

Asymptotically Efficient Lattice-Based Digital Signatures ^{*}

Vadim Lyubashevsky¹ and Daniele Micciancio²

¹ IBM Research – Zurich

² University of California, San Diego

Abstract. We present a general framework that converts certain types of linear collision-resistant hash functions into one-time signatures. Our generic construction can be instantiated based on both general and ideal (e.g. cyclic) lattices, and the resulting signature schemes are provably secure based on the worst-case hardness of approximating the shortest vector (and other standard lattice problems) in the corresponding class of lattices to within a polynomial factor. When instantiated with ideal lattices, the time complexity of the signing and verification algorithms, as well as key and signature size is almost linear (up to poly-logarithmic factors) in the dimension n of the underlying lattice. Since no sub-exponential (in n) time algorithm is known to solve lattice problems in the worst case, even when restricted to ideal lattices, our construction gives a digital signature scheme with an essentially optimal performance/security trade-off.

1 Introduction

Digital signature schemes, initially proposed in Diffie and Hellman’s seminal paper [DH76] and later formalized by Goldwasser, Micali and Rivest, [GMR88], are among the most important and widely used cryptographic primitives. Still, our understanding of these intriguing objects is somehow limited. The definition of digital signatures clearly fits within the public key cryptography framework, yet their existence can be shown to be equivalent to the existence of symmetric cryptographic primitives like pseudorandom generators, one-way hash functions, private key encryption, or even just one-way functions [NY89,Rom90].

When efficiency is taken into account, however, digital signatures seem much closer to public key primitives than to symmetric ones. In the symmetric setting, functions are often expected to run in time which is linear or almost linear in the security parameter k . However, essentially all known digital signatures with a supporting proof of security are based on algebraic functions that take at least $\Omega(k^2)$ time to compute, where 2^k is the conjectured hardness of the underlying problem. For example, all factoring-based schemes must use keys of size approximately $O(k^3)$ to achieve k bits of security to counter the best known sub-exponential time factoring algorithms, and modular exponentiation raises the time complexity to over $\omega(k^4)$ even when restricted to small k -bit exponents and implemented with an asymptotically fast integer multiplication algorithm.

Digital signatures based on arbitrary one-way hash functions have also been considered, due to the much higher speed of conjectured one-way functions (e.g., instantiated with common block ciphers as obtained from ad-hoc constructions) compared to the cost of modular squaring or exponentiation operations typical of number theoretic schemes. Still, the performance advantage of

^{*} To appear in Journal of Cryptology. A preliminary version of this work appeared in *Theory of Cryptography Conference* – Proceedings of TCC 2008. This is an improved, extended, and simplified version of that paper. Research supported in part by NSF grants CCF-0634909, CNS-1117936, SNSF ERC Transfer Grant CRETP2-166734 – FELICITY, and the H2020 Project Safecrypto. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

one-way functions is often lost in the process of transforming them into digital signature schemes: constructions of signature schemes from non-algebraic one-way functions almost invariably rely on Lamport and Diffie’s [DH76] one-time signature scheme (and variants thereof) which requires a number of one-way function applications essentially proportional to the security parameter. So, even if the one-way function can be computed in linear time $O(k)$, the complexity of the resulting signature scheme is again at least quadratic $\Omega(k^2)$.

Therefore, a question of great theoretical and practical interest, is whether digital signature schemes can be realized at essentially the same cost as symmetric key cryptographic primitives. While a generic construction that transforms any one-way function into a signature scheme with similar efficiency seems unlikely, one may wonder if there are specific complexity assumptions that allow to build more efficient digital signature schemes than currently known. Ideally, are there digital signature schemes with $O(k)$ complexity, which can be proved as hard to break as solving a computational problem which is believed to require $2^{\Omega(k)}$ time?

1.1 Results and techniques

The main result in this paper is a construction of a provably-secure digital signature scheme with key size and computation time almost linear (up to poly-logarithmic factors) in the security parameter. In other words, we give a new digital signature scheme with complexity $O(k \log^c k)$ which can be proved to be as hard to break as a problem which is currently conjectured to require $2^{\Omega(k)}$ time to solve. The signature scheme is a particular instantiation inside of a general framework that we present for constructing one-time signatures from certain types of *linear* collision-resistant hash functions.

We show how to instantiate our general framework with signature scheme constructions based on standard lattice and coding problems. The lattice problem underlying our most efficient scheme is that of approximating the shortest vector in a lattice with “cyclic” or “ideal” structure, as already used in [Mic07] for the construction of efficient lattice-based one-way functions, and subsequently extended to collision resistant functions in [PR06,LM06]. As in most previous work on lattices, our scheme can be proved secure based on the *worst case* complexity of the underlying lattice problems.

Since one-way functions are known to imply the existence of many other cryptographic primitives (e.g., pseudorandom generators, digital signatures, private key encryption, etc.), the efficient lattice-based one-way functions of [Mic07] immediately yield corresponding cryptographic primitives based on the complexity of cyclic lattices. However, the known generic constructions of cryptographic primitives from one-way functions are usually very inefficient. So, it was left as an open problem in [Mic07] to find *direct* constructions of other cryptographic primitives from lattice problems with performance and security guarantees similar to those of [Mic07]. For the case of collision resistant hash functions, the problem was resolved in [PR06,LM06], which showed that various variants of the one-way function proposed in [Mic07] are indeed collision resistant. In this paper we build on the results of [Mic07,PR06,LM06] to build an asymptotically-efficient lattice-based digital signature scheme.

Theorem 1.1. *There exists a signature scheme (with security parameter k) such that the signature of an n -bit message (for any message size $n = k^{O(1)}$) is of length $\tilde{O}(k)$ and both the signing and verification algorithms take time $\tilde{O}(n+k)$. The scheme is strongly unforgeable in the chosen message attack model, assuming the hardness of approximating the shortest vector problem in all ideal lattices of dimension k to within a factor $\tilde{O}(k^2)$.*

Our signature scheme is based on a standard transformation from one-time signatures (i.e., signatures that allow to securely sign a single message) to general signature schemes, together with a novel construction of a lattice-based one-time signature. We remark that the same transformation from one-time signatures to unrestricted signature schemes was also employed by virtually all previous constructions of digital signatures from arbitrary one-way functions (e.g., [Mer89,NY89,Rom90]). This transformation, which combines one-time signatures together with a tree structure, is relatively efficient and allows one to sign messages with only a logarithmic number of applications of a hash function and a one-time signature scheme [Szy04]. The bottleneck in one-way function based signature schemes is the construction of one-time signatures from one-way functions. The reason for the slowdown is that the one-way function is typically used to sign a k -bit message one bit at a time, so that the entire signature requires k evaluations of the one-way function. In this paper we give a direct construction of one-time signatures, where each signature just requires two applications of the lattice-based collision-resistant function of [Mic07,PR06,LM06]. The same lattice-based hash function can then be used to efficiently transform the one-time signature into an unrestricted signature scheme with only a logarithmic loss in performance.

One-time signature. The high level structure of our general framework is easily explained (see Figure 1). The underlying hardness assumption is the collision resistance of a certain linear hash function family mapping a subset \mathcal{S} of R^m to R^n , where R is some ring. The linear hash function can be represented by a matrix $\mathbf{H} \in R^{n \times m}$ and the secret key is a matrix $\mathbf{K} \in R^{m \times k}$. The public key consists of the function \mathbf{H} and the image $\hat{\mathbf{K}} = \mathbf{H}\mathbf{K}$. To sign a message $\mathbf{m} \in R^k$, we simply compute $\mathbf{s} = \mathbf{K}\mathbf{m}$. To verify that \mathbf{s} is the signature of \mathbf{m} , the verifier checks that \mathbf{s} is in \mathcal{S} and that $\mathbf{H}\mathbf{s} = \hat{\mathbf{K}}\mathbf{m}$. To make sure that the scheme is complete (i.e. valid signatures are accepted), we need to choose the domain of the secret keys and messages so that $\mathbf{K}\mathbf{m}$ is always in \mathcal{S} .

Depending on the choice of the ring R , we obtain one-time signatures based on different complexity assumptions. Choosing $R = \mathbb{Z}_p$ results in schemes based on the SIS problem, $R = \mathbb{Z}_2$ gives us a scheme based on the Small Codeword Problem, and setting $R = \mathbb{Z}[x]/(x^n + 1)$ produces the most efficient scheme based on the Ring-SIS problem.

Security proof. The security of our general framework relies on the assumption that for a random $\mathbf{H} \in R^{n \times m}$ it is hard to find two distinct elements $\mathbf{s}, \tilde{\mathbf{s}} \in \mathcal{S}$ such that $\mathbf{H}\mathbf{s} = \mathbf{H}\tilde{\mathbf{s}}$. In the security proof, when given a random \mathbf{H} by the challenger, the simulator picks a valid secret key \mathbf{K} and outputs $\mathbf{H}, \hat{\mathbf{K}} = \mathbf{H}\mathbf{K}$ as the public key. Since the simulator knows the secret key, she is able to compute the signature, $\mathbf{K}\mathbf{m}$, of any message \mathbf{m} . If an adversary is then able to produce a valid signature $\tilde{\mathbf{s}}$ of some message $\tilde{\mathbf{m}}$, he will satisfy the equation $\mathbf{H}\tilde{\mathbf{s}} = \hat{\mathbf{K}}\tilde{\mathbf{m}} = \mathbf{H}\mathbf{K}\tilde{\mathbf{m}}$. Thus, unless $\tilde{\mathbf{s}} = \mathbf{K}\tilde{\mathbf{m}}$, we will have found a collision for \mathbf{H} . The main technical part of our proof (Theorem 3.2) clarifies the necessary condition so that the probability of $\tilde{\mathbf{s}} \neq \mathbf{K}\tilde{\mathbf{m}}$ is non-negligible. Towards this end, we define a condition called (ϵ, δ) -Hiding and then prove that if the domains of the hash function, key space, and message space satisfy this requirement for a constant ϵ and a δ close to 1, then the one-time signature scheme will be secure based on the hardness of finding collisions in a random \mathbf{H} . We remark that the (ϵ, δ) -Hiding property is purely combinatorial, and so to prove security of different instantiations based on SIS, Ring-SIS, or coding problems, we simply need to show that the sets used in the instantiations of these schemes satisfy this condition.

1.2 Related work

Lamport showed the first construction of a one-time signature based on the existence of one-way functions. In that scheme, the public key consists of the values $f(x_0), f(x_1)$, where f is a one-way function and x_0, x_1 are randomly chosen elements in its domain. The elements x_0 and x_1 are kept secret, and in order to sign a bit i , the signer reveals x_i . This construction requires one application of the one-way function for every bit in the message. Since then, more efficient constructions have been proposed [Mer87,BC92,BM84,EGM96,BM96,HM02], but there was always an inherent limitation in the number of bits that could be signed efficiently with one application of the one-way function [GGKT05].

