

A preliminary version of this papers appears in the proceedings of PQCrypto 2017. This is the full version of the paper, updating the IACR ePrint Archive Record 2015/755.

Revisiting TESLA in the quantum random oracle model

Erdem Alkim¹, Nina Bindel², Johannes Buchmann², Özgür Dagdelen³, Edward Eaton^{4,5}, Gus Gutoski⁴, Juliane Krämer², and Filip Pawlega^{4,5}

¹ Ege University, Turkey erdemalkim@gmail.com

² Technische Universität Darmstadt, Germany

{nbindel, buchmann, jkraemer}@cdc.informatik.tu-darmstadt.de

³ BridgingIT GmbH, Germany oezdagdelen@googlemail.com

⁴ ISARA Corporation, Canada

{ted.eaton, gus.gutoski, filip.pawlega}@isara.com

⁵ University of Waterloo, Canada

May 4, 2017

Abstract. We study a scheme of Bai and Galbraith (CT-RSA'14), also known as TESLA. TESLA was thought to have a tight security reduction from the learning with errors problem (LWE) in the random oracle model (ROM). Moreover, a variant using chameleon hash functions was lifted to the quantum random oracle model (QROM). However, both reductions were later found to be flawed and hence it remained unresolved until now whether TESLA can be proven to be tightly secure in the (Q)ROM. In the present paper we provide an entirely new, tight security reduction for TESLA from LWE in the QROM (and thus in the ROM). Our security reduction involves the adaptive re-programming of a quantum oracle. Furthermore, we propose parameter sets targeting 128 bits of security against both classical and quantum adversaries and compare TESLA's performance with state-of-the-art signature schemes.

Keywords: Quantum Random Oracle, Post Quantum Cryptography, Lattice-Based Cryptography, Signature Scheme, Tight Security Reduction

1 Introduction

Our interest in the present paper is in a quantum-resistant signature scheme proposed by Bai and Galbraith [8]. Those authors argue the security of their scheme via reductions from the *learning with errors (LWE)* and the *short integer solutions (SIS)* problems in the random oracle model (ROM). This scheme

was subsequently studied by Alkim, Bindel, Buchmann, Dagdelen, and Schwabe under the name TESLA [4], who provided an alternate security reduction from the LWE problem only.

Since then, there have been several follow-up works on the Bai-Galbraith scheme [2, 4, 10, 61]. Most notably, a version of the scheme called ring-TESLA, whose security is based on the ring-LWE problem [2], has the potential to evolve into a practical, quantum-resistant signature scheme that might one day see widespread use as replacement for contemporary signature schemes such as ECDSA.

In what follows, we review the concepts of tightness and the quantum random oracle model as they relate to TESLA. We then list the contributions of the present paper and discuss related work by others.

1.1 Background

Security reduction and parameter choice. The security of digital signature schemes is often argued by reduction. A reductionist security argument typically proves a claim of the form, “any attacker \mathcal{A} who can break the scheme can be used to build an algorithm \mathcal{B} that solves some underlying hard computational problem”. Hence, the security gap can be determined; it measures how much extra work \mathcal{B} must perform in order to convert \mathcal{A} into solving the underlying hard problem. If the run-time and probability of success of \mathcal{B} are close to those of \mathcal{A} , *i.e.*, if the security gap is approximately 1, then the reduction is called *tight*. Achieving a small security gap, ideally a tight security reduction, is of theoretical interest in its own right, but it should also be an important consideration when selecting parameters for a concrete instantiation of a scheme. Specifically, the parameters of a signature scheme ought to be selected so that both (i) the effort needed to solve the underlying hard computational problem, and (ii) the security gap are taken into account. Hence, a tight security reduction is of advantage.

The need to instantiate schemes according to their security reductions and the role tight reductions play in these instantiations have been well argued by numerous authors. We refer the reader to [1, 21, 34] for a representative sample of these arguments.

The quantum random oracle model. Security arguments for the most efficient signature schemes—which therefore enjoy the most widespread real-world use—are typically presented in the ROM. (We refer to [37] by Kobitz and Menezes for discussion on why this might be the case.) The ROM postulates a truly random function that is accessible to attackers only through “black box” queries to an oracle for it—a random oracle. Any concrete proposal for a signature scheme must substitute a specific choice of hash function for the random oracle. An attacker armed with a quantum computer can be expected to evaluate that hash function in quantum superposition. Arguments that establish security even against such quantum-enabled attackers are said to hold in the quantum random oracle model (QROM).

It is conceivable that a signature scheme shown to be secure in the ROM may not be secure in the QROM. Thus, it is important that security arguments for quantum-resistant signature schemes hold not merely in the ROM, but also in the QROM.

Boneh *et al.* have proven that a security reduction in the ROM also holds in the QROM if it is *history-free* [16]. Unfortunately, many signature schemes have security reductions in the ROM that involve the *re-programming* of a random oracle; these reductions are not history-free. For these schemes, there remains a need to precisely clarify under what conditions these security reductions remain meaningful in the QROM.

Tightness in the QROM for TESLA. The security reduction presented by Bai and Galbraith for their signature scheme employs the Forking Lemma [53]. As such, it is non-tight and it involves re-programming, so it holds in the ROM but is not known to hold in the QROM.

As mentioned above, Alkim *et al.* presented an alternate security analysis for the Bai-Galbraith scheme, which they call TESLA. Their reduction is a tight reduction from LWE in the ROM. Moreover, those authors observed that their reduction can be made history-free at the cost of replacing a generic hash function with a chameleon hash function. It then follows from [16] that the history-free security reduction for TESLA holds also in the QROM. (Unfortunately, the use of a chameleon hash function would likely render any signature scheme too inefficient for widespread practical use.)

Unfortunately, a flaw in the original TESLA security reduction has been identified by the present authors. (The flaw was independently discovered by Chris Peikert.) This flaw is also present in several TESLA follow-up works, including ring-TESLA. As such, the status of the TESLA signature scheme and its derivative works has been open until now.

1.2 Our contribution

Our primary contributions are as follows:

New security reduction. We present a new security reduction from LWE to TESLA. Our new reduction is tight. It seems that the flaw in the original tight security reduction of TESLA does not admit a fix without a huge increase in the parameters; our new reduction is a significant re-work of the entire proof.

Security in the QROM with re-programming. Our new security reduction involves the adaptive re-programming of a random oracle and hence it is not history-free. Nevertheless, we show that it holds in the QROM by applying a seminal result from quantum query complexity due to Bennet, Bernstein, Brassard, and Vazirani [12]. It is possible that our approach can be abstracted so as to yield a general result on security reductions with re-programming in the QROM.

Our secondary contributions are as follows:

Parameter selection. We propose three sets of parameters for the concrete instantiation of TESLA: TESLA-0 and TESLA-1 targeting 96 and 128 bit security against a classical adversary, respectively; and TESLA-2, targeting 128 bits of security against a quantum adversary. All three parameter sets are chosen according to our (tight) security reduction.

The concrete parameter space admitted by our new security reduction is worse than that of previous reductions, but those previous reductions are either flawed or non-tight. Consequently, our proposed parameter sets lead to concrete instantiations of TESLA that are less efficient than previous proposals given in [4, 8, 61] that were not chosen according to the given security reduction.

Implementation. We provide a software implementation for the parameter sets TESLA-0 and TESLA-1. Our implementation targets Intel Haswell CPUs to provide a comparison of TESLA’s performance with other signature schemes with different security levels. Unfortunately, the TESLA-2 parameter set does not seem to admit an implementation that can take advantage of the same fast parallel arithmetic instructions available on modern processors that were used in our implementations of TESLA-0 and TESLA-1, and so we do not provide a software implementation for TESLA at this parameter set. See Section 7 for details.

Gaussian Heuristic. All previous analyses of the scheme from [8] relied on an assumption known as the “*Gaussian heuristic*”, which provides an estimate for the number of lattice points contained within a bounding region. This estimate has been well studied, and is not accurate in general — particularly for small bounding bodies. However, its validity has been conjectured for simple bodies and lattices on which LWE and SIS are defined without any published proof. Although our approach does not rely on the Gaussian heuristic, we prove in Section 8 that its use in other works is in fact justified.

1.3 Related work

Tightness from “lossy” keys. In order to avoid the non-tightness inherent in the use of the Forking Lemma, we take an approach that was introduced by Katz and Wang to obtain tightly-secure signatures from the decisional Diffie-Hellman problem [34].

The idea is to use the underlying hardness assumption to show that “real”, properly-formed public keys for the signature scheme are indistinguishable from “lossy”, malformed public keys. The task of forging a signature for a lossy key is then somehow proven to be intractable.

Any attacker must therefore fail to forge when given a lossy public key. Thus, any attacker who succeeds in forging a signature when given a real public key can be used to distinguish real keys from lossy keys, contradicting the underlying hardness assumption.

In the case of TESLA, the real keys are matrices A and $T = AS + E$ for some matrices S, E with small entries. (See Section 2.2 for a proper definition of these

matrices and the LWE problem.) We call these real keys LWE yes-instances. The lossy keys are LWE no-instances: matrices A, T selected uniformly at random, so that the existence of S, E as above occurs with only negligible probability. We prove that the task of forging a TESLA signature for lossy keys is intractable, so that any TESLA forger must be able to solve the decisional LWE problem.

A Fiat-Shamir transform for “lossy” identification schemes. The TESLA signature scheme could be viewed as the result of applying the Fiat-Shamir transform to a “lossy” identification scheme based on LWE. A tight security reduction for TESLA then follows from a general theorem of Abdalla, Fouque, Lyubashevsky, and Tibouchi (AFLT theorem) on the tight security of any signature scheme obtained in this way [1].

In order to leverage the AFLT theorem, one must propose an identification scheme and prove that it is lossy. Such a proof could be obtained by excerpting the relevant parts of our security reduction to establish the simulatability and lossiness properties of a suitably chosen identification scheme. Such an exercise might make our rather monolithic security reduction easier to digest by modularizing it and phrasing it in a familiar framework.

However, security reductions obtained by applying the AFLT theorem are guaranteed to hold only in the ROM. In order to fully recover our security reduction from this framework, one must first re-prove the AFLT theorem in the QROM. This limitation is due to the fact that the proof of the AFLT theorem involves adaptively re-programming a hash oracle. As such, it does not meet any known conditions for lifting a given proof from the ROM into the QROM.

Given that our security reduction in the QROM also involves the adaptive re-programming of a hash oracle, perhaps our approach could be mined for insights to establish the AFLT theorem in the QROM.

Other tightly-secure LWE or SIS signature schemes. Gentry, Peikert, and Vaikuntanathan present a signature scheme with a tight security reduction from SIS in the ROM using a trapdoor construction based on possessing a secret short basis of a lattice [29]. Boneh *et al.* observed that the security reduction for this scheme is history-free, and thus holds in the QROM [16].

Boyer and Li present a signature scheme with a tight security reduction from SIS in the *standard model* [18], also using a short basis trapdoor. Since standard model security reductions do not rely on any assumptions about a random oracle, these reductions hold in the QROM.

The use of a short-basis trapdoor in a signature scheme imposes an additional constraint on the concrete parameter space admitted by that scheme’s security reduction. This additional constraint on the parameters of short-basis trapdoor schemes seems to render them too inefficient for practical use. Since TESLA and its derivatives do not use a trapdoor construction, they do not suffer from this impediment.

Other than TESLA, we are aware of only one example of a signature scheme based on the Fiat-Shamir transform with a tight security reduction from LWE

or SIS. Prior to Bai and Galbraith, a variant of a scheme by Lyubashevsky [42] was shown to admit a tight security reduction in the ROM by Abdalla *et al.* as part of an illustration of the aforementioned AFLT theorem [1]. An artifact of this reduction required Abdalla *et al.* to increase the parameters of the scheme, rendering it too inefficient for practical use. As mentioned earlier, security reductions produced via the AFLT theorem are not known to hold in the QROM.

Re-programming a quantum oracle. Adaptive reprogramming of a quantum oracle has been addressed in some specific cases. Unruh considered a re-programmed quantum oracle in order to establish the security of a quantum position verification scheme [58]. It is not clear whether Unruh’s results apply to our setting.

Eaton and Song present an asymptotic result on re-programming in the QROM [28] in a context quite different from ours. Since their result is asymptotic, it does not allow for concrete parameter selection, for which the tightness of the reduction needs to be explicit.

Our approach to re-programming is independent of these previous works, though some works—such as [16,28]—do draw upon the same result by Bennet *et al.* [12] that we employ. To our knowledge we are the first to present progress on re-programming in the QROM in the context of a cryptographic scheme with potential for quantum-resistant standardization.

A note on “lattice-based” cryptography. Part of the allure of cryptosystems based on LWE or SIS is that those problems enjoy worst-case to average-case reductions from fundamental problems about lattices such as the *approximate shortest independent vectors problem (SIVP)* or the *gap shortest vector problem (GapSVP)*. (See Regev [54] or the survey of Peikert [51] and the references therein.)

These reductions suggest that the ability to solve LWE or SIS on randomly chosen instances implies the ability to solve SIVP or GapSVP, even on the hardest instances. Indeed, cryptosystems based on LWE or SIS are often referred to as *lattice-based* cryptosystems, suggesting that the security of these cryptosystems ultimately rests upon the worst-case hardness of these lattice problems.

However, as observed by Chatterjee, Kobitz, Menezes, and Sarkar, existing worst-case to average-case reductions for LWE and SIS are highly non-tight [21]. We are not aware of a proposal for a concrete instantiation of a cryptosystem based on LWE or SIS with the property that the proposed parameters were selected according to such a reduction. Instead, it is common to instantiate such cryptosystems based on the best known algorithms for solving LWE or SIS. (In addition to TESLA, see for example [5,17].)

For TESLA, we take care to instantiate the scheme according to its security reduction from LWE. However, we are unable to instantiate TESLA according to reductions from underlying lattice problems, due to the non-tightness of these reductions.

2 Preliminaries

In this section we clarify our notation used throughout the paper. Furthermore, we recall the learning with errors problem and give a short overview of used terms in quantum computing.

2.1 Notation

Integer scalars are denoted using Roman letters and if not stated otherwise, q is a prime integer in this paper. For any positive integer n the set \mathbb{Z}_n of integers modulo n is represented by $\{-\lfloor(n-1)/2\rfloor, \dots, \lfloor n/2\rfloor\}$. Fix a positive integer d and define the functions $[\cdot], [\cdot]_L : \mathbb{Z} \rightarrow \mathbb{Z}$ as follows. For any integer x let $[x]_L$ denote the representative of x in \mathbb{Z}_{2^d} , *i.e.*, $x = [x]_L \pmod{2^d}$, and let $[x] = (x - [x]_L)/2^d$. Informally, $[x]_L$ is viewed as the *least significant bits* of x and $[x]$ is viewed as the *most significant bits* of x . The definitions are easily extended to vectors by applying the operators for each component. An integer vector y is *B-short* if each entry is at most B in absolute value.

