Improved Short Lattice Signatures in the Standard Model

Léo Ducas and Daniele Micciancio

University of California, San Diego
{lducas,daniele}@eng.ucsd.edu

Abstract. We present a signature scheme provably secure in the standard model (no random oracles) based on the worst-case complexity of approximating the Shortest Vector Problem in ideal lattices within polynomial factors. The distinguishing feature of our scheme is that it achieves *short* signatures (consisting of a single lattice vector), and *relatively short* public keys (consisting of $O(\log n)$ vectors.) Previous lattice schemes in the standard model with similarly *short* signatures, due to Boyen (PKC 2010) and Micciancio and Peikert (Eurocrypt 2012), had substantially longer public keys consisting of $\Omega(n)$ vectors (even when implemented with ideal lattices).

1 Introduction

Lattice based cryptography [3,4], originally an area of primarily theoretical interest, has seen a tremendous growth during the last decade, due both to substantial efficiency improvements obtainable using lattices with algebraic structure [16,28], and to the enormous versatility afforded by the Learning with Errors (LWE) problem [33]. One of the problems that has received most attention so far, is that of lattice based signatures [24,13,21,9,35,14,22,12,6]. From a theoretical point of view, digital signatures can be constructed from any one-way function [34,19]. So, the existence of digital signature schemes based on the hardness of lattice problems directly follows from Ajtai's seminal work [3]. But generic constructions are rather inefficient. Inputs and outputs of lattice based cryptographic functions typically consist of one or more $\tilde{\Omega}(n)$-dimensional vectors, where n is the security parameter. Generic digital signature constructions require n parallel applications of a one-way function. So, even if each one-way function takes as input a single vector, the resulting digital siguatures consist of n vectors, and require $\tilde{\Omega}(n^2)$ storage even when using algebraic lattices [28]. So, finding efficient constructions of signatures directly based on hard lattice problems has been an important problem since the early days of lattice cryptography, with the main goal of finding "short" signatures, i.e., lattice signatures consisting of a single lattice vector.

The first direct constructions of lattice signatures were given in [24] and [13]. Both schemes achieved "short" signatures, consisting of a single lattice vector, but each work had its own pros and cons. On the one hand [24] gave a

scheme provably secure in the standard model of computation, and with very simple signing/verification procedures, but only provided a direct construction of one-time signatures: digital signature schemes that can be used to sign a single message. Such schemes can be turned into general purpose signature schemes with only a logarithmic loss in efficiency using standard tree constructions. However, these transformations can be quite expensive in practice, because they lead to signatures consisting of $O(\log n)$ vectors. Given that signature size is often the most critical efficiency parameter affecting the practicality of a scheme, such signatures can no longer be considered "short". On the other hand, [13] gave a scheme that allowed to produce short signatures for arbitrarily many messages, but only offered heuristic security in the random oracle model. Moreover, the scheme of [13] was not entirely practical, involving a rather complex signing algorithm based on sampling lattice vectors with gaussian distribution, a problem that only recently has found more satisfactory solutions [29].

Two lines of research have evolved from [13], trying to address either the security or efficiency limitations of that work:

- A first line of research [21,22,14,12,6,15] kept investigating lattice signature in the random oracle model, with the goal of achieving the highest possible levels of performance, and schemes that are efficient enough to be used in practice.
- A second line of work, [11,9,29] kept pursuing the important goal of obaining security in the standard model of computation (no random oracles) while at the same time improving the efficiency and potential practicality of previous schemes. Our work is part of this second line of research, which we describe in more detail.

The current state of the art, when it comes to short lattice signatures in the standard model, is given by the scheme of Boyen [9], with additional security and efficiency improvements described in [29]. This scheme achieved the important goal of "short" lattice signatures (consisting of a single lattice vector), without resorting to the random oracle model. The main drawback of this scheme was the huge public key involved. Lattice public keys, even in the random oracle model [13,21,22,14,12,6], consist of one or more $n \times m$ matrices, each of which typically requires $\tilde{\Omega}(n^2)$ storage. For the sake of comparison, we consider natural adaptations of [11,9,29] to the algebraic/ring setting, where $n \times m$ matrices can be implicitly described by a single m-dimensional vector. Going back to the signature scheme of [9,29], public keys consist of $\Omega(n)$ matrices, and therefore require at least quadratic $\tilde{\Omega}(n^2)$ total storage even when using "compact" algebraic lattices. We remark that digital signature schemes can be efficiently constructed out of identity based encryption (IBE) by using ciphertexts as signatures, and lattice based IBE with short ciphertexts are also known [11,2,1]. However, lattice IBE schemes are built on top of the signature techniques from [11,9], and bear the same limitations when it comes to public key size: lattice IBE [11,2,1] use public keys consisting of $\Omega(n)$ matrices, and result in $\tilde{\Omega}(n^2)$ or even $\tilde{\Omega}(n^3)$ pubic key size depending on the type of lattices employed.

Reducing the size of, not only the signatures, but also the public key, was the main open problem left by [11,9,29,2,1]. We remark that the last few years have seen major efficiency progress on lattice signatures in the random oracle model [13,21,22,14,12,6], leaving a substantial gap between random oracle and standard model signatures. Still, designing efficient signature schemes without random oracles is an important and well established problem, both for the theory and practice of cryptography. A recent work in this direction is the paper of Bohl *et al.* [7,8,36], which formalized[1] a general "confined guessing" technique applicable to a variety of (not only lattice) settings. Here we describe their results, limited to the case of lattice signatures, and specialized to algebraic/ring lattices. Among other things, [7] gives a standard model lattice signature with public keys consisting of a single matrix, and therefore requiring only $O(m) = \tilde{\Omega}(n)$ storage when using algebraic/ring lattices. However, this comes at a substantial cost in terms of signature size: the digital signatures of [7] consist of $O(\log n)$ vectors. While a $O(\log n)$ increase may not seem much, it is quite a high cost when it comes to signature size, both in theory and in practice. In fact, a similar trade-off was already known since the very first direct construction of lattice signatures [24], which, as alredy discussed, produced general signatures consisting of $O(\log n)$ vectors (as well as short public keys). In other words, just like [24], the lattice signatures of [7] are not "short". (The main contribution of [7] over the classic scheme of [24], is that the results of [7] also apply to general lattices.)

Our results. We present the first standard model construction of short signatures based on (algebraic/ring) lattices with relatively small public keys: Similarly to [9,29], we achieve signatures consisting of a single vector without resorting to random oracles. At the same time, we substantially reduce the public key size from the $\Omega(n)$ vectors[2] of previously best short lattice signatures [9,29] to just $O(\log n)$ vectors. Our scheme is stateless, i.e., all signatures can be produced independently by running the signing algorithm on input the secret key and message to be signed. We also give an even more efficient scheme that further improves the public key size from $O(\log n)$ to just $O(\log \log n)$ vectors (and at the same time also improves the tightness of the reduction,) almost matching the *asymptotic* performance of schemes in the random oracle model [13,21,22,14,12,6]. This last improvement comes at the cost of statefulness: the signer has to keep some state information between signatures. However the state information is extremely simple: all that the signer has to do is to maintain a counter keeping track of how many signatures have already been produced.

We remark that it is always possible to reduce the public key size by increasing the size of the signatures, simply by compressing the public key using a collision resistant hash function (which is easily built from lattices [26,5,23,31]),

[1] The technique first appeared in the work of Hohenberger and Waters [18,17] and was also used in [10].

[2] Remember we are in the ring setting, so only one vector is required to represent each matrix.

Scheme	Pub. Key $\mathcal{R}_q^{1\times k}$ mat.	Secret Key $\mathcal{R}_q^{k\times k}$ mat.	Signature $\mathcal{R}_q^k$ vec.	Reduction loss	SIS parameter β
[13](ROM)	1	1	1	1	$\tilde{\Omega}(n)$
[24](Trees)	1	1	$\log n$	Q	$\tilde{\Omega}(n^2)$
[11]	n	n	n	Q	$\tilde{\Omega}(n^{3/2})$
[9,29]	n	n	1	Q	$\tilde{\Omega}(n^{7/2})$, $\tilde{\Omega}(n^{5/2})$
[7]	1	1	$\log_c n$	$O(Q^2/\epsilon)^c$	$\tilde{\Omega}(n^{5/2})$
Stateless (Sec. 3)	$\log_c n$	$\log_c n$	1	$O(Q^2/\epsilon)^c$	$\tilde{\Omega}(n^{7/2})$
Stateful[3]	$2\log_c(\log n)$	$2\log_c(\log n)$	1	$2Q^c$	$\tilde{\Omega}(n^{3/2})$

$\mathcal{R}_q = \mathbb{Z}_q[X]/f(X)$ for some (cyclotomic) polynomial f of degree n, $q = n^{O(1)}$, and $k = O(\log q)$. Q denotes the number of signature queries made by the attacker and ϵ is its success probability. The value $c > 1$ is an arbitrary constant that governs the security/efficiency trade off. The reduction loss is the ratio ϵ'/ϵ between the success probability ϵ' of the reduction and the success probability ϵ of the attacker.

Fig. 1. Comparison to previous work on lattice signatures in the ring setting.

and including the original public key in each signature. So, our first scheme (with $O(\log n)$ vectors in the public key and short signatures) subsumes the results of [7] in the algebraic/ring lattice setting with $O(\log n)$ vectors per signatures.

The efficiency of our lattice constructions, compared to previous schemes (all adapted to the ring setting), is detailed in Figure 1.The trick leading to our stateful signature scheme can also be applied to improve the generic construction of [7]. The description of our generic results is deferred to the full version of our paper.

