

Syndrome Decoding in the Head: Shorter Signatures from Zero-Knowledge Proofs

Thibault Feneuil^{1,2}, Antoine Joux³, and Matthieu Rivain¹

¹ CryptoExperts, Paris, France

² Sorbonne Université, CNRS, INRIA, Institut de Mathématiques
de Jussieu-Paris Rive Gauche, Ouragan, Paris, France

³ CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

{thibault.feneuil,matthieu.rivain}@cryptoexperts.com
joux@cispa.de

Abstract. Zero-knowledge proofs of knowledge are useful tools to design signature schemes. The ongoing effort to build a quantum computer urges the cryptography community to develop new secure cryptographic protocols based on quantum-hard cryptographic problems. One of the few directions is code-based cryptography for which the strongest problem is the *syndrome decoding* (SD) for random linear codes. This problem is known to be NP-hard and the cryptanalysis state of the art has been stable for many years. A zero-knowledge protocol for this problem was pioneered by Stern in 1993. Since its publication, many articles proposed optimizations, implementation, or variants.

In this paper, we introduce a new zero-knowledge proof for the syndrome decoding problem on random linear codes. Instead of using permutations like most of the existing protocols, we rely on the MPC-in-the-head paradigm in which we reduce the task of proving the low Hamming weight of the SD solution to proving some relations between specific polynomials. Specifically, we propose a 5-round zero-knowledge protocol that proves the knowledge of a vector x such that $y = Hx$ and $\text{wt}(x) \leq w$ and which achieves a soundness error closed to $1/N$ for an *arbitrary* N . While turning this protocol into a signature scheme, we achieve a signature size of 11-12 KB for 128-bit security when relying on the hardness of the SD problem on binary fields. Using larger fields (like \mathbb{F}_{2^8}), we can produce fast signatures of around 8 KB. This allows us to outperform Picnic3 and to be competitive with SPHINCS⁺, both post-quantum signature candidates in the ongoing NIST standardization effort. Moreover, our scheme outperforms all the existing code-based signature schemes for the common “signature size + public key size” metric.

1 Introduction

Zero-knowledge proofs are an important tool for many cryptographic protocols and applications. Such proofs enable a *prover* to prove a statement by interacting with a *verifier* without revealing anything more than the statement itself. Zero-knowledge proofs find application in many contexts. Thanks to the Fiat-Shamir transform [FS87], we can convert such proofs into signature schemes.

In this article, we aim to build an efficient code-based signature scheme using this methodology. To do so, we will focus on the *generic decoding* problem, a.k.a. the (computational) *syndrome decoding* (SD) problem: given a matrix $H \in \mathbb{F}_q^{(m-k) \times m}$ and a vector $y \in \mathbb{F}_q^{m-k}$, recover a *small-weight* vector $x \in \mathbb{F}_q^m$ such that $Hx = y$. For random linear codes –i.e. for a random matrix H – this problem is known to be NP-hard and widely believed to be robust for practical sets of parameters.

In a pioneering work from three decades ago, Stern proposed a zero-knowledge protocol to prove the knowledge of a syndrome decoding solution [Ste94]. This protocol achieves a *soundness error* of $2/3$ which means that a malicious prover can fool the verifier with a $2/3$ probability. Although an arbitrary security of $(2/3)^\tau$ can be achieved by repeating the protocol τ times, the induced communication cost for standard security levels (e.g. 128 bits) becomes significant, which is partly due to this high soundness error. Since the work of Stern, a few papers have proposed optimizations and implementations of this protocol (see for instance [Vér96,GG07,AGS11,ACBH13]) but the communication cost was still heavy for random linear codes with standard security levels.

In 2007, Ishai, Kushilevitz, Ostrovsky and Sahai proposed a new technique to build zero-knowledge proofs from secure multi-party computation (MPC) protocols, which is known as the *MPC-in-the-Head* (MPCitH) paradigm [IKOS07]. While this construction was mainly considered of theoretical interest at first, it has been increasingly applied to build practical schemes over the last years. In particular, the Picnic post-quantum signature scheme [CDG⁺], which is a third-round alternate candidate of the ongoing NIST standardization effort, is based on the MPCitH principle. Recently, new zero-knowledge protocols for the SD problem have been inspired by this principle [?,?,BGKM22]. In particular, these protocols achieve an arbitrary soundness error $1/N$ instead of the $2/3$ (or $1/2$) of Stern protocol and variants. These protocols result in smaller proof/signature sizes at the cost of computational overheads.

Our contribution. In this article, we build a new zero-knowledge protocol to prove the knowledge of a syndrome decoding solution using the MPCitH paradigm. We further turn this protocol into an efficient code-based signature scheme.

While proving that $y = Hx$ is communication-free in this paradigm, the hard part consists in proving that x is a small-weight vector. We propose here an efficient way to prove that $\text{wt}(x) \leq w$ through a multi-party computation which is simulated by the prover (“in her head”). The key idea is to prove the equality $x \circ v = 0$ where \circ is the component-wise multiplication and where the coefficients of the vector v are the evaluations of a polynomial Q of degree w . By definition, v has at most w zero coordinates, so the relation $x \circ v = 0$ proves that x has at most w non-zero coordinates (i.e. $\text{wt}(x) \leq w$). **The roots of the polynomial Q give the non-zero positions of the vector x .** In order to prove the relation $x \circ v = 0$, we use techniques borrowed from the Banquet signature scheme [BdK⁺21] with further adaptations. To check that all $x_j \cdot v_j$ are equal to zero, we arrange the input x into a polynomial S , provide a product polynomial $F \cdot P$ as part of the witness, and check that $(F \cdot P)(\cdot)$ indeed equals

the product of $S(\cdot)$ and $Q(\cdot)$. This can be done efficiently by only verifying a few products of these polynomials evaluated at some random points. However, instead of revealing the multiplication operands like in [BdK⁺21], we rely on the product checking protocol proposed in [LN17,BN20] and its batch version recently introduced in [KZ21].

Let us note that the idea of encoding the non-zero positions in a polynomial has been already used in [?] in order to have compact proofs proving that a secret vector is of Hamming weight lower or equal to some constant. However, the proposed zero-knowledge protocol relies on a linearly homomorphic commitment scheme, and such schemes do not exist yet in the post-quantum cryptography.

Thanks to the Fiat-Shamir transform [FS87], we convert our protocol into a signature scheme. Our scheme outperforms all the existing code-based signatures for the “signature size + public key size” metric. When relying on the hardness of the syndrome decoding problem over \mathbb{F}_{256} , our scheme is below 10 KB for this metric, which makes it competitive with Picnic3 [KZ20b] and SPHINCS⁺ [BHK⁺19]. Compared to other code-based signature schemes (such as Wave [DST19] and Durandal [ABG⁺19]), our scheme has the significant advantage of relying on a non-structured NP-hard decoding problem which has been widely studied over the last decades.

To provide more flexibility, we introduce a parameter d in the definition of the syndrome decoding problem. The idea is, instead of having a constraint for the global weight of the secret vector x , to split x into d chunks $x := (x_1 \mid \dots \mid x_d)$ and to have a constraint on the weight of each chunk. By taking $d = 1$, this *d-split* version is equivalent to the standard syndrome decoding problem. We provide a security reduction from this variant to the standard problem which allows us to compensate the security loss by a slight increase of the parameters. This so-called *d-split syndrome decoding problem* offers us more flexibility to find better size-performance trade-offs for our signature scheme.

Paper organization. The paper is organized as follows: In Section 2, we introduce the necessary background on the syndrome decoding problem, zero-knowledge proofs, and the MPC-in-the-Head paradigm. We present our protocol in Section 3 and the signature scheme obtained through the Fiat-Shamir transform in Section 4. To conclude, we provide implementation results and compare our construction with other signature schemes from the state of the art in Section 5.

2 Preliminaries

Throughout the paper, \mathbb{F} shall denote a finite field. For any vector $x \in \mathbb{F}^m$, the *Hamming weight* of x , denoted $\text{wt}(x)$, is the number of non-zero coordinates of x . For two vectors $x_1 \in \mathbb{F}^{m_1}$ and $x_2 \in \mathbb{F}^{m_2}$, we denote $(x_1 \mid x_2) \in \mathbb{F}^{m_1+m_2}$ their concatenation. We denote \circ the component-wise multiplication between two vectors. For any $m \in \mathbb{N}^*$, the integer set $\{1, \dots, m\}$ is denoted $[m]$. For a probability distribution D , the notation $s \leftarrow D$ means that s is sampled from D . For a finite set S , the notation $s \leftarrow S$ means that s is uniformly sampled

at random from S . When the set S is clear from the context, we sometimes denote $s \leftarrow \$$ for a uniform random sampling of s from S . For an algorithm \mathcal{A} , $out \leftarrow \mathcal{A}(in)$ further means that out is obtained by a call to \mathcal{A} on input in (using uniform random coins whenever \mathcal{A} is probabilistic). Along the paper, probabilistic polynomial time is abbreviated PPT.

A function $\mu : \mathbb{N} \rightarrow \mathbb{R}$ is said *negligible* if, for every positive polynomial $p(\cdot)$, there exists an integer $N_p > 0$ such that for every $\lambda > N_p$, we have $|\mu(\lambda)| < 1/p(\lambda)$. When not made explicit, a negligible function in λ is denoted $\text{negl}(\lambda)$ while a polynomial function in λ is denoted $\text{poly}(\lambda)$. We further use the notation $\text{poly}(\lambda_1, \lambda_2, \dots)$ for a polynomial function in several variables.

Two distributions $\{D_\lambda\}_\lambda$ and $\{E_\lambda\}_\lambda$ indexed by a security parameter λ are (t, ε) -*indistinguishable* (where t and ε are $\mathbb{N} \rightarrow \mathbb{R}$ functions) if, for any algorithm \mathcal{A} running in time at most $t(\lambda)$, we have

$$|\Pr[\mathcal{A}(x) = 1 \mid x \leftarrow D_\lambda] - \Pr[\mathcal{A}(x) = 1 \mid x \leftarrow E_\lambda]| \leq \varepsilon(\lambda) .$$

The two distributions are said

- *computationally indistinguishable* if $\varepsilon \in \text{negl}(\lambda)$ for every $t \in \text{poly}(\lambda)$;
- *statistically indistinguishable* if $\varepsilon \in \text{negl}(\lambda)$ for every (unbounded) t ;
- *perfectly indistinguishable* if $\varepsilon = 0$ for every (unbounded) t .

In this paper, we shall use the standard cryptographic notions of secure pseudo-random generator (PRG), tree PRG, collision-resistant hash function, and (hiding and binding) commitment scheme. Those notions are formally defined in the full version [?].

2.1 Syndrome Decoding Problems

Definition 1 (Syndrome Decoding Problem). *Let \mathbb{F} be a finite field. Let m, k and w be positive integers such that $m > k$ and $m > w$. The syndrome decoding problem with parameters (\mathbb{F}, m, k, w) is the following problem:*

Let H, x and y be such that:

1. H is uniformly sampled from $\mathbb{F}^{(m-k) \times m}$,
2. x is uniformly sampled from $\{x \in \mathbb{F}^m : \text{wt}(x) = w\}$,
3. y is defined as $y := Hx$.

From (H, y) , find x .

In the following, a pair (H, y) generated as in the above definition is called an *instance* of the syndrome decoding problem for parameters (\mathbb{F}, m, k, w) . The syndrome decoding problem is known to be NP-hard. For a weight parameter w lower than the Gilbert-Varshamov radius $\tau_{\text{GV}}(m, k)$, which is defined as:

$$w < \tau_{\text{GV}}(m, k) \Leftrightarrow \sum_{j=0}^{w-1} \binom{m}{j} (q-1)^j < q^{m-k} \quad \text{with } q = |\mathbb{F}| ,$$

we know that there exists a unique solution x such that $y = Hx$ with overwhelming probability. Otherwise, an instance has several solutions on average.

There exists two main families of algorithms to solve the syndrome decoding problem: the *information set decoding* (ISD) algorithms and *generalized birthday algorithms* (GBA) [TS16,BBC⁺19]. To obtain a λ -bit security, the parameters of the syndrome decoding problem are hence chosen in a way to ensure that both kind of algorithms run in time greater than 2^λ .

Instead of working on the standard syndrome decoding problem, we will consider an alternative version that we shall call the *d-split syndrome decoding* problem, where the secret x is split into d chunks of same Hamming weights.