Vectors with entries in \mathbb{Z}_q are viewed as column vectors and denoted with lowercase Roman letters in sans-serif font, *e.g.*, y, z, w . Matrices with entries in \mathbb{Z}_q are denoted with uppercase Roman letters in sans-serif font, *e.g.*, A, S, E . The transpose of a vector or a matrix is denoted by v^T or M^T , respectively. We denote by $\|v\|$ the Euclidean norm of a vector v , and by $\|v\|_\infty$ its infinity norm. All logarithms are base 2. A function is called negligible in the security parameter λ , denoted by $\text{negl}(\lambda)$, if it decreases faster than the inverse of every polynomial in λ , for sufficiently large λ .

The centered discrete Gaussian distribution for $x \in \mathbb{Z}$ with standard deviation σ is defined to be $\mathcal{D}_\sigma = \rho_\sigma(x)/\rho_\sigma(\mathbb{Z})$, where $\sigma > 0$, $\rho_\sigma(x) = \exp(-\frac{x^2}{2\sigma^2})$, and $\rho_\sigma(\mathbb{Z}) = 1 + 2 \sum_{x=1}^{\infty} \rho_\sigma(x)$.

For a finite set S , we denote sampling the element s uniformly from S with $s \leftarrow_{\S} \mathcal{U}(S)$ or simply $s \leftarrow_{\S} S$. Let χ be a distribution over \mathbb{Z} , then we write $x \leftarrow \chi$ if x is sampled according to χ . Moreover, we denote sampling each coordinate of a matrix $A \in \mathbb{Z}^{m \times n}$ with distribution χ by $A \leftarrow \chi^{m \times n}$ with $m, n \in \mathbb{Z}_{>0}$. For an algorithm \mathcal{A} , the value $y \leftarrow \mathcal{A}(x)$ denotes the output of \mathcal{A} on input x ; if \mathcal{A} uses randomness then $\mathcal{A}(x)$ is a random variable. \mathcal{A}^χ denotes that \mathcal{A} can request samples from the distribution χ .

2.2 The Learning with Errors Problem

In the following we recall the decisional learning with errors problem (LWE) and define its matrix variant (M-LWE) ⁶.

Definition 1 (Learning with Errors Problem). *Let $n, m, q > 0$ be integers and χ be a distribution over \mathbb{Z} . For $s \leftarrow_{\S} \mathcal{U}(\mathbb{Z}_q^n)$ define $\mathcal{A}_{s, \chi}$ to be the distribution*

⁶ Note that, against common terminology in papers about module lattices, we do not mean *Module LWE* by the abbreviation M-LWE, but *Matrix LWE*.

that samples $\mathbf{a} \leftarrow_{\S} \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow \chi$ and then returns $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + \mathbf{e}) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. The decisional learning with errors problem $LWE_{n,m,q,\chi}$ is (t, ε) -hard if for any algorithm \mathcal{D} , running in time t and making at most m queries to the distribution $\mathcal{A}_{\mathbf{s},\chi}$, it holds that

$$\left| \Pr \left[\mathcal{D}^{\mathcal{A}_{\mathbf{s},\chi}}(\cdot) = 1 \right] - \Pr \left[\mathcal{D}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q)}(\cdot) = 1 \right] \right| \leq \varepsilon.$$

We can also write m LWE instances to a secret $\mathbf{s} \in \mathbb{Z}_q^n$ as $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q})$ with $\mathbf{A} \leftarrow_{\S} \mathbb{Z}_q^{m \times n}$ and $\mathbf{e} \leftarrow \chi^m$.

The security of the signature scheme covered in this paper is based on the matrix version of LWE (M-LWE) defined in the following.

Definition 2 (Matrix Learning with Errors Problem). Let $n, n', m, q > 0$ be integers and χ be a distribution over \mathbb{Z} . Define $\mathcal{A}_{\mathbf{S},\chi}$ to be the distribution that, for $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_{n'})$ with $\mathbf{s}_1, \dots, \mathbf{s}_{n'} \leftarrow_{\S} \mathcal{U}(\mathbb{Z}_q^n)$, samples $\mathbf{a} \leftarrow_{\S} \mathbb{Z}_q^n$ and $e_1, \dots, e_{n'} \leftarrow \chi$ and then returns $(\mathbf{a}^T, \mathbf{a}^T \mathbf{S} + (e_1, \dots, e_{n'})) \in \mathbb{Z}_q^{1 \times n} \times \mathbb{Z}_q^{1 \times n'}$. The matrix decisional learning with errors problem $M-LWE_{n,n',m,q,\chi}$ is (t, ε) -hard if for any algorithm \mathcal{D} , running in time t and making at most m queries to the distribution $\mathcal{A}_{\mathbf{S},\chi}$, it holds that

$$\left| \Pr \left[\mathcal{D}^{\mathcal{A}_{\mathbf{S},\chi}}(\cdot) = 1 \right] - \Pr \left[\mathcal{D}^{\mathcal{U}(\mathbb{Z}_q^n \times \mathbb{Z}_q^{n'})}(\cdot) = 1 \right] \right| \leq \varepsilon.$$

As before, m M-LWE samples to the secret matrix $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_{n'}) \in \mathbb{Z}_q^{n \times n'}$ can be written as $(\mathbf{A}, \mathbf{A}\mathbf{S} + \mathbf{E}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^{m \times n'}$ with $\mathbf{A} \leftarrow_{\S} \mathbb{Z}_q^{m \times n}$ and $\mathbf{E} \leftarrow \chi^{m \times n'}$.

We call $(\mathbf{A}, \mathbf{T}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^{m \times n'}$ a **yes-instance** if there exists an $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_{n'})$ with $\mathbf{s}_1, \dots, \mathbf{s}_{n'} \in \mathbb{Z}_q^n$ and (\mathbf{A}, \mathbf{T}) are m M-LWE samples from the distribution $\mathcal{A}_{\mathbf{S},\chi}$. Otherwise, *i.e.*, when $(\mathbf{A}, \mathbf{T}) \leftarrow_{\S} \mathcal{U}(\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^{m \times n'})$, we call (\mathbf{A}, \mathbf{T}) a **no-instance**.

Theorem 1. If $LWE_{n,m,q,\chi}$ is $(\varepsilon/n', t)$ -hard then $M-LWE_{n,n',m,q,\chi}$ is (ε, t) -hard.

Intuitively, the reduction loss exists since an adversary that can solve LWE has n' possibilities to solve M-LWE (see also [8, 17, 52]). The proof follows similar arguments as given in [52].

We note that the hardness of LWE (resp., M-LWE) is retained even if all coordinates of the secret vector \mathbf{s} are sampled according to the error distribution χ , known as the “normal form” [6, 43]. We use the notation $LWE_{n,m,q,\sigma}$ if χ is distributed according to \mathcal{D}_{σ} . The LWE assumption comes with a worst-to-average-case reduction [20, 50, 54]; breaking certain average instances of LWE allows one to break all instances of certain standard lattice problems (namely GapSVP and SIVP).

2.3 Quantum Information and Quantum Oracles

We assume familiarity with the fundamentals of quantum information, such as the Dirac ket notation $|\cdot\rangle$ for pure quantum states and the density matrix formalism for mixed quantum states. (Recall that a mixed state can be viewed as

a probabilistic mixture of pure states.) For background on quantum information the reader is referred to the books [35, 49].

The *trace norm* $\|\cdot\|_{\text{Tr}}$ (also known as the *Schatten-1 norm* $\|\cdot\|_1$) is a matrix norm that generalizes the concept of statistical distance from probability distributions to quantum states. For example, the quantity $\|\rho - \rho'\|_{\text{Tr}}$ captures the physically observable difference between two quantum states represented by density matrices ρ, ρ' .

A *channel* is a physically realizable map on mixed quantum states - the quantum analogue of a stochastic map applied to a classical probability distribution. Examples of channels include (i) *unitary channels*, which preserve purity of input states and act on classical basis states “in superposition”, and (ii) *classical channels*, which measure input states in the standard basis and whose output is a probabilistic mixture of classical basis states. In this paper, hash oracles are unitary channels and sign oracles are classical channels.

Standard formalism for a quantum oracle for a hash function $H : X \rightarrow Y$ specifies that the user has access to a unitary channel that applies the linear map $|x\rangle|y\rangle \mapsto |x\rangle|y + H(x)\rangle$ on standard basis states, where $x \in X$ and $y \in Y$. It is understood that the range Y of $H(\cdot)$ has a ‘+’ operation defined on it with the property that $y + y' + y' = y$ for all $y, y' \in Y$, so that the unitary channel is its own inverse. Typically, one simply imagines that elements of y are written as binary strings, and the ‘+’ operation is the bitwise exclusive-OR.

3 The Signature Scheme TESLA

In this section, we present the LWE-based signature scheme TESLA. Its original construction was proposed in 2014 by Bai and Galbraith [8]. It was later revisited by Dagdelen *et al.* [61] and by Alkim *et al.* [4].

TESLA’s key generation, sign, and verify algorithms are listed informally in Algorithms 1, 2, and 3. More formal listings of these algorithms are given in Figure 1 in Section 5. Our proposed concrete parameter sets are derived in Section 5 and listed in Table 1.

Algorithm 1 KeyGen

Input: A .

Output: Public key T , secret key (S, E) .

- 1: Choose entries of $S \in \mathbb{Z}_q^{n \times n'}$ and $E \in \mathbb{Z}_q^{m \times n'}$ from \mathcal{D}_σ .
 - 2: If E has a row whose h largest entries sum to L or more then retry at step 1.
 - 3: If S has a row whose h largest entries sum to L_S or more then retry at step 1.
 - 4: $T \leftarrow AS + E$.
 - 5: Return public key T and secret key (S, E) .
-

Parameters and Notation. TESLA is parameterized by positive integers $q, m, n, n', h, d, B, L, L_S, U$, a positive real σ , a hash oracle $H(\cdot)$, and the publicly available matrix $A \leftarrow_{\S} \mathbb{Z}_q^{m \times n}$. Let \mathbb{H} denote the set of vectors $\mathbf{c} \in \{-1, 0, 1\}^{n'}$ with exactly h nonzero entries. For simplicity we assume that the hash oracle $H(\cdot)$ has range \mathbb{H} , *i.e.*, we ignore the encoding function F , cf. Table 1. We call an integer vector \mathbf{w} *well-rounded* if \mathbf{w} is $(\lfloor q/2 \rfloor - L)$ -short and $\lceil \mathbf{w} \rceil$ is $(2^d - L)$ -short.

Deterministic signatures. We recommend that any implementation of TESLA employ standard techniques to achieve deterministic signatures. In particular, the key generation algorithm should produce a random seed as part of the secret key. The sign algorithm should use this seed, along with the input message μ , to derive a deterministic sequence of pseudorandom data that dictate the “random” choices of Algorithm 2. See Section 6 on implementation for further details.

Fixed-weight hash outputs. Any implementation of the hash oracle $H(\cdot)$ will require an encoding function that embeds the output of a concrete hash function such as SHA-256 into the set \mathbb{H} (see [31] for more information on embedding functions of this type). Naturally, the output length of the underlying hash function should be large enough so as to preclude collision attacks.

Additional checks in KeyGen and Sign. In contrast to earlier proposals [8, 61], we add two additional checks. The first one is the check in Line 3 in Algorithm 1. It ensures that no coefficient of the matrix S is too large, which allows for more concrete bounds during the security reduction. The parameter L_S is chosen such that the probability of rejecting S is smaller than $2^{-\lambda}$, cf. Section 5. The second additional check is in Line 5 in Algorithm 2. To ensure correctness of the scheme, it checks that the absolute value of each coordinate of $A\mathbf{y} - E\mathbf{c}$ is less or equal than $\lfloor q/2 \rfloor - L$.

Algorithm 2 Sign

Input: Message μ , secret key (S, E) .

Output: Signature (\mathbf{z}, \mathbf{c}) .

- 1: Choose \mathbf{y} uniformly at random among B -short vectors from \mathbb{Z}_q^n .
 - 2: $\mathbf{c} \leftarrow H(\lceil A\mathbf{y} \rceil, \mu)$.
 - 3: $\mathbf{z} \leftarrow \mathbf{y} + S\mathbf{c}$.
 - 4: If \mathbf{z} is not $(B - U)$ -short then retry at step 1.
 - 5: If $A\mathbf{y} - E\mathbf{c}$ is not well-rounded then retry at step 1.
 - 6: Return signature (\mathbf{z}, \mathbf{c}) .
-

Algorithm 3 Verify

Input: Message μ , public key (A, T) , purported signature (z, c) .

Output: “Accept” or “reject”.

- 1: If z is not $(B - U)$ -short then reject.
 - 2: If $H([Az - Tc], \mu) \neq c$ then reject.
 - 3: Accept.
-

4 Security Reduction for TESLA

Our main theorem on the security of TESLA informally states that as long as M-LWE can not be solved in time t and with success probability ε then no adversary \mathcal{A} exists that can forge signatures of TESLA in time t' and with success probability ε' , if \mathcal{A} is allowed to make at most q_h hash and q_s sign queries. The main theorem is as follows.

Theorem 2 (Security of TESLA). *Let $q, m, n, n', h, d, B, L, L_S, U, \sigma, \lambda, \kappa$ be TESLA parameters that are convenient⁷ (according to Definition 3 in Appendix D) and that satisfy the bounds in Table 1.*

If M-LWE is (t, ε) -hard then TESLA is existentially $(t', \varepsilon', q_h, q_s)$ -unforgeable against adaptively chosen message attacks with $t' \approx t$ in (i) the quantum random oracle model with

$$\varepsilon' < \varepsilon + \frac{3}{2^\lambda} + \frac{2^{m(d+1)+3\lambda+1}}{q^m} (q_h + q_s)^2 q_s^3 + 2(q_h + 1) \sqrt{\frac{1}{2^h \binom{n'}{h}}}, \quad (1)$$

and in (ii) the classical random oracle model with

$$\varepsilon' < \varepsilon + \frac{3}{2^\lambda} + \frac{2^{m(d+1)+3\lambda+1}}{q^m} (q_h + q_s)^2 q_s^3 + q_h \frac{1}{2^h \binom{n'}{h}}. \quad (2)$$

The proof of Theorem 2 is given in appendices A, B, C, D. Here we present a sketch of this proof and a selection of some intermediate results we feel are the most significant technical contributions of the present manuscript.

Let \mathcal{F} be a forger that forges signatures of the TESLA scheme with probability $\Pr[\text{forge}(A, T)]$, where $\text{forge}(A, T)$ denotes the event that \mathcal{F} forges a signature on input (A, T) , which is a yes- or a no-instance of LWE. We build an LWE-solver \mathcal{S} whose run time is close to that of \mathcal{F} and who solves LWE with success bias close to $\Pr[\text{forge}(A, T)]$. It then follows from the presumed hardness of LWE that $\Pr[\text{forge}(A, T)]$ must be small.

Given an LWE input (A, T) , the LWE-solver \mathcal{S} treats (A, T) as a TESLA public key; \mathcal{S} runs \mathcal{F} on input (A, T) and outputs “yes” if and only if \mathcal{F} succeeds in forging a TESLA signature.

⁷ It is not necessary that TESLA parameters be convenient in order to derive negligibly small upper bounds on ε' ; the definition of convenience merely facilitates a simplified statement of those bounds.

