



HAL
open science

Faster Signatures from MPC-in-the-Head

Dung Bui, Eliana Carozza, Geoffroy Couteau, Dahmun Goudarzi, Antoine Joux

► **To cite this version:**

Dung Bui, Eliana Carozza, Geoffroy Couteau, Dahmun Goudarzi, Antoine Joux. Faster Signatures from MPC-in-the-Head. ASIACRYPT 2024 - International Conference on the Theory and Application of Cryptology and Information Security, Dec 2024, Kolkata, India. hal-04770548

HAL Id: hal-04770548

<https://hal.science/hal-04770548v1>

Submitted on 7 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Faster Signatures from MPC-in-the-Head

Dung Bui¹, Eliana Carozza¹, Geoffroy Couteau², Dahmun Goudarzi³, and Antoine Joux⁴

¹ IRIF, Université Paris Cité, Paris, France.
lastname@irif.fr

² CNRS, IRIF, Université Paris Cité, Paris, France.
couteau@irif.fr

³ Quarkslab.

dahmun.goudarzi@gmail.com

⁴ CISA Helmholtz Center for Information Security, Saarbrücken, Germany.
joux@cisca.de

Abstract. We revisit the construction of signature schemes using the MPC-in-the-head paradigm. We obtain two main contributions:

- We observe that previous signatures in the MPC-in-the-head paradigm must rely on a salted version of the GGM puncturable pseudorandom function (PPRF) to avoid collision attacks. We design a new efficient PPRF construction that is provably secure in the multi-instance setting. The security analysis of our PPRF, in the ideal cipher model, is quite involved and forms a core technical contribution of our work. While previous constructions had to rely on a hash function, our construction uses only a fixed-key block cipher and is considerably more efficient as a result: we observe a $12\times$ to $55\times$ speed improvement for a recent signature scheme (Joux and Huth, Crypto’24). Our improved PPRF can be used to speed up many MPC-in-the-head signatures.
- We introduce a new signature scheme from the regular syndrome decoding assumption, based on a new protocol for the MPC-in-the-head paradigm, which significantly reduces communication compared to previous works. Our scheme is conceptually simple, though its security analysis requires a delicate and nontrivial combinatorial analysis.

1 Introduction

Zero-Knowledge, Post-quantum Signatures from Syndrome Decoding Problem. Zero-knowledge proofs allow a prover to demonstrate knowledge of a solution to an NP problem without revealing any additional information. These proofs have numerous cryptographic applications, including their transformation into signature schemes via the Fiat-Shamir heuristic [FS87]. With the growing threat of quantum computing, there’s been a push to find post-quantum signatures, particularly in the area of code-based signatures. The syndrome decoding problem, which involves finding a vector of a small Hamming weight that satisfies a linear equation, is a fundamental code-based problem in post-quantum cryptography and is widely used in constructing efficient post-quantum signatures [FJR22, AGH⁺23, CCJ23].

The MPC in-the-Head Paradigm and PPRF. Recent advancements in Fiat-Shamir code-based signatures have utilized the MPC in the head paradigm [IKOS07]. This approach allows the prover to simulate an MPC protocol in-the-Head, where virtual parties are given shares of the witness, and the target function verifies that the witness is correct. Later, the prover commits the views of all parties, then the verifier checks the consistency of the views from requesting the prover to open $N - 1$ views of parties. This paradigm has proven effective in creating efficient zero-knowledge proofs with low communication overhead [AHIV17], leading to some of the most competitive post-quantum signature schemes [FJR22, AGH⁺23, CCJ23]. To commit the views of all virtual parties and later efficiently open all-but-one views, a common primitive called puncturable pseudorandom function (PPRF) is used in all MPCitH-based signatures to open the views of $N - 1$ parties while only need to reveal $\log N$ values.

Our Contribution. In this work, we revisit signature schemes constructed from the MPC-in-the-head (MPCitH) paradigm. We make two contributions. Our first contribution appeals to all MPCitH signatures, while our second contribution is in the context of code-based MPCitH signatures:

- We introduce the notion of multi-instance puncturable pseudorandom function (PPRF) together with an extremely efficient instantiation from the AES block cipher, which we formally prove secure in the ideal cipher model. Our construction can be used as a drop-in replacement to the hash-based PPRF used in most previous MPCitH signatures, and yields significant improvements in both signing time and verification time (e.g., from $12\times$ to $55\times$ in our experiments with the recent scheme of [HJ24]).
- We introduce a new MPCitH signature from the regular syndrome decoding (RSD) problem. We formally prove that its unforgeability tightly reduces to the multi-instance security of the underlying PPRF, showcasing how our new primitive results in better security reduction (hence better efficiency for a target security level). Our new signature scheme relies on a non-trivial combinatorial analysis, and significantly improves over the state-of-the-art scheme of [CCJ23] on all aspects (signature size, efficiency, security).

1.1 Faster MPCitH signatures from a new multi-instance PPRF

All state-of-the-art MPCitH signature schemes rely at their core on a puncturable pseudorandom function, which allows to generate a large number n of pseudorandom strings such that given an index i , the signer can reveal all pseudorandom values except the i -th one using an “opening” of size $\lambda \cdot \log n$ (where λ is a security parameter). The *de facto* PPRF originally used in prominent works such as Picnic [ZCD⁺20] was the GGM PPRF [GGM86], where the PRF evaluation on input i with key K is the i -th leave of a full binary tree with root labeled with K . In GGM, the labels of the two children of a node x are computed as $F(x) = (F_0(x), F_1(x))$ using a length-doubling pseudorandom generator (PRG) $x \mapsto (F_0(x), F_1(x))$. The PRG is typically instantiated as $x \mapsto (x \oplus \pi_0(x), x \oplus \pi_1(x))$ for a pair of random invertible permutations (π_0, π_1) : this instantiation yields a provably secure construction in the random permutation model [GKWY20] and an extremely fast expansion when instantiating the permutations using the AES block cipher with a fixed key (using commonly available hardware acceleration for AES). However, it was observed in [DN19] that in the context of signatures, this choice allows for a devastating multi-target attack: after 2^t signature queries, an attacker can find the root of one of the GGM trees using on average 2^{128-t} work, and recover the secret signing key as soon as a collision is found.

In reaction to this attack, Picnic [ZCD⁺20] and most recent works on MPCitH signatures, including BBQ [DDOS19], Banquet [BDK⁺21], all MPCitH candidates of the recent NIST call for additional post-quantum signatures (SDitH [AFG⁺23], MIRA [ABB⁺23c], MiRitH [ARV⁺23], MQOM [FR23], PERK [ABB⁺23a], Ryde [ABB⁺23b], and Biscuit [BKPV23]) and more, all implement the PRG using a hash function (such as SHAKE), as follows: $F_b(x, i, \text{salt}) \leftarrow H(x \| i \| b \| \text{salt})$, where i is the index of the parent node and salt is some random salt (included in the signature). Unfortunately, because of the hardware support for AES, replacing AES with a hash function is considerably slower (up to $50\times$ slower according to [GKWY20]).

A faster construction from ideal ciphers. The reason for choosing the hash-based construction over the existing AES-based construction comes from the need to add a per-signature salt at every node computation to thwart the multi-target attack; however, AES can only take a fixed 128-bit input. We make the following simple observation: in the AES-based instantiation, instead of relying on a global fixed-key that remains identical across all instances (as was originally done in Picnic [ZCD⁺20]), the signer can use a per-signature random AES key that will play the role of the salt (crucially, this key needs not be changed at every node of the tree, avoiding costly re-keying of AES), without having to increase the block size.

A multi-instance PPRF in the ideal cipher model. The idea in itself is surprisingly simple, and we do not claim it to be particularly novel: the idea of rotating the key has been mentioned in the literature in other contexts, for example, in [Roy22]. Our main contributions here are twofold:

- We introduce the notion of multi-instance PPRFs, capturing the exact security requirement that the PPRF must satisfy to avoid multi-target attacks (in contrast, previous works made a direct proof of the full signature construction by modeling the hash function in the ROM; our approach is much more modular)
- We formally prove that the PPRF instantiated with $F(x, K_0, K_1) = (x \oplus \text{AES}_{K_0}(x), x \oplus \text{AES}_{K_1}(x))$ is a multi-instance PPRF when AES is modeled as an ideal cipher. Our security analysis is non-trivial and forms a core technical contribution of our work. It relies on the H -coefficient technique of Patarin [Pat09, CS14].

We further expect our new PPRF to find other applications beyond the scope of MPCitH signatures. We sketch one application in Appendix A.

On the security loss of our construction. Our analysis induces a loss of $\log D + 3$ bits of security, where D is the depth of the GGM tree. For example, if $D = 8$, this translates to a loss of 6 bits of security. Before we proceed, we comment on this loss:

- The loss of 3 bits of security comes from bounding the worst-case runtime of the attacker (*i.e.* proving that with probability $2^{-\lambda}$, the attacker must run in time 2^λ to break the scheme with high probability). If bound instead the *expected* runtime of the attacker, the loss reduces to 1 bit.
- the $\log_2 D$ loss comes from the D hybrids in the PRG-to-PRF reduction, similar to [GGM86]. This loss is inherent to the construction (*i.e.*, the analysis is perfectly tight on this aspect).

To mitigate the second loss, we suggest the following variant of our construction: instead of relying on a global pair of keys (K_0, K_1) for the full GGM tree, use instead D pairs of keys $(K_0^i, K_1^i)_{i \leq D}$ (stretched from some 2λ -bit signature salt using a PRG), one for each level of the tree. Then, evaluate the GGM tree using the key pair (K_0^i, K_1^i) for nodes that are the i -th level of the tree. This induces a slight increase in the number of AES rekeying ($2D$ instead of 2), but the cost is completely negligible compared to the rest (the reader can think of D as an integer between 8 and 16). We put forth the conjecture that in the ideal cipher model, this yields a multi-instance PPRF with $\lambda - 3$ bits of security (or $\lambda - 1$ if we bound the expected runtime of the attacker instead), where λ is the key length of the block cipher.

Unfortunately, it is not straightforward to adapt our security analysis to this variant: our analysis proceeds via a multi-instance PPRF to multi-instance PRG reduction, followed by a proof of security of the multi-instance PRG in the ideal cipher model. The $\log_2 D$ loss stems from the first reduction, and our variant does not generically remove this loss for an arbitrary multi-instance PRG. We conjecture that it does, however, remove this loss when the PRG is instantiated in the ideal cipher model. Proving this conjecture appears to require analyzing directly the full multi-instance PPRF in the ideal cipher model. We believe that this is not out of reach, but it goes beyond the scope of this work, and we leave it to future work.

Case analysis: the signature scheme of [HJ24]. We expect that replacing the hash-based PPRF with our AES-based construction in existing MPCitH signature schemes will yield to significant efficiency improvements. The improvement is especially noticeable if the cost of expanding the GGM tree represents a large fraction of the overall runtime. We note that if the most optimized state-of-the-art MPCitH signature, the GGM tree expansion indeed tends to account for most of the running time (for both signing and verification).

To illustrate the concrete efficiency improvements one might expect from using our PPRF, we focus on the recent scheme of [HJ24], an extremely efficient MPCitH signature based on the subfield bilinear collision problem. We run the signature scheme using both the hash-based multi-instance PPRF used in previous works, and our AES-based multi-instance PPRF. In both cases, for fairness of comparison, we include all relevant algorithmic optimizations (this includes an algorithmic optimization introduced in [HJ24] subsequently to our work). We report the results on Table 1 for two sets of parameters: $D = 8$ (fast signing, larger signatures) and $D = 16$ (slower signing, short signatures). We observe a very significant improvement in runtime, from $12\times$ when $D = 8$ up to $55\times$ when $D = 16$. We note that the authors of [HJ24] already integrated our improved PPRF in the latest version of their scheme, to be presented at CRYPTO'24.

1.2 A new MPCitH signature from regular syndrome decoding

We now turn to our second main contribution, a new construction of MPCitH-based signature from the regular syndrome decoding assumption. This contribution is essentially independent of our multi-instance PPRF⁵; however, it also allows us to formally illustrate on a concrete signature scheme how multi-instance PPRFs yield tight and simple proofs of existential unforgeability.

Our starting point is the recent work of [CCJ23], that relies on a share conversion technique to build a signature scheme from the regular syndrome decoding problem (RSD). The RSD problem was originally introduced in 2003 in [AFS03] as the assumption underlying the FSB candidate to the NIST

⁵ We initially introduced our improved PPRF and its analysis as an optimization of our new scheme, but since it impacts all state-of-the-art MPCitH signatures, we singled it out as an independent contribution.

Table 1. Case analysis of the impact of using our faster AES-based multi-instance PPRF on the signature scheme of [HJ24] for two sets of parameters: $D = 8$ (fast signing) and $D = 16$ (short signatures) compared to the standard hash-based construction. D denotes the depth of the GGM tree (equivalently, 2^D corresponds to the number of virtual parties in the MPC protocol run “in the head”), and τ to the number of repetitions to achieve 128 bits of security. All schemes run on one core of an AMD EPYC 9374F processor clocked at 3.85GHz.

[HJ24]	D	τ	$ \sigma $	signing	verification
hash-based PPRF	8	16	6.2 kB	9.24 ms	9.11 ms
	16	8	4.1 kB	1.1 s	1.1 s
AES-based PPRF (this work)	8	16	6.2 kB	0.80 ms	0.71 ms
	16	8	4.1 kB	19.5 ms	19.2 ms

hash function competition, and subsequently analyzed in [FGS07, MDCYA11, BLPS11], among others (it has also been used and analyzed in many recent works on secure computation, such as [HOSS18, BCGI18, BCG⁺19b, BCG⁺19a, BCG⁺20b, YWL⁺20, WYKW21, RS21, CRR21, BCG⁺22]). It states that given a syndrome $H \cdot x$, where x is a *regular* vector (*i.e.*, a concatenation of unit vectors) and H a random compressive matrix, it is infeasible to recover x . To obtain improved performances compared to [CCJ23], we encode the regular syndrome decoding instances using a sparse representation on top of the dense representation used in [CCJ23]. Encoding regular syndrome decoding instances in a sparse manner is quite natural and relies on the use of an indicator vector to locate the non-zero positions. However, such a representation is not compatible with the secret sharing techniques that are used to split the key between the virtual parties that are introduced by the MPC-in-the-head paradigm: in order to use sparse representations, we need to develop new conversion techniques involving both types of representations. Along the way, we rely on a mechanism to prevent cheating behavior in the conversion, which requires a highly non-trivial combinatorial analysis. Overall, our signature scheme is more than 30% shorter compared to the scheme of [CCJ23] and can use significantly more conservative parameter sets, for similar runtimes.

Results and comparison. We provide a full implementation of our signature scheme. Our implementation is a proof-of-concept implementation, and did not use any optimizations such as batching, vectorization, or bit slicing. In particular, we note that our implementation does not include the folding optimization that was subsequently introduced in [HJ24], that yields significant efficiency improvements when the AES-based multi-instance PPRF is used (in contrast, it only yields mild improvements when using a hash-based PPRF, in essence because it shaves a cost which is dominant in the AES-based PPRF, but dominated by the cost of hashing in the hash-based PPRF). Nevertheless, our implementation confirms that our scheme exhibits excellent performance. Since [CCJ23] does not provide an implementation, we compare our scheme to SDitH, the state-of-the-art MPCitH signature from syndrome decoding [AFG⁺23].

Even when compared to a fully-optimized scheme such as SDitH [AFG⁺23] that makes use of batching techniques and advanced hardware instructions, we observe that our scheme performs particularly well. In addition, we provide tight estimates of the performance improvements that results from integrating the fast folding optimization of [HJ24] to our scheme and to SDitH (while we use the implementation of [HJ24] to obtain runtimes for the faster folding, we note that we do not yet have a full-fledged implementation of our scheme integrating the folding optimization). We leave a fully-optimized implementation of our scheme, integrating the optimization of [HJ24], to future works. We outline below a sample of parameters:

- (fast) signature size 7.8kB, signing time 1.65 ms
- (medium) signature size 6.1 kB, signing time 19.1 ms
- (compact) signature size 5.4 kB, signing time 140 ms.

We refer the reader to Table 3 for more details on our parameters and implementation. We compared our scheme to SDitH [AGH⁺23], the fastest known code-based signature scheme to date, by running both schemes on the same hardware and for comparable parameter sets. To better isolate the effect of our improved PPRF, we also benchmark SDitH with their PPRF replaced by our improved construction,⁶ as well as our scheme using the hash-based PPRF of SDitH. For both, we also measure

⁶ In the conference version of their work, the construction of [AGH⁺23] initially used an *unsalted* GGM tree (instantiated using AES), which we show in Section 3 to be insecure (with a concrete attack that

the runtime improvement that comes from the folding optimization of [HJ24]. We summarize our benchmarks on Table 2. Even when comparing our unoptimized implementation to the optimized implementation of [AFG⁺23], we observe 3× to 4× runtime improvements for $D = 8$ (with signatures of comparable size).

Table 2. Comparison of the new signature scheme with SDitH for $D = 8$ and $D = 12$, with and without our improved multi-instance puncturable pseudorandom function (denoted AES-PPRF and hash-PPRF respectively). AES-PPRF+ denotes the AES-PPRF integrating the folding optimization of [HJ24]. All schemes were run on one core of an Intel Core i7 processor 14700KF.

	D	τ	$ \sigma $	signing time
SDitH (hash-PPRF)	8	17	8.2 kB	6.82 ms
	12	11	6.0 kB	46.8 ms
SDitH (AES-PPRF)	8	17	8.2 kB	6.05 ms
	12	11	6.0 kB	37.9 ms
SDitH (AES-PPRF++) (using [HJ24])	8	17	8.2 kB	6.03 ms
	12	11	6.0 kB	31.5 ms
Our scheme (hash-PPRF)	8	17	7.8 kB	4.07 ms
	12	11	6.1 kB	43.83 ms
Our scheme (AES-PPRF)	8	17	7.8 kB	1.65 ms
	12	11	6.1 kB	19.1 ms
Our scheme (AES-PPRF++) (using [HJ24])	8	17	7.8 kB	1.63 ms
	12	11	6.1 kB	12.8 ms

Another advantage of our signature scheme is its simplicity: while [AGH⁺23] requires fast polynomial operations over large fields, our signature uses only very simple operations on strings such as XORs and cyclic shifts. Eventually, we note that our work and [AGH⁺23] use incomparable variants of syndrome decoding: we use regular syndrome decoding over \mathbb{F}_2 , while [AGH⁺23] uses syndrome decoding over \mathbb{F}_{256} . Both variants have received much less attention than the standard syndrome decoding assumption over \mathbb{F}_2 (though we note that RSD over \mathbb{F}_2 seems to have received significantly more attention than the variant of [AGH⁺23] over the past two decades).

Concurrent work. A concurrent and independent work [CLY⁺24] recently introduced another signature scheme based on the Regular Syndrome Decoding assumption. On a technical level, our approaches differ significantly: [CLY⁺24] combines the VOLE-in-the-Head technique from [BBdSG⁺23] with a sketching method of [BGI16] to reduce the check of the noise structure to a system of degree-2 equations, which are then proven using the Quicksilver VOLE-based zero-knowledge proof [YSWW21]. We use the MPC-in-the-Head methodology with a dedicated share-conversion technique. The signatures of [CLY⁺24] are shorter than ours, e.g., 4 kB versus 5.4 kB for comparable runtimes. Since our techniques are incomparable, we nevertheless expect that they could prove useful in future improved constructions of RSD-based signature, and leave as future work the question of exploring whether our combinatorial techniques could be used to further improve the scheme of [CLY⁺24]. We note that our improved PPRF can be used as a drop-in replacement for the one used in [CLY⁺24] (though it uses VOLE-in-the-Head, the methodology still relies on a similar use of a GGM-style PPRF under the hood) and its use should improve the performances of [CLY⁺24].

1.3 Organization

We introduce some preliminaries in Section 2, and provide a technical overview of our main two contributions in Section 3 (the improved GGM construction) and Section 4 (the new signature scheme) respectively. We then formally introduce and prove the security of our main results, in Section 5 and Section 6. This order of the presentation guarantees a better flow of the explanation because the security analysis of the signature scheme (and the statement of the theorem) relies on our new notion

breaks the scheme using 2^{40} signatures in time 2^{69}). The authors later fixed this issue in their NIST submission [AFG⁺23], using a proper salted GGM tree instantiated with a hash function.

of multi-instance secure PPRF, which we, therefore, introduce first. Section 7 explains how to select parameters for our signature scheme, which requires some careful counting arguments.

2 Preliminaries

Given a set S , we write $s \leftarrow_r S$ to indicate that s is uniformly sampled from S . Given a probabilistic Turing machine \mathcal{A} and an input x , we write $y \leftarrow_r \mathcal{A}(x)$ to indicate that y is sampled by running \mathcal{A} on x with a uniform random tape, or $y \leftarrow \mathcal{A}(x; r)$ when we want to make the random coins explicit. Given an integer $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\}$. We use $\lambda = 128$ for the computational security parameter.

Vectors and matrices. We use bold lowercase for vectors and uppercase for matrices. We write $A||B$ to denote the horizontal concatenation of matrices A, B , and $A//B$ to denote their vertical concatenation. We denote by Id_n the $n \times n$ identity matrix. By default, we always view vectors as columns. Given a vector \mathbf{v} , we will often write $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_n)$ to indicate that \mathbf{v} is a (vertical) concatenation of n subvectors \mathbf{v}_i . We use this slight abuse of notation to avoid the (more precise, but cumbersome) notation $\mathbf{v} = (\mathbf{v}_1^\top, \dots, \mathbf{v}_n^\top)^\top$. Given $\mathbf{u}, \mathbf{v} \in \{0, 1\}^n$, we write $\mathbf{u} \oplus \mathbf{v}$ for the bitwise-XOR of \mathbf{u} and \mathbf{v} , and $\text{HW}(\mathbf{u})$ for the Hamming weight of \mathbf{u} , i.e., its number of nonzero entries.

Permutations. We let $\text{Perm}(w)$ denote the set of all permutations of $[w]$. In this work, we typically use permutations over $[w]$ to shuffle the entries of a length- w vector, or even to shuffle the *blocks* of a vector which is the concatenation of w blocks. For example, given a vector $\mathbf{v} \in [\text{bs}]^w$ and a permutation $\pi : [w] \mapsto [w]$, we write $\pi(\mathbf{v})$ to denote the vector $(v_{\pi(1)}, v_{\pi(2)}, \dots, v_{\pi(w)})$. Given a vector \mathbf{v} which is the concatenation of w subvectors $(\mathbf{v}_1, \dots, \mathbf{v}_w)$, we write $\pi(\mathbf{v})$ to denote the vector $(\mathbf{v}_{\pi(1)}, \dots, \mathbf{v}_{\pi(w)})$. We will typically apply this to vectors over \mathbb{F}_2^K , seen as the concatenation of w vectors over \mathbb{F}_2^{bs} .

Code parameters. In this work, K always denotes the number of columns in the parity-check matrix H , and k denote the number of its rows. Equivalently, K is the codeword length, and $K - k$ is the dimension of the code. We let w denote the weight of the noise, which will always divide K . We let $\text{bs} \leftarrow K/w$ denote the block size: a w -regular noise vector is sampled as a concatenation of w random unit vectors (the *blocks*) of length bs . We write Reg_w to denote the set of all length- K w -regular vectors.

Compact and expanded forms. Given an index $i \in [n]$, we let $\mathbf{e}_i \in \mathbb{F}_2^n$ denote the length- n unit \mathbb{F}_2 -vector whose i -th entry is 1. Given w indices $(i_1, \dots, i_w) \in [n]^w$, we extend the previous notation to $\mathbf{e}_i = (\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_w})$, the concatenation of w unit vectors. We typically manipulate noise vectors represented in *compact form*, i.e., as elements (i_1, \dots, i_w) of $[\text{bs}]^w$, where each entry $i_j \in [\text{bs}]$ indicates the position of the 1 in the j -th length- bs unit vector. We let Expand denote the mapping which, given a noise vector $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_w) \in [\text{bs}]^w$, outputs the vector $\mathbf{e}_\mathbf{x} = (\mathbf{e}_{\mathbf{x}_1}, \dots, \mathbf{e}_{\mathbf{x}_w}) \in \mathbb{F}_2^K$. We call $\mathbf{e}_\mathbf{x}$ the *expanded form* of \mathbf{x} .

Cyclic shifts. Given a vector $\mathbf{u} \in \mathbb{F}_2^n$ and $i \in [n]$, we write $\mathbf{u} \downarrow i$ to denote the vector \mathbf{u} cyclically shifted by i steps (in other words, $\mathbf{u} \downarrow i$ is the convolution of \mathbf{u} and \mathbf{e}_i). We also use notation $\text{Shift}(\mathbf{u}, i)$ to denote $\mathbf{u} \downarrow i$. We extend the notation to a *block-by-block* cyclic shift of vectors: given a vector $\mathbf{u} \in \mathbb{F}_2^K$, viewed as a sequence of blocks $(\mathbf{u}_1, \dots, \mathbf{u}_w) \in \mathbb{F}_2^{K/w} \times \dots \times \mathbb{F}_2^{K/w}$, and a vector of shifts $\mathbf{x} \in [\text{bs}]^w$, we write $\mathbf{u} \downarrow \mathbf{x}$ to denote the vector obtained by shifting the blocks of \mathbf{u} according to \mathbf{x} . That is $\mathbf{u} \downarrow \mathbf{x} = (\mathbf{v}_1, \dots, \mathbf{v}_w)$ where each \mathbf{v}_i is the vector obtained by cyclically shifting (downward) the vector \mathbf{u}_i by \mathbf{x}_i steps. To avoid abusing parenthesis, we view \downarrow as a ‘‘top priority’’ operator: by default, for any other operation op , $\mathbf{u} \downarrow \mathbf{x} \text{ op } \mathbf{v}$ means $(\mathbf{u} \downarrow \mathbf{x}) \text{ op } \mathbf{v}$ and not $\mathbf{u} \downarrow (\mathbf{x} \text{ op } \mathbf{v})$.