Techniques. Our results are obtained by combining several techniques previously used in the construction of lattice-based signatures. Most notably, we use the "vanishing trapdoor" technique from [9], and the more recent "confined guessing" method of [7,18,17]. In fact, the key generation, signing and verification algorithms bear strong similarities with previously proposed schemes. However, the combination appears to be novel and nontrivial. In particular, while both the results in [9] and those in [7] are presented for general lattices, the way they are combined in our work makes essential use of the commutativity properties of ring/algebraic lattices. More specifically, our proof of security exploits a key homomorphic property of lattice trapdoors (see Lemma 6) which requires certain matrix products to commute. This is trivially verified in the ring setting, where one of the matrices corresponds to a ring scalar, but glamorously fails when the construction is adapted to arbitrary lattices.

Open problems. Interestingly, the methods employed in this paper to obtain short lattice signatures with small public key seem specific to the ring/algebraic lattice setting. Only our generic result (see the full version of this article) with signatures of $\log\log n$ many vectors applies to arbitrary lattices. We remark

[3] See full version of this article.

that the question of reducing the public key size is mostly important in the ring setting: when using general lattices, even a single matrix takes quadratic storage, so there is little hope to reduce the public key size to linear or quasilinear in the security parameter. Still, it would be nice to achieve results similar to those in our paper, but for general lattices: is there a standard model signature scheme based on general lattices with short signatures (consisting of a single vector) and small public keys (consisting of $O(\log n)$ matrices)?

Another important open problem is to further improve the efficiency of our scheme, and obtain short signatures where the public key is just $O(1)$ matrices (or vectors, in the ring setting). Indeed, schemes offering both short public key and short signatures[4] in the standard model have been constructed based on the Computational Diffie-Hellman (CDH) and RSA problems [18,17].

2 Preliminaries

2.1 Signatures

Definition 1. *A signature scheme* SS *is a triple* $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verif})$ *of PPT (probabilistic polynomial time) algorithms, together with message spaces* $\mathcal{M}_n$*. It is correct if, for all messages* $\mu \in \mathcal{M}_n$*,* $\mathsf{Verif}(pk, \mu, \sigma) = 1$ *holds true, except with negligible probability (in* n*) over the choice of* $(sk, pk) \leftarrow \mathsf{KeyGen}(1^n)$ *and* $\sigma \leftarrow \mathsf{Sign}(sk, \mu)$*.*

The standard definitions of security for digital signature schemes (under adaptive and non-adaptive attacks) is given in Figure 2.

EUF-naCMA$_{\mathsf{SS}}(n, \mathcal{A})$

$\mathcal{A}$ chooses q messages $(\mu^{(j)}) \in \mathcal{M}_n$
$(sk, pk) \leftarrow \mathsf{KeyGen}(1^n)$
For all $j = 0 \dots Q-1$:
  $\sigma^{(j)} \leftarrow \mathsf{Sign}(sk, \mu^{(j)})$.
$\mathcal{A}$ receives $pk, \sigma^{(0)} \dots \sigma^{(Q-1)}$.
$\mathcal{A}$ sends an attempted forgery $(\mu^\diamond, \sigma^\diamond)$
$\mathcal{A}$ wins if $\mathsf{Verif}(pk, \mu^\diamond, \sigma^\diamond) = 1$ and $\mu^\diamond \notin \{\mu^{(j)}\}$.

EUF-CMA$_{\mathsf{SS}}(n, \mathcal{A})$

$(sk, pk) \leftarrow \mathsf{KeyGen}(1^n)$, $\mathcal{A}$ receives pk
For $j = 0 \dots Q-1$:
  $\mathcal{A}$ chooses $\mu^{(j)}$
  $\mathcal{A}$ receives $\sigma^{(j)} \leftarrow \mathsf{Sign}(sk, \mu^{(j)})$
$\mathcal{A}$ sends an attempted forgery $(\mu^\diamond, \sigma^\diamond)$
$\mathcal{A}$ wins if $\mathsf{Verif}(pk, \mu^\diamond, \sigma^\diamond) = 1$ and $\mu^\diamond \notin \{\mu^{(j)}\}$.

A signature scheme $\mathsf{SS} = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verif})$ is **EUF-naCMA**-secure (or Existentially Unforgeable under non-adaptative Chosen Message Attacks) if no PPT adversary $\mathcal{A}$ wins the **EUF-naCMA$_{\mathsf{SS}}$** game (left) with non-negligible probability $n^{-O(1)}$. The scheme is **EUF-CMA**-secure (or Existentially Unforgeable under adaptative Chosen Message Attacks) if no PPT adversary $\mathcal{A}$ wins the **EUF-CMA$_{\mathsf{SS}}$** game (right) with non-negligible probability $n^{-O(1)}$.

Fig. 2. Definition of security for digital signature schemes.

[4] Here by "short" we mean consisting of $O(1)$ group elements.

From Non-Adaptive to Full Security There are two standard techniques to transform non adaptively-secure signature schemes to fully secure ones: Chameleon Hashing and One Time Signatures both of which can be implemented using lattices [25,13]. For a description of the solution based on Chameleon Hashing see the full version of this article.

2.2 Lattices and Gaussian Distributions

A (full rank) n-dimensional *lattice* is the set $\Lambda = \mathcal{L}(\mathbf{B}) = \{\mathbf{Bz} \colon \mathbf{z} \in \mathbb{Z}^n\}$ of all integer linear combinations of n basis vectors $\mathbf{B} = [\mathbf{b}_1, \ldots, \mathbf{b}_n] \in \mathbb{R}^{n \times n}$. We use notation $(x_1, \ldots, x_n)$ for *column* vectors, and similarly write $(\mathbf{A}, \mathbf{B})$ for the result of vertically stacking two matrices. The dual lattice Λ^* is the set of all $\mathbf{v} \in \mathbb{R}^n$ such that $\langle \mathbf{v}, \mathbf{x} \rangle \in \mathbb{Z}$ for every $\mathbf{x} \in \Lambda$. If $\mathbf{B}$ is a basis of Λ, then $\mathbf{B}^* = \mathbf{B}^{-t}$ is a basis of Λ^*. Many cryptographic applications use a particular family of so-called *q-ary* integer lattices, which contain $q\mathbb{Z}^m$ as a sublattice for some (typically small) integer q. For positive integers n, and q, let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be arbitrary and define the following full-rank m-dimensional q-ary lattices:

$$\Lambda^{\perp}(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{Az} = \mathbf{0} \bmod q\}$$
$$\Lambda(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m : \exists\, \mathbf{s} \in \mathbb{Z}_q^n \text{ s.t. } \mathbf{z} = \mathbf{A}^t\mathbf{s} \bmod q\}.$$

It is easy to check that $\Lambda^{\perp}(\mathbf{A})$ and $\Lambda(\mathbf{A})$ are dual lattices, up to a q scaling factor: $q \cdot \Lambda^{\perp}(\mathbf{A})^* = \Lambda(\mathbf{A})$, and vice-versa. For any $\mathbf{u} \in \mathbb{Z}_q^n$ admitting an integral solution to $\mathbf{Ax} = \mathbf{u} \bmod q$, define the coset (or "shifted" lattice) $\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}) = \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{Az} = \mathbf{u} \bmod q\} = \Lambda^{\perp}(\mathbf{A}) + \mathbf{x}$. In the Small Integer Solution problem ($\mathrm{SIS}_{p,n,m,\beta}$), one is given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and is asked to find a nonzero vector $\mathbf{s} \in \Lambda^{\perp}(\mathbf{A})$ such that $\|\mathbf{s}\| \leq \beta$ where $\|\mathbf{s}\| = \sqrt{\sum_i s_i^2}$ is the euclidean norm. The geometric quality of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is measured by its spectral norm $s_1(\mathbf{A}) = \sup_{\mathbf{x}} \|\mathbf{Ax}\|/\|\mathbf{x}\|$.

The n-dimensional Gaussian function $\rho_s \colon \mathbb{R}^n \to (0, 1]$ is defined as $\rho_s(\mathbf{x}) = \exp(-\pi \cdot \|\mathbf{x}/s\|^2)$. For any (countable) set $X \subseteq \mathbb{R}^n$, let $\rho_s(X) = \sum_{\mathbf{x} \in X} \rho_s(\mathbf{x})$. The *smoothing parameter* of a lattice $\eta_\epsilon(\Lambda)$ [30] is the smallest s such that $\rho_{1/s}(\Lambda^*) \leq 1 + \epsilon$. The discrete gaussian distribution $D_{\Lambda,s}$ over a lattice Λ is defined as $D_{\Lambda,s}(\mathbf{x}) = \rho_s(\mathbf{x})/\rho_s(\Lambda)$ for all $\mathbf{x} \in \Lambda$.

We say that a random variable X over $\mathbb{R}$ is *subgaussian* with parameter $s > 0$ if for all $t \in \mathbb{R}$, the (scaled) moment-generating function satisfies $\mathrm{E}[\exp(2\pi t X)] \leq \exp(\pi s^2 t^2)$. If X is subgaussian, then its tails are dominated by a Gaussian of parameter s, i.e., $\Pr[|X| \geq t] \leq 2\exp(-\pi t^2/s^2)$ for all $t \geq 0$. More generally, we say that a random matrix $\mathbf{X}$ is subgaussian (of parameter s) if all its one-dimensional marginals $\mathbf{u}^t\mathbf{Xv}$ for unit vectors $\mathbf{u}, \mathbf{v}$ are subgaussian (of parameter s). It follows immediately from the definition that the concatenation of independent subgaussian vectors with common parameter s, interpreted either as a vector or as a matrix, is subgaussian with parameter s. For any lattice $\Lambda \subset \mathbb{R}^n$ and $s > 0$, the distribution $D_{\Lambda,s}$ is subgaussian with parameter s.