In order to run \mathcal{F} , the LWE-solver \mathcal{S} must respond in some way to \mathcal{F} 's quantum queries to the hash oracle and to \mathcal{F} 's classical queries to the sign oracle. Our description of \mathcal{S} includes a procedure for responding to these queries.

That \mathcal{S} solves LWE with success bias close to $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$ is a consequence of the following facts:

1. For yes-instances of LWE, the probability with which \mathcal{S} outputs “yes” is close to $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$.
2. For no-instances of LWE, \mathcal{F} successfully forges (and hence \mathcal{S} outputs “yes”) with only negligible probability.

4.1 Yes-Instances of LWE

We argue that \mathcal{S} 's responses to \mathcal{F} 's oracle queries are indistinguishable from the responses \mathcal{F} would receive from real oracles, from which it follows that \mathcal{S} reports “yes” with probability close to $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$.

Each time \mathcal{S} simulates a call to the sign oracle, it must “re-program” its simulated hash oracle on one input. Because \mathcal{F} is permitted to make quantum queries to the hash oracle, we must show that \mathcal{F} is unlikely to notice when a quantum random oracle has been re-programmed.

To this end, let \mathbb{Y} denote the set of vectors $\mathbf{y} \in \mathbb{Z}_q^n$ such that \mathbf{y} is B -short and define the following quantities for each choice of TESLA keys $(\mathbf{A}, \mathbf{T}), (\mathbf{S}, \mathbf{E})$:

$\text{nwr}(\mathbf{A}, \mathbf{E})$: The probability over $(\mathbf{y}, \mathbf{c}) \in \mathbb{Y} \times \mathbb{H}$ that $\mathbf{A}\mathbf{y} - \mathbf{E}\mathbf{c}$ is not well-rounded.
 $\text{coll}(\mathbf{A}, \mathbf{E})$: The maximum over all $\mathbf{w} \in \{[\mathbf{x}] : \mathbf{x} \in \mathbb{Z}_q^m\}$ of the probability over $(\mathbf{y}, \mathbf{c}) \in \mathbb{Y} \times \mathbb{H}$ that $[\mathbf{A}\mathbf{y} - \mathbf{E}\mathbf{c}] = \mathbf{w}$.

We prove the following in Appendix B.6, cf. Proposition 7.

Proposition 1 (Re-Programming in TESLA, Informal Statement). *The following holds for each choice of TESLA keys $(\mathbf{A}, \mathbf{T}), (\mathbf{S}, \mathbf{E})$, each hash oracle $\mathbf{H}(\cdot)$, and each $\gamma > 0$.*

Suppose the quantum state $\rho_{\mathbf{H}}$ was prepared by some party \mathcal{D} using t quantum queries to $\mathbf{H}(\cdot)$. Let $\mathbf{H}'(\cdot)$ be a hash oracle that agrees with $\mathbf{H}(\cdot)$ except on a small number of randomly chosen inputs (\cdot, μ) for each possible message μ . Let $\rho_{\mathbf{H}'}$ be the state prepared when \mathcal{D} uses hash oracle $\mathbf{H}'(\cdot)$ instead of $\mathbf{H}(\cdot)$.

Then $\|\rho_{\mathbf{H}'} - \rho_{\mathbf{H}}\|_{\text{Tr}} < \gamma$ except with probability at most

$$\frac{t^2}{\gamma^2} \cdot \frac{\text{coll}(\mathbf{A}, \mathbf{E})}{1 - \text{nwr}(\mathbf{A}, \mathbf{E})} \tag{3}$$

over the choice of inputs upon which $\mathbf{H}(\cdot)$ and $\mathbf{H}'(\cdot)$ differ.

We also prove bounds on $\text{nwr}(\mathbf{A}, \mathbf{E})$ and $\text{coll}(\mathbf{A}, \mathbf{E})$ that hold with high probability over the choice of TESLA keys $(\mathbf{A}, \mathbf{T}), (\mathbf{S}, \mathbf{E})$.

4.2 No-Instances of LWE

We argue that, except with negligibly small probability over the choice of hash oracle $H(\cdot)$ and LWE no-instance (A, T) , a TESLA forger cannot forge a signature for (A, T) without making an intractably large number of queries to the hash oracle.

To forge a signature for message μ , a forger must find a hash input (w, μ) whose output $c = H(w, \mu)$ has the property that there exists a $(B - U)$ -short $z \in \mathbb{Z}_q^n$ for which $[Az - Tc] = w$. Let $\mathbb{H}(w, A, T) \subset \mathbb{H}$ denote the set of all such c . A hash input (w, μ) is called *good* for $H(\cdot)$ and (A, T) if $H(w, \mu) \in \mathbb{H}(w, A, T)$. (Once a good hash input has been found, the forger must then somehow *find* the vector z witnessing this fact. For our purpose, we assume that the forger gets it for free.)

For each LWE no-instance (A, T) , a given hash input (w, μ) is good for $H(\cdot)$ and (A, T) with probability

$$\frac{\#\mathbb{H}(w, A, T)}{\#\mathbb{H}} \quad (4)$$

over the choice of hash oracle $H(\cdot)$. In Appendix C we argue that, except with negligibly small probability over the choice of $H(\cdot)$ and (A, T) , the fraction of hash inputs that are good is at most the expectation over LWE no-instances (A, T) of the ratio (4), maximized over all $w \in \{[x] : x \in \mathbb{Z}_q^m\}$. We then prove the following (*cf.* Proposition 9 in Appendix C).

Proposition 2 (Good Hash Inputs are Rare). *If the TESLA parameters are convenient (according to Definition 3 in Appendix D) then*

$$\mathbb{E}_{(A, T)} \left[\max_w \left\{ \frac{\#\mathbb{H}(w, A, T)}{\#\mathbb{H}} \right\} \right] \leq \frac{1}{\#\mathbb{H}}. \quad (5)$$

Thus, the fraction of good hash inputs is at most $1/\#\mathbb{H}$ except with vanishingly small probability over the choice of hash oracle $H(\cdot)$ and LWE no-instance (A, T) .

Since each hash input is good with a fixed probability independent of other hash inputs, the only way to discover a good input is via search through an unstructured space. It then follows from known lower bounds for quantum search over an unstructured space that the forger cannot find a good hash input—and thus a TESLA forgery—using only q_h quantum queries to the hash oracle.

5 Selecting Parameters for TESLA

In this section we propose parameter sets for TESLA. Table 1 illustrates our concrete choice of parameters and Table 2 gives the hardness of the corresponding LWE instances. We propose three parameter sets: TESLA-0 that targets the same (classical) bit security of 96 bit as the instantiation proposed in [61], called DEG^+ . TESLA-1 targets 128 bit of classical security and TESLA-2 targets 128 bit of security against quantum adversaries. Note that the parameter set DEG^+ was originally proposed to give 128 bit of security, *i.e.*, $\lambda = 128$, but due to new methods to estimate the bit security its bit security is now only 96 bit.

Table 1. Concrete instantiation TESLA-2 of 128 bit of security against classical and quantum adversaries, and TESLA-0 of 96 bit and TESLA-1 of 128 bit of security against classical adversaries; comparison with the instantiation proposed in [61], called DEG⁺, of 96 bit security (classically); sizes are given in kilo byte [KB]; sizes are theoretic sizes for fully compressed keys and signatures; for sizes used by our software see Table 3.

Parameter	Bound	DEG ⁺	TESLA-0	TESLA-1	TESLA-2
λ		128	96	128	128
κ		256	256	256	256
n		532	644	804	1300
n'		532	390	600	1036
m		840	3156	4972	4788
σ	$> 2\sqrt{n}$	43	55	57	73
L	$3h\sigma$ or $2.8h\sigma$, see Sec. 5.1	2322	5082	6703	17987
L_S	$14\sigma h$	-	25410	33516	89936
h	$2^h \binom{n'}{h} \geq 2^{3\lambda}$ (classically) $2^h \binom{n'}{h} \geq 2^{5\lambda}$ (quantumly)	18 -	33 -	42 -	- 88
B	$\geq 14n\sqrt{h}\sigma$	$2^{21} - 1$	$2^{22} - 1$	$2^{22} - 1$	$2^{24} - 1$
U	$\lceil 14\sqrt{h}\sigma \rceil$	2554	4424	5172	9588
d	$(1 - 2L/2^d)^m \geq 0.3$	23	25	26	27
q	satisfying the bound in Eq. 152, $\geq (2^{m(d+1)+4\lambda+1}(q_h + q_s)^2 q_s^3)^{1/m}$	$2^{29} - 3$	$2^{31} - 99$	$2^{31} - 19$	40582171961 $\approx 2^{35.24}$
δ_{KeyGen}	empirically, see Sec. 5.1	0.99	1	1	future work
δ_{Sign}		0.314	0.307	0.154	future work
H	$\{0, 1\}^* \rightarrow \{0, 1\}^\kappa$	SHA-256			
F	$\{0, 1\}^\kappa \rightarrow \mathbb{H}_{n', \omega}$	see [31]			
PRF_1	$\{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$	-	SHA-256		
PRF_2	$\{0, 1\}^\kappa \times \mathbb{Z} \rightarrow [-B, B]^n$	-	ChaCha20		
public-key size	$mn' \lceil \log_2(q) \rceil$	1 582	4 657	11 288	21 799
secret-key size	$(nn' + mn') \lceil \log_2(14\sigma) \rceil$	891	1 809	4 230	7 700
signature size	$n \lceil \log_2(2(B - U)) \rceil + \kappa$	1.4	1.8	2.3	4.0

5.1 Derivation of System Parameters

Our security reduction for TESLA minimizes the underlying assumptions which allows us to choose secure parameters from a greater set of choices compared to [8, 61]. More precisely, our parameters do not have to involve a hard instance of the SIS assumption as it was done by Bai and Galbraith [8] before. We summarize the bounds and conditions of each parameter in Table 1 and explicate the computation of some of the listed parameters in the following. Furthermore, we state the resulting key and signature sizes in the table.

Compared to [8, 61], we introduce the parameter n' as the column dimension of the secret matrices \mathbf{S} and \mathbf{E} to get more flexibility in the choice of parameters. The value of n' influences the parameters h (and hence B , U , q , and the encoding function F) and the size of the secret key.

Another important parameter of the signature scheme is the value L . In the original work [8], it is set to $L = 7h\sigma$, whereas it is set to $L = 3h\sigma$ in [61]. We choose L to be roughly $L = 2.8h\sigma$. We note that the smaller the value L , the higher the probability of acceptance in the signature algorithm (Line 9, Figure 1) becomes.

We add `checkS` to the key generation algorithm and the corresponding parameter L_S to bound $\|\text{Sc}\| \leq L_S$ in the security reduction. We determine the value L_S such that `S` is rejected only with negligibly small (in the security parameter λ) probability. Hence, we do not decrease the size of the key space further. We choose L_S to be $14h\sigma$.

The acceptance probabilities of a signature δ_{sign} and of a secret key (S, E) in Table 1 are determined experimentally.

To ensure both correctness and security of our signature scheme, we choose parameters with respect to our reduction, hence, we choose parameters such that $\epsilon' \approx \epsilon$ in Equation (1) and (2). We propose to choose $q_h \leq 2^\lambda$ and $q_s \leq 2^{\lambda/2}$, since a hash query is merely the evaluation of a publicly available function and hence the adversary can use all its computational power to make hash queries. The number of sign queries is somewhat limited since it involves more complicated operations. We refer to [36] (especially, Section 7) for further discussion.

5.2 Concrete Bit Security of TESLA

The security reduction given in Section 4 provides a reduction from the hardness of M-LWE and bounds *explicitly* the forging probability with the success probability of the reduction. More formally, let $\epsilon_{\mathcal{A}}$ and $t_{\mathcal{A}}$ denote the success probability and the run time of a forger \mathcal{A} against our signature scheme and let $\epsilon_{\mathcal{B}}$ and $t_{\mathcal{B}}$ denote analogous quantities for the reduction \mathcal{B} presented in the proof of Theorem 2. We say that M-LWE is η -bit hard if $t_{\mathcal{B}}/\epsilon_{\mathcal{B}} \geq 2^\eta$; and we say that the signature scheme is λ -bit secure if $t_{\mathcal{A}}/\epsilon_{\mathcal{A}} \geq 2^\lambda$.

Since we choose parameters such that $\epsilon_{\mathcal{B}} \approx \epsilon_{\mathcal{A}}$ and $t_{\mathcal{B}} \approx t_{\mathcal{A}}$, the bit hardness of the M-LWE instance is about the same as the bit security of our signature scheme instantiated as described below.

However, we lose $\lceil \log(n') \rceil$ bits of security due to the reduction from LWE to M-LWE. Hence, we have to choose LWE instances of $\lambda + \lceil \log(n') \rceil$ bit hardness. We explain how we choose those LWE instances in the following.

Estimation of the Bit Hardness of LWE Against Classical Adversaries. Albrecht, Player, and Scott [3] presented the *LWE-Estimator*, a software to estimate the *classical* hardness of LWE given the matrix dimension n , the modulus q , and the relative error rate $\alpha = \frac{\sqrt{2\pi}\sigma}{q}$. The original LWE-Estimator does not take the number of given LWE samples into account, leading to a more conservative choice of parameters for applications where only a restricted number of LWE samples is given. Recently, a fork of the original LWE-Estimator was presented [56, 57] that takes a given number of samples m into consideration. Since for TESLA only a restricted number of m LWE samples is given, we use [57] to estimate the classical bit hardness of our proposed LWE instances.

We summarize the classical estimations of the most efficient LWE solvers for our instances in Table 2, namely, the standard embedding or approach by Kannan [33], the dual embedding [8], and the decoding attack [7, 40, 41]. We also update the estimation of the bit hardness of the LWE instances proposed in [61].

Table 2. Estimation of the hardness of LWE instances given in TESLA-0, TESLA-1, and TESLA-2 against the decoding attack and the (dual and standard) embedding approach, in comparison to the parameter sets proposed by Dagdelen et al. [61], called DEG⁺; estimations are computed using the LWE-Estimator with a restricted number of samples [3, 57].

Attack	DEG ⁺	TESLA-0	TESLA-1	TESLA-2
Classical Hardness [bit]				
Decoding	156	110	142	204
Dual Embedding	96	110	142	205
Standard Embedding	164	111	143	205
Post-Quantum Hardness [bit]				
Decoding	73	74	98	146
Dual Embedding	61	71	94	142
Standard Embedding	111	71	95	142

Estimation of the Bit Hardness of LWE Against Quantum Adversaries. Recently the first proposals on how to estimate the hardness of LWE against quantum adversaries were published [5, 10, 17]. However, those do not follow the same methodology as the LWE-Estimator we used to estimate the classical hardness of LWE. Although the LWE-Estimator originally considers only classical hardness estimations, the first changes towards quantum estimations were made. In the latest version (from commit-id b929691 on) the estimates for a quantumly enhanced sieving algorithm [38] are included as a subroutine to solve the lattice reduction algorithm BKZ2.0 [23].