Definition 2 (*d-Split Syndrome Decoding Problem*). *Let \mathbb{F} be a finite field. Let m, k, w be positive integers such that $m > k, m > w, d \mid w$ and $d \mid m$. The *d-split syndrome decoding problem* with parameters (\mathbb{F}, m, k, w) is the following problem:*

Let H, x and y be such that:

1. H is uniformly sampled from $\mathbb{F}^{(m-k) \times m}$,
2. x is uniformly sampled from

$$\left\{ (x_1 \mid \dots \mid x_d) \in \mathbb{F}^m : \forall i \in [d], x_i \in \mathbb{F}^{m/d}, \text{wt}(x_i) = \frac{w}{d} \right\},$$

3. y is defined as $y := Hx$.

From (H, y) , find x .

By taking $d = 1$, we get the standard syndrome decoding problem. The following theorem gives a way to estimate the difficulty to solve the *d-split syndrome decoding* problem.

Theorem 1. *Let \mathbb{F} be a finite field. Let m, k, w be positive integers such that $m > k, m > w, d \mid w$ and $d \mid m$. Let \mathcal{A}_d be an algorithm which solves a random (\mathbb{F}, m, k, w) -instance of the *d-split syndrome decoding* problem in time t with success probability ε_d . Then there exists an algorithm \mathcal{A}_1 which solves a random (\mathbb{F}, m, k, w) -instance of the *standard syndrome decoding* problem in time t with probability ε_1 , where*

$$\varepsilon_1 \geq \frac{\binom{m/d}{w/d}^d}{\binom{m}{w}} \cdot \varepsilon_d.$$

Informally, an instance of the *standard syndrome decoding* problem is an instance of the *d-split syndrome decoding* problem with probability $\binom{m/d}{w/d}^d / \binom{m}{w}$. Moreover, a *standard syndrome decoding* instance can be “randomized” and input to the *d-split adversary* as much as desired. A formal proof of the above theorem is provided in the full version [?].

Let us note that the *d-split syndrome decoding* problem can be seen as a generation of the *regular syndrome decoding* problem introduced by [?], for which the ratio w/d is equal to 1.

2.2 Zero-Knowledge Proofs of Knowledge

We will focus on a special kind of two-party protocol called an *interactive proof* which involves a *prover* \mathcal{P} and a *verifier* \mathcal{V} . In such a protocol, \mathcal{P} tries to prove a statement to \mathcal{V} . The first message sent by \mathcal{P} is called a *commitment*, denoted COM . From this commitment \mathcal{V} produces a first *challenge* CH_1 to which \mathcal{P} answers with a response RSP_1 , followed by a next challenge CH_2 from \mathcal{V} , and so on. After receiving the last response RSP_n , \mathcal{V} produces a binary output: either 1, meaning that she was convinced by \mathcal{P} , or 0 otherwise.

The sequence of exchanged messages is called the *transcript* of the execution, which is denoted

$$\text{View}(\langle \mathcal{P}(in_{\mathcal{P}}), \mathcal{V}(in_{\mathcal{V}}) \rangle) := (\text{COM}, \text{CH}_1, \text{RSP}_1, \dots, \text{CH}_n, \text{RSP}_n)$$

where $in_{\mathcal{P}}$ and $in_{\mathcal{V}}$ respectively denote the prover and verifier inputs. An execution producing an output out is further denoted

$$\langle \mathcal{P}(in_{\mathcal{P}}), \mathcal{V}(in_{\mathcal{V}}) \rangle \rightarrow out .$$

Definition 3 (Proof of Knowledge). *Let x be a statement of language L in NP , and $W(x)$ the set of witnesses for x such that the following relation holds:*

$$\mathcal{R} = \{(x, w) : x \in L, w \in W(x)\} .$$

A proof of knowledge for relation \mathcal{R} with soundness error ε is a two-party protocol between a prover \mathcal{P} and a verifier \mathcal{V} with the following two properties:

- *(Perfect) Completeness: If $(x, w) \in \mathcal{R}$, then a prover \mathcal{P} who knows a witness w for x succeeds in convincing the verifier \mathcal{V} of his knowledge. More formally:*

$$\Pr[\langle \mathcal{P}(x, w), \mathcal{V}(x) \rangle \rightarrow 1] = 1,$$

i.e. given the interaction between the prover \mathcal{P} and the verifier \mathcal{V} , the probability that the verifier is convinced is 1.

- *Soundness: If there exists a PPT prover $\tilde{\mathcal{P}}$ such that*

$$\tilde{\varepsilon} := \Pr[\langle \tilde{\mathcal{P}}(x), \mathcal{V}(x) \rangle \rightarrow 1] > \varepsilon,$$

then there exists an algorithm \mathcal{E} (called an extractor) which, given rewindable black-box access to $\tilde{\mathcal{P}}$, outputs a witness w' for x in time $\text{poly}(\lambda, (\tilde{\varepsilon} - \varepsilon)^{-1})$ with probability at least $1/2$.

Informally, a proof of knowledge has soundness error ε if a prover $\tilde{\mathcal{P}}$ without knowledge of the witness cannot convince the verifier with probability greater than ε assuming that the underlying problem (recovering a witness for the input statement) is hard. Indeed, if a prover $\tilde{\mathcal{P}}$ can succeed with a probability greater than ε , then the existence of the extractor (algorithm \mathcal{E}) implies that $\tilde{\mathcal{P}}$ can be used to compute a witness $w' \in W(x)$.

Remark 1. In the present article, we focus on proof of knowledge for a syndrome decoding instance defined by a matrix H and a vector y . The problem parameters m, k and w will be considered to be defined by the security parameter λ . In this context, the syndrome decoding instance (H, y) is the *statement*. A *witness* for this statement is a small-weight vector x such that $y = Hx$.

We now recall the notion of honest-verifier zero-knowledge proof:

Definition 4 (Honest-Verifier Zero-Knowledge Proof). *A proof of knowledge is $\{computationally, statistically, perfectly\}$ honest-verifier zero-knowledge (HVZK) if there exists a PPT algorithm \mathcal{S} (called simulator) whose output distribution is $\{computationally, statistically, perfectly\}$ indistinguishable from the distribution $\text{View}(\langle \mathcal{P}(x, w), \mathcal{V}(x) \rangle)$ obtained with an honest \mathcal{V} .*

Informally, the previous definition says a genuine execution of the protocol can be simulated without any knowledge of the witness. In other words, the transcript of an execution between the prover and an honest verifier does not reveal any information about the witness.

2.3 Sharings and Multi-Party Computation

In the scope of this article, all the *sharings* are additive. Specifically, an N -sharing of an element $x \in \mathbb{F}^m$ is an N -tuple

$$\llbracket x \rrbracket = (\llbracket x \rrbracket_1, \dots, \llbracket x \rrbracket_N) \in (\mathbb{F}^m)^N \text{ such that } x = \sum_{i=1}^N \llbracket x \rrbracket_i.$$

Each $\llbracket x \rrbracket_i$ is called a *share* of x . For a polynomial $P \in \mathbb{F}[X]$ of degree at most d , we define its sharing $\llbracket P \rrbracket$ as a N -tuple of $(\mathbb{F}[X])^N$ such that $P = \sum_{i=1}^N \llbracket P \rrbracket_i$, where each $\llbracket P \rrbracket_i$ is of degree at most d . In particular, a sharing of a degree- d polynomial can be seen as the sharing of the d -tuple of its coefficients.

In the context of multi-party computation (MPC), an N -sharing is usually distributed to N parties, meaning that each party gets one of the N shares. From those shares, the parties can perform distributed computation. Let assume that each party $i \in [N]$ receives the shares $\llbracket x \rrbracket_i, \llbracket y \rrbracket_i$ and $\llbracket P \rrbracket_i$ corresponding to shared values $x, y \in \mathbb{F}$ and polynomial $P \in \mathbb{F}[X]$. They can perform the following operations:

- **Addition:** the parties locally compute $\llbracket x + y \rrbracket$ by adding their respective shares:

$$\forall i, \llbracket x + y \rrbracket_i := \llbracket x \rrbracket_i + \llbracket y \rrbracket_i.$$

This process is denoted $\llbracket x + y \rrbracket = \llbracket x \rrbracket + \llbracket y \rrbracket$.

- **Addition with a constant:** for a given constant α , the parties locally compute $\llbracket x + \alpha \rrbracket$ by doing:

$$\begin{cases} \llbracket x + \alpha \rrbracket_1 := \llbracket x \rrbracket_1 + \alpha \\ \llbracket x + \alpha \rrbracket_i := \llbracket x \rrbracket_i \text{ for } i \neq 1 \end{cases}$$

This process is denoted $\llbracket x + \alpha \rrbracket = \llbracket x \rrbracket + \alpha$.

- **Multiplication by a constant:** for a given constant α , the parties locally compute $\llbracket \alpha \cdot x \rrbracket$ by multiplying their respective shares:

$$\forall i, \llbracket \alpha \cdot x \rrbracket_i := \alpha \cdot \llbracket x \rrbracket_i .$$

This process is denoted $\llbracket \alpha \cdot x \rrbracket = \alpha \cdot \llbracket x \rrbracket$.

- **Polynomial evaluation:** for a given r , the parties can locally compute $\llbracket P(r) \rrbracket$ by:

$$\forall i, \llbracket P(r) \rrbracket_i := \llbracket P \rrbracket_i(r) = \sum_{j=0}^d \llbracket P_j \rrbracket_i \cdot r^j ,$$

where $\{\llbracket P_j \rrbracket_i\}_j$ denotes the coefficients of $\llbracket P \rrbracket_i$. This process is denoted $\llbracket P(r) \rrbracket = \llbracket P \rrbracket(r)$.

2.4 The MPC-in-the-Head Paradigm

The MPC-in-the-Head (MPCitH) paradigm introduced in [IKOS07] offers a way to build zero-knowledge proofs from secure multi-party computation (MPC) protocols. Let us assume we have an MPC protocol in which N parties $\mathcal{P}_1, \dots, \mathcal{P}_N$ securely and correctly evaluate a function f on a secret input x with the following properties:

- the secret x is encoded as a sharing $\llbracket x \rrbracket$ and each \mathcal{P}_i takes a share $\llbracket x \rrbracket_i$ as input;
- the function f outputs ACCEPT or REJECT;
- the views of t parties leak no information about the secret x .

We can use this MPC protocol to build a zero-knowledge proof of knowledge of an x for which $f(x)$ evaluates to ACCEPT. The prover proceeds as follows:

- she builds a random sharing $\llbracket x \rrbracket$ of x ;
- she simulates locally (“in her head”) all the parties of the MPC protocol;
- she sends commitments to each party’s view, *i.e.* party’s input share, secret random tape and sent and received messages, to the verifier;
- she sends the output shares $\llbracket f(x) \rrbracket$ of the parties, which should correspond to ACCEPT.

Then the verifier randomly chooses t parties and asks the prover to reveal their views. After receiving them, the verifier checks that they are consistent with an honest execution of the MPC protocol and with the commitments. Since only t parties are opened, revealed views leak no information about the secret x , while the random choice of the opened parties makes the cheating probability upper bounded by $(N-t)/N$, thus ensuring the soundness of the zero-knowledge proof.

In this article, we shall only consider the case $t = N - 1$, *i.e.* when the verifier asks to open all the parties except one. We shall further consider that the function f computed by the MPC protocol might be non-deterministic. Specifically, if the protocol takes what we shall call a *good witness* x as input then the protocol

returns ACCEPT with probability 1. Otherwise, the protocol shall reject most of the time but might still accept with some *false positive* probability p . To summarize, we consider a setting in which the output of the protocol has a probability distribution of the form described in Table 1.

| | Output of f | |
|----------------------------------|---------------|---------|
| | ACCEPT | REJECT |
| x is a good witness | 1 | 0 |
| x is <i>not</i> a good witness | p | $1 - p$ |

Table 1. Probability distribution of the output of the MPC protocol

While moving to the MPCitH setting, the randomness for f is then provided by the verifier and the “pre-randomness” view of each party (input share, random tape, initial message) must be committed before receiving the randomness from the verifier. If the prover is honest (*i.e.* knows a “good witness” x), it will always convince the verifier. On the other hand, a malicious prover might successfully cheat with probability $1/N$ (by corrupting the computation of one party) or make the MPC protocol produce a false positive with probability p . Thus, the resulting zero-knowledge protocol has a soundness error of

$$1 - \left(1 - \frac{1}{N}\right) (1 - p) = \frac{1}{N} + p - \frac{1}{N} \cdot p .$$

2.5 Multi-Party Product Verification

A triple of sharings $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket)$ of three elements $a, b, c \in \mathbb{F}$ is called a *multiplication triple* (or Beaver triples [Bea92]) if the shared values satisfy $a \cdot b = c$. The ability to check the correctness of a multiplication triple is instrumental in many MPC (in the Head) protocols.

