret sharing to the rescue of MPC-in-the-head. In *ASIACRYPT 2023, Part I*, *LNCS* 14438, pages 441–473. Springer, 2023.

FR25. T. Feneuil and M. Rivain. Threshold computation in the head: Improved framework for post-quantum signatures and zero-knowledge arguments. *Journal of Cryptology*, 38(3):28, 2025.

FS86. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, Lecture Notes in Computer Science, pages 186–194, Berlin, 1986. Springer.

GMR85. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *STOC*, pages 291–304. ACM, 1985.

HBD+22. A. Hülsing, D. J. Bernstein, C. Dobraunig, M. Eichlseder, S. Fluhrer, S.-L. Gazdag, P. Kampanakis, S. Kölbl, T. Lange, M. M. Lauridsen, F. Mendel, R. Niederhagen, C. Rechberger, J. Rijneveld, P. Schwabe, J.-P. Aumasson, B. Westerbaan, and W. Beullens. SPHINCS+. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

IKOS07. Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Zero-knowledge from secure multiparty computation. In *STOC*, pages 21–30. ACM, 2007.

KMRP19. E. Koussa, G. Macario-Rat, and J. Patarin. On the complexity of the permuted kernel problem. Cryptology ePrint Archive, Paper 2019/412, 2019.

LDK+22. V. Lyubashevsky, L. Ducas, E. Kiltz, T. Lepoint, P. Schwabe, G. Seiler, D. Stehlé, and S. Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

LN22. V. Lyubashevsky and N. K. Nguyen. BLOOM: Bimodal lattice one-out-of-many proofs and applications. In *ASIACRYPT 2022, Part IV*, *LNCS* 13794, pages 95–125. Springer, Cham, December 2022.

LNS21. V. Lyubashevsky, N. K. Nguyen, and G. Seiler. SMILE: Set membership from ideal lattices with applications to ring signatures and confidential transactions. In *CRYPTO 2021, Part II*, *LNCS* 12826, pages 611–640, Virtual Event, August 2021. Springer, Cham.

MBB+25. C. A. Melchor, S. Bettaieb, L. Bidoux, T. Feneuil, P. Gaborit, N. Gama, S. Gueron, J. Howe, A. Hülsing, D. Joseph, A. Joux, M. Kulkarni, E. Persichetti, T. H. Randrianarisoa, M. Rivain, and D. Yue. SDitH. NIST's Post-Quantum Cryptography Standardization of Additional Digital Signature Schemes Project (Round 2), https://sdith.org/, 2025.

MGH+23. C. A. Melchor, N. Gama, J. Howe, A. Hülsing, D. Joseph, and D. Yue. The return of the SDitH. In *EUROCRYPT 2023, Part V*, *LNCS* 14008, pages 564–596. Springer, 2023.

NIS22. NIST. Call for additional digital signature schemes for the post-quantum cryptography standardization process. https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/call-for-proposals-dig-sig-sept-2022.pdf, 2022.

NIS26. NIST. Post-quantum cryptography: Additional digital signature schemes – round 3 additional signatures. https://csrc.nist.gov/Projects/pqc-dig-sig/round-3-additional-signatures, 2026.

OTX24. Y. Ouyang, D. Tang, and Y. Xu. Code-based zero-knowledge from VOLE-in-the-head and their applications: Simpler, faster, and smaller. In *ASIACRYPT 2024, Part V*, *LNCS* 15488, pages 436–470. Springer, 2024.

PFH+22. T. Prest, P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

SBC24. P. Santini, M. Baldi, and F. Chiaraluce. Computational hardness of the permuted kernel and subcode equivalence problems. *IEEE Transactions on Information Theory*, 70(3):2254–2270, 2024.

XLAZ24. Y. Xue, X. Lu, M. H. Au, and C. Zhang. Efficient linkable ring signatures: New framework and post-quantum instantiations. In *ESORICS 2024, Part IV*, *LNCS* 14985, pages 435–456. Springer, Cham, September 2024.

A Supporting Proofs for the Encoding Framework

A.1 Proof of Lemma 5

Proof. Let $S = \{\boldsymbol{s}_1, \ldots, \boldsymbol{s}_\mu\}$. For each point $\boldsymbol{s}_q$, define its support monomial $m_q(X_1, \ldots, X_\eta) = \prod_{u \in \mathrm{supp}(\boldsymbol{s}_q)} X_u$. The degree of m_q is $\mathsf{wt}(\boldsymbol{s}_q) \le d$, so m_q appears among the columns of $\boldsymbol{G}_{S,d}$.