As mentioned in [5], a recently published quantum algorithm [47] can be applied to enumeration methods, possibly giving a quadratic speed-up on the run time of enumeration. We apply the quadratic speed-up to the currently fastest enumeration estimations by Micciancio and Walter [46] and add the resulting estimations as a subroutine to be used in BKZ2.0. We summarize the estimations using quantum sieving and quantum enumeration in Table 2.

6 Software Implementation

To evaluate the performance of our proposed parameter sets we present a software implementation targeting the Intel Haswell microarchitecture. The starting point for our implementation is the software presented by Dagdelen et al. [61],

which we obtained from the authors. Our software offers the same level of protection against timing attacks as the software presented in [61]. The software makes use of the fast AVX2 instructions on vectors of 4 double-precision floating-point numbers. In this section we first give a description of TESLA with technical details necessary for implementations. Furthermore, we present important modifications to the software proposed in [61].

Algorithm KeyGen	Algorithm Sign
<p>INPUT: $1^\lambda; A, n, n', m, q, \sigma$ OUTPUT: $(S, E, s), T$</p> <ol style="list-style-type: none"> 1. $S \leftarrow_{\S} \mathcal{D}_\sigma^{n \times n'}$ 2. $E \leftarrow_{\S} \mathcal{D}_\sigma^{m \times n'}$ 3. if $\text{checkE}(E) = 0 \vee \text{checkS}(S) = 0$ 4. then Restart 5. $s \leftarrow_{\S} \{0, 1\}^\kappa$ 6. $T \leftarrow AS + E \pmod{q}$ 7. $\text{sk} \leftarrow (S, E, s), \text{pk} \leftarrow T$ 8. return (sk, pk) <p style="border-top: 1px solid black; margin-top: 10px;">Algorithm Verify</p> <p>INPUT: μ, q, z, c, A, T OUTPUT: $\{0, 1\}$</p> <ol style="list-style-type: none"> 1. $c \leftarrow F(c)$ 2. $w' \leftarrow Az - Tc \pmod{q}$ 3. $c' \leftarrow H([w'], \mu)$ 4. if $c' = c \wedge \ z\ _\infty \leq B - U$ 5. then return 1 6. return 0 	<p>INPUT: μ, q, A, S, E, s OUTPUT: (z, c)</p> <ol style="list-style-type: none"> 1. $j \leftarrow 0$ 2. $k \leftarrow PRF_1(s, \mu)$ 3. $y \leftarrow PRF_2(k, j)$ 4. $v \leftarrow Ay \pmod{q}$ 5. $c \leftarrow H([v], \mu)$ 6. $c \leftarrow F(c)$ 7. $z \leftarrow y + Sc$ 8. $w \leftarrow v - Ec \pmod{q}$ 9. if $\ [w]_L \ _\infty > 2^{d-1} - L_E$ $\vee \ w\ _\infty > \lfloor q/2 \rfloor - L_E \vee \ z\ _\infty > B - U$ 10. then $j \leftarrow j + 1$ and go to Step 1 11. return (z, c)

Fig. 1. Specification of the signature scheme TESLA = (KeyGen, Sign, Verify); for details of the functions checkE and checkS see the explanation of the public parameters and definition of functions.

Public Parameters and Definition of Functions. TESLA is parameterized by the integers n, n', m, α, κ , and the security parameter λ with $m > n > \kappa \geq \lambda$; by the matrix $A \leftarrow_{\S} \mathbb{Z}_q^{m \times n}$; by the hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$, by the encoding function $F : \{0, 1\}^\kappa \rightarrow \mathbb{H}$ (see [31] for more information), by the pseudo-random function $PRF_1 : \{0, 1\}^\kappa \times \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$, and the pseudo-random generator $PRF_2 : \{0, 1\}^\kappa \times \mathbb{Z} \rightarrow [-B, B]^n$. The remaining values, *i.e.*, σ, h, d, B, U, q, L , and L_S , are derived as shown in Table 1 and described in Sec. 5.1.

Moreover, we define the functions checkE , introduced in [61, Section 3.2], as follows: for a matrix E , define E_i to be the i -th row of E . The function $\max_k(\cdot)$ returns the k -th largest entry of a vector. The matrix E is rejected if for any

row of \mathbf{E} it holds that $\sum_{k=1}^h \max_k(\mathbf{E}_i)$ is greater than some bound L . We apply a similar check checkS to \mathbf{S} : The matrix \mathbf{S} is rejected if for any row of \mathbf{S} it holds that $\sum_{k=1}^h \max_k(\mathbf{S}_i)$ is greater than some bound L_S .

Remark 1 (Deterministic signature). Note, that signing is deterministic for each message μ since the randomness is determined by the vector \mathbf{y} which is deterministically computed by the secret key and the message to-be-signed. In the original scheme by Bai and Galbraith [8] the vector \mathbf{y} was sampled uniformly random in $[-B, B]^n$. As long as we assume that PRF_1 and PRF_2 are pseudo-random functions, the prf-advantage of $PRF_2 \circ PRF_1$ is negligible in the security parameter. Hence, the reduction given in Theorem 2 is tight with our choice of parameters. The idea to use a pseudo-random function to generate signatures deterministically was deployed several times before [11, 13, 34, 48, 59].

Parallel Matrix-Vector Multiplication. The most costly operation during signing is the computation of $\mathbf{A}\mathbf{y}$, which requires nm multiplications in \mathbb{Z}_q , most of those followed by accumulation. Intel Haswell processors can perform two multiply-accumulate instructions on 256-bit vectors of double-precision floating point numbers every cycle. As we represent elements of \mathbb{Z}_q as double-precision floating point numbers, one obtains a lower bound of $nm/8$ cycles per matrix multiplication. This simple lower bound would translate to 254058 cycles for TESLA-0; however, this is ignoring the fact that each partial product has up to 53 bits (the size of the double-precision mantissa) and requires reduction before accumulation. Reductions are done through a multiplication by an approximate inverse of q , rounding, multiplication by q and subtraction.

Another cause for not even getting close to the lower bound derived from multiply-accumulate throughput is that each coefficient from \mathbf{A} needs to be loaded from (at best) L2 cache, because the whole matrix \mathbf{A} does not fit into the 32KB of level-1 cache. This was already pointed out as an additional bottleneck by Dagdelen et al. [61]. However, on average, signing computes multiple of those matrix-vector multiplications, all with the constant matrix \mathbf{A} . Our software always samples $k = 3$ vectors \mathbf{y} , then performs 3 matrix-vector multiplications and then proceeds to investigate whether one of the 3 results is usable for the signature. The disadvantage of this technique is that most of the time we compute some matrix-vector products that are not used in the end. The advantage is that matrix coefficients, once loaded from L2 cache are used k times instead of just once. The optimal value of k depends on the parameter set. We wrote scripts that generate an optimized assembly routine for different parameters and different values of k . With those scripts we generated and benchmarked code for many different combinations that all offered the targeted security and then picked the fastest one. Computing 3 matrix-vector multiplications of the form $\mathbf{A}\mathbf{y}$ with this parallel approach takes 11875355 cycles an average for TESLA-0.

Signature verification cannot use the k -times-parallel matrix-vector multiplication for the computation of $\mathbf{A}\mathbf{z}$. For this task we use an approach similar to [61].

7 Results and Comparison

Table 3 gives benchmarking results for TESLA-0 and TESLA-1, and compares those benchmarks to state-of-the-art results from the literature. Due to the large values q and B of the parameter set TESLA-2, certain elements do not fit into the 53-bit mantissa of a double-precision floating point variable. Hence, we do not compare the performance of TESLA-2 in Table 3.

We obtain our benchmarks on an Intel Core-i7 4770K (Haswell) processor while disabling Turbo Boost and hyperthreading. Benchmarks of TESLA for signing are averaged over 100,000 signatures; benchmarks of TESLA for verification are the median of 100 verifications. The reason for not reporting the median for TESLA signing performance is that because of the rejection sampling, it would be overly optimistic. For all software results we report the sizes of keys and signatures actually produced by the software, not the theoretically smallest possible sizes with full compression.⁸

As can be seen in Table 3, TESLA is several magnitudes faster and sizes are smaller than the only other lattice-based signature scheme that is also proven tightly secure in the quantum random oracle model for the same (classical) security of 96 bits. However, the signature generation and verification algorithms of TESLA-0 are much slower than the implementation of [61] for the same level of security. This is due to the large difference of the parameters chosen, *e.g.*, the matrix dimension m in TESLA-0 is 3156, while $m = 840$ in the parameter set DEG⁺ proposed by Dagdelen *et al.* [61]. Note that the parameter set TESLA-0 is chosen according to our security reduction, while the set DEG⁺ is not chosen according to the (non-tight) security reduction given in [8].

In the (as of yet quite small) realm of signatures that offer 128 bits of post-quantum security, TESLA-2 offers an alternative to SPHINCS. Public and secret keys of TESLA-2 are much larger than SPHINCS keys, but signatures are several magnitudes smaller. The post-quantum multivariate-based signature scheme Rainbow5640 [22, 25] performs best among all listed schemes but unfortunately, comes with no security reduction to its underlying problem.

8 The Gaussian Heuristic

Informally speaking, the *Gaussian heuristic* is a formula asserting that the number of points the intersection $L \cap S$ of a lattice $L \subset \mathbb{R}^m$ and a measurable set $S \subset \mathbb{R}^m$ is approximated by the ratio $\text{vol}(S)/\det(L)$. Use of the Gaussian heuristic has appeared in prior works relating to TESLA [2, 4, 8]. It seems, however, that a formal statement and proof of the Gaussian heuristic as used in these works has not appeared in any literature. In this section we address this oversight by presenting a formal statement and proof of the Gaussian heuristic for such applications. We emphasize that the present manuscript avoids all use of

⁸ We make an exception for BLISS. The authors of the software obviously did not spend any effort on reducing the size of signatures and keys; we report sizes with “trivial” compression through choosing native data types of appropriate sizes.

the Gaussian heuristic, so the results in this section are of independent interest to the rest of this document.

8.1 History

The study of counting lattice points within a bounded region was studied by Gauss in the two-dimensional case, and has been explored extensively in very generalized cases in literature on geometry [24,30,39]. The expectation intuitively follows from Minkowski’s theorem, but can deviate arbitrarily for many types of lattices. Due to this fact, the results on general lattices in the literature assert error bounds that are not useful in a cryptographic setting. (For example, the deviation from $\text{vol}(S)/\det(L)$ may be as large as $O(r^{n-1})$ for a set S containing a sphere of radius r).

A more useful bound is presented in [44,45], in terms of the covering radius of a lattice.⁹ However, this bound is not defined for sets S whose volume $\text{vol}(S)$ is smaller than the determinant $\det(L)$ of the lattice L . Unfortunately, prior works related to TESLA seek to apply the Gaussian heuristic to precisely this setting.

Previous applications of the Gaussian heuristic for cryptographic purposes have largely focused on so-called “ q -ary lattices”, which are lattices L with the property that $q\mathbb{Z}^m \subseteq L \subseteq \mathbb{Z}^m$. In these applications, the set S is a finite subset of \mathbb{Z}_q^m and the lattices L have the form $\Lambda_q(\mathbf{A})$ or $\Lambda_q^\perp(\mathbf{A})$ given by

$$\Lambda_q(\mathbf{A}) \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{Z}^m : \exists \mathbf{x} \in \mathbb{Z}^n \text{ with } \mathbf{y} = \mathbf{A}\mathbf{x}\} \quad (6)$$

$$\Lambda_q^\perp(\mathbf{A}) \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{Z}^m : \mathbf{A}^\top \mathbf{y} = \mathbf{0} \pmod{q}\} \quad (7)$$

for arbitrary $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ with $m \geq n$. We refer to $\Lambda_q^\perp(\mathbf{A})$ as the *primal* lattice and $\Lambda_q(\mathbf{A})$ as the *dual* lattice. The following discussion immediately generalizes to *shifted* primal lattices of the form

$$\Lambda_{\mathbf{u},q}^\perp(\mathbf{A}) \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{Z}^m : \mathbf{A}^\top \mathbf{y} = \mathbf{u} \pmod{q}\} \quad (8)$$

for arbitrary $\mathbf{u} \in \mathbb{Z}_q^n$.

In these applications we have that the volume $\text{vol}(S)$ is simply the number $\#S$ of elements in S . Moreover, if \mathbf{A} has full rank ($\text{rank}(\mathbf{A}) = n$) then the determinants of the lattices $\Lambda_q(\mathbf{A})$, $\Lambda_{\mathbf{u},q}^\perp(\mathbf{A})$ are given by

$$\det(\Lambda_q(\mathbf{A})) = \frac{1}{q^{m-n}} \quad \det(\Lambda_{\mathbf{u},q}^\perp(\mathbf{A})) = \frac{1}{q^n} \quad (9)$$

The Gaussian heuristic is used in these applications to bound the expectation over a uniformly random \mathbf{A} of the size of the intersections $\Lambda_q(\mathbf{A}) \cap S$ and $\Lambda_{\mathbf{u},q}^\perp(\mathbf{A}) \cap S$. Since a random matrix \mathbf{A} has full rank except with only negligibly small

⁹ Some recent work related to Minkowski’s theorem and estimating the covering radius can be found in [55].

probability, the Gaussian heuristic suggests that

$$\mathbb{E}_A [\# (A_q(\mathbf{A}) \cap S)] \approx \frac{\#S}{q^{m-n}} \quad (10)$$

$$\mathbb{E}_A [\# (A_{u,q}^\perp(\mathbf{A}) \cap S)] \approx \frac{\#S}{q^n} \quad (11)$$

We provide rigorous statements and proofs of these claims below.

8.2 Proofs of the Gaussian Heuristic for q -ary Lattices

Because S is a subset of \mathbb{Z}_q^m , the size of the intersection $A_q(\mathbf{A}) \cap S$ is really a question of the intersection of the image of \mathbf{A} with S over the finite field \mathbb{Z}_q . Similarly, the intersection $A_{u,q}^\perp(\mathbf{A}) \cap S$ is merely the intersection with S of the preimage of \mathbf{u} under \mathbf{A} over the finite field \mathbb{Z}_q .

With this observation in mind, for any vector $\mathbf{u} \in \mathbb{Z}_q^n$ and any matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ define the following subsets of \mathbb{Z}_q^m :

$$\text{Im}(\mathbf{A}) \stackrel{\text{def}}{=} \{\mathbf{A}\mathbf{x} : \mathbf{x} \in \mathbb{Z}_q^n, \mathbf{x} \neq 0\}, \quad (12)$$

$$\text{PreIm}_u(\mathbf{A}) \stackrel{\text{def}}{=} \{\mathbf{y} \in \mathbb{Z}_q^m : \mathbf{A}^\top \mathbf{y} = \mathbf{u} \pmod{q}, \mathbf{y} \neq 0\}. \quad (13)$$

Observe that

$$A_q(\mathbf{A}) \cap S = \text{Im}(\mathbf{A}) \cap S, \quad (14)$$

$$A_{u,q}^\perp(\mathbf{A}) \cap S = \text{PreIm}_u(\mathbf{A}) \cap S. \quad (15)$$

We prove the following.