We will need the following standard result from the non-asymptotic theory of random matrices; for further details, see [37].

Lemma 1. *Let $\mathbf{X} \in \mathbb{R}^{n \times m}$ be a subgaussian random matrix with parameter s. There exists a universal constant $C \approx 1/\sqrt{2\pi}$ such that for any $t \geq 0$, we have $s_1(\mathbf{X}) \leq C \cdot s \cdot (\sqrt{m} + \sqrt{n} + t)$ except with probability at most $2\exp(-\pi t^2)$.*

2.3 Rings and Ideal Lattices

We consider lattice problems restricted to ideal lattices [28,23,32]. Most of our results apply to ideal/module lattices over arbitrary cyclotomic rings, but for simplicity we focus our presentation on so-called "SWIFFT" rings [26,5]. These are rings of the form $\mathcal{R} = \mathbb{Z}[X]/(\Phi_{2n}(X))$ or $\mathcal{R}_q = (\mathcal{R}/q\mathcal{R})$, where n is a power of 2, q is an integer, and $\Phi_{2n}(X) = X^n + 1$ is the cyclotomic polynomial of degree n. For our construction we will require that $\Phi_{2n}(X)$ does not split into low degree polynomials modulo the prime factors of q. More concretely we choose $q = 3^k$ and rely on the following.

Fact 1 (Irreducible factors of $\Phi_{2^k}(X)$ modulo 3. Corollary of [20, Theorem 2.47])**.** *For any $k \geq 3$ and $2n = 2^k$ we have $\Phi_{2n}(X) \equiv (X^{n/2} + X^{n/4} - 1) \cdot (X^{n/2} - X^{n/4} - 1) \bmod 3$ and both factors are irreducible in $\mathbb{F}_3[X]$.*

Lemma 2 (Hensel Lemma). *Let $\mathcal{R} = \mathbb{Z}[X]/(F(X))$ for some monic polynomial $F \in \mathbb{Z}[X]$. For any prime p, if $u \in \mathcal{R}_{p^e}$ is invertible* mod$p$ *(*i.e. it is invertible in $\mathcal{R}_p$) then u is also invertible in $\mathcal{R}_{p^e}$.*

Corollary 1. *let $n \geq 4$ be a power of 2, $q \geq 3$ a power of 3, and set $\mathcal{R}_q = \mathbb{Z}[X]/(\Phi_{2n}(X), q)$. Then, any nonzero polynomial $t \in \mathcal{R}_q$ of degree $d < n/2$ and coefficients in $\{0, \pm 1\}$ is invertible in $\mathcal{R}_q$.*

Elements in $\mathcal{R}$ have a natural representation as polynomials of degree $n-1$ with coefficients in $\mathbb{Z}$, and $\mathcal{R}$ can be identified (as an additive group) with the integer lattice $\mathbb{Z}^n$, where each ring element $\mathbf{a} = a_0 + a_1 x + \ldots + a_{n-1}x^{n-1} \in \mathcal{R}$ is associated with the coefficient vector $(a_0, \ldots, a_{n-1}) \in \mathbb{Z}^n$. We use the identification $\mathcal{R} = \mathbb{Z}^n$ to define standard lattice quantities like the euclidean length of a ring element $\|\mathbf{a}\| = \sqrt{\sum_i |a_i|^2}$, or the spectral norm of a ring element $s_1(r) = \sup_x \|r \cdot x\|/\|x\|$. The ring $\mathcal{R}$ is also identified with the sub-ring of anti-circulant square matrices of dimension n by regarding each ring element $r \in \mathcal{R}$ as a linear transformation $x \mapsto r \cdot x$ over (the coefficient embedding) of $\mathcal{R}$. Notice that the definition of spectral norm of a ring element is consistent with the definition of spectral norm of the corresponding anticirculant matrix. The following lemma provides a useful bound on the spectral norm of ring elements.

Lemma 3. *For any ring element $r \in \mathcal{R}$, we have $s_1(r) \leq \|r\|_1 = \sum_i |r_i|$.*

Proof. Let $\omega_k = e^{\pi\imath(2k-1)/n}$ (for $k = 1, \ldots, n$) be the complex roots of the cyclotomic polynomial Φ_{2n}. Consider the image of r under the canonical embedding $\sigma\colon \mathcal{R} \to \mathbb{C}^n$, which is defined as $\sigma(r) = (r(\omega_1), \ldots, r(\omega_n))$. Using the fact that $\sigma\colon \mathcal{R} \to \mathbb{C}^n$ is a ring homomorphism (with the product $\odot$ in $\mathbb{C}^n$ defined componentwise) and a scaled isometry (satisfying $\|\sigma(r)\| = \sqrt{n} \cdot \|r\|$) we get

$$s_1(r) = \sup_x \frac{\|r \cdot x\|}{\|x\|} = \sup_x \frac{\|\sigma(r \cdot x)\|}{\|\sigma(x)\|} = \sup_x \frac{\|\sigma(r) \odot \sigma(x)\|}{\|\sigma(x)\|} \leq \|\sigma(r)\|_\infty.$$

Since for any i, $|\omega_i| = 1$, we have $|r(\omega_i)| = \left|\sum_j r_j \omega_i^j\right| \leq \sum |r_j| = \|r\|_1$. It follows that $s_1(r) \leq \|\sigma(r)\|_\infty = \max_i |\sigma(r)_i| \leq \|r\|_1$.

The discrete Gaussian distribution over the ring $D_{\mathcal{R},s} \equiv D^n_{\mathbb{Z},s}$ is defined as usual by identifying the ring $\mathcal{R}$ with $\mathbb{Z}^n$ under the coefficient embedding. It follows that the discrete gaussian distribution over the ring $x \leftarrow D_{\mathcal{R},s}$ is subgaussian of parameter s when x is regarded as a vector. For the anti-circulant matrix representation, we have the following fact, (proof in App. A).

Fact 2. *If $\mathbf{R} \leftarrow D^{w\times k}_{\mathcal{R},s}$, then with overwhelming probability we have $s_1(\mathbf{R}) \leq s\sqrt{n} \cdot O(\sqrt{w} + \sqrt{k} + \omega(\sqrt{\log n}))$.*

The euclidean length of vectors in $\mathcal{R}_q^k$ is defined similarly by identifying $\mathbb{Z}_q$ with the set of representatives $\{-(q-1)/2, \ldots, +(q-1)/2\}$. Similarly, we define the q-ary lattices $\Lambda(\mathbf{A})$ and $\Lambda^\perp(\mathbf{A})$ when $\mathbf{A} \in \mathcal{R}_q^{n\times m}$ is a matrix over the ring $\mathcal{R}_q$ using the standard isomorphism of $\mathcal{R}_q$ and the sub-ring of anticirculant matrices in $\mathbb{Z}_q^{n\times n}$.

Definition 2. *In the Small Integer Solution over Rings problem* ($\mathrm{RingSIS}_{q,n,m,\beta}$)*, one is given a row vector $\mathbf{A} \in \mathcal{R}_q^{1\times m}$, and is asked to find a nonzero vector $\mathbf{x} \in \Lambda_q^\perp(\mathbf{A})$ such that $\|\mathbf{x}\| \leq \beta$.*

Let $\mathcal{U}_m$ be the uniform distribution over m-dimensional row vectors of ring elements $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \ldots, \mathbf{a}_m] \in \mathcal{R}_q^{1\times m}$. The smoothness proof from [13] can be adapted to this specific ring case (proof in App. A). A more general ring regularity result can be found [27, Theorem 7.4], but unfortunately it gives a larger bound (by a factor n) on required standard deviation s than our specialized lemma.

Lemma 4 (Smoothness Lemma). *Let $\mathcal{R}_q = \mathbb{Z}[X]/(\Phi_{2n}(X), q)$ for $n \geq 4$ a power of 2 and $q = 3^k$ a power of 3. Let $w \geq 2\lceil \log_2 q \rceil + 2$ and $s \geq \omega(\sqrt{\ln nw})$. With overwhelming probability over the choice of $\mathbf{A} \leftarrow \mathcal{U}_w$, if $\mathbf{x}_i \leftarrow D_{\mathcal{R},s}$ (for $i = 1, \ldots, w$) are chosen independently at random, then the sum $\sum_i \mathbf{a}_i \cdot \mathbf{x}_i$ is within negligible statistical distance from the uniform distribution over $\mathcal{R}$.*

A handy corollary used several time in our proof is the following.

Corollary 2 (Min-entropy bound). *Set $\mathcal{R}_q$ as above, and let $w \geq 2\lceil \log_2 q \rceil + 3$, $s \geq \omega(\sqrt{\ln nw})$. With overwhelming probability over the choice of $\mathbf{A} \leftarrow \mathcal{U}_w$, if $\mathbf{x}_i \leftarrow D_{\mathcal{R},s}$ (for $i = 1, \ldots, w$) are chosen independently at random, then for any nonzero vector $\mathbf{V} \in \mathcal{R}^w \setminus \{\mathbf{0}\}$ the conditional min-entropy of $\sum_i \mathbf{v}_i \cdot \mathbf{x}_i$ given $\sum_i \mathbf{a}_i \cdot \mathbf{x}_i$ is at least $\Omega(n)$.*

2.4 Lattice Trapdoors

We use the strong lattice trapdoor construction and algorithms of [29]. For a modulus $q = 3^k$ and integer dimension n, define the gadget matrix $\mathbf{G} = \left[\mathbf{I}_n \,|\, 3 \cdot \mathbf{I}_n \,|\, \ldots \,|\, 3^{k-1} \cdot \mathbf{I}_n\right] \in \mathbb{Z}_q^{n\times kn}$.