Let $\boldsymbol{M}_S \in \mathbb{F}_2^{\mu \times \mu}$ be the sub-matrix of $\boldsymbol{G}_{S,d}$ formed by the columns indexed by $m_1, \ldots, m_\mu$, where the q-th selected column is the column indexed by m_q. By the definition of the evaluation matrix $\boldsymbol{G}_{S,d}$, the (p,q)-entry of $\boldsymbol{M}_S$ is the value of m_q evaluated at the point $\boldsymbol{s}_p$. Hence, we have $(\boldsymbol{M}_S)_{p,q} = m_q(\boldsymbol{s}_p) = \prod_{u \in \mathrm{supp}(\boldsymbol{s}_q)} (\boldsymbol{s}_p)_u$. Since all coordinates are binary, this product is equal to 1 exactly when $(\boldsymbol{s}_p)_u = 1$ for every $u \in \mathrm{supp}(\boldsymbol{s}_q)$. This condition is equivalent to saying that every nonzero coordinate of $\boldsymbol{s}_q$ is also a nonzero coordinate of $\boldsymbol{s}_p$, namely $\mathrm{supp}(\boldsymbol{s}_q) \subseteq \mathrm{supp}(\boldsymbol{s}_p)$. Therefore, $(\boldsymbol{M}_S)_{p,q} = 1$ if and only if $\mathrm{supp}(\boldsymbol{s}_q) \subseteq \mathrm{supp}(\boldsymbol{s}_p)$.

Now we re-order the points $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_\mu$ so that their Hamming weights are non-decreasing: $\mathsf{wt}(\boldsymbol{s}_1) \le \mathsf{wt}(\boldsymbol{s}_2) \le \cdots \le \mathsf{wt}(\boldsymbol{s}_\mu)$. We order the selected columns of $\boldsymbol{M}_S$ in the same way, meaning that the q-th column is still the column corresponding to the monomial m_q associated with the point $\boldsymbol{s}_q$.

We claim that, under this ordering, $\boldsymbol{M}_S$ is lower triangular with all diagonal entries equal to 1. First, consider a diagonal entry $(\boldsymbol{M}_S)_{q,q}$. By the containment characterization above, $(\boldsymbol{M}_S)_{q,q} = 1 \Leftrightarrow \mathrm{supp}(\boldsymbol{s}_q) \subseteq \mathrm{supp}(\boldsymbol{s}_q)$, which is always true. Hence every diagonal entry of $\boldsymbol{M}_S$ is equal to 1.

It remains to show that all entries above the diagonal are zero. We prove this by contradiction. Take any pair $p < q$. Suppose, for contradiction, that $(\boldsymbol{M}_S)_{p,q} = 1$. By the containment characterization, this means that $\mathrm{supp}(\boldsymbol{s}_q) \subseteq \mathrm{supp}(\boldsymbol{s}_p)$. Therefore, $\mathsf{wt}(\boldsymbol{s}_q) = |\mathrm{supp}(\boldsymbol{s}_q)| \le |\mathrm{supp}(\boldsymbol{s}_p)| = \mathsf{wt}(\boldsymbol{s}_p)$. On the other hand, since $p < q$ and the points were ordered by non-decreasing Hamming weight, we also have $\mathsf{wt}(\boldsymbol{s}_p) \le \mathsf{wt}(\boldsymbol{s}_q)$. Thus the two weights must be equal: $\mathsf{wt}(\boldsymbol{s}_p) = \mathsf{wt}(\boldsymbol{s}_q)$. However, we also have $\mathrm{supp}(\boldsymbol{s}_q) \subseteq \mathrm{supp}(\boldsymbol{s}_p)$. A containment between two finite sets of the same size must be equality, so $\mathrm{supp}(\boldsymbol{s}_q) = \mathrm{supp}(\boldsymbol{s}_p)$. Since the vectors are binary, equality of supports implies $\boldsymbol{s}_q = \boldsymbol{s}_p$, contradicting the fact that $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_\mu$ are distinct points of S. Therefore our assumption was false, and $(\boldsymbol{M}_S)_{p,q} = 0$ for every $p < q$.

We conclude that $\boldsymbol{M}_S$ has zeros above the diagonal and ones on the diagonal. That is, $\boldsymbol{M}_S$ is a lower triangular matrix with unit diagonal. Such a matrix is non-singular. Therefore, $\boldsymbol{M}_S$ has rank μ.