Proposition 3 (Gaussian heuristic, “dual” q -ary lattices). *Let $S \subseteq \mathbb{Z}_q^m$ be any set. It holds that*

$$\mathbb{E}_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} [\# (\text{Im}(\mathbf{A}) \cap S)] \leq \frac{\#S}{q^{m-n}}. \quad (16)$$

Proof. By definition, the expectation in (16) is given by

$$\frac{1}{\#\mathbb{Z}_q^{m \times n}} \sum_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} \sum_{\mathbf{v} \in \mathbb{Z}_q^m} \text{bool}[\mathbf{v} \in \text{Im}(\mathbf{A})] \text{bool}[\mathbf{v} \in S] \quad (17)$$

$$= \sum_{\mathbf{v} \in \mathbb{Z}_q^m} \text{bool}[\mathbf{v} \in S] \Pr_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} [\mathbf{v} \in \text{Im}(\mathbf{A})]. \quad (18)$$

For each \mathbf{v} let us bound the probability

$$\Pr_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} [\mathbf{v} \in \text{Im}(\mathbf{A})]. \quad (19)$$

By definition, the probability (19) is given by

$$\frac{1}{\#\mathbb{Z}_q^{m \times n}} \# \{ \mathbf{A} \in \mathbb{Z}_q^{m \times n} : \exists \mathbf{x} \in \mathbb{Z}_q^n \text{ with } \mathbf{A}\mathbf{x} = \mathbf{v}, \mathbf{x} \neq \mathbf{0} \} \quad (20)$$

$$\leq \frac{1}{\#\mathbb{Z}_q^{m \times n}} \sum_{\substack{\mathbf{x} \in \mathbb{Z}_q^n \\ \mathbf{x} \neq \mathbf{0}}} \# \{ \mathbf{A} : \mathbf{A}\mathbf{x} = \mathbf{v} \}. \quad (21)$$

Observe that

$$\# \{ \mathbf{A} : \mathbf{A}\mathbf{x} = \mathbf{v} \} = q^{mn-m} \quad (22)$$

when $\mathbf{x} \neq \mathbf{0}$ and so the probability (19) is at most

$$\frac{1}{q^{mn}} q^n \cdot q^{mn-m} = \frac{1}{q^{m-n}}. \quad (23)$$

Substituting this bound into (18) we find that the expectation in (16) is at most

$$\sum_{\mathbf{v} \in \mathbb{Z}_q^m} \text{bool}[\mathbf{v} \in S] \frac{1}{q^{m-n}} = \frac{\#S}{q^{m-n}} \quad (24)$$

as desired.

Proposition 4 (Gaussian heuristic, “primal” q -ary lattices). *Let $S \subseteq \mathbb{Z}_q^m$ be any set and let $\mathbf{u} \in \mathbb{Z}_q^n$ be any vector. It holds that*

$$\mathbb{E}_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} [\# (\text{PreIm}_{\mathbf{u}}(\mathbf{A}) \cap S)] = \frac{\#S}{q^n}. \quad (25)$$

Proof. By definition, the expectation in (25) is given by

$$\frac{1}{\#\mathbb{Z}_q^{m \times n}} \sum_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \text{bool}[\mathbf{y} \in \text{PreIm}_{\mathbf{u}}(\mathbf{A})] \text{bool}[\mathbf{y} \in S] \quad (26)$$

$$= \sum_{\mathbf{y} \in \mathbb{Z}_q^m} \text{bool}[\mathbf{y} \in S] \Pr_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} [\mathbf{y} \in \text{PreIm}_{\mathbf{u}}(\mathbf{A})] \quad (27)$$

For each $\mathbf{y} \neq \mathbf{0}$ let us compute probability

$$\Pr_{\mathbf{A} \in \mathbb{Z}_q^{m \times n}} [\mathbf{y} \in \text{PreIm}_{\mathbf{u}}(\mathbf{A})] = \frac{1}{\#\mathbb{Z}_q^{m \times n}} \# \{ \mathbf{A} \in \mathbb{Z}_q^{m \times n} : \mathbf{A}^T \mathbf{y} = \mathbf{u} \pmod{q} \}. \quad (28)$$

Observe that

$$\# \{ \mathbf{A} : \mathbf{A}^T \mathbf{y} = \mathbf{u} \pmod{q} \} = q^{mn-n} \quad (29)$$

since $\mathbf{y} \neq \mathbf{0}$ and so the probability (28) equals

$$\frac{1}{q^{mn}} q^{mn-n} = \frac{1}{q^n}. \quad (30)$$

Substituting this bound into (27) we find that the expectation in (25) equals

$$\sum_{\mathbf{y} \in \mathbb{Z}_q^m} \text{bool}[\mathbf{y} \in S] \frac{1}{q^n} = \frac{\#S}{q^n} \quad (31)$$

as desired.

Acknowledgments

We are especially grateful to Peter Schwabe for contributions to our software implementation and to the presentation of the paper. We thank Chris Peikert for pointing out a flaw in previous security reductions for TESLA in the random oracle model. We thank Steven Galbraith and anonymous reviewers for valuable feedback on an earlier version of this manuscript.

This work has been supported by the German Research Foundation (DFG) as part of project P1 within the CRC 1119 CROSSING, by TÜBITAK under 2214-A Doctoral Research Program Grant and 2211-C PhD Scholarship, by Ege University under project 2014-FEN-065, and by CryptoWorks21.

References

1. Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-Secure Signatures from Lossy Identification Schemes. In *EUROCRYPT 2012*, LNCS. Springer, 2012.
2. Sedat Akleyek, Nina Bindel, Johannes A. Buchmann, Juliane Krämer, and Giorgia Azzurra Marson. An Efficient Lattice-Based Signature Scheme with Provably Secure Instantiation. In *AFRICACRYPT 2016*, volume 9646 of *LNCS*. Springer, 2016.
3. Martin R. Albrecht, Rachel Player, and Sam Scott. On the Concrete Hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9, 2015.
4. Erdem Alkim, Nina Bindel, Johannes Buchmann, Özgür Dagdelen, and Peter Schwabe. TESLA: Tightly-Secure Efficient Signatures from Standard Lattices. Cryptology ePrint Archive, Report 2015/755, version 20161117:055833, 2015.
5. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-Quantum Key Exchange – a New Hope. In *25th USENIX Security Symposium*. USENIX Association, 2016.
6. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, 2009.
7. László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
8. Shi Bai and Steven D. Galbraith. An Improved Compression Technique for Signatures Based on Learning with Errors. In *Topics in Cryptology – CT-RSA 2014*, volume 8366 of *LNCS*. Springer, 2014.
9. Rachid El Bansarkhani and Johannes Buchmann. Improvement and Efficient Implementation of a Lattice-Based Signature Scheme. In *Selected Areas in Cryptography*, volume 8282 of *LNCS*. Springer, 2013.
10. Paulo S. L. M. Barreto, Patrick Longa, Michael Naehrig, Jefferson E. Ricardini, and Gustavo Zanon. Sharper Ring-LWE Signatures. Cryptology ePrint Archive, Report 2016/1026, 2016.
11. George Barwood. Digital Signatures Using Elliptic Curves. message 32f519ad.19609226@news.dial.pipex.com posted to sci.crypt, 1997. <http://groups.google.com/group/sci.crypt/msg/b28aba37180dd6c6>.

12. Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and Weaknesses of Quantum Computing. *SIAM J. Comput.*, 26(5), 1997.
13. Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-Speed High-Security Signatures. In *Cryptographic Hardware and Embedded Systems – CHES 2011*, volume 6917 of *LNCS*. Springer, 2011.
14. Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Peter Schwabe, and Zooko Wilcox-O’Hearn. SPHINCS: Practical Stateless Hash-Based Signatures. In *EUROCRYPT 2015*, volume 9056 of *LNCS*. Springer, 2015.
15. Daniel J. Bernstein and Tanja Lange. eBACS: ECRYPT Benchmarking of Cryptographic Systems. <http://bench.cr.yp.to> (accessed 2015-05-19).
16. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random Oracles in a Quantum World. In *ASIACRYPT 2011*, volume 7073 of *LNCS*. Springer, 2011.
17. Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the Ring! Practical, Quantum-Secure Key Exchange from LWE. In *CCS 2016*. ACM, 2016.
18. Xavier Boyen and Qinyi Li. Towards Tightly Secure Lattice Short Signature and Id-Based Encryption. In *ASIACRYPT 2016*, volume 10032 of *LNCS*. Springer, 2016.
19. Michel Boyer, Gilles Brassard, Peter Hoyer, and Alain Tapp. Tight bounds on quantum searching, 1996.
20. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing (STOC 2013)*, pages 575–584. ACM, 2013.
21. Sanjit Chatterjee, Neal Koblitz, Alfred Menezes, and Palash Sarkar. Another Look at Tightness II: Practical Issues in Cryptography. Cryptology ePrint Archive, Report 2016/360, 2016.
22. Anna Inn-Tung Chen, Ming-Shing Chen, Tien-Ren Chen, Chen-Mou Cheng, Jintai Ding, Eric Li-Hsiang Kuo, Frost Yu-Shuang Lee, and Bo-Yin Yang. SSE Implementation of Multivariate PKCs on Modern x86 CPUs. In *CHES 2009*, volume 5747 of *LNCS*. Springer, 2009.
23. Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better Lattice Security Estimates. In *ASIACRYPT 2011*, volume 7073 of *LNCS*. Springer, 2011.
24. Pete L. Clark. Geometry of numbers with applications to number theory, 2012. <http://math.uga.edu/~pete/geometryofnumbers.pdf>.
25. Jintai Ding and Dieter Schmidt. Rainbow, a New Multivariable Polynomial Signature Scheme. In *Applied Cryptography and Network Security*, volume 3531 of *LNCS*. Springer, 2005.
26. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice Signatures and Bimodal Gaussians. In *CRYPTO 2013*, volume 8042 of *LNCS*. Springer, 2013.
27. Léo Ducas. Accelerating Bliss: the Geometry of Ternary Polynomials. Cryptology ePrint Archive, Report 2014/874, 2014.
28. Edward Eaton and Fang Song. Making Existential-Unforgeable Signatures Strongly Unforgeable in the Quantum Random-Oracle Model. In *10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015*, 2015.

29. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for Hard Lattices and New Cryptographic Constructions. In *STOC 2008*. ACM, 2008.
30. P.M. Gruber and J.M. Wills. *Handbook of Convex Geometry*. Number v. 2 in Convex Geometry. North-Holland, 1993.
31. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical Lattice-Based Cryptography: A Signature Scheme for Embedded Systems. In *CHES 2012*, volume 7428 of *LNCS*. Springer, 2012.
32. Tim Güneysu, Tobias Oder, Thomas Pöppelmann, and Peter Schwabe. Software Speed Records for Lattice-Based Signatures. In *Post-Quantum Cryptography*, volume 7932 of *LNCS*. Springer, 2013.
33. Ravi Kannan. Minkowski’s convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, August 1987.
34. Jonathan Katz and Nan Wang. Efficiency Improvements for Signature Schemes with Tight Security Reductions. In *CCS 2003*. ACM, 2003.
35. Phillip Kaye, Raymond Laflamme, and Michele Mosca. *An Introduction to Quantum Computing*. Oxford University Press, Inc., New York, NY, USA, 2007.
36. Neal Koblitz and Alfred Menezes. Another Look at “Provable Security”. II. In *INDOCRYPT 2006*, volume 4329 of *LNCS*. Springer, 2006.
37. Neal Koblitz and Alfred Menezes. The Random Oracle Model: a Twenty-Year Retrospective. *Designs, Codes and Cryptography*, 77(2), 2015.
38. Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Finding Shortest Lattice Vectors Faster Using Quantum Search. *Designs, Codes and Cryptography*, 2015.
39. Mordechay B. Levin. On the lower bound in the lattice point remainder problem for a parallelepiped, 2013.
40. Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, 2011.
41. Mingjie Liu and PhongQ. Nguyen. Solving BDD by enumeration: An update. In *Topics in Cryptology – CT-RSA 2013*, volume 7779 of *LNCS*, pages 293–309. Springer, 2013.
42. Vadim Lyubashevsky. Lattice Signatures without Trapdoors. In *EUROCRYPT 2012*, volume 7237 of *LNCS*. Springer, 2012.
43. Daniele Micciancio. Improving lattice based cryptosystems using the Hermite normal form. In *Cryptography and Lattices*, volume 2146 of *LNCS*, pages 126–145. Springer, 2001.
44. Daniele Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai’s connection factor. *SIAM Journal on Computing*, 34(1):118–169, 2004. Preliminary version in STOC 2002.
45. Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, March 2002.
46. Daniele Micciancio and Michael Walter. Fast Lattice Point Enumeration with Minimal Overhead. In *SODA 2015*. SIAM, 2015.
47. Ashley Montanaro. Quantum Walk Speedup of Backtracking Algorithms. arXiv preprint arXiv:1509.02374, 2016.
48. David M’Raïhi, David Naccache, David Pointcheval, and Serge Vaudenay. Computational Alternatives to Random Number Generators. In *Selected Areas in Cryptography*, volume 1556 of *LNCS*. Springer, 1998.
49. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, New York, 2000.

50. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Proceedings of the forty-first annual ACM symposium on Theory of computing (STOC 2009)*, pages 333–342. ACM, 2009.
51. Chris Peikert. A Decade of Lattice Cryptography. Cryptology ePrint Archive, Report 2015/939, 2015.
52. Chris Peikert and Brent Waters. Lossy Trapdoor Functions and Their Applications. In *STOC 2008*. ACM, 2008.
53. David Pointcheval and Jacques Stern. Security Proofs for Signature Schemes. In *EUROCRYPT 1996*, volume 1070 of *LNCS*. Springer, 1996.
54. Oded Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *STOC 2005*. ACM, 2005.
55. Oded Regev and Noah Stephens-Davidowitz. A reverse minkowski theorem, 2016.
56. Markus Schmidt. Estimation of the Hardness of the Learning with Errors Problem with a Restricted Number of Samples. GitHub at https://bitbucket.org/Ma_Schmidt/lwe-estimator, 2017.
57. Markus Schmidt and Nina Bindel. Estimation of the Hardness of the Learning with Errors Problem with a Restricted Number of Samples. Cryptology ePrint Archive, Report 2017/140, 2017.
58. Dominique Unruh. Quantum Position Verification in the Random Oracle Model. In *CRYPTO 2014*, volume 8617 of *LCNS*. Springer, 2014.
59. John Wigley. Removing need for RNG in Signatures. message 5gov5dpad@wapping.ecs.soton.ac.uk posted to sci.crypt, 1997. <http://groups.google.com/group/sci.crypt/msg/a6da45bcc8939a89>.
60. Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer, 2012.
61. Özgür Dagdelen, Rachid El Bansarkhani, Florian Göpfert, Tim Güneysu, Tobias Oder, Thomas Pöppelmann, Ana Helena Sánchez, and Peter Schwabe. High-Speed Signatures from Standard Lattices. In *LATINCRYPT 2014*, volume 8895 of *LNCS*. Springer, 2015.