Provably secure cryptography based on lattice problems was pioneered by Ajtai in [Ajt96], and attracted considerable attention within the complexity theory community because of a remarkable worst-case/average-case connection: it is possible to show that breaking the cryptographic function on the average is at least as hard as solving the lattice problem in the worst-case. Unfortunately, functions related to k -dimensional lattices typically involve a k -dimensional matrix/vector multiplication, and therefore require k^2 time to compute (as well as k^2 storage for keys). A fundamental step towards making lattice-based cryptography more attractive in practice, was taken by Micciancio [Mic07] who proposed a variant of Ajtai’s function which is much more efficient to compute (thanks to the use of certain lattices with a special cyclic structure) and still admits a worst-case/average-case proof of security. The performance improvement in [Mic07] (as well as in subsequent work [PR06,LM06],) comes at a cost: the resulting function is as hard to break as solving the shortest vector problem in the worst case over lattices with a cyclic structure. Still, since the best known algorithms do not perform any better on these lattices than on general ones, it seems reasonable to conjecture that the shortest vector problem is still exponentially hard. It was later shown in [PR06,LM06] that, while the function constructed in [Mic07] was only one-way, it is possible to construct efficient collision-resistant hash functions based on the hardness of problems in lattices with a similar algebraic structure.

1.3 Comparison to the proceedings version of this work

In the proceedings version of this work [LM08], we gave a direct construction of a one-time signature scheme based on the hardness of the Ring-SIS problem. The major difference of that scheme with the Ring-SIS scheme in this paper is the key generation algorithm. In the current work, the secret key is simply chosen according to the uniform distribution from some set. In [LM08], however, choosing a secret key first involved selecting a “shell” with a geometrically degrading probability and then picking a uniformly-random element from it. The security proof in the current paper is also much more modular. In particular, we first present an abstract framework for constructing one-time signatures of a particular type, and then show how this framework can be satisfied with instantiations based on various problems such as SIS, Ring-SIS over the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$, and the Small Codeword Problem. Essentially, this paper is a simpler, more modular, and more general version of [LM08].

We also showed, in the proceedings version, constructions of a Ring-SIS signature scheme that worked over rings $\mathbb{Z}[x]/\langle f(x) \rangle$ for an arbitrary monic, irreducible polynomial $f(x)$. Since the main focus of the current paper is on abstracting out the properties needed for constructions of one-time signatures from linear collision-resistant hash functions, we choose not to complicate matters by also presenting the various manners in which one could do these constructions based on different

forms of the Ring-SIS problem (some of which would require first presenting some background from algebraic number theory). Below, we sketch the different manners in which one could proceed to define and instantiate the one-time signature using different rings. The main difference lies in the manner in which the length of polynomials is defined and the domain and range of the hash function \mathbf{H} .

The simplest definition of length is the “coefficient embedding”, where it is defined by taking the norm of the vector formed by the coefficients of the polynomial. This is the approach taken in [LM08] and involves the use of the “expansion factor” [LM06] which gives an upper bound on the size of the norm of the product compared to the norm of the multiplicands. A different way to define the norm of elements in $\mathbb{Z}[x]/\langle f(x) \rangle$ is the “canonical embedding”, which is the norm of a vector formed by evaluating the polynomial on the n (complex) roots of $f(x)$. The advantage of this latter approach is that bounding the product of the norm is very simple, and does not depend on the modulus $f(x)$, because multiplication is component-wise in the canonical embedding.

If one uses the canonical embedding to define the norm, then one also has a choice as to the domain and range of the hash function \mathbf{H} . Instead of being restricted to the ring $\mathbb{Z}[x]/\langle f(x) \rangle$, one may follow the approach taken in [PR07] and define collision-resistant hash functions over the *ring of integers* of number fields $\mathbb{Q}(\zeta)$ where ζ is a primitive root of $f(x)$. In the case that $f(x)$ is a cyclotomic polynomial and ζ is one of its roots (i.e. some root of unity), the ring of integers of $\mathbb{Q}(\zeta)$ is exactly $\mathbb{Z}[x]/\langle f(x) \rangle$, but in other cases, the ring of integers may be a superset of $\mathbb{Z}[x]/\langle f(x) \rangle$ and more “compact”. Since keys need to be sampled from the domain of \mathbf{H} , it is important that the ring of integers of $\mathbb{Q}(\zeta)$ is efficiently samplable in practice - which is not known to be the case for particularly compact choices. Another choice for the domain (and range) of \mathbf{H} , most applicable when $f(x)$ is a cyclotomic polynomial, is the *dual* of the ring of integers (see [LPR13a,LPR13b]). The idea here would be to have \mathbf{H} and \mathbf{m} be elements of the primal ring, while having \mathbf{K} come from the dual one, which is sometimes a little bit more compact.

We point out that in the case of an irreducible $f(x)$ of the form $f(x) = x^n + 1$, the coefficient and canonical embeddings are simply rigid rotations (and scalings) of each other. Also, the ring of integers of $\mathbb{Q}(\zeta)$, where ζ is a root of $x^n + 1$, is exactly $\mathbb{Z}[x]/\langle x^n + 1 \rangle$, and the dual of this ring is the same ring scaled by an integer. Therefore if we choose to work modulo $x^n + 1$, all the above choices are exactly equivalent.

2 Preliminaries

2.1 Signatures

We recall the definitions of signature schemes and what it means for a signature scheme to be secure. In the next definition, G is called the key-generation algorithm, S is the signing algorithm, V is the verification algorithm, and s and $G(s)$ are, respectively, the signing and verification keys.

Definition 2.1. *A signature scheme consists of a triplet of polynomial-time algorithms (G, S, V) such that for any n -bit message m and secret key s (of length polynomial in n), we have*

$$V(G(s), m, S(s, m)) = 1$$

i.e., $S(s, m)$ is a valid signature for message m with respect to public key $G(s)$.

Notice that, for simplicity, we have restricted our definition to signature schemes where the key generation and signing algorithms are deterministic, given the scheme secret key as input. This is without loss of generality because any signature scheme can be made to satisfy these properties by using the key generation randomness as secret key, and derandomizing the signing algorithm using a pseudorandom function.

A signature scheme is said to be strongly unforgeable (under chosen message attacks) if there is only a negligible probability that any (efficient) adversary, after seeing any number of message/signature pairs for adaptively chosen messages of his choice, can produce a new message/signature pair. This is a stronger notion of unforgeability than the standard one [GMR88], which requires the adversary to produce a signature for a new message. In this paper we focus on strong unforgeability because this stronger property is required in some applications, and all our schemes are easily shown to satisfy this stronger property. A one-time signature scheme is a signature scheme that is meant to be used to sign only a single message, and is only required to satisfy the above definition of security under properly restricted adversaries that receive only one signature/message pair. The formal definition is given below.

Definition 2.2. *A one-time signature scheme (G, S, V) is said to be strongly unforgeable if for every polynomial-time (possibly randomized) adversary \mathcal{A} , the success probability of the following experiment is negligible: choose s uniformly at random, compute $v = G(s)$, pass the public key to the adversary to obtain a query message $\mathbf{m} \leftarrow \mathcal{A}(v)$, produce a signature for the message $\mathbf{s} = S(s, \mathbf{m})$, pass the signature to the adversary to obtain a candidate forgery $(\tilde{\mathbf{m}}, \tilde{\mathbf{s}}) \leftarrow \mathcal{A}(v, \mathbf{s})$, and check that the forgery is valid, i.e., $(\mathbf{m}, \mathbf{s}) \neq (\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$ and $V(v, \tilde{\mathbf{m}}, \tilde{\mathbf{s}}) = 1$.*

2.2 Lattices and the SIS Problem

An n -dimensional integer lattice \mathcal{L} is a subgroup of \mathbb{Z}^n . A lattice \mathcal{L} can be represented by a set of linearly independent generating vectors, called a basis.

Definition 2.3. *For an n -dimensional lattice \mathcal{L} and all $1 \leq i \leq n$, $p \in \{\mathbb{Z}^+, \infty\}$, the positive real numbers $\lambda_i^p(\mathcal{L})$ are defined as*

$$\lambda_i^p(\mathcal{L}) = \arg \min_{x \in \mathbb{R}} (\exists i \text{ linearly independent vectors in } \mathcal{L} \text{ of } \ell_p\text{-norm at most } x).$$

Definition 2.4. *The approximate search Shortest Vector Problem, $\text{SVP}_\gamma^p(\mathcal{L})$ asks to find a vector $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\|_p \leq \gamma \cdot \lambda_1^p(\mathcal{L})$.*

Definition 2.5. *For an n -dimensional lattice, the approximate search Shortest Independent Vector Problem, $\text{SIVP}_\gamma^p(\mathcal{L})$ asks to find n linearly independent vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathcal{L}$ such that $\max_i \|\mathbf{v}_i\|_p \leq \gamma \cdot \lambda_n^p(\mathcal{L})$.*

Definition 2.6. *In the Small Integer Solution problem ($\text{SIS}_{p,n,m,\beta}^\infty$), one is given a matrix $\mathbf{H} \in \mathbb{Z}_p^{n \times m}$ and is asked to find a non-zero vector $\mathbf{s} \in \mathbb{Z}^m$ such that $\|\mathbf{s}\|_\infty \leq \beta$ and $\mathbf{H}\mathbf{s} = 0 \pmod{p}$.*

Ajtai's breakthrough result [Ajt96] and its subsequent improvements (e.g. [MR07]) showed that if one can solve SIS in the average case, then one can also solve the approximate Shortest Independent Vector Problem (SIVP) in every lattice.

Theorem 2.7 ([MR07,GPV08,MP13]). *For any $\beta > 0$, modulus $p \geq \beta\sqrt{mn}^{\Omega(1)}$ with at most $n^{O(1)}$ factors less than β , solving the $\text{SIS}_{p,n,m,\beta}^\infty$ problem (on the average, with nonnegligible probability $n^{-\Omega(1)}$) is at least as hard as solving SIVP_γ in the worst case on any n -dimensional lattice within a factor $\gamma = \max\{1, \beta^2\sqrt{m}/p\} \cdot \tilde{O}(\beta\sqrt{nm})$.*

In particular, for any constant $\epsilon > 0$, $\beta \leq n^\epsilon$, and $p \geq \beta\sqrt{mn}^\epsilon$, $\text{SIS}_{p,n,m,\beta}^\infty$ is hard on average under the assumption that SIVP_γ is hard in the worst case for $\gamma = \tilde{O}(\beta\sqrt{nm})$.