Finally, recall that $\boldsymbol{M}_S$ is a $\mu \times \mu$ sub-matrix of $\boldsymbol{G}_{S,d}$. Since $\boldsymbol{G}_{S,d}$ contains a nonsingular $\mu \times \mu$ submatrix, the row rank of $\boldsymbol{G}_{S,d}$ is at least μ. But $\boldsymbol{G}_{S,d}$ has exactly μ rows, so its row rank is exactly μ. Hence $\boldsymbol{G}_{S,d}$ has full row rank. □

The proof above also implies that not all columns of $\boldsymbol{G}_{S,d}$, that is, not all monomials, are needed for polynomial reconstruction. Instead, it suffices to use the μ columns indexed by the following supporting monomials:

$$m_q(X_1, \ldots, X_\eta) = \prod_{u \in \mathrm{supp}(\boldsymbol{s}_q)} X_u, \qquad q \in [\mu].$$

These columns form the non-singular matrix $\boldsymbol{M}_S$ constructed above. Hence, for each $j \in [\mu]$, there exist coefficients $c_{1,j}, \ldots, c_{\mu,j} \in \mathbb{F}_2$ such that the polynomial

$$f_j(X_1, \ldots, X_\eta) = \sum_{q=1}^{\mu} c_{q,j} \cdot m_q(X_1, \ldots, X_\eta) \tag{8}$$

satisfies $f_j(\boldsymbol{s}_i) = 1$ if $i = j$, and $f_j(\boldsymbol{s}_i) = 0$ otherwise. Therefore, we have $(f_1(\boldsymbol{s}_i), \ldots, f_\mu(\boldsymbol{s}_i)) = \boldsymbol{e}_i$ for every $i \in [\mu]$. Equivalently, in the full coefficient matrix $\boldsymbol{A}$, the coefficients of all monomials other than these μ supporting monomials would be set to zero.

A.2 Proof of the Universal Upper Bound

Proof. Choose a basis of $\ker(\boldsymbol{G}_{S,d})$, and let $q_1, \ldots, q_c \in \mathcal{P}_{\eta,d}$ be the corresponding polynomials. They vanish on S. Moreover, the polynomials used to prove membership in S belong to the same kernel and are therefore linear combinations of $q_1, \ldots, q_c$. Hence, their common zero set is exactly S, whose indicator polynomial can be written as $\mathbf{1}_S = \prod_{i=1}^{c}(1 + q_i)$, with $\deg(\mathbf{1}_S) \le cd$.

Suppose that $d(c+1) < \eta$. For every monomial $g^{\langle h \rangle}$ of degree at most d, the polynomial $\mathbf{1}_S g^{\langle h \rangle}$ has degree less than η. The parity of the truth table of a Boolean polynomial equals the coefficient of $X_1 \cdots X_\eta$ in its multilinear representation. Therefore, for every such $g^{\langle h \rangle}$,

$$\sum_{\boldsymbol{s} \in S} g^{\langle h \rangle}(\boldsymbol{s}) = \sum_{\boldsymbol{x} \in \mathbb{F}_2^\eta} \mathbf{1}_S(\boldsymbol{x}) g^{\langle h \rangle}(\boldsymbol{x}) = 0.$$

Thus, the sum of all rows of $\boldsymbol{G}_{S,d}$ is zero, contradicting its full row rank. We must have $d(c+1) \geq \eta$, which gives the claimed bound.

The bound need not be tight for arbitrary (η, d). When $(\eta, d) = (5, 3)$, it permits $c = 1$ and hence $|S| = k_3(5) - 1 = 25$. In that case, however, $\mathbf{1}_S = 1 + q_1$ would have degree at most $3 < 5$ and therefore even Hamming weight, contradicting $|S| = 25$.

B Optimal Representative Sets

B.1 Optimal Representative Sets for $d = 4$

The construction is motivated by the full Hamming ball: we exchange a small number of points across its boundary so that the modified set is cut out by two degree-four polynomials while retaining full reconstruction rank. Write $\eta = 8 + r$, where $1 \leq r \leq 4$, and fix a partition $[\eta] = U_0 \,\dot\cup\, X \,\dot\cup\, Z_0$, where $\dot\cup$ denotes disjoint union, with $|U_0| = 4$, $|X| = r - 1$, and $|Z_0| = 5$. Let $H := \mathcal{H}_{\eta,4}$ and write $\mathbf{1}_R$ for the incidence vector of $R \subseteq [\eta]$. Define

$$\begin{aligned} C &:= \{\mathbf{1}_R : R \subseteq X\} \cup \{\mathbf{1}_{U_0}\}, \qquad & D &:= \{\mathbf{1}_{U_0 \cup R} : \varnothing \neq R \subseteq X\}, \\ S_\eta &:= (H \setminus C) \,\dot\cup\, D. \end{aligned}$$

Since $|C| = 2^{r-1} + 1$ and $|D| = 2^{r-1} - 1$, we have $|S_\eta| = k_4(\eta) - 2$.

