Fiat-Shamir signatures without aborts using Ring-and-Noise assumptions*

Dipayan Das, Antoine Joux and Anand Kumar Narayanan

CISPA – Helmholtz Center for Information Security, Saarbrücken, Germany.
dipayan.das@cispa.de, joux@cispa.de, anand.narayanan@cispa.de

Abstract. Lattice and code based hard problems such as Learning With Errors (LWE) or syndrome decoding (SD) form cornerstones of post-quantum cryptography. However, signature schemes built on these assumptions remain rather complicated. Indeed, signature schemes from LWE problems are built on the Fiat-Shamir with abort paradigm with no apparent means for knowledge extraction. On the code side, signature schemes mainly stem from Stern's zero-knowledge identification scheme. However, because of its large soundness error of 2/3, it is costly to turn into a signature scheme. The latest developments rely on complicated cut-and-choose and multiparty-in-the-head techniques. As a consequence, they apply the Fiat-Shamir transformation on protocols with at least 5 rounds, leading to additional complexity and degraded security parameters. In the present paper, we propose an alternative approach to build a simple zero-knowledge Σ-protocol with a small soundness error, based on the hardness of Ring-and-Noise assumptions, a general family of assumptions that encompasses both lattices and codes. With such a Σ-protocol at hand, signatures can directly be derived by invoking the standard Fiat-Shamir transform, without the need for aborts. The main novel tool that allows us to achieve this is the use of specifically tailored locality sensitive hash functions. We outline our schemes for general Ring-and-Noise assumptions and present them in detail for the ring of residues modulo Mersenne numbers endowed with the Hamming metric. This Mersenne setting is ideal to illustrate our schemes, since it is close in spirit to both lattice and code based assumptions.

1 Introduction

Lattice and code based problems are among the most prominent computational hardness assumptions that post-quantum cryptography is built on. They are also the earliest conceived, most well studied and accepted post-quantum hardness assumptions [25,28]. Lattice, code based and various other hardness assumptions may be viewed as part of a unified framework called the Ring-and-Noise assumptions family [21]. Informally, we have a commutative ring $\mathcal{R}$ endowed with a relaxed notion of distance that behaves well with respect to the ring arithmetic. For random ring elements, the weight induced by the distance to zero needs to be at most additive (up to a small constant) with respect to addition . It also needs to be at most multiplicative with respect to multiplication. Fix a secret: $S \in \mathcal{R}^k$ of dimension $k > 1$ over the ring. The generic hardness assumption is that despite having access to an arbitrary large number of independent samples, the distribution $(U, \langle U, S\rangle + \epsilon)$ with a uniform $U \in \mathcal{R}^k$ and an independent random noise $\epsilon \in \mathcal{R}$ of appropriately small weight is computationally indistinguishable from the uniform distribution. To instantiate the LWE problem in this framework, consider the ring

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

of residues modulo an odd number q with the Euclidean distance between representatives in $\{(-q+1)/2, \ldots, (q+1)/2\}$. To instantiate code assumption, use binary vectors together with the Hamming distance. Another example is the Mersenne low Hamming weight assumption [2]. It takes the ring of residues modulo a Mersenne prime $p = 2^n - 1$. The distance between two elements is defined as the Hamming distance between the corresponding n-bit strings of representatives in $\{0, 1, \ldots, p-1\}$.

There are elegant encryption schemes, conceptually simple to state and implement, whose security is reduced to lattice or code based problems such as LWE. However, the state of signature scheme design based on such problems is not as pleasant. The main difficulty seems to be a simple adaptation of Fiat-Shamir paradigm to fit such problems . For instance, known Fiat-Shamir signatures from LWE problems involve an additional *aborting* step [23,24] that is crucial for security. Aborts serve to enforce the independence of the signature distribution from the secret key. There have been a sequence of signature constructions following this path, with significant improvements and based on varied assumptions [14,6,15,5]. An attempt to remove the need for aborts was proposed recently from "ad-hoc" lattice based assumptions [7].

The central caveat to signatures with aborts is the difficultly of knowledge extraction. Extracting the knowledge of the secret from an honest prover (using a rewinding technique) in a zero knowledge identification protocol is critical in security arguments for the Fiat-Shamir transformation. To complicate things further, it is not known how to apply Fiat-Shamir with abort on code based assumptions involving the Hamming distance. However, it was used in a recent code based signature scheme in the rank metric [4].

In the code based world, most signature schemes derive from Stern's code-based Σ protocol [29], which has been adapted to lattices in [22]. However, its large soundness error is often a problem. To bypass it, subsequent work, such as [30,18,26,3,16], use more complicated cut-and-choose and multiparty-in-the-head techniques [20]. Another issue with these post-quantum signature schemes is that they don't easily accommodate more advanced applications like blind signatures [11], multi signatures [10], and threshold signatures [12,13], which can all be instantiated from a classical hardness assumption like the discrete logartihm problem.

In this paper, we present an honest verifier computational zero-knowledge identification protocol for the language of two (or more) noisy linear equations over commutative rings. The noise here is specified as a part of a Ring-and-Noise hardness assumption. Our protocol is a Σ-protocol, following the canonical commit-challenge-respond motif. Soundness and zero knowledge are both guaranteed by a variation of the Ring-and-Noise assumption, which we call the Small Multiplier Ring-and-Noise Assumption. Uniformly choose a secret vector $S \in \mathcal{R}^k$ of dimension $k \geq 2$ over the ring. The hardness assumption is that despite having access to an arbitrary number of samples, the distribution $(U, \langle U, S\rangle + \epsilon)$ with a noise vector $U \in \mathcal{R}^k$ (composed of independent noise coordinates) and an independent noise $\epsilon \in \mathcal{R}$ is computationally indistinguishable from the distribution arising as the direct product of a noise vector U and a uniform vector. Invoking the Fiat-Shamir transformation, we turn this protocol into a simple signature scheme that is unforgeable under the same hardness assumption – in the random oracle model.

The main novelty in our protocol is the use of a locality sensitive hash function as an extractor, to strongly bind the prover during his initial commitment while preventing the undesired

leakage that led to protocols with aborts. In particular, this allows for seamless knowledge extraction. An outline of our identification protocols and signature schemes is first presented in § 2 for Ring-and-Noise assumptions in broad generality. The exposition is at a high level, to facilitate the translation of our scheme to any particular instantiation of the Ring-and-Noise assumption. We then fully demonstrate one such translation in the subsequent two sections. The necessary extractor function tailored to the Hamming metric is described in § 3. The identification protocol and signature scheme are detailed for the ring of residues modulo a Mersenne number in the Hamming metric in § 4. We choose the Mersenne ring to illustrate our schemes for two reasons. First, taking the Mersenne ring with the Hamming metric is close in spirit to both lattice and code based assumptions. Further, the resulting signature scheme works over the Hamming metric, where we do not know how to use the Fiat-Shamir with abort approach.

2 Signature schemes from Ring-and-Noise assumptions

Let $\mathcal{R}$ denote a finite commutative ring with a 'distance' function $d : \mathcal{R} \times \mathcal{R} \to \mathbb{R}_{\geq 0}$. Our notion of distance is weaker than usual. It needs to be symmetric and enforce the identity of indiscernibles. That is, for $X, Y \in \mathcal{R}$, $d(X, Y) = d(Y, X)$ and $d(X, Y) = 0 \Rightarrow X = Y$. Let $\omega : \mathcal{R} \to \mathbb{R}_{\geq 0}$ denote the weight function induced by the distance as $\omega(X) := d(X, 0), \forall X \in \mathcal{R}$. The weight function must behave well with ring arithmetic involving random elements. That is, for uniform $(X, Y) \in \mathcal{R}^2$, with probability close to one,

$$\omega(X + Y) \leq a(\omega(X) + \omega(Y)) \text{ and } \omega(XY) \leq b\omega(X)\omega(Y)$$

for absolute positive constants a and b. The protocols and assumptions will involve 'vectors' over $\mathcal{R}$, by which we mean elements in some finite cartesian power $\mathcal{R}^c$ seen as a free module, with a fixed basis in mind. We use capital latin alphabets to denote generic ring elements and capital greek alphabets for ring elements designed to be of low weight. We use bold face for vectors; latin bold face for generic vectors and greek bold face for vectors of low weight elements.