2.3 Codes and the Small Codeword problem

Definition 2.8. *In the Small Codeword ($\text{SC}_{n,m,\beta}$) problem, one is given a matrix $\mathbf{H} \in \mathbb{Z}_2^{n \times m}$ and a positive integer β , and is asked to find a non-zero vector $\mathbf{s} \in \mathbb{Z}_2^m$ such that $\|\mathbf{s}\|_1 \leq \beta$ and $\mathbf{H}\mathbf{s} = 0 \pmod{2}$.*

In this paper we will be interested in the above problem where m is a small polynomial in n and $\beta = \Theta(n)$. If β is too big (e.g. $n/2$), then the problem is trivially solved by Gaussian elimination, but if $\beta < n/4$ (or really $\beta < n/c$ for any constant $c > 2$), the best algorithm seems to be the Generalized Birthday attack where one only has few samples, and so it runs in time $2^{\Omega(n/\log \log n)}$.

2.4 RING-SIS in the Ring $\mathbb{Z}_p[x]/\langle x^n + 1 \rangle$

Let R be the ring $\mathbb{Z}_p[x]/\langle x^n + 1 \rangle$ where n is a power of 2. Elements in R have a natural representation as polynomials of degree $n - 1$ with coefficients in the range $\left[-\frac{p-1}{2}, \frac{p-1}{2}\right]$. For an element $\mathbf{a} = a_0 + a_1x + \dots + a_{n-1}x^{n-1} \in R$, we define $\|\mathbf{a}\|_\infty = \max_i(|a_i|)$. Similarly, for a tuple $(\mathbf{a}_1, \dots, \mathbf{a}_m) \in R^m$, we define $\|(\mathbf{a}_1, \dots, \mathbf{a}_m)\|_\infty = \max_i(\|\mathbf{a}_i\|_\infty)$. Notice that $\|\cdot\|_\infty$ is not exactly a norm because $\|\alpha\mathbf{a}\|_\infty \neq \alpha\|\mathbf{a}\|_\infty$ for all integers α (because of the reduction modulo p), but it still holds true that $\|\mathbf{a} + \mathbf{b}\|_\infty \leq \|\mathbf{a}\|_\infty + \|\mathbf{b}\|_\infty$ and $\|\alpha\mathbf{a}\|_\infty \leq \alpha\|\mathbf{a}\|_\infty$. It can also be easily checked that for any $\mathbf{a}, \mathbf{b} \in R$, we have $\|\mathbf{a}\mathbf{b} \pmod{x^n + 1}\|_\infty \leq n\|\mathbf{a}\|_\infty \cdot \|\mathbf{b}\|_\infty$ and if \mathbf{a} only had w non-zero coefficients, then $\|\mathbf{a}\mathbf{b} \pmod{x^n + 1}\|_\infty \leq w\|\mathbf{a}\|_\infty\|\mathbf{b}\|_\infty$.

Definition 2.9. *Let R be the ring $\mathbb{Z}_p[x]/\langle x^n + 1 \rangle$. In the Small Integer Solution over Rings problem ($\text{RING-SIS}_{p,n,m,\beta}$), one is given a matrix $\mathbf{H} \in R^{1 \times m}$ and is asked to find a non-zero vector $\mathbf{s} \in R^m$ such that $\|\mathbf{s}\|_\infty \leq \beta$ and $\mathbf{H}\mathbf{s} = 0 \pmod{p}$.*

Theorem 2.10 ([LM06]). *For $m > \log p / \log(2\beta)$, $\gamma = 16\beta \cdot m \cdot n \log^2 n$, and $p \geq \frac{\gamma\sqrt{n}}{4\log n}$, solving the $\text{RING-SIS}_{p,n,m,\beta}$ problem in uniformly random matrices in $R^{1 \times m}$ is at least as hard as solving SVP_γ^∞ in any ideal in the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$.*

3 The One-Time Signature Scheme

In this section we present our one-time signature scheme. The security of the scheme is based on the collision resistance properties of a linear (e.g., lattice or coding based) hash function. The scheme can be instantiated with a number of different hash functions, leading to digital signature schemes that are ultimately based on the worst-case hardness of approximating lattice problems in various lattice families (ranging from arbitrary lattices, to ideal lattices,) or similar (average-case) problems from coding theory.

The scheme is parametrized by

Key Generation	Sign(\mathbf{m})	Verify(\mathbf{m}, \mathbf{s})
Secret Key: $\mathbf{K} \in \mathcal{K}$ Public Key: $\mathbf{H} \in \mathcal{H}$, $\hat{\mathbf{K}} = \mathbf{H}\mathbf{K}$	Signature $\mathbf{s} = \mathbf{K}\mathbf{m}$	Check that $\mathbf{m} \in \mathcal{M}$, $\mathbf{s} \in \mathcal{S}$, and $\mathbf{H}\mathbf{s} = \hat{\mathbf{K}}\mathbf{m}$

Fig. 1. The One-Time Signature Scheme

- integers m, k, n ,
- a ring R
- Subsets of matrices $\mathcal{H} \subseteq R^{n \times m}$, $\mathcal{K} \subseteq R^{m \times k}$, and vectors $\mathcal{M} \subseteq R^k$, $\mathcal{S} \subseteq R^m$.

The parameters should satisfy certain properties for the scheme to work and be secure, but before stating the properties, we describe how the sets of matrices are used to define the one-time signature scheme.

The scheme is defined by the following procedures (also see Figure 1):

- **Setup:** A random matrix $\mathbf{H} \in \mathcal{H} \subseteq R^{n \times m}$ is chosen and can be shared by all users. The matrix \mathbf{H} will be used as a hash function mapping (a subset of) R^m to R^n , and extended to matrices in $R^{m \times k}$ in the obvious way.¹
- **Key Generation:** A secret key $\mathbf{K} \in \mathcal{K} \subseteq R^{m \times k}$ is chosen uniformly at random. The corresponding public key $\hat{\mathbf{K}} = \mathbf{H}\mathbf{K} \in \hat{\mathcal{K}} = R^{n \times k}$ is obtained by hashing the secret key using \mathbf{H} .
- **Signing:** Messages are represented as vectors $\mathbf{m} \in \mathcal{M} \subseteq R^k$. On input secret key \mathbf{K} and message $\mathbf{m} \in \mathcal{M}$, the signing algorithm outputs $\mathbf{s} = \mathbf{K}\mathbf{m} \in R^m$.
- **Verification:** The verification algorithm, on input public key $\hat{\mathbf{K}}$, message \mathbf{m} and signature \mathbf{s} , checks that $\mathbf{s} \in \mathcal{S}$ and $\mathbf{H}\mathbf{s} = \hat{\mathbf{K}}\mathbf{m}$.

The correctness and security of the scheme is based on the following three properties:

1. (*Closure*) $\mathbf{K}\mathbf{m} \in \mathcal{S}$ for all $\mathbf{K} \in \mathcal{K}$ and $\mathbf{m} \in \mathcal{M}$.
2. (*Collision Resistance*) The function family $\{\mathbf{H}: \mathcal{S} \rightarrow R^n \mid \mathbf{H} \in \mathcal{H}\}$ is collision resistant, i.e., any efficient adversary, on input a randomly chosen \mathbf{H} , outputs a collision ($\mathbf{s} \neq \tilde{\mathbf{s}}, \mathbf{H}\mathbf{s} = \mathbf{H}\tilde{\mathbf{s}}$) with at most negligible probability.
3. ((ϵ, δ) -*Hiding*) For any $\mathbf{H} \in \mathcal{H}$, $\mathbf{K} \in \mathcal{K}$ and $\mathbf{m} \in \mathcal{M}$, let

$$\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) = \{\tilde{\mathbf{K}} \in \mathcal{K} : \mathbf{H}\tilde{\mathbf{K}} = \mathbf{H}\mathbf{K} \wedge \tilde{\mathbf{K}}\mathbf{m} = \mathbf{K}\mathbf{m}\}$$

be the set of secret keys that are consistent with the public key $\mathbf{H}\mathbf{K}$ and \mathbf{m} -signature $\mathbf{K}\mathbf{m}$ associated to \mathbf{K} . The scheme is (ϵ, δ) -*Hiding* if for any $\mathbf{H} \in \mathcal{H}$,

$$\Pr_{\mathbf{K} \in \mathcal{K}} [\forall \mathbf{m} \neq \tilde{\mathbf{m}}, |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}})| \leq \epsilon |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})|] \geq \delta.$$

In the analysis of the schemes in this paper we will only use the (ϵ, δ) -*Hiding* property with $\epsilon = 1/2$ and $\delta \approx 1$. For notational simplicity, if a scheme is (ϵ, δ) -*Hiding* for some $\delta = 1 - n^{-\omega(1)}$

¹ To make sure that someone does not choose \mathbf{H} with a planted trapdoor, it could be demanded that $\mathbf{H} = \text{XOF}(x)$ where XOF is some extendable output function (e.g. SHAKE) and x is a public seed.




Fig. 2. (ϵ, δ) -*Hiding* Property. If $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$ (respectively $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}})$) is defined to be the set of secret keys consistent with the public key \mathbf{HK} and signature \mathbf{Km} (respectively $\mathbf{K}\tilde{\mathbf{m}}$), then we do not want the grey region to be an overwhelming fraction of $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$.

overwhelmingly close to 1, then we simply say that it is $(\epsilon$ -*Hiding*). So, the signature schemes analyzed in this paper can be described as being $(\frac{1}{2}$ -*Hiding*).

The (*Closure*) and (*Collision Resistance*) properties are self-explanatory, whereas the (ϵ, δ) -*Hiding* one could use some motivation. For concreteness, let us use $(\frac{1}{2}$ -*Hiding*) as an example. Recall from our proof sketch in Section 1.1 that we can find a collision to the challenge hash function \mathbf{H} if the adversary returns a signature $\tilde{\mathbf{s}}$ of a message $\tilde{\mathbf{m}}$ such that $\tilde{\mathbf{s}} \neq \mathbf{K}\tilde{\mathbf{m}}$, where \mathbf{K} is our chosen secret key with which we signed the message \mathbf{m} . If the adversary is to output a signature $\tilde{\mathbf{s}}$ such that $\tilde{\mathbf{s}} = \mathbf{K}\tilde{\mathbf{m}}$, then \mathbf{K} must be in the grey intersection in Figure 2. The $(\frac{1}{2}$ -*Hiding*) condition says that with probability ≈ 1 , this grey region will be at most half the size of the set $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$. Since after seeing the signature of \mathbf{m} , the secret key is equally likely to be anywhere in $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$, it can be shown that even an all-powerful adversary has at most an $\frac{1}{2}$ chance of producing a signature $\tilde{\mathbf{s}}$ which equals $\mathbf{K}\tilde{\mathbf{m}}$. Thus the reduction's probability of outputting a valid collision is $1 - \frac{1}{2} = \frac{1}{2}$.