For $W \subseteq [\eta]$, let $e_j(W)$ be the j-th elementary symmetric polynomial in the variables indexed by W, and set $P_4(W) := \sum_{j=0}^4 e_j(W)$. Let $Z := [\eta] \setminus X = U_0 \,\dot\cup\, Z_0$ and $X_J := \prod_{j \in J} X_j$. The two degree-4 membership polynomials are

$$q_0 := P_4(Z), \qquad q_1 := P_4(Z) + P_4([\eta]) + X_{U_0}.$$

Proposition 9. *For every $9 \leq \eta \leq 12$, the set S_η defined above is admissible under degree bound 4. More precisely,*

$$\{\boldsymbol{x} \in \mathbb{F}_2^\eta : q_0(\boldsymbol{x}) = q_1(\boldsymbol{x}) = 0\} = S_\eta, \qquad \operatorname{rank}(\boldsymbol{G}_{S_\eta,4}) = |S_\eta| = \mu_4(\eta) = k_4(\eta) - 2.$$

Proof. Fix $\eta = 8 + r$ with $1 \leq r \leq 4$ and use the sets U_0, X, Z_0, Z, C, and D defined above. Since $|C| = 2^{r-1} + 1$ and $|D| = 2^{r-1} - 1$, we have $|S_\eta| = |\mathcal{H}_{\eta,4}| - |C| + |D| = k_4(\eta) - 2$.

We first prove that q_0 and q_1 vanish simultaneously exactly on S_η. For $0 \leq n \leq 12$, define $\mathit{\Pi}_4(n) := \sum_{j=0}^4 \binom{n}{j} \bmod 2$. Pascal's identity and Lucas' theorem give

$$\mathit{\Pi}_4(n) = 1 \quad \Longleftrightarrow \quad n = 0 \;\text{ or }\; 5 \leq n \leq 8.$$

Indeed, $\mathit{\Pi}_4(0) = 1$, while $\mathit{\Pi}_4(n) = 2^n = 0$ in $\mathbb{F}_2$ for $1 \leq n \leq 4$; for $n \geq 5$, Pascal's identity gives $\mathit{\Pi}_4(n) = \binom{n-1}{4} \bmod 2$, whose value follows from Lucas' theorem.

Write an arbitrary point as $\boldsymbol{x} = \mathbf{1}_Y$, and set $w := |Y|$, $t := |Y \cap X|$, $z := |Y \cap Z| = w - t$, and $u := \mathbf{1}[U_0 \subseteq Y]$. Since $P_4(W)(\mathbf{1}_Y) = \mathit{\Pi}_4(|Y \cap W|)$, we have

$$q_0(\mathbf{1}_Y) = \mathit{\Pi}_4(z), \qquad q_1(\mathbf{1}_Y) = \mathit{\Pi}_4(z) + \mathit{\Pi}_4(w) + u.$$

For $w \leq 4$, the values on $\mathbf{1}_\varnothing$, $\mathbf{1}_{U_0}$, and $\mathbf{1}_R$ with $\varnothing \neq R \subseteq X$ are $(1,0)$, $(0,1)$, and $(1,1)$, respectively, whereas every point of $\mathcal{H}_{\eta,4} \setminus C$ has value $(0,0)$.

It remains to consider $w \geq 5$. If $5 \leq w \leq 8$, then $\mathit{\Pi}_4(w) = 1$, so common vanishing requires $\mathit{\Pi}_4(z) = 0$ and $u = 1$. The latter gives $z \geq 4$, and the displayed characterization of $\mathit{\Pi}_4$ forces $z = 4$. Hence $Y = U_0 \cup R$ for some nonempty $R \subseteq X$, which gives exactly the points in D. If $9 \leq w \leq \eta \leq 12$, then $\mathit{\Pi}_4(w) = 0$, so common vanishing requires $u = 0$. Since $t \leq |X| \leq 3$, we have $z \geq 6$; the equality

$\Pi_4(z) = 0$ and the bound $z \le |Z| = 9$ then force $z = 9$. This implies $Z \subseteq Y$ and hence $U_0 \subseteq Y$, contradicting $u = 0$. Therefore, $q_0(\boldsymbol{x}) = q_1(\boldsymbol{x}) = 0$ if and only if $\boldsymbol{x} \in (\mathcal{H}_{\eta,4} \setminus C) \dot\cup D = S_\eta$. Both membership polynomials have degree at most 4.