Definition 3. *For any* $\mathbf{A} \in \mathbb{Z}_q^{n\times(m+kn)}$*, and (invertible)* $\mathbf{H} \in \mathbb{Z}_q^{n\times n}$*, a* $\mathbf{G}$*-trapdoor for* $\mathbf{A}$ *with tag* $\mathbf{H}$ *is a matrix* $\mathbf{R} \in \mathbb{Z}_q^{m\times kn}$ *such that* $\mathbf{A}(\mathbf{R}, \mathbf{I}) = \mathbf{HG}$*. The definition is extended to trapdoors* $\mathbf{R} \in \mathbb{Z}_q^{m'\times kn}$ *with* $m' \leq m$ *by padding them with zero columns so that* $[\mathbf{R}, \mathbf{O}] \in \mathbb{Z}_q^{m\times kn}$*.*

The quality of a trapdoor $\mathbf{R}$ is measured by the spectral norm $s_1(\mathbf{R})$, and [29] gives efficient algorithms to generate uniformly random matrices $\mathbf{A}$ together with high quality trapdoors, and to sample cosets $\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A})$ with Gaussian distribution D_s for sufficiently large s. Notice that the tag $\mathbf{H}$ can immediately be recovered from $\mathbf{A}$ and $\mathbf{R}$ as the first block of $\mathbf{HG}$, and does not need to be specified explicitly. But when one says that $\mathbf{R}$ is a trapdoor, it is usually assumed that the associated tag $\mathbf{H}$ is an invertible matrix.

Theorem 3 ([29]). *There is an efficient algorithm* $\mathsf{SampleD}(\mathbf{A}, \mathbf{u}, \mathbf{R}, s)$ *that on input a matrix* $\mathbf{A} \in \mathbb{Z}_q^{n\times(m+kn)}$*, a syndrome* $\mathbf{u} \in \mathbb{Z}_q^n$*, a* $\mathbf{G}$*-trapdoor* $\mathbf{R} \in \mathbb{Z}_q^{m\times kn}$ *for* $\mathbf{A}$*, and parameter* $s > \omega(\sqrt{\log n}) \cdot s_1(\mathbf{R})$*, produces a sample from the distribution* $D_{\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}),s}$*.*

The efficient trapdoor generation algorithm of [29] follows immediately from the definition of $\mathbf{G}$-trapdoor: one simply chooses $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$ uniformly at random, samples a trapdoor matrix $\mathbf{R} \in \mathbb{Z}_q^{m\times nk}$ with small entries, and outputs $\mathbf{A} = [\mathbf{A}', \mathbf{HG} - \mathbf{A}'\mathbf{R}]$. As pointed out in [29], the algorithm is immediately adapted to ideal lattices, using the observation that the identity matrix $\mathbf{I}_n$ is precisely the matrix corresponding to the ring element $1 \in \mathcal{R}$, so the gadget matrix $\mathbf{G}$ can be regarded as a row vector of ring elements $[1, 3, 9, \ldots, 3^{k-1}] \in \mathcal{R}^{1\times k}$. The trapdoor generation algorithm is then analyzed using Theorem 4, and the trapdoor quality is bounded applying Fact 2 to the concatenation of subgaussian random variables $\mathbf{r}_i \leftarrow D_{\mathcal{R},s} \equiv D_{\mathbb{Z},s}^n$. The formal result is stated below.

Theorem 4. *There is a polynomial time algorithm* $\mathsf{GenTrap}(\mathbf{A}', \mathbf{H}, s)$ *that on input a matrix* $\mathbf{A}' \in \mathcal{R}_q^{1\times w}$*, tag* $\mathbf{H} \in \mathcal{R}_q$*, and parameter* $s > \omega(\sqrt{\ln nw})$*, outputs a matrix* $\mathbf{A}'' \in \mathcal{R}_q^{1\times k}$ *and a* $\mathbf{G}$*-trapdoor* $\mathbf{R} \in \mathcal{R}_q^{w\times k}$ *for* $\mathbf{A} = [\mathbf{A}', \mathbf{A}'']$ *with tag* $\mathbf{H}$ *such that* $s_1(\mathbf{R}) = s \cdot O(\sqrt{w} + \sqrt{k} + \omega(\sqrt{\log n}))$*. Moreover, if* $w \geq 2(\lceil \log_2 q \rceil + 1)$ *then with overwhelming probability over the choice of* $\mathbf{A}' \leftarrow \mathcal{U}_w$*, the distribution of* $\mathbf{A}''$ *is statistically close to uniform.*

In order to allow for the generation of trapdoors for multiple matrices that share the same $\mathbf{A}'$, we made $\mathbf{A}'$ an explicit input to the trapdoor generation algorithm. When $\mathbf{A}' \leftarrow \mathcal{U}_w$ is chosen freshly at random, we simply write $\mathsf{GenTrap}(w, \mathbf{H}, s)$ and let $\mathsf{GenTrap}$ output the whole $\mathbf{A} = [\mathbf{A}', \mathbf{A}'']$.

Notice that $\mathbf{G}$-trapdoors generated in the ring setting also satisfy the definition of $\mathbf{G}$-trapdoor for general lattices. So, Theorem 3 can be used as it is, simply by viewing ring trapdoors $\mathbf{R} \in \mathcal{R}_q^{w\times k}$ as matrices $\mathbf{R} \in \mathbb{Z}^{wn\times kn}$ under the standard embedding from $\mathcal{R}$ to the subring of anticirculant matrices. For convenience, we reformulate Theorem 3 as a corollary specialized to the ring setting.

Corollary 3. *There is an efficient algorithm* $\mathsf{SampleD}(\mathbf{A}, \mathbf{u}, \mathbf{R}, s)$ *that on input a matrix* $\mathbf{A} \in \mathcal{R}_q^{1\times(w+k)}$*, a syndrome* $\mathbf{u} \in \mathcal{R}_q$*, a* $\mathbf{G}$*-trapdoor* $\mathbf{R} \in \mathcal{R}_q^{w\times k}$ *for* $\mathbf{A}$ *with invertible tag* $\mathbf{H} \in \mathcal{R}$*, and parameter* $s > \omega(\sqrt{\log n}) \cdot s_1(\mathbf{R})$*, produces a sample statistically close to the distribution* $D_{\Lambda^{\perp}_{\mathbf{u}}(\mathbf{A}),s}$.

We remark that $\mathsf{GenTrap}$ can be called with arbitrary (not necessarily invertible) tags $\mathbf{H}$. The algorithm still outputs a uniformly random $\mathbf{A}$ and small $s_1(\mathbf{R})$, but the inversion algorithm of Corollary 3 cannot be used with such invalid trapdoors.

3 Our Scheme

The scheme is parametrized by an integer n which we assume is a power of 2, and a modulus $q = 3^k$ which we assume to be a power of 3. (Other parameter settings are possible, but we consider these specific values for concreteness.) These parameters define the ring $\mathcal{R}_q = \mathbb{Z}[X]/(\Phi_{2n}(X), q)$, where (for n a power of 2) $\Phi_{2n}(X) = X^n + 1$ is the cyclotomic polynomial of degree n. The scheme also uses the parameters $w = 2\lceil \log_2 q \rceil + 2$, $m = w + k$, $s = n^{3/2} \cdot \omega(\log n)^{3/2}$, and a collection of tags defined below. We recall that the polynomial $\Phi_{2n}(X)$ is irreducible in $\mathbb{Z}[X]$, but it can be factored in $\mathbb{F}_p[X]$ for some primes p. Our choice of $q = 3^k$ ensures that, in $\mathbb{F}_3[X]$, the polynomial $\Phi_{2n}(X)$ factors into the product of just 2 irreducible polynomials of degree $n/2$. (See Fact 1.) In particular, by Corollary 1, any nonzero polynomial of degree less than $n/2$ with coefficients in $\{0, \pm 1\}$ is invertible in $\mathcal{R}_q$.

Tags For any real constants $c > 1$ and $\alpha \geq \frac{1}{c-1}$ (fixed throughout the rest of this section) define the sets of *tag prefixes* $\mathcal{T}_i = \{0,1\}^{c_i}$ of (strictly increasing) lengths $c_0 = 0$, $c_i = \lfloor \alpha c^i \rfloor$ for $i \in \{1, \ldots, d\}$ where $d = \lfloor \log_c(n/(2\alpha)) \rfloor = O(\log n)$. We identify each tag prefix $t = [t_0, \ldots, t_{c_i-1}] \in \mathcal{T}_i$ with a corresponding ring element $t(X) = \sum_{j<c_i} t_j X^j \in \mathcal{R}_q$ with binary coefficients $t_j \in \{0,1\}$ and degree less than $c_i \leq c_d \leq n/2$. It follows that for any two distinct tag prefixes $t, t' \in \mathcal{T}_i$, the difference $(t(X) - t'(X))$ is invertible in $\mathcal{R}_q$. For any *full tag* $t \in \mathcal{T} = \mathcal{T}_d$ and $i \leq d$, we write $t_{\leq i} \in \mathcal{T}_i$ for its prefix of length c_i, and $t_{[i]}$ for the (ring) difference $t_{\leq i}(X) - t_{\leq i-1}(X) \in \mathcal{R}_q$.

Unlike previous work using tags [29,11,9], our construction relies not only on the algebraic (invertibility) properties of tags, but also on their geometric properties, described in the following lemma.