Examples. A guiding example (realizing the LWE problem) is to take $\mathcal{R}$ to be the ring $\mathbb{Z}/q\mathbb{Z}$ of residues modulo a (possibly composite) odd number q with d being the Euclidean distance between representatives in $\{(-q + 1)/2, \ldots, 0, \ldots, (q - 1)/2\}$. In this case, the constants a and b may be taken to be one. Another example realizes the Mersenne low Hamming weight combination assumption. Take $\mathcal{R}$ to be the ring $\mathbb{Z}/p\mathbb{Z}$ of residues modulo a Mersenne prime $p = 2^n - 1$. Abusing notation, identify residues modulo p with the n-bit strings of representatives in $\{0, 1, \ldots, p\}$. For $X, Y \in \mathbb{Z}/p\mathbb{Z}$, take $d(X, Y)$ to be the Hamming distance between the n-bit vector representatives of X and Y. The distance measure behaves well with ring arithmetic involving random elements with constants $a = 2$ and $b = 1$ [2]. Unlike LWE, randomness is needed to ensure the distance behaves well with arithmetic in this case.

Noise. For a real number ϵ, let $\mathcal{R}_\epsilon$ denote the set of elements of $\mathcal{R}$ of weight at most ϵ. Informally, we call an element of weight much lower than that of a generic element as noise. A vector/matrix all of whose coordinates are noise is called as a noise vector/matrix. Given ϵ, the protocols will assume an appropriate noise distribution to draw elements from $\mathcal{R}_\epsilon$. When the protocols draw random noise vectors/matrices, we mean that the coordinates are chosen independently, each according to the aforementioned noise distribution.

Ring-and-Noise assumptions. Fix a small positive integer parameter $k > 1$. To build intuition, it is convenient to think of k as 2. For technical reasons, such as the proof of uniformity of certain extractor outputs, one may require a slightly larger k. Fix a weight bound σ such that the product of two elements of weight sigma is still smaller than that of a generic element. In the Mersenne example, $\sigma = \Theta((\log(p))^{1/3})$ suffices; since a generic element has weight $\Theta(p)$.

Ring-and-Noise assumption: Let $\mathbf{S} \in \mathcal{R}^k$ be a uniformly random secret vector. Despite access to arbitrarily many samples, $(\mathbf{U}, \langle \mathbf{U}, \mathbf{S} \rangle + \Gamma)$ drawn with uniform $\mathbf{U} \in \mathcal{R}^k$ and random noise $\Gamma \in \mathcal{R}_\sigma$ is computationally indistinguishable from the uniform distribution over in $\mathcal{R}^k \times \mathcal{R}$.

Small Multiplier Ring-and-Noise assumption: Let $\mathbf{S} \in \mathcal{R}^k$ be a uniformly random secret vector. Despite access to arbitrarily many samples, $(\mathbf{\Theta}, \langle \mathbf{\Theta}, \mathbf{S} \rangle + \Gamma)$ drawn with a random noise vector $\mathbf{\Theta} \in \mathcal{R}_\sigma^k$ and a random noise $\Gamma \in \mathcal{R}_\sigma$ is computationally indistinguishable from the product of the noise vector distribution over $\mathcal{R}_\sigma^k$ and the uniform distribution over $\mathcal{R}$.

We present an honest verifier public coin computational zero knowledge identification protocol for the language

$$\big(\text{description of } \mathcal{R}, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta} \mid \mathbf{R} \in \mathcal{R}^k, \text{noise vector } \boldsymbol{\Delta} \in \mathcal{R}_\sigma^k, A \in \mathcal{R}\big)$$

conditioned on the Small Multiplier Ring-and-Noise assumption.

Extractor. A key ingredient in our protocols is an extractor function

$$E : \mathcal{K} \times \mathcal{R} \to \mathcal{S}$$

where $\mathcal{K}$ is a set of size comparable to that of $\mathcal{R}$ and $\mathcal{S}$ is a small finite set. The first argument to the extractor will typically be a public random string from our protocols. We will be interested in the function values while the first argument is drawn at random. Denote the image of $(W, X) \in \mathcal{K} \times \mathcal{R}$ under E by $E_W(X)$. We require for uniform $W \in K$ and appropriate weight thresholds σ_1, σ_2 that

$$\begin{cases} E_W(X) = E_W(Y) \text{ with probability close to 1}, & \text{if } d(X,Y) \le \sigma_1 \\ E_W(X) \text{ and } E_W(Y) \text{ are independent and uniform}, & \text{if } d(X,Y) > \sigma_2. \end{cases}$$

Let ℓ denote the number of samples, chosen large enough to assure soundness. Typically, $\ell = \Theta(\log(|\mathcal{R}|))$ suffices. Identifying $\mathcal{K}^\ell \times \mathcal{R}^\ell$ with $(\mathcal{K} \times \mathcal{R})^\ell$ and applying the extractor E coordinate wise, we get

$$\mathrm{Ext}^\ell : \mathcal{K}^\ell \times \mathcal{R}^\ell \to \mathcal{S}^\ell.$$

The image of $(\mathbf{W}, \mathbf{X}) \in \mathcal{K}^\ell \times \mathcal{R}^\ell$ under Ext^ℓ is denoted by $\mathrm{Ext}^\ell_{\mathbf{W}}(\mathbf{X})$.

Identification schemes. As public randomness, draw a random ℓ by k noise matrix $\mathbf{\Lambda} \in \mathcal{R}_\sigma^{\ell \times k}$ and a uniform $\mathbf{W} \in \mathcal{R}^\ell$.

- Prover: Independently draw random noise vectors $\boldsymbol{\Phi} \in \mathcal{R}_\sigma^k, \boldsymbol{\Psi} \in \mathcal{R}_\sigma^\ell$ and a uniform element $Y \in \mathcal{R}$. Using a commitment protocol, commit to $E := \mathrm{Ext}^\ell_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right)$.

– Verifier: Send a noise $\Gamma \in \mathcal{R}_{\sigma'}$ as a challenge. The challenge weight bound σ' may be much smaller than σ. Informally, the product of two elements of weight σ with another element of weight σ' should be smaller than the weight of a generic element.

– Prover: Send the response $Y + A\Gamma$ to the challenge and open the commitment to reveal E.

– Verifier: Reject if the revealed E is inconsistent with its commitment. Draw a random noise vector $\boldsymbol{\Psi}' \in \mathcal{R}_\sigma^\ell$ and reject if $E' := \mathrm{Ext}_{\mathbf{W}}^\ell\left(\boldsymbol{\Lambda}\left((Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \boldsymbol{\Delta})\right) + \boldsymbol{\Psi}'\right)$ is far from E. Else, accept.

Identification protocol for $\left(\text{description of } \mathcal{R}, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta} \mid \mathbf{R} \in \mathcal{R}^k, \boldsymbol{\Delta} \in \mathcal{R}_\sigma^k, A \in \mathbf{R}\right).$

Public randomness: $\boldsymbol{\Lambda} \in_\$ \mathcal{R}_\sigma^{\ell \times k}$ and $\mathbf{W} \in_\$ \mathcal{R}^\ell$

Prover		**Verifier**
Draw $\boldsymbol{\Phi} \in_\$ \mathcal{R}_\sigma^k, \boldsymbol{\Psi} \in_\$ \mathcal{R}_\sigma^\ell$ and $Y \in_\$ \mathcal{R}$.		
Compute $E := \mathrm{Ext}_{\mathbf{W}}^\ell\left(\boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right)$.		
	$\xrightarrow{\text{Commitment}(E)}$	
	$\xleftarrow{\text{Challenge } \Gamma \in_\$ \mathcal{R}_{\sigma'}}$	
	$\xrightarrow{Y + A\Gamma,\ \text{Opening}(E)}$	
		Draw $\boldsymbol{\Psi}' \in_\$ \mathcal{R}_\sigma^\ell$ and compute
		$E' := \mathrm{Ext}_{\mathbf{W}}^\ell\left(\boldsymbol{\Lambda}\left((Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \boldsymbol{\Delta})\right) + \boldsymbol{\Psi}'\right)$.
		Accept if and only if E and E' are close.

