An Efficient Lattice-Based Signature Scheme
with Provably Secure Instantiation

Sedat Akleylek[1], Nina Bindel[2], Johannes Buchmann[2], Juliane Krämer[2], and Giorgia Azzurra Marson[2]

[1] Ondokuz Mayis University, Turkey
`sedat.akleylek@bil.omu.edu.tr`
[2] Technische Universität Darmstadt, Germany
`nbindel@cdc.informatik.tu-darmstadt.de`, `buchmann@cdc.informatik.tu-darmstadt.de`,
`jkraemer@cdc.informatik.tu-darmstadt.de`, `giorgia.marson@cased.de`

February 22, 2016

Abstract. In view of the expected progress in cryptanalysis it is important to find alternatives for currently used signature schemes such as RSA and ECDSA. The most promising lattice-based signature schemes to replace these schemes are BLISS (CRYPTO 2013) and GLP (CHES 2012). Both come with a security reduction from a lattice problem and have high performance. However, their parameters are not chosen according to their provided security reduction, i.e., the *instantiation* is not provably secure. In this paper, we present the first lattice-based signature scheme with good performance when provably secure instantiated. To this end, we provide a tight security reduction for the new scheme from the ring learning with errors problem which allows for provably secure and efficient instantiations. We present experimental results obtained from a software implementation of our scheme. They show that our scheme, when provably secure instantiated, performs comparably with BLISS and the GLP scheme.

Keywords. lattice-based cryptography, tightness, ideal lattices, signatures, ring learning with errors

1 Introduction

Electronic signatures are essential for cybersecurity. For example, they provide authenticity proofs for billions of software downloads daily on the Internet. In recent years, lattice-based signatures such as BLISS [22] or the GLP [27] signature scheme have become an interesting alternative to the schemes that are currently being used in practice, like RSA and ECDSA. Providing such alternatives is very important in view of the expected progress in cryptanalysis of RSA and ECDSA, in particular by quantum computers.

The lattice-based signature schemes BLISS and GLP have two important properties. They have *good performance*, i.e., they can compete with RSA and ECDSA. Also, they are *provably secure*: they allow for security reductions from lattice problems that are expected to be hard even in the presence of quantum computers.

Provable security is a very strong security argument. In this paper, we go one step further and present an R-LWE-based signature scheme which has a security property which we consider to be even stronger: *good performance with provably secure instantiation*. By this property we mean that the parameters are chosen according to the security reduction and at the same time allow for good performance. This implies the following: suppose that parameters are constructed for a certain security level. By virtue of the security reduction these parameters correspond to an instance of the ring learning with errors problem (R-LWE). Since the parameters were chosen according to the security reduction, this reduction provably guarantees that our scheme has the selected security level as long as the corresponding R-LWE

instance is intractable. In other words, hardness statements for R-LWE instances have a provable consequence for the security levels of our scheme. Currently, both BLISS and GLP do not allow for good performance and provably secure instantiation at the same time. Choosing parameters according to the security reductions for these schemes reduces their performance significantly (see for example [10, 16]).

We note that our scheme has another potential advantage over BLISS. BLISS uses Gaussian sampling, which is generally assumed to be vulnerable to timing attacks [13, 19], while GLP and our scheme use uniform sampling during signature generation which appears to not have this vulnerability.

Our signature scheme is based on the design of Bai and Galbraith [9] and its optimizations by Dagdelen *et al.* [19]. The reason why our scheme allows for good performance with provably secure instantiation is that we are able to give a tight security reduction from the R-LWE problem to our scheme. The proof of this result is an optimized adaption of the tightness proof in [5] to the R-LWE setting which allows for better tightness bounds. To demonstrate that our scheme has good performance, we present experimental results which are based on a software implementation. These results show that our scheme, when provably secure instantiated, performs comparably with BLISS and the GLP scheme without provably secure instantiation.

Related Work. The first lattice-based signature scheme with tight security reduction is the GPV signature scheme [26]. Its instantiations are provably secure, but not efficient. Most of the recent lattice-based signature schemes [9,19,22,27,37] come neither with a tight reduction nor with provably secure instantiation. The security of all those schemes was proven by applying the powerful Forking Lemma [41], which inherently results in a non-tight security reduction.

Abdalla *et al.* [1] circumvent the Forking Lemma and use an approach inspired by the proof idea introduced by Katz and Wang [31]. However, their tight reduction demands an impractically large choice of the modulus. Recently, Alkim *et al.* [5] also used the approach by Katz and Wang [31] to provide a tight security reduction from the learning with errors problem over standard lattices (LWE) to an improved variant of the Bai-Galbraith signature scheme [9, 19]. Instantiations of their scheme are provably secure, but they yield larger key sizes and worse run times than the BLISS and GLP signature scheme.

Organization. After stating notations and definitions in Section 2, we describe the signature scheme in Section 3. In Section 4, we analyze the hardness of R-LWE and we explain the derivation of the parameter sets. Our implementation is described in Section 5. We give our experimental results and compare them with BLISS and GLP in Section 6.

2 Preliminaries

2.1 Notation

Let $k \in \mathbb{N}$. Throughout this paper we define $n = 2^k \in \mathbb{N}$. Let $q \in \mathbb{N}$ be a prime with $q = 1$ (mod $2n$). We denote by $\mathbb{Z}_q$ the finite field $\mathbb{Z}/q\mathbb{Z}$ and identify an element in $\mathbb{Z}_q$ with its representative in $[-\lceil q/2 \rceil, \lfloor q/2 \rfloor]$, and we write (mod q) to denote the unique representative in $\mathbb{Z}_q$. We define the ring $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$ and denote the set of its units by $\mathcal{R}^\times$. Further, we define $\mathcal{R}_q = \mathbb{Z}_q[x]/(x^n + 1)$, $\mathcal{R}_{q,[B]} = \{\sum_{i=0}^{n-1} a_i x^i \in \mathcal{R}_q \mid i \in [0, n-1], a_i \in [-B, B]\}$ for $B \in [0, q/2]$, and $\mathbb{B}_{n,\omega} = \left\{ \mathbf{v} \in \{0, 1\}^n \mid ||\mathbf{v}||^2 = \omega \right\}$ for $\omega \in [0, n]$. We denote polynomials by

lower case letters (e.g., p) and (column) vectors by bold lower case letters (e.g., $\mathbf{v}$). Without further mentioning, we use the symbol $\mathbf{p}$ to denote the coefficient vector of a polynomial p. We denote matrices by bold upper case letters (e.g., $\mathbf{M}$) and the transpose of a matrix $\mathbf{M}$ by $\mathbf{M}^T$. We indicate the Euclidean norm of a vector $\mathbf{v} \in \mathbb{R}^n$ by $\|\mathbf{v}\|$. All logarithms are in base 2.

Rounding Operators. Let $d \in \mathbb{N}$ and $c \in \mathbb{Z}$. We denote by $[c]_{2^d}$ the unique representative of c modulo 2^d in the set $(-2^{d-1}, 2^{d-1}] \subset \mathbb{Z}$. Let $\lfloor \cdot \rceil_d$ be the rounding operator defined as $\lfloor \cdot \rceil_d : \mathbb{Z} \to \mathbb{Z}, c \mapsto (c - [c]_{2^d})/2^d$. We naturally extend these definitions to vectors and polynomials by applying $\lfloor \cdot \rceil_d$ and $[\cdot]_{2^d}$ to each component of the vector and to each coefficient of the polynomial, respectively. We abbreviate $\lfloor v \pmod{q} \rceil_d$ by $\lfloor v \rceil_{d,q}$.

Algorithms and Distributions. If $\mathcal{A}$ is a randomized algorithm we denote by $y \leftarrow \mathcal{A}(x)$ the output of $\mathcal{A}$ on input x and randomly chosen (internal) coins. For an oracle $\mathcal{O}$ we write $\mathcal{A}^{\mathcal{O}}$ to indicate that $\mathcal{A}$ has access to that oracle. Let $\sigma \in \mathbb{R}_{>0}$. The centered discrete Gaussian distribution $\mathcal{D}_\sigma$ on $\mathbb{Z}$ with standard deviation σ is defined as follows: for every $z \in \mathbb{Z}$ the probability of z is given by $\rho_\sigma(z)/\rho_\sigma(\mathbb{Z})$, where $\rho_\sigma(z) = \exp(\frac{-z^2}{2\sigma^2})$ and $\rho_\sigma(\mathbb{Z}) = 1 + 2\sum_{z=1}^{\infty} \rho_\sigma(z)$. We denote by $d \leftarrow \mathcal{D}_\sigma$ the operation of sampling an element d with Gaussian distribution $\mathcal{D}_\sigma$. When writing $\mathbf{v} \leftarrow \mathcal{D}_\sigma^n$ we mean sampling each component of the vector $\mathbf{v}$ with Gaussian distribution. To simplify the notation we indicate sampling all coefficients of a polynomial $a \in \mathcal{R}$ with Gaussian distribution by $a \leftarrow \mathcal{D}_\sigma^n$ as well. Similarly, for a finite set S we write $s \leftarrow \mathcal{U}(S)$, or simply $s \leftarrow_\$ S$, to indicate that an element s is sampled uniformly at random from S.

Lattices and Gaussian Heuristic. Let $n \geq k > 0$. A k-dimensional lattice Λ is a discrete additive subgroup of $\mathbb{R}^n$ containing all integer linear combinations of k linearly independent vectors $\{\mathbf{b}_1, \ldots, \mathbf{b}_k\} = \mathbf{B}$, i.e., $\Lambda = \Lambda(\mathbf{B}) = \{\mathbf{B}\mathbf{x} \mid \mathbf{x} \in \mathbb{Z}^k\}$. The determinant of a lattice is defined by $\det(\Lambda(\mathbf{B})) = \sqrt{\det(\mathbf{B}^\top \mathbf{B})}$.