Also note that the (*Hiding*) property precludes the message space \mathcal{M} from containing both \mathbf{m} and $c \cdot \mathbf{m}$, for any $c \in R$. Intuitively, this should be disallowed because otherwise an adversary who sees the signature \mathbf{s} of message \mathbf{m} could output a forgery $\tilde{\mathbf{s}} = c \cdot \mathbf{s}$ on the message $\tilde{\mathbf{m}} = c \cdot \mathbf{m}$. And indeed, this cannot happen if the scheme satisfies the (ϵ, δ) -*Hiding* property for any $\epsilon < 1$ and $\delta > 0$. In fact, if \mathbf{m} and $\tilde{\mathbf{m}} = c \cdot \mathbf{m}$ are both in \mathcal{M} , then one can see that $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \subseteq \mathcal{D}_{\mathbf{H}}(\mathbf{K}, c \cdot \mathbf{m})$. Therefore $|\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, c \cdot \mathbf{m})| = |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})|$ and the (ϵ, δ) -*Hiding* property cannot hold for $\epsilon < 1$ and $\delta > 0$. Since \mathbf{m} is a vector, the most natural way to enforce that $c \cdot \mathbf{m}$ cannot be in \mathcal{M} (which is a necessary condition a secure scheme needs to have) is to force all vectors in \mathcal{M} to have 1 as their last component. This is in fact how the message space is constructed in the examples in Section 4.

Lemma 3.1. *If the (*Closure*) property holds, then the scheme is correct, i.e., the verification algorithm always accepts signatures produced by the legitimate signer.*

Proof. It immediately follows from the definition of the (*Closure*) property and the signature verification algorithm.

Theorem 3.2. *Assume the signature scheme satisfies the (ϵ, δ) -*Hiding* and (*Closure*) properties. If there is an adversary \mathcal{A} that succeeds in breaking the strong unforgeability of the one-time signature scheme with probability γ , then there exists an algorithm that can break the (*Collision Resistance*)*

property with probability at least $(\gamma + \delta - 1) \cdot (1 - \epsilon)/(2 - \epsilon)$ in essentially the same running time as the forgery attack.

In particular, if the (**Closure**), (**Collision Resistance**) and (ϵ -**Hiding**) properties hold true for any constant $\epsilon < 1$, then the one-time signature scheme is strongly unforgeable.

Proof. Let \mathcal{A} be an efficient forger that can break the one-time signature scheme with probability γ . We use \mathcal{A} to build an attacker to the collision resistance of \mathbf{H} that works as follows:

1. Given an $\mathbf{H} \in \mathcal{H}$, pick a uniformly-random secret key $\mathbf{K} \in \mathcal{K}$.
2. Send the public key $(\mathbf{H}, \mathbf{HK})$ to \mathcal{A} .
3. Obtain query message $\mathbf{m} \leftarrow \mathcal{A}(\mathbf{H}, \mathbf{HK})$.
4. Check that $\mathbf{m} \in \mathcal{M}$ and send the signature $\mathbf{s} = \mathbf{K}\mathbf{m}$ to \mathcal{A} .
5. Obtain a candidate forgery $(\tilde{\mathbf{m}}, \tilde{\mathbf{s}}) \leftarrow \mathcal{A}(\mathbf{H}, \mathbf{HK}, \mathbf{s})$.
6. Output $(\mathbf{K}\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$ as a candidate collision to \mathbf{H} .

By the (**Closure**) property, we may assume that $\mathbf{s}, \mathbf{K}\tilde{\mathbf{m}} \in \mathcal{S}$ are valid signatures. In the rest of the proof we assume without loss of generality that \mathcal{A} always outputs syntactically valid messages $\mathbf{m}, \tilde{\mathbf{m}} \in \mathcal{M}$ and a valid signature $\tilde{\mathbf{s}} \in \mathcal{S}$ satisfying $\mathbf{H}\tilde{\mathbf{s}} = \mathbf{H}\mathbf{K}\tilde{\mathbf{m}}$. (An adversary can always be modified to achieve this property, while preserving the success probability of the attack, by checking that $(\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$ is a valid message/signature pair, and if not, output (\mathbf{m}, \mathbf{s}) .) Under these conventions, the collision finding algorithm always outputs a valid collision, and it is successful if and only if the collision is nontrivial, i.e., the following event

$$\mathbf{K}\tilde{\mathbf{m}} \neq \tilde{\mathbf{s}} \quad (\mathbf{Collision})$$

is satisfied. Similarly, the forger \mathcal{A} always outputs a valid message-signature pair and it is successful if and only if the pair is nontrivial, i.e., the condition

$$(\mathbf{m}, \mathbf{s}) \neq (\tilde{\mathbf{m}}, \tilde{\mathbf{s}}) \quad (\mathbf{Forgery})$$

holds true.

We know by assumption that this event has probability $\Pr\{(\mathbf{Forgery})\} = \gamma$. We need to bound the probability of (**Collision**). To this end, we replace step 6. in the above experiment with the following additional steps

7. Choose a random bit $b \in \{0, 1\}$ with $\Pr\{b = 0\} = (1 - \epsilon)/(2 - \epsilon)$, and $\Pr\{b = 1\} = 1 - \Pr\{b = 0\} = 1/(2 - \epsilon)$.
8. If $b = 0$, then set $\tilde{\mathbf{K}} = \mathbf{K}$, and otherwise choose $\tilde{\mathbf{K}}$ uniformly at random from the set $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$.
9. Output $(\tilde{\mathbf{K}}\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$ as an candidate collision to \mathbf{H} .

Notice that the set $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$ is always nonempty because it contains \mathbf{K} . So, step 8. is well defined. The success of the extended experiment is defined by the event

$$\tilde{\mathbf{K}}\tilde{\mathbf{m}} \neq \tilde{\mathbf{s}}. \quad (\mathbf{Collision}')$$

Notice that this condition is identical to (**Collision**), except for the use of the new key $\tilde{\mathbf{K}}$ instead of the original one \mathbf{K} . We remark that these additional steps are just part of a mental experiment used in the analysis, and they are not required to be efficiently computable.

We observe that the output of \mathcal{A} only depends on its random coins and the messages $\mathbf{H}, \mathbf{HK}, \mathbf{Km}$ received from the challenger. Moreover, by definition, $\mathcal{D}_{\mathbf{H}}$ is precisely the set of keys $\tilde{\mathbf{K}}$ that are consistent with these messages $\mathbf{H}, \mathbf{H}\tilde{\mathbf{K}} = \mathbf{HK}, \tilde{\mathbf{K}}\mathbf{m} = \mathbf{Km}$. So, the conditional distribution of \mathbf{K} given $\mathbf{H}, \mathbf{HK}, \mathbf{Km}$ is precisely the uniform distribution over $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$. This proves that the $(\tilde{\mathbf{K}}\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$ is distributed identically to the output $(\mathbf{K}\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$ of the original collision finding algorithm. In particular, the original and modified experiments have exactly the same success probability $\Pr\{(\mathbf{Collision}')\} = \Pr\{(\mathbf{Collision})\}$ at finding a nontrivial collision. So, in what follows, we will bound the probability of $(\mathbf{Collision}')$ rather than $(\mathbf{Collision})$.

In order to bound the probability of $(\mathbf{Collision}')$, we break the corresponding event into three components:

$$\begin{aligned} \Pr\{(\mathbf{Collision}')\} &= \Pr\{(\mathbf{Collision}') \wedge (\mathbf{m} = \tilde{\mathbf{m}})\} \\ &\quad + \Pr\{(\mathbf{Collision}') \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge (\mathbf{Collision})\} \\ &\quad + \Pr\{(\mathbf{Collision}') \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision})\} \end{aligned}$$

and observe that the bit b is chosen independently of $\mathbf{m}, \tilde{\mathbf{m}}, \mathbf{s}, \tilde{\mathbf{s}}$ and \mathbf{K} , because only $\tilde{\mathbf{K}}$ depends on b . In particular, the events $(b = 0)$ and $(b = 1)$ are statistically independent from $(\mathbf{m} = \tilde{\mathbf{m}})$, $(\mathbf{m} \neq \tilde{\mathbf{m}})$, the original $(\mathbf{Collision})$ event $\mathbf{K}\tilde{\mathbf{m}} \neq \tilde{\mathbf{s}}$, and the $(\mathbf{Forgery})$ event $(\mathbf{m}, \mathbf{s}) \neq (\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$.

First we consider the simple case when $\mathbf{m} = \tilde{\mathbf{m}}$, i.e., the adversary attempts to forge a different signature $\tilde{\mathbf{s}} \neq \mathbf{s}$ for the same message $\tilde{\mathbf{m}} = \mathbf{m}$. Formally, if $(\mathbf{Forgery}) \wedge (\mathbf{m} = \tilde{\mathbf{m}}) \wedge (b = 0)$ holds true, then it must be that $\mathbf{s} \neq \tilde{\mathbf{s}}, \tilde{\mathbf{K}} = \mathbf{K}$ and²

$$\tilde{\mathbf{K}}\tilde{\mathbf{m}} = \mathbf{Km} = \mathbf{s} \neq \tilde{\mathbf{s}}.$$

But $\tilde{\mathbf{K}}\tilde{\mathbf{m}} \neq \tilde{\mathbf{s}}$ is precisely the definition of $(\mathbf{Collision}')$. So, $(\mathbf{Forgery}) \wedge (\mathbf{m} = \tilde{\mathbf{m}}) \wedge (b = 0)$ implies $(\mathbf{Collision}') \wedge (\mathbf{m} = \tilde{\mathbf{m}})$, and

$$\begin{aligned} \Pr\{(\mathbf{Collision}') \wedge (\mathbf{m} = \tilde{\mathbf{m}})\} &\geq \Pr\{(\mathbf{Forgery}) \wedge (\mathbf{m} = \tilde{\mathbf{m}}) \wedge (b = 0)\} \\ &= \Pr\{(\mathbf{Forgery}) \wedge (\mathbf{m} = \tilde{\mathbf{m}})\} \cdot \frac{1 - \epsilon}{2 - \epsilon}. \end{aligned}$$

We now move on to the case where $\mathbf{m} \neq \tilde{\mathbf{m}}$ and the $(\mathbf{Collision})$ nontriviality property $\tilde{\mathbf{s}} \neq \mathbf{K}\tilde{\mathbf{m}}$ are satisfied, i.e., the adversary produces a forgery on a different message $\tilde{\mathbf{m}}$ that leads to a collision in the original game. If $(\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge (\mathbf{Collision}) \wedge (b = 0)$, then $\tilde{\mathbf{K}} = \mathbf{K}$, and the $(\mathbf{Collision}')$ property holds true because $(\mathbf{Collision})$ and $(\mathbf{Collision}')$ are the same for $\tilde{\mathbf{K}} = \mathbf{K}$. Therefore,

$$\begin{aligned} &\Pr\{(\mathbf{Collision}') \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge (\mathbf{Collision})\} \\ &\geq \Pr\{(\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge (\mathbf{Collision}) \wedge (b = 0)\} \\ &= \Pr\{(\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge (\mathbf{Collision})\} \cdot \Pr\{b = 0\} \\ &\geq \Pr\{(\mathbf{Forgery}) \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge (\mathbf{Collision})\} \cdot \frac{1 - \epsilon}{2 - \epsilon}. \end{aligned}$$

We remark that the last inequality is actually an equality because $\mathbf{m} \neq \tilde{\mathbf{m}}$ implies the $(\mathbf{Forgery})$ property $(\mathbf{m}, \mathbf{s}) \neq (\tilde{\mathbf{m}}, \tilde{\mathbf{s}})$, but this makes no difference in our proof.