Signatures. We transform the identification scheme into a signature scheme using Fiat-Shamir [17]. Fix a public cryptographic hash function H whose image is restricted to $\mathcal{R}_{\sigma'}$. Draw a public random matrix $\boldsymbol{\Lambda}$ and an extractor seed $\mathbf{W}$; just as in the identification protocol. The public verification key is the instance

$$(\text{description of } \mathcal{R}, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta}, \boldsymbol{\Lambda}, \mathbf{W})$$

with independently chosen uniform $\mathbf{R} \in \mathcal{R}^k$, uniform $A \in \mathcal{R}$ and noise vector $\boldsymbol{\Delta} \in \mathcal{R}_\sigma^k$. To reduce the key size in practice, the public randomness $\boldsymbol{\Lambda}, \mathbf{W}$ may be generated using a fixed public pseudorandom generator and a public seed. The public seed is then made part of the public verification key. The signer knows the private key A (and hence also knows $\boldsymbol{\Delta}$). Let M denote the message. The signer first extracts $\mathrm{Ext}_{\mathbf{W}}^\ell\left(\boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right)$ as in the identification scheme. Without interaction, the signer then prepares the challenge Γ as the hash of the extraction concatenated with the message

$$\Gamma := H\left(\mathrm{Ext}_{\mathbf{W}}^\ell\left(\boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right), M\right).$$

The signer appends the response $Y + A\Gamma$ and the extraction as the signature

$$\left(Y + A\Gamma, \mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \mathbf{\Phi}) + \mathbf{\Psi}\right)\right)$$

of M. The verifier checks that

$$\mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \mathbf{\Phi}) + \mathbf{\Psi}\right) \text{ and } \mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\mathbf{\Lambda}\left((Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \mathbf{\Delta})\right)\right)$$

are close in agreement. The left hand side was revealed in the signature, which also allows the verifier to compute the challenge Γ using the hash H. The right hand side is computed from the knowledge of the challenge Γ, the response $Y + A\Gamma$ and the randomness $\mathbf{\Lambda}$.

Signature scheme.

Key Generation: $\mathbf{R} \in_{\$} \mathcal{R}^k, A \in_{\$} \mathcal{R}, \mathbf{\Delta} \in_{\$} \mathcal{R}^k_{\sigma}$.

Private signing key: A.

Public verification key: $\mathbf{R}, A\mathbf{R} + \mathbf{\Delta}, \mathbf{\Lambda} \in_{\$} \mathcal{R}^{\ell \times k}_{\sigma}, \mathbf{W} \in_{\$} \mathcal{R}^{\ell}$.

Public hash function H with range $\mathcal{R}_{\sigma'}$.

Signer	**Verifier**
Generate a message M.	
Draw $\mathbf{\Phi} \in_{\$} \mathcal{R}^k_{\sigma}, \mathbf{\Psi} \in_{\$} \mathcal{R}^{\ell}_{\sigma}, Y \in_{\$} \mathcal{R}$ and extract	
$E := \mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \mathbf{\Phi}) + \mathbf{\Psi}\right)$.	
Generate the challenge $\Gamma := H\left(E, M\right)$.	
$\xrightarrow{M,\ \text{Signature}= (Y + A\Gamma, E)}$	
	Compute the challenge $\Gamma := H\left(E, M\right)$.
	Verify if $\mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\mathbf{\Lambda}\left((Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \mathbf{\Delta})\right)\right)$ is close to E.

We next sketch arguments for the completeness, soundness and zero knowledge of the Σ-protocol. The unforgeability of the signature scheme in the random oracle model follows from the now standard Fiat-Shamir arguments [27,1].

Completeness. Consider an instance (description of $\mathcal{R}, \mathbf{R}, A\mathbf{R} + \mathbf{\Delta}$) that belongs to the language in question and an honest prover with knowledge of a witness A of membership. To claim completeness, the verifier must accept with a constant probability. Assuming compliance to the protocol, the prover's honesty ensures passing of the commitment protocol. By construction,

$$(Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \mathbf{\Delta}) = Y\mathbf{R} - \Gamma\mathbf{\Delta}$$

$$\Rightarrow \mathbf{\Lambda}(Y\mathbf{R} + \mathbf{\Phi}) - \mathbf{\Lambda}((Y + \Gamma A)\mathbf{R} - \Gamma\mathbf{R}') = \mathbf{\Lambda}(\mathbf{\Phi} - \Gamma\mathbf{\Delta}).$$

The righthand side makes it evident that this difference is a vector whose entries are the result of ring arithmetic involving small weight elements. The choice of weight thresholds σ, σ' along

with the fact that k is a small constant ensures that the vector on the right has small entries. Therefore, most coordinates of $\mathbf{\Lambda}(Y\mathbf{R}+\boldsymbol{\Phi})$ and $\mathbf{\Lambda}((Y+\Gamma A)\mathbf{R}-\Gamma(A\mathbf{R}+\boldsymbol{\Delta}))$ are close enough for the extractor to ensure the prover passes verification. The additive noise vectors $\boldsymbol{\Psi}, \boldsymbol{\Psi}'$ do not contribute much to the respective extractions.

Soundness. Consider an instance (description of $\mathcal{R}, \mathbf{R}, \mathbf{R}'$) and a prover $\mathcal{P}$ who convinces an honest verifier $\mathcal{V}$ with probability greater than half. We claim that the instance must belong to the language. That is, there is an $A \in \mathcal{R}$ such that $\mathbf{R}' - A\mathbf{R} \in \mathcal{R}_\sigma^k$. To this end, first challenge $\mathcal{P}$ with a Γ and proceed with the protocol to successful completion. Let Z denote the response of $\mathcal{P}$ in the final step, in place of $Y + A\Gamma$. Rewind the protocol up to the instant after $\mathcal{P}$ commits to E. Now challenge the prover with a new $\widehat{\Gamma}$ distinct from Γ and proceed with the protocol to successful completion. Let $\widehat{Z}$ denote the response of $\mathcal{P}$ in the final step, in place of $Y + A\widehat{\Gamma}$.

Since both runs of the protocol succeeded and were tied to the same E,

$$\mathrm{Ext}_{\mathbf{W}}^{\ell}\left(\mathbf{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}') + \boldsymbol{\Psi}'\right) \text{ and } \mathrm{Ext}_{\mathbf{W}}^{\ell}\left(\mathbf{\Lambda}(\widehat{Z}\mathbf{R} - \widehat{\Gamma}\mathbf{R}') + \boldsymbol{\Psi}'\right) \tag{1}$$

agree at a great fraction of the ℓ coordinates. Consider the maps

$$\begin{array}{ccccc} \mathcal{R} \times \mathcal{R}_{\sigma'} & \xrightarrow{\qquad f_{\mathbf{R},\mathbf{\Lambda}} \qquad} & \mathcal{R}^{\ell} & \xrightarrow{\qquad g_{\mathbf{W}} \qquad} & \mathcal{S}^{\ell} \\ (X, \Omega) & \longmapsto & \mathbf{\Lambda}(X\mathbf{R} - \Omega\mathbf{R}') + \boldsymbol{\Psi}' & \longmapsto & \mathrm{Ext}_{\mathbf{W}}^{\ell}\left(\mathbf{\Lambda}(X\mathbf{R} - \Omega\mathbf{R}') + \boldsymbol{\Psi}'\right) \end{array}$$

parametrised by the $\mathbf{R}$ part of the instance and the shared public randomness $\mathbf{\Lambda}, \mathbf{W}$. The prover found two distinct codewords in the code $g_{\mathbf{W}}(f_{\mathbf{R},\mathbf{\Lambda}}(\mathcal{R} \times \mathcal{R}_{\sigma'})) \subset \mathcal{S}^{\ell}$ that are very close, namely $g_{\mathbf{W}}(f_{\mathbf{R},\mathbf{\Lambda}}(Z, \Gamma))$ and $g_{\mathbf{W}}(f_{\mathbf{R},\mathbf{\Lambda}}(\widehat{Z}, \widehat{\Gamma}))$.