Lemma 5. *For any* $i \leq d$ *and tags* $t, t' \in \mathcal{T}$*, one has* $s_1((t-t')_{[i]}) \leq c_i - c_{i-1}$.

Proof. Since the difference $(t-t')_{[i]}$ is a trinary polynomial with at most $c_i - c_{i-1}$ nonzero coefficients, we have $\|(t-t')_{[i]}\|_1 \leq c_i - c_{i-1}$. It follows from Lemma 3 that $s_1((t-t')_{[i]}) \leq \|(t-t')_{[i]}\|_1 \leq c_i - c_{i-1}$.

3.1 Our Scheme

Key Generation $\mathsf{naSS.KeyGen}(n)$*:* The key generation algorithm runs $(\mathbf{A}, \mathbf{R}) \leftarrow \mathsf{GenTrap}(w, \mathbf{I}, \sigma)$ with $\sigma = \omega(\sqrt{\log n})$, and chooses $\mathbf{A}_{[0]}, \mathbf{A}_{[1]}, \dots \mathbf{A}_{[d]}, \mathbf{U} \in \mathcal{R}_q^{1\times k}$ and $\mathbf{v} \in \mathcal{R}_q$ uniformly at random. It then outputs the secret key $sk = \mathbf{R}$, and public key $pk = (\mathbf{A}, \mathbf{A}_{[0]}, \mathbf{A}_{[1]}, \dots \mathbf{A}_{[d]}, \mathbf{U}, \mathbf{v})$. The public key implicitly defines a collection of matrices $\mathbf{A}_t = [\mathbf{A}|\mathbf{A}_{[0]} + \sum_{i=1}^{d} t_{[i]} \cdot \mathbf{A}_{[i]}]$ indexed by the tags $t \in \mathcal{T}$.

Since $\sigma = \omega(\sqrt{\log n})$, by Theorem 4 and Lemma 2, the distribution of $\mathbf{A} \in \mathcal{R}_q^{1\times m}$ is statistically close to $\mathcal{U}_m$, and $\mathbf{R}$ is a $\mathbf{G}$-trapdoor for $\mathbf{A}$ (and therefore also for all $\mathbf{A}_t$) with invertible tag $\mathbf{I}$ and quality $s_1(\mathbf{R}) \leq \sqrt{n} \cdot \omega(\log n)$.

Signature $\mathsf{naSS.Sign}(sk = \mathbf{R}, \boldsymbol{\mu} \in \{0,1\}^{nk} \subset \mathcal{R}_q^k)$*:* Parse $\boldsymbol{\mu}$ as a vector of $\mathcal{R}_q^k$ splitting the nk bits into k binary polynomials. Choose a uniformly random tag $t \in \mathcal{T}$, and compute the matrix $\mathbf{A}_t$ and ring element $\mathbf{u} = \mathbf{U} \cdot \boldsymbol{\mu} + \mathbf{v}$. Then, use the $\mathbf{G}$-trapdoor $\mathbf{R}$ to sample a vector $\mathbf{s} \leftarrow \mathsf{SampleD}(\mathbf{A}, \mathbf{u}, \mathbf{R}, s)$. Output the pair $\sigma = (t, \mathbf{s})$ as the signature.

Verification $\mathsf{naSS.Verif}(pk, \boldsymbol{\mu} \in \{0,1\}^{nk} \subset \mathcal{R}_q^k, \sigma = (t, \mathbf{s}))$*:* Compute $\mathbf{A}_t$ and $\mathbf{u} = \mathbf{U} \cdot \boldsymbol{\mu} + \mathbf{v}$ as in the signing algorithm. Then, check that $\|\mathbf{s}\| \leq s\sqrt{nm}$ and that $\mathbf{A}_t \cdot \mathbf{s} = \mathbf{u}$.

Correctness The correctness of the scheme is easily verified: Since $s > \omega(\sqrt{\log n}) \cdot s_1(\mathbf{R})$, by Corollary 3 the vector $\mathbf{s}$ produced during the signature generation process follows the distribution $D_{\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A}_t), s}$ and has length at most $s\sqrt{nm} = O(s\sqrt{nk})$ with overwhelming probability. So, the signature $(t, \mathbf{s})$ is accepted by the verification algorithm.

3.2 Security Proof

The security of the scheme is based on the following homomorphic property of $\mathbf{G}$-trapdoors over rings. We remark that the property makes essential use of the commutativity of matrices corresponding to ring elements in $\mathcal{R}_q$, so it does not trivially adapts to general lattices, unless one restricts the set of tags to scalar matrices.

Lemma 6. *For $i = 0, \dots, d$, let $\mathbf{R}_{[i]} \in \mathcal{R}^{w\times k}$ be a $\mathbf{G}$-trapdoor for $[\mathbf{A}, \mathbf{A}_{[i]}]$ with tag $\mathbf{H}_{[i]} \in \mathcal{R}_q$, where $\mathbf{A}_{[i]} \in \mathcal{R}_q^{1\times k}$. Then, any linear combination $\mathbf{R} = \sum_i c_i \cdot \mathbf{R}_{[i]}$ with $c_i \in \mathcal{R}_q$ is a $\mathbf{G}$-trapdoor for $[\mathbf{A}, \sum_i c_i \mathbf{A}_{[i]}]$ with tag $\mathbf{H} = \sum_i c_i \mathbf{H}_{[i]}$.*

Proof. By definition of $\mathbf{G}$-trapdoor, we know that $[\mathbf{A}, \mathbf{A}_{[i]}](\mathbf{R}_{[i]}, \mathbf{I}) = \mathbf{H}_{[i]}\mathbf{G}$ for all i. Therefore

$$\begin{aligned}\left[\mathbf{A}, \sum\nolimits_i c_i \mathbf{A}_{[i]}\right](\mathbf{R}, \mathbf{I}) &= \mathbf{AR} + \sum\nolimits_i c_i \mathbf{A}_{[i]} = \sum\nolimits_i c_i(\mathbf{AR}_{[i]} + \mathbf{A}_{[i]}) \\ &= \sum\nolimits_i c_i [\mathbf{A}, \mathbf{A}_{[i]}](\mathbf{R}_{[i]}, \mathbf{I}) = \sum\nolimits_i c_i \mathbf{H}_{[i]}\mathbf{G} \quad = \mathbf{HG}.\end{aligned}$$

Therefore $\mathbf{R}$ is a $\mathbf{G}$-trapdoor with tag $\mathbf{H}$.

Theorem 5 (EUF-naCMASecurity). *Under the* $\mathrm{RingSIS}_{n,m,q,\beta}$ *assumption for* $\beta = \tilde{O}(n^{7/2})$*, the above scheme* naSS *is* **EUF-naCMA** *secure. More precisely, if there exists an attacker* $\mathcal{A}$ *against* $\mathbf{EUF\text{-}naCMA}_{\mathsf{naSS}}$ *that runs in time* T*, makes at most* Q *queries where* $1 \leq Q \leq 2^{o(n)}$ *and succeeds with probability* $\epsilon \geq 2^{-o(n)}$*, then, there exists an algorithm* $\mathcal{S}^{\mathcal{A}}$ *that runs in time* $T' = T + \mathrm{poly}(n)$*, and solves* $\mathrm{SIS}(n, w, q, \beta)$ *with probability* $\epsilon' \geq \Omega\left(\frac{\epsilon^{1+c}}{Q^{2c}}\right)$.

The rest of the section is devoted to the proof of the theorem.

Confined Guessing Stage We assume we have an attacker $\mathcal{A}$ against the EUF-naCMA security of naSS that makes at most $Q = 2^{o(n)}$ signature queries, and succeeds with probability $\epsilon \geq 2^{-o(n)}$. Let $i^\star$ the smallest index such that $2Q^2/\epsilon \leq |\mathcal{T}_{i^\star}|$. (Notice that such index exists because $2Q^2/\epsilon = 2^{o(n)} \leq 2^{\lfloor \frac{n}{2c} \rfloor} \leq |\mathcal{T}|$.) This guarantees that, if one chooses Q tags at random in $\mathcal{T}_{i^\star}$, then they will be all distinct except with probability at most $\epsilon/2$.

The simulator $\mathcal{S}$ receives Q non-adaptive signature queries $\boldsymbol{\mu}^{(0)} \dots \boldsymbol{\mu}^{(Q-1)}$ from $\mathcal{A}$. For each message $\boldsymbol{\mu}^{(j)}$, the simulator $\mathcal{S}$ chooses a uniformly random tag $t^{(j)} \in \mathcal{T}$. If a collision of prefixes happens (i.e., if $t^{(j)}_{\leq i^\star} = t^{(k)}_{\leq i^\star}$ for some $j \neq k$) the simulator aborts. (This happens with probability at most $\epsilon/2$.) Otherwise, $\mathcal{S}$ chooses a prefix $t^\star_{\leq i^\star} \in \mathcal{T}_{i^\star}$ uniformly at random. (The rest of the tag $t^\star$ will be specified later on.) The hope is that the adversary will output a forgery $(t^\diamond, \mathbf{s}^\diamond)$ such that $t^\diamond_{\leq i^\star} = t^\star_{\leq i^\star}$. We will make the adversary's view statistically independent from the choice of $t^\star_{\leq i^\star} \in \mathcal{T}_{i^\star}$, so that $t^\diamond_{\leq i^\star} = t^\star_{\leq i^\star}$ will hold true with probability $1/|\mathcal{T}_{i^\star}|$.