The authors of [LN17, BN20] propose an MPC protocol to verify the correctness of a multiplication triple by “sacrificing” another one. Specifically, given a random triple $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket)$, the protocol simultaneously verifies the correctness of $(\llbracket x \rrbracket, \llbracket y \rrbracket, \llbracket z \rrbracket)$ and $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket)$, *i.e.* verifies that $c = a \cdot b$ and $z = x \cdot y$, without revealing any information on (x, y, z) in the following way:

1. The parties get a random $\varepsilon \in \mathbb{F}$ (from the verifier in the MPCitH paradigm),
2. The parties locally set $\llbracket \alpha \rrbracket = \varepsilon \llbracket x \rrbracket + \llbracket a \rrbracket$ and $\llbracket \beta \rrbracket = \llbracket y \rrbracket + \llbracket b \rrbracket$.
3. The parties broadcast $\llbracket \alpha \rrbracket$ and $\llbracket \beta \rrbracket$ to obtain α and β .
4. The parties locally set $\llbracket v \rrbracket = \varepsilon \llbracket z \rrbracket - \llbracket c \rrbracket + \alpha \cdot \llbracket b \rrbracket + \beta \cdot \llbracket a \rrbracket - \alpha \cdot \beta$.
5. The parties broadcast $\llbracket v \rrbracket$ to obtain v .
6. The parties output ACCEPT if $v = 0$ and REJECT otherwise.

Observe that if both triples are correct multiplication triples (*i.e.*, $z = xy$ and $c = ab$) then the parties will always accept since

$$\begin{aligned} v &= \varepsilon \cdot z - c + \alpha \cdot b + \beta \cdot a - \alpha \cdot \beta \\ &= \varepsilon \cdot x \cdot y - a \cdot b + (\varepsilon \cdot x + a) \cdot b + (y + b) \cdot a - (\varepsilon \cdot x + a) \cdot (y + b) = 0 \end{aligned}$$

In contrast, if one or both triples are incorrect, then the parties will accept with probability at most $1/|\mathbb{F}|$ as shown in Lemma 1.

Lemma 1 ([BN20]). *If $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket)$ or $(\llbracket x \rrbracket, \llbracket y \rrbracket, \llbracket z \rrbracket)$ is an incorrect multiplication triple then the parties output ACCEPT in the sub-protocol above with probability $\frac{1}{|\mathbb{F}|}$.*

The authors of [KZ21] propose a variant of the above protocol to batch the verification of the d multiplication triples $(\llbracket x_j \rrbracket, \llbracket y_j \rrbracket, \llbracket z_j \rrbracket)$ by sacrificing a random *dot-product tuple* $(\llbracket a_j \rrbracket, \llbracket b_j \rrbracket)_{j \in [d]}, \llbracket c \rrbracket$ verifying $c = \langle a, b \rangle$.

1. The parties gets a random $\varepsilon \in \mathbb{F}^d$ (from the verifier in the MPCitH paradigm),
2. The parties locally set $\llbracket \alpha \rrbracket = \varepsilon \circ \llbracket x \rrbracket + \llbracket a \rrbracket$ and $\llbracket \beta \rrbracket = \llbracket y \rrbracket + \llbracket b \rrbracket$.
3. The parties broadcast $\llbracket \alpha \rrbracket$ and $\llbracket \beta \rrbracket$ to obtain α and β .
4. The parties locally set $\llbracket v \rrbracket = -\llbracket c \rrbracket + \langle \varepsilon, \llbracket z \rrbracket \rangle + \langle \alpha, \llbracket b \rrbracket \rangle + \langle \beta, \llbracket a \rrbracket \rangle - \langle \alpha, \beta \rangle$.
5. The parties broadcast $\llbracket v \rrbracket$ to obtain v .
6. The parties output ACCEPT if $v = 0$ and REJECT otherwise.

Lemma 2 ([KZ21]). *If $(\llbracket x_j \rrbracket, \llbracket y_j \rrbracket, \llbracket z_j \rrbracket)_{j \in [d]}$ contains an incorrect multiplication triple or if $(\llbracket a_j \rrbracket, \llbracket b_j \rrbracket)_{j \in [d]}, \llbracket c \rrbracket$ form an incorrect dot product, then the parties output ACCEPT in the sub-protocol above with probability at most $\frac{1}{|\mathbb{F}|}$.*

This variant requires less communication for c and v , compared to the case where we repeat d times the original protocol. But depending on the context, repeating d times the original protocol might be preferred to lower the false positive probability (*i.e.* $1/|\mathbb{F}|^d$ against $1/|\mathbb{F}|$).

3 A Zero-Knowledge Protocol for Syndrome Decoding

Let us consider an instance (H, y) of the (d -split) syndrome decoding problem, and let us denote x a solution of this instance. We denote \mathbb{F}_{SD} the field on which the instance is defined.

Without loss of generality, we assume that H is in the standard form, *i.e.* that $H = (H' | I_{m-k})$ for some $H' \in \mathbb{F}_{\text{SD}}^{(m-k) \times k}$. Thus the solution x can be written as $(x_A | x_B)$ such that we have the linear relation

$$y = H'x_A + x_B . \tag{1}$$

This implies that one simply needs to send x_A ($k \cdot \log |\mathbb{F}_{\text{SD}}|$ bits) to reveal the solution of the instance (H, y) .

In the following sections, we first build an MPC protocol that takes a sharing of $\llbracket x_A \rrbracket$, builds the corresponding $\llbracket x \rrbracket$ thanks to Equation (1), and checks that $\llbracket x \rrbracket$ corresponds to a vector with a Hamming weight of at most w/d on each chunk. Since $\llbracket x \rrbracket$ would verify $y = Hx$ by construction, this MPC protocol verifies that $\llbracket x_A \rrbracket$ corresponds to a solution of the syndrome decoding instance (H, y) . Then, in Section 3.3, we transform it into a zero-knowledge protocol which proves the knowledge of a solution of the syndrome decoding instance (H, y) thanks to the MPC-in-the-Head paradigm (described in Section 2.4).

3.1 Standard Case ($d = 1$)

We first focus on the case where (H, y) is an instance of the standard syndrome decoding problem (*i.e.* we have $d = 1$). We will then show how to extend the protocol to the general case of any d . We consider a field extension $\mathbb{F}_{\text{poly}} \supseteq \mathbb{F}_{\text{SD}}$ such that $|\mathbb{F}_{\text{poly}}| \geq m$ (we recall that m is the length of the secret x , *i.e.* $x \in \mathbb{F}_{\text{SD}}^m$). We denote $\phi : \mathbb{F}_{\text{SD}} \rightarrow \mathbb{F}_{\text{poly}}$ the canonical inclusion of \mathbb{F}_{SD} into \mathbb{F}_{poly} . Let us take a bijection γ between $\{1, \dots, |\mathbb{F}_{\text{poly}}|\}$ and \mathbb{F}_{poly} . Then, to ease the notation, we denote γ_i for $\gamma(i)$.

The protocol must check that $y = Hx$ and $\text{wt}(x) \leq w$. As explained in the introduction of the section, the input for the MPC protocol will be $\llbracket x_A \rrbracket$, then it will build the sharing $\llbracket x \rrbracket$ using the linear relation (1). Then we directly have that $y = Hx$. It remains to check that $\text{wt}(x) \leq w$.

To prove that $\text{wt}(x) \leq w$, the prover build the three following polynomials:

- The polynomial $S \in \mathbb{F}_{\text{poly}}[X]$ satisfying

$$\forall i \in [m], S(\gamma_i) = \phi(x_i) ,$$

as well as $\deg S \leq m - 1$. This S is unique and can be computed by interpolation.

- The polynomial $Q \in \mathbb{F}_{\text{poly}}[X]$ defined as

$$Q(X) := \prod_{i \in E} (X - \gamma_i)$$

for some $E \subset [m]$ such that $|E| = w$ and $\{i \in [m] : x_i \neq 0\} \subset E$, implying $\deg Q = w$.

- The polynomial $P \in \mathbb{F}_{\text{poly}}[X]$ defined as

$$P := (Q \cdot S)/F \text{ with } F(X) := \prod_{i=1}^m (X - \gamma_i) .$$

We stress some useful properties of these polynomials:

- The polynomial Q is a monic polynomial of degree w . Moreover, for every $i \in [m]$, we have

$$x_i \neq 0 \Rightarrow i \in E \Rightarrow Q(\gamma_i) = 0 .$$

- The polynomial F divides $Q \cdot S$. Indeed, for every $i \in [m]$, we have

$$(Q \cdot S)(\gamma_i) = 0$$

since $S(\gamma_i) \neq 0 \Rightarrow x_i \neq 0 \Rightarrow Q(\gamma_i) = 0$. The polynomial P is hence well defined.

- The polynomial P has degree $\deg P \leq w - 1$.

If the prover convinces the verifier that there exists two polynomials P (with $\deg P \leq w - 1$) and Q (with $\deg Q = w$) such that $Q \cdot S - P \cdot F = 0$ where S and F are built as described above, then the verifier can deduce the following:

$$\begin{aligned} \forall i \in [m], (Q \cdot S)(\gamma_i) &= P(\gamma_i) \cdot F(\gamma_i) = 0 \\ &\Rightarrow \forall i \in [m], Q(\gamma_i) = 0 \text{ or } S(\gamma_i) = \phi(x_i) = 0 \end{aligned}$$

Since Q has at most w roots, the verifier concludes that $\phi(x_i) \neq 0$ in at most w positions. Thus $\text{wt}(x) \leq w$.

We now explain how to prove this statement in the MPCitH paradigm. For this purpose, we describe an MPC protocol, which on input x , P and Q outputs ACCEPT if the above condition is verified and REJECT otherwise, except with a small false positive probability. The parties' inputs are defined as the shares of $\llbracket x_A \rrbracket$, $\llbracket Q \rrbracket$ and $\llbracket P \rrbracket$. Let us recall that a sharing of a polynomial is naturally defined as a sharing of its coefficients (see Section 2.3). However, for the sharing of Q , we share all of its coefficients except the leading one. Indeed since Q is monic, its leading coefficient is publicly known and is equal to 1. Moreover, it enables to convince the verifier that Q is of degree *exactly* w , which is important since otherwise, a malicious prover could take Q as the zero polynomial.

From its inputs, the MPC protocol first builds the polynomial S from x_A . Then, to verify $Q \cdot S = P \cdot F$, it evaluates the two sides of the relation on t random points r_1, \dots, r_t (sampled by the verifier in the MPCitH setting). If the relation is not verified, the probability to observe $Q(r_j) \cdot S(r_j) = P(r_j) \cdot F(r_j)$ for all $j \in [t]$ will be low, which stems from the Schwartz-Zippel Lemma (see the full version [?]). The larger the set from which the evaluation points r_j are sampled, the smaller the false positive probability p . For this reason, we take these evaluation points in a field extension $\mathbb{F}_{\text{points}}$ of \mathbb{F}_{poly} . Such a field extension allows us to have more points and so to detect more efficiently when $Q \cdot S \neq P \cdot F$. In practice, given an evaluation point r_j , the parties of the MPC protocol verify the relations $Q(r_j) \cdot S(r_j) = (P \cdot F)(r_j)$ by sacrificing multiplication triples as described in Section 2.5. To proceed, the prover must previously build t multiplication triples $(\llbracket a_j \rrbracket, \llbracket b_j \rrbracket, \llbracket c_j \rrbracket)$ for random elements $a_j, b_j, c_j \in \mathbb{F}_{\text{points}}$ satisfying $a_j \cdot b_j = c_j$ for $j \in [t]$ and include them to the parties' inputs (each party getting its corresponding share from $\llbracket a_j \rrbracket$, $\llbracket b_j \rrbracket$ and $\llbracket c_j \rrbracket$).