Table 3. Overview of state-of-the-art post-quantum signature schemes; signature sizes are given in byte [B], key sizes are given in kilo byte [KB]; the column “ROM?, tight?” states whether the scheme has a security reduction in the random oracle model and whether this reduction is tight; “QROM?, tight?” states the same for the quantum random oracle model; “Security (PreQ)” lists the claimed pre-quantum security level; “Security (PostQ)” lists the claimed post-quantum security level, if available

Scheme/ Software	Comp. Assum.	ROM? Tight?	QROM? Tight?	Security (PreQ) (PostQ)	CPU	Key Size [KB]	Sig. Size [B]	Cycle counts
Selected signature schemes over standard lattices								
GPV [9, 29]	SIS	yes yes	yes yes	96 59	AMD Opteron 8356 (Barcelona)	vk: 27,840 sk: 12,064	30,105	sign: 312,800,000 verify: 50,600,000
DEG ⁺ [8, 61]	SIS, LWE	yes no	–	96 ?	Intel Core i7-4770K (Haswell)	vk: 1,581 sk: 891	1,495	sign: 1,203,924 verify: 335,072
TESLA-0 (this paper)	LWE	yes yes	yes yes	96 ?	Intel Core-i7-4770K (Haswell)	vk: 4,808 sk: 2,895	1,964	sign: 27,243,747 verify: 5,374,884
TESLA-1 (this paper)	LWE	yes yes	yes yes	128 ?	Intel Core-i7-4770K (Haswell)	vk: 11,653 sk: 6,769	2,444	sign: 143,402,231 verify: 19,284,672
Selected signatures schemes over ideal lattices								
GPV-poly [9, 29]	R-SIS	yes yes	yes yes	96 59	AMD Opteron 8356 (Barcelona)	vk: 55 sk: 26	32,972	sign: 80,500,000 verify: 11,500,000
GLP [31, 32, 61] ^a	DCK	yes no	–	75–80 ?	Intel Core i5-3210M (Ivy Bridge)	vk: 1.5 sk: 0.25	1,186	sign: 452,223 verify: 34,004
BLISS-BI [26, 27] ^b	R-SIS, NTRU	yes no	–	128 ?	“Intel Core @ 3.4GHz”	vk: 7 sk: 2	1,559	sign: ≈358,400 verify: 102,000
Selected other post-quantum signature schemes								
SPHINCS-256 [14]	Hash collisions 2nd preimage	yes no ^c	yes no ^c	>128 128	Intel Xeon E3-1275 (Haswell)	vk: 1 sk: 1	41,000	sign: 51,636,372 verify: 1,451,004
Rainbow5640 [22, 25] ^d	MQ, EIP ^e	–	–	80 ?	Intel Xeon E3-1275 (Haswell)	vk: 43 sk: 84	37	sign: 42,700 verify: 36,072

^a In the benchmarks we include the improvements by Dagdelen et al. presented in [61].

^b We report sizes of keys and signatures with “trivial” compression as explained in the text.

^c The security of SPHINCS is reduced tightly from the hardness of finding hash collisions and non-tightly from the hardness of finding 2nd preimages in the standard model. Hence the reduction also holds in the ROM and QROM.

^d Benchmark on Haswell CPU from [15].

^e The security of Rainbow5640 is based on the Multivariate Quadratic polynomial (MQ) and the Extended Isomorphism of Polynomials (EIP) problem, but no security reduction has been given yet.

A Overview of Security Proof

As in Section 4, let \mathcal{F} be a forging algorithm that, on input a TESLA public key (\mathbf{A}, \mathbf{T}) , makes no more than q_h quantum queries to a hash oracle $H(\cdot)$ and no more than q_s classical queries to a TESLA sign oracle for public key (\mathbf{A}, \mathbf{T}) . Let $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$ denote the probability that \mathcal{F} produces a valid TESLA forgery. We build a LWE-solver \mathcal{S} whose run time is close to that of \mathcal{F} and who solves LWE with success bias close to $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$. It then follows from the presumed hardness of LWE that $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$ must be small.

The LWE-solver \mathcal{S} is described in Algorithm 4. Classical queries made by \mathcal{F} to the sign oracle are simulated by \mathcal{S} as specified in Simulated sign (Algorithm 5). Quantum queries made by \mathcal{F} to the hash oracle are simulated by \mathcal{S} according to the construction of Zhandry based on $2q_h$ -wise independent functions [60]. (Alternately, if performance of the simulator is a concern then one could instead use a quantum-resistant pseudorandom function.)

Algorithm 4 LWE-solver \mathcal{S} using a forger \mathcal{F} .

Input: A LWE instance (\mathbf{A}, \mathbf{T}) .

Output: “Yes” or “no”.

- 1: Invoke the forger \mathcal{F} with public key (\mathbf{A}, \mathbf{T}) .
Whenever \mathcal{F} makes a hash or sign query, simulate that query as follows:
 - Classical sign queries.** Execute Simulated sign (Algorithm 5).
 - Quantum hash queries.** Apply a quantum circuit that implements a random but fixed $2q_h$ -wise independent function, except on inputs that have been re-programmed by Simulated sign.
 - 2: Eventually, \mathcal{F} produces a purported forgery.
If that forgery is legit then output “yes”, otherwise output “no”.
-

Algorithm 5 Simulated sign

Input: Message μ , public key (\mathbf{A}, \mathbf{T}) .

Output: Signature (\mathbf{z}, \mathbf{c}) .

- 1: Choose \mathbf{z} uniformly at random among $(B - U)$ -short vectors from \mathbb{Z}_q^n .
 - 2: Choose $\mathbf{c} \in \mathbb{H}$ uniformly at random.
 - 3: If $\mathbf{Az} - \mathbf{Tc}$ is not well-rounded then retry at step 1.
 - 4: Re-program the hash oracle $H(\cdot)$ so that $H([\mathbf{Az} - \mathbf{Tc}], \mu) = \mathbf{c}$.
 - 5: Return (\mathbf{z}, \mathbf{c}) .
-

That \mathcal{S} solves LWE with success bias close to $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$ is a consequence of the following facts, which are proven in subsequent sections:

Appendix B: For yes-instances of LWE, the probability with which S outputs “yes” is close to $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$.

Appendix C: For no-instances of LWE, \mathcal{F} successfully forges (and hence S outputs “yes”) with only negligible probability.

A.1 Notation and Sizes for Various Sets of Vectors

Discussion in appendices B, C, D refers to the following sets of integer vectors:

\mathbb{Y} : The set of vectors $\mathbf{y} \in \mathbb{Z}_q^n$ such that \mathbf{y} is B -short.

$\Delta\mathbb{Y}$: $\{\mathbf{y} - \mathbf{y}' : \mathbf{y}, \mathbf{y}' \in \mathbb{Y}\}$.

\mathbb{S} : The set of vectors $\mathbf{z} \in \mathbb{Z}_q^n$ such that \mathbf{z} is $(B - U)$ -short.

$\Delta\mathbb{S}$: $\{\mathbf{z} - \mathbf{z}' : \mathbf{z}, \mathbf{z}' \in \mathbb{S}\}$.

\mathbb{H} : The set of vectors $\mathbf{c} \in \{-1, 0, 1\}^{n'}$ with exactly h nonzero entries.

$\Delta\mathbb{H}$: $\{\mathbf{c} - \mathbf{c}' : \mathbf{c}, \mathbf{c}' \in \mathbb{H}\}$.

\mathbb{W} : The set $\{\mathbf{w} : \mathbf{w} \in \mathbb{Z}_q^m\}$ of integer vectors obtained from the high bits of a vector in \mathbb{Z}_q^m .

$\Delta\mathbb{L}$: $\{\mathbf{x} - \mathbf{x}' : \mathbf{x}, \mathbf{x}' \in \mathbb{Z}_q^m \text{ and } [\mathbf{x}] = [\mathbf{x}']\} = [-(2^d - 1), 2^d - 1]^m$.

The sizes of some of these sets are listed below:

$$\#\mathbb{Y} = (2B - 1)^n \qquad \#\Delta\mathbb{Y} = (4B - 1)^n \qquad (32)$$

$$\#\mathbb{S} = (2(B - U) - 1)^n \qquad \#\Delta\mathbb{S} = (4(B - U) - 1)^n \qquad (33)$$

$$\#\mathbb{H} = \binom{n'}{h} 2^h \qquad (34)$$

$$\#\Delta\mathbb{L} = (2^{d+1} - 1)^m \qquad (35)$$

The size of $\Delta\mathbb{H}$ is computed as follows.

Lemma 1 (Size of $\Delta\mathbb{H}$).

$$\#\Delta\mathbb{H} = \sum_{k=0}^h \sum_{i=0}^{h-k} \binom{n'}{2i} 2^{2i} \binom{n' - 2i}{k} 2^k \qquad (36)$$

Proof. For each $k = 0, \dots, h$ let $\Delta\mathbb{H}_k \subset \Delta\mathbb{H}$ denote the set of vectors in $\Delta\mathbb{H}$ with exactly k entries in $\{-2, 2\}$ and exactly $n' - k$ entries in $\{-1, 0, 1\}$. Observe that $\#\Delta\mathbb{H} = \sum_{k=0}^h \#\Delta\mathbb{H}_k$.

Fix k and for each $i = 0, \dots, h - k$ let $\Delta\mathbb{H}_{k,i} \subset \Delta\mathbb{H}_k$ denote the set of vectors in $\Delta\mathbb{H}_k$ with exactly $2i$ entries in $\{-1, 1\}$. Observe that $\#\Delta\mathbb{H}_k = \sum_{i=0}^{h-k} \#\Delta\mathbb{H}_{k,i}$.

One can count the number of elements in each $\Delta\mathbb{H}_{k,i}$ as

$$\#\Delta\mathbb{H}_{k,i} = \binom{n'}{2i} 2^{2i} \binom{n' - 2i}{k} 2^k, \qquad (37)$$

from which the lemma follows.

B Yes-Instances of LWE

In this section we establish a lower bound on the probability

$$\Pr[\mathcal{S} \text{ output "yes"} \mid (\mathbf{A}, \mathbf{T}) \text{ yes-instance of LWE}] \quad (38)$$

in terms of $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$. This is accomplished by proving that the simulated oracles are indistinguishable from the real oracles.

To this end we consider an arbitrary distinguisher \mathcal{D} who, like the forger \mathcal{F} , makes at most q_h quantum queries to the hash oracle and at most q_s classical queries to the sign oracle. Unlike \mathcal{F} , however, \mathcal{D} 's goal is merely to distinguish the real oracles from the simulated oracles.

B.1 Adaptively Chosen Queries

An arbitrary distinguisher could adaptively and arbitrarily interleave its hash and sign queries. To facilitate our analysis we wish to model the distinguisher in such a way that sign queries occur at fixed, predictable points throughout the protocol. This goal is accomplished by a continuous accounting method for hash queries that we describe presently.

Standard formalism specifies that a quantum oracle for $H : X \rightarrow Y$ be implemented by a unitary channel $|x\rangle|y\rangle \mapsto |x\rangle|y + H(x)\rangle$. We modify this formalism so that the unitary channel for $H(\cdot)$ is a *controlled* unitary channel. Specifically, the channel acts on an additional qubit, the state of which dictates whether the unitary channel is applied:

$$|\text{off}\rangle|x\rangle|y\rangle \mapsto |\text{off}\rangle|x\rangle|y\rangle \quad (39)$$

$$|\text{on}\rangle|x\rangle|y\rangle \mapsto |\text{on}\rangle|x\rangle|y + H(x)\rangle \quad (40)$$

Consider a new type of distinguisher with the following properties:

1. The distinguisher makes $q_h q_s$ hash queries instead of just q_h hash queries.
2. Exactly one sign query occurs after every q_h hash queries.
3. For each choice of hash oracle $H(\cdot)$, the distinguisher's total "query magnitude" (in the BBBV sense - see Section B.6) on query states with the control qubit set to $|\text{on}\rangle$ over all $q_h q_s$ hash queries does not exceed q_h .

A distinguisher of this form is called a *live-switch distinguisher*. For later convenience, we refer to the query magnitude on states with the control qubit set to $|\text{on}\rangle$ as the *query magnitude on the live-switch*.

Intuition: a live-switch distinguisher can make "partial" queries to the hash oracle. If its query state has only α amplitude on the live-switch then the distinguisher is "charged" for only a $|\alpha|^2$ -fraction of a query.

It is clear that any ordinary distinguisher who makes q_h hash queries and q_s sign queries, adaptively chosen and interleaved, could be simulated by a live-switch distinguisher. Thus, live-switch distinguishers are at least as powerful

as ordinary distinguishers, and possibly more powerful. We will prove indistinguishability against a live-switch distinguisher, which is more security than is strictly necessary.

The benefit of the live-switch distinguisher is that sign queries occur at fixed points throughout the protocol - one sign query after every q_h hash queries. This property allows us to partition the interaction into q_s blocks. Each block consists of q_h quantum queries to the hash oracle, followed by a single classical query to the sign oracle. We prove security of each block and then claim security of the entire interaction inductively.

B.2 The Distinguisher's State, a First Look

To begin, consider the state of \mathcal{D} 's system immediately prior to the sign oracle query in the first block. At this point in the interaction the real and simulated oracles are perfectly identical - both respond to the first q_h hash queries in accordance with some fixed choice of hash oracle $H(\cdot)$. Let ρ_H denote the state of \mathcal{D} 's system at this point in the interaction, conditioned on $H(\cdot)$. The sign oracle (both real and simulated) acts as follows on \mathcal{D} 's system:

1. Measure the message register, resulting in outcome μ .
2. Select a signature (z, c) for message μ .
3. Prepare an output register in the classical basis state $|\mu\rangle|(z, c)$.

These actions can be viewed as a quantum channel. If the sign oracle is a real sign oracle then the signature (z, c) is a function of private randomness and the hash oracle $H(\cdot)$. In this case, the channel is denoted $\Psi_{\text{real}, H}$. If the sign oracle is a simulated sign oracle then the signature (z, c) is a function only of private randomness. In this case, the channel is denoted Ψ_{sim} .

Thus, the state of \mathcal{D} 's system at the end of the first block, conditioned on the choice of $H(\cdot)$, is either $\Psi_{\text{sim}}(\rho_H)$ or $\Psi_{\text{real}, H}(\rho_H)$. We will argue that the state $\Psi_{\text{sim}}(\rho_H)$ is δ -close to a probabilistic mixture over re-programmed hash oracles $H'(\cdot)$ of states of the form $\Psi_{\text{real}, H'}(\rho_{H'})$.

This δ -closeness is preserved by the hash queries in the second block of the interaction, since both the real and simulated hash oracles remain consistent with $H'(\cdot)$ in this block. Let $\rho_{2, H'}$ denote the state of \mathcal{D} 's system immediately prior to the sign oracle query in the second block. As above, we have that $\Psi_{\text{sim}}(\rho_{2, H'})$ is δ -close to a mixture of states of the form $\Psi_{\text{real}, H''}(\rho_{2, H''})$.

Continuing inductively, we see that the state of \mathcal{D} 's system at the end of an interaction with simulated oracles is $q_s \delta$ -close to a probabilistic mixture over hash oracles of states of \mathcal{D} 's system at the end of an interaction with real oracles. Averaging over the choice of initial hash oracle $H(\cdot)$, we then see that the simulated oracles are indistinguishable from the real oracles.