Simulating Key Generation and Signatures The simulator also receives a RingSIS challenge, the row vector $\mathbf{A} \leftarrow \mathcal{U}_m$, from which it will build the public key. This is done by running $(\mathbf{A}_{[i]}, \mathbf{R}_{[i]}) \leftarrow \mathsf{GenTrap}(\mathbf{A}, \mathbf{H}_{[i]}, \sigma')$ with $\sigma' = \omega(\sqrt{\log n})$) for $i = 0, \dots, d$ and

$$\mathbf{H}_{[i]} = \begin{cases} 0 \in \mathcal{R}_q & \text{if } \ i > i^\star \\ 1 \in \mathcal{R}_q & \text{if } \ 1 \leq i \leq i^\star \\ -t^\star_{\leq i^\star} & \text{if } \ i = 0. \end{cases}$$

Since $\omega(\sqrt{\log n}) \leq \sigma'$, by Theorem 4 the matrices $\mathbf{A}_{[i]}$ are statistically close to uniform. Moreover, by Fact 2, each $\mathbf{R}_{[i]} \in \mathcal{R}^{m \times k}$ is a $\mathbf{G}$-trapdoor for $[\mathbf{A}, \mathbf{A}_{[i]}]$ with $s_1(\mathbf{R}_{[i]}) \leq \sqrt{n} \cdot \omega(\log n)$. Therefore, by Lemma 6, $\mathbf{R}_t = \mathbf{R}_{[0]} + \sum_{i=1}^d t_{[i]} \cdot \mathbf{R}_{[i]}$ is a $\mathbf{G}$-trapdoor for $\mathbf{A}_t = [\mathbf{A}, \mathbf{A}_{[0]} + \sum_i t_{[i]} \cdot \mathbf{A}_{[i]}]$ with tag $\mathbf{H}_t = t_{\leq i^\star} - t^\star_{\leq i^\star}$. The quality of this trapdoor is

$$\begin{aligned} s_1(\mathbf{R}_t) &\leq s_1(\mathbf{R}_{[0]}) + \sum\nolimits_i s_1(t_{[i]} \cdot \mathbf{R}_{[i]}) \leq \left(1 + \sum\nolimits_i s_1(t_{[i]})\right) \max_i s_1(\mathbf{R}_{[i]}) \\ &\leq \left(1 + \sum\nolimits_i (c_i - c_{i-1})\right) \sqrt{n} \cdot \omega(\log n) = n^{3/2} \cdot \omega(\log n) \end{aligned}$$

where we have used the geometric bound $s_1(t_{[i]}) \leq c_i - c_{i-1}$ from Fact 5. So, the simulator can use $\mathbf{R}_t$ as a trapdoor to sign messages with tag t as long as $\mathbf{H}_t$ is invertible. We observe that $\mathbf{H}_t = \mathbf{0}$ whenever $t^\star_{\leq i^\star} = t_{\leq i^\star}$ (i.e., when $t^\star_{\leq i^\star}$ is a

prefix of t), and it is invertible otherwise. So, the simulator can efficiently answer all signature queries except at most for one index j such that $t^{(j)}_{\leq i^\star} = t^\star_{\leq i^\star}$. If such index exists, set $\boldsymbol{\mu}^\star = \boldsymbol{\mu}^{(j)}$ and $t^\star = t^{(j)}$ (recall that we've only chosen the prefix $t^\star_{\leq i^\star}$ of $t^\star$ at the confined guessing stage), otherwise $\mathcal{S}$ chooses a random $\boldsymbol{\mu}^\star$ and a random $t^\star$ extension of $t^\star_{\leq i^\star}$. We will use our last degree of freedom $\mathbf{v}$ to "program" a signature for this only message $\boldsymbol{\mu}^\star$: choose a signature $\mathbf{s}^\star \leftarrow D^m_{\mathcal{R},s}$, and set $\mathbf{v} = \mathbf{A}_{t^\star}\mathbf{s}^\star - \mathbf{U}\boldsymbol{\mu}^\star$. Applying Lemma 4, we check that $\mathbf{v}$ is close to uniform and independent of $\mathbf{A}_{t^\star}$, $\mathbf{U}$ and $\boldsymbol{\mu}$. This shows how to efficiently simulate public key and signatures that are indistinguishable from a real attack.

Notice that we have not specified how to choose $\mathbf{U}$ yet. In order to for the simulator to exploit the forgery, we want $\mathbf{U} = \mathbf{A}\mathbf{R}_\mathbf{U}$ for some $\mathbf{R}_\mathbf{U}$ with small entries. We can set $\mathbf{R}_\mathbf{U} \leftarrow D_{\mathcal{R},\sigma'}$ so that, by Lemma 4, $\mathbf{U} = \mathbf{A}\mathbf{R}_\mathbf{U}$ is statistically close to uniform, and $s_1(\mathbf{R}_\mathbf{U}) = \sqrt{n} \cdot \omega(\log n)$.

Exploiting the forgery After all those shenanigans from the simulators $\mathcal{S}$, with probability at least $\epsilon/2$, the adversary outputs a forgery $(t^\diamond, \mathbf{s}^\diamond)$ for some message $\boldsymbol{\mu}^\diamond$ of his choice. The simulator's secret hope that $t^\diamond_{\leq i^\star} = t^\star_{\leq i^\star}$ is fulfilled with probability $1/|\mathcal{T}_{i^\star}|$; if not, $\mathcal{S}$ aborts. Otherwise we have

$$\mathbf{A}_{t^\star} \cdot \mathbf{s}^\star = \mathbf{U} \cdot \boldsymbol{\mu}^\star + \mathbf{v} \qquad \text{and} \qquad \mathbf{A}_{t^\diamond} \cdot \mathbf{s}^\diamond = \mathbf{U} \cdot \boldsymbol{\mu}^\diamond + \mathbf{v}$$

Recall that for any tag $t \in \mathcal{T}$ we have $\mathbf{A}_t = [\mathbf{A}|\mathbf{H}_t\mathbf{G} - \mathbf{A}\mathbf{R}_t]$ ($\mathbf{R}_t$ is a $\mathbf{G}$-trapdoor of $\mathbf{A}_t$ with tag $\mathbf{H}_t$); additionally the condition $t^\diamond_{\leq i^\star} = t^\star_{\leq i^\star}$ ensures $\mathbf{H}_{t^\star} = \mathbf{H}_{t^\diamond} = \mathbf{0}$. We derive

$$[\mathbf{A}|-\mathbf{A}\mathbf{R}_{t^\star}|-\mathbf{A}\mathbf{R}_\mathbf{U}] \cdot \begin{bmatrix} \mathbf{s}^\star_1 \\ \mathbf{s}^\star_2 \\ \boldsymbol{\mu}^\star \end{bmatrix} = \mathbf{v} = [\mathbf{A}|-\mathbf{A}\mathbf{R}_{t^\diamond}|-\mathbf{A}\mathbf{R}_\mathbf{U}] \cdot \begin{bmatrix} \mathbf{s}^\diamond_1 \\ \mathbf{s}^\diamond_2 \\ \boldsymbol{\mu}^\diamond \end{bmatrix}.$$

In particular we obtain $\mathbf{A}\mathbf{w} = \mathbf{0}$ for

$$\mathbf{w} = (\mathbf{s}^\star_1 - \mathbf{s}^\diamond_1 - (\mathbf{R}_{t^\star} \cdot \mathbf{s}^\star_2 - \mathbf{R}_{t^\diamond} \cdot \mathbf{s}^\diamond_2) - \mathbf{R}_\mathbf{U}(\boldsymbol{\mu}^\star - \boldsymbol{\mu}^\diamond)).$$

Quite obviously, $\mathbf{w}$ is small (we will quantify below). Less obviously, it is nonzero, except with negligible probability. We split our analysis into 4 different cases, corresponding to different types of forgeries $(\mu^\star, t^\star, \mathbf{s}^\star) \neq (\mu^\diamond, t^\diamond, \mathbf{s}^\diamond)$:

case 1 $\mathbf{s}^\star_2 \neq \mathbf{s}^\diamond_2$. Even revealing $\mathbf{R}_\mathbf{U}$ and all $\mathbf{R}_{[i]}$ for $i > 0$, one has that $\mathbf{R}_{[0]} \cdot (\mathbf{s}^\star_1 - \mathbf{s}^\diamond_1)$ conditioned on the knowledge of $\bar{\mathbf{A}}$ and $\mathbf{A}_{[0]} = \mathbf{A}\mathbf{R}_{[0]}$ contains at least $\Omega(n)$ bits of min-entropy, using Corollary 2. In particular the probability that $\mathbf{w} = \mathbf{0}$ is less than $2^{-\Omega(n)}$.

case 2 $\boldsymbol{\mu}^\star \neq \boldsymbol{\mu}^\diamond$. Even revealing all $\mathbf{R}_{[i]}$ for $i \geq 0$, one has that $\mathbf{R}_\mathbf{U} \cdot (\mathbf{s}^\star_1 - \mathbf{s}^\diamond_1)$ conditioned on the knowledge of $\bar{\mathbf{A}}$ and $\mathbf{U} = \mathbf{A}\mathbf{R}_\mathbf{U}$ contains at least $\Omega(n)$ bits of min-entropy, using Corollary 2. In particular the probability that $\mathbf{w} = \mathbf{0}$ is less than $2^{-\Omega(n)}$.

case 3 $\mathbf{s}^\star_1 = \mathbf{s}^\diamond_1$, $t^\star \neq t^\diamond$. Choose some i such that $t^\star_{[i]} \neq t^\diamond_{[i]}$. Even revealing $\mathbf{R}_\mathbf{U}$ and all $\mathbf{R}_{[j]}$ for $j \neq i$, one has that $\mathbf{R}_{[i]} \cdot \mathbf{s}^\star_1$ conditioned on the knowledge of $\bar{\mathbf{A}}$ and

$\mathbf{A}_{[i]} = \mathbf{A}\mathbf{R}_{[i]}$ contains at least $\Omega(n)$ bits of min-entropy, using Corollary 2. So does $(t^\star_{[i]} - t^\diamond_{[i]}) \cdot \mathbf{R}_{[i]} \cdot \mathbf{s}^\star_1$ since $t^\star_{[i]} - t^\diamond_{[i]}$ is an invertible element of $\mathcal{R}_q$ (Corollary 1). In particular the probability that $\mathbf{w} = \mathbf{0}$ is less than $2^{-\Omega(n)}$.

case 4 $\mathbf{s}^\star_2 = \mathbf{s}^\diamond_2, \boldsymbol{\mu}^\star = \boldsymbol{\mu}^\diamond, t^\star = t^\diamond, \mathbf{s}^\star_1 \neq \mathbf{s}^\diamond_1$. In this case one notices that $\mathbf{w} = \mathbf{s}^\star_1 - \mathbf{s}^\diamond_1 \neq \mathbf{0}$ and concludes.