² Notice that the following equality holds true also when $b = 1$, because $\tilde{\mathbf{K}}\mathbf{m} = \mathbf{Km}$ for all $\tilde{\mathbf{K}} \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$. But this is not used in this step of the proof.

For the last component, consider the set $\mathcal{X}_{\mathbf{H}} \subseteq \mathcal{K}$ of all secret keys \mathbf{K} satisfying the (ϵ -**Hiding**) property

$$\mathcal{X}_{\mathbf{H}} = \{\mathbf{K} \in \mathcal{K} : \forall \mathbf{m} \neq \tilde{\mathbf{m}}, |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}})| \leq \epsilon |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})|\}.$$

We know that, by the (ϵ, δ -**Hiding**) assumption, for all \mathbf{H} we have $\Pr\{\mathbf{K} \in \mathcal{X}_{\mathbf{H}}\} \geq \delta$. Using the independence of b , and a union bound, we see that the event

$$(\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision}) \wedge (\mathbf{K} \in \mathcal{X}_{\mathbf{H}}) \wedge (b = 1) \tag{\mathcal{X}}$$

has probability

$$\begin{aligned} \Pr\{(\mathcal{X})\} &= \Pr\{b = 1\} \cdot \Pr\{(\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision}) \wedge (\mathbf{K} \in \mathcal{X}_{\mathbf{H}})\} \\ &\geq \frac{\Pr\{(\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision})\} - \Pr\{\mathbf{K} \notin \mathcal{X}_{\mathbf{H}}\}}{2 - \epsilon} \\ &\geq \frac{\Pr\{(\mathbf{Forgery}) \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision})\} - 1 + \delta}{2 - \epsilon}. \end{aligned}$$

Next, notice that the event (\mathcal{X}) implies $\neg(\mathbf{Collision})$, i.e., $\tilde{\mathbf{s}} = \mathbf{K}\tilde{\mathbf{m}}$. So, given (\mathcal{X}) , the $(\mathbf{Collision}')$ event $\tilde{\mathbf{K}}\tilde{\mathbf{m}} \neq \tilde{\mathbf{s}}$ is equivalent to $\tilde{\mathbf{K}}\tilde{\mathbf{m}} \neq \mathbf{K}\tilde{\mathbf{m}}$. Therefore, for all $\tilde{\mathbf{K}}$ such that $\mathbf{H}\tilde{\mathbf{K}} = \mathbf{H}\mathbf{K}$ (in particular, for all $\tilde{\mathbf{K}} \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})$), and conditioned on (\mathcal{X}) , the $(\mathbf{Collision}')$ property is satisfied if and only if $\tilde{\mathbf{K}} \notin \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}})$, i.e.,

$$\begin{aligned} \Pr\{(\mathbf{Collision}') \mid (\mathcal{X})\} &= \Pr\{\tilde{\mathbf{K}} \notin \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}}) \mid (\mathcal{X})\} \\ &= 1 - \Pr\{\tilde{\mathbf{K}} \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}}) \mid (\mathcal{X})\} \\ &\geq 1 - \max_{\mathbf{H}, \mathbf{K} \in \mathcal{X}_{\mathbf{H}}, \mathbf{m} \neq \tilde{\mathbf{m}}} \frac{|\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \tilde{\mathbf{m}})|}{|\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})|} \\ &\geq 1 - \epsilon \end{aligned}$$

where, in the last inequality we have used the definition of $\mathcal{X}_{\mathbf{H}}$. We can now compute

$$\begin{aligned} &\Pr\{(\mathbf{Collision}') \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision})\} \\ &\geq \Pr\{(\mathbf{Collision}') \wedge (\mathcal{X})\} \\ &= \Pr\{(\mathcal{X})\} \cdot \Pr\{(\mathbf{Collision}') \mid (\mathcal{X})\} \\ &\geq (\Pr\{(\mathbf{Forgery}) \wedge (\mathbf{m} \neq \tilde{\mathbf{m}}) \wedge \neg(\mathbf{Collision})\} - 1 + \delta) \cdot \frac{1 - \epsilon}{2 - \epsilon}. \end{aligned}$$

Adding up the three bounds gives

$$\Pr\{(\mathbf{Collision}')\} \geq \left(\Pr\{(\mathbf{Forgery})\} - 1 + \delta \right) \cdot \frac{1 - \epsilon}{2 - \epsilon} = (\gamma - 1 + \delta) \cdot \frac{1 - \epsilon}{2 - \epsilon}.$$

Finally, we observe that for any $\delta = 1 - n^{-\omega(1)}$ overwhelmingly close to 1 and constant $\epsilon < 1$, we have $(\gamma - 1 + \delta)(1 - \epsilon)/(2 - \epsilon) = O(\gamma - n^{-\omega(1)})$. So, if the $(\mathbf{Closure})$, (ϵ -**Hiding**) and $(\mathbf{Collision Resistance})$ properties hold true, then $\Pr\{(\mathbf{Collision}')\}$ and γ are both negligible, and the signature scheme is strongly unforgeable. \square

4 Instantiation With Lattices and Codes

In this section we describe instantiations of our general one-time signature scheme based on various classes of lattices and linear codes over finite fields. All schemes are proved secure showing that they satisfy the [Closure], [$\frac{1}{2}$ -Hiding] and [Collision Resistance] properties, and then using Theorem 3.2. Throughout this section, λ is a statistical security parameter that can be set, for example, to $\lambda = 128$. The following simple lemma is used in the analysis of all schemes.

Lemma 4.1. *Let $h : X \rightarrow Y$ be a deterministic function where X and Y are finite sets and $|X| \geq 2^\lambda |Y|$. If x is chosen uniformly at random from X , then with probability at least $1 - 2^{-\lambda}$, there exists another $x' \in X$ such that $h(x) = h(x')$.*

Proof. There are at most $|Y| - 1$ elements x in X for which there is no x' such that $h(x) = h(x')$. Therefore the probability that a randomly chosen x does have a corresponding x' for which $h(x) = h(x')$ is at least $(|X| - |Y| + 1)/|X| = 1 - |Y|/|X| + 1/|X| > 1 - 2^{-\lambda}$. \square

4.1 One-time signature as hard as SIS

The lattice based signature scheme is defined by the sets in Figure 3 parametrized by integers

$$\begin{aligned}
 R &= \mathbb{Z}_p \\
 \mathcal{H} &= R^{n \times m} \\
 \mathcal{K} &= \{\mathbf{K} \in R^{m \times k} : \|\mathbf{K}\|_\infty \leq b\} \\
 \mathcal{M} &\subseteq \{\mathbf{m} \in \{0, 1\}^k : \|\mathbf{m}\|_1 = w\} \\
 \mathcal{S} &= \{\mathbf{s} \in R^m : \|\mathbf{s}\|_\infty \leq wb\}.
 \end{aligned}$$

Fig. 3. Instantiation of the one-time signature scheme based on general lattices. The sets are parametrized by the integers n, m, k, p, w, b .

n, m, k, p, w , and b which should satisfy certain relationships. The size of the message space is $\binom{k}{w}$, and so we need to set k and w so that this number is large enough. The choice of k and w offers a trade-off between security and efficiency. Specifically, the size of both secret and public keys is linear in k , so smaller values of k result in more efficient schemes. On the other hand, larger values of w result in stronger security assumptions. For proving the security of our scheme based on the SIS problem, we also need to have $b = \left\lceil \frac{p^{n/m} 2^{\lambda/m} - 1}{2} \right\rceil$. For concreteness, the reader may assume $m = \lceil (\lambda + n \log_2 p) / \log_2 3 \rceil$, which allows to set $b = 1$. In practice, larger values of b may also be interesting, as they allow for smaller values of m . Again, this offers a trade-off between security and efficiency, where smaller values of m result in shorter signatures, while smaller values of b give better security guarantees.

Additionally, if we would like to preserve the connection between average-case SIS and the worst-case SIVP problem from Theorem 2.7, then we will also need to have $p \geq 2wb\sqrt{mn}^{\Omega(1)}$.

We now proceed to show that as defined above, our scheme satisfies the [Closure], [Collision Resistance], and [$\frac{1}{2}$ -Hiding] properties defined in Section 3.

Lemma 4.2. *The [Closure] property holds.*

Proof. It is clear that for any secret key \mathbf{K} and message \mathbf{m} , we have $\|\mathbf{K}\mathbf{m}\|_\infty \leq \|\mathbf{K}\|_\infty \cdot \|\mathbf{m}\|_1 \leq wb$, and therefore $\mathbf{K}\mathbf{m} \in \mathcal{S}$.

Lemma 4.3. *The function family $\{\mathbf{H} : \mathcal{S} \rightarrow R^n \mid \mathbf{H} \in \mathcal{H}\}$ satisfies the [Collision Resistance] property based on the average-case hardness of the $\text{SIS}_{n,m,p,2wb}^\infty$ problem. Furthermore, if $p \geq 2wb\sqrt{mn}^{\Omega(1)}$, then the property is satisfied based on the worst-case hardness of SIVP_γ in n -dimensional lattices for $\gamma = \tilde{O}(wb\sqrt{nm}) \cdot \max\{1, 4w^2b^2\sqrt{m}/p\}$.*

Proof. The first part of the claim follows simply because if one can find $\mathbf{x} \neq \mathbf{x}' \in \mathcal{S}$ for a random \mathbf{H} from \mathcal{H} such that $\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{x}'$, then one has that $\mathbf{H}(\mathbf{x} - \mathbf{x}') = 0$ and $\|\mathbf{x} - \mathbf{x}'\|_\infty \leq 2wb$. The connection to SIVP_γ follows directly from Theorem 2.7. \square

Before analyzing the $[\frac{1}{2}\text{-Hiding}]$ property, we prove a simple lemma that states that with very high probability, for a randomly-chosen secret key $\mathbf{K} \in \mathcal{K}$, there are other “similar-looking” possible secret keys \mathbf{K}' such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$.