The MPC protocol runs as follows:

1. The parties sample t random points r_1, \dots, r_t of $\mathbb{F}_{\text{points}}$.

2. The parties locally compute $\llbracket x \rrbracket$ from $\llbracket x_A \rrbracket$ using Equation (1).
3. The parties locally compute $\llbracket S(r_j) \rrbracket$, $\llbracket Q(r_j) \rrbracket$ and $\llbracket (F \cdot P)(r_j) \rrbracket$ for all $j \in [t]$. Let us remark that $\llbracket S(r_j) \rrbracket$ can be computed from $\llbracket x \rrbracket$ by the parties without any interaction thanks to the linearity of Lagrange interpolation formula:

$$\llbracket S(r_j) \rrbracket = \sum_{i \in [m]} \llbracket x_i \rrbracket \prod_{\ell \in [m], \ell \neq i} \frac{r_j - \gamma_\ell}{\gamma_i - \gamma_\ell}.$$

On the other hand $\llbracket (F \cdot P)(r_j) \rrbracket$ is computed as $F(r_j) \cdot \llbracket P(r_j) \rrbracket$ since F is publicly known.

4. For every $j \in [t]$, the parties run an MPC verification of the multiplication triple $(\llbracket S(r_j) \rrbracket, \llbracket Q(r_j) \rrbracket, \llbracket (F \cdot P)(r_j) \rrbracket)$ by sacrificing the triple $(\llbracket a_j \rrbracket, \llbracket b_j \rrbracket, \llbracket c_j \rrbracket)$:
 - The parties sample a random $\varepsilon_j \in \mathbb{F}_{\text{points}}$.
 - The parties locally set

$$\llbracket \alpha_j \rrbracket = \varepsilon_j \cdot \llbracket Q(r_j) \rrbracket + \llbracket a_j \rrbracket \quad \text{and} \quad \llbracket \beta_j \rrbracket = \llbracket S(r_j) \rrbracket + \llbracket b_j \rrbracket.$$

- The parties broadcast $\llbracket \alpha_j \rrbracket$ and $\llbracket \beta_j \rrbracket$ to obtain α_j and β_j .
- The parties locally set

$$\llbracket v_j \rrbracket = \varepsilon_j \cdot \llbracket (F \cdot P)(r_j) \rrbracket - \llbracket c_j \rrbracket + \alpha_j \cdot \llbracket b_j \rrbracket + \beta_j \cdot \llbracket a_j \rrbracket - \alpha_j \cdot \beta_j.$$

- The parties broadcast $\llbracket v_j \rrbracket$ to obtain v_j .

5. The parties output ACCEPT if $v = 0$ and REJECT otherwise.

Note that we do not need to specify how the random values r_j 's and ε_j 's are sampled by the parties since they will be provided as challenges from the verifier while turning to the zero-knowledge setting.

The above MPC protocol computes a non-deterministic function f which takes x , Q and P (and t multiplication triples) as input and which outputs ACCEPT or REJECT. The randomness of this function comes from the random evaluations points r_1, \dots, r_t and from the random challenges $\varepsilon_1, \dots, \varepsilon_t$ used by the product checking protocol. Whenever x indeed satisfies $\text{wt}(x) \leq w$ and the polynomials P and Q are genuinely computed as described above, the protocol outputs ACCEPT with probability one. Whenever the protocol input is not of this form, the protocol shall output REJECT except with a small false positive probability p . In other words, the output of the above protocol follows the distribution depicted in Table 1 where a good witness here means an x of weight at most w and polynomials P and Q which are correctly built.

Let us make explicit the false positive probability p . We shall denote $\Delta := |\mathbb{F}_{\text{points}}|$. Whenever the protocol input is not a good witness, *i.e.* $\text{wt}(x) > w$, P or Q are not correctly built, we have $Q \cdot S \neq F \cdot P$. In the above protocol, both sides of the relation are evaluated in t random points. The probability to have the equality for i evaluation points among the t points is at most

$$\frac{\max_{\ell \leq m+w-1} \left\{ \binom{\ell}{i} \binom{\Delta-\ell}{t-i} \right\}}{\binom{\Delta}{t}}$$

since $Q \cdot S - F \cdot P$ is a polynomial of degree at most $m + w - 1$. This holds from a simple extension of the Schwartz-Zippel Lemma that we provide in the full version [?]. When this event occurs, the probability to obtain ACCEPT as output is

$$\left(\frac{1}{\Delta}\right)^{t-i},$$

which corresponds to the probability to get the $t - i$ false positives in the verification of multiplication triples (for the $t - i$ remaining evaluation points r_j for which $Q(r_j) \cdot S(r_j) \neq F(r_j) \cdot P(r_j)$). Thus, the global false positive probability p satisfies

$$p \leq \sum_{i=0}^t \frac{\max_{\ell \leq m+w-1} \left\{ \binom{\ell}{i} \binom{\Delta-\ell}{t-i} \right\}}{\binom{\Delta}{t}} \left(\frac{1}{\Delta}\right)^{t-i}. \quad (2)$$

3.2 General case (any d)

Let us now assume that (H, y) is an instance of a d -split syndrome decoding problem for some $d \geq 1$. We can easily adapt our protocol in that case. Instead of having a unique polynomial Q of degree w , we will have d polynomials Q_1, \dots, Q_d of degree exactly w/d to prove the weight bound $\text{wt}(x_j) \leq w/d$ for each chunk x_j of the SD solution. We then have d polynomials S_j (of degree $m/d - 1$) and d polynomials P_j (of degree $w/d - 1$) satisfying the d relations $Q_j \cdot S_j = F \cdot P_j$ with $F := \prod_{j=1}^{m/d} (X - \gamma_j)$. To prove those d relations we evaluate each of them on t random points r_1, \dots, r_t . We stress that the same t random points can be used for each chunk, *i.e.* for every $j \in [d]$.

A malicious prover might try to cheat on a single relation (*i.e.* on a single chunk of the SD solution), in such a way that there exists $j_0 \in [d]$ with

$$\begin{cases} Q_{j_0} \cdot S_{j_0} \neq F \cdot P_{j_0}, \\ \forall j \neq j_0, Q_j \cdot S_j = F \cdot P_j. \end{cases}$$

So for a given point r , we use the dot-product checking of [KZ21] (described in Section 2.5) to check all the equalities $Q_j(r) \cdot S_j(r) = F(r) \cdot P_j(r)$ at once. This saves communication without impacting the soundness error compared to independent checks of the d relations.

Whenever the input x , $\{P_j\}$, $\{Q_j\}$ is not a good witness (*i.e.* whenever one x_j has a weight greater than w/d or one polynomial P_j or Q_j is not correctly built), at least one of the relations $Q_j \cdot S_j = F \cdot P_j$ is not verified. Since $Q_j \cdot S_j - F \cdot P_j$ is a polynomial of degree at most $(m+w)/d - 1$, the global false positive probability for the d -split variant becomes

$$p \leq \sum_{i=0}^t \frac{\max_{\ell \leq (m+w)/d-1} \left\{ \binom{\ell}{i} \binom{\Delta-\ell}{t-i} \right\}}{\binom{\Delta}{t}} \left(\frac{1}{\Delta}\right)^{t-i} \quad (3)$$

with $\Delta := |\mathbb{F}_{\text{points}}|$. (This upper bound is equivalent to (2) where the max degree $m + w - 1$ is replaced by $(m + w)/d - 1$).

The constraint on the size of \mathbb{F}_{poly} now becomes

$$|\mathbb{F}_{\text{poly}}| \geq \frac{m}{d}$$

since we only need w/d points for the interpolation of the polynomials S_1, \dots, S_d . Thus using the d -split version allows us to use smaller fields for \mathbb{F}_{poly} and $\mathbb{F}_{\text{points}}$.

Let us note that in practice the new communication is not smaller than before, but rather equivalent or higher, since we need to use bigger syndrome decoding instances to compensate the security loss of the d -split version. The main benefit to introduce the d -split version is to work on polynomials of smaller degree and/or on specific fields which provides better performance trade-offs (see Section 4.5).

3.3 Description of the Protocol

We now give the formal description of our zero-knowledge protocol (general case) in Protocol 1. For the sake of clarity in the protocol description, we denote \mathbf{Q} the tuple of polynomials (Q_1, \dots, Q_d) . Same for the polynomials \mathbf{P} and \mathbf{S} . The additions, subtractions and polynomial evaluations of these tuples are component-wise defined. For example, for a point $r \in \mathbb{F}_{\text{points}}$, $\mathbf{Q}(r)$ means $(Q_1(r), \dots, Q_d(r))$. We also use this bold notation for $\mathbf{a}_j, \mathbf{b}_j, \mathbf{\alpha}_j, \mathbf{\beta}_j$ and $\mathbf{\varepsilon}_j$ which shall represent vectors of $\mathbb{F}_{\text{points}}^d$. Let us recall that \circ denotes the component-wise multiplication. In the scope of this protocol, the polynomial F is defined as $F(X) := \prod_{i=1}^{m/d} (X - \gamma_i)$ with $\mathbb{F}_{\text{poly}} = \{\gamma_1, \gamma_2, \dots\}$.

3.4 Security Proofs

The following theorems state the completeness, zero-knowledge and soundness of Protocol 1. The proofs of Theorems 3 and 4 are provided in the full version [?].

Theorem 2 (Completeness). *Protocol 1 is perfectly complete, i.e. a prover \mathcal{P} who knows a solution x to the syndrome decoding instance (H, y) and who follows the steps of the protocol always succeeds in convincing the verifier \mathcal{V} .*

Proof. For any sampling of the random coins of \mathcal{P} and \mathcal{V} , if the computation described in Protocol 1 is genuinely performed then all the checks of \mathcal{V} pass. \square

Theorem 3 (Honest-Verifier Zero-Knowledge). *Let the PRG used in Protocol 1 be $(t, \varepsilon_{\text{PRG}})$ -secure and the commitment scheme Com be $(t, \varepsilon_{\text{Com}})$ -hiding. There exists an efficient simulator \mathcal{S} which, given random challenge i^* outputs a transcript which is $(t, \varepsilon_{\text{PRG}} + \varepsilon_{\text{Com}})$ -indistinguishable from a real transcript of Protocol 1.*

Theorem 4 (Soundness). *Suppose that there is an efficient prover $\tilde{\mathcal{P}}$ that, on input (H, y) , convinces the honest verifier \mathcal{V} on input H, y to accept with probability*

$$\tilde{\varepsilon} := \Pr[\langle \tilde{\mathcal{P}}, \mathcal{V} \rangle(H, y) \rightarrow 1] > \varepsilon$$

Inputs: Both parties have $H = (H'|I_{m-k}) \in \mathbb{F}_{\text{SD}}^{(m-k) \times m}$ and $y \in \mathbb{F}_{\text{SD}}^{m-k}$, the prover also holds $x := (x_1 | x_2 | \dots | x_d) \in \mathbb{F}_{\text{SD}}^m$ such that $y = Hx$ and $\text{wt}(x_j) \leq w$ for $j \in [d]$.

Round 1: The prover computes the proof witness: for all chunk $j \in [d]$,

1. Choose a set $E_j \subset [\frac{m}{d}]$ s.t. $|E_j| = \frac{w}{d}$ and $\{\ell : (x_j)_\ell \neq 0\} \subset E_j$.
2. Compute $Q_j(X) = \prod_{\ell \in E_j} (X - \gamma_\ell) \in \mathbb{F}_{\text{poly}}[X]$.
3. Compute $S_j(X) \in \mathbb{F}_{\text{poly}}[X]$ by interpolation s.t. $\deg S_j \leq \frac{m}{d} - 1$ and $\forall \ell \in [\frac{m}{d}], S_j(\gamma_\ell) = (x_j)_\ell$.
4. Compute $P_j(X) = S_j(X)Q_j(X)/F(X) \in \mathbb{F}_{\text{poly}}[X]$.