Binary tree. Given a tree of size 2^D , for each leaf $i \in [2^D]$, we define $\text{CoPath}(i)$ as co-path to i in the tree, i.e., the set of intermediate nodes that can be used to recover all leaves except the i -th one. Denote bit-decompose i as $\sum_{j=1}^D 2^{j-1} \cdot i_j$ for $i_j \in \{0, 1\}$, the associated value of i -th leaf is defined as $X_i := X_{i_1, \dots, i_D}$.

2.1 Basic Cryptographic Definitions

We cover a few additional standard preliminaries. Recall that two distributions X, Y are (t, ϵ) -indistinguishable if for an algorithm $D : \{0, 1\}^m \rightarrow \{0, 1\}$ running in time t , we have $|\Pr[D(X) = 1] - \Pr[D(Y) = 1]| \leq \epsilon$.

Definition 1 ((t, ϵ)-secure PRG). Let $G : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and let $l(\cdot)$ be a polynomial such that for any input $s \in \{0, 1\}^\lambda$ we have $G(s) \in \{0, 1\}^{l(\lambda)}$. Then, G is a (t, ϵ) -secure pseudorandom generator if

- Expansion: $l(\lambda) > \lambda$;
- $\{G(s) | s \leftarrow \{0, 1\}^\lambda\}$ and $\{r | r \leftarrow \{0, 1\}^{l(\lambda)}\}$ are (t, ϵ) -indistinguishable.

Definition 2 (Collision-Resistant Hash Functions). A family of functions $\text{Hash}_k : \{0, 1\}^* \rightarrow \{0, 1\}^{l(\lambda)}$; $k \in \{0, 1\}^{\kappa(\lambda)}$ indexed by a security parameter λ is collision-resistant if there exists a negligible function v such that, for any PPT algorithm \mathcal{A} , we have:

$$\Pr \left[\begin{array}{c} x \neq x' \\ \cap \text{Hash}_k(x) = \text{Hash}_k(x') \end{array} \middle| \begin{array}{c} k \in \{0, 1\}^{\kappa(\lambda)} \\ (x, x') \leftarrow \mathcal{A}(k) \end{array} \right] \leq v(\lambda)$$

2.2 Regular Syndrome Decoding Problem

Given a weight parameter w , the syndrome decoding problem asks to find a solution of Hamming weight w (under the promise that it exists) to a random system of linear equations $H \cdot \mathbf{x}$ over \mathbb{F}_2 . There exist several well-established variants of the syndrome decoding problem, with different matrix distributions, underlying fields, or noise distributions. In this work, we focus on a relatively well-studied variant known as the *regular syndrome decoding* (RSD) problem, introduced in 2003 in [AFS03] as the assumption underlying the FSB candidate to the NIST hash function competition. In RSD, the solution \mathbf{x} is sampled randomly from the set Reg_w of w -regular vectors (*i.e.*, \mathbf{x} is a concatenation of w unit vectors of length K/w). This variant has been used (and analyzed) quite often in the literature [AFS03, FGS07, MDCYA11, BLPS11, HOSS18, BCGI18, BCG⁺19b, BCG⁺19a, BCG⁺20b, YWL⁺20, WYKW21, RS21, CRR21, BCG⁺22].

Definition 3 (Regular Syndrome Decoding Problem). Let K, k, w be three integers, with $K > k > w$. The syndrome decoding problem with parameters (K, k, w) is defined as follows:

- (Problem generation) Sample $H \leftarrow_r \mathbb{F}_2^{k \times K}$ and $\mathbf{x} \leftarrow_r \text{Reg}_w$. Set $\mathbf{y} \leftarrow H \cdot \mathbf{x}$. Output (H, \mathbf{y}) .
- (Goal) Given (H, \mathbf{y}) , find $\mathbf{x} \in \text{Reg}_w$ such that $H \cdot \mathbf{x} = \mathbf{y}$.

2.3 The MPC-in-the-Head Paradigm

The MPC-in-the-head paradigm was initiated by the work of Ishai et al [IKOS07] and provided a compiler that can build honest-verifier zero-knowledge (HVZK) proofs for arbitrary circuits from secure MPC protocols. Assume we have an MPC protocol with the following properties:

- N parties (P_1, \dots, P_N) securely and jointly evaluate a function $f : \{0, 1\}^* \rightarrow \{0, 1\}$ on \mathbf{x} while each party possess an additive share $\llbracket \mathbf{x} \rrbracket_i$ of input \mathbf{x} ,
- Secure against passive corruption of $N - 1$ parties *i.e.* any $(N - 1)$ parties can not recover any information about the secret \mathbf{x} .

Then the HVZK proof of knowledge of \mathbf{x} such that $f(\mathbf{x}) = 1$ is constructed as:

- Prover generates the additively shares of the witness \mathbf{x} into $(\llbracket \mathbf{x}_1 \rrbracket, \dots, \llbracket \mathbf{x}_N \rrbracket)$ among N virtual parties (P_1, \dots, P_N) and emulate the MPC protocol "in-the-head".
- Prover commits to the view of each party and sends commitments to the verifier.
- Verifier chooses randomly $(N - 1)$ parties and asks the prover to reveal the view of these parties except one. Verifier later accepts if all the views are consistent with an honest execution of MPC protocol with output 1 and agrees with the commitments.

Security of MPC protocol implies that the verifier learns nothing about the input \mathbf{x} from the $N - 1$ shares, and MPC correctness guarantees that the Prover can only cheat with probability $1/N$. Security can then be amplified with parallel repetitions.

3 Technical Overview: Optimized GGM Trees for Faster MPCitH Signatures

Our starting point is the GGM puncturable pseudorandom function [KPTZ13, BW13, BGI14, GGM86], which is used in all modern MPC-in-the-head signatures. At a high level, all MPC-in-the-head protocols start by letting the prover generate shares of the witness, possibly together with shares of some appropriate preprocessing material, to be distributed among the n virtual parties. Then, in the last round, the prover will reveal $n-1$ out of n shares to the verifier. Since each share appears random, the share of each party P_i can be locally stretched from a short seed seed_i . To maintain correctness, an auxiliary “correction word” aux_n is appended to the seed seed_n of the last party P_n (e.g. to guarantee that the stretched shares XOR to the correct witness).

Puncturable PRFs allow us to significantly optimize this step. A *puncturable pseudorandom function* (PPRF) is a PRF F such that given an input x , and a PRF key k , one can generate a *punctured* key, denoted $k\{x\}$, which allows evaluating F at every point except for x , and does not reveal any information about the value $F.\text{Eval}(k, x)$. Using a PPRF, the prover can define all seeds seed_i as outputs of the PRF, using a master seed seed^* as the PRF key. Then, revealing the key seed^* punctured at a point i suffices to succinctly reveal all seeds $(\text{seed}_j)_{j \neq i}$ while hiding seed_i . Concretely, using the GGM PPRF [GGM86], the prover generates n seeds $\text{seed}_1, \dots, \text{seed}_n$ as the leaves of a binary tree of depth $\lceil \log_2 n \rceil$, where the two children of each node are computed using length-doubling pseudorandom generators. This way, revealing all seeds except seed_i requires only sending the seeds on the nodes along the co-path from the root to the i -th leaf, which reduces the communication from $\lambda \cdot (n-1)$ to $\lambda \cdot \lceil \log_2 n \rceil$.

3.1 On the use of salt to avoid collisions

As shown in [DN19], MPC-in-the-head can suffer from collision attacks if the GGM PPRF is used *as is*: after about $2^{\lambda/2}$ signature queries, the adversary is likely to observe two signatures computed with the same master seed seed^* , an event which leaks the secret signing key. To circumvent this issue, previous works have relied in one way or another on a random 2λ -bit salt. However, the use of salt within the GGM PPRF is inconsistent across existing works. As a result, some constructions are either poorly specified or even insecure. Specifically:

- In Banquet [BDK⁺21], and in the more recent work of [AGH⁺23], the seeds $(\text{seed}_1, \dots, \text{seed}_n)$ are generated from an *unsalted* GGM PPRF, and the salt is only used at the leaves, when stretching the share of each party P_i from its seed as $\text{PRG}(\text{seed}_i, \text{salt})$.
- In [FJR22] and [CCJ23], the signature description loosely states that $(\text{seed}_1, \dots, \text{seed}_n)$ are generated in a tree-based fashion using the master seed seed^* and the salt salt . However, the way the salt is used within the GGM construction is not specified precisely. In particular, the definition of the GGM tree in [FJR22] does not include the salt, and their implementation results only mention that it has been implemented “using AES in counter mode”. The work of [CCJ23] does not have an implementation.

We observe that using the salt only at the leaves, as in [BDK⁺21, AGH⁺23], does not shield the signature from collision attacks. The attack is relatively simple:

- The attacker queries m signatures. Each signature will contain some number τ of $\lceil \log_2 n \rceil$ -tuples of intermediate PRG evaluations (corresponding to the seeds on co-path to the unopened leaf; τ corresponds to the number of repetitions of the underlying identification scheme). Let $(\text{seed}^1, \dots, \text{seed}^k)$ denote all seeds received this way, where $k = m \cdot \tau \cdot \lceil \log_2 n \rceil$.
- The attacker locally samples random seeds seed and evaluates its two children $(\text{seed}_0, \text{seed}_1) \leftarrow_r \text{PRG}(\text{seed})$, until one of the seed_b collides with some seed^i .
- Since it knows the preimage of seed_b , it recovers the parent seed of seed^i , from which it can compute the seed associated with the unopened leaf in one of the signatures.
- Given this seed, and using the salt salt associated to the signature (which is public), the attacker reconstructs all virtual parties’ shares, and reconstructs the secret witness (the AES secret key in [BDK⁺21], the syndrome decoding solution in [AGH⁺23]). Using the witness, the attacker can now forge arbitrary signatures.

As should be clear from the above description, we note that adding salt to the leaves has absolutely *no* effect on the security of the signature against this collision attack. Efficiency-wise, after receiving m signatures, the attacker finds a collision in time $2^\lambda / (m \cdot \tau \cdot \lceil \log_2 n \rceil)$. For example, using $\lambda = 128$, $n = 2^{16}$, and $\tau = 9$ (this is a parameter set from [AGH⁺23]), after seeing only $m = 2^{40}$ signatures, the attacker can break the scheme in time $\approx 2^{69}$.

3.2 On the efficiency of salted GGM trees

We believe that the attack pointed above is mostly an issue of the presentation in the respective papers, and that the authors are generally aware of this issue. For example, we observe that the implementation of Banquet⁷ correctly fixes the issue, by adding salt within all intermediate computations of the tree. As for [AGH⁺23], while their original implementation suffers from the attack above, the authors recently included their scheme in a NIST submission [AFG⁺23] whose implementation properly deals with this issue. However, the state of affairs remains quite unsatisfying on the *efficiency* front: unsalted GGM trees can be instantiated very efficiently using fixed-key AES [GKWY20], which enables the use of Intel’s AES-NI instruction set. Unfortunately, the fixed block size of AES makes it hard to add salt. And indeed, existing implementations such as Picnic [CDG⁺20], BBQ [DDOS19], Banquet [BDK⁺21], and the recent NIST submissions based on [AGH⁺23], all implement the PRG using a hash function (such as SHAKE), as follows: $\text{seed}_b \leftarrow H(\text{seed} \| i \| j \| b \| \text{salt})$, where i is the index of the parent node, and $j \leq \tau$ is a counter for the repetitions of the identification scheme. Unfortunately, because of the hardware support for AES, replacing AES with a hash function is up to $50\times$ slower. This is especially problematic in recent protocols that use the hypercube technique [AGH⁺23], where the cost of generating the tree dominates the signing time.

3.3 A fast salted GGM tree in the ideal cipher model

We now turn to our contribution: we introduce a new construction of salted GGM tree which matches the efficiency of the fastest *unsalted* GGM trees, but which yields much stronger security guarantees. We provide formal security proof that our new construction is a *multi-instance secure* PPRF, a notion that we introduce. Multi-instance PPRFs can be used as a drop-in replacement for PPRFs in MPCitH signatures. In contrast with standard PPRFs, whose use incurs a security loss proportional to the number of signature queries (as illustrated by our attack), the unforgeability of MPCitH signatures tightly reduces to the multi-instance security of the PPRF.

In essence, our multi-instance PPRF is based on a very simple idea: use the previous top-performing GGM construction from a fixed-key block-cipher, and *use the cipher key as the salt*. While the intuition is very natural, formally proving security is actually quite challenging. Our full proof of security, in the ideal cipher model, is one of the core technical contributions of this work. It relies on the H-coefficient technique of Patarin [Pat09, CS14] and combines it with a balls-and-bins analysis to measure the number of seed and cipher key collisions, and tightly estimate their impact on security.

Starting Point: a PPRF in the Random Permutation Model. Our starting point is a PPRF construction from [GKWY20]. The construction of [GKWY20] is a tweak on the original GGM construction, where the PRG is instantiated with the following “Davies-Meyer” function:

$$G : x \rightarrow (\pi_0(x) \oplus x, \pi_1(x) \oplus x).$$

In this construction, (π_0, π_1) are two fixed pseudorandom permutations. Using this PRG, the construction of PPRF proceeds in a tree-based fashion: sample a PPRF key $K \leftarrow_r \{0, 1\}^\lambda$. On input $x = (x_1, \dots, x_n)$, the PPRF F_K returns $G_{x_n}(G_{x_{n-1}}(\dots G_{x_1}(K)\dots))$, where G_0, G_1 denote the left and right half of the output of G , respectively. Puncturing x is done by computing all values on the co-path to x in the tree, *i.e.*, the values $G_{\bar{x}_i}(G_{x_{i-1}}(\dots G_{x_1}(K)\dots))$ for $i = 2$ to n : knowing the values on the co-path allows reconstructing the entire tree except for $F_K(x)$, whose values are pseudorandom under the security of G . To prove the security of the construction, the authors of [GKWY20] rely on the *random permutation model*, where (π_0, π_1) are modeled as two independent random permutations.

In [GKWY20], the motivation for introducing the construction is that in practice, π_0, π_1 can be instantiated using the AES block cipher with two fixed keys (K_0, K_1) . This allows to evaluate G

⁷ <https://github.com/dkales/banquet>

using two calls to AES, which is extremely fast using the AES-NI hardware instruction set (encrypting with AES using AES-NI takes as little as *1.3 cycle per Byte* according to [MSY21]). Furthermore, the entire construction requires only two executions of the AES key schedule. This GGM construction is to date, by a significant margin, the fastest known PPRF, and it has been featured extensively in recent works on functions secret sharing [GI14, BGI15, BGI16, BGI19, BCG⁺21], pseudorandom correlation generators [BCGI18, BCG⁺19b, BCG⁺19a, BCG⁺20b, YWL⁺20, WYKW21, CRR21, BCG⁺22], and many more. It is also the construction suggested in [AGH⁺23], though as we saw above it is insecure in the context of signatures.

Observing that this fast PPRF construction is typically instantiated using a block cipher suggests the following idea, which is very natural in retrospect: use the above construction, but instantiate (π_0, π_1) using a block cipher (such as AES) and *use the block cipher keys (K_0, K_1) as a random salt*. This means that in each instance, the pair (K_0, K_1) will be sampled at random. When using AES, this changes nothing to the efficiency of the construction, since in each instance, one still only has to execute the AES key schedule twice. Yet, now, there is some hope that the use of fresh cipher keys in distinct instances can prevent the collision attack.

Multi-instance PPRF and PRGs. To formalize this idea, we introduce the primitive of multi-instance PPRF. At a high level, we define an N -instance PPRF as a PPRF that additionally takes as input a random salt. In the N -instance security game, N keys (k_1, \dots, k_N) , inputs (x_1, \dots, x_N) , and salts $(\text{salt}_1, \dots, \text{salt}_N)$ are sampled randomly. The game also samples a bit $b \leftarrow_{\tau} \{0, 1\}$. Then, the adversary receives $((x_1, \text{salt}_1), \dots, (x_N, \text{salt}_N))$ and the N punctured keys $(k_1 \setminus \{x_1\}, \dots, k_N \setminus \{x_N\})$. If $b = 0$, the adversary additionally receives $(F_k(x_1, \text{salt}_1), \dots, F_k(x_N, \text{salt}_N))$; else, if $b = 1$, the adversary receives N random outputs (y_1, \dots, y_N) instead. The adversary outputs a guess b' and wins if $b' = b$. The PPRF is said to be N -instance (t, ε) -secure if the advantage of any t -time adversary in this game is at most ε . Since our constructions use τ parallel calls to a PPRF with the same salt, we generalize the notion to (N, τ) -instance security to capture the setting where N instances of τ repetitions of a PPRF are used, where the salt differ across instances, but not across internal repetitions.

As a first step toward proving the security of our construction, we also introduce the similar (but simpler) notion of (N, τ) -instance (t, ε) -secure PRG, which is a PRG $G : (\text{seed}, \text{salt}) \rightarrow (G_0(\text{seed}, \text{salt}), G_1(\text{seed}, \text{salt}))$ that additionally takes some random salt. In the N -instance security game, the adversary attempts to distinguish $(G_0(\text{seed}_i, \text{salt}_i), G_1(\text{seed}_i, \text{salt}_i))_{i \leq N}$ from random given the salts $(\text{salt}_1, \dots, \text{salt}_N)$ (the game extends to (N, τ) -instance security in a straightforward way, but the description is more tedious). We show that the standard GGM reduction extends to the multi-input setting: an (N, τ) -input (t, ε) -secure PRG implies an (N, τ) -input $(t, D \cdot \varepsilon)$ -secure PPRF on input domain $[2^D]$ via a straightforward sequence of hybrids.

A multi-instance PRG in the ideal cipher model. The crux of the analysis is then to show that our PRG is indeed (N, τ) -instance (t, ε) -secure (for a suitable choice of N, τ, t, ε). Since the PRG now explicitly uses a block cipher, we cannot rely on the random permutation model anymore; instead, we prove security in the *ideal cipher model*, where each key $K \in \{0, 1\}^\lambda$ defines a truly random permutation π_K , and all parties are given oracle access to π_K and π_K^{-1} for all K (we measure the running time t of the attacker as its number of queries q to the oracles). Using the H -coefficient technique of Patarin, we formally prove that our construction is an (N, τ) -instance (q, ε) -secure PRG for any N up to $2^{\lambda-1}$, with $\varepsilon \leq \frac{4\tau \cdot \lambda}{\ln \lambda} \cdot \frac{q}{2^\lambda}$, where the term $4\tau \lambda / \ln \lambda$ can be replaced by 8τ when $N \leq 2^{\lambda/2}$ (the above inequality is an approximation, see Theorem 9 for the formal inequality). Our analysis is non-trivial, and the bound stems from a careful analysis of the influence of the number of collisions among seeds on the adversarial advantage. We show that this number can be bounded using standard lemmas on the maximum load of a bin when $2N$ balls are thrown into 2^λ bins.

Concretely, this means that one can use our new multi-instance PPRF construction as a drop-in replacement for previous (much slower) hash-based construction, at the (small) cost of a security loss of $4\tau D \lambda / \ln \lambda$ (or simply $8\tau D$ when we bound the number of signature queries by $2^{\lambda/2}$). For $D = 16$, $\tau = 8$, and $\lambda = 128$, this translates to a loss of 14 bits of security (when the number of queries is up to 2^{127}) or 10 bits of security (for up to 2^{64} queries). Furthermore, we can reduce this loss to 8 bits at the (mild) cost of only guaranteeing that the *expected* runtime of the adversary is above 2^λ .

Additionally, we introduce an optimization that converts (N, τ) -instance (t, ε) -secure PRG to $(\tau \cdot N, 1)$ -instance (t, ε) -secure by using a pseudorandom generator to sample the τ salts $(\text{salt}^{i,e})_{e \leq \tau}$ in a given instance from a global salt salt_i for each $i \leq N$. This shaves a factor τ from security loss, which

is reduced to 5 bits for $D = 16, \tau = 8, \lambda = 128$. We believe that this is a very reasonable tradeoff in exchange for the benefits of using a much faster AES-based construction. Eventually, we suggest a final optimization that further reduces the security loss to 3 bits (independently of D): using a pseudorandom generator to generate $(\tau \cdot D)$ salts $(\text{salt}^{i,e})_{e \leq \tau, i \leq D}$ from a global salt salt , and evaluating each *level* of each GGM tree with a different salt. Now all salts are sampled randomly, it leads to collisions among $(\text{salt}^{i,e}, \text{seed}^{i,e})_{e \leq \tau, i \leq N}$ happening with a negligible probability. We conjecture that this variant can be proven secure with only 1 bit of loss in the ideal cipher model. Under the hood, we expect the proof of this conjecture to be similar to the proof of our multi-instances PRG based on the H -coefficient technique, however, it requires a considerably more cumbersome direct analysis of the full multi-instance PPRF in the ideal cipher model (without reducing it first to a multi-instance PRG, which is a much simpler object). We leave proving this last conjecture for future work.

4 Technical Overview: New Signature from RSD

We now move to our second main contribution, a new signature scheme from the regular syndrome decoding assumption. We start with a brief high-level overview of the RSD-based signature scheme from [CCJ23], since it serves as a starting point for our scheme. Let $H \in \mathbb{F}_2^{k \times K}$ be a matrix and $\mathbf{x} \in \mathbb{F}_2^K$ be a w -regular vector (*i.e.*, a concatenation of w unit vectors). We let $\text{bs} \leftarrow K/w$ denote the *block size* of \mathbf{x} . The signature builds upon an efficient n -party protocol which, on input shares of \mathbf{x} , checks that (1) \mathbf{x} is a regular vector, and (2) $H \cdot \mathbf{x} = \mathbf{y}$. This n -party protocol is then compiled into a zero-knowledge proof via the MPC-in-the-head paradigm (which we sometimes abbreviate MPCitH), and the proof is further compiled into a signature scheme via Fiat-Shamir. The main idea underlying the protocol of [CCJ23] is that each of (1) and (2) above can be checked very efficiently, provided that the parties are given a suitable sharing of \mathbf{x} in each case:

- Given (entry-wise) shares of \mathbf{x} over \mathbb{Z}_{bs} , checking that a block of coordinates $x_1, \dots, x_{\text{bs}}$ has weight 1 boils down to checking that $\sum_{i=1}^{\text{bs}} x_i = 1 \pmod{\text{bs}}$, which is a linear equation over \mathbb{Z}_{bs} .
- Given shares of \mathbf{x} over \mathbb{F}_2 , checking $H \cdot \mathbf{x} = \mathbf{y}$ simply amounts to checking a linear equation over \mathbb{F}_2 .

Since in the MPC-in-the-head paradigm, checking linear equations is for free, the task of building the protocol reduces to the task of designing a *sharing conversion* protocol, which converts \mathbb{F}_2 -shares of \mathbf{x} into \mathbb{Z}_{bs} -shares. The next observation of [CCJ23] is that converting shares mod-2 of some value x into shares mod- bs can be done very efficiently given precomputed shares mod-2 and mod- bs of the same random bit r , which the prover can generate by themselves. The only missing ingredient is a mechanism to ensure that the prover honestly computes mod-2 / mod- bs pairs of the *same* identical random bit. The last, and most involved, observation of [CCJ23] is that the verifier can completely dispense with the need to perform this check, by picking a random permutation π of $[K]$ and instructing the prover to shuffle the pairs according to π before running the protocol. Using a careful and non-trivial combinatorial analysis, [CCJ23] showed that whenever \mathbf{x} is sufficiently far from being a regular vector (meaning that it has many non-unit blocks), a malicious prover using \mathbf{x} has negligible success probability over the choice of π , even if they use incorrect mod-2 / mod- bs pairs. Of course, this does not prevent a malicious prover from using an incorrect but close-to-regular witness. However, by choosing the parameters (K, k, w) in a highly injective setting it can be guaranteed that the *only* close-to-regular solution to $H \cdot \mathbf{x} = \mathbf{y}$ is a regular vector.

4.1 An Alternative Share-Conversion Approach

The approach of [CCJ23] yields a competitive signature scheme, but has its shortcomings. Its main efficiency bottleneck stems from the use of shares over \mathbb{Z}_{bs} : because of that, the signature includes several (one for each of the τ repetitions of the basic proof) length- K vectors over \mathbb{Z}_{bs} (using a CRT trick, this can be reduced to $\mathbb{Z}_{\text{bs}/2}$ whenever $\text{bs}/2$ is odd and ≥ 3). This yields a $O(K \cdot \text{bs})$ communication cost, which is (by a significant margin) the dominant cost of their protocol. To mitigate this cost, the authors set the block size bs to be the smallest possible value $\text{bs} = 6$ (such that $\text{bs}/2 = 3$). In turn, this forces them to rely on RSD with very high weight $w = K/6$, which requires significantly increasing the parameters to compensate for the security loss.

Our first observation is that all of these shortcomings can be eliminated at once by relying on an alternative share conversion approach. Because \mathbf{x} is w -regular, it admits a compressed representation

as a list of w integers in $[\text{bs}]$, which indicates the position of the nonzero entry in each of the w unit vectors. Now, observe that if the parties hold shares of w integers (i_1, \dots, i_w) modulo bs , these can always be interpreted as representing some regular vector \mathbf{x} ; in other words, given such shares, condition (1) is satisfied by default. The crux of our protocol is a conversion procedure that turns shares of this compressed representation into shares modulo 2 of the “decompressed” regular vector (with which the parties can check the linear equation $H \cdot \mathbf{x} = \mathbf{y}$ for free). Furthermore, this share conversion can again be implemented very efficiently if the parties are given shares of pairs of the same random unit vector in compressed representation and in standard representation. Concretely, given an integer $r \in [\text{bs}]$, let \mathbf{e}_r denote the length- bs unit vector with a 1 at position r . Assume that the n parties, holding shares of some $i \in [\text{bs}]$, are given shares of r modulo bs , and shares of \mathbf{e}_r over \mathbb{F}_2 . Consider the following simple protocol:

- All parties broadcast their shares of $z = i - r \bmod \text{bs}$ and reconstruct z .
- All parties locally shift cyclically their share of \mathbf{e}_r by z .