Assume there is no $\mathcal{R}$-linear equation of the form $\mathbf{R}' - A\mathbf{R} \in \mathcal{R}_\sigma^k$. By the Small Multiplier Ring-and-Noise assumption, for each choice of $(X, \Omega) \in \mathcal{R} \times \mathcal{R}_{\sigma'}$ that is not a pair of zeroes, $\mathbf{\Lambda}(X\mathbf{R} - \Omega\mathbf{R}') + \boldsymbol{\Psi}'$ is indistinguishable from uniform in $\mathcal{R}^{\ell}$. Therefore, the code $f_{\mathbf{R},\mathbf{\Lambda}}(\mathcal{R} \times \mathcal{R}_{\sigma'})$ is indistinguishable from a random subset of the same size. The size of the code $f_{\mathbf{R},\mathbf{\Lambda}}(\mathcal{R} \times \mathcal{R}_{\sigma'})$ is small (closer to $|\mathcal{R}|$ than $|\mathcal{R}|^2$); therefore it has a large minimum distance.

The extractor map $g_{\mathbf{W}}$ is locality sensitive. Finding a collision in $\mathcal{R}^{\ell}$ under $g_{\mathbf{W}}$ is hence easy, merely pick two vectors that differ by a noise vector in $\mathcal{R}_{\sigma_1}^{\ell}$. But the prover found something stronger; a collision in the code $f_{\mathbf{R},\mathbf{\Lambda}}(\mathcal{R} \times \mathcal{R}_{\sigma'})$ under $g_{\mathbf{W}}$. A union bound argument shows that there are likely no collisions in a similarly small random subset of $\mathcal{R}^{\ell}$ under $g_{\mathbf{W}}$. By finding a collision, the prover distinguishes the code from a random subset; a contradiction. Therefore our assumption was wrong. There is indeed an $\mathcal{R}$-linear equation of the form $\mathbf{R}' - A\mathbf{R} \in \mathcal{R}_\sigma^k$; meaning our instance is in the language and the protocol is sound.

When the instance is in the language, the collision the prover finds differs by a noise vector

$$\mathbf{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}') - \mathbf{\Lambda}(\widehat{Z}\mathbf{R} - \widehat{\Gamma}\mathbf{R}') \in \mathcal{R}_{\sigma_2}^{\ell}.$$

From the smallness of the coefficients of $\mathbf{\Lambda}$ and the randomness (which ensures that the linear system $\mathbf{\Lambda}$ defines has sufficient rank), we can conclude that

$$Z\mathbf{R} - \Gamma\mathbf{R}' - (\widehat{Z}\mathbf{R} - \widehat{\Gamma}\mathbf{R}') \in \mathcal{R}_\sigma^k \Rightarrow \mathbf{R}' - (\widehat{\Gamma} - \Gamma)^{-1}(Z - \widehat{Z})\mathbf{R} \in \mathcal{R}_\sigma^k.$$

Therefore, the prover knows a witness $(\widehat{\Gamma}-\Gamma)^{-1}(Z-\widehat{Z})$. In this argument, we inverted elements and linear systems as if $\mathcal{R}$ were a field. In the special cases of interest, informally $\mathcal{R}$ will behave as a field for random elements. Otherwise, something else unlikely will happen, such as finding two large factors of a Mersenne number.

Zero Knowledge. Let $\mathcal{V}'$ be an arbitrary verifier. We describe a Simulator $\mathcal{S}$ for $\mathcal{V}'$. The simulator $\mathcal{S}$ is a non interactive probabilistic polynomial time algorithm that for instances in the language produces a distribution that is indistinguishable from the transcript of interaction between an honest prover $\mathcal{P}$ with a witness A of membership and $\mathcal{V}'$. In addition to the instance and shared randomness, $\mathcal{S}$ draws/knows the challenge Γ.

Simulator $\mathcal{S}$: Uniformly draw a challenge $\Gamma \in \mathcal{R}_{\sigma'}$. Independently draw a uniform $Z \in \mathcal{R}$ and a random noise vector $\widehat{\boldsymbol{\Psi}} \in \mathcal{R}^\ell$. Extract $E^s := \mathrm{Ext}^\ell_{\mathbf{W}}\left(\widehat{\boldsymbol{\Psi}} + \boldsymbol{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}')\right)$ and output

$$(\text{a commitment to } E^s, \text{challenge } \Gamma, Z, \text{revelation of } E^s)\,.$$

We claim that the output of $\mathcal{S}$ is indistinguishable from a transcript

$$(\text{a commitment to } E, \text{challenge } \Gamma, Y + A\Gamma, \text{revelation of } E)$$

of the protocol between an honest prover $\mathcal{P}$ (who knows A) and $\mathcal{V}'$. We claim something stronger: that the transcripts are indistinguishable even prior to the extraction. That is,

$$\left(\Gamma, Z, \widehat{\boldsymbol{\Psi}} + \boldsymbol{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}')\right) \text{ and } (\Gamma, Y + A\Gamma, \boldsymbol{\Psi} + \boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}))$$

are indistinguishable. If not, either $\widehat{\boldsymbol{\Psi}} + \boldsymbol{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}')$ or $\boldsymbol{\Psi} + \boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi})$ is distinguishable from the uniform distribution, contradicting the noise multiplier Ring-and-Noise assumption.

3 Extractor design in the Hamming metric

For a positive integer n, define the deterministic function

$$E^n : \{0,1\}^n \times \{0,1\}^n \longrightarrow \{0,1\}$$

$$(U, S) \longmapsto \begin{cases} 1, & \text{if } \sum_i (-1)^{x_i} \geq 0 \\ 0, & \text{else} \end{cases}$$

where $(x_1, x_2, \ldots, x_n) := U \oplus S \in \{0,1\}^n$ denotes the pointwise xor of U and S.

When used in the cryptographic protocols, we will move one of the arguments to the subscript and denote $E(U, S)$ by $E_U(S)$. The reason being that U will often be shared randomness in the protocols and the interest will be in the function values $E_U(S)$ as U is fixed and S varies. The functions $E^n_U()$ are extractors modulo proximity in the Hamming metric. Say U is drawn uniformly at random. When two strings S and $\widehat{S}$ are close in Hamming distance, their extractions $E^n_U(S), E^n_U(\widehat{S})$ will likely be identical. When they are far, their extractions will be independent and uniform. The rest of this section quantifies these two properties. These functions are close in spirit to the locality sensitive hash functions of Indyk and Motwani [19]. Our

extractors however are not linear projections, due to the non linearity thresholding.

To extract more than one bit, we devise a sequence of deterministic functions

$$\mathrm{Ext}^{n,m} : (\{0,1\}^n)^m \times \{0,1\}^n \longrightarrow \{0,1\}^m$$

$$(W, S) \longmapsto \mathrm{Ext}_W^{n,m}(S)$$

indexed by positive integers n and m. The first coordinate $W \in (\{0,1\}^n)^m$ of the argument is presented as a sequence of m n-bit strings $W = (W_1, W_2, \ldots, W_m)$. Define

$$\mathrm{Ext}_W^{n,m}(S) := (E^n(W_1, S), E^n(W_2, S), \ldots, E^n(W_m, S)) .$$

Lemma 1. *Let n be an odd positive integer greater than 1. Let $S, \widehat{S} \in \{0,1\}^n$ be two n-bit strings with Hamming distance $d_H(S, \widehat{S}) \leq d$ bounded by some positive integer d. For $U \in \{0,1\}^n$ drawn uniformly at random,*

$$E^n(U, S) = E^n(U, \widehat{S})$$

with probability at least

$$1 - \frac{1}{\pi}\sqrt{\frac{d}{2n}} + o\left(\frac{d}{n}\right).$$

Proof. Let $X = (x_1, x_2, \ldots, x_n) := U \oplus S \in \{0,1\}^n$ and $Y = (y_1, y_2, \ldots, y_n) := U \oplus \widehat{S} \in \{0,1\}^n$ respectively denote the pointwise xors. Since U is uniformly random in $\{0,1\}^n$, X and Y are (dependent but) both uniformly random in $\{0,1\}^n$. Let $I \subseteq \{1, 2, \ldots, n\} \in \{0,1\}^n$ denote the set of indices where S and $\widehat{S}$ agree and let $k := |I|$ be its cardinality. By construction, I is also the set of indices where X and Y agree. Fixing orderings of I and its complement in $\{0,1\}^n$, the strings $(x_i)_{i \in I} \in \{0,1\}^k$ and $(x_i)_{i \notin I} \in \{0,1\}^{n-k}$ are independent and uniformly distributed. For $E^n(U, S)$ to not equal $E^n(U, \widehat{S})$, it is necessary

$$\left|\sum_{i \in I} (-1)^{x_i}\right| < \left|\sum_{i \notin I} (-1)^{x_i}\right|. \tag{2}$$