Throughout this paper we are mostly concerned with q-ary lattices. $\Lambda \in \mathbb{Z}^n$ is called a *q-ary* lattice if $q\mathbb{Z} \subset \Lambda$ for some $q \in \mathbb{Z}$. Let $\mathbf{A} \leftarrow_\$ \mathbb{Z}_q^{m \times n}$. We define the q-ary lattices $\Lambda_q^{\perp}(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^n \mid \mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}\}$ and $\Lambda_q(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^n \mid \exists \mathbf{s} \in \mathbb{Z}^m \text{ s.t. } \mathbf{x} = \mathbf{A}^\top \mathbf{s} \pmod{q}\}$. Furthermore, for $\mathbf{u} \in \mathbb{Z}_q^m$ we define cosets $\Lambda_{\mathbf{u},q}^{\perp}(\mathbf{A}) = \{x \in \mathbb{Z}^n \mid \mathbf{A}\mathbf{x} = \mathbf{u} \pmod{q}\}$, i.e., $\Lambda_q^{\perp}(\mathbf{A}) = \Lambda_{\mathbf{0},q}^{\perp}(\mathbf{A})$. One can consider $\Lambda_{\mathbf{u},q}^{\perp}(\mathbf{A})$ as a *shifted lattice* by a vector $\mathbf{u}$, i.e., $\Lambda_{\mathbf{u},q}^{\perp}(\mathbf{A}) = \Lambda_q^{\perp}(\mathbf{A}) + \mathbf{y}$ where $\mathbf{y} \in \mathbb{Z}^m$ is an integer solution of $\mathbf{A}\mathbf{x} = \mathbf{u} \pmod{q}$.

Let S be a measurable set and let $\Lambda \subset \mathbb{Z}^n$ be a lattice. The *Gaussian heuristic* approximates the number of lattice points in the set S by $|S \cap \Lambda| = \frac{vol(S)}{det(\Lambda)}$.

2.2   The Learning with Errors Problem over Rings

Given the isomorphism $\Phi_q : \mathbb{Z}_n \to \mathcal{R}_q$ with $(a_0, \ldots, a_{n-1}) \mapsto a_0 + a_1 x + \ldots + a_{n-1} x^{n-1}$, $\mathcal{R}_q$ is isomorphic to $\mathbb{Z}_q^n$ as a $\mathbb{Z}$-module. Therefore, we can identify a polynomial $a = a_0 + a_1 x + \ldots + a_{n-1} x^{n-1} \in \mathcal{R}_q$ with its coefficient vector $\mathbf{a} = (a_0, \ldots, a_{n-1})^T$. We define the rotation of a vector $\mathbf{a} = (a_0, \ldots, a_{n-1})^T$ to be the coefficient vector of $ax \in \mathcal{R}_q$, i.e., $\mathrm{rot}(\mathbf{a}) = (-a_{n-1}, a_0, \ldots, a_{n-2})^T$. Furthermore, we define the rotation matrix of a polynomial a as $\mathrm{Rot}(a) = (\mathbf{a}, \mathrm{rot}(\mathbf{a}), \mathrm{rot}^2(\mathbf{a}), \ldots, \mathrm{rot}^{n-1}(\mathbf{a})) \in \mathbb{Z}_q^{n \times n}$. Polynomial multiplication of $a, b \in \mathcal{R}_q$ is equivalent to the matrix-vector multiplication $\mathrm{Rot}(a)\mathbf{b}$ in $\mathbb{Z}_q$. It can be easily shown that $a \in \mathcal{R}_q$ is invertible, i.e., $a \in \mathcal{R}_q^\times$, if and only if $\mathrm{rank}(\mathrm{Rot}(a)) = n$.

We define the learning with errors distribution and the ring learning with errors problem (R-LWE) in the following.

Definition 1 (Learning with Errors Distribution). *Let $n, q > 0$ be integers, $s \in \mathcal{R}_q$, and χ be a distribution over $\mathcal{R}$. We define by $\mathcal{D}_{s,\chi}$ the R-LWE distribution which outputs $(a, \langle a, s \rangle + e) \in \mathcal{R}_q \times \mathcal{R}_q$, where $a \leftarrow_\$ \mathcal{R}_q$ and $e \leftarrow \chi$.*

Since our signature scheme is based on the decisional R-LWE problem, we omit the definition of the search version and state only the decisional learning with errors problem.

Definition 2 (Ring Learning with Errors Problem). *Let $n, q > 0$ be integers and $q = 2^k$ for some $k \in \mathbb{N}_{>0}$ and χ be a distribution over $\mathcal{R}$. Moreover, define $\mathcal{O}_\chi$ to be an oracle, which upon input polynomial $s \in \mathcal{R}_q$ returns samples from the learning with errors distribution $\mathcal{D}_{s,\chi}$. The ring learning with errors problem R-LWE$_{n,m,q,\chi}$ is (t, ε)-hard if for any probabilistic polynomial time (PPT) algorithm $\mathcal{A}$, running in time t and making at most m queries to its oracle, it holds that*

$$Adv_{n,q,\chi}^{R\text{-}LWE}(\mathcal{A}) = \left| \Pr\left[\mathcal{A}^{\mathcal{O}_\chi(s)(\cdot)} = 1 \right] - \Pr\left[\mathcal{A}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)(\cdot)} = 1 \right] \right| \leq \varepsilon ,$$

where the probabilities are taken over the random choices of $s \leftarrow \mathcal{U}(\mathcal{R}_q)$, the random choice of the distribution $\mathcal{D}_{s,\chi}$, as well as the random coins of $\mathcal{A}$.

The R-LWE assumption comes with a worst-case to average-case reduction to problems over ideal lattices [38]. Furthermore, it was shown in [6] that the learning with errors problem remains hard if one chooses the secret distribution to be the same as the error distribution. We write R-LWE$_{n,m,q,\sigma}$ if χ is the discrete Gaussian distribution with standard deviation σ.

3 Description and Security of the Signature Scheme

In this section, we present our signature scheme and we prove it to be unforgeable against a chosen-message attack—shortly ufcma-secure (cf. Appendix A, Figure 3)—as long as R-LWE is computationally hard. We recall basic definitions and notations about signatures schemes in Appendix A. We name our scheme ring-TESLA since it is based on the signature scheme TESLA by Alkim *et al.* [5].

Our signature scheme is parametrized by the integers $n \in \mathbb{N}_{>0}$, ω, d, B, q, U, L, κ, and the security parameter λ with $n > \kappa > \lambda$, by the Gaussian distribution $\mathcal{D}_\sigma$ with standard deviation σ, by the hash function $H : \{0,1\}^* \to \{0,1\}^\kappa$, and by the encoding function $F : \{0,1\}^\kappa \to \mathbb{B}_{n,\omega}$. The encoding function F takes the (binary) output of the hash function H and maps it to a vector of length n and weight ω. For more information about the encoding function see [27]. Furthermore, let $a_1, a_2 \in \mathcal{R}_q^\times$ be two uniformly sampled polynomials which are publicly known as global constants. They can be shared among arbitrarily many signers.

Key Generation. On input polynomials a_1 and a_2 the algorithm returns a key pair $(\mathsf{sk}, \mathsf{pk})$ derived as follows. First, it samples three polynomials s, e_1, e_2 according to the discrete Gaussian distribution. The elements e_1 and e_2 have to fulfill certain properties to ensure that signatures are short and verified correctly. Those properties are checked during checkE, similar to the procedure introduced by Dagdelen *et al.* [19]. The function $max_k(\cdot)$ takes as input a vector and returns its k-th largest entry. The key polynomials e_1, e_2 are rejected if $\sum_{k=1}^\omega max_k(\mathbf{e_i})$ is greater then L for at least one of e_1 or e_2. Otherwise e_1, e_2 are accepted and the key generation algorithm returns the signing key $\mathsf{sk} = (s, e_1, e_2)$ and the verification key $\mathsf{pk} = (t_1, t_2) = (a_1 s + e_1, a_2 s + e_2)$.

KeyGen($1^\lambda; a_1, a_2$) :

1. $s, e_1, e_2 \leftarrow D_\sigma^n$
2. If checkE(e_1) $= 0 \vee$ checkE(e_2) $= 0$
3. Restart
4. $t_1 \leftarrow a_1 s + e_1 \pmod{q}$
5. $t_2 \leftarrow a_2 s + e_2 \pmod{q}$
6. sk $\leftarrow (s, e_1, e_2)$
7. pk $\leftarrow (t_1, t_2)$
8. Return (sk, pk)

Verify($\mu; z, c'; a_1, a_2, t_1, t_2$)

20. $c \leftarrow F(c')$
21. $w_1' \leftarrow a_1 z - t_1 c \pmod{q}$
22. $w_2' \leftarrow a_2 z - t_2 c \pmod{q}$
23. $c'' \leftarrow H\left(\lfloor w_1' \rceil_{d,q}, \lfloor w_2' \rceil_{d,q}, \mu\right)$
24. If $c' = c'' \wedge z \in \mathcal{R}_{B-U}$:
25. Return 1
26. Else: Return 0

Sign($\mu; a_1, a_2, s, e_1, e_2$) :

9. $y \leftarrow_\$ \mathcal{R}_{q,[B]}$
10. $v_1 \leftarrow a_1 y \pmod{q}$
11. $v_2 \leftarrow a_2 y \pmod{q}$
12. $c' \leftarrow H\left(\lfloor v_1 \rceil_{d,q}, \lfloor v_2 \rceil_{d,q}, \mu\right)$
13. $c \leftarrow F(c')$
14. $z \leftarrow y + sc$
15. $w_1 \leftarrow v_1 - e_1 c \pmod{q}$
16. $w_2 \leftarrow v_2 - e_2 c \pmod{q}$
17. If $[w_1]_{2^d}, [w_2]_{2^d} \notin \mathcal{R}_{2^d - L} \vee z \notin \mathcal{R}_{B-U}$:
18. Restart
19. Return (z, c')

Fig. 1: Specification of the signature scheme ring-TESLA

Signing Algorithm. During signing a message $\mu \in \{0,1\}^*$, first a polynomial y is sampled uniformly random in $\mathcal{R}_{q,[B]}$ and the polynomials $v_i = a_i y \pmod{q}$ for $i = 1, 2$ are computed. Afterwards, μ and the rounded values of v_1 and v_2 are hashed to $c' = H(\lfloor v_1 \rceil_{d,q}, \lfloor v_2 \rceil_{d,q}, \mu)$ and encoded to $c = F(c')$. Further on, the polynomial $z = y + sc \in \mathcal{R}$ is computed. Now, rejection sampling is applied. To this end, two polynomials $w_i = v_i - e_i c \pmod{q}$ for $i = 1, 2$ are computed and it is checked whether the d least significant bits in the bit representation of w_1, w_2 are small enough, i.e., $[w_i]_{2^d} \in \mathcal{R}_{2^d - L}$ for $i = 1, 2$. If this is satisfied and $z \in \mathcal{R}_{q,[B-U]}$, the signature (z, c') is returned. Otherwise, the signing algorithm discards (z, c') and restarts.