After this protocol, all parties end up with shares of the vector \mathbf{e}_r shifted by z , which we denote $\mathbf{e}_r \downarrow z$ (we view vectors as columns, hence the shift by z is downward). Observe that $\mathbf{e}_r \downarrow z = \mathbf{e}_r \downarrow (i - r) = (\mathbf{e}_r \uparrow r) \downarrow i = \mathbf{e}_{\text{bs} \downarrow i} = \mathbf{e}_i$. As in [CCJ23], we will let the prover generate w random pairs (r, \mathbf{e}_r) and share them between the virtual parties. To dispense with the need to check that the pairs were honestly generated, we rely on the same strategy and let the verifier sample a random permutation π of $[w]$, and instruct the prover to shuffle the pairs according to π before using them in the protocol. The high-level structure of the MPCitH-compiled zero-knowledge proof (without optimizations) is below:

- **Parameters and input:** let (K, k, w) be parameters for the syndrome decoding problem, and let $\text{bs} \leftarrow K/w$. The prover holds a w -regular witness $\mathbf{x} \in [\text{bs}]^w$ (in compressed representation) for the relation $H \cdot \mathbf{x} = \mathbf{y}$, where $H \in \mathbb{F}_2^{k \times K}$ and $\mathbf{y} \in \mathbb{F}_2^k$ are public. Let n be the number of virtual parties.
- **Round 1:** the prover samples w pairs (r_i, \mathbf{e}_{r_i}) where $r_i \leftarrow_r [\text{bs}]$. We denote $(\mathbf{r}, \mathbf{e}_{\mathbf{r}})$ the vector of pairs. The prover generates n shares of $\mathbf{e}_{\mathbf{r}}$ (over \mathbb{F}_2) and of \mathbf{x}, \mathbf{r} (modulo bs) distributed between the virtual parties, and commits to the local state of each party.
- **Round 2:** the verifier samples and sends to the prover a random permutation $\pi \leftarrow_r \text{Perm}(w)$. We write $\pi(\mathbf{r})$ (resp. $\pi(\mathbf{e}_{\mathbf{r}})$) for the vector $(r_{\pi(1)}, \dots, r_{\pi(w)})$ (resp. $(\mathbf{e}_{r_{\pi(1)}}, \dots, \mathbf{e}_{r_{\pi(w)}})$).
- **Round 3:** the prover runs in their head the following protocol and commits to the views of all parties:
 - All parties reconstruct $\mathbf{z} = \mathbf{x} - \pi(\mathbf{r})$ and shift their shares of $\pi(\mathbf{e}_{\mathbf{r}})$, getting shares of $\pi(\mathbf{e}_{\mathbf{r}}) \downarrow \mathbf{z}$ (the shifting is done blockwise: each $\mathbf{e}_{r_{\pi(i)}}$ is cyclically shifted by z_i). Note that $\pi(\mathbf{e}_{\mathbf{r}}) \downarrow \mathbf{z} = \mathbf{e}_{\mathbf{x}}$ (*i.e.* the “uncompressed” representation of the witness \mathbf{x}).
 - All parties compute a share of $H \cdot (\pi(\mathbf{e}_{\mathbf{r}}) \downarrow \mathbf{z})$ and broadcast them. All parties check that the shares reconstruct to \mathbf{y} .
- **Round 4:** the verifier picks $i \leftarrow_r [n]$ and challenges the prover to open the views of all parties except i .
- **Round 5:** the prover sends the $n - 1$ openings to the verifier, who checks that the views are consistent with the commitments, with each other, and with the output of the protocol being \mathbf{y} .

The soundness of the scheme is $\varepsilon = \mathbf{p} + (1/n) \cdot (1 - \mathbf{p})$, where $\mathbf{p} = \mathbf{p}(K, k, w)$ is an upper bound on the probability (over the choice of the random permutation π) that a cheating prover, that commits in the first round to an incorrect witness (*i.e.* a compressed vector \mathbf{x}^* such that $H \cdot \mathbf{e}_{\mathbf{x}^*} \neq \mathbf{y}$), manages to generate a valid MPC transcript (*i.e.* finds —possibly incorrect— pairs (\mathbf{r}, \mathbf{u}) such that $H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{z}) = \mathbf{y}$, where $\mathbf{z} = \mathbf{x}^* - \pi(\mathbf{r})$). The crux of our analysis lies in computing a tight evaluation of \mathbf{p} .

In our final signature, we incorporate multiple optimizations on top of this basic template, including the usual optimization of generating the shares in a tree-based fashion using the GGM puncturable pseudorandom function [KPTZ13, BW13, BGI14, GGM86], but also the more recent hypercube technique from [AGH⁺23], and a number of additional optimizations tailored to our scheme.

In terms of signature size, the dominant cost stems from the size of a share of \mathbf{x} and of w pairs (r, \mathbf{e}_r) (using standard optimizations, all shares except one can be compressed, hence the communication is dominated by the size of a single share, ignoring for now the number of repetitions of the identification

scheme). The size of a share of \mathbf{x} together with w pairs (r, \mathbf{e}_r) is $2w \log \text{bs} + K$ bits⁸, whereas the size of \mathbf{x} (now shared as a vector over \mathbb{F}_2^K) and of the pairs in [CCJ23] is $K \cdot (2 + \text{bs}/2)$ bits. This directly incurs a significant reduction in the signature size. Furthermore, with this alternative conversion, using a very small block size is not advantageous anymore, which allows us to explore a much wider range of parameters, resulting in further savings.

4.2 Combinatorial Analysis

Although the high-level strategy —shuffling the random pairs— is the same as in [CCJ23], the security analysis is *entirely* different and forms a core technical contribution of our work. Shuffling the prover-generated correlated randomness is a highly non-generic technique, where each new protocol requires a new and dedicated combinatorial analysis.⁹ The crux of the proof lies in bounding the success probability of a cheating adversary \mathcal{A} in the following game:

- \mathcal{A} holds a vector $\mathbf{x}^* \in [\text{bs}]^w$ and chooses $\mathbf{r} \in [\text{bs}]^w$ and $\mathbf{u} \in \mathbb{F}_2^K$, such that \mathbf{u} is *not* a regular vector.
- A uniformly random permutation π is sampled from $\text{Perm}(w)$.
- \mathcal{A} wins iff $H \cdot (\pi(\mathbf{u}) \downarrow (\mathbf{x}^* - \pi(\mathbf{r}) \bmod \text{bs})) = \mathbf{y}$.

Given a bound on \mathcal{A} 's winning probability in this game, the rest of the proof follows in a relatively standard way and is similar to previous security proofs of code-based signatures schemes in the MPCitH paradigm, such as [CCJ23] (we still provide a full proof in the paper for completeness). Above, note that for any vector $\mathbf{s} \in [\text{bs}]^w$, $\pi(\mathbf{u}) \downarrow \mathbf{s}$ is a regular vector if and only if \mathbf{u} is a regular vector. Note also that whether \mathbf{x}^* is actually a correct witness or not (*i.e.* whether $H \cdot \mathbf{e}_{\mathbf{x}^*}$) does not matter: as long as \mathbf{u} is regular, if \mathcal{A} wins the game above, then an extractor can recover a valid regular solution $\pi(\mathbf{u}) \downarrow (\mathbf{x}^* + \mathbf{r} \bmod \text{bs})$ to the syndrome decoding problem (hence \mathcal{A} “knew” a solution to the problem in the first place). Eventually, note that

$$\pi(\mathbf{u}) \downarrow (\mathbf{x}^* - \pi(\mathbf{r}) \bmod \text{bs}) = \pi(\mathbf{u} \uparrow \mathbf{r}) \downarrow \mathbf{x}^*,$$

hence, the game above simplifies to the following: \mathcal{A} chooses $\mathbf{x}^* \in [\text{bs}]^w$ and $\mathbf{u} \in \mathbb{F}_2^K \setminus \text{Reg}_w$, and wins iff $H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^*) = \mathbf{y}$ holds over the choice of a random permutation π .

Eliminating spurious solutions. An immediate issue with the above game is that an adversary *might* win with a very high probability, if the system of equations $H \cdot \mathbf{x} = \mathbf{y}$ admits solutions that are mostly invariant by blockwise permutation. Concretely, assume that there exists a vector \mathbf{u}^* which satisfies $H \cdot \mathbf{u}^* = \mathbf{y}$, and such that \mathbf{u}^* is not a regular vector, yet \mathbf{v}^* is a concatenation of w *identical* vectors from \mathbb{F}_2^{bs} . If this happens, then there is an easy winning strategy: \mathcal{A} sets $\mathbf{u} \leftarrow \mathbf{u}^*$ and $\mathbf{x}^* \leftarrow 0^w$. Since $H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^*) = H \cdot \pi(\mathbf{u}) = H \cdot \mathbf{u}^* = \mathbf{y}$, \mathcal{A} is guaranteed to win. More generally, if $H \cdot \mathbf{x} = \mathbf{y}$ admits a solution \mathbf{u} whose blocks are *mostly* identical, then the equation $H \cdot \pi(\mathbf{u}^*) = \mathbf{y}$ has a relatively large chance to hold simply because $\pi(\mathbf{u}^*)$ has a relatively large chance to be equal to \mathbf{u}^* .

Setting up some notations. Given a vector \mathbf{u} , we let $\text{pn}(\mathbf{u})$ denote $|\{\pi(\mathbf{u}) \mid \pi \in \text{Perm}([w])\}|$. That is, $\text{pn}(\mathbf{u})$ is the number of distinct vectors in \mathbb{F}_2^K which can be obtained by shuffling \mathbf{u} blockwise; we call $\text{pn}(\mathbf{u})$ the *permutation number* of \mathbf{u} . Then, given a bound B , we define $\text{PN}_B = \{\mathbf{u} \mid \text{pn}(\mathbf{u}) > B\}$, the set of vectors with a large permutation number. We let X denote the set $\{\mathbf{v} \in \mathbb{F}_2^K : \exists \mathbf{u} \in \mathbb{F}_2^K \setminus \text{PN}_B, \exists \mathbf{x}^* \in [\text{bs}]^w, \mathbf{v} = \mathbf{u} \downarrow \mathbf{x}^*\}$. The set X captures exactly the possible spurious solutions: it contains the vectors \mathbf{v} such that there exists some choice of the shift \mathbf{x}^* such that $\mathbf{v} \uparrow \mathbf{x}^*$ has a small permutation number ($\text{pn}(\mathbf{v} \uparrow \mathbf{x}^*) \leq B$). Denoting $\text{Ker}(H) \oplus \mathbf{y}$ the solutions to $H \cdot \mathbf{x} = \mathbf{y}$, if there is a vector $\mathbf{v} \in X \cap \text{Ker}(H) \oplus \mathbf{y}$, then \mathcal{A} can pick \mathbf{u}, \mathbf{x}^* such that $\mathbf{v} = \mathbf{u} \downarrow \mathbf{x}^*$ with $\text{pn}(\mathbf{u}) \leq B$. This guarantees that with probability at least $1/B$, a random permutation π will satisfy $\pi(\mathbf{u}) = \mathbf{u}$, hence $H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^*) = H \cdot (\mathbf{u} \downarrow \mathbf{x}^*) = H \cdot \mathbf{v} = \mathbf{y}$.

⁸ As in [CCJ23], this number is multiplied by a number τ of repetition, but since it is the same in both works, we ignore it in this discussion for simplicity.

⁹ To give a sense of how specific the analysis of [CCJ23] was, not only does it work only for their type of pairs: it works exclusively for $\text{bs} = 6$, corresponding to pairs of bits shared modulo 2 and modulo 3.

Sampling highly-injective instances. Fix some bound B . To eliminate spurious solutions in X , which an adversary could use to win with probability at least $1/B$, we choose parameters (K, k, w) such that when sampling the regular syndrome decoding instance $(H, \mathbf{y} = H \cdot \mathbf{x})$ (for some $\mathbf{x} \in \text{Reg}_w$), it holds with probability $1 - 1/2^\lambda$, the *only* element of X that also belongs to $\text{Ker}(H) \oplus \mathbf{y}$ is the w -regular solution \mathbf{x} . It follows from a standard analysis that this is the case as soon as $\log_2 k \geq \log_2 |X| + \lambda$. To select k , we therefore compute a tight upper bound on $|X|$ (see Lemma 24). Counting the number of elements of X is not entirely straightforward due to the fact that we count “up to some blockwise shift”, but a closed formula can be established using known bounds for counting k -necklaces (*i.e.* bitstrings counted up to cyclic shifts) by leveraging Pólya’s enumeration theorem [Red27]. Given the formula, we use a short Python program to compute explicitly the bound on $|X|$ and select a suitable parameter k (for a fixed choice of K, w). This also faces some challenges: the formula of Lemma 24 requires summing binomial coefficients over all *integer partitions* of the weight parameter w (*i.e.*, the number of tuples of distinct positive integers that sum to w). Because w is around 120, its number of integer partitions is too large to simply enumerate. With some careful considerations, we observe that many of these partitions can be eliminated from the counting procedure and leverage this observation to reduce the runtime of the program.

Bounding the success probability. We now turn to the crux of the analysis: showing that if \mathcal{A} picks $(\mathbf{u}, \mathbf{x}^*)$ where $\text{pn}(\mathbf{u}) > B$, then their probability of winning the game is at most $O(1/B)$ over the choice of the permutation π . What makes the analysis challenging is that in principle, it could be that some vector \mathbf{u} has a high permutation number, yet *many of its permutations belong to* $\text{Ker}(H) \oplus \mathbf{y}$. The core technical component of the analysis is a proof that with very high probability over the choice of a random syndrome decoding instance (H, \mathbf{y}) , it will simultaneously hold for *all* vectors \mathbf{u} with $\text{pn}(\mathbf{u}) > B$ that for any choice of shift \mathbf{x}^* , $\Pr_\pi[H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^*) = \mathbf{y}] \leq 4/B$. To state the result formally, we define “good” syndrome decoding instances below:

Definition 4 (GOOD $_B$). Given a bound B , GOOD $_B$ is defined as the set of syndrome decoding instances $(H, \mathbf{y}) \in \mathbb{F}_2^{k \times K} \times \mathbb{F}_2^k$ such that for every $\mathbf{u} \in \text{PN}_B \setminus \text{Reg}_w$ and for all $\mathbf{x}^* \in [\text{bs}]^w$,

$$\Pr_{\pi \leftarrow_r \text{Perm}_w} [H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^*) = \mathbf{y}] \leq 4/B.$$

Our main technical result of the analysis is stated below:

Lemma 5 (Most syndrome decoding instances are good).

$$\Pr_{H, \mathbf{y}} [(H, \mathbf{y}) \in \text{GOOD}_B] > 1 - \binom{2B}{5} \cdot \frac{2^{K+1}}{B \cdot 2^{3k}} \cdot \left(10 + \frac{(K/w)^w}{2^k}\right).$$

To parse the above, the reader can consider that $(K/w)^w \ll 2^k$ will hold for our selection of parameters, hence the probability that $(H, \mathbf{y}) \in \text{GOOD}_B$ is of the order of $1 - B^4 \cdot 2^{K-3k}$. For concreteness, the reader can think of K as being around 1550, k as being around 820, w being around 200, and B as being around 70, resulting in the above being around $1 - 2^{-630}$.

Key intuition. We outline the main idea of the proof. Given a vector \mathbf{u} with $\text{pn}(\mathbf{u}) = N$, fix some ordering $\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)}$ of its distinct blockwise permutations, and let $\mathbf{x}^* \in [\text{bs}]^w$ denote some shift. Sample a random matrix $H \leftarrow_r \mathbb{F}_2^{k \times K}$, a random regular vector $\mathbf{x} \leftarrow_r \text{Reg}_w$, and set $\mathbf{y} \leftarrow H \cdot \mathbf{x}$. Let $(\mathbf{v}_1, \dots, \mathbf{v}_N) \leftarrow ((\mathbf{u}^{(1)} \downarrow \mathbf{x}^*) \oplus \mathbf{x}, \dots, (\mathbf{u}^{(N)} \downarrow \mathbf{x}^*) \oplus \mathbf{x})$ (note that $H \cdot \mathbf{v}_i = 0$ iff $H \cdot (\mathbf{u}^{(i)} \downarrow \mathbf{x}^*) = \mathbf{y}$). Observe that the \mathbf{v}_i are random variables, but they are set independently of H (since \mathbf{x} is sampled independently from H). Then, for any subset S of t linearly independent vectors \mathbf{v}_i , it holds that

$$\Pr_{H \leftarrow_r \mathbb{F}_2^{k \times K}} [H \cdot \mathbf{v}_i = 0 \text{ for all } i \in S] = 2^{-k \cdot t}.$$

In other words, whenever the \mathbf{v}_i ’s are linearly independent, the binary random variables X_i equal to 1 if $H \cdot \mathbf{v}_i = 0$ are independent. Building upon this observation, we show the following: fix an arbitrary subset S of five indices. Then

- S contains a size-3 linearly independent subset with probability 1, and
- S contains a size-4 linearly independent subset, except with probability at most $10 \cdot (K/w)^{-w}$.

Together with the previous bound on the probability that $H \cdot \mathbf{v}_i = 0$ for linearly independent vectors, this yields a probability bound of $10 \cdot (K/w)^{-w}/2^{3 \cdot k} + 1/2^{4 \cdot k}$ that $H \cdot \mathbf{v}_i = 0$ for all $i \in S$. To see why this bound holds, observe that:

- The \mathbf{v}_i are pairwise distinct and nonzero by construction (because \mathbf{u} is assumed to be nonregular, so $\pi(\mathbf{u}) \downarrow \mathbf{x}^*$ is never $\mathbf{0}$, and the $\mathbf{u}^{(i)}$ are distinct by definition).
- If e.g. $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ are linearly dependent, they therefore need to satisfy $\mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_3 = \mathbf{0}$. But then, $\mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_4 \neq \mathbf{0}$ (otherwise, we would have $\mathbf{v}_3 = \mathbf{v}_4$, contradicting the fact that the vectors are pairwise distinct). Hence, we are guaranteed to find a size-3 independent subset of vectors in S .
- By the same reasoning, S contains necessarily a 4-tuple of \mathbf{v}_i 's that does not XOR to 0, say, $(\mathbf{v}_1, \dots, \mathbf{v}_4)$ (since if both $(\mathbf{v}_1, \dots, \mathbf{v}_4)$ and $(\mathbf{v}_1, \dots, \mathbf{v}_3, \mathbf{v}_5)$ XOR to 0, then $\mathbf{v}_4 = \mathbf{v}_5$). Then, either $(\mathbf{v}_1, \dots, \mathbf{v}_4)$ is linearly independent (in which case we are done, since we found a 4-independent subset), or it must contain a size-3 subset that XORs to 0.
- For any subset of 3 \mathbf{v}_i 's, the probability that they XOR to 0 is at most $(K/w)^{-w}$. This follows from the fact that the \mathbf{v}_i 's are equal to $(\mathbf{a} \oplus \mathbf{x}, \mathbf{b} \oplus \mathbf{x}, \mathbf{c} \oplus \mathbf{x})$ for some *fixed* vectors $(\mathbf{a}, \mathbf{b}, \mathbf{c})$, and a uniformly random regular vector $\mathbf{x} \in [\mathbf{bs}]^w$. But then, $\mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_3 = \mathbf{0}$ rewrites to $\mathbf{a} \oplus \mathbf{b} \oplus \mathbf{c} = \mathbf{x}$, which happens with probability at most $\mathbf{bs}^{-w} = (K/w)^{-w}$ over the choice of \mathbf{x} .

Since there are 10 size-3 subsets of S , the bound follows. To summarize, we fixed a vector \mathbf{u} with $\text{pn}(\mathbf{u}) = N > B$ and a shift \mathbf{x}^* , and showed that for every size-5 subset S of $[N]$, the probability that $H \cdot (\mathbf{u}^{(i)} \downarrow \mathbf{x}^*) = \mathbf{y}$ holds simultaneously for all $i \in S$ is at most $10 \cdot (K/w)^{-w}/2^{3 \cdot k} + 1/2^{4 \cdot k}$.

A careful union bound. To finish the proof of Lemma 5, it remains to compute a union bound over all possible vectors \mathbf{u} , shifts \mathbf{x}^* , and size-5 subsets S . However, a quick calculation shows that a naive union bound does not suffice: first, the number of subsets is $\binom{N}{5}$, but since we only know that $N > B$ is the permutation number of \mathbf{u} , we can only bound it by $w!$, which is way too large. Second, the number of vectors \mathbf{u} is 2^K , which is also too large for the union bound to yield a nontrivial result.

We overcome this issue by providing a more careful union bound. First, we divide the distinct blockwise permutations of \mathbf{u} , $(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)})$, into size- B blocks of vectors. We apply the previous bound to all size-5 subsets inside each *block* of vectors, which reduces the factor resulting from the union bound to $(N/B) \cdot \binom{B}{5}$. This suffices to guarantee that in each size- B block, at most 4 vectors \mathbf{v}_i can simultaneously satisfy $H \cdot \mathbf{v}_i = 0$, hence guaranteeing a success probability for \mathcal{A} of at most $4/B$ over the random choice of π . Second, instead of enumerating over all vectors \mathbf{u} , we enumerate over all *equivalence classes* of vectors \mathbf{u} which generate the same list $(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)})$. Each equivalence class contains exactly N vectors, and all equivalence classes are disjoint, and we save a factor N this way from the union bound. Eventually, we finish the union bound by summing over all possible values of $N = \text{pn}(\mathbf{u})$ from $B + 1$ to $w!$. This finishes the proof of Lemma 5.

5 Multi-Instance PPRFs in the Ideal Cipher Model

In this section, we introduce the notion of *multi-instance* puncturable pseudorandom function. We describe an efficient construction from a block cipher, and formally prove its security in the ideal cipher model.

5.1 Defining Multi-Instance Puncturable PRF

Pseudorandom functions [GGM86], are families of keyed functions F_k such that no adversary can distinguish between a black-box access to F_k for a random key k and access to a truly random function. A puncturable pseudorandom function (PPRF) [KPTZ13, BW13, BGI14] is a PRF F such that given an input x , and a PRF key k , one can generate a *punctured* key, denoted $k\{x\} = F.\text{Punc}(K, x)$, which allows evaluating F at every point except for x (i.e., there is an algorithm $F.\text{Eval}$ such that $F.\text{Eval}(k\{x\}, x') = F_K(x')$ for all $x' \neq x$), and such that $F_k(x)$ is indistinguishable from random given $k\{x\}$. Then,

Definition 6 ((N, τ) -instance (t, ϵ) -secure PPRF). *A function family $F = \{F_K\}$ with input domain $[2^D]$, salt domain $\{0, 1\}^s$, and output domain $\{0, 1\}^\lambda$, is an (N, τ) -instance (t, ϵ) -secure PPRF*

if it is a PPRF which additionally takes as input a salt salt , and for every non-uniform PPT distinguisher \mathcal{D} running in time at most t , it holds that for all sufficiently large λ ,

$$\text{Adv}^{\text{PPRF}}(\mathcal{D}) = |\Pr[\text{Exp}_{\mathcal{D}}^{\text{rw-pprf}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{D}}^{\text{iw-pprf}}(\lambda) = 1]| \leq \epsilon(\lambda)$$

where the experiments $\text{Exp}_{\mathcal{D}}^{\text{rw-pprf}}(\lambda)$ and $\text{Exp}_{\mathcal{D}}^{\text{iw-pprf}}(\lambda)$ are defined below.

$\text{Exp}_{\mathcal{D}}^{\text{rw-pprf}}(\lambda) :$ <ul style="list-style-type: none"> - $((K_{j,e})_{j \leq N, e \leq \tau} \leftarrow_r (\{0, 1\}^\lambda)^{N \cdot \tau}$ - $\text{salt} := (\text{salt}_1, \dots, \text{salt}_N) \leftarrow_r \{0, 1\}^s$ - $\mathbf{i} := ((i_{1,e})_{e \leq \tau}, \dots, (i_{N,e})_{e \leq \tau}) \leftarrow_r [2^D]^{N \cdot \tau}$ - $\forall j \leq N, e \leq \tau : K_{j,e}^{i_{j,e}} \leftarrow F.\text{Punc}(K_{j,e}, i_{j,e})$ - $(y_{j,e})_{j \leq N, e \leq \tau} \leftarrow (F_{K_{j,e}}^{i_{j,e}}(\text{salt}_j))_{j \leq N, e \leq \tau}$ <p>Output $b \leftarrow \mathcal{D}(\text{salt}, \mathbf{i}, (K_{j,e}^{i_{j,e}}, y_{j,e})_{j \leq N, e \leq \tau})$</p>	$\text{Exp}_{\mathcal{D}}^{\text{iw-pprf}}(\lambda) :$ <ul style="list-style-type: none"> - $((K_{j,e})_{j \leq N, e \leq \tau} \leftarrow_r (\{0, 1\}^\lambda)^{N \cdot \tau}$ - $\text{salt} := (\text{salt}_1, \dots, \text{salt}_N) \leftarrow_r \{0, 1\}^s$ - $\mathbf{i} := ((i_{1,e})_{e \leq \tau}, \dots, (i_{N,e})_{e \leq \tau}) \leftarrow_r [2^D]^{N \cdot \tau}$ - $\forall j \leq N, e \leq \tau : K_{j,e}^{i_{j,e}} \leftarrow F.\text{Punc}(K_{j,e}, i_{j,e})$ - $(y_{j,e})_{j \leq N, e \leq \tau} \leftarrow_r (\{0, 1\}^\lambda)^{N \cdot \tau}$ <p>Output $b \leftarrow \mathcal{D}(\text{salt}, \mathbf{i}, (K_{j,e}^{i_{j,e}}, y_{j,e})_{j \leq N, e \leq \tau})$</p>
---	--

The motivation for adding the parameter τ in Definition 6 stems from our use of PPRFs in signatures: our signature construction uses τ parallel instances of the PPRF using the same salt, while distinct salts are used across distinct signature queries.