In essence, we have two independent simple symmetric random walks of length k and $n-k$ on $\mathbb{Z}$. We are interested in the regime where $n - k \leq d \ll k$ and claim the probability that the shorter random walk is farther from 0 than the longer random walk is small.

Since n is odd, precisely one of either k or $n-k$ is even. Consider first, the case when k is even and $n-k$ is odd. The probability that inequality 2 holds is at most

$$2 \sum_{j=1}^{\frac{n-k-1}{2}} \mathrm{Prob}\left(\sum_{i \in I} (-1)^{x_i} = 2j\right) \left[1 - 2 \sum_{\ell=1}^{j-1} \mathrm{Prob}\left(\sum_{i \notin I} (-1)^{x_i} = 2\ell + 1\right)\right]$$

$$= 2 \sum_{j=1}^{\frac{n-k-1}{2}} \mathrm{Prob}\left(\sum_{i \in I} (-1)^{x_i} = 2j\right) \left[2 \sum_{\ell=j}^{\frac{n-k-1}{2}} \mathrm{Prob}\left(\sum_{i \notin I} (-1)^{x_i} = 2\ell + 1\right)\right]. \tag{3}$$

Since the mode of $\sum_{i\in I}(-1)^{x_i}$ is 0 and

$$\text{Prob}\left(\sum_{i\in I}(-1)^{x_i}=0\right)=2^{-k}\binom{k}{k/2}=\frac{1}{\sqrt{2\pi k}}+o(1),$$

the quantity 3 is upper bounded by

$$\sqrt{\frac{2}{\pi k}}\sum_{j=1}^{\frac{n-k-1}{2}}\left[2\sum_{\ell=j}^{\frac{n-k-1}{2}}\text{Prob}\left(\sum_{i\notin I}(-1)^{x_i}=2\ell+1\right)\right]+o(1).$$

$$=\sqrt{\frac{2}{\pi k}}\sum_{\ell=1}^{\frac{n-k-1}{2}}\left(\frac{n-k+1}{2}-\ell\right)\text{Prob}\left(\sum_{i\notin I}(-1)^{x_i}=2\ell+1\right)+o(1). \tag{4}$$

Since $\frac{n-k+1}{2}+\sum_{i\notin I}(-1)^{x_i}$ follows the symmetric Binomial distribution $\text{Bin}(n-k,1/2)$, the summation in equation 4 is recognised as the absolute mean deviation

$$\mathbb{E}\left[|\text{Bin}(n-k,1/2)-\mathbb{E}\left(\text{Bin}(n-k,1/2)\right)|\right]$$

which is a constant fraction

$$\sqrt{\frac{2}{\pi}}\sqrt{\frac{n-k}{8}}+o\left(\sqrt{\frac{1}{n-k}}\right)$$

of the standard deviation $\sqrt{\frac{n-k}{8}}$ of $\text{Bin}(n-k,1/2)$ [9](see also [8]). In summary, equation 4 is upper bounded as

$$\frac{1}{\pi}\sqrt{\frac{n-k-1}{2k}}+o\left(\sqrt{\frac{n-k}{k}}\right)\leq\frac{1}{\pi}\sqrt{\frac{d}{2n}}+o\left(\sqrt{\frac{d}{n}}\right).$$

The proof translates to the remaining case (where k is odd and $n-k$ is even) *mutatis mutandis*.

Lemma 2. *Let n be an odd positive integer. Let $S,\widehat{S}\in\{0,1\}^n$ be two n-bit strings with Hamming distance bounded as $d_H(S,\widehat{S})\leq d$. For $W\in(\{0,1\}^n)^m$ drawn uniformly at random,*

$$Ext_W^{n,m}(S)=Ext_W^{n,m}(\widehat{S})$$

with probability at least

$$1-\frac{1}{\pi}\sqrt{\frac{dm^2}{2n}}+o\left(\sqrt{\frac{dm^2}{n}}\right).$$

Proof. Write $W=(W_1,W_2,\ldots,W_m)$ as presented to $\text{Ext}^{n,m}$. Since W is drawn uniformly at random from $(\{0,1\}^n)^m$, $W_1,W_2,\ldots,W_m$ are independent uniformly random strings in $\{0,1\}^n$. Therefore,

$$\text{Prob}\left(E^n(W_i,S)=E^n(W_i,\widehat{S}),\forall i\in\{1,2,\ldots,m\}\right)=\prod_{i=1}^{m}\text{Prob}\left((E^n(W_i,S)=E^n(W_i,\widehat{S})\right),$$

which, by lemma 1 is

$$\geq 1-\frac{1}{\pi}\sqrt{\frac{dm^2}{2n}}+o\left(\sqrt{\frac{dm^2}{n}}\right).$$

The extractors have a complementary assurance too, that if two strings are far in Hamming distance, then their extractions are different. Ideally, we want extractions of two strings that are very far to be independent.

Lemma 3. *Let n be an odd positive integer greater than 1. Let $S, \widehat{S} \in \{0,1\}^n$ be two n-bit strings with Hamming distance $d_H(S, \widehat{S}) = (1/2 - \zeta)n$ for some constant $\zeta \in [0, 1/100)$. For $U \in \{0,1\}^n$ drawn uniformly at random, $E^n(U, S) \neq E^n(U, \widehat{S})$ with probability at least $\frac{2}{9}$.*

Proof. We begin similar to the proof of lemma 1. Let $X = (x_1, x_2, \ldots, x_n) := U \oplus S \in \{0,1\}^n$ and $Y = (y_1, y_2, \ldots, y_n) := U \oplus \widehat{S} \in \{0,1\}^n$ respectively denote the pointwise xors. Since S is uniformly random in $\{0,1\}^n$, X and Y are (dependent but) both uniformly random in $\{0,1\}^n$. Let $I \subseteq \{1, 2, \ldots, n\} \in \{0,1\}^n$ denote the set of indices where S and $\widehat{S}$ agree and let $k := |I|$ be its cardinality. By construction, I is also the set of indices where X and Y agree. Fixing orderings of I and its complement in $\{0,1\}^n$, the strings $(x_i)_{i \in I} \in \{0,1\}^k$ and $(x_i)_{i \notin I} \in \{0,1\}^{n-k}$ are independent and uniformly distributed. If

$$\left| \sum_{i \in I} (-1)^{x_i} \right| < \left| \sum_{i \notin I} (-1)^{x_i} \right| \tag{5}$$

then $E^n(U, S) \neq E^n(U, \widehat{S})$ with probability at least half. In essence, we have two independent simple symmetric random walks of length $k = \left(\frac{1}{2} - \zeta\right) n$ and $n - k = \left(\frac{1}{2} + \zeta\right) n$ on $\mathbb{Z}$. Condition 5 identifies the scenario when the shorter random walk is farther from 0 than the longer random walk. Therefore the sign of the shorter random walk, which remains uniform conditioned on 5, determines if $E^n(U, S) \neq E^n(U, \widehat{S})$. The following are very crude Gaussian approximations. The longer random walk concentrates as

$$\text{Prob}\left(\left| \sum_{i \in I} (-1)^{x_i} \right| \leq \frac{\sqrt{n}}{2\sqrt{2}} \right) \geq \frac{2}{3} \tag{6}$$

since the standard deviation of $\sum_{i \in I} (-1)^{x_i}$ is $\sqrt{n(\frac{1}{2} + \zeta)}/2$. The shorter random walk deviates as

$$\text{Prob}\left(\left| \sum_{i \notin I} (-1)^{x_i} \right| > \frac{\sqrt{n}}{2\sqrt{2}} \right) \geq \frac{1}{3} \tag{7}$$

since the standard deviation of $\sum_{i \notin I} (-1)^{x_i}$ is $\sqrt{n(\frac{1}{2} - \zeta)}/2$. From inequalities 6,7 and the independence of $\sum_{i \in I} (-1)^{x_i}$ and $\sum_{i \notin I} (-1)^{x_i}$, condition 5 holds with probability at least $2/9$. □.