We next construct degree-4 reconstruction polynomials, which directly proves full row rank. For $B \subseteq [\eta]$ with $|B| \le 4$, write $X_B := \prod_{i\in B} X_i$ and define

$$L_B(\boldsymbol{X}) := X_B \sum_{j=0}^{4-|B|} e_j([\eta] \setminus B).$$

For $|Y| \le 4$, we have $L_B(\mathbf{1}_Y) = 0$ when $B \not\subseteq Y$, while for $B \subseteq Y$ its value is $\sum_{j=0}^{4-|B|} \binom{|Y|-|B|}{j} = 2^{|Y|-|B|}$ in $\mathbb{F}_2$. Hence $L_B(\mathbf{1}_Y) = 1$ if and only if $B = Y$. For each nonempty $R \subseteq X$, put $\mathbf{d}_R := \mathbf{1}_{U_0 \cup R}$. If $R, R' \subseteq X$ are nonempty, then $L_R(\mathbf{d}_{R'}) = 0$ when $R \not\subseteq R'$; for $R \subseteq R'$, its value is $\binom{3+|R'|-|R|}{4-|R|} \bmod 2$. This coefficient is odd when $|R| = |R'|$ and even for the three possible strict inclusions $(|R|, |R'|) \in \{(1,2), (1,3), (2,3)\}$. Thus $L_R(\mathbf{d}_{R'}) = \mathbf{1}[R = R']$.

Associate to each inserted point $\mathbf{d}_R \in D$ and each retained point $\mathbf{1}_B \in \mathcal{H}_{\eta,4} \setminus C$ the polynomials

$$p_{\mathbf{d}_R} := L_R, \qquad p_{\mathbf{1}_B} := L_B + \sum_{\varnothing \ne R \subseteq X} L_B(\mathbf{d}_R) L_R.$$

All these polynomials have degree at most 4. The preceding evaluations show that $p_{\mathbf{d}_R}$ is one at $\mathbf{d}_R$ and zero at every other point of S_η. For a retained point $\mathbf{1}_B$, the correction term vanishes on $\mathcal{H}_{\eta,4} \setminus C$ and cancels L_B at every inserted point. Hence $p_{\boldsymbol{s}}(\boldsymbol{t}) = \mathbf{1}[\boldsymbol{s} = \boldsymbol{t}]$ for all $\boldsymbol{s}, \boldsymbol{t} \in S_\eta$. Let $\boldsymbol{A}$ contain the coefficient vectors of these polynomials. Then $\boldsymbol{G}_{S_\eta,4}\boldsymbol{A} = \boldsymbol{I}_{|S_\eta|}$, so $\boldsymbol{G}_{S_\eta,4}$ has full row rank and S_η is admissible under degree bound 4.

Finally, the universal upper bound gives $k_4(\eta) - |S| \ge \lceil \eta/4 \rceil - 1 = 2$ for every admissible $S \subseteq \mathbb{F}_2^\eta$ and $9 \le \eta \le 12$. Since $|S_\eta| = k_4(\eta) - 2$, the construction attains this bound and is optimal.

Products of Hamming balls are constrained by both $\sum_{j=1}^{\ell} r_j \le d$ and $r_j + 1 \le d$. By contrast, the construction above loses only $|C| - |D| = 2$ representatives relative to the full Hamming ball, while the two polynomials q_0, q_1 cut out the resulting set and its evaluation matrix retains full row rank.

B.2 Useful Optimal Representative Sets for $d = 3$

For $d = 3$, we give four optimal representative sets covering the parameter range needed by our applications. The constructions for $\eta = 7, 8, 9$ arise from structured algebraic ansätze: a two-slice Pascal decomposition, a block-Pascal basis exchange, and a cyclic two-cubic graph, respectively. For polynomials $q_1, \ldots, q_c$, write $Z(q_1, \ldots, q_c) := \{\boldsymbol{x} : q_1(\boldsymbol{x}) = \cdots = q_c(\boldsymbol{x}) = 0\}$. Their parameters are

η	$k_3(\eta)$	$k_3(\eta) - \lvert S_\eta\rvert$	$\lvert S_\eta\rvert$
7	64	2	62
8	93	3	90
9	130	2	128
12	299	3	296

The case $\eta = 7$. Take coordinates $(z_1, \ldots, z_6, y)$, and let $a = e_1(z_2, \ldots, z_6)$, $b = e_2(z_2, \ldots, z_6)$, and $L = z_3(z_2 + z_4 + z_5)$. Define

$$\begin{aligned} q_1 &= ab + L + y(1 + a + b + L), \\ q_2 &= z_1 b + z_4 z_5 z_6 + z_2 z_3 + y\big(z_1(1 + a) + z_2 z_3\big), \end{aligned} \qquad S_7 = Z(q_1, q_2).$$

For $w = \mathsf{wt}(z_2, \ldots, z_6)$, Lucas' theorem gives $a = w \bmod 2$, $b = \lfloor w/2 \rfloor \bmod 2$, and $ab = e_3(z_2, \ldots, z_6)$ as Boolean functions; hence q_1, q_2 have degree at most three. After the change of variables $x_i = 1 + z_i$, the two Pascal slices of $\mathbb{F}_2^7 \setminus S_7$ have sizes 42 and 24, and their degree-three and degree-two evaluation matrices have ranks 42 and 22, respectively. The induced block decomposition gives $\operatorname{rank}(\boldsymbol{G}_{\mathbb{F}_2^7 \setminus S_7, 3}) = 64$, and Reed–Muller duality then gives $|S_7| = 62$ and $\operatorname{rank}(\boldsymbol{G}_{S_7,3}) = 62$.