Verification Algorithm. The verification algorithm upon input a message μ and a signature (z, c') first computes the encoding $c = F(c')$. Afterwards, the polynomials $w_i' = a_i z - t_i c \pmod{q}$ for $i = 1, 2$ are computed. The algorithm returns 1 if $c' = H\left(\lfloor w_1' \rceil_{d,q}, \lfloor w_2' \rceil_{d,q}, \mu\right)$ and $z \in \mathcal{R}_{q,[B-U]}$ are both satisfied; otherwise, it returns 0.

3.1 Security Reduction

In this subsection we prove the signature scheme ring-TESLA presented in this paper to be unforgeable against a chosen-message attack (ufcma), described in Figure 3 in Appendix A, as long as the learning with errors problem over rings is computationally hard. In our security reduction we follow an idea introduced by Katz and Wang [31] that can be summarized at follows: assume there exists an algorithm $\mathcal{A}$ that forges a signature given a valid public key, i.e., an LWE tuple. In contrast, given a random key $\mathcal{A}$ forges a signature only with very small

probability. Hence, the security reduction distinguishes whether its own challenge tuple is an LWE tuple or not by the different behavior of the algorithm $\mathcal{A}$.

Theorem 1. *Let n, ω, d, B, q, U, L, and σ be arbitrary parameters satisfying the constraints described in Section 4. Assume that the Gaussian heuristic holds for lattice instances defined by the parameters above. For every ufcma-adversary $\mathcal{A}$ that runs in time $t_{\mathcal{A}}$, asks at most q_s and q_h queries to the signing oracle and the hash oracle, respectively, and forges a valid signature of the signature scheme ring-TESLA with probability $\varepsilon_{\mathcal{A}}$, there exists a distinguisher $\mathcal{D}$ that runs in time $t_{\mathcal{D}} = t_{\mathcal{A}} + \mathcal{O}(q_s \kappa^2 + q_h)$ and breaks the R-LWE$_{n,2,q,\sigma}$ problem (in the random oracle model) with success probability*

$$\varepsilon_{\mathcal{D}} \geq \varepsilon_{\mathcal{A}} \left(1 - \frac{q_s q_h 2^{(d+1)2n}}{(2B+1)^n q^n} \right) - \frac{q_h 2^{dn}(2B - 2U + 1)^n + (28\sigma + 1)^{3n}}{q^{2n}} .$$

Proof sketch. We show how to turn any successful forger $\mathcal{A}$ against the signature scheme ring-TESLA into a distinguisher $\mathcal{D}$ for the R-LWE problem. The distinguisher obtains two R-LWE samples from its sampling oracle $\mathcal{O}_\chi(s)$ (cf. Definition 2) and embeds them into a public key pk. Thus, $\mathcal{D}$ simulates the ufcma game (cf. Figure 3, Appendix A). When $\mathcal{A}$ returns a forgery (μ, σ), $\mathcal{D}$ checks whether σ is a valid signature for message μ under key pk: if so, it outputs 1 as a guess that $\mathcal{O}_\chi(s)$ presented two R-LWE tuples, otherwise it outputs 0. To derive the explicit relation between $\mathcal{D}$ and $\mathcal{A}$'s success probabilities $\varepsilon_{\mathcal{D}}$ and $\varepsilon_{\mathcal{A}}$ as indicated in the theorem statement, we show that (i) $\mathcal{D}$ provides a good simulation of the ufcma game for $\mathcal{A}$. In particular, we show that the simulated signatures look like genuine ones. And we prove, (ii) $\mathcal{D}$'s simulation does not abort too often. Formal proofs of both facts, (i) and (ii), require several technical lemmas that we state and prove in Appendix B. For proving fact (i), we observe that $\mathcal{D}$ simulates signatures $\sigma = (z, c')$ by choosing z and c' uniformly at random from appropriate spaces. By applying rejection sampling (cf. Lemma 1, Appendix B) and the fact that c' is the output of a random oracle, we show that simulated signatures are statistically indistinguishable from genuine ones. Concerning fact (ii), we first note that $\mathcal{D}$'s signing simulation needs to program the random oracle H, which may lead to inconsistencies in case one of $\mathcal{A}$'s signature requests results in programming a hash value $H(x)$ for which x was already queried. Such an occurrence causes a premature termination of the simulation. Applying Lemma 2, we prove that the latter happens only with small probability. We give a full proof of Theorem 1 and of the above-mentioned lemmata in Appendix B. $\qquad \square$

As described in [23, Section 3.3], the probability that a polynomial chosen uniformly random in $\mathcal{R}_q$ is in the subset of multiplicative invertible elements of $\mathcal{R}_q$ is given by $\Pr\left[a \in \mathcal{R}_q^\times \right] = (1 - 1/q)^n$, where the probability is taken over random choices of $a \leftarrow_\$ \mathcal{R}_q$. This probability is overwhelming for our choices of q and n in the signature scheme presented in this paper. Thus, it is justified to sample the polynomials a_1 and a_2 uniformly random in $\mathcal{R}_q^\times$ instead of $\mathcal{R}_q$ as defined in the R-LWE problem.

Relation to Former Security Reductions. The scheme ring-TESLA is based on the signature scheme by Bai and Galbraith [9] with a tight security reduction by Alkim *et al.* [5]. Essentially, we convert the scheme by Bai and Galbraith to a scheme over ideal lattices. Our security reduction follows the proof strategy of [5]. We emphasize that lifting the security statements for the original (lattice-based) scheme to our (ideal lattice-based) scheme is not

trivial. For example, it is unclear whether distributions remain the same when lemmata are applied on rotation matrices instead of matrices chosen uniformly random; in some cases even improvements can be made. Indeed, we could sharpen the bound given in [5, Lemma 2]. Our corresponding result is stated in Lemma 3. Moreover, we formulate and prove a similar lemma to [9, Lemma 3] for ideal lattices and we state explicitly which property related to the Gaussian heuristic is necessary to prove the statement. Likewise, Bai and Galbraith make use of the Gaussian heuristic in their corresponding proof. The methods used in our security reduction resemble those formalized by Abdalla *et al.* [1]. Abdalla *et al.* define four properties of identification schemes for which they give a black-box-transformation to signature schemes with tight security reduction. Applying their black-box-transformation to a lattice-based signature scheme led to inefficiently large parameters as stated by the authors [1]. Hence, we prove unforgeability of ring-TESLA more directly—without passing through an intermediate identification scheme—by following the proof technique introduced by Katz and Wang [31]. This yields practical instantiation as we show in Section 4.

4 Selecting Parameters

The reductionist approach to prove security of a given cryptosystem essentially consists in building an efficient reduction that turns any successful adversary against the cryptosystem into one that solves some computationally hard problem. The hardness of breaking the cryptosystem and of solving the underlying problem are often expressed asymptotically. When a scheme is to be deployed in the real world, however, for a security analysis to be realistic it is essential that run times and success probabilities are estimated in a more explicit way. Moreover, given a (concrete and) tight security reduction, the security of the scheme is about the same as the hardness of the underlying computational assumption when the scheme is instantiated according to the reduction. In contrast, if only a non-tight reduction is available, larger security parameters shall be used in order to achieve a specific level of security. As a consequence, it is often hard to tell whether a provably secure scheme with a non-tight reduction effectively provides the claimed level of security and performance.

In this section, we propose our choice of provably secure parameters for different levels of bit-security for the signature scheme presented in this paper and we explain how we estimate the hardness of the ring learning with errors problem.

4.1 Derivation of Parameters for Different Security Levels

The security reduction given in Section 3 provides a *tight* reduction to the hardness of R-LWE and bounds *explicitly* the forging probability with the success probability of the reduction. More formally, let $\varepsilon_{\mathcal{A}}$ and $t_{\mathcal{A}}$ denote the success probability and the runtime of a forger $\mathcal{A}$ against our signature scheme and let $\varepsilon_{\mathcal{D}}$ and $t_{\mathcal{D}}$ denote analogous quantities for the reduction $\mathcal{D}$ presented in the proof of Theorem 1. We can write the explicit relations $\varepsilon_{\mathcal{D}} \geq c_1 \varepsilon_{\mathcal{A}} + c_2$ and $t_{\mathcal{D}} \leq c_3 t_{\mathcal{A}} + c_4$, where c_1, c_2, c_3, c_4 are constants which are fixed for a concrete instantiation of the signature scheme. We say that R-LWE is *n-bit hard* if $t_{\mathcal{D}}/\varepsilon_{\mathcal{D}} \geq 2^n$; similarly, we say that the signature scheme is *m-bit secure* if $t_{\mathcal{A}}/\varepsilon_{\mathcal{A}} \geq 2^m$.

Given an explicit security reduction and the assumed bit-hardness of R-LWE, we can compute the bit-security of the signature scheme. In our case, we instantiate the signature scheme such that the constants c_1, c_2, and c_3 are roughly equal to 1. Thus, the bit-hardness of the R-LWE instance is the same as the bit-security of our signature instantiated as described

below. To ensure both correctness and security of our signature, the following dependencies must hold.

Let λ be the security parameter. We choose a hash function $H : \{0,1\}^* \to \{0,1\}^\kappa$ with $\kappa > \lambda$ to ensure that the hash function gives at least a bit-hardness of λ. We instantiate the hash function for our parameter sets with SHA-256. Furthermore, security relies on the encoding function $F : \{0,1\}^\kappa \to \mathbb{B}_{n,\omega}$. Following Bai and Galbraith [9], we require F to be close to an injective function. That means that the probability of mapping two different values to the same output is smaller than or equal to $2^{-\lambda}$. We choose ω such that $2^\kappa \geq |\mathbb{B}_{n,\omega}| = 2^\omega \binom{n}{\omega}$. To use efficient polynomial multiplication, i.e., the number theoretic transform (NTT) in the ring $\mathcal{R}_q$, we restrict ourselves to a polynomial degree of a power of 2, i.e, $n = 2^k$ for $k \in \mathbb{N}$. Choosing the Gaussian parameter σ, we can compute the system parameters to give a concrete instantiation of ring-TESLA with λ-bit security.