4 Signature schemes over Mersenne rings with Hamming metric

We next tailor the identification protocols and signature schemes from Ring-and-Noise assumptions to the ring $\mathbb{Z}/p\mathbb{Z}$ of residues modulo a Mersenne prime $p = 2^n - 1$ endowed with the Hamming metric. The primality constraint may be relaxed to allow for Mersenne numbers without small factors [2]. However, we insist on p being prime for ease of exposition. Therefore, n is necessarily a prime. We further assume n is an odd prime thereby excluding the smallest Mersenne prime $p = 3$. Assume there are infinitely many Mersenne primes, for otherwise our

claims for the identification protocol(while still true) are vacuous.

Abusing notation, we often identify residues modulo p with the n-bit strings of canonical representatives in $\{0, 1, \ldots, p-1\}$. In particular, for an $X \in \mathbb{Z}/p\mathbb{Z}$, $\omega_H(X)$ will denote the Hamming weight of the corresponding n-bit string. For $X, X' \in \mathbb{Z}/p\mathbb{Z}$, $d_H(X, X')$ will denote the Hamming distance between the corresponding bit strings. For an $X \in \mathbb{Z}/p\mathbb{Z}$, when we call the extractor function with X as an argument, the corresponding n-bit string is meant as the argument.

We will keep the vector notation from § 2 specialized to $\mathcal{R} = \mathbb{Z}/p\mathbb{Z}$. In particular, for a positive bound ϵ, $(\mathbb{Z}/p\mathbb{Z})_\epsilon$ will denote the residues of Hamming weight at most ϵ. The noise distribution on $(\mathbb{Z}/p\mathbb{Z})_\epsilon$, will be the uniform distribution. Random noise vectors and matrices are drawn with independent noise coordinates.

Set the weight bound $\sigma := n^{1/3}$ for the noise. The one third exponent is chosen to simplify narration and not optimized for performance. Specialized to Mersenne rings with Hamming metric, for every integer k greater than one, our identification protocol is for the language

$$\left(n, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta} \mid \mathbf{R} \in \mathcal{R}^k, \text{noise vector } \boldsymbol{\Delta} \in (\mathbb{Z}/p\mathbb{Z})_\sigma^k, A \in \mathbb{Z}/p\mathbb{Z}\right)$$

of k noisy linear congruences modulo Mersenne primes.

4.1 Identification protocol

We propose a three round commit-challenge-respond Σ-protocol with shared randomness. The instance published is

$$(n, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta}).$$

We will use the extractor function designed in the previous section, whose outputs are bit strings. Set the total number of bits extracted to $\ell := 34n$. Set $\epsilon := 1/10$ and take the challenge weight bound to be $\sigma' := n^\epsilon$. Again, these parameters are chosen to facilitate exposition and are not optimized for performance. The protocol will be sound and zero knowledge conditioned on the following assumption, which is a small weight multiplier twist on Mersenne low Hamming combination assumptions [2].

Small Multiplier Mersenne Low Hamming Combination assumption: Let $\mathbf{S} \in (\mathbb{Z}/p\mathbb{Z})^k$ be a uniform secret vector. Despite access to arbitrarily many samples, $(\boldsymbol{\Theta}, \langle \boldsymbol{\Theta}, \mathbf{S} \rangle + \Gamma)$ drawn with a uniform noise vector $\boldsymbol{\Theta} \in (\mathbb{Z}/p\mathbb{Z})_\sigma^k$ and uniform noise $\Gamma \in (\mathbb{Z}/p\mathbb{Z})_\sigma$ is computationally indistinguishable from the the uniform distribution over $(\mathbb{Z}/p\mathbb{Z})_\sigma^k \times (\mathbb{Z}/p\mathbb{Z})$.

As public randomness, draw a uniform ℓ by k noise matrix $\boldsymbol{\Lambda} \in (\mathbb{Z}/p\mathbb{Z})_\sigma^{\ell \times k}$ and a uniform $\mathbf{W} \in (\mathbb{Z}/p\mathbb{Z})^\ell$.

- Prover: Independently draw uniform noise vectors $\boldsymbol{\Phi} \in (\mathbb{Z}/p\mathbb{Z})_\sigma^k, \boldsymbol{\Psi} \in (\mathbb{Z}/p\mathbb{Z})_\sigma^\ell$ and a uniform $Y \in \mathbb{Z}/p\mathbb{Z}$. Using a commitment protocol, commit to $E := \mathrm{Ext}_{\mathbf{W}}^\ell\left(\boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right)$.
- Verifier: Send a uniform challenge $\Gamma \in (\mathbb{Z}/p\mathbb{Z})_{\sigma'}$.

– Prover: Send the response $Y + A\Gamma$ to the challenge and open the commitment to reveal E.

– Verifier: Reject if the revealed E is inconsistent with its commitment. Draw a uniform noise vector $\boldsymbol{\Psi}' \in (\mathbb{Z}/p\mathbb{Z})^{\ell}_{\sigma}$ and reject if $E' := \mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\boldsymbol{\Lambda}\left((Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \boldsymbol{\Delta})\right) + \boldsymbol{\Psi}'\right)$ is far from E. Else, accept.

We next argue for completeness and soundness; taking into consideration the weight bounds σ, σ' and the distribution guarantees of the extractors from § 3. The proof of zero knowledge is omitted, as it translate readily from the general case to the Mersenne case, without much regard to the parameters chosen.

Completeness. Consider an instance (description of $\mathcal{R}, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta}$) that belongs to the language in question and an honest prover with knowledge of a witness A of membership. To claim completeness, the verifier must accept with a constant probability. Assuming compliance to the protocol, the prover's honesty ensures passing of the commitment protocol. By construction,

$$(Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \boldsymbol{\Delta}) = Y\mathbf{R} - \Gamma\boldsymbol{\Delta}$$

$$\Rightarrow \boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) - \boldsymbol{\Lambda}((Y + \Gamma A)\mathbf{R} - \Gamma\mathbf{R}') = \boldsymbol{\Lambda}(\boldsymbol{\Phi} - \Gamma\boldsymbol{\Delta}).$$

Theorem 3 in [2] together with the triangle inequality implies

$$\omega_H(\boldsymbol{\Lambda}(\boldsymbol{\Phi} - \Gamma\boldsymbol{\Delta})) = n^{2/3+o(1)}$$

by the choice of weight thresholds σ, σ', uniformity of $\boldsymbol{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi})$ and k being a constant. By lemma 1, $E \oplus E'$ is a Bernoulli distribution with $\ell = 34n$ trials and probability of 1 in each trial bounded by $n^{-1/3+o(1)}$. Therefore with probability close to one, the prover passes the last step of verification.