Furthermore, we observe our actual construction satisfies a stronger property, in which indistinguishability is preserved even the ideal world experiment does not only sample (y_1, \dots, y_N) uniformly at random, but also samples “fake” punctured keys $K_j^{x_k}$ uniformly at random over an appropriate domain. This stronger notion is not strictly necessary in our signature construction, but its use simplifies the analysis. Below, we state the definition explicitly for the punctured key domain that corresponds to our (GGM-based) construction, but the notion extends naturally to arbitrary domains.

Definition 7 ((N, τ)-instance strongly (t, ϵ)-secure PPRF). A function family $F = \{F_K\}$ with input domain $[2^D]$, salt domain $\{0, 1\}^s$, output domain $\{0, 1\}^\lambda$, and punctured key domain $(\{0, 1\}^\lambda)^D$ is an (N, τ)-instance (t, ϵ)-secure PPRF if it is a PPRF which additionally takes as input a salt salt , and for every non-uniform PPT distinguisher \mathcal{D} running in time at most t , it holds that for all sufficiently large λ ,

$$\text{Adv}^{\text{PPRF}}(\mathcal{D}) = |\Pr[\text{Exp}_{\mathcal{D}}^{\text{rw-spprf}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{D}}^{\text{iw-spprf}}(\lambda) = 1]| \leq \epsilon(\lambda),$$

where the experiment $\text{Exp}_{\mathcal{D}}^{\text{iw-spprf}}(\lambda)$ is defined as $\text{Exp}_{\mathcal{D}}^{\text{iw-pprf}}(\lambda)$, except that the line $\forall j \leq N, e \leq \tau : K_{j,e}^{i_{j,e}} \leftarrow F.\text{Punc}(K_{j,e}, i_{j,e})$ is replaced by $\forall j \leq N, e \leq \tau : K_{j,e}^{i_{j,e}} \leftarrow_r (\{0, 1\}^\lambda)^D$.

5.2 Constructing Multi-Instance Puncturable PRFs

In this section, we introduce the notion of (N, τ)-instance (t, ϵ)-secure pseudorandom generator, which extends the notion of pseudorandom generators to the multi-instance setting (with salt) analogously to our definition of multi-instance PPRFs. Then, we show that the standard GGM construction extends immediately to the multi-instance setting: (length-doubling) (N, τ)-instance (t, ϵ)-secure PRGs imply (N, τ)-instance strongly ($t, D \cdot \epsilon$)-secure PPRFs with input domain $[2^D]$ and punctured key domain $(\{0, 1\}^\lambda)^D$. We start by defining (N, τ)-instance (t, ϵ)-secure length-doubling PRGs. Below, to interface more easily with the tree-based GGM construction of PPRFs, we use (F_0, F_1) to denote functions that compute the left half and right half of the length-doubling PRG output.

Definition 8 ((N, τ)-instance (t, ϵ)-secure PRG). A PRG $\text{PRG} = (F_0, F_1)$ with $F_b : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ is an (N, τ)-instance (t, ϵ)-secure length-doubling PRG if for every non-uniform PPT distinguisher \mathcal{D} running in time at most t , it holds that for all sufficiently large λ ,

$$\text{Adv}^{\text{PRG}}(\mathcal{D}) = |\Pr[\text{Exp}_{\mathcal{D}}^{\text{rw-prg}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{D}}^{\text{iw-prg}}(\lambda) = 1]| \leq \epsilon(\lambda),$$

where $\text{Exp}_{\mathcal{D}}^{\text{rw-prg}}(\lambda)$ and $\text{Exp}_{\mathcal{D}}^{\text{iw-prg}}(\lambda)$ are defined below.

$\text{Exp}_{\mathcal{D}}^{\text{rw-prg}}(\lambda) :$ <ul style="list-style-type: none"> - $(\text{salt}_1, \text{salt}_2, \dots, \text{salt}_{2N}) \leftarrow_r \{0, 1\}^\lambda$ - $(\text{seed}_{i,e})_{i \leq N, e \leq \tau} \leftarrow_r (\{0, 1\}^\lambda)^{N \cdot \tau}$ - $\forall i \leq N, e \leq \tau :$ <ul style="list-style-type: none"> • $y_{2i-1,e} \leftarrow F_0(\text{seed}_{i,e}, \text{salt}_{2i-1})$ • $y_{2i,e} \leftarrow F_1(\text{seed}_{i,e}, \text{salt}_{2i})$ <p>Output $b \leftarrow \mathcal{D}((\text{salt}_i, (y_{i,e})_{e \leq \tau})_{i \leq 2N})$</p>	$\text{Exp}_{\mathcal{D}}^{\text{iw-prg}}(\lambda) :$ <ul style="list-style-type: none"> - $(\text{salt}_1, \text{salt}_2, \dots, \text{salt}_{2N}) \leftarrow_r \{0, 1\}^\lambda$ - $(y_{i,e})_{i \leq 2N, e \leq \tau} \leftarrow_r (\{0, 1\}^\lambda)^{2N \cdot \tau}$ <p>Output $b \leftarrow \mathcal{D}((\text{salt}_i, (y_{i,e})_{e \leq \tau})_{i \leq 2N})$</p>
--	--

We note that the definition extends immediately to PRGs that stretch their seeds by a larger factor. We also remark that in the definition above, we assumed that each of F_0 and F_1 takes a distinct λ -bit salt. The definition can be extended to more general salting procedures, but we defined multi-instance PRG with respect to the way we use salt in our actual construction for notational convenience. Looking ahead, the fact that each F_b takes only λ bits of salt is actually a crucial byproduct of our use of block ciphers, and the main reason why the security analysis becomes highly non-trivial.

Now, given a seed $\text{seed} \leftarrow_r \{0, 1\}^\lambda$, salt $\text{salt} := (\text{salt}_0, \text{salt}_1) \leftarrow_r \{0, 1\}^{2\lambda}$, and a multi-instance secure PRG $F_0, F_1 : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$, we recursively define a PPRF $\text{PPRF}(\text{seed}, \text{salt}) = \text{PPRF}(\text{seed}, \text{salt}, 2^D)$ over input domain $\{0, 1\}^D$ (which we later identify with $[2^D]$) in a tree-based fashion as follows:

- The first layer includes two nodes $X_0 := F_0(\text{seed}, \text{salt}_0)$, $X_1 := F_1(\text{seed}, \text{salt}_1)$.
- Each layer of the tree is constructed from the nodes of the previous layer similarly, as follows:

$$\begin{aligned} \text{PPRF}_{\text{seed}}(\text{salt}, i) &= F_{i_D}(\text{PPRF}_{\text{seed}}(\text{salt}, i_1, \dots, i_{D-1}), \text{salt}) \\ &= F_{i_D}(F_{i_{D-1}}(\dots(F_{i_1}(\text{seed}, \text{salt}), \text{salt}), \text{salt}), \text{salt}), \end{aligned}$$

where i_1, \dots, i_D denote the bits of i .

As with the standard GGM construction, a punctured key at i is just the co-path to i in the tree, *i.e.*, the set of intermediate nodes that can be used to recover all leaves except the i -th one: $\text{CoPath}_{\text{seed}}(\text{salt}, i) = \text{PPRF}_{\text{seed}}(\text{salt}, i_1, \dots, \bar{j})_{j=1, \dots, D}$. The formal construction is presented in Figure 1 and the proof of security is shown in Theorem 9. We note that the proof is a natural extension of the security analysis of the GGM construction [GGM86].

Theorem 9 (PPRF security). *Assume that $\text{PRG} = (F_0, F_1)$ with $F_b : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ is an (N, τ) -instance (t, ϵ) -secure length-doubling PRG. Then the construction $\text{PPRF}(\text{seed}, \text{salt}, 2^D)$ described in Figure 1 is an (N, τ) -instance strongly $(t, D \cdot \epsilon)$ -secure PPRF with input domain $[2^D]$ and punctured key domain $(\{0, 1\}^\lambda)^D$.*

Proof. We proceed in a sequence of hybrids where each hybrid relies on the (N, τ) -instance security of F_0, F_1 .

First, recall that for each leaf $i^{(j,e)} \in \{0, 1\}^D$ in each tree $\text{PPRF}(\text{seed}_{j,e}, \text{salt}_j, 2^D)$, the value assigned to this leaf $i^{(j,e)}$ is denoted $X_{i_1^{(j,e)}, \dots, i_D^{(j,e)}}$. The secret path from the root $(\text{seed}_{j,e}, \text{salt}_j)$ to the leaf $i^{(j,e)}$ is the tuple of intermediate nodes $\{X_{i_1^{(j,e)}}, X_{i_1^{(j,e)}, i_2^{(j,e)}}, \dots, X_{i_1^{(j,e)}, \dots, i_D^{(j,e)}}\}$.

- **Experiment 0 (Exp⁰).** All trees of the N instances are obtained through the actual scheme described in Figure 1, which is run at each level to generate the leaves of the next level. More in detail: for each $j \leq N, e \leq \tau$ the construction of the (j, e) -th tree is carried out starting from a random master $(\text{seed}_{j,e}, \text{salt}_j)$ and using, for all 2^D levels, F_0 and F_1 to generate the right child and the left child.

Experiment 1 (Exp¹). Same as the previous experiment, the only difference is at the first level of each tree. For all $j = 1, \dots, N, e \leq \tau$, the leaves at the first level $(X_{1^{(j,e)}}, X_{0^{(j,e)}})$ are not generated with F_0, F_1 , but are instead randomly sampled. Since F_0, F_1 is an (N, τ) -instance (t, ϵ) -secure PRG, then

$$|\Pr[\text{Exp}^0(\lambda) = 1] - \Pr[\text{Exp}^1(\lambda) = 1]| \leq \epsilon(\lambda)$$

Experiment 2 (Exp²). The difference with the previous experiment is in the second level of each tree: all the leaves $(X_{i_1^{(j,e)}, 0}, X_{i_1^{(j,e)}, 1})$ previously computed by using F_0 and F_1 are now randomly chosen for each $j = 1, \dots, N, e \leq \tau$. As before, using the secure property of F_0 and F_1 , we obtain

$$|\Pr[\text{Exp}^1(\lambda) = 1] - \Pr[\text{Exp}^2(\lambda) = 1]| \leq \epsilon(\lambda)$$

As it is easy to guess, traversing along the secret path of each tree, this mechanism of replacing the two leaves $(X_{i_1^{(j,e)}, \dots, i_{k-1}^{(j,e)}, 0}, X_{i_1^{(j,e)}, \dots, i_{k-1}^{(j,e)}, 1})$ at each level $k \in [1, D]$ by uniformly random values




Fig. 1. New construction $\text{PPRF}(\text{seed}, \text{salt}, 2^D)$ of Puncturable PRF

can continue for the whole depth D of the tree. This way, we have D experiments and applying the same hypothesis about the security of the F_0 and F_1 used, we will get:

$$|\Pr[\text{Exp}^{i-1}(\lambda) = 1] - \Pr[\text{Exp}^i(\lambda) = 1]| \leq \epsilon(\lambda)$$

for all $i = 1, \dots, D$. Furthermore, when we traverse the path this way, we simultaneously replace all values on the *co-path* to the leaves $i^{(j,e)}$ by uniformly random values.

Experiment D (Exp^D). In the last experiment, all nodes on the co-path to $i^{(j,e)}$ as well as the leaf $i^{(j,e)}$ are picked uniformly at random, for $j = 1$ to N and $e = 1$ to τ . We obtain the final bound

$$|\Pr[\text{Exp}^0(\lambda) = 1] - \Pr[\text{Exp}^D(\lambda) = 1]| \leq D \cdot \epsilon(\lambda),$$

which concludes the proof.

5.3 A Multi-Instance PRG in the Ideal Cipher Model

In this section, we describe the construction of multi-instance PRG in the ideal cipher model. Our construction itself is not really new, but is a tweak on a construction of [GKWY20]. The work of [GKWY20] gives a construction of PPRF in the random permutation model, which is obtained by applying the GGM reduction to the following ‘‘Davies-Meyer’’ construction of a length-doubling PRG $G : x \rightarrow (\pi_0(x) \oplus x, \pi_1(x) \oplus x)$, where (π_0, π_1) are pseudorandom permutations. The PRG is proven secure in the random permutation model (in the analysis, all parties are given oracle access to π_0, π_1 , and their inverses). Our core observation, which is quite simple in hindsight, is that the most efficient instantiation of this construction implements the permutations π_0, π_1 by fixing two keys (K_0, K_1) and defining $\pi_b := E_{K_B}$, where E_{K_B} is a *block cipher* (such as AES). This suggests the following

idea: instead of fixing the keys (K_0, K_1) , sample them randomly and use them as a salt for the PRG in the multi-instance setting. The candidate multi-instance PRG becomes $G = (F_0, F_1) : (x, \text{salt}) \rightarrow (E_{\text{salt}_0}(x) \oplus x, E_{\text{salt}_1}(x) \oplus x)$. The formal construction is given in Figure 2. While the high-level intuition is straightforward, the formal analysis turns out to be considerably more involved. The remainder of this section is devoted to a formal proof that the above construction is an (N, τ) -instance (t, ϵ) -secure PRG, for parameters (N, τ, t, ϵ) which will be specified later. The proof is in the *ideal cipher model*: in this model, each key $K \in \{0, 1\}^\lambda$ defines an independent uniformly random permutation π_K . All parties are given access to an oracle which, on input $(0, K, x)$, outputs $\pi_K(x)$, and on input $(1, K, y)$, outputs $\pi_K^{-1}(y)$. And the proof security is shown in Theorem 11, it relies on a careful analysis using Patarin’s H -coefficient technique [Pat09, CS14] and forms one of the core technical contributions of this work.

Definition 10 (Ideal Cipher Oracle). For every $K \in \{0, 1\}^\lambda$, let $\pi_K : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ be a uniformly random permutation over $\{0, 1\}^\lambda$. The ideal cipher oracle \mathcal{O}_π is defined as follows:

- On input $(x, K) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$, outputs $\pi_K(x)$.
- On input (inv, x, K) , outputs $\pi_K^{-1}(x)$.

<p>Parameters:</p> <ul style="list-style-type: none"> - For each $K \in \{0, 1\}^\lambda$, $\pi_K : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ is a uniformly random permutation. <p>Construction:</p> <ul style="list-style-type: none"> - Sample $\text{salt} \leftarrow_r \{0, 1\}^{2\lambda}$. parse $\text{salt} := (K_0, K_1)$. - $F_b : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\lambda$ is defined as $F_b(\text{seed}, \text{salt}_b) = \pi_{K_b}(\text{seed}) \oplus \text{seed}$ for $b \in \{0, 1\}$ and $\text{seed} \in \{0, 1\}^\lambda$.
--

Fig. 2. Multi-instance PRG F_0, F_1 in the ideal cipher model

Theorem 11. Let F_0, F_1 be the functions defined in Figure 2. Let q be the number of queries to the oracle \mathcal{O}_π . Then (F_0, F_1) is an (N, τ) -instance (q, ϵ) -secure PRG in the ideal cipher model (where the parties are given oracle access to \mathcal{O}_π from Definition 10), where

$$\epsilon \leq f_N(\lambda) \cdot q \cdot \left(\frac{1}{2^{\lambda-1}} + \frac{1}{2^\lambda - q} \right) + \frac{4\tau N}{2^{2\lambda}},$$

for some function f_N such that if $N \leq 2^{\lambda-1}$, $f_N(\lambda) \leq \frac{3\tau\lambda \cdot \ln 2}{\ln \lambda + \ln \ln 2}$, and if $N \leq 2^{\lambda/2}$, $f_N(\lambda) \leq 4\tau$.

Proof. Fix a number of instances N and a number of repetitions τ . We consider a distinguisher \mathcal{D} that receives $(\text{salt}_i, (y_{i,e})_{e \leq \tau})_{i \leq 2N}$ according to either the real world experiment $\text{Exp}_{\mathcal{D}}^{\text{rw-prg}}$ or the ideal world experiment $\text{Exp}_{\mathcal{D}}^{\text{iw-prg}}$ of Definition 8, interacts with the ideal cipher oracle \mathcal{O}_π , and outputs a guess b . Let q be a bound on the number of queries of \mathcal{D} to \mathcal{O}_π . To simplify the discussion, we assume that the $N \cdot \tau$ seeds $(\text{seed}^{(1,e)}, \dots, \text{seed}^{(N,e)})_{e \leq \tau}$ are also sampled (but not used) in the experiment $\text{Exp}_{\mathcal{D}}^{\text{iw-prg}}$. We also write salt_i as (K_0^i, K_1^i) .

Reformulating the experiment. Now, sample $(\text{seed}^{(1,e)}, \dots, \text{seed}^{(N,e)})_{e \leq \tau}$ and pairs of keys $(K_0^i, K_1^i)_{i \leq N}$. If $N \cdot \tau$ is large, with a high probability there will be some collisions among the seeds. Let $M \leq N \cdot \tau$ denote the number of distinct seeds. To simplify the analysis, we reorder and rename the seeds and the keys as follows:

- $\text{seed}_1, \dots, \text{seed}_M$ are the M distinct seeds from the set of $N \cdot \tau$ sampled seeds $\text{seed}^{(j,e)}$. For each seed seed_i , define $S_i \subseteq \{0, 1\}^\lambda$ to be the set of indices such that $K \in S_i$ if there is an index (j, e) such that $\text{seed}^{(j,e)} = \text{seed}_i$ and either $K = K_0^j$ or $K = K_1^j$ (that is, seed_i was sampled at least once together with a salt that contains K). Note that S_i corresponds to all keys K such that π_K is queried on seed_i in $\text{Exp}_{\mathcal{D}}^{\text{rw-prg}}$.

- For each π_K , define $S'_K := \{i : K \in S_i\} \subseteq [M]$ to be the set of indices of seeds that will be queried to π_K .

With the above notations, the distinguisher \mathcal{D} receives the sets S_1, \dots, S_M , and for each $i \leq M$, it gets either $\pi_K(\text{seed}_i) \oplus \text{seed}_i$ for all $K \in S_i$ (experiment $\text{Exp}_{\mathcal{D}}^{\text{rw-prg}}$), or a set of random values $(y_{K,i})_{K \in S_i}$ (experiment $\text{Exp}_{\mathcal{D}}^{\text{iw}}$). These alternative experiments only differ from the original experiments if it happens that two seeds $\text{seed}^{(i,e)}$, $\text{seed}^{(j,f)}$ collide, and two of their keys (K_0^i, K_1^i) and (K_0^j, K_1^j) also collide: in this case, the original experiments would return distinct values y in the ideal world, but identical values in the real world, making them trivially distinguishable. However, the probability of this even happening is very small:

$$\Pr[\exists(i, e) \neq (j, f), \text{seed}^{(i,e)} = \text{seed}^{(j,f)} \wedge \exists(b_i, b_j) \in \{0, 1\}^2, K_{b_i}^i = K_{b_j}^j] \leq \frac{4N \cdot \tau}{2^{2\lambda}}.$$

Condition on this even not happening, the new experiments become perfectly equivalent to the original experiments. We therefore raise a flag if the above condition occurs, abort if a flag is raised, and focus on bounding the distinguishing advantage in these new experiments.

Bounding the size of S'_K . We start by bounding the maximum size of S'_K for any K . We will need a standard lemma on the maximum load of a bin when tossing m balls into n bins:

Lemma 12 (balls-and-bins). *Consider tossing m balls into n bins. For $m \leq n$, denoting max_load as the maximum number of balls that end up in any single bin, we have*

$$\Pr \left[\text{max_load} \geq \frac{3 \ln n}{\ln \ln n} \right] \leq \frac{1}{n}.$$

By definition, the maximum size of S'_K is reached for the permutation π_K that is invoked on the largest number of distinct seeds. A tight upper bound on this number follows from a simple balls-and-bins analysis: each time τ new seeds $(\text{seed}^{(i,e)})_{e \leq \tau}$ are sampled, two keys (K_0^i, K_1^i) are sampled, which we view as throwing two balls to two random bins, sampled randomly from 2^λ possible bins. After N steps of this experiment (hence after throwing $2N$ balls at random), denoting max_load the maximum load of any bin, $\tau \cdot \text{max_load}$ is an upper bound on $\max_K |S'_K|$.¹⁰ We get:

Claim. Whenever $2N \leq 2^\lambda$, the maximum load $\max_K |S'_K|$ is bounded by $\frac{3\tau \cdot \ln 2^\lambda}{\ln \ln 2^\lambda}$ with probability $1 - 2^{-\lambda}$. Furthermore, if $2N \leq 2^{\lambda/2}$, $\max_K |S'_K|$ is bounded by 4τ with probability $1 - 2^{-\lambda}$.

The first part of the claim follows directly from the balls-and-bins lemma 12. The last part of the claim follows from the fact that when $2N \leq 2^{\lambda/2}$, the probability of having 4 balls in any given bin is at most $1/2^{2\lambda}$, and the claim follows by a union bound over the 2^λ bins.

Bounding the advantage of \mathcal{D} . We now move to the crux of the analysis, where we bound the advantage of \mathcal{D} in distinguishing the real world and the ideal world experiments. We formally define below the *transcript* of the interaction of \mathcal{D} in the experiments:

Definition 13 (Transcript). *We define a transcript of \mathcal{D} 's interaction by*

$$Q = ((y_{i,j})_{i \leq M, j \in S_i}, Q_\pi, (\text{seed}_i)_{i \leq M})$$

where $Q_\pi = (z, j, \pi_j(z))$ records all \mathcal{D} 's queries/answers to/from the permutation oracle \mathcal{O}_π (the queries for the inverse of permutation can be considered as $(\pi_b^{-1}(z), b, z)$). Note that here, $(\text{seed}_i)_{i \leq M}$ is included to facilitate the analysis but the distinguisher \mathcal{D} does not get them: in the real-world, $(\text{seed}_i)_{i \leq M}$ are used to compute $(y_{i,j})_{i \leq M, j \in S_i}$ otherwise in the ideal-world, $(y_{i,j})_{i \leq M, j \in S_i}$ are sampled uniformly random from $\{0, 1\}^\lambda$.

We say that a transcript Q is **attainable** for some fixed \mathcal{D} if there exist some oracles \mathcal{O}_π such that the interaction of \mathcal{D} with those oracles would lead to transcript Q .

¹⁰ It is a very tight upper bound: because $|S'_K|$ counts only *distinct* seeds, we are overcounting whenever it happens that a new seed $\text{seed}^{(j,e)}$ is sampled that collides with one of the previous seeds $(\text{seed}^{(j^{ne})} = \text{seed}^{(i,f)})$ for some $i < j$ and K_0^j or K_1^j also collides with one of the two keys (K_0^i, K_1^i) . But the chance that this happens is at most $\tau \cdot N/2^{2\lambda}$.

In the game of distinguishing between the ideal world and the real world, we have

$$\text{Adv}(\mathcal{D}_\pi^{\mathcal{O}}) = |\Pr_{\text{rw}}[\mathcal{D}^{\mathcal{O}\pi} = 1] - \Pr_{\text{iw}}[\mathcal{D}^{\mathcal{O}\pi} = 1]|$$

Our proof will crucially rely on Patarin’s H -coefficient technique [Pat09, CS14], which we recall below. The H -coefficient theorem allows to bound the advantage of distinguisher by classifying the set of attainable transcripts into “good” and “bad” transcripts:

Theorem 14 (H-coefficient). *Fix some distinguisher \mathcal{D} . Let \mathcal{T} denote the set of attainable transcripts Q and \Pr_{rw} and \Pr_{iw} denote the probabilities of events in the real and ideal world, respectively. Let \mathcal{T}_{bad} denote a set of “bad” transcripts, and $\mathcal{T}_{\text{good}} = \mathcal{T} \setminus \mathcal{T}_{\text{bad}}$ be the set of “good” transcripts, suppose that:*

- $\Pr_{\text{iw}}[Q \in \mathcal{T}_{\text{bad}}] \leq \nu$.
- $\left| \frac{\Pr_{\text{rw}}[Q]}{\Pr_{\text{iw}}[Q]} - 1 \right| \leq \mu$ for all $Q \in \mathcal{T}_{\text{good}}$.

Then $\text{Adv}(\mathcal{D}_\pi^{\mathcal{O}}) \leq \nu + \mu$.

One key insight of the H -coefficient technique is that the ratio $\frac{\Pr_{\text{rw}}[Q]}{\Pr_{\text{iw}}[Q]}$ is equal to the ratio between the probability that the real-world oracles are consistent with Q and the probability that the ideal-world oracles are consistent with Q . We denote $\Pr[\text{RW is consistent with } Q]$ and $\Pr[\text{IW is consistent with } Q]$ as $\Pr_{\text{rw}}(Q)$ and $\Pr_{\text{iw}}(Q)$ respectively. Then

$$\forall Q \in \mathcal{T}_{\text{good}}, \quad \frac{\Pr_{\text{rw}}[Q]}{\Pr_{\text{iw}}[Q]} = \frac{\Pr_{\text{rw}}(Q)}{\Pr_{\text{iw}}(Q)}$$

Our goal now is to use the H -coefficient theorem to prove Theorem 11. Define \mathcal{T}_{bad} and $\mathcal{T}_{\text{good}}$ from the sets of distinct seeds and permutations.

- \mathcal{T}_{bad} contains transcripts $Q = ((y_{i,K})_{i \leq M, K \in S_i}, Q_\pi, (\text{seed}_i)_{i \leq M}) \in \mathcal{T}$ s.t.
 - $\exists(\text{seed}_i, K, *) \in Q_\pi$ with $K \in S_i$.
 - $\exists(*, K, \text{seed}_i \oplus y_{i,K})$ with $K \in S_i$.
- $\mathcal{T}_{\text{good}} = \mathcal{T} \setminus \mathcal{T}_{\text{bad}}$.