To apply rejection sampling (see Lemma 1, Appendix B), we choose $U = 14\sqrt{\omega}\sigma$ and $B \geq 14(n-1)\sqrt{\omega}\sigma$. The rejection probability is given by $M = \left(\frac{2(B-U)+1}{2B+1}\right)^n$. We select the rounding value d to be larger than $\log(B)$ and such that the acceptance probability in the first part of Step 17 in Figure 1 is greater than or equal to 0.4, i.e., $(1 - 2L/2^d)^{2n} \geq 0.4$. The bound L is important during the key generation as well as during the sign procedure. We choose L such that we reject only very few of the possible key pairs in checkE. For example, we achieve an acceptance probability of almost 100% in KeyGen and an acceptance probability of 0.34 in Sign for parameter ring-TESLA-II. At last, the modulus q has to be greater than or equal to $\left(\frac{2^{(d+1)2n+\kappa}}{(2B)^n}\right)^{1/n}$ and greater than or equal to $4B$. The theoretical size of the secret key is given by $3n\lceil\log(14\sigma)\rceil$ bits. The public key is represented by $2n\lceil\log(q)\rceil$ bits and the length of the signature is $n\lceil\log(2B-2U)\rceil + \kappa$ bits. Given the concrete instantiations in Table 1, we get a signature size of 1,488 byte, a public key size of 3,328 byte, and a secret key size of 1,920 byte for parameters chosen such that the signature scheme is 128-bit secure. In Table 1 we also propose instantiations for 80 bit of security. For comparison, we depict our signature and key sizes together with the corresponding values of BLISS [22] and the GLP [27] signature scheme in Table 2.

Table 1: Parameter sets for our signature scheme in comparison; the hardness of the LWE instance is defined by the dimension n, the modulus q, and the Gaussian parameter σ; derivation of L, ω, B, U, d is explained in Section 4.1; pk and sk denote the public and private key, resp.

Parameter selection									
Parameter Set	Security (bit)	n	σ	L	ω	B	U	d	q
ring-TESLA-I	80	512	30	814	11	$2^{21}-1$	993	21	8399873
ring-TESLA-II	128	512	52	2766	19	$2^{22}-1$	3173	23	39960577

		Acceptance prob.		pk Size	sk Size	Signature Size
		KeyGen	Sign	(byte)	(byte)	(byte)
ring-TESLA-I	80	0.5	0.23	3,072	1,728	1,418
ring-TESLA-II	128	0.99	0.34	3,328	1,920	1,488

4.2 Hardness Estimation of the R-LWE Problem

Since the introduction of the learning with errors problem over rings [38], it is an open question whether the R-LWE is as hard as the LWE problem. Recently, the cyclic structure of ideal lattices has been exploited by Garg *et al.* [25], by Campbell *et al.* [15], by Cramer *et al.* [18], and by Elias *et al.* [24]. However, up to now, these novel results are not known to be directly applicable to most of the proposed ideal-lattice-based signature schemes. Hence, as the R-LWE problem can be seen as an instantiation of the LWE problem, we estimate the hardness of R-LWE via state-of-the-art attacks against LWE. We explain four basic attacks on LWE: the embedding approach, the decoding attack, the algorithm by Blum, Kalai, and Wassermann [12], and the Arora-Ge-Algorithm [7]. We briefly describe the algorithms next. The most efficient practical approaches to solve LWE are the embedding approach and the decoding attack.

During the *decoding attack*, an LWE instance $(\mathbf{A}, \mathbf{As} + \mathbf{e})$ is seen as an instance of the bounded distance decoding problem (BDD). The idea of the attack is to reduce the lattice by algorithms such as the BKZ algorithm [17] first, and to find the closest lattice vector to a target vector via the nearest plane algorithm by Babai [8] (or improved variants such as by Linder and Peikert [35] or Liu and Nguyen [36]) afterwards. The closest vector corresponds to $\mathbf{As}$ of the LWE instance, such that the secret can be easily discovered.

The *embedding approach* is to solve an LWE instance by reducing it to an instance of the (unique) shortest vector problem. There are different ways to define a lattice that contains the error term of an LWE instance (e.g., [3, 9, 11]). In the end, the short error term is found as a shortest vector of the constructed lattice via basis reductions such as BKZ [17] and LLL [17, 34], or directly via sieving algorithms [33, 39] or enumeration [4]. Recent results [14, 30, 42] exploit the cyclic structure of ideal lattices to improve sieving algorithms. However, the improved sieving algorithms are still slower than the enumeration approach on instances currently used for signatures.

Further, there are two non-lattice approaches to solve LWE, namely the attack based on the algorithm by *Blum, Kalai*, and *Wassermann* (BKW) [12] and the algorithm by *Arora and Ge* [7]. Both algorithms require a (very) large number of LWE samples to be applied efficiently. Although the number of required samples was crucially reduced, for both BKW [20,29,32] and the Arora-Ge algorithm [2], our proposed instances give far less LWE samples than required for the attacks. Hence, we only take the decoding attack and the embedding approach into account when estimating the bit-security of our instances.

We estimate the hardness of our chosen LWE instances based on [3, 9, 35]. We propose parameters for two different levels of security: 80-bit security (ring-TESLA-I) and 128-bit security (ring-TESLA-II). The embedding attack yields 166 bit of security and the result of the decoding attack is a bit security of 139 on the instances in ring-TESLA-II.

5 Software Implementation

The implementation of the proposed scheme targets the Intel Haswell micro architecture. We perform benchmarks on a machine with an Intel Core i7-5820K (Haswell) CPU at 3300MHz and 16GB of RAM. In our software we use the benefits of AVX2 instructions, where multiplication, addition, and subtraction instructions have one cycle throughput for eight doubles. The software is compiled with gcc-4.7 with optimization code. The experimental results are obtained by using gcc-4.7 with "-Ofast" optimization since it enables all "-O3" optimizations

together with turning on "-ffast-math". This optimization helps us to reduce the timing results significantly. The performance of our implementation mainly depends on the number of rejections during Sign and KeyGen and on the time a single polynomial multiplication takes. The derivation of the number of rejections is explained in Section 4.1. We optimized the time for multiplication by choosing the most suitable multiplication algorithm for different cases as it is explained below.

Polynomial Multiplication. In the presented scheme two types of polynomial multiplication occur: standard and sparse polynomial multiplication. For standard polynomial multiplication we use the number theoretic transform (NTT) since NTT performs polynomial multiplication with a quasilinear complexity, i.e., $O(n \log n)$. Thus, the parameter sets are selected in such a way that NTT is applicable, i.e., $q = 1 \pmod{2n}$, where n is a power of 2. In our implementation, we store the integer in double format in a word. Then, after arithmetic operations in NTT, it is expected to fit in a double, i.e., $\log(\log(n)q) + \log(q) < 54$. To avoid an overflow one needs to make extra reduction operations when using ring-TESLA-II because $\log(q)$ is represented by 26 bits. This, of course, results in a drawback of the performance. NTT with extra modulo q reduction would need almost 28383 cycles for $n = 512$ and $\omega = 19$ as chosen in ring-TESLA-II. Without extra reductions, the average cycle count of NTT developed for ring-TESLA-I is 10625. Barrett reduction is preferred over reducing the coefficients because of the modular structure. The hybrid approach of using NTT and sparse polynomial multiplication requires more inverse NTT operations since sparse polynomial multiplication is applicable only in the integer domain.

Input: array $d = [i_1, ..., i_\omega]$, poly $a(x) = \sum_{i=0}^{n-1} a_i x^i$, poly $b(x) = \sum_{i=0}^{n-1} b_i x^i$; with $a_i \in \mathbb{Z}_q$, $b_i \in \{0, 1\}$, $d[k] = i_k$ such that $b_{i_k} = 1$
Output: poly $c(x) = a(x)b(x)$

1   Set all coefficients of $c(x)$ to 0
2   for $i = 0, ..., \omega - 1$:
3   for $j = 0, ..., n - 1$:
4       $c_{j+d[i]} \leftarrow c_{j+d[i]} + a_j$
5   for $i = 0, ..., n - 1$
6     $c_i \leftarrow c_i - c_{i+n} \pmod{q}$
7   Return $c(x)$

Fig. 2: Sparse Polynomial Multiplication

Recall that the weight of c, i.e., the number of 1's, is ω. Then, the multiplication operations in the signature generation phase (Step 14, 15, and 16: sc, e_1c, and e_2c) and in the signature verification phase (Step 20 and 21: t_1c and t_2c) can be considered as sparse polynomial multiplications because of the number of nonzero elements in c. In order to speed up, we use the sparse polynomial multiplication given in Algorithm 2. The complexity of Algorithm 2 depends on the nonzero coefficients of $b(x)$. Note that polynomial multiplication is performed by using only additions if one of the multiplicands is sparse. The required number of additions and subtractions is $(\omega n + n)$. The last for-loop is designed for polynomial reduction modulo $x^n + 1$. There is only one reduction modulo q of the coefficients. This improves the runtime and complexity. Sparse multiplication requires almost 3650 cycles.

We place our implementation of ring-TESLA in public domain. It can be found under https://www.cdc.informatik.tu-darmstadt.de/cdc/personen/nina-bindel.

6 Performance Analysis

We performed our benchmarks on a machine with an Intel Core i7-5820K (Haswell) CPU at 3300MHz and 16 GB of RAM, while disabling Turbo Boost and hyper threading. In our measurement we considered two parameter sets: ring-TESLA-I and ring-TESLA-II with 80 and 128 bits of security, respectively. Our benchmarks are averaged[3] over 10,000 runs of Sign and Verify. We summarize benchmarks for our proposed parameter sets and state-of-the-art ideal-lattice-based signature schemes in Table 2. We emphasize once more that our parameter sets are the only ones in Table 2 which are chosen according to the given security reduction, cf. Section 4. Nevertheless, we achieve good performance with respect to time and space. In the following, we compare sizes and run times for 80 and 128 bits of security.

Table 2: Comparison of our results with the software implementations of the signature schemes BLISS [21, 22] and GLP [19, 27, 28]. To indicate the considered platforms Intel Core i5-3210M (Ivy Bridge), Intel Core i7-5820K (Haswell), and Intel Core 3.4 GHz we use shortcuts A, B, and C, respectively. Sizes of signatures, signing and verification keys are indicated in Bytes. We abbreviate 'Decisional Compact Knapsack problem' by DCK. In the benchmarks of GLP we include the improvements by Dagdelen et al. presented in [19]. In the benchmarks of BLISS we include the improvements by Ducas presented in [21].