Soundness. Consider an instance $(n, \mathbf{R}, \mathbf{R}')$ and a prover $\mathcal{P}$ who convinces an honest verifier $\mathcal{V}$ with probability greater than half. We claim that the instance must belong to the language. That is, there is an $A \in \mathcal{R}$ such that $\mathbf{R}' - A\mathbf{R} \in \mathbf{R}^{k}_{\sigma}$. To this end, first challenge $\mathcal{P}$ with a Γ and proceed with the protocol to successful completion. Let Z denote the response of $\mathcal{P}$ in the final step, in place of $Y + A\Gamma$. Rewind the protocol up to the instant after $\mathcal{P}$ commits to E. Now challenge the prover with a new $\widehat{\Gamma}$ distinct from Γ and proceed with the protocol to successful completion. Let $\widehat{Z}$ denote the response of $\mathcal{P}$ in the final step, in place of $Y + A\widehat{\Gamma}$.

Since both runs of the protocol succeeded and were tied to the same E,

$$\mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\boldsymbol{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}') + \boldsymbol{\Psi}'\right) \text{ and } \mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\boldsymbol{\Lambda}(\widehat{Z}\mathbf{R} - \widehat{\Gamma}\mathbf{R}') + \boldsymbol{\Psi}'\right) \tag{8}$$

agree at a great fraction of the ℓ coordinates. Consider the maps

$$\mathbb{Z}/p\mathbb{Z} \times (\mathbb{Z}/p\mathbb{Z})_{\sigma'} \xrightarrow{\quad f_{\mathbf{R},\boldsymbol{\Lambda}} \quad} (\mathbb{Z}/p\mathbb{Z})^{\ell} \xrightarrow{\quad g_{\mathbf{W}} \quad} \{0,1\}^{\ell}$$

$$(X, \Omega) \longmapsto \boldsymbol{\Lambda}(X\mathbf{R} - \Omega\mathbf{R}') + \boldsymbol{\Psi}' \longmapsto \mathrm{Ext}^{\ell}_{\mathbf{W}}\left(\boldsymbol{\Lambda}(X\mathbf{R} - \Omega\mathbf{R}') + \boldsymbol{\Psi}'\right)$$

parametrised by the $\mathbf{R}$ part of the instance and the shared public randomness $\boldsymbol{\Lambda}, \mathbf{W}$. The prover found two distinct codewords in the code $g_{\mathbf{W}}(f_{\mathbf{R},\boldsymbol{\Lambda}}((\mathbb{Z}/p\mathbb{Z}) \times (\mathbb{Z}/p\mathbb{Z})_{\sigma'})) \subset \{0,1\}^{\ell}$ that are very

close. Namely $g_{\mathbf{W}}(f_{\mathbf{R},\mathbf{\Lambda}}(Z,\Gamma))$ and $g_{\mathbf{W}}(f_{\mathbf{R},\mathbf{\Lambda}}(\widehat{Z},\widehat{\Gamma}))$.

Assume there is no $\mathbb{Z}/p\mathbb{Z}$-linear equation of the form $\mathbf{R}' - A\mathbf{R} \in (\mathbb{Z}/p\mathbb{Z})^k_\sigma$. By the Small Multiplier Mersenne Low Hamming Combination assumption, for each choice of $(X,\Omega) \in \mathbb{Z}/p\mathbb{Z} \times (\mathbb{Z}/p\mathbb{Z})_{\sigma'}$ that is not a pair of zeroes, $\mathbf{\Lambda}(X\mathbf{R} - \Omega\mathbf{R}') + \boldsymbol{\Psi}'$ is indistinguishable from uniform in $(\mathbb{Z}/p\mathbb{Z})^\ell$. Therefore, the code $f_{\mathbf{R},\mathbf{\Lambda}}(\mathbb{Z}/p\mathbb{Z} \times (\mathbb{Z}/p\mathbb{Z})_{\sigma'})$ is indistinguishable from a random subset of the same size. The size of the code $f_{\mathbf{R},\mathbf{\Lambda}}(\mathbb{Z}/p\mathbb{Z} \times (\mathbb{Z}/p\mathbb{Z})_{\sigma'})$ is small (close to $p^{11/10} \approx 2^{11n/10}$); therefore it has a large minimum distance. In particular, by the Gilbert-Varshamov bound, the relative distance is at least $1/2 - \delta$, for a small $\delta \in (0, 1/100]$.

The extractor map $g_{\mathbf{W}}$ is locality sensitive. Finding a collision in $(\mathbb{Z}/p\mathbb{Z})\ell$ under $g_{\mathbf{W}}$ is hence easy. By lemma 2, merely picking two vectors that differ by a noise vector in $(\mathbb{Z}/p\mathbb{Z})^\ell_\sigma$ suffices. But the prover found something stronger; a collision in the code $f_{\mathbf{R},\mathbf{\Lambda}}(\mathbb{Z}/p\mathbb{Z} \times (\mathbb{Z}/p\mathbb{Z})_{\sigma'})$ under $g_{\mathbf{W}}$. In the ensuing paragraph, we present a union bound argument showing that there are likely no collisions in a similarly small random subset of $(\mathbb{Z}/p\mathbb{Z})^\ell$ under $g_{\mathbf{W}}$. If fact, the union bound argument will only use the fact that such a small random subset has relative distance at least $1/2 - \delta$. By finding a collision, the prover distinguishes the code from a random subset; a contradiction. Therefore our assumption was wrong. There is indeed an $\mathbb{Z}/p\mathbb{Z}$-linear equation of the form $\mathbf{R}' - A\mathbf{R} \in (\mathbb{Z}/p\mathbb{Z})^k_\sigma$; meaning our instance is in the language and the protocol is sound.

Let $\delta \in (0, 1/100]$. By lemma 5, for two distinct $\mathbf{C}, \mathbf{C}' \in f_{\mathbf{R},\mathbf{\Lambda}}(\mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}) \subset \{0,1\}^{n\ell}$ with Hamming distance $d_H(\mathbf{C}, \mathbf{C}') > (1/2 - \delta)n\ell$, $h_{\mathbf{W}}(\mathbf{C}) \oplus h_{\mathbf{W}}(\mathbf{C}')$ is a Bernoulli distribution with probability of 1 for a trial at least $2/9$. Therefore, the probability that the encodings of $\mathbf{C}$ and $\mathbf{C}'$ under $h_{\mathbf{W}}$ collapse in relative distance is bounded by

$$\text{Prob}_{\mathbf{W}}\left(d_H(h_{\mathbf{W}}(\mathbf{C}), h_{\mathbf{W}}(\mathbf{C}')) \leq \frac{\ell}{9}\right) \leq e^{-\ell D(\frac{1}{9}||\frac{2}{9})} \approx e^{-0.04167\ \ell} \approx 2^{-0.06\ \ell}$$

where $D(\frac{1}{9}||\frac{2}{9})$ is the Kullback-Leibler divergence between Bernoulli distributions with probabilities $1/9$ and $2/9$. By the union bound,

$$\text{Prob}_{\mathbf{W}}\left(\exists \text{ distinct } \mathbf{C}, \mathbf{C}' \in f_{\mathbf{R},\mathbf{\Lambda}}(\mathbb{Z}/p\mathbb{Z}) \mid d_H(h_{\mathbf{W}}(\mathbf{C}), h_{\mathbf{W}}(\mathbf{C}')) \leq \frac{\ell}{9}\right) \leq \binom{2^{11n/10}}{2} 2^{-0.06\ \ell}$$

which is exponentially small for $\ell \geq 34n$. This completes the proof of membership.

When the instance is in the language, the collision the prover finds differs by a noise vector

$$\mathbf{\Lambda}(Z\mathbf{R} - \Gamma\mathbf{R}') - \mathbf{\Lambda}(\widehat{Z}\mathbf{R} - \widehat{\Gamma}\mathbf{R}') \in (\mathbb{Z}/p\mathbb{Z})^\ell_\sigma.$$

From the smallness of the coefficients of Λ and the randomness (which ensures that the linear system Λ defines has sufficient rank), we can conclude that

$$Z\mathbf{R} - \Gamma\mathbf{R}' - (\widehat{Z}\mathbf{R} - \widehat{\Gamma}\mathbf{R}') \in (\mathbb{Z}/p\mathbb{Z})^k_\sigma \Rightarrow \mathbf{R}' - (\widehat{\Gamma} - \Gamma)^{-1}(Z - \widehat{Z})\mathbf{R} \in (\mathbb{Z}/p\mathbb{Z})^k_\sigma.$$

Therefore, the prover knows a witness $(\widehat{\Gamma} - \Gamma)^{-1}(Z - \widehat{Z})$.