Bounding $\Pr_{\text{iw}}[Q \in \mathcal{T}_{\text{bad}}]$. Let denote $|Q_\pi| = q = \sum_{K \in \{0,1\}^\lambda} q_K$ where $q_K := |Q_{\pi_K}| := |\{(*, K, *) \in Q_\pi\}|$ for $K \in \{0,1\}^\lambda$. In the ideal-world, $(\text{seed}_i)_{i \leq M}$ are independent of $((y_{i,K})_{i \leq M, K \in S_i})$, and we have:

$$\begin{aligned} \Pr_{\text{iw}}[Q \in \mathcal{T}_{\text{bad}}] &\leq \sum_{K \in \{0,1\}^\lambda} (\Pr_{\text{iw}}[\exists(\text{seed}_i, K, *) \in Q_\pi | i \in S'_K] \\ &\quad + \Pr_{\text{iw}}[\exists(*, K, y_{i,K} \oplus \text{seed}_i) \in Q_\pi | i \in S'_K]) \\ &= \sum_{K \in \{0,1\}^\lambda} \frac{2q_K \cdot |S'_K|}{|2^\lambda|} = \frac{1}{2^{\lambda-1}} \cdot \sum_{K \in \{0,1\}^\lambda} q_K \cdot |S'_K| \\ &\leq \frac{1}{2^{\lambda-1}} \cdot q \cdot \max_K |S'_K|. \end{aligned}$$

Bounding $\Pr_{\text{rw}}[Q]/\Pr_{\text{iw}}[Q]$ for $Q \in \mathcal{T}_{\text{good}}$. First, we compute the probability $\Pr_{\text{iw}}(Q)$ that the ideal-world oracle is consistent with Q . Denote $((\text{seed}'_i)_{i \leq M}, (y'_{i,K})_{i \leq M, K \in S_i}, (\pi_K)_{K \in \{0,1\}^\lambda})$ some arbitrary setting of the ideal world experiment, where $(\text{seed}'_i)_{i \leq M}, (y'_{i,K})_{i \leq M, K \in S_i}$ are sampled as in $\text{Exp}_D^{\text{iw}}(\lambda)$, and $\pi_K : \{0,1\}^\lambda \rightarrow \{0,1\}^\lambda$ are fixed random permutations. Let $\pi \vdash Q_\pi$ denote the event that permutation π is consistent with the queries/answers in Q_π . Let us write $(\pi_K)_K \vdash Q_\pi$ to indicate that random permutations π_k are consistent with all queries in the transcript Q_π . Since in the ideal

world all these values are sampled independently, denoting $p_\pi = \Pr_{\pi_K}[(\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi]$, we have:

$$\begin{aligned}
\Pr_{\text{iw}}(Q) &= \Pr_{\text{seed}'_i, y'_{i,j}, \pi_j} [(\text{seed}'_i = \text{seed}_i) \wedge (\forall K \in S_i, y_{i,K} = y'_{i,K}) \wedge ((\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi)] \\
&= \Pr_{\text{seed}'_i} [\text{seed}'_i = \text{seed}_i] \cdot \Pr_{y'_{i,K}} [\forall K \in S_i, y_{i,K} = y'_{i,K}] \cdot p_\pi \\
&= \left(\frac{1}{2^\lambda}\right)^M \cdot \left(\frac{1}{2^\lambda}\right)^{\sum_{i=1}^M |S_i|} \cdot \prod_{K \in \{0,1\}^\lambda} \frac{1}{(2^\lambda)_{q_K}} \\
&= 2^{-\lambda \cdot \sum_{K \in \{0,1\}^\lambda} |S'_K|} \cdot \prod_{K \in \{0,1\}^\lambda} \frac{1}{(2^\lambda)_{q_K}} \cdot 2^{-\lambda \cdot M},
\end{aligned}$$

where for $1 \leq b \leq a$, $(a)_b := a \cdot (a-1) \cdot (a-2) \cdots (a-b+1)$. Note that the last equality comes from the fact that $\sum_{i=1}^M |S_i| = \sum_{K \in \{0,1\}^\lambda} |S'_K|$.

We next compute the probability $\Pr_{\text{rw}}(Q)$ that the real-world oracle is consistent with Q . We also denote by $((\text{seed}'_i)_{i \leq M}, (y'_{i,K})_{i \leq M, K \in S_i}, (\pi_K)_{K \in \{0,1\}^\lambda})$ a setting of the real world. The main difference is that $(y'_{i,K})_{i \leq M, K \in S_i}$ are now dependent on $(\text{seed}'_i)_{i \leq M}$. Denoting $p_\pi = \Pr_{\pi_K}[(\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi]$, we have

$$\begin{aligned}
\Pr_{\text{rw}}(Q) &= \Pr[(\text{seed}'_i = \text{seed}_i) \wedge (\forall K \in S_i, y_{i,K} = y'_{i,K}) \wedge ((\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi)] \\
&= \Pr_{\text{seed}'_i} [(\text{seed}'_i = \text{seed}_i) \wedge (\forall K \in S_i, y_{i,K} = \pi_K(\text{seed}_i) \oplus \text{seed}_i)] \cdot p_\pi \\
&= \Pr_{\text{seed}'_i} [(\text{seed}'_i = \text{seed}_i) \wedge (\forall i \in S'_K, y_{i,K} = \pi_K(\text{seed}_i) \oplus \text{seed}_i)] \cdot p_\pi \\
&= \Pr_{\pi_K} [y_{i,K} = \pi_K(\text{seed}_i) \oplus \text{seed}_i \mid (\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi] \cdot p_\pi \cdot \Pr_{\text{seed}'_i} [\text{seed}'_i = \text{seed}_i] \\
&= \frac{1}{2^{\lambda \cdot M}} \cdot \prod_{K \in \{0,1\}^\lambda} \frac{1}{(2^\lambda)_{q_K}} \cdot \Pr_{\pi_K} [y_{i,K} = \pi_K(\text{seed}_i) \oplus \text{seed}_i \mid (\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi]
\end{aligned}$$

Since $Q \in \mathcal{T}_{\text{good}}$ then $\nexists(\text{seed}_i, K, *) \in Q_\pi$ with $K \in S_i \wedge \nexists(*, K, \text{seed}_i \oplus y_{i,K})$ with $K \in S_i$. This leads to

$$\begin{aligned}
&\Pr_{\pi_K} [y_{i,K} = \pi_K(\text{seed}_i) \oplus \text{seed}_i \mid (\pi_K)_{K \in \{0,1\}^\lambda} \vdash Q_\pi] \\
&= \prod_{K \in \{0,1\}^\lambda} \Pr_{\pi_K} [\pi_K(\text{seed}_i) = y_{i,K} \oplus \text{seed}_i] = \prod_{K \in \{0,1\}^\lambda} \frac{1}{(2^\lambda - q_K)_{|S'_K|}}.
\end{aligned}$$

Putting equations together, we obtain

$$\Pr_{\text{rw}}(Q) = \frac{1}{2^{\lambda \cdot M}} \cdot \prod_{K \in \{0,1\}^\lambda} \frac{1}{(2^\lambda)_{q_K}} \cdot \prod_{K \in \{0,1\}^\lambda} \frac{1}{(2^\lambda - q_K)_{|S'_K|}}$$

and eventually

$$\begin{aligned}
\forall Q \in \mathcal{T}_{\text{good}}, \frac{\Pr_{\text{rw}}[Q]}{\Pr_{\text{iw}}[Q]} &= \frac{\Pr_{\text{rw}}(Q)}{\Pr_{\text{iw}}(Q)} = \prod_{K \in \{0,1\}^\lambda} \frac{2^{\lambda \cdot \sum_{K \in \{0,1\}^\lambda} |S'_K|}}{(2^\lambda - q_K)_{|S'_K|}} \\
&= \prod_{K \in \{0,1\}^\lambda} \frac{2^{2N \cdot \lambda}}{(2^\lambda - q_K)_{|S'_K|}}.
\end{aligned}$$

Distinguishing advantage. Equipped with the above calculations, we can finally bound the distinguishing advantage of $\mathcal{D}^{\mathcal{O}^\pi}$. To upper bound $\text{Adv}(\mathcal{D}^{\mathcal{O}^\pi})$, we upper bound the ratio $\frac{\text{Pr}_{\text{rw}}[Q]}{\text{Pr}_{\text{iw}}[Q]}$, which translates to computing a lower bound on $\prod_{K \in \{0,1\}^\lambda} (2^\lambda - q_K)^{|S'_K|}$. Denote $K_{\max} \in K \in \{0,1\}^\lambda$ the index of the set among all $\{S'_K\}_{K \in \{0,1\}^\lambda}$ that has $\max_K |S'_K|$ elements. Then we have

$$\begin{aligned} \prod_{K \in \{0,1\}^\lambda} (2^\lambda - q_K)^{|S'_K|} &\geq \prod_{K \neq K_{\max}} (2^\lambda)^{|S'_K|} \cdot (2^\lambda - q)^{\max_K |S'_K|} \\ &= (2^\lambda)^{\sum_{K \in \{0,1\}^\lambda} |S'_K|} \cdot \frac{(2^\lambda - q)^{\max_K |S'_K|}}{(2^\lambda)^{\max_K |S'_K|}} \\ &= 2^{2N \cdot \lambda} \cdot \frac{(2^\lambda - q)^{\max_K |S'_K|}}{(2^\lambda)^{\max_K |S'_K|}} \geq 2^{2N \cdot \lambda} \cdot \left(\frac{2^\lambda - q}{2^\lambda}\right)^{\max_K |S'_K|} \\ \implies \frac{\text{Pr}_{\text{rw}}[Q]}{\text{Pr}_{\text{iw}}[Q]} &\leq \left(\frac{2^\lambda}{2^\lambda - q}\right)^{\max_K |S'_K|} = \left(1 + \frac{q}{2^\lambda - q}\right)^{\max_K |S'_K|}. \end{aligned}$$

The above yields

$$\frac{\text{Pr}_{\text{rw}}[Q]}{\text{Pr}_{\text{iw}}[Q]} \leq 1 + \frac{q \cdot \max_K |S'_K|}{2^\lambda - q}.$$

Then, using the H -coefficient theorem (Theorem 5.3), we get:

$$\text{Adv}(\mathcal{D}^{\mathcal{O}^\pi}) = \frac{1}{2^{\lambda-1}} \cdot q \cdot \max_K |S'_K| + \frac{q \cdot \max_j |S'_K|}{2^\lambda - q}.$$

Plugging the bound on $|S'_K|$ from the claim finishes the proof. \square

6 A Signature scheme from Regular Syndrome Decoding

In this section, we introduce a new signature scheme from the regular syndrome decoding assumption. A signature scheme is given by three algorithms (KeyGen , Sign , Verify). KeyGen returns a key pair (pk, sk) where pk and sk are the public and private key. Sign on an input a message m and the secret key sk , produces a signature σ . Verify , on input a message m , a public key pk and a signature σ , returns 0 or 1. Standard security notions for signature schemes are existential unforgeability against key-only attacks (EUF-KO, Definition 16) and against chosen-message attacks (EUF-CMA, Definition 15).

Definition 15 (EUF-CMA security). *Given a signature scheme $\text{Sig} = (\text{Setup}, \text{Sign}, \text{Verify})$ and security parameter λ , we say that Sig is EUF-CMA-secure if any PPT algorithm \mathcal{A} has negligible advantage in the EUF-CMA game, defined as*

$$\text{Adv}_{\mathcal{A}}^{\text{EUF-CMA}} = \Pr \left[\begin{array}{l} \text{Verify}(\text{pk}, \mu^*, \sigma^*) = 1 \\ \wedge \mu^* \notin Q \end{array} \middle| \begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{Setup}(\{0,1\}^\lambda) \\ (\mu^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(\text{pk}) \end{array} \right],$$

where $\mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}$ denotes \mathcal{A} 's access to a signing oracle with private key sk and Q denotes the set of messages μ that were queried to $\text{Sign}(\text{sk}, \cdot)$ by \mathcal{A} .

Definition 16 (EUF-KO security). *Given a signature scheme $\text{Sig} = (\text{Setup}, \text{Sign}, \text{Verify})$ and security parameter λ , we say that Sig is EUF-KO-secure if any PPT algorithm \mathcal{A} has negligible advantage in the EUF-KO game, defined as*

$$\text{Adv}_{\mathcal{A}}^{\text{EUF-KO}} = \Pr \left[\text{Verify}(\text{pk}, \mu^*, \sigma^*) = 1 \middle| \begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{Setup}(\{0,1\}^\lambda) \\ (\mu^*, \sigma^*) \leftarrow \mathcal{A}(\text{pk}) \end{array} \right].$$

6.1 Description of the Signature Scheme

The key generation algorithm (Figure 3) randomly samples a syndrome decoding instance (H, \mathbf{y}) with solution \mathbf{x} . The signing algorithm with secret key $\text{sk} = (H, \mathbf{y}, \mathbf{x})$ and message $m \in \{0,1\}^*$ is described on Figure 5. The verification algorithm with public key $\text{pk} = (H, \mathbf{y})$ (matrix H can be computed from PRG with a random seed, the public key size is around 0.09kB), message $m \in \{0,1\}^*$, and signature σ , is described in Figure 4.

Inputs: A security parameter λ .

1. Sample $\text{seed} \leftarrow \{0, 1\}^\lambda$;
2. Set $H \leftarrow \text{PRG}(\text{seed})$ with $H \in \mathbb{F}_2^{k \times K}$;
3. Sample $\mathbf{x} \leftarrow_r [\text{bs}]^w$ and set $\mathbf{y} \leftarrow H \cdot \text{Expand}(\mathbf{x})$ and $\text{sk} \leftarrow (\text{seed}, \mathbf{x})$.

Fig. 3. Key generation algorithm of the signature scheme

An optimization. For readability, the description of the signing and verification algorithms ignores an optimization that slightly reduces the signature size, but significantly complexifies the description. Concretely, because we know that the vectors u^e should be regular vectors, it suffices to share the $\text{bs} - 1$ first entries $(u_1, \dots, u_{\text{bs}-1})$ of each block of \mathbf{u}^e , since the last one can be reconstructed as $\bigoplus u_i \oplus 1$. This reduces the size of \mathbf{u} from $K = w \cdot \text{bs}$ to $w \cdot (\text{bs} - 1) = K - w$ bits. Consequently, the share u_n^e of u^e need also only be shared over \mathbb{F}_2^{K-w} . This reduces by w the size of aux_n^e for each $e \leq \tau$, hence overall by $\tau \cdot w$ the size of the signature. An additional byproduct of this optimization is that it reduces the number of possible “cheating” vectors u^e that a malicious prover could choose, which has some positive repercussions on the size of the RSD parameters (K, k, w) which we can choose (we elaborate in Section 7).

Inputs: A public key $\text{pk} = (H, \mathbf{y})$, a message $m \in \{0, 1\}^*$ and a signature σ .

1. Split the signature as follows:

$$\sigma = \left(\text{salt}, h_1, h_2, (\text{CoPath}_{\text{salt}}(i^e, \text{seed}^e), z^e, \text{com}_{i^e}^e, \text{aux}_n^e)_{e \leq \tau} \right);$$

2. Recompute $\pi_{e \in \tau}^e$ where $\pi^e \in \text{Perm}([w])$ via a pseudorandom generator using h_1 ;
3. Recompute $(b_1^e, \dots, b_D^e)_{e \leq \tau}$ via a pseudorandom generator using h_2 and define $i^e \leftarrow \sum_{d=1}^D b_d^e \cdot 2^{d-1}$;
4. For each iteration $e \in [\tau]$,
 - For $d = 1$ to D :
 - Denote $b = 1 - b_d^e$;
 - Set $(X_{d,b}^e, R_{d,b}^e, U_{d,b}^e) \leftarrow (0, 0, 0) \in [\text{bs}]^w \times [\text{bs}]^w \times \{0, 1\}^K$;
 - For each $i \neq i^e$:
 - * Recompute seed_i^e from the $\text{CoPath}_{\text{salt}}(i^e, \text{seed}^e)$;
 - * If $i \neq n$, recompute $(x_i^e, r_i^e, u_i^e, \text{com}_i^e) \leftarrow \text{PRG}(\text{seed}_i^e)$; else, parse aux_n^e as (x_n^e, u_n^e) , and compute $r_n^e \leftarrow \text{PRG}(\text{seed}_n^e)$;
 - * If $i[d] = b$, update:
 - $X_{d,b}^e \leftarrow X_{d,b}^e + x_i^e \bmod \text{bs}$;
 - $R_{d,b}^e \leftarrow R_{d,b}^e + r_i^e \bmod \text{bs}$;
 - $U_{d,b}^e \leftarrow U_{d,b}^e \oplus u_i^e$;
 - Recompute $(y_{d,b}^e, z_{d,b}^e)$ by simulating the Phase 3 of the signing algorithm as below:
 - $y_{d,b}^e \leftarrow H \cdot \text{Shift}(\pi^e(U_{d,b}^e), z^e)$;
 - $z_{d,b}^e \leftarrow X_{d,b}^e - \pi^e(R_{d,b}^e) \bmod \text{bs}$;
 - Recompute $(y_{d,1-b}^e, z_{d,1-b}^e)$ as below:
 - $y_{d,1-b}^e \leftarrow y_{d,b}^e \oplus \mathbf{y}$;
 - $z_{d,1-b}^e \leftarrow z^e - z_{d,b}^e \bmod \text{bs}$;
5. Check if $h_1 \leftarrow H_1(m, \text{salt}, \text{com}_1^1, \dots, \text{com}_1^1, \dots, \text{com}_1^\tau, \dots, \text{com}_1^\tau)$;
6. Check if $h_2 \leftarrow H_2(m, \text{salt}, h_1, (y_{d,b}^e, z_{d,b}^e)_{d \leq D, b \in \{0,1\}, e \leq \tau})$;
7. Output ACCEPT if both conditions are satisfied.

Fig. 4. Verification algorithm of the signature scheme

Theorem 17. Assume that PPRF is a (q_s, τ) -instance $(t, \epsilon_{\text{PPRF}})$ -secure PPRF, that PRG is a (q_s, τ) -instance $(t, \epsilon_{\text{PRG}})$ -secure PRG, and that any adversary running in time t has at advantage at most ϵ_{SD} against the regular syndrome decoding problem. Model the hash functions H_1, H_2 as random oracles with output of length 2λ -bit and the pseudorandom generator PRG_2 as a random oracle. Then chosen-message adversary against the signature scheme depicted in Figure 5, running in time t , making q_s signing queries, and making q_1, q_2, q_3 queries, respectively, to the random oracles H_1, H_2 and PRG_2 , succeeds in outputting a valid forgery with probability

$$\Pr[\text{Forge}] \leq \frac{q_s (q_s + q_1 + q_2 + q_3)}{2^{2\lambda}} + \epsilon_{\text{PPRF}} + \epsilon_{\text{PRG}} + \epsilon_{\text{SD}} + \Pr[X + Y = \tau] + \epsilon_{\text{G}} + \frac{1}{2^\lambda},$$

Inputs: A secret key sk and a message $m \in \{0, 1\}^*$.

Initialization. Parse sk as $(\text{seed}, \mathbf{x})$

- Let $H \leftarrow \text{PRG}(\text{seed})$ and $\mathbf{y} \leftarrow H \cdot \text{Expand}(\mathbf{x})$; // $H \in \mathbb{F}_2^{k \times K}$ is a (pseudo)random matrix in systematic form.
- Sample $(K_0, K_1) \leftarrow_r \{0, 1\}^\lambda \times \{0, 1\}^\lambda$. Set $\text{salt} \leftarrow (K_0, K_1)$.

Phase 1. For each iteration $e \in [\tau]$:

- Sample $\text{seed}^e \leftarrow_r \{0, 1\}^\lambda$;
- For $d = 1$ to D , set $(X_{d,0}^e, R_{d,0}^e, U_{d,0}^e) \leftarrow (0, 0, 0) \in [\text{bs}]^w \times [\text{bs}]^w \times \{0, 1\}^K$;
- Set $x_n^e \leftarrow \mathbf{x}$, $u_n^e \leftarrow 0$, and $r^e \leftarrow 0$;
- **For** $i = 1$ **to** $n - 1$:
 1. Compute $\text{seed}_i^e \leftarrow \text{PPRF}_{\text{salt}}(\text{seed}^e, i)$; // Can be computed efficiently by always storing the path to the current node: to move from i to $i + 1$, start from the closest ancestor of $i + 1$ in the path to leave i .
 2. Set $\text{state}_i^e \leftarrow \text{seed}_i^e$;
 3. $(x_i^e, r_i^e, u_i^e, \text{com}_i^e) \leftarrow \text{PRG}(\text{seed}_i^e)$; // $(x_i^e, r_i^e, u_i^e, \text{com}_i^e) \in [\text{bs}]^w \times [\text{bs}]^w \times \{0, 1\}^K \times \{0, 1\}^\lambda$.
 4. $x_n^e \leftarrow x_n^e - x_i^e \bmod \text{bs}$, $u_n^e \leftarrow u_n^e \oplus u_i^e$, and $r^e \leftarrow r^e + r_i^e \bmod \text{bs}$;
 5. For all $d \leq D$ such that $i[d] = 0$, set: // $i[d]$ is the d -th bit of the integer i .
 - $X_{d,0}^e \leftarrow X_{d,0}^e + x_i^e \bmod \text{bs}$;
 - $R_{d,0}^e \leftarrow R_{d,0}^e + r_i^e \bmod \text{bs}$;
 - $U_{d,0}^e \leftarrow U_{d,0}^e \oplus u_i^e$;
- **On node** n :
 1. Compute $\text{seed}_n^e \leftarrow \text{PPRF}_{\text{salt}}(\text{seed}^e, n)$;
 2. Compute $r_n^e \leftarrow \text{PRG}(\text{seed}_n^e)$;
 3. $r^e \leftarrow r^e + r_n^e \bmod \text{bs}$, $u^e \leftarrow \text{Expand}(r^e)$, and $u_n^e \leftarrow u_n^e \oplus u^e$; // The $(x_i^e)_i$ form n pseudorandom shares of $\mathbf{x} \in [\text{bs}]^w$, the $(r_i^e)_i$ form n pseudorandom shares of $r^e \in [\text{bs}]^w$, and the $(u_i^e)_i$ form n pseudorandom shares of $u^e = \text{Expand}(r^e) \in \{0, 1\}^K$.
 4. Define $\text{aux}_n^e \leftarrow (x_n^e, u_n^e)$;
 5. Set $\text{state}_n^e \leftarrow \text{aux}_n^e \parallel \text{seed}_n^e$ and $\text{com}_n^e \leftarrow H(\text{state}_n^e)$.

Phase 2.

1. $h_1 \leftarrow H_1(m, \text{salt}, \text{com}_1^1, \dots, \text{com}_n^1, \dots, \text{com}_1^\tau, \dots, \text{com}_n^\tau)$; // Accumulate the commitments inside the hash rather than storing and hashing all at once.
2. $\pi_{\{e \in \tau\}}^e \leftarrow \text{PRG}_1(h_1)$. // $\pi^e \in \text{Perm}([w])$.

Phase 3. For each iteration $e \in [\tau]$:

1. $z^e \leftarrow x - \pi^e(r^e) \bmod \text{bs}$;
2. For $d = 1$ to D , set:
 - $y_{d,0}^e \leftarrow H \cdot \text{Shift}(\pi^e(U_{d,0}^e), z^e)$ and $y_{d,1}^e \leftarrow y_{d,0}^e \oplus \mathbf{y}$;
 - $z_{d,0}^e \leftarrow X_d^e - \pi^e(R_{d,0}^e) \bmod \text{bs}$ and $z_{d,1}^e \leftarrow z^e - z_{d,0}^e \bmod \text{bs}$.

Phase 4.

1. $h_2 \leftarrow H_2(m, \text{salt}, h_1, (y_{d,b}^e, z_{d,b}^e)_{d \leq D, b \in \{0,1\}, e \leq \tau})$;
2. Set $(b_1^e, \dots, b_D^e)_{e \leq \tau} \leftarrow \text{PRG}_2(h_2)$ and let $i^e \leftarrow \sum_{d=1}^D b_d^e \cdot 2^{d-1}$.

Phase 5. Output $\sigma = (\text{salt}, h_1, h_2, (\text{CoPath}_{\text{salt}}(i^e, \text{seed}^e), z^e, \text{com}_{i^e}^e, \text{aux}_{i^e}^e)_{e \leq \tau})$. // $\text{aux}_{i^e}^e$ is not included if $i^e = n$.

Fig. 5. Signing algorithm of the signature scheme

where $\epsilon = \mathbf{p} + \frac{1}{n} - \frac{\mathbf{p}}{n}$, with $\mathbf{p} = 4/B$ and $\epsilon_G = \epsilon_G(K, k, w, B)$ is $\Pr[(H, \mathbf{y}) \notin \text{GOOD}_B]$, which is defined on Lemma 19, $X = \max_{\alpha \in Q_1} \{X_\alpha\}$ and $Y = \max_{\beta \in Q_2} \{Y_\beta\}$ with $X_\alpha \sim \text{Binomial}(\tau, \mathbf{p})$ and $Y_\beta \sim \text{Binomial}(\tau - X, \frac{1}{n})$ where Q_1 and Q_2 are sets of all queries to oracles H_1 and H_2 .

Computing the bound \mathbf{p} from Theorem 17 requires a dedicated and involved combinatorial analysis which forms a core technical contribution of this work. We cover it extensively in section 6.2. The proof of Theorem 17 is deferred to next section 6.3.

6.2 Combinatorial Analysis of the Construction

In this section, we provide bounds on the probability that a random regular syndrome decoding instance (H, \mathbf{y}) are *bad*, in a sense that we formally define below. The bounds obtained in this section form a core component of the security analysis of our scheme in Section 6.3.

Bounding the Number of Distinct $\pi(\mathbf{u}) \downarrow \mathbf{x}$ Solutions Let $\text{Perm}_w := \text{Perm}(w)$ denote the set of all permutations $\pi : [w] \mapsto [w]$. Given a vector $\mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_w) \in \mathbb{F}_2^K$, where $(\mathbf{u}_1, \dots, \mathbf{u}_w)$ forms a decomposition of \mathbf{u} into w blocks $\mathbf{u}_i \in \mathbb{F}_2^{K/w}$, we write $\pi(\mathbf{u})$ to denote the vector $(\mathbf{u}_{\pi(1)}, \dots, \mathbf{u}_{\pi(w)})$. That is, $\pi(\mathbf{u})$ is the vector obtained by shuffling the w blocks of \mathbf{u} according to the permutation π .