80-bit security	GLP [19, 27, 28]	ring-TESLA-I (this paper)
Assumption	DCK	R-LWE
CPU	A	B
Signing key size	256	1,728
Verification key size	1,536	3,072
Signature size	1,186	1,568
Sign cycle counts	452,223	370,880
Verify cycle counts	34,004	94,124

128-bit security	BLISS [21, 22] BLISS-I	BLISS-II	ring-TESLA-II (this paper)
Assumption	R-SIS, NTRU		R-LWE
CPU	C		B
Signing key size	2,048	2,048	1,920
Verification key size	7,168	7,168	3,328
Signature size	1,559	1,514	1,568
Sign cycle counts	351,333	582,857	510,981
Verify cycle counts	102,000	102,000	167,791

---

[3] Sometimes benchmarks are given as the median instead of the average value. Due to the rejection sampling, taking the median value of our experiments would be overly optimistic for Sign.

For *low security of 80-bit*, key and signature sizes of GLP-I are smaller than those of our proposed parameters. Our run time of Sign is a factor of 1.19 faster than GLP. As Table 2 indicates, the software implementations of ring-TESLA and of the GLP signature scheme are optimized for micro architectures.

For *medium security of 128-bit* the instantiation of our scheme gives smallest key sizes. Signature sizes are comparably good. We emphasize that we report the signature size used in the publicly available software implementation of BLISS-I and BLISS-II [4]. Those sizes differ from the theoretical signature sizes presented in [22], which are 700 and 625 bytes for BLISS-I and BLISS-II, respectively, because signatures are not compressed in the BLISS software. To our knowledge, there is no implementation of BLISS available that compresses the signature sizes. The signature size of ring-TESLA are also obtained from our implementation.

The time-optimized implementation of BLISS-I by Ducas [21] is only a factor of 1.45 faster than our implementation. We note that our signature scheme uses uniform sampling during Sign. In contrast, BLISS uses Gaussian sampling, which might be vulnerable to timing attacks [13,19]. Up to now, available implementations of BLISS do not protect against timing-attacks. It would be very interesting to compare our implementation with an optimized and timing-attack-protected implementation of BLISS.

In summary, our signature scheme has good performance compared to state-of-the-art ideal-lattice-based signature schemes, while it is instantiated provably secure. Hence, when real world security matters our presented scheme is a very interesting choice.

Acknowledgment

References

1. M. Abdalla, P.-A. Fouque, V. Lyubashevsky, and M. Tibouchi. Tightly-secure signatures from lossy identification schemes. In Pointcheval and Johansson [40], pages 572–590. 2, 7
2. M. Albrecht, C. Cid, J.-C. Faugère, R. Fitzpatrick, and L. Perret. Algebraic algorithms for LWE problems. Cryptology ePrint Archive, Report 2014/1018, 2014. `http://eprint.iacr.org/2014/1018/`. 9
3. M. R. Albrecht, R. Fitzpatrick, and F. Göpfert. On the efficacy of solving LWE by reduction to unique-SVP. In H.-S. Lee and D.-G. Han, editors, *ICISC 13*, volume 8565 of *LNCS*, pages 293–310, Seoul, Korea, Nov. 27–29, 2014. Springer, Heidelberg, Germany. 9
4. M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of learning with errors. Cryptology ePrint Archive, Report 2015/046, 2015. `http://eprint.iacr.org/`. 9
5. E. Alkim, N. Bindel, J. Buchmann, and O. Dagdelen. Tesla: Tightly-secure efficient signatures from standard lattices. Cryptology ePrint Archive, Report 2015/755, 2015. `http://eprint.iacr.org/`. 2, 4, 6, 7
6. B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618, Santa Barbara, CA, USA, Aug. 16–20, 2009. Springer, Heidelberg, Germany. 4
7. S. Arora and R. Ge. New algorithms for learning in presence of errors. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP 2011, Part I*, volume 6755 of *LNCS*, pages 403–415, Zurich, Switzerland, July 4–8, 2011. Springer, Heidelberg, Germany. 9
8. L. Babai. A las vegas-NC algorithm for isomorphism of graphs with bounded multiplicity of eigenvalues. In *27th FOCS*, pages 303–312, Toronto, Ontario, Canada, Oct. 27–29, 1986. IEEE Computer Society Press. 9

---

[4] `bliss.di.ens.fr`

9. S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In J. Benaloh, editor, *CT-RSA 2014*, volume 8366 of *LNCS*, pages 28–47, San Francisco, CA, USA, Feb. 25–28, 2014. Springer, Heidelberg, Germany. 2, 6, 7, 8, 9, 15

10. M. Bellare and P. Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In U. M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 399–416, Saragossa, Spain, May 12–16, 1996. Springer, Heidelberg, Germany. 2

11. D. J. Bernstein, J. Buchmann, and E. Dahmen, editors. *Post-quantum cryptography*. Mathematics and Statistics Springer-11649; ZDB-2-SMA. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. 9

12. A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. In *32nd ACM STOC*, pages 435–440, Portland, Oregon, USA, May 21–23, 2000. ACM Press. 9

13. J. W. Bos, C. Costello, M. Naehrig, and D. Stebila. Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In *Proc. IEEE Symposium on Security and Privacy (S&P) 2015*, pages 553–570. IEEE, May 2015. 2, 12

14. J. W. Bos, M. Naehrig, and J. van de Pol. Sieving for shortest vectors in ideal lattices: a practical perspective. Cryptology ePrint Archive, Report 2014/880, 2014. http://eprint.iacr.org/2014/880. 9

15. P. Campbell, M. Groves, and D. Shepherd. SOLILOQUY: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf. 9

16. S. Chatterjee, A. Menezes, and P. Sarkar. Another look at tightness. In A. Miri and S. Vaudenay, editors, *SAC 2011*, volume 7118 of *LNCS*, pages 293–319, Toronto, Ontario, Canada, Aug. 11–12, 2012. Springer, Heidelberg, Germany. 2

17. Y. Chen and P. Q. Nguyen. BKZ 2.0: Better lattice security estimates. In D. H. Lee and X. Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 1–20, Seoul, South Korea, Dec. 4–8, 2011. Springer, Heidelberg, Germany. 9

18. R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. Cryptology ePrint Archive, Report 2015/313, 2015. http://eprint.iacr.org/2015/313. 9

19. Ö. Dagdelen, R. E. Bansarkhani, F. Göpfert, T. Güneysu, T. Oder, T. Pöppelmann, A. H. Sánchez, and P. Schwabe. High-speed signatures from standard lattices. In D. F. Aranha and A. Menezes, editors, *LAT-INCRYPT 2014*, volume 8895 of *LNCS*, pages 84–103, Florianópolis, Brazil, Sept. 17–19, 2015. Springer, Heidelberg, Germany. 2, 4, 11, 12

20. A. Duc, F. Tramèr, and S. Vaudenay. Better algorithms for LWE and LWR. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015*, volume 9056 of *LNCS*, pages 173–202. Springer, 2015. 9

21. L. Ducas. Accelerating bliss: the geometry of ternary polynomials. Cryptology ePrint Archive, Report 2014/874, 2014. http://eprint.iacr.org/2014/874. 11, 12

22. 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, Heidelberg, Germany. 1, 2, 8, 11, 12

23. R. El Bansarkhani and J. Buchmann. High performance lattice-based CCA-secure encryption. Cryptology ePrint Archive, Report 2015/042, 2015. http://eprint.iacr.org/2015/042. 6

24. Y. Elias, K. E. Lauter, E. Ozman, and K. E. Stange. Ring-LWE cryptography for the number theorist. Cryptology ePrint Archive, Report 2015/758, 2015. http://eprint.iacr.org/2015/758. 9

25. S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany. 9

26. 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. 2

27. T. Güneysu, V. Lyubashevsky, and T. Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In E. Prouff and P. Schaumont, editors, *Cryptographic Hardware and Embedded Systems – CHES 2012*, volume 7428 of *Lecture Notes in Computer Science*, pages 530–547. Springer, 2012. 1, 2, 4, 8, 11

28. T. Güneysu, T. Oder, T. Pöppelmann, and P. Schwabe. Software speed records for lattice-based signatures. In *Post-Quantum Cryptography*, volume 7932 of *LNCS*, pages 67–82. Springer, 2013. 11

29. Q. Guo, T. Johansson, and P. Stankovski. Coded-bkw: Solving lwe using lattice codes. In R. Gennaro and M. Robshaw, editors, *CRYPTO 2015*, volume 9215 of *LNCS*, pages 23–42. Springer, 2015. 9

30. T. Ishiguro, S. Kiyomoto, Y. Miyake, and T. Takagi. Parallel gauss sieve algorithm: Solving the SVP challenge over a 128-dimensional ideal lattice. In H. Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 411–428, Buenos Aires, Argentina, Mar. 26–28, 2014. Springer, Heidelberg, Germany. 9

31. J. Katz and N. Wang. Efficiency improvements for signature schemes with tight security reductions. In S. Jajodia, V. Atluri, and T. Jaeger, editors, *ACM CCS 03*, pages 155–164, Washington D.C., USA, Oct. 27–30, 2003. ACM Press. 2, 5, 7

32. P. Kirchner and P.-A. Fouque. An improved bkw algorithm for lwe with applications to cryptography and lattices. In R. Gennaro and M. Robshaw, editors, *CRYPTO 2015*, volume 9215 of *LNCS*, pages 43–62. Springer, 2015. 9

33. T. Laarhoven. Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In R. Gennaro and M. Robshaw, editors, *CRYPTO 2015*, volume 9215 of *Lecture Notes in Computer Science*, pages 3–22. Springer Berlin Heidelberg, 2015. 9

34. A. Lenstra, H. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:515–534, 1982. 9

35. R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In A. Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339, San Francisco, CA, USA, Feb. 14–18, 2011. Springer, Heidelberg, Germany. 9

36. M. Liu and P. Nguyen. Solving BDD by enumeration: An update. In E. Dawson, editor, *CT-RSA 2013*, volume 7779 of *LNCS*, pages 293–309. Springer, 2013. 9

37. V. Lyubashevsky. Lattice signatures without trapdoors. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, 2012. 2, 15, 16

38. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany. 4, 9

39. D. Micciancio and P. Voulgaris. Faster exponential time algorithms for the shortest vector problem. In M. Charika, editor, *21st SODA*, pages 1468–1480, Austin, Texas, USA, Jan. 17–19, 2010. ACM-SIAM. 9

40. D. Pointcheval and T. Johansson, editors. *EUROCRYPT 2012*, volume 7237 of *LNCS*, Cambridge, UK, Apr. 15–19, 2012. Springer, Heidelberg, Germany. 12

41. D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000. 2

42. M. Schneider. Sieving for shortest vectors in ideal lattices. In A. Youssef, A. Nitaj, and A. E. Hassanien, editors, *AFRICACRYPT 13*, volume 7918 of *LNCS*, pages 375–391, Cairo, Egypt, June 22–24, 2013. Springer, Heidelberg, Germany. 9

A   Extended Definitions and Security Notions

A.1   Syntax, Functionality, and Security of Signature Schemes

A signature scheme with key space $\mathcal{K}$, message space $\mathcal{M}$, and signature space $\mathcal{S}$, is a tuple $\Sigma = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ of algorithms defined as follows.