The case $\eta = 8$. Fix a partition $[8] = A \dot\cup B \dot\cup \{\alpha, \beta\}$, where $|A| = |B| = 3$, and let $H = \mathcal{H}_{8,4}$. Define

$$\begin{aligned} &U_1 = A \cup B \cup \{\alpha\}, \quad U_2 = A \cup B \cup \{\beta\}, \quad U_3 = A \cup \{\alpha, \beta\}, \\ &U_4 = B \cup \{\alpha, \beta\}, \quad U_5 = A \cup B, \quad U_6 = [8]. \end{aligned}$$

With $D = \{\varnothing, A \cup \{\alpha\}, B \cup \{\beta\}\}$, set

$$S_8 = \big((\mathbb{F}_2^8 \setminus H) \setminus \{\mathbf{1}_{U_1}, \ldots, \mathbf{1}_{U_6}\}\big) \dot\cup \{\mathbf{1}_P : P \in D\}.$$

Thus, $|S_8| = (256 - 163) - 6 + 3 = 90$. For $|U| \geq 5$, let

$$Q_U(\boldsymbol{x}) = \prod_{i \notin U} (1 + x_i) \sum_{j=0}^{|U|-5} e_j(1 + x_i : i \in U), \qquad \begin{aligned} q_1 &= Q_{U_1} + Q_{U_3}, \\ q_2 &= Q_{U_2} + Q_{U_4}, \\ q_3 &= Q_{U_3} + Q_{U_4} + Q_{U_5} + Q_{U_6}. \end{aligned}$$

Each Q_U has degree at most three. Pascal's identity shows that $Z(q_1, q_2, q_3) = S_8$: the six points $\mathbf{1}_{U_i}$ are excluded, while precisely the three points indexed by D are retained from H. The coefficient matrix $(Q_{U_i}(\mathbf{1}_P))_{P \in D,\, i \in [6]}$ has rank 3, so $T_8 := \mathbb{F}_2^8 \setminus S_8$ satisfies $\operatorname{rank}(\boldsymbol{G}_{T_8,4}) = 163$. Reed–Muller duality therefore gives $\operatorname{rank}(\boldsymbol{G}_{S_8,3}) = 90$.

The case $\eta = 9$. Take $X = (x_0, x_1, x_2)$, $Y = (y_0, y_1, y_2)$, and t, with indices interpreted modulo three. Define

$$\begin{aligned} F_1 &= y_0 y_1 y_2 + \sum_{i=0}^{2} x_i x_{i+1}(y_i + t), \\ F_2 &= \sum_{i=0}^{2} y_i\big(x_{i+1}(x_{i+2} + y_{i+1}) + t x_{i+2}\big), \\ S_9 &= \{(X, Y, t, z_1, z_2) : z_1 = F_1,\ z_2 = F_2\}. \end{aligned}$$

The two cubics $z_1 + F_1$ and $z_2 + F_2$ give the membership proof, while the graph structure immediately gives $|S_9| = 2^7 = 128$. A cyclic interpolation argument gives

$$\mathcal{P}_{7,3} + F_1\mathcal{P}_{7,2} + F_2\mathcal{P}_{7,2} + F_1F_2\mathcal{P}_{7,1} = \mathbb{F}_2^{\mathbb{F}_2^7}.$$

Thus, every function on the seven free coordinates is obtained by restricting a degree-three polynomial to S_9, which proves $\operatorname{rank}(\boldsymbol{G}_{S_9,3}) = 128$.

The case $\eta = 12$. This denser construction was found using a dedicated search script. For $a = \sum_{j=1}^{12} a_j 2^{j-1}$, let $m_a(\boldsymbol{x}) = \prod_{j:a_j=1} x_j$. Define

$$q_i = \sum_{a \in \mathcal{C}_i} m_a, \qquad S_{12} = Z(q_1, q_2, q_3).$$

Here the three coefficient-support sets are the following:

$\mathcal{C}_1 = \{$0, 3, 4, 5, 12, 13, 18, 19, 20, 21, 25, 26, 32, 36, 37, 38, 40, 41, 42, 44, 49, 52, 67, 68, 72, 74, 76, 81, 82, 84, 97, 100, 128, 131, 133, 140, 148, 164, 168, 176, 192, 208, 258, 260, 262, 266, 268, 273, 274, 288, 289, 290, 292, 320, 322, 336, 352, 385, 388, 416, 513, 514, 515, 516, 521, 528, 532, 544, 548, 552, 560, 577, 580, 592, 640, 641, 642, 644, 704, 770, 772, 800, 896, 1025, 1026, 1034, 1040, 1041, 1042, 1056, 1058, 1060, 1072, 1088, 1104, 1120, 1153, 1156, 1160, 1168, 1184, 1216, 1280, 1282, 1296, 1408, 1538, 1552, 1568, 1600, 2048, 2049, 2050, 2051, 2056, 2057, 2058, 2080, 2081, 2084, 2088, 2114, 2120, 2128, 2144, 2180, 2184, 2192, 2240, 2306, 2308, 2320, 2368, 2432, 2560, 2568, 2576, 2624, 2816, 3073, 3076, 3088, 3104, 3136, 3584$\}$.

$\mathcal{C}_2 = \{$0, 1, 3, 4, 5, 6, 7, 10, 11, 12, 13, 17, 26, 32, 33, 34, 37, 38, 40, 41, 42, 48, 56, 66, 69, 73, 74, 76, 80, 81, 82, 96, 98, 100, 104, 112, 133, 137, 138, 140, 144, 148, 152, 161, 164, 176, 192, 194, 196, 200, 208, 224, 256, 258, 260, 264, 272, 276, 321, 324, 328, 336, 385, 386, 392, 416, 513, 514, 516, 517, 518, 520, 522, 530, 536, 545, 546, 548, 577, 592, 608, 641, 642, 644, 672, 769, 770, 772, 800, 832, 1024, 1032, 1033, 1034, 1042, 1044, 1060, 1064, 1088, 1090, 1092, 1096, 1104, 1120, 1152, 1154, 1168, 1216, 1281, 1284, 1296, 1312, 1344, 1408, 1538, 1540, 1544, 1552, 1568, 2050, 2051, 2052, 2056, 2060, 2066, 2082, 2088, 2096, 2112, 2113, 2114, 2116, 2176, 2178, 2184, 2192, 2208, 2240, 2304, 2305, 2308, 2368, 2432, 2561, 2564, 2576, 2592, 3076, 3104, 3136, 3328, 3584$\}$.

$\mathcal{C}_3 = \{$4, 6, 7, 8, 9, 11, 14, 16, 20, 21, 22, 32, 33, 34, 35, 36, 37, 41, 44, 56, 64, 66, 67, 73, 74, 80, 81, 88, 97, 98, 100, 104, 128, 131, 133, 136, 137, 138, 144, 146, 148, 152, 160, 162, 164, 176, 192, 194, 200, 224, 258, 260, 261, 262, 265, 268, 274, 280, 288, 296, 304, 336, 352, 400, 416, 448, 513, 514, 515, 517, 521, 528, 529, 530, 532, 536, 544, 548, 552,

580, 584, 608, 640, 642, 644, 656, 672, 768, 770, 772, 776, 784, 800, 1025, 1026, 1027, 1032, 1033, 1041, 1042, 1056, 1058, 1060, 1064, 1072, 1089, 1090, 1096, 1120, 1152, 1156, 1160, 1168, 1184, 1216, 1288, 1296, 1312, 1536, 1537, 1538, 1568, 1600, 1664, 1792, 2048, 2049, 2051, 2052, 2053, 2056, 2060, 2064, 2065, 2072, 2081, 2084, 2112, 2113, 2114, 2128, 2144, 2177, 2184, 2240, 2308, 2320, 2336, 2560, 2564, 2624, 2688, 3072, 3074, 3076, 3080, 3088, 3104, 3136, 3328, 3584}.

Every listed mask has Hamming weight at most three, so q_1, q_2, q_3 are degree-three polynomials. The resulting set satisfies $|S_{12}| = 296$ and $\mathrm{rank}(\boldsymbol{G}_{S_{12},3}) = 296$. Hence q_1, q_2, q_3 give the membership proof, while full row rank gives the degree-three reconstruction polynomials. The later application with $\mu = 295$ uses a 295-element subset of S_{12}; deleting one representative preserves full row rank, and a degree-three interpolation polynomial excludes the deleted point.

Finally, let $c = k_3(\eta) - |S|$. The universal upper bound $c \geq \lceil \eta/3 \rceil - 1$ proves optimality for $\eta = 7, 9, 12$. For $\eta = 8$, $c = 2$ would make the indicator of S a polynomial of degree at most six and therefore of even Hamming weight, whereas $k_3(8) - 2 = 91$ is odd. Hence $c \geq 3$, proving that all four constructions are optimal.