For every $\mathbf{u} \in \mathbb{F}_2^K$, define $\text{pn}(\mathbf{u}) = |\{\pi(\mathbf{u}) \mid \pi \in \text{Perm}([w])\}|$. That is, $\text{pn}(\mathbf{u})$ is the number of distinct vectors in \mathbb{F}_2^K which can be obtained by permuting \mathbf{u} blockwise. Given a bound B , define $\text{PN}_B = \{\mathbf{u} \mid \text{pn}(\mathbf{u}) > B\}$.

Definition 18 (GOOD_B). *Given a bound B , GOOD_B is defined as the set of syndrome decoding instances $(H, \mathbf{y}) \in \mathbb{F}_2^{k \times K} \times \mathbb{F}_2^k$ such that for every $\mathbf{u} \in \text{PN}_B \setminus \text{Reg}_w$ and for all $\mathbf{x}^* \in [\text{bs}]^w$,*

$$\Pr_{\pi \leftarrow_r \text{Perm}_w} [H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^*) = \mathbf{y}] \leq \frac{4}{B}.$$

In other words, GOOD_B is the set of syndrome decoding instances (H, \mathbf{y}) such that for every $\mathbf{u} \notin \text{Reg}_w$ with at least B distinct blockwise permutations, at most a fraction $4/B$ of all blockwise permutations $\pi(\mathbf{u})$ are close to being solutions to $H \cdot \mathbf{x} = \mathbf{y}$, where we say that $\pi(\mathbf{u})$ is “close” to a solution if there exists a suitable cyclic shift of its block $\pi(\mathbf{u}) \downarrow \mathbf{x}^$ which is a solution.*

Equipped with this definition, we have the following lemma:

Lemma 19 (Most syndrome decoding instances are good).

$$\Pr_{H, \mathbf{y}} [(H, \mathbf{y}) \in \text{GOOD}_B] > 1 - \varepsilon_G,$$

where

$$\varepsilon_G = \binom{2B}{5} \cdot \frac{2^{K+1}}{B \cdot 2^{3k}} \cdot \left(10 + \frac{K^w}{w^w \cdot 2^k}\right).$$

Proof. The proof hinges upon a small technical lemma which we state below:

Claim. For any integer $t \leq K$ and every t -tuple of linearly-independent vectors $(\mathbf{v}_1, \dots, \mathbf{v}_t)$, it holds that

$$\Pr_{H \leftarrow_r \mathbb{F}_2^{k \times K}} [H \cdot \mathbf{v}_i = \mathbf{0} \text{ for } i = 1 \text{ to } t] = \frac{1}{2^{k-t}}.$$

Proof. Let V denote the matrix $(\mathbf{v}_1 \parallel \dots \parallel \mathbf{v}_t)$. Write $V = V^\top // V^\perp$, where $V^\top \in \mathbb{F}_2^{t \times t}$ denotes the invertible square matrix formed by the first t rows of V , and V^\perp denotes the bottom $K - t$ rows. Given a matrix H , we write $H = H_L // H_R$, where H_L denotes the t leftmost columns of H , and H_R its remaining columns. We have:

$$\begin{aligned} H \cdot V = \mathbf{0} &\iff H \cdot [V^\top // V^\perp] = \mathbf{0} \\ &\iff H \cdot [\text{Id}_t // V^\perp \cdot (V^\top)^{-1}] \cdot V^\top = \mathbf{0} \\ &\iff (H_L \cdot \text{Id}_t + H_R \cdot V^\perp \cdot (V^\top)^{-1}) \cdot V^\top = \mathbf{0} \\ &\iff H_R \cdot V^\perp \cdot (V^\top)^{-1} = H_L. \end{aligned}$$

Therefore, when H is sampled as a uniformly random matrix, we have $\Pr[H \cdot V = \mathbf{0}] = \Pr[H_R \cdot V^\perp \cdot (V^\top)^{-1} = H_L] = 1/2^{k-t}$, since the right hand side is a uniformly random matrix $H_L \leftarrow_r \mathbb{F}_2^{k \times t}$, sampled independently of the left hand side. The claim follows. \square

Now, fix $\mathbf{u} \in \text{PN}_B \setminus \text{Reg}_w$ and $\mathbf{x}^* \in [\text{bs}]^w$. Let $N \leftarrow \text{pn}(\mathbf{u})$ and $\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)}$ be the lexical ordering of all distinct vectors of the form $\pi(\mathbf{u})$ for some $\pi \in \text{Perm}_w$. Fix any subset $S = \{i_1, \dots, i_5\} \subset [N]$ of five indices. In the following, we will bound the probability

$$p(S) = \Pr_{H, \mathbf{y}} [H \cdot (\mathbf{u}^{(i_1)} \downarrow \mathbf{x}^*) = \mathbf{y} \wedge \dots \wedge H \cdot (\mathbf{u}^{(i_5)} \downarrow \mathbf{x}^*) = \mathbf{y}].$$

Recall that a regular syndrome decoding instance (H, \mathbf{y}) is sampled by picking a uniformly random matrix $H \leftarrow_r F_2^{k \times K}$, a uniformly random regular vector $\mathbf{x} \leftarrow_r \text{Reg}_w$, and setting $\mathbf{y} \leftarrow H \cdot \mathbf{x}$. When making the sampling of \mathbf{x} explicit, the probability $p(S)$ rewrites to

$$p(S) = \Pr_{H, \mathbf{x}}[H \cdot (\mathbf{u}^{(i_1)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0} \wedge \dots \wedge H \cdot (\mathbf{u}^{(i_5)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}].$$

Now, write $(\mathbf{v}_1, \dots, \mathbf{v}_5) \leftarrow (\mathbf{u}^{(i_1)} \downarrow \mathbf{x}^* \oplus \mathbf{x}, \dots, \mathbf{u}^{(i_5)} \downarrow \mathbf{x}^* \oplus \mathbf{x})$, which are random variables defined over the sampling of \mathbf{x} , and let Z_S denote the event (defined over the sampling of both \mathbf{x} and H) that $H \cdot \mathbf{v}_i = \mathbf{0}$ for $i = 1$ to 5 (in other words, $p(S) = \Pr[Z_S]$). If the vectors $(\mathbf{v}_1, \dots, \mathbf{v}_5)$ were guaranteed to be linearly independent, we would immediately get $p(S) = \Pr[Z_S] = 1/2^{5k}$ by the previous claim; however, they are not necessarily independent, and a more fine-grained approach is required. To bound $p(S)$, we make a few simple observations:

- Since $\mathbf{u} \notin \text{Reg}_w$, it also holds that for any permutation π and shifts \mathbf{x}^* , $\pi(\mathbf{u}) \downarrow \mathbf{x}^* \notin \text{Reg}_w$ (since shuffling the blocks and cyclically shifting each block yields an invertible mapping that preserves regularity). This implies that $\mathbf{v}_j \neq \mathbf{0}$ holds with probability 1 for $j = 1$ to 5 (since $\mathbf{v}_j = \mathbf{0} \iff \mathbf{u}^{(i_1)} \downarrow \mathbf{x}^* = \mathbf{x}$, and $\mathbf{x} \in \text{Reg}_w$).
- Because the $\mathbf{u}^{(i)}$ are pairwise distinct (by definition), the \mathbf{v}_j are pairwise distinct.

Equipped with these observations, let us denote E_S the event that there exist three integers $\alpha \neq \beta \neq \gamma \in [5]$ such that $\mathbf{v}_\alpha \oplus \mathbf{v}_\beta \oplus \mathbf{v}_\gamma = \mathbf{0}$. Observe that

$$\begin{aligned} \Pr[E_S] &= \Pr_{\mathbf{x}}[\exists \alpha \neq \beta \neq \gamma \in [5] : (\mathbf{u}^{(i_\alpha)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) \oplus (\mathbf{u}^{(i_\beta)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) \oplus (\mathbf{u}^{(i_\gamma)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}] \\ &= \Pr_{\mathbf{x}}[\exists \alpha \neq \beta \neq \gamma \in [5] : (\mathbf{u}^{(i_\alpha)} \downarrow \mathbf{x}^*) \oplus (\mathbf{u}^{(i_\beta)} \downarrow \mathbf{x}^*) \oplus (\mathbf{u}^{(i_\gamma)} \downarrow \mathbf{x}^*) = \mathbf{x}] \\ &\leq \sum_{\alpha \neq \beta \neq \gamma} \Pr_{\mathbf{x}}[(\mathbf{u}^{(i_\alpha)} \downarrow \mathbf{x}^*) \oplus (\mathbf{u}^{(i_\beta)} \downarrow \mathbf{x}^*) \oplus (\mathbf{u}^{(i_\gamma)} \downarrow \mathbf{x}^*) = \mathbf{x}] \\ &\leq \binom{5}{3} \cdot \left(\frac{K}{w}\right)^{-w} = 10 \cdot \left(\frac{K}{w}\right)^{-w}, \end{aligned}$$

which follows from a union bound over all possible size-3 subsets of $[5]$ and because there are $(K/w)^w$ vectors in Reg_w , hence a $(K/w)^{-w}$ probability (at most) that a random vector $\mathbf{x} \leftarrow_r \text{Reg}_w$ is equal to the fixed vector $(\mathbf{u}^{(i_\alpha)} \downarrow \mathbf{x}^*) \oplus (\mathbf{u}^{(i_\beta)} \downarrow \mathbf{x}^*) \oplus (\mathbf{u}^{(i_\gamma)} \downarrow \mathbf{x}^*)$. Now, we have

$$\begin{aligned} \Pr[Z_S] &= \Pr[Z_S \mid E_S] \cdot \Pr[E_S] + \Pr[Z_S \mid \neg E_S] \cdot \Pr[\neg E_S] \\ &\leq 10 \cdot (K/w)^{-w} \cdot \Pr[Z \mid E_S] + \Pr[Z \mid \neg E_S]. \end{aligned}$$

We now bound $\Pr[Z_S \mid E_S]$. For simplicity and without loss of generality, assume that after sampling \mathbf{x} , we have $\mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_3 = \mathbf{0}$ (this is without loss of generality because we can always reorder the \mathbf{v}_i 's after sampling \mathbf{x} ; note that the event E_S is defined only over the sampling of \mathbf{x}). Then, because $\mathbf{v}_4 \neq \mathbf{v}_3$, it necessarily holds that $\mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_4 \neq \mathbf{0}$. Furthermore, since the \mathbf{v}_i are all pairwise distinct, and all nonzero, this implies that $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_4)$ are linearly independent. Then, using the claim:

$$\Pr[Z_S \mid E_S] \leq \Pr[H \cdot \mathbf{v}_1 = \mathbf{0} \wedge H \cdot \mathbf{v}_2 = \mathbf{0} \wedge H \cdot \mathbf{v}_4 = \mathbf{0} \mid E_S] = \frac{1}{2^{3k}}.$$

We now bound $\Pr[Z_S \mid \neg E_S]$. By a similar reasoning, after sampling \mathbf{x} , it necessarily holds that there is a 4-tuple of the \mathbf{v}_i 's that does not XOR to 0 (since if all 4-tuples of the \mathbf{v}_i 's XOR to 0, we have $\mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_3 \oplus \mathbf{v}_4 = \mathbf{v}_1 \oplus \mathbf{v}_2 \oplus \mathbf{v}_3 \oplus \mathbf{v}_5 = \mathbf{0}$, which implies $\mathbf{v}_4 = \mathbf{v}_5$, contradicting the fact that the \mathbf{v}_i 's are pairwise distinct). Without loss of generality, assume that $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$ do not XOR to 0. Because we condition on $\neg E_S$, it also holds that no 3-tuple of vectors from $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$ XOR to 0, and because the \mathbf{v}_i 's are pairwise distinct (*i.e.* no two-tuple XOR to 0) and nonzero, it follows that $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$ are linearly independent. By the previous claim:

$$\Pr[Z_S \mid \neg E_S] \leq \Pr[H \cdot \mathbf{v}_1 = \mathbf{0} \wedge H \cdot \mathbf{v}_2 = \mathbf{0} \wedge H \cdot \mathbf{v}_3 = \mathbf{0} \wedge H \cdot \mathbf{v}_4 = \mathbf{0} \mid E] = \frac{1}{2^{4k}}.$$

Eventually, we get

$$\begin{aligned} p(S) &= \Pr[Z_S] \leq 10 \cdot (K/w)^{-w} \cdot \Pr[Z_S \mid E_S] + \Pr[Z_S \mid \neg E_S] \\ &\leq \frac{1}{2^{3k}} \cdot \left(10 \cdot \left(\frac{K}{w}\right)^{-w} + \frac{1}{2^k}\right). \end{aligned}$$

We now finish the proof of Lemma 19 by a careful union bound. Given $\mathbf{u} \in \text{PN}_B \setminus \text{Reg}_w$ and $\mathbf{x}^* \in [\text{bs}]^w$, let us partition the $N = \text{pn}(\mathbf{u})$ vectors $\mathbf{u}^{(i)} \downarrow \mathbf{x}^*$ into $m \leq N/B$ blocks of at most $2B$ vectors each. Let N_1, \dots, N_m denote the m disjoint subsets $N_i \subset [N]$ of size $|N_i| \leq B$ corresponding to this partition. We first use a union bound over all possible blocks N_i , and all possible size-5 subsets of N_i :

$$\begin{aligned} & \Pr_{H, \mathbf{x}} [\exists i \leq m, \exists S_i \subset N_i \subset [N] \text{ with } |S_i| = 5 : H \cdot (\mathbf{u}^{(j)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0} \text{ for all } j \in S_i] \\ & \leq m \cdot \binom{2B}{5} \cdot \frac{1}{2^{3k}} \cdot \left(10 \cdot \left(\frac{K}{w} \right)^{-w} + \frac{1}{2^k} \right). \end{aligned}$$

In particular, this implies that for any fixed $\mathbf{u} \in \mathbb{F}_2^K$ with $\text{pn}(\mathbf{u}) = N$, and any fixed $\mathbf{x}^* \in [\text{bs}]^w$, there are at most $4 \cdot m$ indices $j \in [N]$ such that $H \cdot (\mathbf{u}^{(j)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}$ with high probability (since with high probability, in each of the m block, there are at most 4 such indices):

$$\begin{aligned} & 1 - m \cdot \binom{2B}{5} \cdot \frac{1}{2^{3k}} \cdot \left(10 \cdot \left(\frac{K}{w} \right)^{-w} + \frac{1}{2^k} \right) \\ & < \Pr_{H, \mathbf{x}} [\forall i \leq m, \forall S_i \subset N_i \subset [N] \text{ with } |S_i| = 5 : \exists j \in S_i, H \cdot (\mathbf{u}^{(j)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}] \\ & = \Pr_{H, \mathbf{x}} [\forall i \leq m : \text{there are at most 4 } j \in N_i \text{ s.t. } H \cdot (\mathbf{u}^{(j)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}] \\ & \leq \Pr_{H, \mathbf{x}} [\exists \leq 4 \cdot m \text{ indices } j \in [N] \text{ such that } H \cdot (\mathbf{u}^{(j)} \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0} \text{ for all } j \in S_i] \\ & = \Pr_{H, \mathbf{x}} \left[\Pr_{\pi \in \text{Perm}_w} [H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^* \oplus \mathbf{x}) \neq \mathbf{0}] \leq \frac{4 \cdot N/B}{N} = \frac{4}{B} \right]. \end{aligned}$$

Next, we compute a union bound over all possible vectors \mathbf{u} with permutation number $\text{pn}(\mathbf{u}) = N$ (where $N \leq w!$, with equality when all blocks of \mathbf{u} are distinct) and all shifts $\mathbf{x}^* \in [\text{bs}]^w$. For any $N \in [w!]$, let $n(N)$ denote the total number of vectors $\mathbf{u} \in \mathbb{F}_2^K$ with $\text{pn}(\mathbf{u}) = N$ (note that $\sum_{i \in [w!]} n(N) = 2^K$). We group all vectors \mathbf{u} with $\text{pn}(\mathbf{u}) = N$ into $n(N)/N$ equivalence classes $U_1, \dots, U_{n(N)/N}$, where two vectors $\mathbf{u}_1, \mathbf{u}_2$ belong to the same equivalence class U_i if and only if there exists $\pi \in \text{Perm}_w$ such that $\mathbf{u}_1 = \pi(\mathbf{u}_2)$ (note that each equivalence class is of size exactly N by definition of pn , and the U_i form a partition of the set $\{\mathbf{u} \in \mathbb{F}_2^K : \text{pn}(\mathbf{u}) = N\}$). An important observation is that, because any two vectors $\mathbf{u}_1, \mathbf{u}_2$ that belong to the same equivalence class U_i generate the exact same N -tuple of distinct permuted vectors $(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(N)})$ (ordered lexically), it suffices to do the union bound over all possible *equivalence classes* $(U_1, \dots, U_{n(N)/N})$, and over all shifts \mathbf{x}^* :

$$\begin{aligned} & \Pr_{H, \mathbf{x}} \left[\exists i \leq n(N)/N, \exists \mathbf{x}^* \in [\text{bs}]^w, \Pr_{\pi \in \text{Perm}_w} [H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}] > \frac{4}{B} \right] \\ & \leq \frac{n(N)}{N} \cdot \frac{N}{B} \cdot \binom{2B}{5} \cdot \left(\frac{K}{w} \right)^w \cdot \frac{1}{2^{3k}} \cdot \left(10 \cdot \left(\frac{K}{w} \right)^{-w} + \frac{1}{2^k} \right), \end{aligned}$$

where the vector \mathbf{u} in the probability denote any representent of the class U_i . Eventually, we use a union bound over all possible values $N \in [w!]$ with $N \geq B$:

$$\begin{aligned} & \Pr_{H, \mathbf{y}} [(H, \mathbf{y}) \notin \text{GOOD}_B] \\ & = \Pr_{H, \mathbf{x}} \left[\exists N \geq B \in [w!], \exists i \leq n(N)/N, \exists \mathbf{x}^* \in [\text{bs}]^w, \Pr_{\pi \in \text{Perm}_w} [H \cdot (\pi(\mathbf{u}) \downarrow \mathbf{x}^* \oplus \mathbf{x}) = \mathbf{0}] > \frac{4}{B} \right] \\ & \leq \binom{2B}{5} \cdot \left(\frac{K}{w} \right)^w \cdot \frac{1}{B \cdot 2^{3k}} \cdot \left(10 \cdot \left(\frac{K}{w} \right)^{-w} + \frac{1}{2^k} \right) \cdot \sum_{N=B}^{w!} n(N) \\ & < \binom{2B}{5} \cdot \frac{2^{K+1}}{B \cdot 2^{3k}} \cdot \left(10 + \frac{K^w}{w^w \cdot 2^k} \right), \end{aligned}$$

which concludes the proof of Lemma 19. \square

6.3 Security Analysis of the Signature Scheme

In this section, we prove Theorem 17.

Reducing to EUF-KO Security We start by proving the following lemma:

Lemma 20 (EUF-KO \implies EUF-CMA).

$$\text{Adv}_{\mathcal{A}}^{\text{EUF-CMA}} \leq \text{Adv}_{\mathcal{A}}^{\text{EUF-KO}} + \frac{q_s(q_s + q_1 + q_2 + q_3)}{2^{2\lambda}} + \epsilon_{\text{PPRF}} + \epsilon_{\text{PRG}}$$

Proof. Let us consider an adversary \mathcal{A} against the EUF-CMA property of the signature scheme. To prove security we will define a sequence of experiments involving \mathcal{A} , where the first corresponds to the experiment in which \mathcal{A} interacts with the real signature scheme, and the last one is an experiment in which \mathcal{A} is using only random element independent from the witness.

Game 1 (Gm^1). This corresponds to the actual interaction of \mathcal{A} with the real signature scheme. We need to bound the probability of what we'll call **Forge**, i.e. the event that \mathcal{A} can generate a valid signature for a message that was not previously queried to the signing oracle.

Game 2 (Gm^2). For this step, we abort if the sampled salt **salt** collides with the value sampled in any of the previous queries to hash functions H_1 or H_2 or if the input of PRG_2 collides with the value obtained in any of the previous queries. Therefore we can bound this probability by

$$|\Pr[\text{Gm}^1(\text{Forge})] - \Pr[\text{Gm}^2(\text{Forge})]| \leq \frac{q_s \cdot (q_s + q_1 + q_2 + q_3)}{2^{2\lambda}}$$

Game 3 (Gm^3). The difference with the previous game is that now before signing a message we choose uniformly random values h_1, h_2 and i^* . Since Phase1, Phase3 and Phase5 are computed as before and the only change compared to the previous game is that we set the output of H_1 as h_1 , the output of H_2 as h_2 and the output of $\text{PRG}_2(h_2)$ as i^* then the difference in forgery probability is due to the event that query to H_1, H_2 or PRG_2 was ever made before but in this scenario Game 2 aborts, so

$$\Pr[\text{Gm}^2(\text{Forge})] = \Pr[\text{Gm}^3(\text{Forge})]$$

Game 4 (Gm^4). In this game we sample at random the i^* -th seed seed_{i^*} and the related co-path CoPath_{i^*} . By using all the seeds $\{\text{seed}_i\}_{i \neq i^*}$ in the CoPath_{i^*} we will proceed by computing all the parties' views as well as the auxiliary material. Therefore, Phase 1 and Phase 3 are executed in the actual way (i.e. by using the real witness) except for i^* , for which the values are obtained randomly instead of using the PPRF. Distinguishing between this game and the previous one is perfectly equivalent to breaking the multi-instance security of the PPRF:

$$|\Pr[\text{Gm}^4(\text{Forge})] - \Pr[\text{Gm}^6(\text{Forge})]| \leq \epsilon_{\text{PPRF}}$$

Game 5 (Gm^5). Now, before signing a message, we choose a uniformly random value to be used as the i^* -th party's view, i.e. $(x_{i^*}, r_{i^*}, u_{i^*})$, and its commitment com_{i^*} . Since in the previous game, these values were computed by using a multi-instance PRG on a random seed, with salt **salt**, we can bound

$$|\Pr[\text{Gm}^4(\text{Forge})] - \Pr[\text{Gm}^6(\text{Forge})]| \leq \epsilon_{\text{PRG}}$$

Game 6 (Gm^6). In this game, we will change Phase 1 and Phase 3 by making the signer use the internal HVZK simulator described in 6. Looking in detail, the only change between the previous game and this one is that the auxiliary material **aux** is now selected as random. Anyway, since in the previous game **aux** was computed by using all real values but one (randomly chosen and never made public), there is essentially no difference between this game and the previous one. Therefore,

$$\Pr[\text{Gm}^5(\text{Forge})] = \Pr[\text{Gm}^6(\text{Forge})]$$

Game 7 (Gm^7). We say that an execution e^* of a query

$$h_2 = \text{H}_2(m, \text{salt}, h_1, (y_{d,b}^e, z_{d,b}^e)_{d \leq D, b \in \{0,1\}, e \leq \tau})$$

defines a correct witness if the following criteria are satisfied:

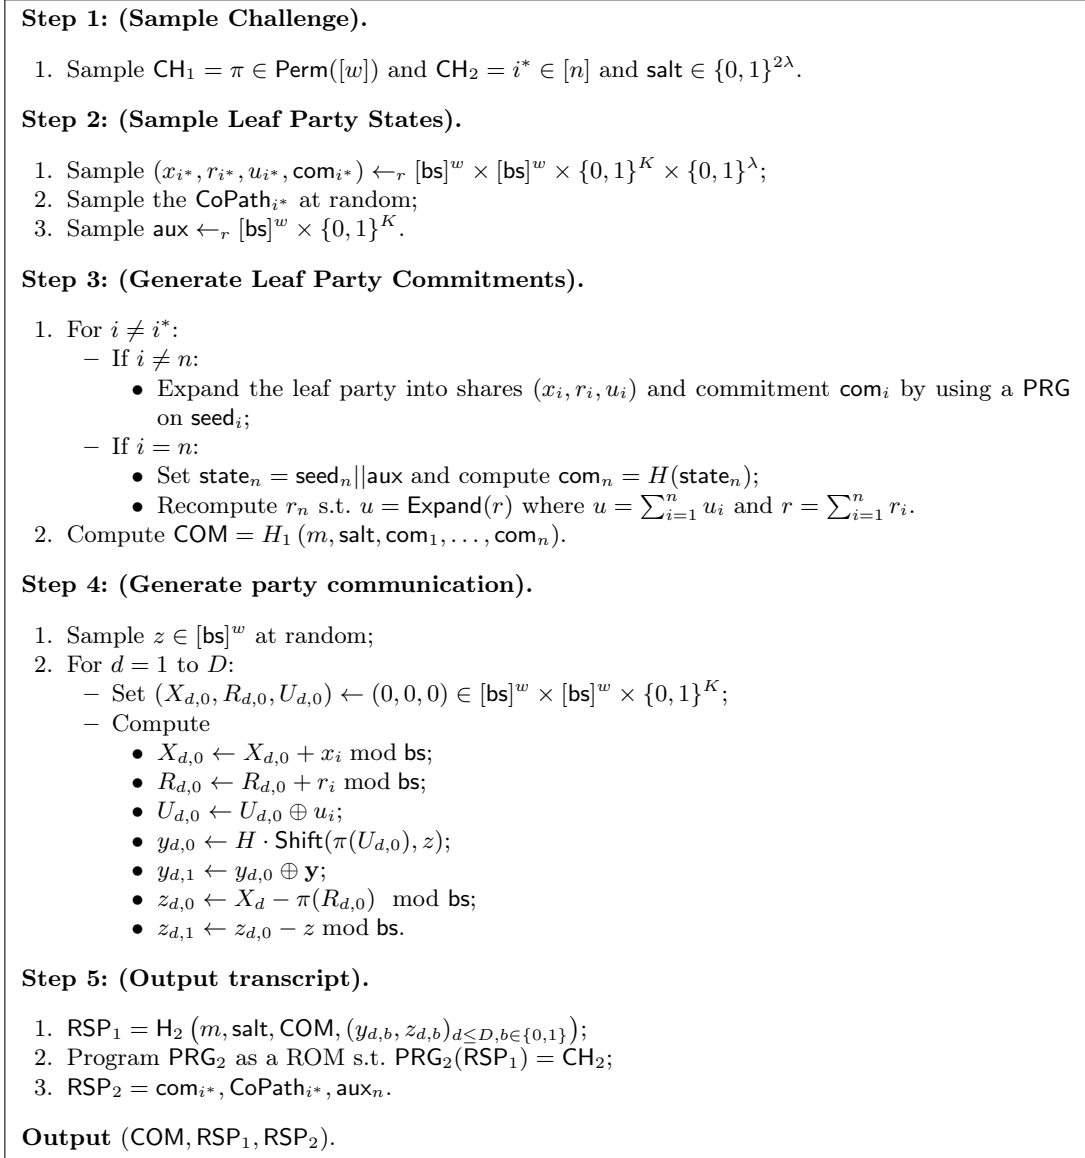


Fig. 6. Internal HVZK simulator for signing algorithm

– h_1 was output by a previous query

$$h_1 \leftarrow H_1(m, \text{salt}, \text{com}_1^1, \dots, \text{com}_n^1, \dots, \text{com}_1^\tau, \dots, \text{com}_n^\tau);$$

– each $\text{com}_i^{e^*}$ in this query was output by a previous query

$$\text{com}_i^{e^*} = \text{PRG}(\text{seed}^{e^*}, i)$$

for each $i \in [N]$;

– The vector \mathbf{x} defined by the leaf party states $\{\text{state}_i\}_{i \in [N]}$ satisfies $\text{HW}(\mathbf{x}) = w$ and $H\mathbf{x} = \mathbf{y}$.