Lemma 4.4. *Let $b = \left\lceil \frac{p^{n/m}2^{\lambda/m} - 1}{2} \right\rceil$. For every $\mathbf{H} \in \mathcal{H}$, if \mathbf{K} is chosen uniformly at random from \mathcal{K} , then with probability at least $1 - k2^{-\lambda}$, there exists a key $\mathbf{K}' \in \mathcal{K}$ such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$ and $\mathbf{K}' \neq \mathbf{K}$ differ in every column.*

Proof. Consider \mathbf{H} as a function mapping from domain $X = \{-b, \dots, b\}^m$ to range $Y = \mathbb{Z}_p^n$. Notice that by our choice of b , we have $|X| = (2b + 1)^m \geq p^n 2^\lambda$; and $|Y|$ is exactly p^n . By Lemma 4.1, we know that for a randomly chosen vector $\mathbf{x} \in X$, with probability at least $1 - 2^{-\lambda}$, there is another vector $\mathbf{x}' \in X$ such that $\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{x}'$. Thus we have that for any particular column \mathbf{K}_j , with probability at least $1 - 2^{-\lambda}$, there exists a column \mathbf{K}'_j such that $\mathbf{H}\mathbf{K}_j = \mathbf{H}\mathbf{K}'_j$ and $\mathbf{K}_j \neq \mathbf{K}'_j$. Applying the union bound, we get that with probability at least $1 - k2^{-\lambda}$ this is true for every column $j = 1, \dots, k$, giving a key \mathbf{K}' such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$ and $\mathbf{K}_j \neq \mathbf{K}'_j$ for all j . \square

Lemma 4.5. *Let $b = \left\lceil \frac{p^{n/m}2^{\lambda/m} - 1}{2} \right\rceil$ as in Lemma 4.4. Then the scheme satisfies the $[\frac{1}{2}\text{-Hiding}]$ property.*

Proof. Fix a hash function $\mathbf{H} \in \mathcal{H}$. We know that with probability at least $1 - k2^{-\lambda}$, a randomly-chosen key \mathbf{K} has the property from Lemma 4.4, i.e., there is another key \mathbf{K}' such that $\mathbf{H}\mathbf{K}' = \mathbf{H}\mathbf{K}$ and $\mathbf{K}'_j \neq \mathbf{K}_j$ for every $j = 1, \dots, k$. We now proceed to show that for any such key \mathbf{K} , and for any $\mathbf{m} \neq \mathbf{m}'$, we have

$$|\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')| \leq |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \setminus \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')|, \quad (1)$$

or, equivalently,

$$|\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')| \leq \frac{1}{2} \cdot |\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m})|,$$

which proves the lemma.

In order to prove (1), we give an injective function f from $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$ to $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \setminus \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$. Since $\mathbf{m}' \neq \mathbf{m}$, there must be a j such that the j^{th} coefficient is 0 in \mathbf{m} and is 1 in \mathbf{m}' . For any $\mathbf{X} \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$, we define $\mathbf{X}' = f(\mathbf{X})$ as follows:

1. $\mathbf{X}'_i = \mathbf{X}_i$ for all $i \neq j$
2. $\mathbf{X}'_j \in \{\mathbf{K}_j, \mathbf{K}'_j\} \setminus \{\mathbf{X}_j\}$. Notice that since $\mathbf{K}_j \neq \mathbf{K}'_j$, at least one of them is different from \mathbf{X}_j . If they are both different, then \mathbf{X}'_j can be chosen between them arbitrarily.

We need to show that $\mathbf{X}' \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \setminus \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$, and that f is injective.

For $\mathbf{X}' \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \setminus \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$, we need to verify the following three conditions: $\mathbf{H}\mathbf{X}' = \mathbf{H}\mathbf{K}$, $\mathbf{X}'\mathbf{m} = \mathbf{K}\mathbf{m}$ and $\mathbf{X}'\mathbf{m}' \neq \mathbf{K}\mathbf{m}'$, under the assumption that $\mathbf{H}\mathbf{X} = \mathbf{H}\mathbf{K}$, $\mathbf{X}\mathbf{m} = \mathbf{K}\mathbf{m}$ and $\mathbf{X}\mathbf{m}' = \mathbf{K}\mathbf{m}'$. For each $i = 1, \dots, k$, we have $\mathbf{X}'_i \in \{\mathbf{X}_i, \mathbf{K}_i, \mathbf{K}'_i\}$. Since $\mathbf{H}\mathbf{X} = \mathbf{H}\mathbf{K}$ and $\mathbf{H}\mathbf{K}' = \mathbf{H}\mathbf{K}$ (by our choice of \mathbf{K}'), we have $\mathbf{H}\mathbf{X}' = \mathbf{H}\mathbf{K}$, proving the first condition. The second condition $\mathbf{X}'\mathbf{m} = \mathbf{K}\mathbf{m}$ follows from the fact that $\mathbf{X}'\mathbf{m} = \mathbf{X}\mathbf{m}$ (because \mathbf{X}' and \mathbf{X} differ only in the j th column and $\mathbf{m}_j = 0$) and $\mathbf{X}\mathbf{m} = \mathbf{K}\mathbf{m}$. Similarly, the third condition $\mathbf{X}'\mathbf{m}' \neq \mathbf{K}\mathbf{m}'$ follows from the fact that $\mathbf{X}'\mathbf{m}' \neq \mathbf{X}\mathbf{m}'$ (because \mathbf{X}' and \mathbf{X} differ only in the j th column and $\mathbf{m}'_j = 1$) and $\mathbf{X}\mathbf{m}' = \mathbf{K}\mathbf{m}'$.

It remains to prove that f is injective. Assume for contradiction that $f(\mathbf{X}) = f(\mathbf{X}')$ for some $\mathbf{X} \neq \mathbf{X}'$ both in $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$. Then, by definition of f , $\mathbf{X}_i = \mathbf{X}'_i$ for all $i \neq j$. Therefore \mathbf{X}_j and \mathbf{X}'_j must differ. But then $\mathbf{X}\mathbf{m}' \neq \mathbf{X}'\mathbf{m}'$ because $\mathbf{m}'_j = 1$, and so they cannot both be in $\mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$. \square

Combining the previous lemmas, and Theorem 3.2, we obtain the following corollary.

Corollary 4.6. *For any $\epsilon > 0$, let $p \geq 2wb\sqrt{mn}^\epsilon$ and $b = \left\lceil \frac{p^{n/m}2^{\lambda/m}-1}{2} \right\rceil$. Then, the one-time signature scheme from Section 3, instantiated with the sets in Figure 3, is strongly unforgeable under the assumption that SIVP_γ is hard in the worst case for $\gamma = \tilde{O}(wb\sqrt{nm}) \max\{1, 2wb/n^\epsilon\}$.*

In particular, for $m = \lceil (\lambda + n \log_2 p) / \log_2 3 \rceil$, $b = 1$ and $p \geq 2w\sqrt{mn}^\epsilon$, the scheme is strongly unforgeable under the assumption that SIVP_γ is hard in the worst case for $\gamma = \tilde{O}(w\sqrt{nm}) \max\{1, 2w/n^\epsilon\}$.

4.2 One-time signature as hard as RING-SIS

Our one-time signature based on the RING-SIS problem from Definition 2.9 is parametrized by integers n, m, p, w , and b that must satisfy certain relationships. The integer n is assumed to be a power of 2, so that the polynomial $x^n + 1$ is irreducible over $\mathbb{Z}[x]$. The size of the message space \mathcal{M} is at most $\sum_{i \leq w} 2^i \binom{n}{i}$, and so we need to set n and w to sufficiently large integers. As usual, the choice of n and w offers a trade-off between efficiency and security. For proving the security of our scheme based on the RING-SIS problem, we also need to have $b = \lfloor (|\mathcal{M}|^{1/n} 2^{\lambda/n} p)^{1/m} \rfloor$ and $p > 8wb$. Notice that by choosing m large enough, one can set $b = 1$, but higher values of b can offer improved efficiency at the cost of stronger security assumptions. Additionally, if we would like to preserve the connection between average-case RING-SIS and the worst-case SVP problem in ideal lattices from Theorem 2.10, then we will also need to have $p = \omega(n^{1.5} m w b)$.

The scheme is parametrized by the sets in Figure 4. The message space is set to an appropriate

$$\begin{aligned}
R &= \mathbb{Z}_p[x] / \langle x^n + 1 \rangle \\
\mathcal{H} &= R^{1 \times m} \\
\mathcal{K} &= \{[\mathbf{k}_1, \mathbf{k}_2] \in R^{m \times 2} : \|\mathbf{k}_1\|_\infty \leq b, \|\mathbf{k}_2\|_\infty \leq wb\} \\
\mathcal{M} &\subseteq \{\mathbf{m} = [m_1, 1]^T \in R^2, \|m_1\|_\infty \leq 1, \|m_1\|_1 \leq w\} \\
\mathcal{S} &= \{\mathbf{s} \in R^m : \|\mathbf{s}\|_\infty \leq 2wb\}.
\end{aligned}$$

Fig. 4. Instantiation of the one-time signature scheme based on ideal lattices.

subset of all vectors with entries bounded by 1 in absolute value, and at most w non-zero entries.

The set \mathcal{M} should be chosen in such a way that messages can be efficiently encoded as elements of \mathcal{M} .

Lemma 4.7. *The function family $\{\mathbf{H} : \mathcal{S} \rightarrow R \mid \mathbf{H} \in \mathcal{H}\}$ satisfies the [Collision Resistance] property based on the average-case hardness of the RING-SIS $_{n,m,p,4wb}$ problem. Furthermore, for $\gamma = 64wbmn \log^2 n$ and $p \geq \frac{\gamma\sqrt{n}}{4 \log n}$, the property is satisfied based on the worst-case hardness of SVP_γ^∞ in all n -dimensional ideals of the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$.*

Proof. The first part of the claim follows simply because if one can find $\mathbf{x} \neq \mathbf{x}' \in \mathcal{S}$ for a random \mathbf{H} from \mathcal{H} such that $\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{x}'$, then one has that $\mathbf{H}(\mathbf{x} - \mathbf{x}') = 0$ and $\|\mathbf{x} - \mathbf{x}'\|_\infty \leq 4wb$. The connection to SVP_γ^∞ follows directly from Theorem 2.10. \square

Lemma 4.8. *The [Closure] property holds true.*

Proof. Notice that for any secret key $\mathbf{K} = [\mathbf{k}_1, \mathbf{k}_2]$ and message $\mathbf{m} = [m_1, 1]^T$,

$$\|\mathbf{K}\mathbf{m}\|_\infty = \|\mathbf{k}_1 m_1 + \mathbf{k}_2\|_\infty \leq \|\mathbf{k}_1 m_1\|_\infty + \|\mathbf{k}_2\|_\infty \leq wb + wb = 2wb.$$

\square

Lemma 4.9. *Let $b = \lfloor (|\mathcal{M}|^{1/n} 2^{\lambda/n} p)^{1/m} \rfloor$. For every $\mathbf{H} \in \mathcal{H}$, if \mathbf{K} is chosen uniformly at random from \mathcal{K} , then with probability at least $1 - 2^{-\lambda}$, for every message $\mathbf{m} \in \mathcal{M}$ there is another $\mathbf{K}' \in \mathcal{K}$ such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$ and $\mathbf{K}\mathbf{m} = \mathbf{K}'\mathbf{m}$.*