– The (probabilistic) key generation algorithm on input the security parameter 1^λ returns a key pair $(\mathsf{sk}, \mathsf{pk}) \in \mathcal{K}$. We write $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$ and call sk the secret or signing key and pk the public or verification key.
– The (probabilistic) signing algorithm takes as input a signing key sk, a message $\mu \in \mathcal{M}$, and outputs a signature $\sigma \in \mathcal{S}$. We write $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, \mu)$.
– The verification algorithm, on input a verification key pk, a message $\mu \in \mathcal{M}$, and a signature $\sigma \in \mathcal{S}$, returns a bit b: if $b = 1$ we say that the algorithm accepts, otherwise we say that it rejects. We write $b \leftarrow \mathsf{Verify}(\mathsf{pk}, \mu, \sigma)$.

We require (perfect) correctness of the signature scheme: for every security parameter λ, every choice of the randomness of the probabilistic algorithms, every key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$, every message $\mu \in \mathcal{M}$, and every signature $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, \mu)$, $\mathsf{Verify}(\mathsf{pk}, \mu, \sigma) = 1$ holds.

We target the standard security requirement for signature schemes, namely *unforgeability under chosen-message attack* (ufcma). The corresponding experiment involving an adversary

$\mathcal{A}$ against a signature scheme Σ is depicted in Figure 3. Since we prove security of the scheme presented in Section 3 in the random oracle model, we reproduce a corresponding ufcma experiment which grants $\mathcal{A}$ access to a random oracle H. Given the experiment, we say that a signature scheme Σ is (t, q_s, q_h, ϵ)-*unforgeable under chosen-message attack* if every adversary $\mathcal{A}$ which runs in time t and poses at most q_s queries to the signing oracle and q_h queries to the random oracle has advantage

$$\mathrm{Adv}_\Sigma^{\mathsf{ufcma}}(\mathcal{A}) = \Pr\left[\mathrm{Expt}_{\Sigma,\mathcal{A}}^{\mathsf{ufcma}} = 1\right] \leq \epsilon .$$

$\mathrm{Expt}_{\Sigma,\mathcal{A}}^{\mathsf{ufcma}}(1^\lambda) :$

  1   $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KeyGen}(1^\lambda)$

  2   $(\mu^*, \sigma^*) \leftarrow \mathcal{A}(1^\lambda, \mathsf{pk})^{\mathcal{O}_{\mathsf{Sign}}(\cdot), H(\cdot)}$

  3   If $\mathsf{Verify}(\mathsf{pk}, \mu^*, \sigma^*) = 1 \land \mu^* \notin \mathcal{Q}_S$:

  4   Return 1

  5   Else: Return 0

If $\mathcal{A}$ queries $\mathcal{O}_{\mathsf{Sign}}(\mu) :$

  6   $\mathcal{Q}_S \leftarrow \mathcal{Q}_S \cup \{\mu\}$

  7   $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, \mu)$

  8   Return σ to $\mathcal{A}$

Fig. 3: Security experiment of unforgeability under chosen-message attack for an adversary $\mathcal{A}$ against a signature scheme $\Sigma = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ in the random oracle model (i.e., all parties including $\mathcal{A}$ have access to a public function H with uniformly distributed output).

B   Security Reduction of our Signature Scheme

In this section, we provide the full proof of Theorem 1. To this end, we state and prove a number of auxiliary statements. First, we recap the rejection sampling. We use the lemma as stated by Lyubashevsky in [37, Lemma 4.6].

Lemma 1 (Rejection sampling). *For $n > 0$ let X be a probability distribution on $\mathbb{Z}^n$ and let Y be a probability distribution on $V \subset \mathbb{Z}^n$ with density functions $f_X : \mathbb{Z}^n \to (0, 1)$ and $f_Y : V \to (0, 1)$. Furthermore, let $\{Z_v\}_{v \in V}$ be a family of probability distributions with index $v \in V$ and density functions $f_{Z_v} : \mathbb{Z}^n \to (0, 1)$, such that for almost all $v \in V$ there exists $M \in \mathbb{R}$ such that*

$$Pr\left[M \cdot f_{Z_v}(z) \geq f_X(z)\right] \geq 1 - 2^{-n}, \tag{1}$$

where the probability is taken over random choice of $z \leftarrow X$. Then the output distributions of the following two algorithms are statistically close:

1. *pick $v \leftarrow Y, z \leftarrow Z_v$, output (z, v) with probability $\min\left(\frac{f_X(z)}{M f_{Z_v}(z)}, 1\right)$.*
2. *pick $v \leftarrow Y, z \leftarrow X$, output (z, v) with probability $\frac{1}{M}$.*

The next theorem and proof are inspired by [9, Lemma 3].

Lemma 2. *Let $n, d, B,$ and q be as in Theorem 1 and let $a_1, a_2 \in \mathcal{R}_q^\times$. Furthermore, assume that $L \cap \left[-2^{d+1}, 2^{d+1}\right]^{2n} \leq 2^{(d+2)2n}/q^n$ holds for the lattice $L = \{\mathbf{v} \in \mathbb{Z}^{2n} \mid \mathbf{v} = \mathbf{A}\mathbf{w} \pmod{}$*

q), *for* $\mathbf{w} \in \mathbb{Z}^n\}$ *for instances as in Theorem 1. Then for all* $y_1 \in \mathcal{R}_{q,[B]}$ *and* $y_2 \leftarrow_\$ \mathcal{R}_{q,[B]}$ *we have*

$$\Pr\left[\lfloor \mathbf{A}\mathbf{y_1}\rceil_{d,q} = \lfloor \mathbf{A}\mathbf{y_2}\rceil_{d,q}\right] \leq \frac{2^{(d+1)2n}}{(2B+1)^n q^n},$$

where the probability is taken over random choices of y_2*. In particular we have* $\left|\{\lfloor \mathbf{A}y\rceil_{d,q} \mid y \in \mathcal{R}_{q,[B]}\}\right| \geq \frac{(2B+1)^n q^n}{2^{(d+2)2n}}.$

Proof: Let $y_1 \in \mathcal{R}_{q,[B]}$. We show that we can upper bound the size of the set $S = \{y_2 \in \mathcal{R}_{q,[B]} \mid \lfloor \mathbf{A}\mathbf{y_1}\rceil_{d,q} = \lfloor \mathbf{A}\mathbf{y_2}\rceil_{d,q}\}$ by $\frac{2^{(d+1)2n}}{q^n}$, which implies the statement of the theorem, since

$$\Pr\left[\lfloor \mathbf{A}\mathbf{y_1}\rceil_{d,q} = \lfloor \mathbf{A}\mathbf{y_2}\rceil_{d,q}\right] = \frac{|S|}{\left|\mathcal{R}_{q,[B]}\right|} = \frac{|S|}{(2B+1)^n},$$

where the probability is taken over $y_2 \leftarrow_\$ \mathcal{R}_{q,[B]}$*.* We define $M = \{y_1 - x \mid x \in S\}$ and $N = \{\mathbf{A}x \pmod{q} \mid x \in M\}$. It holds $|M| = |S|$. Furthermore, since $\mathbf{A}$ has trivial kernel (over $\mathbb{Z}_q$), $M \subseteq \mathcal{R}_{q,[2B]}$, and $\frac{q}{2} > 2B$ it holds that $|N| = |M|$. Hence, it suffices to bound $|N|$. We claim that $N \subseteq [-2^{d+1}, 2^{d+1}]^{2n}$. In order to see this, let $\mathbf{x} = \mathbf{A}(\mathbf{y_1} - \mathbf{z}) \pmod{q}$ with $\mathbf{z} \in S$ be an arbitrary element in N. Then by definition of S we have $\lfloor \mathbf{A}\mathbf{y_1}\rceil_{d,q} = \lfloor \mathbf{A}\mathbf{z}\rceil_{d,q}$ and thus $\lfloor \mathbf{x}\rceil_{d+1,q} = \lfloor \mathbf{A}\mathbf{y_1} - \mathbf{A}\mathbf{z}\rceil_{d+1,q} = \mathbf{0}$, which means $\mathbf{x} \in [-2^{d+1}, 2^{d+1}]^{2n}$.
In addition, by definition N is a subset of the lattice L. Therefore, $N \subseteq L \cap [-2^{d+1}, 2^{d+1}]^{2n}$ and thus by the assumptions of the theorem we obtain $|N| \leq \left|L \cap [-2^{d+1}, 2^{d+1}]^{2n}\right| \leq \frac{2^{(d+1)2n}}{q^n}$.
∎

In the security reduction of Theorem 1 we need to bound the probability that for uniformly chosen polynomials a_1, a_2, t_1, t_2 there exist polynomials $s, e_1,$ and e_2 with "small" entries such that $a_i s + e_i = t_i \pmod{q}$, for $i = 1, 2$. In the following lemma we make this bound explicit.

Lemma 3. *Let* $n, m, q,$ *and* σ *be chosen as in Theorem 1. Furthermore, let* $a_1, a_2 \leftarrow_\$ \mathcal{R}_q^\times$ *and* $t_1, t_2 \leftarrow_\$ \mathcal{R}_q$*. It holds that*

$$\Pr\left[\forall i \in \{1,2\} \exists s, e_1, e_2 \in \mathcal{R}_{q,[k\sigma]} : a_i s + e_i = t_i \pmod{q}\right] \leq \frac{(2k\sigma+1)^{3n}}{q^{2n}},$$

where the probability is taken over random choices of $a_1, a_2, t_1,$ *and* t_2*.*

Remark 1. By [37, Lemma 4.4] we know that $\forall k > 0$ and $z \leftarrow D_\sigma$ it holds that $\Pr\left[|z| > k\sigma\right] \leq 2e^{-k^2/2}$. Hence, the probability that all coefficients of $s \leftarrow D_\sigma^n$ and $e_1, e_2 \leftarrow D_\sigma^n$ are in $[-k\sigma, k\sigma]$ is at least $\left(1 - 2e^{-k^2/2}\right)^{3n}$. That means, for $k = 14$ the probability $\left(1 - 2e^{-k^2/2}\right)^{3n} \geq 1 - 2^{-\lambda}$ is overwhelming for our parameter choice given in Table 1.