In this game, for each query of H_2 made by the adversary, we will check if there is an execution e^* that defines a correct witness. Calling this event **Solve** of course, since if it occurs then the states $\{\text{state}_i^{e^*}\}$ define a solution for the RSD, the probability $\Pr[\text{Solve} \leq \epsilon_{\text{SD}}]$.

EUF-KO Security We prove the following lemma:

Lemma 21 (EUf-KO security).

$$\text{Adv}_{\mathcal{A}}^{\text{EUf-KO}} \leq \epsilon_{\text{SD}} + \Pr[X + Y = \tau] + \epsilon_{\text{G}} + \frac{1}{2^\lambda}.$$

Together with lemma 20, this completes the proof of Theorem 17. To prove EUF-KO security of our signature scheme, we first analyze the soundness of the underlying identification scheme, and then apply the standard reduction to EUF-KO security after compiling the scheme with the Fiat-Shamir transform. Concretely, our signature scheme is obtained by applying the Fiat-Shamir transform to the τ -fold parallel repetition of the identification scheme defined in Figure 7.

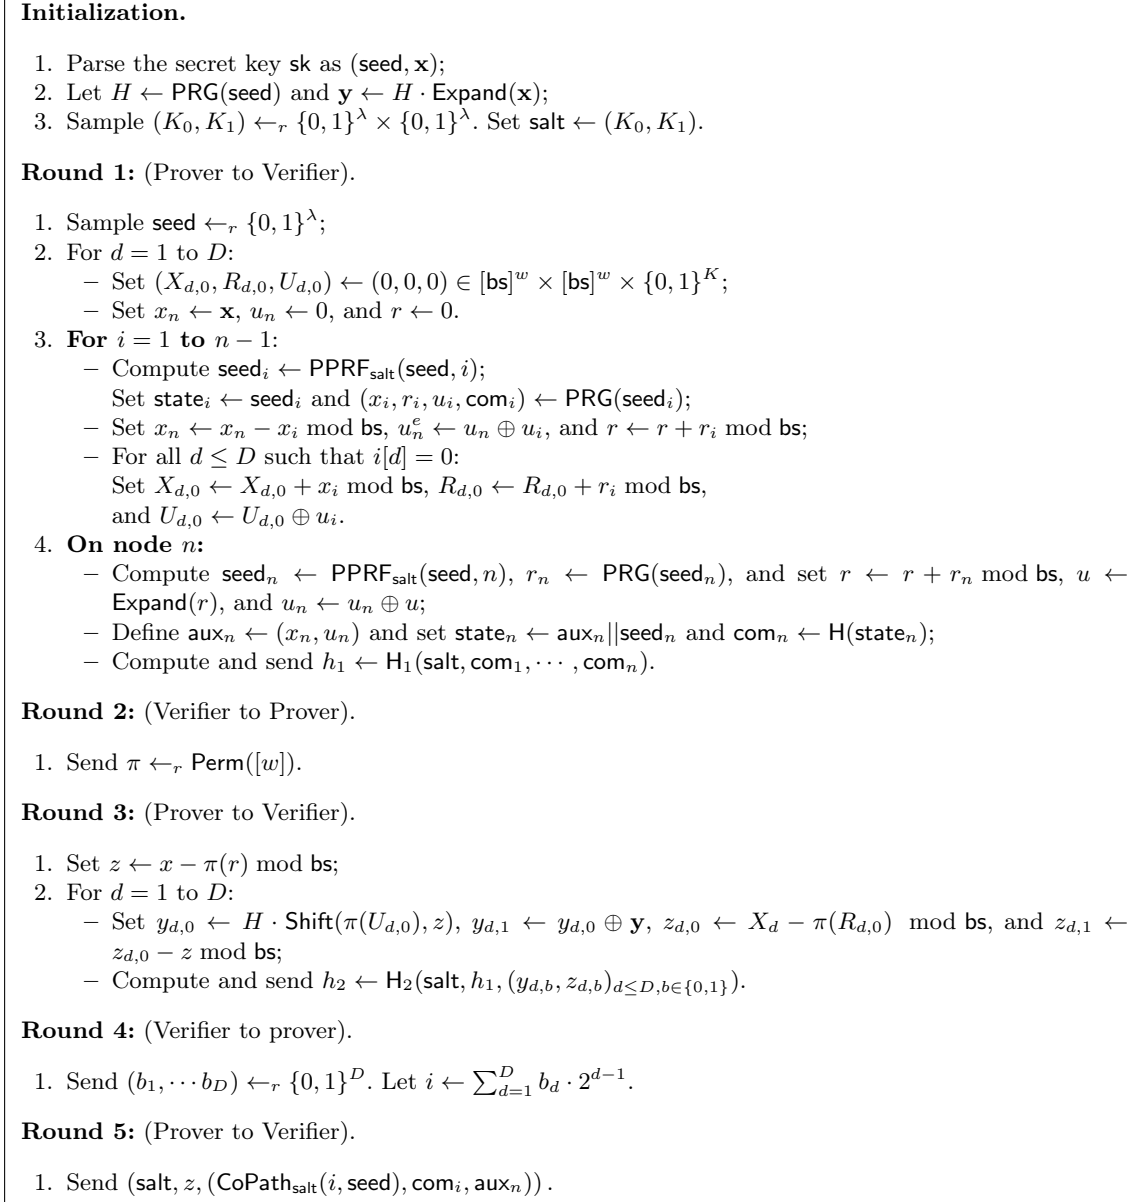


Fig. 7. A five-round identification scheme with secret key $\text{sk} = (\text{seed}, \mathbf{x})$ for the relation $\mathbf{y} = H \cdot \mathbf{x}$ with $\mathbf{y} \in \text{Reg}_w$ and $H = \text{PRG}(\text{seed})$. The scheme has soundness $\varepsilon = p + (1 - p)/n$.

Soundness of the identification scheme. The core of our analysis will be dedicated to show that from any cheating prover, one can extract a *weakly valid witness*. Concretely, a weakly valid witness is a pair $(\mathbf{v}, \mathbf{x}^*)$ where \mathbf{v} is a solution to $H \cdot \mathbf{v} = \mathbf{y}$ which might not be regular, but which satisfies $\text{pn}(\mathbf{v} \downarrow \mathbf{x}^*) \leq B$. In other words, this means that \mathbf{v} contains mostly identical blocks “up to shift”. Note that a regular vector contains only copies of the unit vector \mathbf{e}_1 “up to shift”, hence this generalizes the class of regular vectors in a specific sense. Formally:

Definition 22. A weakly valid witness to a syndrome decoding instance (H, \mathbf{y}) is a pair $(\mathbf{v}, \mathbf{x}^*)$ such that $H \cdot \mathbf{v} = \mathbf{y}$ and $\mathbf{v} \in X$, where X is defined as:

$$X = \{\mathbf{v} \in \mathbb{F}_2^K : \exists \mathbf{u} \in \mathbb{F}_2^K \setminus \text{PN}_B, \exists \mathbf{x}^* \in [\text{bs}]^w, \mathbf{v} = \mathbf{u} \downarrow \mathbf{x}^*\}.$$

Additionally, the second term \mathbf{x}^* of the pair is a shift that satisfies $\mathbf{v} \downarrow \mathbf{x}^* \in \mathbb{F}_2^K \setminus \text{PN}_B$.

Lemma 23 (Soundness of the identification scheme in Figure 7). Assume that H_1, H_2 are collision-resistant hash functions, that the mapping $\text{PRG}(\text{seed})_{1..\lambda}$ (i.e. the first λ bits of the output of PRG on an input seed) is computationally binding, and that the PRG used during key generation in $H \leftarrow \text{PRG}(\text{seed})$ is modeled as a random oracle (hence $\text{PRG}(\text{seed})$ selects a truly random matrix H). Then with probability at least $1 - 1/2^\lambda - \varepsilon_G$ over the random choice of the RSD instance (H, \mathbf{y}) , there exists an expected polynomial time extractor algorithm which, given rewinding access to a prover \tilde{P} which generates an accepting proof with probability at least $\tilde{\varepsilon} > \mathbf{p} + 1/n - \mathbf{p}/n$, extracts a weakly valid witness \mathbf{x} for the relation $H \cdot \mathbf{x} = \mathbf{y}$.

Looking ahead, our soundness proof does not prevent a cheating prover from coming up with a weakly valid witness which is not a true regular witness. This will be guaranteed by the fact that with probability $1 - 1/2^\lambda$ over the choice of a random instance (H, \mathbf{y}) , the system of equations $H \cdot \mathbf{v} = \mathbf{y}$ does not have *any* solution in X beyond the regular solution. This holds for a suitable choice of the parameters (K, k, w) , which we cover in detail in Section 7. The term $1/2^\lambda$ in the bound of Lemma 21 reflects the probability that (H, \mathbf{y}) admits weakly valid solutions which are not regular.

Proof of Lemma 23. Let \tilde{P} be a prover which manages to generate an accepting proof with probability $\tilde{\varepsilon} > \varepsilon$. We exhibit an *extractor* which finds a witness \mathbf{x} such that $H \cdot \mathbf{x} = \mathbf{y}$, where \mathbf{x} is guaranteed to be a *weakly valid* witness (see Definition 22). Let R denote the randomness used by \tilde{P} to generate the commitment h of the first round, and by R^* a possible realization of R . Let $\text{Succ}_{\tilde{P}}$ denote the event that \tilde{P} succeeds in convincing V . By hypothesis

$$\Pr[\text{Succ}_{\tilde{P}}] = \tilde{\varepsilon} > \varepsilon = \mathbf{p} + \frac{1}{n} - \frac{\mathbf{p}}{n}.$$

Let us fix an arbitrary value $\alpha \in \{0, 1\}$ such that $(1 - \alpha)\tilde{\varepsilon} > \varepsilon$, which exists since $\tilde{\varepsilon} > \varepsilon$. We say that a realization R^* of the prover randomness for the first flow is *good* if it holds that

$$\Pr[\text{Succ}_{\tilde{P}} | R = R^*] \geq (1 - \alpha)\tilde{\varepsilon}.$$

Furthermore, by the Splitting Lemma (see e.g. [FJR22]), we have $\Pr[R \text{ good} | \text{Succ}_{\tilde{P}}] \geq \alpha$. Assume now that T_0 is the transcript of a successful execution of the zero-knowledge proof with \tilde{P} . Let R^* denote the random coin used by \tilde{P} in the first round, and let i_0 denote the Round 4 message of the verifier. If R^* is *good*, then

$$\Pr[\text{Succ}_{\tilde{P}} | R = R^*] \geq (1 - \alpha)\tilde{\varepsilon} > \varepsilon > \frac{1}{n},$$

which implies that there necessarily exists a second successful transcript T_1 with a different Round 4 message $i_1 \neq i_0$.

Consistency of (T_0, T_1) . Let (π_0, i_0) and (π_1, i_1) be the verifier challenges in the successful transcripts T_0 and T_1 respectively, with $i_0 \neq i_1$. Let us denote $(\text{state}_{i \neq i_0}^0, \text{com}_{i_0}^0)$ and $(\text{state}_{i \neq i_1}^1, \text{com}_{i_1}^1)$ the states (recomputed from the co-path included in the transcript) and the commitment in the transcripts T_0 and T_1 respectively. Suppose that $\exists i \in [n] \setminus \{i_0, i_1\}$ such that $\text{state}_i^0 \neq \text{state}_i^1$. Then there are two possibilities:

- The commitments are different:

$$\text{com}_i = \text{PRG}(\text{state}_i)_{1..\lambda} \neq \text{PRG}(\text{state}'_i)_{1..\lambda} = \text{com}'_i.$$

But since T_0 and T_1 are accepting transcripts, this implies in particular that $h = H_1(\text{com}_1, \dots, \text{com}_n)$ and $h = H_1(\text{com}'_1, \dots, \text{com}'_n)$ which contradicts the collision resistance of H_1 .

- The commitments are equal:

$$\text{com}_i = \text{PRG}(\text{state}_i)_{1..\lambda} = \text{PRG}(\text{state}'_i)_{1..\lambda} = \text{com}'_i.$$

This directly contradicts the binding property of PRG.

Therefore, it necessarily holds that the states are mutually consistent (that is $\text{state}_{i \neq i_0, i_1}^0 = \text{state}_{i \neq i_0, i_1}^1$). Since $i_0 \neq i_1$, they jointly define a unique tuple $(\text{state}_i)_{i \in [n]}$, from which we can recompute $\mathbf{x} = \sum_i x_i \bmod \text{bs}$, $u = \bigoplus_i u_i$, and $r = \sum_i r_i \bmod \text{bs}$. Let us denote $v \leftarrow u \uparrow r$.

Claim. The vector v belongs to $\mathbb{F}_2^K \setminus \text{PN}_B$.

To prove the claim, we show that if $v \in \text{PN}_B$, then $\Pr[\text{Succ}_{\bar{p}} | R = R^*] \leq \varepsilon$, contradicting our assumption that R^* is good. Let us denote $\text{BadPerm} = \text{BadPerm}_{v, \mathbf{x}}$ the event (defined over the random choice of a permutation π , and for the fixed value of $(v, \mathbf{x}, H, \mathbf{y})$) that $\mathbf{y} = H \cdot (\pi(v) \downarrow \mathbf{x})$. Let ε_G denote the bound of Lemma 19.

By Lemma 19, it holds with probability $1 - \varepsilon_G$ over the random choice of H that $\Pr[\text{BadPerm}] \leq \mathbf{p}$ with $\mathbf{p} = 4/B$ (here, we use the fact that in the random oracle model, $H = \text{PRG}(\text{seed})$ is uniformly random). Now,

$$\begin{aligned} \Pr[\text{Succ}_{\bar{p}} | R = R^*] &= \Pr[\text{Succ}_{\bar{p}} \wedge \text{BadPerm} | R = R^*] + \Pr[\text{Succ}_{\bar{p}} \wedge \neg \text{BadPerm} | R = R^*] \\ &\leq \mathbf{p} + (1 - \mathbf{p}) \cdot \Pr[\text{Succ}_{\bar{p}} | R = R^* \wedge \neg \text{BadPerm}]. \end{aligned}$$

We now bound $\Pr[\text{Succ}_{\bar{p}} | R = R^* \wedge \neg \text{BadPerm}]$. Assume for the sake of contradiction that $\Pr[\text{Succ}_{\bar{p}} | R = R^* \wedge \neg \text{BadPerm}] > 1/n$. This implies that given any successful transcript T'_0 with fourth-round i'_0 , there necessarily exists a second successful transcript T'_1 with the same first three rounds and a different fourth-round $i'_1 \neq i'_0$. Let us fix two such transcripts (T'_0, T'_1) , and let π' be the (common) permutation sent in Round 2 of these transcripts.

By the same argument as before, T'_0 and T'_1 are necessarily consistent, and uniquely define a tuple $(\text{state}'_i)_{i \in [n]}$. Furthermore, since we condition on $R = R^*$, meaning that the first flow h'_1 is the same as the first flow h_1 in T_0, T_1 , it must hold that $(\text{state}'_i)_{i \in [n]} = (\text{state}_i)_{i \in [n]}$, the states uniquely defined by (T_0, T_1) (else, we contradict either the collision-resistance of H or the binding of PRG, as already shown).

Let $d \leq D$ be a position such that $i'_0[d] \neq i'_1[d]$. Without loss of generality (since we can always swap the roles of T'_0 and T'_1), let us assume that $i'_0[d] = 0$ and $i'_1[d] = 1$. Reconstruct the values $(y_{b,d}^{(0)}, z_{b,d}^{(0)})_{b \in \{0,1\}}$ using the seeds $(\text{seed}_i)_{i \neq i'_0}$ and the permutation π' from transcript T'_0 , using the same procedure as the verification procedure. Similarly, reconstruct the values $(y_{b,d}^{(1)}, z_{b,d}^{(1)})_{d \leq D, b \in \{0,1\}}$ using the seeds $(\text{seed}_i)_{i \neq i'_1}$ and the permutation π' from the transcript T'_1 (which are the same as in T'_0). This yields

$$\begin{aligned} y_{d,0}^{(0)} &= H \cdot (\pi'(U_{d,0}^{(0)}) \downarrow z^{(0)}) \\ y_{d,1}^{(1)} &= H \cdot (\pi'(U_{d,1}^{(1)}) \downarrow z^{(1)}), \end{aligned}$$

where $z^{(0)}$ and $z^{(1)}$ are the Round 5 vectors included in the transcripts T'_0 and T'_1 , and

$$\begin{aligned} z_{d,0}^{(0)} &= X_{d,0}^{(0)} - \pi'(R_{d,0}^{(0)}) \\ z_{d,1}^{(1)} &= X_{d,1}^{(1)} - \pi'(R_{d,1}^{(1)}). \end{aligned}$$

Now, because T'_0 and T'_1 share the same states $(\text{state}'_i)_{i \leq n}$, it holds by construction that $U_{d,0}^{(0)} + U_{d,1}^{(1)} = u$, $X_{d,0}^{(0)} + X_{d,1}^{(1)} = \mathbf{x}$, $R_{d,0}^{(0)} + R_{d,1}^{(1)} = r$, and $\mathbf{y} = y_{d,0}^{(0)} + y_{d,1}^{(1)}$. Furthermore, by the collision-resistance of H_2 , it must hold that $y_{d,b}^{(0)} = y_{d,b}^{(1)}$ and $z_{d,b}^{(0)} = z_{d,b}^{(1)}$ for every $b \in \{0,1\}$. The latter equality implies that $z^{(0)} = z_{d,0}^{(0)} + z_{d,1}^{(0)} = z^{(1)}$ (we denote z this value from now on). This gives

$$z = z_{d,0}^{(0)} + z_{d,1}^{(1)} = X_{d,0}^{(0)} - \pi'(R_{d,0}^{(0)}) + X_{d,1}^{(1)} - \pi'(R_{d,1}^{(1)}) = \mathbf{x} - \pi(r).$$

Furthermore,

$$\mathbf{y} = y_{d,0}^{(0)} + y_{d,1}^{(1)} = H \cdot (\pi'(U_{d,0}^{(0)}) + \pi'(U_{d,1}^{(1)}) \downarrow z) = H \cdot (\pi'(u) \downarrow z).$$

We conclude by observing that $\pi(u) \downarrow z = \pi'(u) \downarrow (\mathbf{x} - \pi(r)) = \pi'(u \uparrow v) \downarrow \mathbf{x} = \pi'(v) \downarrow \mathbf{x}$, hence we have $H \cdot (\pi'(v) \downarrow \mathbf{x}) = \mathbf{y}$, which is a contradiction since the sampling on π' is conditioned on

$\neg\text{BadPerm}$. Hence, assuming the collision-resistance of H_2 , it necessarily holds that $\Pr[\text{Succ}_{\tilde{P}}|R = R^* \wedge \neg\text{BadPerm}] \leq 1/n$. Finishing the proof:

$$\begin{aligned} \Pr[\text{Succ}_{\tilde{P}}|R = R^*] &\leq \mathfrak{p} + (1 - \mathfrak{p}) \cdot \Pr[\text{Succ}_{\tilde{P}}|R = R^* \wedge \neg\text{BadPerm}] \\ &\leq \mathfrak{p} + (1 - \mathfrak{p}) \cdot \frac{1}{n} = \varepsilon, \end{aligned}$$

contradicting our assumption that R^* is good. Therefore, we have extracted a vector v , a tuple \mathbf{x} , and a permutation π' such that $H \cdot (\pi'(v) \downarrow \mathbf{x}) = \mathbf{y}$, yet $v \in \mathbb{F}_2^K \setminus \text{PN}_B$.

The extractor. Equipped with the above analysis, we describe an extractor \mathcal{E} which is given rewindable black-box access to a prover \tilde{P} . Define $N \leftarrow \ln(2)/((1 - \alpha)\tilde{\varepsilon} - \varepsilon)$. \mathcal{E} works as follows:

- Run \tilde{P} and simulate a honest verifier V to get a transcript T_0 . Restart until T_0 is a successful transcript.
- Do N times:
 - Run \tilde{P} with a honest V and the same randomness as in T_0 to get a transcript T_1 .
 - If T_1 is a successful transcript with $i_0 \neq i_1$, extract the tuple (\mathbf{x}, u, r) and the permutation π . Output $\pi(v) \downarrow \mathbf{x}$.

The end of the proof is perfectly identical to the analysis in [FJR22, Appendix F]: given that \mathcal{E} found a first successful transcript T_0 , we have

$$\begin{aligned} \Pr[\text{Succ}_{\tilde{P}}^{T_1} \wedge i_1 \neq i_0 | R \text{ good}] &= \Pr[\text{Succ}_{\tilde{P}}^{T_1} | R \text{ good}] - \Pr[\text{Succ}_{\tilde{P}}^{T_1} \wedge i_1 = i_0 | R \text{ good}] \\ &\geq (1 - \alpha)\tilde{\varepsilon} - 1/n \geq (1 - \alpha)\tilde{\varepsilon} - \varepsilon, \end{aligned}$$

hence by definition of N , \mathcal{E} gets a second successful transcript with probability at least $1/2$. From there, the analysis of the expected number of calls $\mathbb{E}[\text{call}]$ of \mathcal{E} to \tilde{P} is identical to [FJR22, Appendix F]:

$$\begin{aligned} \mathbb{E}[\text{call}] &\leq 1 + (1 - \Pr[\text{Succ}_{\tilde{P}}]) \cdot \mathbb{E}[\text{call}] + \Pr[\text{Succ}_{\tilde{P}}] \cdot (N + (1 - \alpha/2) \cdot \mathbb{E}[\text{call}]) \\ \implies \mathbb{E}[\text{call}] &\leq \frac{2}{\alpha\tilde{\varepsilon}} \cdot \left(1 + \tilde{\varepsilon} \cdot \frac{\ln(2)}{(1 - \alpha)\tilde{\varepsilon} - \varepsilon} \right), \end{aligned}$$

which gives an expected number of calls $\text{poly}(\lambda, (\tilde{\varepsilon} - \varepsilon)^{-1})$ by setting $\alpha \leftarrow (1 - \varepsilon/\tilde{\varepsilon})/2$ (corresponding to $(1 - \alpha)\tilde{\varepsilon} = (\varepsilon + \tilde{\varepsilon})/2$). This concludes the proof.

From soundness to EUF-KO security. Given a five-round identification protocol where the probability of sampling a “bad” Round 3 challenge is bounded by \mathfrak{p} , and the probability of sampling a “bad” Round 5 challenge is bounded by $1/n$, it follows from a standard application of the Fiat-Shamir methodology (adapted to 5-round protocols) to the τ -fold parallel repetition of the identification scheme given in Figure 7 that, when modeling H_1 and H_2 with random oracles, there exists an extractor which extracts a weakly valid witness $\mathbf{x}^* \in X$ given any adversary that succeeds with probability at least $\Pr[X + Y = \tau]$, where with probability at least $1 - \varepsilon_G$ over the random choice of (H, \mathbf{y}) , it holds that $X = \max_{\alpha \in Q_1} \{X_\alpha\}$ and $Y = \max_{\beta \in Q_2} \{Y_\beta\}$ with $\mathfrak{p} = 4/B$, $X_\alpha \sim \text{Binomial}(\tau, \mathfrak{p})$, and $Y_\beta \sim \text{Binomial}(\tau - X, \frac{1}{n})$ where Q_1 and Q_2 are sets of all queries to the oracles H_1 and H_2 . With probability at least $1 - 1/2^\lambda$, this weakly valid witness is necessarily a regular witness. This concludes the proof.

7 Parameter Selection and Efficiency

7.1 Parameters Selection Process

In this section, we explain how to select parameters for our new signature scheme. The first goal is to pick parameters that minimize the number of repetitions τ of the underlying identification scheme, since this parameter has a large impact on the signature size. Concretely, as in previous works, we choose τ such that the cost of the forgery attack on the Fiat-Shamir-compiled signature is at least 2^{128} , where cost is given by the formula below (which comes from the attack of Kales and Zaverucha [KZ20]):

$$\text{cost} = \min_{\tau_1, \tau_2: \tau_1 + \tau_2 = \tau} \left\{ \frac{1}{\sum_{i=\tau_1}^{\tau} \binom{\tau}{i} \mathbf{p}^i (1-\mathbf{p})^{\tau-i}} + n^{\tau_2} \right\}. \quad (1)$$

We observe that setting $B = 2^{-134}$ in $\mathbf{p} = 4/B$ suffices to guarantee that τ is always the smallest possible for any given number of leaves $n = 2^D$ (i.e. reducing \mathbf{p} further does not reduce τ). The choice of the number of leaves, 2^D , is a tradeoff parameter: larger values of D yield smaller signature size, at the expense of a larger runtime.