Then, the prover prepares the inputs for the multi-party computation as follows:

1. Sample a root seed: $\text{seed} \xleftarrow{\$} \{0, 1\}^\lambda$.
2. Compute parties' seeds and commitment randomness $(\text{seed}_i, \rho_i)_{i \in [N]}$ with $\text{TreePRG}(\text{seed})$.
3. For each party $i \in \{1, \dots, N\}$,
 - $[\mathbf{a}_j]_i, [\mathbf{b}_j]_i \leftarrow \text{PRG}(\text{seed}_i)$, for each $j \in [t]$
 - If $i \neq N$,
 - $\{[c_j]_i\}_{j \in [t]}, [x_A]_i, [\mathbf{Q}]_i, [\mathbf{P}]_i \leftarrow \text{PRG}(\text{seed}_i)$
 - $\text{state}_i = \text{seed}_i$
 - Else,
 - $[x_A]_N = x_A - \sum_{\ell \neq N} [x_A]_\ell$
 - $[\mathbf{Q}]_N = \mathbf{Q} - \sum_{\ell \neq N} [\mathbf{Q}]_\ell$.
 - $[\mathbf{P}]_N = \mathbf{P} - \sum_{\ell \neq N} [\mathbf{P}]_\ell$.
 - $[c_j]_N = \langle \mathbf{a}_j, \mathbf{b}_j \rangle - \sum_{\ell \neq N} [c_j]_\ell$, for each $j \in [t]$
 - $\text{aux} = ([x_A]_N, [\mathbf{Q}]_N, [\mathbf{P}]_N, \{[c_j]_N\}_{j \in [t]})$
 - $\text{state}_N = \text{seed}_N \parallel \text{aux}$
 - Commit the party's state: $\text{com}_i = \text{Com}(\text{state}_i; \rho_i)$.

The prover builds $h = \text{Hash}(\text{com}_1, \dots, \text{com}_N)$ and sends it to the verifier.

Round 2: The verifier uniformly samples, for each $j \in [t]$, an evaluation point $r_j \leftarrow \mathbb{F}_{\text{points}}$ and a vector $\boldsymbol{\varepsilon}_j \leftarrow \mathbb{F}_{\text{points}}^d$, and sends them to the prover.

Round 3: The prover simulates the MPC protocol:

1. The parties locally set $[x_B] = y - H'[x_A]$.
2. The parties locally compute $[\mathbf{S}]$ by interpolation using $[x] := ([x_A] | [x_B])$.
3. Then for all $j \in [t]$,
 - The parties locally compute $[\mathbf{S}(r_j)], [\mathbf{Q}(r_j)]$ and $[\mathbf{P}(r_j)]$.
 - They locally set $[\boldsymbol{\alpha}_j] = \boldsymbol{\varepsilon}_j \circ [\mathbf{Q}(r_j)] + [\mathbf{a}_j]$.
 - They locally set $[\boldsymbol{\beta}_j] = [\mathbf{S}(r_j)] + [\mathbf{b}_j]$.
 - The parties open $[\boldsymbol{\alpha}_j]$ and $[\boldsymbol{\beta}_j]$ to get $\boldsymbol{\alpha}_j$ and $\boldsymbol{\beta}_j$.
 - The parties locally set

$$[v_j] = -[c_j] + \langle \boldsymbol{\varepsilon}_j, F(r_j) \cdot [\mathbf{P}(r_j)] \rangle + \langle \boldsymbol{\alpha}_j, [\mathbf{b}_j] \rangle + \langle \boldsymbol{\beta}_j, [\mathbf{a}_j] \rangle - \langle \boldsymbol{\alpha}_j, \boldsymbol{\beta}_j \rangle.$$

The prover builds $h' = \text{Hash}([\boldsymbol{\alpha}_1], [\boldsymbol{\beta}_1], [v_1], \dots, [\boldsymbol{\alpha}_t], [\boldsymbol{\beta}_t], [v_t])$ and sends it to the verifier.

Round 4: The verifier uniformly samples $i^* \leftarrow [N]$ and sends it to the prover.

Round 5: The prover sends $(\text{state}_i, \rho_i)_{i \neq i^*}, \text{com}_{i^*}, \{[\boldsymbol{\alpha}_j]_{i^*}\}_{j \in [t]}$ and $\{[\boldsymbol{\beta}_j]_{i^*}\}_{j \in [t]}$.

Verification: The verifier accepts iff all the following checks succeed:

1. For each $i \neq i^*$, she computes all the commitments to the parties' states: $\text{com}_i = \text{Com}(\text{state}_i; \rho_i)$. Then she checks that $h \stackrel{?}{=} \text{Hash}(\text{com}_1, \dots, \text{com}_N)$.
2. Using $\{\text{state}_i\}_{i \neq i^*}$, she simulates all the parties except for i^* . From the recomputed shares, she checks that $h' \stackrel{?}{=} \text{Hash}([\boldsymbol{\alpha}_1], [\boldsymbol{\beta}_1], [v_1], \dots, [\boldsymbol{\alpha}_t], [\boldsymbol{\beta}_t], [v_t])$ where $[v_j]_{i^*} := -\sum_{i \neq i^*} [v_j]_i$.

Protocol 1: Zero-knowledge proof for syndrome decoding.

where the soundness error ε is equal to

$$p + \frac{1}{N} - p \cdot \frac{1}{N}$$

with p defined in Equation (3). Then, there exists an efficient probabilistic extraction algorithm \mathcal{E} that, given rewindable black-box access to $\tilde{\mathcal{P}}$, produces with either a witness x such that $y = Hx$ and $\text{wt}(x) \leq w$, or a commitment collision, by making an average number of calls to $\tilde{\mathcal{P}}$ which is upper bounded by

$$\frac{4}{\tilde{\varepsilon} - \varepsilon} \cdot \left(1 + \tilde{\varepsilon} \cdot \frac{2 \cdot \ln(2)}{\tilde{\varepsilon} - \varepsilon} \right).$$

By adapting the parameters t and Δ , we can produce a protocol with soundness error arbitrarily close to $1/N$.

3.5 Performance

In the following analysis, we exclude the challenges from the communication cost since they are of very moderate impact (and do not count whenever making the protocol non-interactive). The communication then consists into

- COM := h ,
- RES₁ := h' and
- RES₂ := $((\text{state}_i, \rho_i)_{i \neq i^*}, \text{com}_{i^*}, \{\llbracket \alpha_j \rrbracket_{i^*}\}_{j \in [t]}, \{\llbracket \beta_j \rrbracket_{i^*}\}_{j \in [t]})$.

For $i \neq N$, state_i simply consists in a seed of λ bits. For $i = N$, state_i contains

- a seed of λ bits,
- the share $\llbracket x_A \rrbracket_N$ of a plaintext,
- the shares $\llbracket Q \rrbracket_N$ and $\llbracket P \rrbracket_N$ which are $2 \cdot d$ polynomials of degree $w/d - 1$,
- and the shares $\{\llbracket c_j \rrbracket_N\}_{j \in [t]}$ of t points of $\mathbb{F}_{\text{points}}$.

Let us recall that seeds are sampled using a tree PRG. Instead of sending the $N - 1$ seeds and commitment randomness of $(\text{state}_i, \rho_i)_{i \neq i^*}$, we can instead send the sibling path from $(\text{state}_{i^*}, \rho_{i^*})$ to the tree root, it costs at most $\lambda \cdot \log_2(N)$ bits (we need to reveal $\log_2(N)$ nodes of the tree). Moreover com_{i^*} is a commitment of 2λ bits, and $\{\llbracket \alpha_j \rrbracket_{i^*}\}_{j \in [t]}, \{\llbracket \beta_j \rrbracket_{i^*}\}_{j \in [t]}$ are elements of $\mathbb{F}_{\text{points}}$. The communication cost (in bits) of the protocol is then

$$\begin{aligned} \text{SIZE} = & 4\lambda + \underbrace{k \cdot \log_2 |\mathbb{F}_{\text{SD}}|}_{\llbracket x_A \rrbracket_N} + \underbrace{(2 \cdot w) \cdot \log_2 |\mathbb{F}_{\text{poly}}|}_{\llbracket Q \rrbracket_N, \llbracket P \rrbracket_N} \\ & + \underbrace{(2 \cdot d + 1) \cdot t \cdot \log_2 |\mathbb{F}_{\text{points}}|}_{\{\llbracket \alpha_j \rrbracket_{i^*}, \llbracket \beta_j \rrbracket_{i^*}, \llbracket c_j \rrbracket_N\}_{j \in [t]}} + \underbrace{\lambda \cdot \log_2(N)}_{(\text{seed}_i)_{i \neq i^*}} + \underbrace{2\lambda}_{\text{com}_{i^*}} \end{aligned}$$

As usual, to achieve a targeted soundness error $2^{-\lambda}$, we can perform τ parallel repetitions of the protocol such that $\varepsilon^\tau \leq 2^{-\lambda}$. And instead of sending τ values

for h and h' , we can merge them together to send a single h and a single h' . The communication cost (in bits) of the protocol with τ repetitions is

$$\begin{aligned} \text{SIZE} = & 4\lambda + \tau \cdot \left(k \cdot \log_2 |\mathbb{F}_{\text{SD}}| + (2 \cdot w) \cdot \log_2 |\mathbb{F}_{\text{poly}}| \right. \\ & \left. + (2 \cdot d + 1) \cdot t \cdot \log_2 |\mathbb{F}_{\text{points}}| + \lambda \cdot \log_2(N) + 2\lambda \right) \end{aligned}$$

and the obtained soundness error is

$$\left(p + \frac{1}{n} - p \cdot \frac{1}{n} \right)^\tau .$$

3.6 Comparison

We compare our new protocol with existing zero-knowledge protocols for syndrome decoding (or equivalently for *message decoding*). We compare these protocols on two SD instances of 128-bit security:

- Instance 1 [?]: Syndrome Decoding on \mathbb{F}_2 with parameters

$$(m, k, w) = (1280, 640, 132);$$

- Instance 2 [CVE11]: Syndrome Decoding on \mathbb{F}_{2^8} with parameters

$$(m, k, w) = (208, 104, 78).$$

The comparison for a soundness error of 2^{-128} is given in the Table 2. For our protocol, we provide two instantiations for each syndrome decoding instance to give the reader an idea of the obtained performance while changing the number of parties. The first instantiation called “short” corresponds to an instantiation which provides small communication cost. The second one called “fast” corresponds to an instantiation with faster computation but higher communication cost. The used parameters $(N, \tau, |\mathbb{F}_{\text{poly}}|, |\mathbb{F}_{\text{points}}|, t)$ for our scheme are

- Instance 1:

Short: $(256, 16, 2^{11}, 2^{22}, 2) \Rightarrow \varepsilon^\tau = 2^{-128.0}$

Fast: $(32, 26, 2^{11}, 2^{22}, 1) \Rightarrow \varepsilon^\tau = 2^{-129.6}$

- Instance 2:

Short: $(256, 16, 2^8, 2^{24}, 2) \Rightarrow \varepsilon^\tau = 2^{-128.0}$

Fast: $(32, 26, 2^8, 2^{24}, 1) \Rightarrow \varepsilon^\tau = 2^{-130.0}$

We can remark that all the previous protocols prove an equality for the Hamming weight by relying on isometries (*i.e.* permutations if $\mathbb{F}_{\text{SD}} = \mathbb{F}_2$). On our side, we only prove the inequality $\text{wt}(w) \leq w$. We stress that both versions (equality or inequality) can be merely equivalent for some SD parameters. Indeed, if w is chosen sufficiently below the Gilbert-Varshamov bound and if we know there exists an SD solution x of Hamming weight w , then proving the knowledge of a solution x' with $\text{wt}(x') \leq w$ amounts to proving the knowledge of x with overwhelming probability.

| Name Protocol | Year | Instance 1 | Instance 2 | Proved statement |
|--------------------|------|------------|------------|-------------------------------|
| [Ste94] | 1993 | 37.4 KB | 46.1 KB | $y = Hx, \text{wt}(x) = w$ |
| [Vér96] | 1997 | 31.7 KB | 38.7 KB | <i>message decoding</i> |
| [CVE11] | 2010 | - | 37.4 KB | $y = Hx, \text{wt}(x) = w$ |
| [AGS11] | 2011 | 24.8 KB | - | $y = Hx, \text{wt}(x) = w$ |
| [?] (short) | 2021 | - | 15.2 KB | $y = Hx, \text{wt}(x) = w$ |
| [?] (fast) | 2021 | - | 19.9 KB | $y = Hx, \text{wt}(x) = w$ |
| [?] (short) | 2021 | 12.9 KB | 15.6 KB | $y = Hx, \text{wt}(x) = w$ |
| [?] (fast) | 2021 | 20.0 KB | 24.7 KB | $y = Hx, \text{wt}(x) = w$ |
| Our scheme (short) | 2022 | 9.7 KB | 6.9 KB | $y = Hx, \text{wt}(x) \leq w$ |
| Our scheme (fast) | 2022 | 14.4 KB | 9.7 KB | $y = Hx, \text{wt}(x) \leq w$ |

Table 2. Comparison of our protocol with state-of-the-art zero-knowledge protocols for syndrome decoding. The formulae for the communication costs of the different protocols and the used parameters are detailed in the full version [?].