Proof: [Lemma 3] Define $\mathbf{A} = \left(\mathrm{Rot}(a_1) \ \mathrm{Rot}(a_2)\right)^T$ and $\mathbf{t} = \left(\mathbf{t_1} \ \mathbf{t_2}\right)^T$. Since the polynomials a_1, a_2, t_1, t_2 are chosen uniformly random (and hence the coefficients are chosen uniformly random) the probability can be bound by the ratio of the number of possible vectors $\mathbf{A}s + \mathbf{e}$ with $\mathbf{s} \in [-k\sigma, k\sigma]^n$ and $\mathbf{e} \in [-k\sigma, k\sigma]^{2n}$, and the number of possible vectors for $\mathbf{t} \leftarrow_\$ \mathbb{Z}_q^m$,

16

i.e.,

$$\Pr\left[\,\exists\, s, e_1, e_2 \in \mathcal{R}_{q,[k\sigma]} \text{ such that } a_i s + e_i = t_i \pmod{q} \text{ for } i = 1, 2 : \right.$$
$$\left. a_1, a_2 \leftarrow_{\$} \mathcal{R}_q^{\times},\ t_1, t_1 \leftarrow_{\$} \mathcal{R}_q\right]$$
$$\leq \frac{|\{\mathbf{As} + \mathbf{e} \mid \mathbf{s} \in [-k\sigma, k\sigma]^n \text{ and } \mathbf{e} \in [-k\sigma, k\sigma]^{2n}\}|}{|\{\mathbf{t} \in \mathbb{Z}_q^{2n}\}|}.$$

It holds that

$$|\{\mathbf{t} \in \mathbb{Z}_q^{2n}\}| = q^{2n},$$
$$|\{\mathbf{As} \mid \mathbf{s} \in [-k\sigma, k\sigma]^n\}| = |\{\mathbf{s} \in [-k\sigma, k\sigma]^n\}| \leq (2k\sigma + 1)^n,$$
$$|\{\mathbf{e} \in [-k\sigma, k\sigma]^{2n}\}| \leq (2k\sigma + 1)^{2n}.$$

Thus,

$$|\{\mathbf{As} + \mathbf{e} \mid \mathbf{s} \in [-k\sigma, k\sigma]^n \text{ and } \mathbf{e} \in [-k\sigma, k\sigma]^{2n}\}| \leq (2k\sigma + 1)^n (2k\sigma + 1)^{2n}$$
$$= (2k\sigma + 1)^{3n}.$$

∎

In the security reduction of Theorem 1 we use the fact that if the reduction $\mathcal{D}$ is given two randomly chosen challenge tuples (a_1, t_1) and (a_2, t_2) then the forger $\mathcal{A}$ can produce a valid signature only with small probability. To prove this we utilize the following lemma and its corollary.

Lemma 4. *Let n, d, B, σ, and q be as in Theorem 1, and let $\delta \in \mathbb{Q}_{>0}$. Furthermore, let $L = \{\bar{\mathbf{z}} \in \mathbb{Z}^n \mid \mathbf{A}\bar{\mathbf{z}} = \mathbf{u} \pmod{q}\}$ be a lattice and assume that $|L \cap [-\delta, \delta]^n| = (2\delta + 1)^n / q^n$. Then for all $v_1, v_2 \in \mathcal{R}_q$, $c' \leftarrow_{\$} \{0, 1\}^{\kappa}$ with $F(c') = c$, $a_1, a_2 \leftarrow_{\$} \mathcal{R}_q^{\times}$, and $t_1, t_2 \leftarrow_{\$} \mathcal{R}_q$ it holds that*

$$\Pr\left[\forall\, i = 1, 2\ \exists\, z \in \mathcal{R}_{q,[\delta]}\ :\ \lfloor a_i z \rceil_{d,q} = v_i + \lfloor t_i c \rceil_{d,q} \pmod{q}\right]$$
$$\leq \frac{2^{dn}(2\delta + 1)^n}{q^{2n}},$$

where the probability is taken over random choices of c', a_1, a_2, t_1, and t_2.

Proof: Recall that the rounding operator $\lfloor \cdot \rceil_d$ essentially drops component-wise the least d bits of each coefficient. Thus, we can upper bound the probability in the theorem statement by the probability that there exists a $z \in \mathcal{R}_{q,[\delta]}$ such that $a_i z = v_i + \lfloor t_i c \rceil_{d,q} \pmod{q}$, for $i = 1, 2$, multiplied by the factor of 2^{dn}. This is because there are at most 2^{dn} possible values for $\left(a_1 z\ \ a_2 z\right)^T$ such that $\lfloor a_i z \rceil_{d,q} = v_i + \lfloor t_i c \rceil_{d,q} \pmod{q}$, for $i = 1, 2$, and the map defined by $\left(\text{Rot}(a_1)\ \ \text{Rot}(a_2)\right)^T$ is injective. It remains to show that for all v_1, v_2 it holds that

$$\Pr\left[\forall\, i = 1, 2\ \exists\, z \in \mathcal{R}_{q,[\delta]}\ :\ a_i z = v_i + t_i c \pmod{q} \text{ with } F(c') = c\right]$$
$$\leq \frac{(2\delta + 1)^n}{q^{2n}}, \tag{2}$$

where the probability is taken over random choices of c' and a_1, a_2, t_1, t_2. Define $\mathbf{u} = \begin{pmatrix} \mathbf{u_1} & \mathbf{u_2} \end{pmatrix}^T$ with $\mathbf{u_1} = \mathbf{v_1} + \lfloor \mathbf{t_1 c} \rfloor_{d,q}$ and $\mathbf{u_2} = \mathbf{v_2} + \lfloor \mathbf{t_2 c} \rfloor_{d,q}$. Furthermore, let $\mathbf{A} = \begin{pmatrix} \mathbf{A_1} & \mathbf{A_2} \end{pmatrix}^T$ with $\mathbf{A_1} = \mathrm{Rot}(a_1)$ and $\mathbf{A_2} = \mathrm{Rot}(a_2)$. Then Equation 2 is the same as

$$\Pr\left[\forall\ i = 1, 2\ \exists\ \mathbf{z} \in [-\delta, \delta]^n : \mathbf{A_i z} = \mathbf{v_i} + \lfloor \mathbf{t_i c} \rfloor_{d,q} \pmod{q}\right]$$
$$\leq \frac{(2\delta + 1)^n}{q^{2n}}, \tag{3}$$

where the probability is taken over random choices of $\mathbf{c'}, a_1, a_2, t_1, t_2$.

Since $a_1, a_2 \leftarrow_\$ \mathcal{R}_q^\times$ it holds that $\mathrm{rank}(\mathbf{A_1}) = \mathrm{rank}(\mathbf{A_2}) = n$. Hence, $\mathrm{rank}(\mathbf{A}) = n$. Thus, there exists a unique $\bar{\mathbf{z}} \in \mathbb{Z}_q^n$ such that $\mathbf{A_1} \bar{\mathbf{z}} = \mathbf{u_1} \pmod{q}$. We prove Equation (2), by pointing on two things: first, due to the injectivity of the matrices $\mathbf{A_1}$ and $\mathbf{A_2}$, the probability that $\mathbf{A_2} \bar{\mathbf{z}} = \mathbf{u_2} \pmod{q}$ is bounded by $1/q^n$. Secondly, the probability that $\bar{\mathbf{z}}$ is in $[-\delta, \delta]^n$ can be bounded by $(2\delta + 1)^n/q^n$ as we show next: assume that $\mathbf{A_2} \bar{\mathbf{z}} = \mathbf{u_2} \pmod{q}$. Define the set $S_u^\perp = \{\bar{\mathbf{z}} \in [-\delta, \delta]^n \mid \mathbf{A} \bar{\mathbf{z}} = \mathbf{u} \pmod{q}\}$. The matrix $\mathbf{A}$ defines the following shifted lattice $\Delta_{u,q}^\perp(\mathbf{A}) = \{\bar{\mathbf{z}} \in \mathbb{Z}^n \mid \mathbf{A} \bar{\mathbf{z}} = \mathbf{u} \pmod{q}\}$, with $\det(\Delta_{u,q}^\perp(\mathbf{A})) = \det(\Delta_{0,q}^\perp(\mathbf{A})) = q^n$. By assumption, we can bound the number of elements in $S_u^\perp$ by

$$|S_u^\perp| = |\Delta_{u,q}^\perp(\mathbf{A}) \cap [-\delta, \delta]^n| = \frac{(2\delta + 1)^n}{q^n}.$$

Thus, the probability that $\bar{\mathbf{z}}$ is in $[-\delta, \delta]^n$ can be bounded by $(2\delta + 1)^n/q^n$. ∎

We reformulate the lemma to obtain the following corollary useful for the proof of our main theorem.

Corollary 1. *Let n, d, B, σ, and q be as in Theorem 1, and let $\delta \in \mathbb{Q}_{>0}$. Then, for $\forall\ v_1, v_2 \in \mathcal{R}_q$ and $c' \leftarrow_\$ \{0,1\}^\kappa$, $a_1, a_2 \leftarrow_\$ \mathcal{R}_q^\times$, and $t_1, t_2 \leftarrow_\$ \mathcal{R}_q$ it holds that*

$$\Pr\left[\exists\ z \in \mathcal{R}_{q,[\delta]} : v_i = \lfloor a_i z \rfloor_{d,q} - t_i c \pmod{q}\ \textit{for}\ i = 1, 2\ \textit{and}\ F(c') = c\right]$$
$$\leq \frac{2^{dn}(2\delta + 1)^n}{q^{2n}},$$

where the probability is taken over random choices of c', a_1, a_2, t_1, t_2.

At last, we give the proof of Theorem 1.