We formalize these arguments in subsequent sections.

B.3 Mid-Sign

Consider the sign oracle Mid-sign of Algorithm 6. Mid-sign should be viewed

Algorithm 6 Mid-sign

Input: Message μ , public key (A, T) , secret key (S, E) .

Output: Signature (z, c) .

- 1: Choose $(y, c) \in \mathbb{Y} \times \mathbb{H}$ uniformly at random.
 - 2: $z \leftarrow y + \mathbf{S}c$.
 - 3: If $z \notin \mathbb{S}$ then retry at step 1.
 - 4: If $Ay - \mathbf{E}c$ is not well-rounded then retry at step 1.
 - 5: Re-program the hash oracle $H(\cdot)$ so that $H([Ay], \mu) = c$.
 - 6: Return (z, c) .
-

as a hybrid of Simulated sign (Algorithm 5) and the real sign oracle Sign (Algorithm 2).

In this section we prove that Mid-sign (Algorithm 6) and Simulated sign (Algorithm 5) are identical. This fact can be stated in terms of quantum channels as follows. Let Ψ_{mid} denote the channel described by Algorithm 6. The claim of this section is that $\Psi_{\text{mid}} = \Psi_{\text{sim}}$.

To this end, define the following sets for each choice of $c \in \mathbb{H}$:

$$\text{good}_{\text{sim}}(c) \stackrel{\text{def}}{=} \{z \in \mathbb{S} : Az - \mathbf{T}c \text{ is well-rounded}\} \quad (41)$$

$$\text{good}_{\text{mid}}(c) \stackrel{\text{def}}{=} \{y \in \mathbb{Y} : y + \mathbf{S}c \in \mathbb{S} \text{ and } Az - \mathbf{T}c \text{ is well-rounded}\} \quad (42)$$

We begin with a simple observation on these sets.

Lemma 2. *The mapping $f : y \mapsto y + \mathbf{S}c$ is a bijection from $\text{good}_{\text{mid}}(c)$ to $\text{good}_{\text{sim}}(c)$ with inverse $f^{-1} : z \mapsto z - \mathbf{S}c$.*

Proof. It is clear that $f^{-1}f$ is the identity function on $\text{good}_{\text{mid}}(c)$. It remains to prove the following:

1. For each $y \in \text{good}_{\text{mid}}(c)$ it holds that $f(y) \in \text{good}_{\text{sim}}(c)$.
2. For each $z \in \text{good}_{\text{sim}}(c)$ it holds that $f^{-1}(z) \in \text{good}_{\text{mid}}(c)$.

To prove item 1 we must show (i) $f(y) \in \mathbb{S}$, and (ii) $Af(y) - \mathbf{T}c$ is well-rounded. Both items are immediate from the definitions of $\text{good}_{\text{mid}}(c)$ and f .

To prove item 2 we must show (i) $f^{-1}(z) \in \mathbb{Y}$, and (ii) $Af^{-1}(z) - \mathbf{E}c$ is well-rounded. Item (i) follows from the fact that z is $(B - U)$ -short and $\mathbf{S}c$ is U -short. Item (ii) is immediate from the definitions of $\text{good}_{\text{sim}}(c)$ and f^{-1} .

We now prove this section's claim.

Proposition 5 (Equivalence of Mid-sign and Simulated sign). *The observable behaviour of Mid-sign (Algorithm 6) is statistically identical to that of Simulated sign (Algorithm 5). In terms of quantum channels, we have $\Psi_{\text{mid}} = \Psi_{\text{sim}}$.*

Proof. The observable effects of both the Simulated sign and Mid-sign algorithms can be summarized as follows. Given a message μ as input, the algorithm

selects (i) a signature (z_μ, c_μ) as output, and (ii) a vector w_μ inducing a hash input (w_μ, μ) upon which the hash oracle is re-programmed. Thus, to establish statistical equivalence between Simulated sign and Mid-sign it suffices to prove that, for each choice of message μ , the joint distribution over (z_μ, c_μ, w_μ) induced by each algorithm is identical.

Fix an arbitrary message μ and let $(Z_{\text{sim}}, C_{\text{sim}}, W_{\text{sim}})$, $(Z_{\text{mid}}, C_{\text{mid}}, W_{\text{mid}})$ denote the joint random variables representing the observable behaviour of Simulated sign and Mid-sign, respectively, on input message μ . We argue that the joint random variables $(Z_{\text{sim}}, C_{\text{sim}})$, $(Z_{\text{mid}}, C_{\text{mid}})$ are identical. The proposition will then follow from the observation that the hash input w to be re-programmed is specified in both algorithms by the same deterministic function of (z, c) . Specifically, in Simulated sign we have $w = [Az - Tc]$, whereas in Mid-sign we have $w = [Ay]$. Since $Ay - Ec$ is well-rounded, we have

$$w = [Ay] = [Ay - Ec] = [A(y + Sc) - Tc] = [Az - Tc] \quad (43)$$

as desired.

To begin, we argue that

$$\Pr [Z_{\text{mid}} = z \mid C_{\text{mid}} = c] = \Pr [Z_{\text{sim}} = z \mid C_{\text{sim}} = c] \quad (44)$$

for each choice of $c \in \mathbb{H}$. In Simulated sign, conditioned on a choice of c , the vector z is chosen uniformly among those $z \in \text{good}_{\text{sim}}(c)$. In Mid-sign, conditioned on a choice of c , the vector y is chosen uniformly among those $y \in \text{good}_{\text{mid}}(c)$ and the vector z is computed as $z \leftarrow y + Sc$. It follows from Lemma 2 that z is uniform on $\text{good}_{\text{sim}}(c)$, as desired.

Next, we argue that

$$\Pr [C_{\text{mid}} = c] = \Pr [C_{\text{sim}} = c] \quad (45)$$

for each choice of $c \in \mathbb{H}$, from which it follows that the joint random variables $(Z_{\text{sim}}, C_{\text{sim}})$, $(Z_{\text{mid}}, C_{\text{mid}})$ are identical. It follows from Lemma 2 that $\#\text{good}_{\text{mid}}(c) = \#\text{good}_{\text{sim}}(c)$ for each $c \in \mathbb{H}$. Thus,

$$\Pr [C_{\text{mid}} = c] = \frac{\#\text{good}_{\text{mid}}(c)}{\sum_{c'} \#\text{good}_{\text{mid}}(c')} = \frac{\#\text{good}_{\text{sim}}(c)}{\sum_{c'} \#\text{good}_{\text{sim}}(c')} = \Pr [C_{\text{sim}} = c] \quad (46)$$

as desired.

B.4 Consistent-Mid-Sign

Broadly speaking, Mid-sign (Algorithm 6) behaves like Sign (Algorithm 2) except that c is selected freshly at random instead of according to some hash oracle $H(\cdot)$. It is tempting to claim that the only difference between Mid-sign and Sign is that repeated invocations of Sign always use the same hash oracle $H(\cdot)$, whereas each invocation of Mid-sign switches to another hash oracle $H'(\cdot)$ that differs from $H(\cdot)$ on a small number of randomly selected inputs.

However, there is a small probability that the random choices in a given execution of Mid-sign are not consistent with *any* hash oracle. To understand how such an inconsistency can occur, observe that each candidate (y, c) selected by Mid-sign induces an associated claim about the underlying hash oracle: that $H([Ay], \mu) = c$. Suppose Mid-sign rejects one candidate pair (y, c) because $Ay - Ec$ is not well-rounded before finally accepting another candidate pair (y', c') . If $[Ay] = [Ay']$ but $c \neq c'$ then these two candidates represent conflicting claims about the underlying hash oracle.

To address this problem we present a new sign oracle Consistent-mid-sign in Algorithm 7 and argue that its observable behaviour is negligibly close to that of Mid-sign (Algorithm 6). This fact can be stated in terms of quantum channels as follows. Let $\Psi_{c\text{-mid}}$ denote the channel described by Algorithm 7. The claim of this section is that $\Psi_{c\text{-mid}} \approx \Psi_{\text{mid}}$, meaning that $\Psi_{c\text{-mid}}(\rho) \approx \Psi_{\text{mid}}(\rho)$ for all input states ρ .

Algorithm 7 Consistent-mid-sign

Input: Message μ , public key (A, T) , secret key (S, E) .

Output: Signature (z, c) .

- 1: Initialize the dictionary $\mathcal{A} \subset (\mathbb{W} \mapsto \mathbb{H})$ to the empty dictionary $\mathcal{A} = \emptyset$.
 - 2: Choose $y \in \mathbb{Y}$ uniformly at random.
 - 3: **if** $[Ay] \in \mathcal{A}$ **then**
 - 4: $c \leftarrow \mathcal{A}[[Ay]]$
 - 5: **else**
 - 6: choose $c \in \mathbb{H}$ uniformly at random
 - 7: add $\mathcal{A}[[Ay]] \leftarrow c$ to the dictionary \mathcal{A} .
 - 8: **end if**
 - 9: $z \leftarrow y + Sc$.
 - 10: If $z \notin \mathbb{S}$ then retry at step 2.
 - 11: If $Ay - Ec$ is not well-rounded then retry at step 2.
 - 12: Re-program the hash oracle $H(\cdot)$ so that $H([Ay], \mu) = c$.
 - 13: Return (z, c) .
-

The only difference between Consistent-mid-sign and Mid-sign is that each invocation of Consistent-mid-sign remembers the random candidate pairs it selected throughout the invocation and alters them as needed so as to maintain consistency with a hash oracle. Thus, in order to prove $\Psi_{c\text{-mid}} \approx \Psi_{\text{mid}}$ it suffices to prove that only a negligibly small fraction of the random choices made by Mid-sign lead to an inconsistency that is corrected in Consistent-mid-sign.

A sequence $r = \{(y_i, c_i)\}_{i=1}^{\infty}$ of random choices made by Mid-sign leads to an inconsistently derived signature only if there exists $k \geq 2$ such that the following conditions hold:

1. $Ay_1 - Ec_1, \dots, Ay_{k-1} - Ec_{k-1}$ are not well-rounded.
2. $Ay_k - Ec_k$ is well-rounded.

3. $[\text{Ay}_k] \in \{[\text{Ay}_1], \dots, [\text{Ay}_{k-1}]\}$.

Consider the event that a random sequence r meets conditions 1–3 for some choice of $k \geq 2$, and let $\text{inconsistent}(r)$ denote the infinite disjunction of these events over all $k \geq 2$.¹⁰ We seek an upper bound on the probability of event $\text{inconsistent}(r)$ over the choice of r . To this end, define the following quantities for each choice of TESLA keys (A, T) , (S, E) :

$\text{nwr}(\text{A}, \text{E})$: The probability over $(y, c) \in \mathbb{Y} \times \mathbb{H}$ that $\text{Ay} - \text{Ec}$ is not well-rounded.
 $\text{coll}(\text{A}, \text{E})$: The maximum over all $w \in \mathbb{W}$ of the probability over $(y, c) \in \mathbb{Y} \times \mathbb{H}$ that $[\text{Ay} - \text{Ec}] = w$.

In symbols, these quantities are written

$$\text{nwr}(\text{A}, \text{E}) \stackrel{\text{def}}{=} \Pr_{(y,c) \in \mathbb{Y} \times \mathbb{H}} [\text{Ay} - \text{Ec} \text{ not well-rounded}] \quad (47)$$

$$\text{coll}(\text{A}, \text{E}) \stackrel{\text{def}}{=} \max_{w \in \mathbb{W}} \left\{ \Pr_{(y,c) \in \mathbb{Y} \times \mathbb{H}} [[\text{Ay} - \text{Ec}] = w] \right\} \quad (48)$$

In Sections B.8 and B.9 we prove bounds on these quantities that hold with high probability over the choice of TESLA keys (A, T) , (S, E) .

Broadly speaking, $\text{nwr}(\text{A}, \text{E})$ should be viewed as a constant that's noticeably smaller than 1 - for example, $\text{nwr}(\text{A}, \text{E}) = 1/2$. By contrast $\text{coll}(\text{A}, \text{E})$ is negligibly small. We prove the following.

Proposition 6 (Probability of inconsistency). *For each choice of TESLA keys (A, T) , (S, E) it holds that*

$$\Pr_r [\text{inconsistent}(r)] \leq \text{coll}(\text{A}, \text{E}) \frac{\text{nwr}(\text{A}, \text{E})}{(1 - \text{nwr}(\text{A}, \text{E}))^2} \quad (49)$$

Proof. For each $k \geq 2$ the probability with which events 1 and 2 hold is

$$\text{nwr}(\text{A}, \text{E})^{k-1} (1 - \text{nwr}(\text{A}, \text{E})). \quad (50)$$

¹⁰ In item 3 it suffices to look for a collision only between $[\text{Ay}_k]$ and any previous $[\text{Ay}_i]$; we don't need to look for a collision among arbitrary $[\text{Ay}_i] = [\text{Ay}_j]$. This is because such a collision among the bad entries is statistically identical to as if $\mathbf{c}_i = \mathbf{c}_j$. Namely, $[\text{Ay}_i]$ is rejected regardless of whether $\mathbf{c}_i = \mathbf{c}_j$. If however, \mathbf{c}_j changes $[\text{Ay}_i]$ from bad to good then that difference will be detected at $k = j$. Thus, there's no need to check for this when $k > j$.

Conditioned on those events, the probability of event 3 is

$$\Pr_{\substack{(y_1, c_1), \dots, (y_k, c_k) \notin \text{WR}(\mathbf{A}, \mathbf{E}) \\ (y_k, c_k) \in \text{WR}(\mathbf{A}, \mathbf{E})}} \left[\bigvee_{i=1}^{k-1} [\mathbf{A}y_k] = [\mathbf{A}y_i] \right] \quad (51)$$

$$\leq (k-1) \max_{w \in \mathbb{W}} \left\{ \Pr_{(y, c) \in \text{WR}(\mathbf{A}, \mathbf{E})} [[\mathbf{A}y] = w] \right\} \quad (52)$$

$$\leq (k-1) \frac{\max_{w \in \mathbb{W}} \left\{ \Pr_{(y, c) \in \mathbb{Y} \times \mathbb{H}} [[\mathbf{A}y - \mathbf{E}c] = w] \right\}}{\Pr_{(y, c) \in \mathbb{Y} \times \mathbb{H}} [\mathbf{A}y - \mathbf{E}c \text{ is well-rounded}]} \quad (53)$$

$$= (k-1) \frac{\text{coll}(\mathbf{A}, \mathbf{E})}{1 - \text{nrw}(\mathbf{A}, \mathbf{E})} \quad (54)$$

(Here we have used the notation $(y, c) \in \text{WR}(\mathbf{A}, \mathbf{E})$ to mean that $\mathbf{A}y - \mathbf{E}c$ is well-rounded.) Thus,

$$\Pr_r [\text{inconsistent}(r)] \leq \sum_{k=2}^{\infty} \text{nrw}(\mathbf{A}, \mathbf{E})^{k-1} (1 - \text{nrw}(\mathbf{A}, \mathbf{E})) (k-1) \frac{\text{coll}(\mathbf{A}, \mathbf{E})}{1 - \text{nrw}(\mathbf{A}, \mathbf{E})} \quad (55)$$

$$= \text{coll}(\mathbf{A}, \mathbf{E}) \sum_{k=2}^{\infty} (k-1) \text{nrw}(\mathbf{A}, \mathbf{E})^{k-1}. \quad (56)$$

The proposition then follows from the formula for the derivative of a geometric progression.