Size of the extracted SIS *solution* Because $\mathbf{s}^\star, \mathbf{s}^\diamond$ are valid signatures, $\|\mathbf{s}^\star\|, \|\mathbf{s}^\diamond\| \leq s\sqrt{m} \leq n^2 w \cdot \omega(\log n)^{3/2}$. Additionally $s_1(\mathbf{R}_t) \leq n^{3/2} \cdot \omega(\log n)$ for any tag $t \in \mathcal{T}$, as proved above, and $\|\boldsymbol{\mu}^\star\|, \|\boldsymbol{\mu}^\diamond\| \leq \sqrt{m} = O(\sqrt{nk})$ and $\mathbf{R}_\mathbf{U} \leq \sqrt{n} \cdot \omega(\log n)$. Combining all those bounds we obtain

$$\|\mathbf{w}\| \leq n^{7/2} \cdot \log n \cdot \omega(\log n)^{5/2}.$$

Success probability of the simulation The success probability ϵ' of the simulator is at least $(\epsilon - \epsilon/2)/|\mathcal{T}_{i^\star}| - 2^{-\Omega(n)}$ where

- ϵ is the success probability of the attacker,
- $\epsilon/2$ bounds the probability of a collision of tags,
- $1/|\mathcal{T}_{i^\star}|$ is the probability that the confined guess is correct, i.e., $t^\diamond_{\leq i^\star} = t^\star_{\leq i^\star}$, and
- $2^{-\Omega(n)}$ bounds the probability that the extracted SIS solution is zero.

Our choice of $i^\star$ (see confined guessing stage) guarantees that $2^{c_{i^\star - 1}} < \frac{2Q^2}{\epsilon} \leq 2^{c_{i^\star}} = |\mathcal{T}_{i^\star}|$. We also have $c_{i^\star} \leq \alpha c^{i^\star} = c(\alpha c^{i^\star - 1}) < c(c_{i^\star - 1} + 1)$. Therefore $|\mathcal{T}_{i^\star}| = 2^{c_{i^\star}} \leq 2^{c \cdot (c_{i^\star - 1} + 1)} \leq \left(\frac{4Q^2}{\epsilon}\right)^c$. Overall the success probability of solving the SIS instance is at least

$$\epsilon' \geq \frac{\epsilon}{2}\left(\frac{\epsilon}{4Q^2}\right)^c - 2^{-\Omega(n)} = \Omega\left(\frac{\epsilon^{1+c}}{Q^{2c}}\right).$$

□

Acknowledgments

The authors wish to thank Sorina Ionica for helpful conversations, as well as the anonymous CRYPTO'14 reviewers for pointing out several issues in a preliminary version of this paper. This research was supported in part by the DARPA PROCEED program and NSF grant CNS-1117936. Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA or NSF.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (h)ibe in the standard model. In H. Gilbert, editor, *Advances in Cryptology EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 553–572. Springer Berlin / Heidelberg, 2010.

2. S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical ibe. In T. Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*. Springer, 2010.
3. M. Ajtai. Generating hard instances of lattice problems. *Complexity of Computations and Proofs, Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC '96.
4. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of STOC '97*, pages 284–293. ACM, May 1997.
5. Y. Arbitman, G. Dogon, V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. Swifftx: A proposal for the sha-3 standard. Submission to NIST, 2008. `http://csrc.nist.gov/groups/ST/hash/sha-3/Round1/documents/SWIFFTX.zip`.
6. S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, pages 28–47, 2014.
7. F. Böhl, D. Hofheinz, T. Jager, J. Koch, J. H. Seo, and C. Striecks. Practical signatures from standard assumptions. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 461–485, Athens, Greece, May 26–30, 2013. Springer, Berlin, Germany.
8. F. Böhl, D. Hofheinz, T. Jager, J. Koch, and C. Striecks. Confined guessing: New signatures from standard assumptions. Cryptology ePrint Archive, Report 2013/171, 2013. `http://eprint.iacr.org/2013/171`.
9. X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In P. Q. Nguyen and D. Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 499–517, Paris, France, May 26–28, 2010. Springer, Berlin, Germany.
10. Z. Brakerski and Y. T. Kalai. A framework for efficient signatures, ring signatures and identity based encryption in the standard model. Cryptology ePrint Archive, Report 2010/086, 2010. `http://eprint.iacr.org/2010/086`.
11. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. *Journal of Cryptology*, 25(4):601–639, Oct. 2012.
12. L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice signatures and bimodal gaussians. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 40–56, Santa Barbara, CA, USA, Aug. 18–22, 2013. Springer, Berlin, Germany.
13. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In R. E. Ladner and C. Dwork, editors, *40th ACM STOC*, pages 197–206, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press.
14. T. Güneysu, V. Lyubashevsky, and T. Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In E. Prouff and P. Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 530–547, Leuven, Belgium, Sept. 9–12, 2012. Springer, Berlin, Germany.
15. J. Hoffstein, J. Pipher, J. Schanck, J. H. Silverman, and W. Whyte. Practical signatures from the partial fourier recovery problem. Cryptology ePrint Archive, Report 2013/757, 2013. `http://eprint.iacr.org/2013/757`.
16. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a ring based public key cryptosystem. In *Proceedings of ANTS-III*, volume 1423 of *LNCS*, pages 267–288. Springer, June 1998.
17. S. Hohenberger and B. Waters. Realizing hash-and-sign signatures under standard assumptions. In A. Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 333–350, Cologne, Germany, Apr. 26–30, 2009. Springer, Berlin, Germany.

18. S. Hohenberger and B. Waters. Short and stateless signatures from the RSA assumption. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 654–670, Santa Barbara, CA, USA, Aug. 16–20, 2009. Springer, Berlin, Germany.
19. L. Lamport. Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, Oct. 1979.
20. R. Lidl and H. Niederreiter. *Finite Fields*. Reading, MA: Addison-Wesley, 1983. Encyclopedia of Mathematics and its Applications, Volume 20.
21. V. Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In *Proceedings of Asiacrypt*, LNCS. Springer, 2009.
22. V. Lyubashevsky. Lattice signatures without trapdoors. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755, Cambridge, UK, Apr. 15–19, 2012. Springer, Berlin, Germany.
23. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proceedings of ICALP*, volume 4052 of *LNCS*, pages 144–155. Springer, July 2006.
24. V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In R. Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 37–54, San Francisco, CA, USA, Mar. 19–21, 2008. Springer, Berlin, Germany.
25. V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In *Proceedings of TCC*, volume 4948 of *LNCS*, pages 37–54. Springer, Mar. 2008.
26. V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: a modest proposal for FFT hashing. In *Fast Software Encryption – Proceedings*, volume 5086 of *LNCS*, pages 54–72. Springer, Feb. 2008.
27. V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54, Athens, Greece, May 26–30, 2013. Springer, Berlin, Germany.
28. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, Dec. 2007. Preliminary version in FOCS '02.
29. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718, Cambridge, UK, Apr. 15–19, 2012. Springer, Berlin, Germany.
30. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007. Preliminary version in FOCS '04.
31. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proceedings of TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, Mar. 2006.
32. C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *Proceedings of STOC*, pages 478–487. ACM, June 2007.
33. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of ACM*, 56(6):34, Sept. 2009. Preliminary version in STOC '05.
34. J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd ACM STOC*, pages 387–394, Baltimore, Maryland, USA, May 14–16, 1990. ACM Press.
35. M. Rückert. Strongly unforgeable signatures and hierarchical identity-based signatures from lattices without random oracles. In *Post Quantum Cryptography – Proceedings*, 2010.

36. J. H. Seo. Short signatures from Diffie-Hellman: Realizing short public key. Cryptology ePrint Archive, Report 2012/480, 2012. http://eprint.iacr.org/2012/480.
37. R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *CoRR*, abs/1011.3027, 2010. http://www-personal.umich.edu/~romanv/papers/non-asymptotic-rmt-plain.pdf.