Proof: Let $\mathcal{A}$ be an algorithm that runs in time $t_\mathcal{A}$, makes q_h hash queries and q_s signing queries, and forges a valid signature with probability $\varepsilon_\mathcal{A}$. We show how to build a distinguisher $\mathcal{D}$, that runs $\mathcal{A}$ internally, solving R-LWE$_{n,2,q,\sigma}$ in time $t_\mathcal{D}$ and with probability $\varepsilon_\mathcal{D}$ as in the theorem statement. To win the R-LWE game, $\mathcal{D}$ simulates the ufcma-experiment for $\mathcal{A}$. Since $\mathcal{A}$ expects a public key as input, $\mathcal{D}$ gets two R-LWE challenge tuples (a_1, t_1) and (a_2, t_2) and hands $\mathsf{pk} = (a_1, a_2, t_1, t_2)$ over to the forger. Then, it answers hash and signing queries as follows.

Hash queries: When a given query was already asked $\mathcal{D}$ returns the corresponding answer consistently, otherwise it chooses a fresh $c' \in \{0,1\}^\kappa$ uniformly at random, registers this value together with its input, hence returns c' to $\mathcal{A}$.

Signing queries: Upon a signature request on a message μ, $\mathcal{D}$ simulates a signature (z, c') on μ by implementing the following steps: it chooses uniformly random $c' \leftarrow_\$ \{0,1\}^\kappa$, a polynomial $z \leftarrow_\$ \mathcal{R}_{q,[B-U]}$, and computes $c = F(c')$. It then computes polynomials $w_i = a_i z - t_i c \pmod{q}$, for $i = 1, 2$, and checks whether $|\lfloor w_{ij} \rceil_{2^d}| < 2^{d-1} - L$ for all $j \in \{1, \ldots, n\}$, i.e., it checks if the d least significant bits of every coefficient of w_1 and w_2 are smaller than $2^{d-1} - L$; if not, then $\mathcal{D}$ repeats the whole process starting with choosing new values c' and z. Moreover, if the hash oracle H was queried before on $(\lfloor w_1 \rceil_{d,q}, \lfloor w_2 \rceil_{d,q} \mu)$, then $\mathcal{D}$ aborts the simulation. Otherwise, $\mathcal{D}$ defines $c' = H(\lfloor w_1 \rceil_{d,q}, \lfloor w_2 \rceil_{d,q}, \mu)$ and returns (z, c').

Eventually $\mathcal{A}$ outputs a forgery $(\tilde{z}, \tilde{c}')$ on a message $\tilde{\mu}$ for which it did not ask a signature to the signing oracle. Now $\mathcal{D}$ checks validity of the forgery: If $\mathsf{Verify}(\mathsf{pk}, \mu, (\tilde{z}, \tilde{c}')) = 1$ it returns 1, else it returns 0. We show that the responses to hash and signing queries provided by the simulation $\mathcal{D}$ are indistinguishable from the random oracle's and signing oracle's responses. The reduction emulates a random oracle via *lazy sampling*, hence its answers are truly random as long as some given hash value was not already set. However, in the latter case the reduction aborts. It remains to show that **(i)** the distributions of the simulated signatures and the ones produced by the signing algorithm (Figure 1) are statistically close and that **(ii)** the simulation does not abort too often.

(i) We apply Lemma 1 to show that the distribution of $z \in \mathcal{R}_{q,[B-U]}$ computed by the signing algorithm in Figure 1 is statistically close to the uniform distribution on $\mathcal{R}_{q,[B-U]}$. To apply Lemma 1 we take X to be the uniform distribution on $\mathcal{R}_{q,[B-U]}$ with $U = 14\sigma\sqrt{\omega}$ and $B = 14\sigma\sqrt{\omega}n$ as in Table 1. Furthermore, Z_v is the uniform distribution on the polynomials with coefficients in $[-B, B]^n + \mathbf{v}$. That means, $z \leftarrow Z_v$ can be written as $z = v + y$, where $y \leftarrow_\$ \mathcal{R}_{q,[B]}$. Each element v is a polynomial with coefficient representation $\mathbf{v} = \mathbf{Sc}$ with entries Gaussian distributed with standard deviation $\sigma\sqrt{\omega}$. Hence, the coefficients of v are bounded by $14\sigma\sqrt{\omega}$ with overwhelming probability. By definition of the density function f_X and $f_{Z_{sc}}$ Equation 1 holds with overwhelming probability for $1/M = 1/e \approx \left(\frac{2(B-U)+1}{2B+1}\right)^n$. Hence, the hypothesis of Lemma 1 are fulfilled.

(ii) Assume that the sign simulation samples an additional value y uniformly at random in $\mathcal{R}_{q,[B]}$. Furthermore, the simulation programs not only $c' = H(\lfloor w_1 \rceil_d, \lfloor w_2 \rceil_d, \mu)$ but also $c' = H(\lfloor a_1 y \rceil_d, \lfloor a_2 y \rceil_d, \mu)$. Sampling y does not influence (z, c'). It is clear that the probability of aborting during the orginial sign simulation is upper bounded by the abort probability during the sign simulation with the changes just described. The latter probability is the same as finding a collision. Hence, by Lemma 2, the probability that $\mathcal{D}$ will abort during the simulation of $\mathcal{A}$'s environment is upper bounded by $q_s(q_h + q_s)\frac{2^{(d+1)2n}}{(2B+1)^n q^n}$.

It remains to show how to lower bound $\mathcal{D}$'s distinguishing advantage in the R-LWE game with $\mathcal{A}$'s forging advantage against the scheme Σ. Assume w.l.o.g. that $\mathcal{A}$ asks the query $(\lfloor a_1\tilde{z} - t_1\tilde{c}\rceil_{d,q}, \lfloor a_2\tilde{z} - t_2\tilde{c}\rceil_{d,q}, \tilde{\mu})$ to the hash oracle. In the following, we distinguish between two cases: both t_1 and t_2 follow the R-LWE distribution—i.e, $t_1 = a_1 s + e_1 \pmod{q}$ and $t_2 = a_2 s + e_2 \pmod{q}$—or they are sampled uniformly at random in $\mathcal{R}_q$.

1st case, $t_1 = a_1 s + e_1$ and $t_2 = a_2 s + e_2$: The only possibilities that $\mathcal{D}$ outputs 0 are that $\mathcal{D}$ aborts when answering a signing query or that the algorithm $\mathcal{A}$ does not output a valid forgery. The latter occurs with probability $1 - \varepsilon_\mathcal{A}$. As described in Appendix B, the probability that $\mathcal{D}$ aborts during the simulation is $q_h \frac{2^{(d+1)2n}}{(2B+1)^n q^n}$. Thus, $\mathcal{D}$ returns 1 with probability at least $\varepsilon_\mathcal{A}\left(1 - q_s q_h \frac{2^{(d+1)2n}}{(2B+1)^n q^n}\right)$.

2nd case, $t_1, t_2 \leftarrow_\$ \mathcal{R}_q$**:** By Remark 1, Appendix B, we can bound the coefficients of the polynomials s, e_1, e_2 by 14σ with overwhelming probability. By Lemma 3 the probability that there exist polynomials with coefficient representation $\mathbf{s} \in [-14\sigma, 14\sigma]^n$ and $\mathbf{e_1}, \mathbf{e_2} \in [-14\sigma, 14\sigma]^n$ such that $t_i = a_i s + e_i \pmod q$ with $t_i \leftarrow_\$ \mathcal{R}_q$, for $i = 1, 2$, is smaller than or equal to $(28\sigma + 1)^{3n}/q^{2n}$. Assume t_1, t_2 are not of this form. By the assumption that a hash query on $(\lfloor a_1 \tilde{z} - t_1 \tilde{c} \rceil_{d,q}, \lfloor a_2 \tilde{z} - t_2 \tilde{c} \rceil_{d,q}, \tilde{\mu})$ is made by $\mathcal{A}$, where $(\tilde{c}', \tilde{z})$ is $\mathcal{A}$'s forgery on message $\tilde{\mu}$, it suffices to show that for every hash query $c' = H(v_1, v_2, \mu)$, $v_1, v_2 \in \mathcal{R}_q$, posed by $\mathcal{A}$ we can bound the probability that there exists a polynomial z such that (z, c') is a valid signature on μ. That means, it is enough to show that for $v_1, v_2 \in \mathcal{R}_q$ we can bound the probability that for $c' \leftarrow_\$ \{0, 1\}^\kappa$ there exists $z \in \mathcal{R}_q$ such that $||\mathbf{z}||_\infty \leq B - U$ and $v_i = \lfloor a_i z - t_i c \rceil_{d,q} \pmod q$, $i = 1, 2$. By Corollary 1 the probability that such a polynomial z exists is smaller than $\frac{2^{dn}(2(B-U)+1)^n}{q^{2n}}$. Thus, the probability that $\mathcal{A}$ forges a valid signature can be bound by $q_h \frac{2^{dn}(2(B-U)+1)^n}{q^{2n}}$. Hence, we can upper bound the probability that $\mathcal{D}$ returns 1 by $q_h \frac{2^{dn}(2(B-U)+1)^n}{q^{2n}} + \frac{(28\sigma+1)^{3n}}{q^{2n}}$.

Finally, $\mathrm{Adv}_{n,q,\chi}^{\mathsf{R\text{-}LWE}}(\mathcal{A}) = \varepsilon_\mathcal{D} \geq \varepsilon_\mathcal{A}\left(1 - \frac{q_s q_h 2^{(d+1)2n}}{(2B+1)^n q^n}\right) - \frac{q_h 2^{dn}(2B-2U+1)^n}{q^{2n}} - \frac{(28\sigma+1)^{3n}}{q^{2n}}$.

It remains to show that $\mathcal{A}$'s runtime $t_\mathcal{A}$ and $\mathcal{D}$'s runtime $t_\mathcal{D}$ are close. As $\mathcal{D}$ executes $\mathcal{A}$ as a subprocedure, we have $t_\mathcal{D} \geq t_\mathcal{A}$. The overhead for $\mathcal{D}$ is due to the extra steps needed to emulate the unforgeability game for $\mathcal{A}$, i.e., two oracle queries to get challenge tuples (a_1, t_1) and (a_2, t_2), plus the time necessary to answer hash and signing oracle queries. We already showed that the distributions of the signatures simulated by $\mathcal{D}$ and those obtained using the actual signing algorithm are statistically close; thus, when emulating the signing procedure $\mathcal{D}$ rejects a pair (c', z) with the same probability as when running algorithm Sign line 18 is executed. More precisely, simulating the signing process essentially consists of a number of polynomial multiplications, on average $\mathcal{O}(\kappa^2)$ per query, which leads to the approximative bound $t_\mathcal{D} \approx t_\mathcal{A} + \mathcal{O}(q_s \kappa^2 + q_h)$. ∎