Although S_{12} is optimal, its dense coefficient representation is parameter-specific and does not yet reveal a concise algebraic structure comparable to the constructions for $\eta = 7, 8, 9$. Finding such a structural description and determining whether it extends to a uniform family are left for future work.

C Implementation Parameters and Detailed Comparisons

This appendix reports the proof-system parameters used by our implementations and the complete comparison data with PERK and Chou25 summarized in Section 5.2.

C.1 VOLEitH Proof-System Parameters

Our SD-based scheme uses the VOLEitH proof-system parameters of SDitH v2.0 [MBB+25], whereas our PKP-based scheme uses those of PERK v2.2 [ABB+25a]. This keeps each implementation aligned with its main comparison target. The SDitH parameter sets use the same depth k for all τ repetitions and satisfy $\tau k + w_{\mathsf{pow}} = \lambda + 2$. PERK follows the mixed-depth parameterization of FAEST v2: τ_0 repetitions have depth κ_0 and τ_1 have depth κ_1, where $\tau = \tau_0 + \tau_1$ and $\tau_0\kappa_0 + \tau_1\kappa_1 + w_{\mathsf{pow}} = \lambda + 2$. In both cases, w_{pow} is the number of proof-of-work bits and T_{open} is the maximum number of opened commitment nodes; "Short" and "Fast" denote the size- and speed-oriented variants, respectively. Table 12 reports the two parameter sets together; in the depth-profile column, $a \times b$ denotes a repetitions of depth b.

Table 12. VOLEitH parameters used by our SD- and PKP-based schemes, following SDitH [MBB+25] and PERK v2.2 [ABB+25a]

Security	Variant	Source	τ	Depth profile	w_{pow}	T_{open}
NIST I	Short	SDitH	11	11×11	9	107
		PERK	11	11×11	9	106
	Fast	SDitH	16	16×8	2	101
		PERK	16	$9 \times 8 + 7 \times 7$	9	110
NIST III	Short	SDitH	16	16×12	2	157
		PERK	16	$8 \times 12 + 8 \times 11$	10	166
	Fast	SDitH	24	24×8	2	153
		PERK	24	$16 \times 8 + 8 \times 7$	10	166
NIST V	Short	SDitH	21	21×12	6	216
		PERK	22	$8 \times 12 + 14 \times 11$	8	222
	Fast	SDitH	32	32×8	2	207
		PERK	32	$26 \times 8 + 6 \times 7$	8	220

C.2 Comparison with PERK

Table 13 compares our PKP-based scheme with PERK v2.2 [ABB+25a] at all three security levels and for both variants. Across these parameter sets, our signatures are 6.7%–12.5% smaller; signing and verification are 1.54×–2.28× and 1.35×–2.37× faster, respectively.

Table 13. Comparison of our PKP-based scheme with PERK v2.2 [ABB+25a]

Scheme	Sizes (bytes)			Runtimes (ms)	
	$\|\mathsf{pk}\|$	$\|\mathsf{sig}\|$	$\|\mathsf{pk}\| + \|\mathsf{sig}\|$	Sign	Verify
Ours-L1-short	**102**	**3239**	**3341**	**6.161**	**5.804**
PERK-L1-short	104	3473	3577	9.469	8.240
Ours-L1-fast	**102**	**3988**	**4090**	**1.493**	**1.254**
PERK-L1-fast	104	4312	4416	2.352	1.695
Ours-L3-short	**148**	**7420**	**7568**	**18.569**	**16.759**
PERK-L3-short	151	8311	8462	41.960	37.374
Ours-L3-fast	**148**	**9116**	**9264**	**5.306**	**3.715**
PERK-L3-fast	151	10423	10574	11.251	7.574
Ours-L5-short	**195**	**13646**	**13841**	**30.071**	**26.160**
PERK-L5-short	195	14830	15025	68.456	61.974
Ours-L5-fast	**195**	**16532**	**16727**	**10.864**	**7.030**
PERK-L5-fast	195	18216	18411	20.905	15.344

C.3 Comparison with Chou25

Table 14 compares our signature sizes with those of Chou25 [Cho25, Secs. 5.2 and 5.3], which does not report implementation results and provides PKP sizes only at NIST Levels I and V.

Table 14. Signature sizes of our schemes and Chou25 [Cho25]

Scheme	Signature Size (bytes)		
	NIST I	NIST III	NIST V
This work, SD (short)	**3431**	**7428**	**13180**
Chou25, SD (short)	3551	7676	13575
This work, SD (fast)	**4100**	**9140**	**16148**
Chou25, SD (fast)	4260	9484	16708
This work, PKP (short)	**3239**	–	**13646**
Chou25, PKP	3257	–	13842

For SD, our signatures are 2.9%–3.8% smaller; for PKP, the Level-I sizes are essentially the same and ours is approximately 1.4% smaller at Level V.