Proof. For any \mathbf{H} and \mathbf{m} , consider (\mathbf{H}, \mathbf{m}) as a function that maps any element \mathbf{K} in \mathcal{K} to the ordered pair $(\mathbf{H}\mathbf{K}, \mathbf{K}\mathbf{m})$. We will first show that the domain size of this function is at least $|\mathcal{M}| \cdot 2^\lambda$ times larger than its range. The domain size of this function is exactly $|\mathcal{K}| = (2b+1)^{mn} \cdot (2wb+1)^{mn}$. To bound the size of the range, we first notice that by Lemma 4.8 we have $\|\mathbf{K}\mathbf{m}\|_\infty \leq 2wb$. Therefore, the number of possibilities for $\mathbf{K}\mathbf{m}$ is at most $(4wb+1)^{mn}$. We then notice that while there are p^{2n} possibilities for $\mathbf{H}\mathbf{K} = [\mathbf{H}\mathbf{k}_1, \mathbf{H}\mathbf{k}_2]$ in general, if we have already fixed \mathbf{H} , \mathbf{m} , $\mathbf{H}\mathbf{k}_1$, and $\mathbf{K}\mathbf{m}$, then $\mathbf{H}\mathbf{k}_2 = \mathbf{H}\mathbf{K}\mathbf{m} - \mathbf{H}\mathbf{k}_1 m_1$ is completely determined. Thus, there are only at most $(4wb+1)^{mn} \cdot p^n$ possibilities for $(\mathbf{H}\mathbf{K}, \mathbf{K}\mathbf{m})$. Therefore the ratio of the sizes of the domain and range of the function (\mathbf{H}, \mathbf{m}) is at least

$$\frac{(2b+1)^{mn} \cdot (2wb+1)^{mn}}{(4wb+1)^{mn} \cdot p^n} > \frac{(2b+1)^{mn} \cdot (2wb+1)^{mn}}{(4wb+2)^{mn} \cdot p^n} = \left(\frac{(b + \frac{1}{2})^m}{p} \right)^n.$$

Using $b = \lfloor (|\mathcal{M}|^{1/n} 2^{\lambda/n} p)^{1/m} \rfloor \geq (|\mathcal{M}|^{1/n} 2^{\lambda/n} p)^{1/m} - \frac{1}{2}$, we get that the ratio is at least $|\mathcal{M}| \cdot 2^\lambda$. Applying Lemma 4.1, we obtain that with probability at least $1 - 2^{-\lambda}/|\mathcal{M}|$ over the random choice of $\mathbf{K} \in \mathcal{K}$, there exists another $\mathbf{K}' \in \mathcal{K}$ such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$ and $\mathbf{K}\mathbf{m} = \mathbf{K}'\mathbf{m}$. Applying the union bound over all messages in \mathcal{M} concludes the proof. \square

Lemma 4.10. *Let $b = \lfloor (|\mathcal{M}|^{1/n} 2^{\lambda/n} p)^{1/m} \rfloor$ and $p > 8wb$. Then the scheme satisfies the $\lfloor \frac{1}{2} \rfloor$ -Hiding property.*

Proof. Fix \mathbf{H} . By Lemma 4.9, we know that with probability of at least $1 - 2^{-\lambda}$ over the random choice of \mathbf{K} , for every message \mathbf{m} , the size of the set $\mathcal{D}_\mathbf{H}(\mathbf{K}, \mathbf{m})$ is at least 2. To complete the proof, we will show that for all $\mathbf{H}, \mathbf{K}, \mathbf{m} \neq \mathbf{m}'$, the size of the set $\mathcal{D}_\mathbf{H}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_\mathbf{H}(\mathbf{K}, \mathbf{m}')$ is at most 1.

We prove that for any $\mathbf{X}, \mathbf{X}' \in \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}) \cap \mathcal{D}_{\mathbf{H}}(\mathbf{K}, \mathbf{m}')$, it must be $\mathbf{X} = \mathbf{X}'$. By the definition of $\mathcal{D}_{\mathbf{H}}$, we know that $\mathbf{X}\mathbf{m} = \mathbf{X}'\mathbf{m}$ and $\mathbf{X}\mathbf{m}' = \mathbf{X}'\mathbf{m}'$. Therefore, $(\mathbf{X} - \mathbf{X}')(\mathbf{m} - \mathbf{m}') = 0$. But $\mathbf{m} - \mathbf{m}' = [m_1, 1]^T - [m'_1, 1]^T = [m_1 - m'_1, 0]^T$, and

$$(\mathbf{x}_1 - \mathbf{x}'_1)(m_1 - m'_1) = (\mathbf{X} - \mathbf{X}')(\mathbf{m} - \mathbf{m}') = 0 \quad (2)$$

in the ring R . Now we observe that, since the product of $\|\mathbf{x}_1 - \mathbf{x}'_1\|_{\infty} \leq 2b$ and $\|m_1 - m'_1\|_1 \leq 2w$ is at most $4wb < p/2$, no reduction modulo p takes place during the multiplication of $(\mathbf{x}_1 - \mathbf{x}'_1)$ by $(m_1 - m'_1)$, and therefore (2) holds over the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$. Since $\mathbb{Z}[x]/\langle x^n + 1 \rangle$ is an integral domain and $m_1 \neq m'_1$, we can conclude that (2) is equivalent to $\mathbf{x}_1 = \mathbf{x}'_1$. This proves that the keys \mathbf{X} and \mathbf{X}' have the same first vector. But if $\mathbf{x}_1 = \mathbf{x}'_1$, then we also have $\mathbf{x}_2 = \mathbf{X}\mathbf{m} - \mathbf{x}_1 m_1 = \mathbf{X}'\mathbf{m} - \mathbf{x}'_1 m_1 = \mathbf{x}'_2$, and so the two keys \mathbf{X}, \mathbf{X}' are identical. \square

Combining the previous lemmas, and Theorem 3.2, we obtain the following corollary.

Corollary 4.11. *Let $b = \lfloor (|\mathcal{M}|^{1/n} 2^{\lambda/n} p)^{1/m} \rfloor$ and $p > 8wb$. Then, the one-time signature scheme from Section 3, instantiated with the sets in Figure 4, is strongly unforgeable based on the assumed average-case hardness of the RING-SIS $_{n,m,p,4wb}$ problem. Furthermore, for $\gamma = 64wbmn \log^2 n$ and $p \geq \frac{\gamma \sqrt{n}}{4 \log n}$, the scheme is secure based on the worst-case hardness of SVP $_{\gamma}^{\infty}$ in all n -dimensional ideals of the ring $\mathbb{Z}[x]/\langle x^n + 1 \rangle$.*

We remark that for the message space \mathcal{M} to be superpolynomial size, we must have $w = \omega(1)$. So, even using RING-SIS average-case hardness assumptions, we must have $p = \omega(1)$. The expression for b can be simplified by setting $|\mathcal{M}| = 2^n$ and $\lambda = n$. This gives $b = \lfloor (4p)^{1/m} \rfloor$, which, for $m > (2 + \log_2 p) / (\log_2 3 - 1) = O(\log p)$ is just $b = 1$. In practice, one may want to use higher values of b (and smaller values of m), to improve the signature size and overall efficiency of the scheme, at the cost of making stronger security assumptions.

When basing the problem on the worst-case hardness of SVP on ideal lattices, one could set $w = O(n/\log n)$, $b = 1$, $m = O(\log p)$, modulus $p = n^{2.5} \log n$, and worst-case approximation factor $\gamma = O(n^2 \log^2 n)$.

4.3 One-time signature as hard as the Small Codeword Problem

The code-based signature scheme is defined by instantiating the abstract construction from Section 3 with the sets in Figure 5 parametrized by integers n, m, k, w , and b which should satisfy certain

$$\begin{aligned} R &= \mathbb{Z}_2 \\ \mathcal{H} &= R^{n \times m} \\ \mathcal{K} &= \{\mathbf{K} \in R^{m \times k} : \|\mathbf{K}\|_1 \leq b\} \\ \mathcal{M} &\subseteq \{\mathbf{m} \in R^k : \|\mathbf{m}\|_1 = w\} \\ \mathcal{S} &= \{\mathbf{s} \in R^m : \|\mathbf{s}\|_1 \leq wb\} \end{aligned}$$

Fig. 5. Code-based instantiation of the one-time signature scheme, parametrized by integers n, m, k, w, b .

relationships. The size of the message space will be $\binom{k}{w}$ and we will prove the security of our scheme based on the hardness of the SC $_{n,m,2wb}$ problem from Definition 2.8.

Unlike for the lattice scheme in the previous section, we do not have as much freedom in how to set the parameters. This is mostly due to the fact that the ring in this scheme is fixed to \mathbb{Z}_2 , whereas in the lattice scheme, we had the freedom to set the parameter p for $R = \mathbb{Z}_p$. For some constants c, c' , we instantiate the scheme with parameters $m = n^{c+1+c\lambda/n}$, $b = n/(c \log n)$, and $w = c' \log n$. These values satisfy the relation

$$\sum_{i=0}^b \binom{m}{i} > \binom{n^{c+1+c\lambda/n}}{\frac{n}{c \log n}} > \left(n^{c(1+\lambda/n)\frac{n}{c \log n}} \right) = 2^{n+\lambda},$$

which will be used to prove the security of the scheme based on the hardness of $\text{SC}_{n,m,2wb}$. Notice that for $k = n^{\Omega(1)}$, the size of the message space size is $|\mathcal{M}| = \binom{k}{w} = 2^{\Omega(c' \log^2 n)}$, which is super-polynomial, but much smaller than the exponential message space size of our lattice based schemes. Finally, for the $\text{SC}_{n,m,2wb}$ problem to be hard (see Lemma 4.13), we need $2wb = 2nc'/c < n/4$. Thus we require $c' < c/8$.

Lemma 4.12. *The [Closure] property holds*

Proof. It's clear that for any secret key \mathbf{K} and message \mathbf{m} , we have $\|\mathbf{K}\mathbf{m}\|_1 \leq wb$.