Finding a bound on k . A crucial aspect of our parameter selection process is that our combinatorial analysis in section 6.2 only guarantees that with very high probability, for any $(\mathbf{u}, \mathbf{x}^*) \in \text{PN}_B \times [\text{bs}]^w$ where \mathbf{u} is not regular, $\pi(\mathbf{u}) \downarrow \mathbf{x}^*$ will not be a valid solution \mathbf{v} to $H \cdot \mathbf{v} = \mathbf{y}$. However, it says nothing about vectors \mathbf{u} outside PN_B , that is, vectors with low permutation number $\text{pn}(\mathbf{u}) \leq B$. Therefore, we must select RSD parameters such that, with overwhelming probability, there will not be *any* solution \mathbf{v} to $H \cdot \mathbf{v} = \mathbf{y}$ of the form $\mathbf{v} = \mathbf{u} \downarrow \mathbf{x}^*$ for $\mathbf{u} \in \mathbb{F}_2^K \setminus \text{PN}_B$. Since we defined the set $X = \{\mathbf{v} \in \mathbb{F}_2^K : \exists \mathbf{u} \in \mathbb{F}_2^K \setminus \text{PN}_B, \exists \mathbf{x}^* \in [\text{bs}]^w, \mathbf{v} = \mathbf{u} \downarrow \mathbf{x}^*\}$, to guarantee that there will not be any solution $\mathbf{v} \in X$ to $H \cdot \mathbf{v} = \mathbf{y}$, it suffices to pick $\log_2 k \geq |X| + \lambda$. This follows from a standard ‘‘Gilbert-Varshamov-style’’ analysis: when sampling a random instance $(H, \mathbf{y} = H \cdot \mathbf{x})$ of the RSD problem, the expected number of solutions in X (beyond \mathbf{x}) is

$$\mathbb{E}_{H, \mathbf{x}} [|\{\mathbf{x}' : H \cdot \mathbf{x}' = H \cdot \mathbf{x} \wedge \mathbf{x}' \in X\}|] = \sum_{\substack{\mathbf{x}' \neq \mathbf{x} \\ \mathbf{x}' \in X}} \Pr_{H, \mathbf{x}}[H \cdot \mathbf{x}' = H \cdot \mathbf{x}] = \frac{|X| - 1}{2^k},$$

and we conclude with a Markov bound. To choose k , we use a bound on $|X|$:

Lemma 24. *Let $\text{P}_{i,w}$ denote the set of integer partitions of w in i parts, i.e., the set of all tuples (k_1, \dots, k_i) with $0 < k_1 \leq k_2 \leq \dots \leq k_i \leq w$ such that $\sum_{j=1}^i k_j = w$. Let T_B denote the function such that $\text{T}_B(x) = x$ when $x \leq B$, and $\text{T}_B(x) = 0$ when $x > B$. Then*

$$|X| \leq \text{bs}^w \cdot \sum_{i=1}^L \left(\binom{L}{i} \cdot i! \cdot \sum_{(k_1, \dots, k_i) \in \text{P}_{i,w}} \text{T}_B \left(\frac{w!}{\prod_{j=1}^i (k_j)!} \right) \right),$$

where L is (using the Euler totient ϕ and denoting $a|b$ for ‘‘ a divides b ’’):

$$L = \frac{1}{\text{bs}} \cdot \sum_{\substack{i \leq \text{bs} \\ i \text{ odd}}} \sum_{d | \gcd(\text{bs}-i, i)} \phi(d) \cdot \binom{\text{bs}/d}{i/d}.$$

The proof of Lemma 24 follows from a counting argument, which we detail below. We enumerate over the bs^w possible vectors $\mathbf{x}^* \in [\text{bs}]^w$, and over all possible \mathbf{u} with $\text{pn}(\mathbf{u}) \geq B$. To count the latter, we proceed in steps:

Counting the number of distinct blocks. We compute the number L of possible distinct blocks. A loose upper bound would be $L \leq 2^{\text{bs}}$ (since a block is a vector in \mathbb{F}_2^{bs}). However, because we already enumerate over all possible shifts \mathbf{x}^* of the w blocks, we must only count the number of distinct blocks *up to cyclic shift*. In combinatorics, this amounts to counting the number of length- bs necklaces with two colors. Additionally, because of the optimization given in Section 6 where the last entry of each block is fixed such that all entries of a block XOR to 1, we only need to enumerate over all necklaces with an odd number of ones. The formula for L in Lemma 24 is a direct application of Pólya’s enumeration theorem [Red27], a classical theorem on the combinatorics of necklaces.

Counting the number of vectors. For $i = 1$ to L , we count the number of vectors which have exactly i distinct blocks. There are $\binom{L}{i}$ ways to select the i distinct blocks out of L possible blocks. Since each vector has w blocks in total, we enumerate over all partitions of the integer w in exactly i parts $0 < k_1 \leq k_2 \leq \dots \leq k_i \leq w$, where k_j denotes the number of copies of the j -th block from the selection. Because we enumerate over *ordered* partitions, we are ordering the i selected blocks by number of copies; hence, we multiply by $i!$ to account for all possible configurations of number of copies (this is a slightly loose upper bound, since some partitions may have equal numbers $k_j = k_{j+1}$:

the right value would be to multiply by the factorial of the number of distinct integers in (k_1, \dots, k_j) , but we ignore this for simplicity). Then, having fixed a choice of i specific distinct blocks and the numbers (k_1, \dots, k_i) of copies of each block, there are $w! / \prod_{j=1}^i (k_j)!$ distinct blockwise permutations of (this is the standard combinatorial formula for counting multisets). Eventually, since we only want to keep vectors whose permutation number is at most B , we keep only in the count the vectors for which $w! / \prod_{j=1}^i (k_j)! \leq B$ (this is the purpose of the threshold function T_B in the formula). This yields the formula stated in Lemma 24.

Computing $|X|$. It remains to compute explicitly the formula of Lemma 24. We used a Python script to perform the calculation. A small nontriviality is that enumerating over all integer partitions of w (which is around 120) would be very slow. Fortunately, we observe that the condition $w! / \prod_{j=1}^i (k_j)! \leq B$, together with the bound L on i , impose a sharp bound on the value of k_i : a quick calculation shows that we need $k_i \geq w/2$ to be such that $\binom{w}{k_i} \leq B$. Given this bound on k_i , we enumerate over all remaining possible values of k_i , and compute the number of partitions of $w - k_i$ into $i - 1$ parts to obtain the rest of the partition. We provide the script used to compute this bound in Appendix B of the Supplementary Material.

Finding (K, w) . To find the RSD parameters (K, k, w) , we proceed iteratively: we fix a choice of K, w and compute the value of k as $|X| + 128$ (note that X depends on (K, w)), using our script to compute the bound on $|X|$ from Lemma 24. Then, we rely on the estimator implemented in state-of-the-art cryptanalysis of [ES23] (which improves over a previous cryptanalysis from [CCJ23]) to compute (an estimate of) the bit security of the instance obtained against all known attacks on RSD. If the bit security is below 128, we increase K by 1 and start over (every time, we also compute the parameters for a list of weight parameters w , since the impact of w on the proof size is slightly subtle). Eventually, after settling for a choice of (K, k, w) , we check that the probability bound of Lemma 19 is overwhelming (with our choice of parameters, it is always above $1 - 2^{-200}$).

7.2 Concrete Parameters and Implementation

We outline below a few parameter sets, for different values of $D \in \{8, \dots, 17\}$. For all values of D , the smallest signature size was achieved by setting $K = 1736$, $k = 960$, $w = 217$, $\mathsf{bs} = 8$.

Table 3. Signature size and signing time for various values of D , using the parameter set $K = 1736$, $k = 960$, $w = 217$, $\mathsf{bs} = 8$. All timings computing on one core of an Intel Core i7 processor 14700KF.

D	τ	$ \sigma $	signing time	verification time
8	16	7.8 kB	1.47 ms	1.51 ms
9	15	7.6 kB	3.53 ms	3.58 ms
10	13	6.8 kB	6.32 ms	6.80 ms
11	12	6.5 kB	12.1 ms	13.3 ms
12	11	6.1 kB	19.1 ms	21.3 ms
13	10	5.7 kB	44.1 ms	50.0 ms
15	9	5.4 kB	140 ms	166 ms
16	8	4.9 kB	-	-

We implemented our signature scheme in C. Our implementation is a proof of concept implementation: we did not use any of the optimizations such as batching, vectorization, or bit slicing, and an optimized implementation can likely achieve significantly faster runtime. We used the AES-NI instruction set to implement our multi-instance PPRF and our multi-instance PRG from Section 5. All our experiments were run on one core of an Intel Core i7 processor 14700KF. The following optimization flags have been used during compilation: `-O3 -flto -mavx2 -mpclmul -msse4.2 -maes -rdnd`. We note that our verification time is slightly slower than our signing time. This is an artifact of our proof-of-concept implementation: a more properly optimized implementation would have verification run slightly faster than signing. We plan to optimize our implementation in the future.

References

- ABB⁺23a. N. Aaraj, S. Bettaieb, L. Bidoux, A. Budroni, V. Dyseryn, A. Esser, P. Gaborit, M. Kulkarni, V. Mateu, M. Palumbi, L. Perin, and J. Tillich. PERK. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- ABB⁺23b. N. Aragon, M. Bardet, L. Bidoux, J.-J. Chi-Domínguez, V. Dyseryn, T. Feneuil, P. Gaborit, A. Joux, M. Rivain, J. Tillich, and A. Vinçotte. RYDE. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- ABB⁺23c. N. Aragon, M. Bardet, L. Bidoux, J. Chi-Domínguez, V. Dyseryn, T. Feneuil, P. Gaborit, R. Neveu, M. Rivain, and J. Tillich. MIRA. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- AFG⁺23. C. Aguilar-Melchor, T. Feneuil, N. Gama, S. Gueron, J. Howe, D. Joseph, A. Joux, E. Persichetti, T. H. Randrianarisoa, M. Rivain, and D. Yue. SDitH — Syndrome Decoding in the Head. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- AFS03. D. Augot, M. Finiasz, and N. Sendrier. A fast provably secure cryptographic hash function. Cryptology ePrint Archive, Report 2003/230, 2003. <https://eprint.iacr.org/2003/230>.
- AGH⁺23. C. Aguilar-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, Heidelberg, April 2023.
- AHIV17. S. Ames, C. Hazay, Y. Ishai, and M. Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *ACM CCS 2017*, pages 2087–2104. ACM Press, October / November 2017.
- ARV⁺23. G. Adj, L. Rivera-Zamarripa, J. Verbel, E. Bellini, S. Barbero, A. Esser, C. Sanna, and F. Zwendinger. MiRitH — MinRank in the Head. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- BBdSG⁺23. C. Baum, L. Braun, C. D. de Saint Guilhem, M. Kloof, E. Orsini, L. Roy, and P. Scholl. Publicly verifiable zero-knowledge and post-quantum signatures from vole-in-the-head. In *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part V, Lecture Notes in Computer Science 14085*, pages 581–615. Springer, 2023.
- BCdSG24. D. Bui, K. Cong, and C. D. de Saint Guilhem. Improved all-but-one vector commitment with applications to post-quantum signatures. Cryptology ePrint Archive, Paper 2024/097, 2024. <https://eprint.iacr.org/2024/097>.
- BCG⁺19a. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, P. Rindal, and P. Scholl. Efficient two-round OT extension and silent non-interactive secure computation. In *ACM CCS 2019*, pages 291–308. ACM Press, November 2019.
- BCG⁺19b. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In *CRYPTO 2019, Part III, LNCS 11694*, pages 489–518. Springer, Heidelberg, August 2019.
- BCG⁺20a. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl. Correlated pseudorandom functions from variable-density LPN. In *61st FOCS*, pages 1069–1080. IEEE Computer Society Press, November 2020.
- BCG⁺20b. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl. Efficient pseudorandom correlation generators from ring-LPN. In *CRYPTO 2020, Part II, LNCS 12171*, pages 387–416. Springer, Heidelberg, August 2020.
- BCG⁺21. E. Boyle, N. Chandran, N. Gilboa, D. Gupta, Y. Ishai, N. Kumar, and M. Rathee. Function secret sharing for mixed-mode and fixed-point secure computation. In *EUROCRYPT 2021, Part II, LNCS 12697*, pages 871–900. Springer, Heidelberg, October 2021.
- BCG⁺22. E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, N. Resch, and P. Scholl. Correlated pseudorandomness from expand-accumulate codes. In *CRYPTO 2022, Part II, LNCS 13508*, pages 603–633. Springer, Heidelberg, August 2022.
- BCGI18. E. Boyle, G. Couteau, N. Gilboa, and Y. Ishai. Compressing vector OLE. In *ACM CCS 2018*, pages 896–912. ACM Press, October 2018.
- BDK⁺21. C. Baum, C. Delpech de Saint Guilhem, D. Kales, E. Orsini, P. Scholl, and G. Zaverucha. Banquet: Short and fast signatures from AES. In *PKC 2021, Part I, LNCS 12710*, pages 266–297. Springer, Heidelberg, May 2021.
- BGI14. E. Boyle, S. Goldwasser, and I. Ivan. Functional signatures and pseudorandom functions. In *PKC 2014, LNCS 8383*, pages 501–519. Springer, Heidelberg, March 2014.

- BGI15. E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing. In *EUROCRYPT 2015, Part II, LNCS 9057*, pages 337–367. Springer, Heidelberg, April 2015.
- BGI16. E. Boyle, N. Gilboa, and Y. Ishai. Function secret sharing: Improvements and extensions. In *ACM CCS 2016*, pages 1292–1303. ACM Press, October 2016.
- BGI19. E. Boyle, N. Gilboa, and Y. Ishai. Secure computation with preprocessing via function secret sharing. In *TCC 2019, Part I, LNCS 11891*, pages 341–371. Springer, Heidelberg, December 2019.
- BKPV23. L. Bettale, D. Kahrobaei, L. Perret, and J. Verbel. Biscuit. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- BLPS11. D. J. Bernstein, T. Lange, C. Peters, and P. Schwabe. Really fast syndrome-based hashing. In *AFRICACRYPT 11, LNCS 6737*, pages 134–152. Springer, Heidelberg, July 2011.
- BW13. D. Boneh and B. Waters. Constrained pseudorandom functions and their applications. In *ASIACRYPT 2013, Part II, LNCS 8270*, pages 280–300. Springer, Heidelberg, December 2013.
- 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, Heidelberg, April 2023.
- CDG⁺20. M. Chase, D. Derler, S. Goldfeder, J. Katz, V. Kolesnikov, C. Orlandi, S. Ramacher, C. Reicherberger, D. Slamanig, X. Wang, et al. The picnic signature scheme. *Submission to NIST Post-Quantum Cryptography project*, 2020.
- 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*, pages 229–258. Springer, Heidelberg, May 2024.
- CRR21. G. Couteau, P. Rindal, and S. Raghuraman. Silver: Silent VOLE and oblivious transfer from hardness of decoding structured LDPC codes. In *CRYPTO 2021, Part III, LNCS 12827*, pages 502–534, Virtual Event, August 2021. Springer, Heidelberg.
- CS14. S. Chen and J. P. Steinberger. Tight security bounds for key-alternating ciphers. In *EUROCRYPT 2014, LNCS 8441*, pages 327–350. Springer, Heidelberg, May 2014.
- DDOS19. C. Delpech de Saint Guilhem, L. De Meyer, E. Orsini, and N. P. Smart. BBQ: Using AES in picnic signatures. In *SAC 2019, LNCS 11959*, pages 669–692. Springer, Heidelberg, August 2019.
- DN19. I. Dinur and N. Nadler. Multi-target attacks on the Picnic signature scheme and related protocols. In *EUROCRYPT 2019, Part III, LNCS 11478*, pages 699–727. Springer, Heidelberg, May 2019.
- ES23. A. Esser and P. Santini. Not just regular decoding: Asymptotics and improvements of regular syndrome decoding attacks. *Cryptology ePrint Archive*, 2023.
- FGS07. M. Finiasz, P. Gaborit, and N. Sendrier. Improved fast syndrome based cryptographic hash functions. In *Proceedings of ECRYPT Hash Workshop*, page 155. Citeseer, 2007.
- 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, Heidelberg, August 2022.
- FR23. T. Feneuil and M. Rivain. MQOM — MQ on my Mind. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>.
- FS87. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO’86, LNCS 263*, pages 186–194. Springer, Heidelberg, August 1987.
- GGM86. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986.
- GI14. N. Gilboa and Y. Ishai. Distributed point functions and their applications. In *EUROCRYPT 2014, LNCS 8441*, pages 640–658. Springer, Heidelberg, May 2014.
- GKWY20. C. Guo, J. Katz, X. Wang, and Y. Yu. Efficient and secure multiparty computation from fixed-key block ciphers. In *2020 IEEE Symposium on Security and Privacy*, pages 825–841. IEEE Computer Society Press, May 2020.
- GYW⁺23. X. Guo, K. Yang, X. Wang, W. Zhang, X. Xie, J. Zhang, and Z. Liu. Half-tree: Halving the cost of tree expansion in COT and DPF. In *EUROCRYPT 2023, Part I, LNCS 14004*, pages 330–362. Springer, Heidelberg, April 2023.
- HJ24. J. Huth and A. Joux. MPC in the head using the subfield bilinear collision problem. In *CRYPTO 2024*, 2024.
- HOSS18. C. Hazay, E. Orsini, P. Scholl, and E. Soria-Vazquez. TinyKeys: A new approach to efficient multi-party computation. In *CRYPTO 2018, Part III, LNCS 10993*, pages 3–33. Springer, Heidelberg, August 2018.
- IKOS07. Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. Zero-knowledge from secure multiparty computation. In *39th ACM STOC*, pages 21–30. ACM Press, June 2007.

- KPTZ13. A. Kiayias, S. Papadopoulos, N. Triandopoulos, and T. Zacharias. Delegatable pseudorandom functions and applications. In *ACM CCS 2013*, pages 669–684. ACM Press, November 2013.
- KZ20. D. Kales and G. Zaverucha. An attack on some signature schemes constructed from five-pass identification schemes. In *CANS 20, LNCS 12579*, pages 3–22. Springer, Heidelberg, December 2020.
- MDCYA11. M. Meziani, Ö. Dagdelen, P.-L. Cayrel, and S. M. E. Yousfi Alaoui. S-fsb: An improved variant of the fsb hash family. In *International Conference on Information Security and Assurance*, pages 132–145. Springer, 2011.
- MSY21. J.-P. Münch, T. Schneider, and H. Yalame. Vasa: Vector aes instructions for security applications. In *Annual Computer Security Applications Conference*, pages 131–145, 2021.
- Pat09. J. Patarin. The “coefficients h” technique. In *Selected Areas in Cryptography*, pages 328–345, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- Red27. J. H. Redfield. The theory of group-reduced distributions. *American Journal of Mathematics*, 49(3):433–455, 1927.
- Roy22. L. Roy. SoftSpokenOT: Quieter OT extension from small-field silent VOLE in the minicrypt model. In *CRYPTO 2022, Part I, LNCS 13507*, pages 657–687. Springer, Heidelberg, August 2022.
- RS21. P. Rindal and P. Schoppmann. VOLE-PSI: Fast OPRF and circuit-PSI from vector-OLE. In *EUROCRYPT 2021, Part II, LNCS 12697*, pages 901–930. Springer, Heidelberg, October 2021.
- WYKW21. C. Weng, K. Yang, J. Katz, and X. Wang. Wolverine: Fast, scalable, and communication-efficient zero-knowledge proofs for boolean and arithmetic circuits. In *2021 IEEE Symposium on Security and Privacy*, pages 1074–1091. IEEE Computer Society Press, May 2021.
- YSWW21. K. Yang, P. Sarkar, C. Weng, and X. Wang. QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In *ACM CCS 2021*, pages 2986–3001. ACM Press, November 2021.
- YWL⁺20. K. Yang, C. Weng, X. Lan, J. Zhang, and X. Wang. Ferret: Fast extension for correlated OT with small communication. In *ACM CCS 2020*, pages 1607–1626. ACM Press, November 2020.
- ZCD⁺20. G. Zaverucha, M. Chase, D. Derler, S. Goldfeder, C. Orlandi, S. Ramacher, C. Rechberger, D. Slamanig, J. Katz, X. Wang, V. Kolesnikov, and D. Kales. Picnic. Technical report, National Institute of Standards and Technology, 2020. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>.

Supplementary Material

A Further Applications of Multi-Instance PPRFs

We expect our tight analysis of fast PPRFs in the multi-instance setting to find applications beyond the realm of MPCitH signatures. As a sample application, pseudorandom correlation functions [BCG⁺20a, BCG⁺22], which are used to efficiently generate correlated randomness in secure computation, are typically constructed using a large number of distributed point functions (DPFs). DPFs are very similar to GGM-style PPRFs, and we expect our analysis to extend almost immediately to multi-instance DPFs. Using to the attack which we describe in Section 3, the concrete security of PCFs using N copies of a DPF with a tree of depth D is $2^\lambda / (N \cdot D)$. In many settings, this turns out to be a significant security loss: for example, using the PCF of [BCG⁺22] to generate 2^{30} degree-2 correlations requires $N = 664^2$ copies of a GGM tree of depth $\log_2(2^{30}/664)^2$ (in fact, a 2-dimensional GGM tree). With the collision attack, this translates to a concrete loss of 27 bits of security. Using our methodology to extend GGM PPRFs to the multi-instance setting, the security loss could be reduced to 5 bits, without any sacrifice on efficiency. We also expect that our proof technique could be used to improve the parameters of other schemes. For example, the recent work of [BCdSG24] also takes advantage of AES using half-tree construction [GYW⁺23] to improve VOLEitH signatures. The scheme of [BCdSG24] uses a direct construction of half-tree based on the circular collision-resistant hash function (CCR), and due to the limitation of the security of CCR hash (to 128-bit blocks and key size) when using an AES-based instantiation, the size of tree leaves have been extended to $2\lambda = 256$ bits to reach 128 bits of security. We expect that the techniques developed in the concrete security analysis of our multi-instance PPRF using the H-coefficient technique can be applied to security proof of [BCdSG24] to get rid of the need to expand the size of the last layer.

B Python Script

Below, we provide the Python script used to compute the value of $k \geq |X| + \lambda$ using the bound on $|X|$ from Lemma 24. The code for listing all partitions of an integer into i parts was taken from Stackoverflow.¹¹

```
import numpy as np
import math
from math import factorial
from math import gcd
from math import comb
from math import log

def phi(n):
    amount = 0
    for k in range(1, n + 1):
        if gcd(n, k) == 1:
            amount += 1
    return amount

def L(blocksize):
    num = 0
    for k in range(1, blocksize+1):
        if k%2 == 1:
            D = gcd(blocksize-k, k)
            val = 0
            for d in range(1, D+1):
                if math.gcd(d, D)==d:
                    val += phi(d)*factorial(int(blocksize/d))/(factorial(int((blocksize-k)/d))*factorial(int(k/d)))
```

¹¹ <https://stackoverflow.com/questions/18503096/python-integer-partitioning-with-given-k-partitions>

```

    num += val
    return int(num/blocksize)

def part(n, k):
    def memoize(f):
        cache = [[None] * n for j in range(k)] for i in range(n)]
        def wrapper(n, k, pre):
            if cache[n-1][k-1][pre-1] is None:
                cache[n-1][k-1][pre-1] = f(n, k, pre)
            return cache[n-1][k-1][pre-1]
        return wrapper
    @memoize
    def _part(n, k, pre):
        if n <= 0:
            return []
        if k == 1:
            if n <= pre:
                return [(n,)]
            return []
        ret = []
        for i in range(min(pre, n), 0, -1):
            ret += [(i,) + sub for sub in _part(n-i, k-1, i)]
        return ret
    return _part(n, k, n)

def max_numblocks(w,B):
    # return the largest N such that u can
    # have N distinct blocks yet generates
    # less than B permutations

    N = 1
    numb = 1
    for j in range(w-N+2,w+1):
        numb *= j
    while numb <= B:
        N += 1
        numb = 1
        for j in range(w-N+2,w+1):
            numb *= j
    return N-1

def min_k(w,B):
    # returns the smallest k (above w/2, by
    # symmetry) such that c = (w choose k)
    # is less than or equal to B

    k = int(w/2)
    c = math.comb(w,k)
    while c > B:
        k += 1
        c = math.comb(w,k)
    return (c,k)

def gvbound(w,bs,B,b=1,secp=128):
    K = w*bs
    l = L(bs)
    N = max_numblocks(w,B)
    res = 0
    for n in range(2,N+1):
        # starts at 2 to avoid the empty
        # partition later when we isolate the
        # biggest element

        val = 0
        (c,k) = min_k(w,B)
        # computes the smallest k (above w/2, by
        # symmetry) such that c = (w choose k)
        # is less than or equal to B

```

```

for i in range(k, w-n+2):          # for all possible values i of the first
                                  term in the partition of w (from k to
                                  w-(N-1))
    partitions = part(w-i,n-1)    # stores all partitions of the integer
                                  w-i into N-1 parts

    fact = 1
    for j in range(i+1, w+1):
        fact *= j                # computes w!/i!, a common factor of the terms
                                  summed over all partitions of w-i

    for partition in partitions:
        term = fact
        for v in partition:
            term = term//factorial(v) # computes w!/(i! * prod_v v!), where v
                                       goes over the elements of a partition
                                       of w-i

        if term <= B:            # checks that the term w!/(i! * prod_v v!) is
                                  below the bound B

            val += term          # val accumulates the terms of the sum accross
                                  all partitions of w-i, then accross
                                  all values of i

    res += factorial(n)*math.comb(l, n)*val # res accumulates the
                                             result accross all possible numbers n
                                             of blocks, from 1 to N = max_numblocks
                                             (w,B)

dim = math.log((res+1)*(bs**w),2) # dim is the base 2 logarithm of (K/w
                                  )^w * (the sum computed so far). /\
                                  will need to add lambda later

                                  # also adds 1 to the result because we excluded n=1
return dim + b*secpair

```