4 The Signature Scheme

A signature scheme is a triplet of PPT algorithms (**KeyGen**, **Sign**, **Verif**). On input 1^λ for security level λ , **KeyGen** outputs a pair (pk, sk) where $pk \in \{0, 1\}^{\text{poly}(\lambda)}$ is a public key and $sk \in \{0, 1\}^{\text{poly}(\lambda)}$ is a private key (a.k.a. secret key). On input a secret key sk and a message $m \in \{0, 1\}^*$, **Sign** produces a signature $s \in \{0, 1\}^{\text{poly}(\lambda)}$. **Verif** is a deterministic algorithm which, on input a public key pk , a signature s and a message m , outputs 1 if s is a valid signature for m under pk (meaning that it is a possible output $s \leftarrow \text{Sign}(sk, m)$ for the corresponding sk) and it outputs 0 otherwise. The standard security property for a signature scheme is the *existential unforgeability* against chosen message attacks: an adversary \mathcal{A} given pk and a oracle access to $\text{Sign}(sk, \cdot)$ should not be able to produce a pair (s, m) satisfying $\text{Verif}(pk, s, m) = 1$ (for a message m which was not queried to the signing oracle).

In this section, we show how to turn our 5-round HVZK protocol into a signature scheme using the Fiat-Shamir transform [FS87, AABN02]. After explaining the transformation, we give the description of the signature scheme and then provide a security proof in the random oracle model (ROM).

4.1 Transformation into a Non-Interactive Scheme

To transform our protocol into a non-interactive scheme, we apply the multi-round variant of the Fiat-Shamir transform [FS87] (see *e.g.* [EDV⁺12, CHR⁺16]). Concretely, we compute the challenge CH_1 and CH_2 as

$$\begin{aligned} h_1 &= \text{Hash}_1(m, \text{salt}, h) \\ \text{CH}_1 &\leftarrow \text{PRG}(h_1) \end{aligned}$$

and

$$\begin{aligned} h_2 &= \text{Hash}_2(m, \text{salt}, h, h') \\ \text{CH}_2 &\leftarrow \text{PRG}(h_2) \end{aligned}$$

where m is the input message, where Hash_1 and Hash_2 are some hash functions (that shall be modeled as random oracles) and where h and h' are the Round 1 and Round 3 hash commitments merged for the τ repetitions. We introduce a value salt called *salt* which is sampled from $\{0, 1\}^{2\lambda}$ at the beginning of the signing process. This value is then used for each commitment to the parties' states. Without it, the security of the signature would be at most $2^{\lambda/2}$ because of the seed collisions between several signatures. Moreover, since the signature security relies on the random oracle model, we can safely replace the commitment scheme Com of the Protocol 1 by a single hash function Hash_0 .

The security of the obtained scheme is lower than the soundness error of Protocol 1. Indeed, in [KZ20a], Kales and Zaverucha describe a forgery attack against signature schemes obtained by applying the Fiat-Shamir transform to 5-round protocols. Adapting this attack to our context yields a forgery cost of

$$\text{cost}_{\text{forge}} := \min_{\tau_1, \tau_2: \tau_1 + \tau_2 = \tau} \left\{ \frac{1}{\sum_{i=\tau_1}^{\tau} \binom{\tau}{i} p^i (1-p)^{\tau-i}} + N^{\tau_2} \right\} \quad (4)$$

with p defined in Equation (3). This is substantially lower than the target forgery cost of $1/\varepsilon$, for ε being the soundness error of Protocol 1 (see Theorem 4). We therefore need to adapt the parameters to fill this gap.

4.2 Description of the Signature Scheme

In our signature scheme, the key generation algorithm randomly samples a syndrome decoding instance (H, y) of the syndrome decoding problem with solution x (i.e. $y = Hx$) with security parameter λ . In order to make the key pair compact, the matrix H is pseudorandomly generated from a λ -bit seed. Specifically, a call to the KeyGen algorithm outputs a pair $(pk, sk) := ((\text{seed}_H, y), \text{mseed})$ generated as follows:

1. $\text{mseed} \leftarrow \{0, 1\}^\lambda$
2. $(\text{seed}_H, x) \leftarrow \text{PRG}(\text{mseed})$ where x is sampled in $\{x \in \mathbb{F}_2^m \mid \text{wt}(x) = w\}$
3. $H \leftarrow \text{PRG}(\text{seed}_H)$
4. $y = Hx$; $pk = (\text{seed}_H, y)$; $sk = \text{mseed}$

For the sake of simplicity, we omit the re-generation of H and x from the seeds in the algorithms below and assume $pk = (H, y)$ and $sk = (H, y, x)$.

Given a secret key $sk = (H, y, x)$ and a message $m \in \{0, 1\}^*$, the algorithm Sign proceeds as described in Figure 1. And given a public key $pk = (H, y)$, a signature σ and a message $m \in \{0, 1\}^*$, the algorithm Verif proceeds as described in Figure 2. For the sake of clarity, as for the protocol description in Section 3.3, we use the bold notation to represent a tuple of d polynomials or of d points.

4.3 Signature Properties

We now state the security of our signature scheme in the following theorem. The proof is provided in the full version [?].



Fig. 1. Code-based signature scheme - Signing algorithm.

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

1. Parse the signature σ as

$$\text{salt} \mid h_1 \mid h_2 \mid \left((\text{state}_i^{[e]})_{i \neq i^*[e]} \mid \text{com}_{i^*[e]}^{[e]} \mid \{ \llbracket \alpha_j^{[e]} \rrbracket_{i^*[e]} \}_{j \in [t]} \mid \{ \llbracket \beta_j^{[e]} \rrbracket_{i^*[e]} \}_{j \in [t]} \right)_{e \in [\tau]} .$$
2. Extend hash $\{r_j^{[e]}, \mathbf{e}_j^{[e]}\}_{e \in [\tau], j \in [t]} \leftarrow \text{PRG}(h_1)$ where $r_j^{[e]} \in \mathbb{F}_{\text{points}}$ and $\mathbf{e}_j^{[e]} \in \mathbb{F}_{\text{points}}^d$.
3. Extend hash $\{i^*[e]\}_{e \in [\tau]} \leftarrow \text{PRG}(h_2)$ where $i^*[e] \in [N]$.
4. For each iteration $e \in [\tau]$,
 - For each $i \neq i^*[e]$, computes $\text{com}_i^{[e]} = \text{Hash}_0(\text{salt}, e, i, \text{state}_i^{[e]})$.
 - Using $\{\text{state}_i^{[e]}\}_{i \neq i^*[e]}$, simulate all the parties except for $i^*[e]$ as in the Phase 3 of the signing algorithm and get $\llbracket \alpha_1 \rrbracket, \dots, \llbracket \alpha_t \rrbracket, \llbracket \beta_1 \rrbracket, \dots, \llbracket \beta_t \rrbracket, \llbracket v \rrbracket$ for all parties except for $i^*[e]$.
 - Compute $\llbracket v_j^{[e]} \rrbracket_{i^*[e]} := -\sum_{i \neq i^*[e]} \llbracket v_j^{[e]} \rrbracket_i$ for all $j \in [t]$.
5. Compute $h'_1 = \text{Hash}_1(m, \text{com}_1^{[1]}, \text{com}_2^{[1]}, \dots, \text{com}_{N-1}^{[\tau]}, \text{com}_N^{[\tau]})$.
6. Compute $h'_2 = \text{Hash}_2(m, \{ \llbracket \alpha_j^{[e]} \rrbracket, \llbracket \beta_j^{[e]} \rrbracket, \llbracket v_j^{[e]} \rrbracket \}_{j \in [t], e \in [\tau]})$.
7. Output ACCEPT iff $h'_1 \stackrel{?}{=} h_1$ and $h'_2 \stackrel{?}{=} h_2$.

Fig. 2. Code-based signature scheme - Verification algorithm.

Theorem 5. *Suppose the PRG used is $(t, \varepsilon_{\text{PRG}})$ -secure and any adversary running in time t has at most an advantage ε_{SD} against the underlying d -split syndrome decoding problem. Model Hash_0 , Hash_1 and Hash_2 as random oracles where Hash_0 , Hash_1 and Hash_2 have 2λ -bit output length. Then chosen-message adversary against the signature scheme depicted in Figure 1, running in time t , making q_s signing queries, and making q_0, q_1, q_2 queries, respectively, to the random oracles, succeeds in outputting a valid forgery with probability*

$$\Pr[\text{Forge}] \leq \frac{(q_0 + \tau N q_s)^2}{2 \cdot 2^{2\lambda}} + \frac{q_s(q_s + q_0 + q_1 + q_2)}{2^{2\lambda}} + q_s \cdot \tau \cdot \varepsilon_{\text{PRG}} + \varepsilon_{\text{SD}} + q_2 \cdot \varepsilon^\tau,$$

where $\varepsilon = p + \frac{1}{N} - p \cdot \frac{1}{N}$ and p defined in Equation (3).

4.4 Parameters

In what follows, we propose three parameter sets which achieve a security level of 128 bits for the signature:

- the first one shall rely on the hardness to solve the SD problem on \mathbb{F}_2 ;
- the second one shall also rely on the hardness to solve the SD problem on \mathbb{F}_2 , but we shall use a d -split version to get polynomials over a chosen field, concretely \mathbb{F}_{256} ;
- the last one shall rely on the hardness to solve the SD problem on \mathbb{F}_{256} .

Choice of the SD parameters. Let us first describe how we estimate the security level of a syndrome decoding instance for a random linear code over \mathbb{F}_2 . The best *practical* attack for our parameters is the algorithm of May, Meurer and

Thomae [MMT11]. As argued in [?], we can lower bound the cost of this attack by only considering the cost of its topmost recursion step:

$$\frac{\binom{m}{w}}{\binom{k+\ell}{p} \binom{m-k-\ell}{w-p}} \cdot \left(L + \frac{L^2}{2^{\ell-p}} \right) \quad \text{with } L := \frac{\binom{k+\ell}{p/2}}{2^p}.$$

As usual in ISD algorithm we need to optimize for the parameters ℓ (a number of rows) and p (a partial Hamming weight). Since we only account for the cost of the topmost level in the algorithm, this yields a slightly conservative estimate for the security level. We use this estimate to choose the parameters of our scheme.

Given these considerations, we suggest the following concrete parameters:

- Variant 1: standard binary syndrome decoding problem. We propose the parameters

$$(q, m, k, w, d) = (2, 1280, 640, 132, 1)$$

which achieve a security level of 128 bits according to the above formula.

- Variant 2: d -split binary syndrome syndrome decoding problem, where d is taken to have $m/d \leq 256$ so that $\mathbb{F}_{\text{poly}} = \mathbb{F}_{256}$. We propose the parameters

$$(q, m, k, w, d) = (2, 1536, 888, 120, 6)$$

which achieve a security of 129 bits. Indeed, the standard SD problem with the same parameters (but $d = 1$) has a security of 145 bits and we know, thanks to the Theorem 2, that there is a security loss of at most 16 bits while switching to $d = 6$. **Let us stress that the reduction we use is very conservative since the current state of the art does not contain attacks which fill the gap of the reduction. The aim was here to build a practical signature scheme with conservative security, but searching more aggressive parameters which fit the state of the art would be interesting for a future work.**

- Variant 3: syndrome decoding instance defined over \mathbb{F}_{256} . The cryptanalysis of the syndrome decoding problem on a field which is larger than \mathbb{F}_2 has been less studied. Previous articles [CVE11,?] propose parameters sets for syndrome decoding instances over \mathbb{F}_{2^s} where the code length m is between 200 and 210. In our case, we choose $m = 256$ in such a way that the polynomial degree is equal to the field size. Besides being more conservative, this choice has the advantage of easing the use of a Fast Fourier Transform. We propose the following parameters⁴ for this variant:

$$(q, m, k, w, d) = (256, 256, 128, 80, 1) .$$

Choice of the MPC parameters. For each variant, we suggest in Table 3 a parameter set for the MPC protocol.