Lemma 4.13. *The function family $\{\mathbf{H} : \mathcal{S} \rightarrow R^n \mid \mathbf{H} \in \mathcal{H}\}$ satisfies the [Collision Resistance] property based on the average-case hardness of the $\text{SC}_{n,m,2wb}$ problem.*

Proof. If one can find $\mathbf{x} \neq \mathbf{x}' \in \mathcal{S}$ for a random \mathbf{H} from \mathcal{H} such that $\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{x}'$, then one has that $\mathbf{H}(\mathbf{x} - \mathbf{x}') = 0$ and $\|\mathbf{x} - \mathbf{x}'\|_1 \leq 2wb$. \square

Lemma 4.14. *For every $\mathbf{H} \in \mathcal{H}$, if \mathbf{K} is chosen uniformly at random from \mathcal{K} , then with probability at least $1 - k2^{-\lambda}$, there exists a key $\mathbf{K}' \in \mathcal{K}$ such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$ and $\mathbf{K}' \neq \mathbf{K}$ differ in every column.*

Proof. Consider \mathbf{H} as a function mapping from domain $X = \{\mathbf{x} \in \mathbb{Z}_2^m : \|\mathbf{x}\|_1 \leq b\}$ to range $Y = \mathbb{Z}_2^b$. Notice that by our setup, $|X| = \sum_{i=0}^b \binom{m}{i} \geq 2^{n+\lambda}$ and $|Y|$ is exactly 2^n . By Lemma 4.1, we know that for a randomly chosen vector $\mathbf{x} \in X$, with probability at least $1 - 2^{-\lambda}$, there is another vector $\mathbf{x}' \in X$ such that $\mathbf{H}\mathbf{x} = \mathbf{H}\mathbf{x}'$. Thus we have that for any particular column \mathbf{K}_j , with probability at least $1 - 2^{-\lambda}$, there exists a column \mathbf{K}'_j such that $\mathbf{H}\mathbf{K}_j = \mathbf{H}\mathbf{K}'_j$ and $\mathbf{K}_j \neq \mathbf{K}'_j$. Applying the union bound, we get that with probability at least $1 - k2^{-\lambda}$ this is true for every column $j = 1, \dots, k$, giving a key \mathbf{K}' such that $\mathbf{H}\mathbf{K} = \mathbf{H}\mathbf{K}'$ and $\mathbf{K}_j \neq \mathbf{K}'_j$ for all j . \square

Lemma 4.15. *The $[\frac{1}{2}$ -Hiding] property holds true.*

Proof. The proof is verbatim the proof of Lemma 4.5 except that references to Lemma 4.4 should be replaced with references to Lemma 4.14.

Combining the previous lemmas, and Theorem 3.2, we obtain the following corollary.

Corollary 4.16. *Let $m = n^{c+1+c\lambda/n}$, $b = n/(c \log n)$ and $w = c' \log n$ for some constants $c > 8c' > 0$. The one-time signature scheme from Section 3, instantiated with the set in Figure 5, is strongly unforgeable based on the assumed average-case hardness of the $\text{SC}_{n,m,p,2wb}$ problem.*

5 Conclusions and open problems

The main technical contribution of this work is a construction of a one-time digital signature scheme that takes $\tilde{O}(k)$ time to compute and has conjectured security of $2^{\Omega(k)}$. Since its original publication, the techniques in this paper were used as a starting point in constructions of more “advanced” lattice primitives such as identification schemes [Lyu08,Lyu09], signature schemes (without the “one-time” restriction) [Lyu09,Lyu12,DDLL13,BG14,DLL⁺17], blind signature schemes [Rüc10], and ring signature schemes [MBB⁺13].³ The main conceptual difference between the one-time signature in this paper and the schemes listed above is that it is fine to leak a little information about the secret key in the one-time construction as long as it does not information-theoretically reveal the secret key. In the latter schemes, however, this leakage occurs with every signature (not just once) and so will eventually reveal the entire key. To prevent leakage while retaining efficiency, one needs to use the “Fiat-Shamir with Aborts” technique introduced in [Lyu08,Lyu09] and refined in subsequent works.

Because the full digital signature schemes mentioned above are fairly compact (signatures and public keys around 2KB for 128 bits of conjectured security against quantum attackers), one might think that the one-time signature in this paper would have even smaller parameters. Unfortunately, this is not the case. Starting from [Lyu12], it was observed that the optimal way to set parameters is to have the secret key \mathbf{K} come from a domain for which there is a unique $\hat{\mathbf{K}}$ satisfying $\mathbf{H}\mathbf{K} = \hat{\mathbf{K}}$.⁴ The signature $\mathbf{s} = \mathbf{K}\mathbf{m}$, on the other hand comes from a domain for which there are multiple possible \mathbf{s}' satisfying $\mathbf{H}\mathbf{s}' = \hat{\mathbf{K}}\mathbf{m}$. The reason for setting parameters in this manner is due to the fact that the hardest knapsack problems have density 1 [IN96] – that is if $\mathbf{H} : D \rightarrow R$ is a linear function and $D' \subset D$ is a subset of D with small coefficients, then finding a pre-image $\mathbf{s} \in D'$ satisfying $\mathbf{H}\mathbf{s} = \mathbf{t}$ is hardest when $|D'| \approx |R|$ and gets progressively easier as $|D'|$ increases or decreases. Positioning both the key and signature parameters around density 1 knapsacks (unlike in this paper where the problem of recovering the key is close to a density 1 problem, whereas recovering the signature is further away) therefore allows us to base the hardness of the scheme on a harder problem.

In our current scheme, we crucially need that there exist multiple secret keys \mathbf{K} for every public key $\hat{\mathbf{K}}$, and so cannot use the smaller secret key domain mentioned above. One may try to overcome this problem (and indeed this is what was done in [Lyu12]) by using the indistinguishability of $(\mathbf{H}, \hat{\mathbf{K}} = \mathbf{H}\mathbf{K})$ from uniform based on the hardness of the Learning with Errors problem to argue that we can substitute a real public by one that comes from the domain we need for the proof. But using this idea, we run into the problem that the reduction is not able to generate a valid signature. In [Lyu12] this was not an issue because the random oracle could be programmed so that valid signatures could be simulated even with an invalid public key. Without a random oracle, we do not see how this step could be accomplished. Even with a random oracle, it is not straightforward to adapt our current construction so that it uses programming. In full-fledged signatures, the distribution of the signature is independent of the secret key, thus one could simulate a valid signature (using standard simulation techniques for Σ -protocols) by first picking a signature from the correct

³ All the signature schemes are proved secure in the random oracle model.

⁴ We are just using the notation from this paper as an analogy. The actual keys in the full-fledged signature are constructed a little differently. In particular, the (secret and public) keys in this work actually comprise both the (secret and public) keys and the “commit” step of the Σ -protocols underlying the full signature schemes. We refer readers to [Lyu09] for a more in-depth discussion about the relationship of collision-resistant hash functions, one-time signatures, Σ -protocols, and full signatures.

distribution and then filling in the other parts. In our case, however, the signature depends on the secret key, and so the same simulation technique does not work. In short, constructing a one-time signature scheme that is more practical than full-fledged signatures in the random oracle model remains an open problem.

We mention that there was also recent work [Lyu16] that showed how to construct digital signatures in the random oracle model based on the simultaneous hardness of the SIVP problem *simultaneously* in all rings $\mathbb{Z}[x]/\langle f(x) \rangle$. The construction was built on top of a collision-resistant hash function defined over the ring $\mathbb{Z}[x]$ in which finding collisions is as hard as solving the SIVP problem in all rings. It is relatively straightforward to adapt our instantiation from Section 4.2 to this collision-resistant hash function.

An interesting question deals with improving the efficiency of the code-based scheme in this paper. We show that it is possible to instantiate our general framework based on the hardness of the Small Codeword Problem, but the resulting scheme is quite inefficient. In particular, to get super-polynomial hardness, we are only able to sign messages of length approximately $\log^2 k$ and base the hardness of our scheme on a problem that is only $2^{\Omega(\log^2 k)}$ hard. Interestingly, the more practical hash-and-sign code-based signature scheme of Courtois et al. [CFS01] is also asymptotically based on the hardness of a problem that is at most $2^{O(\log^2 k)}$ hard. Furthermore, technical reasons prevent us from instantiating the code-based scheme based on a problem allowing for a more structured public key, analogous to Ring-SIS. Thus, the problem of constructing efficient code-based one-time signatures without using random oracles remains open.

It would also be interesting to see if our general framework can be instantiated using different assumptions, such as those from multivariate cryptography.

6 Acknowledgements.

We would like to thank the anonymous referees for their insightful comments and corrections that helped to improve the presentation of this article.

References

- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
- [BC92] Jurjen N. Bos and David Chaum. Provably unforgeable signatures. In *CRYPTO*, pages 1–14, 1992.
- [BG14] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, pages 28–47, 2014.
- [BM84] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo-random bits. *SIAM J. Comput.*, 13(4):850–864, 1984. Prelim. version in FOCS 1982.
- [BM96] Daniel Bleichenbacher and Ueli M. Maurer. On the efficiency of one-time digital signatures. In *ASIACRYPT*, pages 145–158, 1996.
- [CFS01] Nicolas Courtois, Matthieu Finiasz, and Nicolas Sendrier. How to achieve a McEliece-based digital signature scheme. In *ASIACRYPT*, pages 157–174, 2001.
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO (1)*, pages 40–56, 2013.
- [DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [DLL⁺17] Léo Ducas, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - dilithium: Digital signatures from module lattices. <http://eprint.iacr.org/2017/633>, 2017.
- [EGM96] Shimon Even, Oded Goldreich, and Silvio Micali. On-line/off-line digital signatures. *J. Cryptology*, 9(1):35–67, 1996. Prelim. version in CRYPTO 1989.

- [GGKT05] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM J. Comput.*, 35(1):217–246, 2005. Prelim. versions in FOCS 2000 and STOC 2003.
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [HM02] Alejandro Hevia and Daniele Micciancio. The provable security of graph-based one-time signatures and extensions to algebraic signature schemes. In *ASIACRYPT*, pages 379–396, 2002.
- [IN96] Russell Impagliazzo and Moni Naor. Efficient cryptographic schemes provably as secure as subset sum. *J. Cryptology*, 9(4):199–216, 1996.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, pages 144–155, 2006.
- [LM08] Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In *TCC*, pages 37–54, 2008.
- [LPR13a] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43:1–43:35, 2013. Prelim. version in Eurocrypt 2010.
- [LPR13b] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for Ring-LWE cryptography. In *EUROCRYPT*, pages 35–54, 2013.
- [Lyu08] Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In *Public Key Cryptography*, pages 162–179, 2008.
- [Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, pages 598–616, 2009.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, pages 738–755, 2012.
- [Lyu16] Vadim Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In *ASIACRYPT*, pages 196–214, 2016.
- [MBB⁺13] Carlos Aguilar Melchor, Slim Bettaieb, Xavier Boyen, Laurent Fousse, and Philippe Gaborit. Adapting Lyubashevsky’s signature schemes to the ring signature setting. In *AFRICACRYPT*, pages 1–25, 2013.
- [Mer87] Ralph C. Merkle. A digital signature based on a conventional encryption function. In *CRYPTO*, pages 369–378, 1987.
- [Mer89] Ralph C. Merkle. A certified digital signature. In *CRYPTO*, pages 218–238, 1989.
- [Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4):365–411, 2007. Prelim. version in FOCS 2002.
- [MP13] Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In *CRYPTO (1)*, pages 21–39, 2013.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Prelim. version in FOCS 2004.
- [NY89] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *STOC*, pages 33–43, 1989.
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166, 2006.
- [PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *STOC*, pages 478–487, 2007.
- [Rom90] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *STOC*, pages 387–394, 1990.
- [Rüc10] Markus Rückert. Lattice-based blind signatures. In *ASIACRYPT*, pages 413–430, 2010.
- [Szy04] Michael Szydło. Merkle tree traversal in log space and time. In *EUROCRYPT*, pages 541–554, 2004.