Signature scheme. We transform the identification scheme into a signature scheme using Fiat-Shamir [17]. Fix a public cryptographic hash function H whose image is restricted to $(\mathbb{Z}/p\mathbb{Z})_{\sigma'}$.

Draw a public random matrix $\mathbf{\Lambda}$ and an extractor seed $\mathbf{W}$; just as in the identification protocol. The public verification key is the instance

$$(n, \mathbf{R}, A\mathbf{R} + \boldsymbol{\Delta}, \mathbf{\Lambda}, \mathbf{W})$$

with independently chosen uniform $\mathbf{R} \in (\mathbb{Z}/p\mathbb{Z})^k$, uniform $A \in \mathbb{Z}/p\mathbb{Z}$ and noise vector $\boldsymbol{\Delta} \in (\mathbb{Z}/p\mathbb{Z})^k_\sigma$. To reduce the key size in practice, the public randomness $\mathbf{\Lambda}, \mathbf{W}$ may be generated using a fixed public pseudorandom generator and a public seed. The public seed is then made part of the public verification key. The signer knows the private key A (and hence also knows $\boldsymbol{\Delta}$). Let M denote the message to be signed. The signer first extracts $\mathrm{Ext}^\ell_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right)$ as in the identification scheme. Without interaction, the signer then prepares the challenge Γ as the hash of the extraction concatenated with the message

$$\Gamma := H\left(\mathrm{Ext}^\ell_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right), M\right).$$

The signer appends the response $Y + A\Gamma$ and the extraction as the signature

$$\left(Y + A\Gamma, \mathrm{Ext}^\ell_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right)\right)$$

of M. The verifier checks that

$$\mathrm{Ext}^\ell_{\mathbf{W}}\left(\mathbf{\Lambda}(Y\mathbf{R} + \boldsymbol{\Phi}) + \boldsymbol{\Psi}\right) \text{ and } \mathrm{Ext}^\ell_{\mathbf{W}}\left(\mathbf{\Lambda}\left((Y + A\Gamma)\mathbf{R} - \Gamma(A\mathbf{R} + \boldsymbol{\Delta})\right)\right)$$

are close in agreement. The left hand side was revealed in the signature, which also allows the verifier to compute the challenge Γ using the hash H. The right hand side is computed from the knowledge of the challenge Γ, the response $Y + A\Gamma$ and the randomness $\mathbf{\Lambda}$. The signature scheme is unforgeable conditioned on the Small Multiplier Low Hamming Combination assumption in the random oracle model.

References

1. 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 Lars R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Amsterdam, The Netherlands, April 28 - May 2, 2002, Proceedings*, volume 2332 of *Lecture Notes in Computer Science*, pages 418–433. Springer, 2002.
2. D. Aggarwal, A. Joux, A. Prakash, and M. Santha. A new public-key cryptosystem via mersenne numbers. *CRYPTO 2018: Advances in Cryptology*, pages 459–482, 2018.
3. Sidi Mohamed El Yousfi Alaoui, Pierre-Louis Cayrel, Rachid El Bansarkhani, and Gerhard Hoffmann. Code-based identification and signature schemes in software. In Alfredo Cuzzocrea, Christian Kittl, Dimitris E. Simos, Edgar R. Weippl, and Lida Xu, editors, *Security Engineering and Intelligence Informatics - CD-ARES 2013 Workshops: MoCrySEn and SeCIHD, Regensburg, Germany, September 2-6, 2013. Proceedings*, volume 8128 of *Lecture Notes in Computer Science*, pages 122–136. Springer, 2013.
4. Nicolas Aragon, Olivier Blazy, Philippe Gaborit, Adrien Hauteville, and Gilles Zémor. Durandal: A rank metric based signature scheme. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 728–758. Springer, 2019.

5. Shi Bai, Dipayan Das, Ryo Hiromasa, Miruna Rosca, Amin Sakzad, Damien Stehlé, Ron Steinfeld, and Zhenfei Zhang. Mpsign: A signature from small-secret middle-product learning with errors. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part II*, volume 12111 of *Lecture Notes in Computer Science*, pages 66–93. Springer, 2020.
6. Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *Topics in Cryptology - CT-RSA 2014 - The Cryptographer's Track at the RSA Conference 2014, San Francisco, CA, USA, February 25-28, 2014. Proceedings*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47. Springer, 2014.
7. Rouzbeh Behnia, Yilei Chen, and Daniel Masny. On removing rejection conditions in practical lattice-based signatures. In Jung Hee Cheon and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography - 12th International Workshop, PQCrypto 2021, Daejeon, South Korea, July 20-22, 2021, Proceedings*, volume 12841 of *Lecture Notes in Computer Science*, pages 380–398. Springer, 2021.
8. D Berend and A Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.
9. C. R. Blyth. Expected absolute error of the usual estimator of the binomial parameter. *The American Statistician*, 34(3):155–157, 1980.
10. Colin Boyd. Digital multisignatures. *Cryptography and coding*, pages 241–246, 1986.
11. David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology: Proceedings of CRYPTO '82, Santa Barbara, California, USA, August 23-25, 1982*, pages 199–203. Plenum Press, New York, 1982.
12. Yvo Desmedt. Society and group oriented cryptography: A new concept. In Carl Pomerance, editor, *Advances in Cryptology - CRYPTO '87, A Conference on the Theory and Applications of Cryptographic Techniques, Santa Barbara, California, USA, August 16-20, 1987, Proceedings*, volume 293 of *Lecture Notes in Computer Science*, pages 120–127. Springer, 1987.
13. Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In Gilles Brassard, editor, *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 307–315. Springer, 1989.
14. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 40–56. Springer, 2013.
15. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-Dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.
16. Thibauld Feneuil, Antoine Joux, and Matthieu Rivain. Shared permutation for syndrome decoding: New zero-knowledge protocol and code-based signature. *IACR Cryptol. ePrint Arch.*, page 1576, 2021.
17. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology - CRYPTO '86, Santa Barbara, California, USA, 1986, Proceedings*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194. Springer, 1986.
18. 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.
19. P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613, 1998.
20. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Zero-knowledge from secure multiparty computation. In David S. Johnson and Uriel Feige, editors, *39th Annual ACM Symposium on Theory of Computing*, pages 21–30, San Diego, CA, USA, June 11–13, 2007. ACM Press.

21. Antoine Joux. NIST workshop on Post-quantum cryptography. `https://csrc.nist.gov/CSRC/media/Presentations/Mersenne-756839/images-media/Mersenne-April2018.pdf`.
22. San Ling, Khoa Nguyen, Damien Stehlé, and Huaxiong Wang. Improved zero-knowledge proofs of knowledge for the ISIS problem, and applications. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *Public-Key Cryptography - PKC 2013 - 16th International Conference on Practice and Theory in Public-Key Cryptography, Nara, Japan, February 26 - March 1, 2013. Proceedings*, volume 7778 of *Lecture Notes in Computer Science*, pages 107–124. Springer, 2013.
23. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 598–616. Springer, 2009.
24. Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.
25. Robert J McEliece. A public-key cryptosystem based on algebraic. *Coding Thv*, 4244:114–116, 1978.
26. Carlos Aguilar Melchor, Philippe Gaborit, and Julien Schrek. A new zero-knowledge code based identification scheme with reduced communication. In *2011 IEEE Information Theory Workshop, ITW 2011, Paraty, Brazil, October 16-20, 2011*, pages 648–652. IEEE, 2011.
27. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *J. Cryptol.*, 13(3):361–396, 2000.
28. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press.
29. Jacques Stern. A new identification scheme based on syndrome decoding. In Douglas R. Stinson, editor, *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 13–21. Springer, 1993.
30. Pascal Véron. Improved identification schemes based on error-correcting codes. *Appl. Algebra Eng. Commun. Comput.*, 8(1):57–69, 1996.