⁴ More cryptanalysis of the SD problem over \mathbb{F}_{256} would be welcome to get more confidence in the choice of the parameters. Such research is out of the scope of present article.

| Scheme | SD Parameters | | | | | MPC Parameters | | | |
|-----------|---------------|------|-----|-----|-----|------------------------------|--------------------------------|-----|-------------------|
| | q | m | k | w | d | $ \mathbb{F}_{\text{poly}} $ | $ \mathbb{F}_{\text{points}} $ | t | p |
| Variant 1 | 2 | 1280 | 640 | 132 | 1 | 2^{11} | 2^{22} | 6 | $\approx 2^{-69}$ |
| Variant 2 | 2 | 1536 | 888 | 120 | 6 | 2^8 | 2^{24} | 5 | $\approx 2^{-79}$ |
| Variant 3 | 2^8 | 256 | 128 | 80 | 1 | 2^8 | 2^{24} | 5 | $\approx 2^{-78}$ |

Table 3. SD and MPC parameters.

To have a short signature, we take the smallest possible field \mathbb{F}_{poly} since a signature transcript includes polynomials on that field. As explained in Section 3, \mathbb{F}_{poly} must be a field extension of \mathbb{F}_{SD} which verifies the relation $|\mathbb{F}_{\text{poly}}| \geq m/d$. Then, it remains to choose $|\mathbb{F}_{\text{points}}|$ and t . These parameters are chosen to make the false positive probability p is negligible compared to $1/N$ such that the optimal forgery strategy of an attacker is to take $\tau_1 = 1$ in the Equation (4). As a result, we just need to increase the number of iterations τ by one compared to the interactive protocol.

4.5 Implementation and Performances

For each repetition in the computation of each party, d polynomial interpolations are involved. Indeed, from $\llbracket x \rrbracket$, the parties must compute

$$\llbracket S_\ell \rrbracket(X) = \sum_{i=1}^{m/d} \llbracket x_{\frac{m}{d}\ell+i} \rrbracket \cdot \prod_{j=1, j \neq i}^{m/d} \frac{X - w_j}{w_i - w_j}$$

for all $\ell \in [d]$. Then, the parties must evaluate $\llbracket S_\ell \rrbracket$ in t random evaluation points sampled by the verifier, for all $\ell \in [d]$. The natural way to implement that is to compute the coefficients of all the polynomials $\{\llbracket S_\ell \rrbracket\}_\ell$ from $\llbracket x \rrbracket$, then to evaluate these polynomials t times. However this implies that the signer must realize $\tau \cdot N \cdot d$ interpolations. Instead, the signer can compute the vector $u(r)$ defined as

$$u(r) = \left(\prod_{j=1, j \neq i}^{m/d} \frac{r - w_j}{w_i - w_j} \right)_{1 \leq i \leq \frac{m}{d}}$$

for each evaluation point r , and then use these vectors in the computation of all the parties as

$$\llbracket S_\ell(r) \rrbracket = \langle \llbracket x_\ell \rrbracket, u(r) \rangle$$

where $\llbracket x_\ell \rrbracket$ is the ℓ^{th} chunk of $\llbracket x \rrbracket$. By proceeding this way, the number of (transposed) interpolations done by the signer is of $\tau \cdot t$.

To reduce the computational cost of the interpolations, we can make use of a Fast Fourier Transform (FFT). We are working on field extensions of \mathbb{F}_2 , so we can use the Additive FFT independently introduced by Wang-Zhu

in 1988 [WZ88] and by Cantor in 1989 [Can89], which was further improved in [vzGG03,GM10]. Although such additive FFT exists for any extension of \mathbb{F}_2 , the algorithms are simpler for a field of size $2^{(2^i)}$ for some i , which is why we define \mathbb{F}_{poly} as \mathbb{F}_{256} . On such a field \mathbb{F} , we indeed have an efficient additive FFT using $\frac{1}{2}|\mathbb{F}| \log_2 |\mathbb{F}|$ multiplications to evaluate a polynomial (of degree lower than $|\mathbb{F}|$) in $|\mathbb{F}|$ points.

We implemented the signature scheme in C. In our implementation, the pseudo-randomness is generated using AES in counter mode and the hash function is instantiated with SHAKE. We benchmarked our scheme on a 3.8 GHz Intel Core i7 CPU with support of AVX2 and AES instructions. All the reported timings were measured on this CPU while disabling Intel Turbo Boost.

Remark 2. Another motivation for using $\mathbb{F}_{\text{poly}} = \mathbb{F}_{256}$ is that some Intel processors have dedicated instructions for \mathbb{F}_{256} arithmetic. We therefore expect substantial speed-ups for the instances of our signature scheme using $\mathbb{F}_{\text{poly}} = \mathbb{F}_{256}$ on these processors. Optimizing and benchmarking such implementations is left for future research.

We instantiate two trade-offs per variant: the first one lowering communication cost to produce short signatures, and the second one lowering computational cost to get a fast signature computation. We obtain the parameters and sizes described in Table 4. We provide the measured computational performances of our signature implementation in Table 5.

| λ | Scheme | Aim | Parameters | | $ \text{pk} $ | $ \text{sk} $ | Signature | |
|-----------|-----------|-------|------------|--------|---------------|---------------|----------------------|---------------------------|
| | | | N | τ | | | $ \text{sgn} $ (max) | $ \text{sgn} $ (avg, std) |
| 128 | Variant 1 | Fast | 32 | 27 | 96 | 16 | 16 422 | 16 006, 446 |
| 128 | | Short | 256 | 17 | 96 | 16 | 11 193 | 11 160, 127 |
| 128 | Variant 2 | Fast | 32 | 27 | 97 | 16 | 17 866 | 17 406, 494 |
| 128 | | Short | 256 | 17 | 97 | 16 | 12 102 | 12 066, 141 |
| 128 | Variant 3 | Fast | 32 | 27 | 144 | 16 | 12 115 | 11 835, 302 |
| 128 | | Short | 256 | 17 | 144 | 16 | 8 481 | 8 459, 86 |

Table 4. Parameters (N, τ) with the achieved communication costs (in bytes).

Future investigations. We tried to optimize the implementation using some algorithmic tricks, but we did not yet investigate the possible software optimizations like vectorialization or bitslicing. Although the variants 1 and 2 are more conservative because they rely on the hardness of the binary syndrome decoding problem, variant 3 is more promising in terms of signature size and computation time. While we have investigated parameter sets where \mathbb{F}_{SD} is a field extension of \mathbb{F}_2 , more cryptanalysis for the SD problem on those fields as well as on non-binary fields would be welcome. An interesting idea would be to

| λ | Variant | Aim | Keygen | Sign | Verify |
|-----------|-----------|-------|---------------------------|-------------------------------|-------------------------------|
| 128 | Variant 1 | Fast | n/a [†] | | |
| 128 | | Short | | | |
| 128 | Variant 2 | Fast | 0.03 ms 114 162 cycles | 13.4 ms 52 463 114 cycles | 12.7 ms 50 306 845 cycles |
| 128 | | Short | 0.03 ms 113 852 cycles | 64.2 ms 251 099 099 cycles | 60.7 ms 243 055 474 cycles |
| 128 | Variant 3 | Fast | 0.01 ms 49 181 cycles | 6.4 ms 25 253 580 cycles | 5.9 ms 23 816 143 cycles |
| 128 | | Short | 0.01 ms 49 057 cycles | 29.5 ms 114 226 505 cycles | 27.1 ms 108 541 768 cycles |

Table 5. Benchmarks of our signature implementation. Timings are averaged over 10 000 measurements. The CPU clock cycles have been measured using SUPERCOP (<https://bench.cr.yp.to/supercop.html>).

[†] We only have a proof of concept implementation with irrelevant timings.

instantiate our scheme with a prime field \mathbb{F}_{SD} for which the Number-Theoretic Transform (NTT) is defined. If \mathbb{F}_{SD} is large enough, we could then take the same field for \mathbb{F}_{poly} than \mathbb{F}_{SD} and we would have fast polynomial interpolations and simpler multiplication operations.

5 Comparison

In this section, we compare our scheme to different code-based and post-quantum signature schemes from the literature.

5.1 Comparison with Other Code-Based Signature Schemes

In the state of the art, there exist two approaches to build signatures. On one hand, there is the hash-and-sign paradigm which relies on the existence of a (code-based) trapdoor permutation. Wave [DST19] is a code-based signature scheme in this paradigm. Such schemes are often more vulnerable to structural attacks. On the other hand, further signature schemes are constructed by applying the Fiat-Shamir transform to (zero-knowledge) identification schemes, which can rely on weaker assumptions (and typically the SD for random linear codes). Historically such schemes, like the famous Stern protocol, give rise to large signatures because of the high soundness error of the underlying identification scheme (2/3 or 1/2). To avoid this issue, a solution consists in relying on different code-based problems. For instance, LESS is a recent scheme which security relies on the hardness of the Linear Code Equivalence problem [BMPS20,?]. Another direction is to find a way to adapt the Schnorr-Lyubashevsky approach to code-based cryptography. Durandal is a recent scheme following this approach [ABG⁺19]. More recently, some works [?,?,BGKM22] have obtained better soundness by relying on the MPC-in-the-Head principle. The proposed

schemes achieve small signature sizes at the cost of slower computation. Depending on the setting, they can produce signatures with different trade-offs between the signature size and the computational cost. The current work follows this approach while achieving better trade-offs than any of these previous works.

| Scheme Name | Year | $ \text{sgn} $ | $ \text{pk} $ | t_{sgn} | t_{verif} | Assumption |
|--------------------|------|----------------|---------------|------------------|--------------------|--------------------------------------------------------------------------|
| Wave | 2019 | 2.07 K | 3.2 M | 300 | - | SD over \mathbb{F}_3 (large weight) ($U, U + V$)-codes indisting. |
| Durandal - I | 2018 | 3.97 K | 14.9 K | 4 | 5 | Rank SD over \mathbb{F}_{2^m} |
| Durandal - II | 2018 | 4.90 K | 18.2 K | 5 | 6 | Rank SD over \mathbb{F}_{2^m} |
| LESS-FM - I | 2020 | 9.77 K | 15.2 K | - | - | Linear Code Equivalence |
| LESS-FM - II | 2020 | 206 K | 5.25 K | - | - | Perm. Code Equivalence |
| LESS-FM - III | 2020 | 11.57 K | 10.39 K | - | - | Perm. Code Equivalence |
| [?]-256 | 2021 | 24.0 K | 0.11 K | - | - | SD over \mathbb{F}_{256} |
| [?]-1024 | 2021 | 19.8 K | 0.12 K | - | - | SD over \mathbb{F}_{1024} |
| [?] (fast) | 2021 | 22.6 K | 0.09 K | 13 | 12 | SD over \mathbb{F}_2 |
| [?] (short) | 2021 | 16.0 K | 0.09 K | 62 | 57 | SD over \mathbb{F}_2 |
| [BGKM22] - Sig1 | 2022 | 23.7 K | 0.1 K | - | - | SD over \mathbb{F}_2 |
| [BGKM22] - Sig2 | 2022 | 20.6 K | 0.2 K | - | - | (QC)SD over \mathbb{F}_2 |
| Our scheme - Var1f | 2022 | 15.6 K | 0.09 K | - | - | SD over \mathbb{F}_2 |
| Our scheme - Var1s | 2022 | 10.9 K | 0.09 K | - | - | SD over \mathbb{F}_2 |
| Our scheme - Var2f | 2022 | 17.0 K | 0.09 K | 13 | 13 | SD over \mathbb{F}_2 |
| Our scheme - Var2s | 2022 | 11.8 K | 0.09 K | 64 | 61 | SD over \mathbb{F}_2 |
| Our scheme - Var3f | 2022 | 11.5 K | 0.14 K | 6 | 6 | SD over \mathbb{F}_{256} |
| Our scheme - Var3s | 2022 | 8.26 K | 0.14 K | 30 | 27 | SD over \mathbb{F}_{256} |

Table 6. Comparison of our scheme with signatures from the literature (128-bit security). The sizes are in bytes and the timings are in milliseconds. Reported timings are from the original publications: Wave has been benchmarked on a 3.5 Ghz Intel Xeon E3-1240 v5, Durandal on a 2.8 Ghz Intel Core i5-7440HQ, while [?] and our scheme on a 3.8 GHz Intel Core i7.