A Missing proofs

Proof (Fact 2). For a vector $\mathbf{v} \in R^n$ over ring R, let $\mathrm{Diag}(\mathbf{v})$ denotes the diagonal matrice with entries $v_1 \dots v_n$. Notice that the component wise product of two vectors $\mathbf{f} \odot \mathbf{g}$ can be written as the matrix-vector product $\mathrm{Diag}(\mathbf{f}) \cdot \mathbf{g}$. This gives the identity $\sigma(f \cdot g) = \mathrm{Diag}(\sigma(f)) \cdot \sigma(g)$ for $f, g \in \mathcal{R}$ with $\sigma : \mathcal{R} \to \mathbb{C}^n$ denoting the canonical embedding:

$$\sigma : f \in \mathcal{R} \mapsto (f(\omega_1), \dots f(\omega_\ell)) \in \mathbb{C}^n \text{ where } \omega_\ell = e^{(2\ell-1)\imath\pi/n}$$

. Let $\mathbf{R} = (r_{i,j}) \leftarrow D_{\mathcal{R},s}^{w \times k}$; and set

$$\mathbf{D} = \begin{bmatrix} \mathbf{D}_{1,1} & \cdots & \mathbf{D}_{1,k} \\ \vdots & & \vdots \\ \mathbf{D}_{w,1} & \cdots & \mathbf{D}_{w,k} \end{bmatrix} \in \mathbb{C}^{nw \times nk} \quad \text{and} \quad \mathbf{D}_{i,j} = \mathrm{Diag}(\sigma(r_{i,j})) \in \mathbb{C}^{n \times n}.$$

We extend the canonical embedding $\sigma : \mathcal{R} \to \mathbb{C}^n$ to vectors in $\mathcal{R}^d$ as its componentwise application; $\sigma(\mathbf{v}) = (\sigma(v_1), \dots \sigma(v_k)) \in \mathbb{C}^{nk}$. With this notation, we have $\sigma(\mathbf{R} \cdot \mathbf{v}) = \mathbf{D} \cdot \sigma(\mathbf{v})$; and because the canonical embedding σ is a scaled isometry, we have $s_1(\mathbf{R}) = s_1(\mathbf{D})$.

Permuting rows and column, $\mathbf{D}$ can be rewritten as the block-diagonal matrix $\mathbf{B} = \mathrm{Diag}(\mathbf{B}_1, \dots \mathbf{B}_n) \in \mathbb{C}^{nw \times nk}$, $\mathbf{B}_\ell \in \mathbb{C}^{w \times k}$ where the coefficients of $\mathbf{B}_\ell$ are all the embeddings $\sigma_\ell(r_{i,j}) = r_{i,j}(\omega_\ell)$ for $(i,j) \in \{1 \dots w\} \times \{1 \dots k\}$. The coefficients of $\mathrm{Re}(\mathbf{B}_\ell)$ (the real part of $\mathbf{B}_\ell$) are independent and sub-gaussian of parameter $s\sqrt{n}$. Indeed

$$\mathrm{Re}(\mathbf{B}_\ell) = \sum_{k=0}^{n-1} \mathrm{Re}(\omega_\ell^k) \cdot (r_{i,j})_k$$

where the $(r_{i,j})_k$ are independent and sub-gaussian of parameter s while $|\mathrm{Re}(\omega_\ell^k)| \leq 1$. Therefore by Lemma 1

$$s_1(\mathrm{Re}(\mathbf{B}_\ell)) \leq s\sqrt{n} \cdot O(\sqrt{w} + \sqrt{k} + \omega(\sqrt{\log n}))$$

with overwhelming probability. The same results hold for the imaginary part $\mathrm{Im}(\mathbf{B}_\ell)$ of $\mathbf{B}_\ell$. We conclude

$$\begin{aligned} s_1(\mathbf{D}) \leq s_1(\mathbf{B}) \leq \max_\ell s_1(\mathbf{B}_\ell) &\leq \max_\ell \sqrt{s_1(\mathrm{Re}(\mathbf{B}_\ell))^2 + s_1(\mathrm{Im}(\mathbf{B}_\ell))^2} \\ &\leq s\sqrt{n} \cdot O(\sqrt{w} + \sqrt{k} + \omega(\sqrt{\log n})). \end{aligned}$$

Proof (Lemma 4). The proof is adapted from [13, Lemma 5.3]. Consider the lattice $\Lambda(\mathbf{A}^\top)$ spanned by the columns of $\mathbf{A}^\top$ and the vectors of $q\mathbb{Z}^{nw}$; it is the (scaled) dual of $\Lambda^\perp(\mathbf{A})$. We will first show that the minimal distance $\lambda_1^\infty(\Lambda(\mathbf{A}^\top))$ is at least $q/12$ with overwhelming probability, and conclude using [13, Lemma 2.6] that $\eta_\epsilon(\Lambda^\perp(\mathbf{A})) \leq \omega(\sqrt{\ln nw})$ for some negligible function $\epsilon(n)$.

Recall that the irreducible factors of $\Phi_{2n}(X) \bmod 3$ are $P_1(X) = X^{n/2} + X^{n/4} - 1$ and $P_2(X) = X^{n/2} - X^{n/4} - 1$. Setting $\mathfrak{p}_1 = (P_1(X))$, $\mathfrak{p}_2 = (P_2(X))$ the nonzero ideals of $\mathcal{R}_q$ are exactly $\mathfrak{p}_1, 3\mathfrak{p}_1 \dots 3^{k-1}\mathfrak{p}_1$; $\mathfrak{p}_2, 3\mathfrak{p}_2 \dots 3^{k-1}\mathfrak{p}_2$ and $(1), (3), (3^2), \dots (3^{k-1})$.

Now, fix some $x \in \mathcal{R} \setminus \{0\}$, et set $\mathfrak{I} = (x)$, it is one of the nonzero ideal listed above. Let $r \geq n/2$ denotes its rank. Our goal is to prove, that over the randomness of $\mathbf{A} \in \mathcal{R}^{1\times w}$, the probability that $\mathbf{A}x$ falls in in the hypercube $\mathcal{C}^w = \{\mathbf{v} \in \mathcal{R}^w | \ \|\mathbf{v}\|_\infty < q/12\}$ is less than $2^{-O(wr)}$. Because x is a generator of $\mathfrak{I}$ the distribution of $\mathbf{A}x$ is uniform over $\mathfrak{I}^w$. We proceed by bounding the ratio $|\mathcal{C} \cap \mathfrak{I}|/|\mathfrak{I}|$.

Case 1: ($\mathfrak{I} = (3^h)$ for $h \in \{0 \dots k-1\}$). Observe that $|\mathcal{C} \cap \mathfrak{I}| \leq |\{3^h\mathbb{Z} \cap (-q/12, q/12)\}^n| \leq \lceil 3^{k-h}/6 \rceil^n$; which leads to

$$|\mathcal{C} \cap \mathfrak{I}|/|\mathfrak{I}| \leq \left(\frac{3^{k-h}/6+1}{3^{k-h}}\right)^n \leq \left(\frac{1}{6} + \frac{1}{3^{k-h}}\right)^n \leq 2^{-n}.$$

Case 2: ($\mathfrak{I} = 3^h\mathfrak{p}_i$ for $h \in \{0 \dots k-1\}$). Start by noting that any element e of $\mathfrak{I}$ can be uniquely written $e = P_i(X) \cdot s$ where $s = \sum_{i=0}^{n/2-1} s_i X^i$ is a polynomial and of degree strictly less than $n/2$ in the ideal (3^h) of $\mathcal{R}$. Also note that $\|e\|_\infty \leq q/12$ implies $\|s\|_\infty \leq q/12$, indeed for $i \in \{0 \dots n/4-1\}$ we have $e_i = -s_i$ and for $i \in \{n/4 \dots n/2-1\}$ we have $e_{i+n/2} = s_i$. Using a similar counting argument on valid values of s we derive

$$|\mathcal{C} \cap \mathfrak{I}|/|\mathfrak{I}| \leq \left(\frac{3^{k-h}/6+1}{3^{k-h}}\right)^{n/2} \leq \left(\frac{1}{6} + \frac{1}{3^{k-h}}\right)^{n/2} \leq 2^{-n/2}.$$

Taking the union bound over all nonzero x we conclude that $\lambda_1^\infty(\Lambda(\mathbf{A}^\top)) \geq q/12$ except with probability $q^n \cdot 2^{-nw/2} \leq 2^{-\Omega(n)}$.

Proof (Corollary 2). Without loss of generality assume that $\mathbf{v}_1 \neq 0$. Applying the previous Lemma 4 on $\sum_{i\geq 2} \mathbf{a}_i \cdot \mathbf{x}_i$, the knowledge of $\sum_i \mathbf{a}_i \cdot \mathbf{x}_i = \mathbf{a}_1 \cdot \mathbf{x}_1 + \sum_{i\geq 2} \mathbf{a}_i \cdot \mathbf{x}_i$ reveals only negligible any information about $\mathbf{x}_1$. Also note that $\mathbf{x}_1 \bmod 3$ is negligibly close to uniform ($\eta_\epsilon(3\mathbb{Z}) \leq \omega(\sqrt{\ln n})$ for some negligible function $\epsilon(n)$).

Setting $\mathfrak{I} = (\mathbf{v}_1) \neq (0)$ we deduce that $\mathbf{v}_1 \cdot \mathbf{x}_1 \bmod 3\mathfrak{I}$ is almost uniform in $\mathfrak{I}/3\mathfrak{I}$. Recall from the previous proof that the only nonzero ideals of $\mathcal{R}_q$ are exactly $\mathfrak{p}_1, 3\mathfrak{p}_1 \dots 3^{k-1}\mathfrak{p}_1$; $\mathfrak{p}_2, 3\mathfrak{p}_2 \dots 3^{k-1}\mathfrak{p}_2$ and $(1), (3), (3^2), \dots (3^{k-1})$ where both $\mathfrak{p}_1$ and $\mathfrak{p}_2$ are ideals of rank $n/2$. This implies that $|\mathfrak{I}/3\mathfrak{I}| = 3^{n/2}$ or 3^n. We conclude that $\mathbf{v}_1 \cdot \mathbf{x}_1$ has at least $\Omega(n)$ bits of entropy and so has $\sum_i \mathbf{v}_i \cdot \mathbf{x}_i$.