Table 6 compares the performances of our scheme with the current code-based signature state of the art, for the 128-bit security level.⁵ We observe that our scheme outperforms all the existing code-based signatures for the $|\text{sgn}| + |\text{pk}|$ metric. Depending on the parameters, it can even produce signatures such that $|\text{sgn}| + |\text{pk}|$ is below the symbolic cap of 10 KB. Regardless of the key size, Wave still achieves the shortest signatures. In terms of security, our scheme has the advantage of relying on the hardness of one of the oldest problems of the code-based cryptography, namely the syndrome decoding for random linear codes in Hamming weight metric.

⁵ We did not include “Sig 3” from [BGKM22] since it is similar to [?] with slight differences (*message decoding* setting) which do not improve the scheme.

5.2 Comparison with other Post-Quantum Signature Schemes

Finally, we compare in Table 7 our construction with other signature schemes aiming at post-quantum security. **First of all, let note that the lattice-based signature scheme (as Dilithium [?] and Falcon [?]) are currently the most efficient post-quantum schemes. They have small signature size and very efficient running times. However, the goal of our construction is to propose a signature scheme based on an alternative problem since it is important to diversify the security assumptions.** All the others schemes have very short public keys and secret keys (less than 150 bytes for 128-bit security), thus it is not a point for comparison. Depending on the chosen parameters, our scheme can be competitive with Picnic3 [KZ20b] and the recently proposed “Picnic4” [KZ21] which also rely on the MPC-in-the-Head paradigm. Like Picnic4, it can produce signatures with a size of around 8 KB. However, our scheme is arguably more conservative in terms of security since Picnic is based on the hardness of inverting LOWMC [ARS⁺15], a cipher with unconventional design choices, while our scheme is based on the hardness of the syndrome decoding problem on linear codes, which has a long cryptanalysis history and is believed to be very robust. Banquet [BdK⁺21] is a signature scheme for which the security is based on the hardness of inverting AES (instead of LowMC), which can also be argued to be a conservative choice. Our scheme over \mathbb{F}_2 is competitive with Banquet: slightly shorter and slightly slower (but the timing could be optimized). On the other hand, our scheme on \mathbb{F}_{256} clearly outperforms Banquet. Our scheme can also be competitive with SPHINCS⁺ [BHK⁺19] depending on the exact criteria. For similar signature sizes, our signature computation is significantly faster while our signature verification is significantly slower than those of SPHINCS⁺.

Acknowledgements. This work has been supported by the European Union’s H2020 Programme under grant agreement number ERC-669891.

References

- AABN02. Michel Abdalla, Jee Hea An, Mihir Bellare, and Chanathip Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Heidelberg, April / May 2002.
- ABG⁺19. Nicolas Aragon, Olivier Blazy, Philippe Gaborit, Adrien Hauteville, and Gilles Zémor. Durandal: A rank metric based signature scheme. In *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 728–758. Springer, Heidelberg, May 2019.
- ACBH13. Sidi Mohamed El Yousfi Alaoui, Pierre-Louis Cayrel, Rachid El Bansarkhani, and Gerhard Hoffmann. Code-Based Identification and Signature Schemes in Software. In *Security Engineering and Intelligence Informatics - CD-ARES 2013 Workshops: MoCrySEn and SeCIHD, Regensburg*,

| Scheme Name | sgn | pk | t_{sgn} | t_{verif} |
|----------------------------|--------|--------|------------------|--------------------|
| SPHINCS ⁺ -128f | 16.7 K | 0.03 K | 14 | 1.7 |
| SPHINCS ⁺ -128s | 7.7 K | 0.03 K | 239 | 0.7 |
| Picnic3 | 12.3 K | 0.03 K | 5.2 | 4.0 |
| Picnic4 | 7.8 K | 0.03 K | ≈ 20 | ≈ 20 |
| Banquet (fast) | 19.3 K | 0.03 K | 6 | 5 |
| Banquet (short) | 13.0 K | 0.03 K | 44 | 40 |
| Dilithium2 | 2.4 K | 1.3 K | 0.065 | 0.024 |
| Falcon-512 | 0.65 K | 0.88 K | 0.168 | 0.036 |
| Our scheme - Var1f | 15.6 K | 0.09 K | - | - |
| Our scheme - Var1s | 10.9 K | 0.09 K | - | - |
| Our scheme - Var2f | 17.0 K | 0.09 K | 13 | 13 |
| Our scheme - Var2s | 11.8 K | 0.09 K | 64 | 61 |
| Our scheme - Var3f | 11.5 K | 0.14 K | 6 | 6 |
| Our scheme - Var3s | 8.3 K | 0.14 K | 30 | 27 |

Table 7. Comparison of our scheme with signatures from the literature. The sizes are in bytes and the timings are in milliseconds. **Reported timings for Falcon have been benchmarked on a 2.3 Ghz Intel Core i5-8259U in [?], and timings for Dilithium and our scheme have been benchmarked on a 3.8 Ghz Intel Core i7.** The benchmarks of the other schemes have been realized on a Intel Xeon W-2133 CPU at 3.60GHz, the values for SPHINCS⁺ and Banquet have been extracted from [BdK⁺21] while the values for Picnic3 have been extracted from its original publication [KZ20b].

- Germany, September 2-6, 2013. *Proceedings*, volume 8128 of *Lecture Notes in Computer Science*, pages 122–136. Springer, 2013.
- AGS11. Carlos Aguilar, Philippe Gaborit, and Julien Schrek. A new zero-knowledge code based identification scheme with reduced communication. In *2011 IEEE Information Theory Workshop*, pages 648–652, 2011.
- ARS⁺15. Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tjessen, and Michael Zohner. Ciphers for MPC and FHE. In *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 430–454. Springer, Heidelberg, April 2015.
- BBC⁺19. Marco Baldi, Alessandro Barenghi, Franco Chiaraluce, Gerardo Pelosi, and Paolo Santini. A Finite Regime Analysis of Information Set Decoding Algorithms. *Algorithms*, 12(10):209, 2019.
- BBPS21. Alessandro Barenghi, Jean-Francois Biase, Edoardo Persichetti, and Paolo Santini. LESS-FM: Fine-tuning Signatures from the Code Equivalence Problem. International Workshop on Post-Quantum Cryptography (PQCrypto), 2021.
- BdK⁺21. Carsten Baum, Cyprien de Saint Guilhem, Daniel Kales, Emmanuela Orsini, Peter Scholl, and Greg Zaverucha. Banquet: Short and fast signatures from AES. In *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 266–297. Springer, Heidelberg, May 2021.
- Bea92. Donald Beaver. Efficient multiparty protocols using circuit randomization. In *CRYPTO’91*, volume 576 of *LNCS*, pages 420–432. Springer, Heidelberg, August 1992.

- BGKM22. Loïc Bidoux, Philippe Gaborit, Mukul Kulkarni, and Victor Mateu. Code-based Signatures from New Proofs of Knowledge for the Syndrome Decoding Problem, 2022.
- BHK⁺19. Daniel J. Bernstein, Andreas Hülsing, Stefan Kölbl, Ruben Niederhagen, Joost Rijneveld, and Peter Schwabe. The SPHINCS⁺ signature framework. In *ACM CCS 2019*, pages 2129–2146. ACM Press, November 2019.
- BMPS20. Jean-François Biasse, Giacomo Micheli, Edoardo Persichetti, and Paolo Santini. LESS is more: Code-based signatures without syndromes. In *AFRICACRYPT 20*, volume 12174 of *LNCS*, pages 45–65. Springer, Heidelberg, July 2020.
- BN20. Carsten Baum and Ariel Nof. Concretely-efficient zero-knowledge arguments for arithmetic circuits and their application to lattice-based cryptography. In *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 495–526. Springer, Heidelberg, May 2020.
- Can89. David G. Cantor. On arithmetical algorithms over finite fields. *Journal of Combinatorial Theory, Series A*, 50:285–300, 1989.
- CDG⁺. Melissa Chase, David Derler, Steven Goldfeder, Jonathan Katz, Vladimir Kolesnikov, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, Xiao Wang, and Greg Zaverucha. The Picnic Signature Scheme – Design Document. Version 2.2 – 14 April 2020.
- CHR⁺16. Ming-Shing Chen, Andreas Hülsing, Joost Rijneveld, Simona Samardjiska, and Peter Schwabe. From 5-pass MQ-based identification to MQ-based signatures. In *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 135–165. Springer, Heidelberg, December 2016.
- CVE11. Pierre-Louis Cayrel, Pascal Véron, and Sidi Mohamed El Yousfi Alaoui. A zero-knowledge identification scheme based on the q-ary syndrome decoding problem. In *SAC 2010*, volume 6544 of *LNCS*, pages 171–186. Springer, Heidelberg, August 2011.
- DST19. Thomas Debris-Alazard, Nicolas Sendrier, and Jean-Pierre Tillich. Wave: A new family of trapdoor one-way preimage sampleable functions based on codes. In *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 21–51. Springer, Heidelberg, December 2019.
- EDV⁺12. Sidi Mohamed El Yousfi Alaoui, Özgür Dagdelen, Pascal Véron, David Galindo, and Pierre-Louis Cayrel. Extended security arguments for signature schemes. In *AFRICACRYPT 12*, volume 7374 of *LNCS*, pages 19–34. Springer, Heidelberg, July 2012.
- FJR21. Thibault Feneuil, Antoine Joux, and Matthieu Rivain. Shared Permutation for Syndrome Decoding: New Zero-Knowledge Protocol and Code-Based Signature. Cryptology ePrint Archive, Report 2021/1576, 2021.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- GG07. Philippe Gaborit and Marc Girault. Lightweight code-based identification and signature. In *IEEE International Symposium on Information Theory, ISIT 2007, Nice, France, June 24-29, 2007*, pages 191–195. IEEE, 2007.
- GM10. Shuhong Gao and Todd Mateer. Additive Fast Fourier Transforms Over Finite Fields. *IEEE Transactions on Information Theory*, 56(12):6265–6272, 2010.
- GPS21. Shay Gueron, Edoardo Persichetti, and Paolo Santini. Designing a Practical Code-based Signature Scheme from Zero-Knowledge Proofs with Trusted Setup. Cryptology ePrint Archive, Report 2021/1020, 2021.

- IKOS07. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In *39th ACM STOC*, pages 21–30. ACM Press, June 2007.
- KZ20a. Daniel Kales and Greg Zaverucha. An attack on some signature schemes constructed from five-pass identification schemes. In *CANS 20*, volume 12579 of *LNCS*, pages 3–22. Springer, Heidelberg, December 2020.
- KZ20b. Daniel Kales and Greg Zaverucha. Improving the performance of the Picnic signature scheme. *IACR TCHES*, 2020(4):154–188, 2020.
- KZ21. Daniel Kales and Greg Zaverucha. Efficient Lifting for Shorter Zero-Knowledge Proofs and Post-Quantum Signatures. Preliminary Draft, October 29, 2021, 2021.
- LN17. Yehuda Lindell and Ariel Nof. A framework for constructing fast MPC over arithmetic circuits with malicious adversaries and an honest-majority. In *ACM CCS 2017*, pages 259–276. ACM Press, October / November 2017.
- MMT11. Alexander May, Alexander Meurer, and Enrico Thomae. Decoding random linear codes in $\tilde{O}(2^{0.054n})$. In *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 107–124. Springer, Heidelberg, December 2011.
- PS00. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
- Ste94. Jacques Stern. A new identification scheme based on syndrome decoding. In *CRYPTO'93*, volume 773 of *LNCS*, pages 13–21. Springer, Heidelberg, August 1994.
- TS16. Rodolfo Canto Torres and Nicolas Sendrier. Analysis of information set decoding for a sub-linear error weight. In *Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016*, pages 144–161. Springer, Heidelberg, 2016.
- Vér96. Pascal Véron. Improved identification schemes based on error-correcting codes. *Appl. Algebra Eng. Commun. Comput.*, 8(1):57–69, 1996.
- vzGG03. J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 2003.
- WZ88. Y. Wang and X. Zhu. A fast algorithm for the Fourier transform over finite fields and its VLSI implementation. *IEEE Journal on Selected Areas in Communications*, 6(3):572–577, 1988.