An immediate corollary of Proposition 6 is that

$$\|\Psi_{\text{c-mid}}(\rho) - \Psi_{\text{mid}}(\rho)\|_{\text{Tr}} < 2 \Pr_r [\text{inconsistent}(r)] \quad (57)$$

for all states ρ .

B.5 Consistent-Mid-Sign is a Mixture of Real Sign Oracles

In the previous section we introduced the sign oracle Consistent-mid-sign (Algorithm 7) and claimed that it behaves exactly like Sign (Algorithm 2) with the following exception: repeated invocations of Sign always use the same hash oracle $\mathbf{H}(\cdot)$, whereas each invocation of Consistent-mid-sign switches to another hash oracle $\mathbf{H}'(\cdot)$ that differs from $\mathbf{H}(\cdot)$ on a small fraction of randomly selected inputs.

Let us formalize this claim. Unfortunately, we must introduce some cumbersome notation. For each message μ define the symbols

- y_μ : A sequence $\{y_{\mu, i}\}_{i=1}^{\infty}$ of elements drawn randomly from \mathbb{Y} .
- $c_\mu(y_\mu)$: A sequence $\{c_{\mu, i}\}_{i=1}^{\infty}$ of elements drawn randomly from \mathbb{H} subject to the constraint that if $[\mathbf{A}y_{\mu, i}] = [\mathbf{A}y_{\mu, j}]$ then $c_{\mu, i} = c_{\mu, j}$.

The output of Consistent-mid-sign on input μ is a deterministic function of the random data $y_\mu, c_\mu(y_\mu)$. Specifically, let $k(\mu)$ denote the minimum index for which $\text{Ay}_{\mu, k(\mu)} - \text{Ec}_{\mu, k(\mu)}$ is well-rounded. Then Consistent-mid-sign outputs the signature $(y_{\mu, k(\mu)} + \text{Sc}_{\mu, k(\mu)}, c_{\mu, k(\mu)})$. For shorthand, write $\tau_\mu = (y_\mu, c_\mu(y_\mu))$.

Let $y = \{y_\mu\}_\mu$ and $c(y) = \{c_\mu(y_\mu)\}_\mu$ denote selections of random data for each possible message μ . For shorthand, write $\tau = (y, c(y))$ so that the behaviour of Consistent-mid-sign on all inputs is completely specified by τ .

For each choice of hash oracle $H(\cdot)$ and random data τ consider the hash oracle $H_\tau(\cdot)$ that agrees with $H(\cdot)$ everywhere except that $H_\tau([\text{Ay}_{\mu, i}], \mu) = c_{\mu, i}$ for each message μ and each $i = 1, \dots, k(\mu)$. In other words,

$$H_\tau(w, \mu) = \begin{cases} c_{\mu, i} & \text{if } w = [\text{Ay}_{\mu, i}] \text{ for some } i \in \{1, \dots, k(\mu)\} \\ H(w, \mu) & \text{otherwise} \end{cases}. \quad (58)$$

The behaviour of Sign (Algorithm 2) with hash oracle $H_\tau(\cdot)$ on all inputs is completely specified by y and $H_\tau(\cdot)$. Moreover, the behaviour of Sign with hash oracle $H_\tau(\cdot)$ and random data y is *identical* to the behaviour of Consistent-mid-sign with random data τ .

For each choice of random data τ define the following quantum channels:

$\Psi_{\text{c-mid}, \tau}$: The quantum channel representing the actions of Consistent-mid-sign (Algorithm 7) with randomness τ .

$\Psi_{\text{real}, H_\tau, y}$: The quantum channel representing the actions of Sign (Algorithm 2) with hash oracle $H_\tau(\cdot)$ and randomness y .

The previous observations establish

$$\Psi_{\text{c-mid}, \tau} = \Psi_{\text{real}, H_\tau, y} \quad (59)$$

for each choice of τ . Because $\Psi_{\text{c-mid}}$ is simply a uniform mixture of channels $\Psi_{\text{c-mid}, \tau}$, it follows that¹¹

$$\Psi_{\text{c-mid}} = \sum_{\tau} \text{Pr}[\tau] \Psi_{\text{real}, H_\tau, y}. \quad (60)$$

B.6 Re-Programming of Hash Oracles is Hard to Detect

Thus far we have proved that Consistent-mid-sign behaves like a mixture of real sign oracles *when viewed in isolation*. That is, for all states ρ we have

$$\Psi_{\text{c-mid}}(\rho) = \sum_{\tau} \text{Pr}[\tau] \Psi_{\text{real}, H_\tau, y}(\rho). \quad (61)$$

¹¹ Strictly speaking, $\text{Pr}[\tau]$ is zero because it represents the uniform distribution over a countably infinite set. We should switch to a probability measure on τ and use an integral instead of a summation over τ . Future versions of this manuscript will address this detail.

But we must extend this proof so that it holds even in the presence of independent information on the underlying hash oracle. In particular, for any hash oracle $H(\cdot)$ and any state ρ_H prepared using only a tractable number of queries to $H(\cdot)$ we must show that

$$\Psi_{c\text{-mid}}(\rho_H) \approx \sum_{\tau} \Pr[\tau] \Psi_{\text{real}, H_{\tau}, y}(\rho_{H_{\tau}}). \quad (62)$$

To establish this claim it suffices to show that $\rho_H \approx \rho_{H_{\tau}}$ with high probability over the choice of τ .

This claim is proven by an application of the BBBV Theorem [12, Theorem 3.3], so let us introduce the formalism necessary to state this theorem. Suppose ρ_H was prepared by some party \mathcal{R} using t queries to some hash oracle $H : X \rightarrow Y$. For each $i = 1, \dots, t$ let ρ_i denote the state of \mathcal{R} 's system immediately prior to the i th query to $H(\cdot)$. For each hash input $x \in X$ let

$$\mathcal{Q}_{\mathcal{R}(H)}(x) \stackrel{\text{def}}{=} \sum_{i=1}^t \text{Tr}(|x\rangle\langle x| \rho_i) \quad (63)$$

denote the *query magnitude* on input x for \mathcal{R} 's interaction with hash oracle $H(\cdot)$. The BBBV Theorem (or rather, a consequence of it) is as follows.

Theorem 3 ([12, Theorem 3.3]). *The following holds for each $\varepsilon > 0$. Suppose ρ_H was prepared by some party \mathcal{R} using t queries to some hash oracle $H : X \rightarrow Y$. Let $H'(\cdot)$ be a hash oracle that agrees with $H(\cdot)$ except on a subset $X' \subset X$ of inputs with the property that*

$$\sum_{x \in X'} \mathcal{Q}_{\mathcal{R}(H)}(x) \leq \frac{\varepsilon^2}{t}. \quad (64)$$

Let $\rho_{H'}$ be the state prepared when \mathcal{R} uses hash oracle $H'(\cdot)$ instead of $H(\cdot)$. It holds that $\|\rho_{H'} - \rho_H\|_{\text{Tr}} \leq \varepsilon$.

We are now ready to prove the claim of this section.

Proposition 7 (Re-Programming in TESLA). *The following holds for each choice of TESLA keys (A, T) , (S, E) and each $\delta > 0$.*

Suppose ρ_H was prepared by some party \mathcal{D} using t queries to hash oracle $H(\cdot)$. Let τ be random data and let $H_{\tau}(\cdot)$ be a hash oracle derived from $H(\cdot)$ and τ as described in Section B.5. Let $\rho_{H'}$ be the state prepared when \mathcal{D} uses hash oracle $H'(\cdot)$ instead of $H(\cdot)$.

Then $\|\rho_{H_{\tau}} - \rho_H\|_{\text{Tr}} < \delta$ except with probability at most

$$\frac{t^2}{\delta^2} \frac{\text{coll}(A, E)}{1 - \text{nwr}(A, E)} \quad (65)$$

over the choice of τ .

Proof. By Theorem 3 it suffices to prove that the quantity

$$\sum_{\mu} \sum_{i=1}^{k(\mu)} \mathcal{Q}_{\mathcal{D}(\mathbb{H})}([\mathbf{A}y_{\mu,i}], \mu) \quad (66)$$

is at most δ^2/t with high probability over the choice of τ . To this end, for each message μ let

$$X_{\mu} = \{([\mathbf{A}y], \mu) : y \in \mathbb{Y}\} \quad (67)$$

denote the set of hash inputs for message μ that are candidates for re-programming, and let

$$t_{\mu} = \sum_{x \in X_{\mu}} \mathcal{Q}_{\mathcal{D}(\mathbb{H})}(x) \quad (68)$$

denote the total query magnitude for message μ . Observe that $t = \sum_{\mu} t_{\mu}$.

The quantity (66) is maximized if for each message μ all the query magnitude t_{μ} allotted to message μ is placed on the element $w \in \mathbb{W}$ most likely to collide with $[\mathbf{A}y]$ when $y \in \mathbb{Y}$ is chosen uniformly at random. In this case, the quantity (66) is at most

$$\sum_{\mu} k(\mu) t_{\mu} \text{coll}(\mathbf{A}, \mathbf{E}). \quad (69)$$

By Markov's inequality we have

$$\Pr_{\tau} \left[\sum_{\mu} k(\mu) t_{\mu} \text{coll}(\mathbf{A}, \mathbf{E}) \geq \frac{\delta^2}{t} \right] \leq \frac{t}{\delta^2} \text{Ex}_{\tau} \left[\sum_{\mu} k(\mu) t_{\mu} \text{coll}(\mathbf{A}, \mathbf{E}) \right] \quad (70)$$

$$= \frac{t}{\delta^2} \text{coll}(\mathbf{A}, \mathbf{E}) \sum_{\mu} t_{\mu} \text{Ex}_{\tau_{\mu}} [k(\mu)] \quad (71)$$

Thus, it suffices to bound the expected number $k(\mu)$ of entries one must view from a given list τ_{μ} before encountering an entry (y, c) for which $\mathbf{A}y - \mathbf{E}c$ is well-rounded. We have

$$\text{Ex}_{\tau_{\mu}} [k(\mu)] = \sum_{k=1}^{\infty} k \text{nr}(\mathbf{A}, \mathbf{E})^{k-1} (1 - \text{nr}(\mathbf{A}, \mathbf{E})) = \frac{1}{1 - \text{nr}(\mathbf{A}, \mathbf{E})} \quad (72)$$

where the final equality follows from the formula for the derivative of a geometric progression. The proposition follows from $\sum_{\mu} t_{\mu} = t$.

B.7 The Distinguisher's State, Revisited

Recall from Section B.2 the state $\rho_{\mathbb{H}}$, which is the state of \mathcal{D} 's system immediately prior to the sign query in the first block. Let $\kappa_1 \leq q_h$ denote the query magnitude on the live-switch for the hash oracle in the first block. We proved

the following in previous sections:

$$\Psi_{\text{sim}}(\rho_{\text{H}}) = \Psi_{\text{mid}}(\rho_{\text{H}}) \quad (73)$$

$$\|\Psi_{\text{mid}}(\rho_{\text{H}}) - \Psi_{\text{c-mid}}(\rho_{\text{H}})\|_{\text{Tr}} < 2 \text{coll}(\mathbf{A}, \mathbf{E}) \frac{\text{nwr}(\mathbf{A}, \mathbf{E})}{(1 - \text{nwr}(\mathbf{A}, \mathbf{E}))^2} \quad (74)$$

$$\Psi_{\text{c-mid}}(\rho_{\text{H}}) = \sum_{\tau} \Pr[\tau] \Psi_{\text{real}, \text{H}_{\tau}, y}(\rho_{\text{H}}) \quad (75)$$

$$\Pr_{\tau} [\|\rho_{\text{H}_{\tau}} - \rho_{\text{H}}\|_{\text{Tr}} > \varepsilon] < \frac{\kappa_1^2}{\varepsilon^2} \frac{\text{coll}(\mathbf{A}, \mathbf{E})}{1 - \text{nwr}(\mathbf{A}, \mathbf{E})} \quad (76)$$

We conclude that

$$\left\| \Psi_{\text{sim}}(\rho_{\text{H}}) - \sum_{\tau} \Pr[\tau] \Psi_{\text{real}, \text{H}_{\tau}, y}(\rho_{\text{H}_{\tau}}) \right\|_{\text{Tr}} < \delta(\kappa_1) \quad (77)$$

where¹²

$$\delta(\kappa) \stackrel{\text{def}}{=} 2 \text{coll}(\mathbf{A}, \mathbf{E}) \frac{\text{nwr}(\mathbf{A}, \mathbf{E})}{(1 - \text{nwr}(\mathbf{A}, \mathbf{E}))^2} + \varepsilon + \frac{\kappa^2}{\varepsilon^2} \frac{\text{coll}(\mathbf{A}, \mathbf{E})}{1 - \text{nwr}(\mathbf{A}, \mathbf{E})}. \quad (78)$$

That is, the state of \mathcal{D} 's system at the end of the first block of an interaction with simulated oracles is $\delta(\kappa_1)$ -close to a probabilistic mixture over states of \mathcal{D} 's system, each of which could have been obtained from an interaction with real hash oracles.

As suggested in Section B.2, we continue inductively throughout the q_s blocks. As with κ_1 , let $\kappa_2, \dots, \kappa_{q_s}$ denote the query magnitude on the live-switch for blocks two through q_s . Define

$$\delta_{\text{yes}} = \sum_{i=1}^{q_s} \delta(\kappa_i) \quad (79)$$

and observe that the state of \mathcal{D} 's system at the end of an interaction with simulated oracles is δ_{yes} -close to a probabilistic mixture over states obtained from an interaction with real hash oracles.

Let us compute an upper bound on δ_{yes} . Because each block includes information from hash queries in previous blocks plus one additional hash query learned from the sign oracle, we have

$$\kappa_{i+1} \geq \kappa_i + 1. \quad (80)$$

Because \mathcal{D} is permitted at most q_h total query magnitude on the live-switch for its hash queries, we have

$$\kappa_i \leq q_h + i - 1. \quad (81)$$

¹² The final two terms of (78) are due to the fact that a $\frac{\kappa^2}{\varepsilon^2} \frac{\text{nwr}(\mathbf{A}, \mathbf{E})}{(1 - \text{nwr}(\mathbf{A}, \mathbf{E}))^2}$ -fraction of τ lead to a hash oracle $\text{H}_{\tau}(\cdot)$ for which $\|\rho_{\text{H}} - \rho_{\text{H}_{\tau}}\|_{\text{Tr}} > \varepsilon$, in which case we assume that $\rho_{\text{H}}, \rho_{\text{H}_{\tau}}$ are perfectly distinguishable. For all other τ it holds that $\rho_{\text{H}}, \rho_{\text{H}_{\tau}}$ are ε -close.

It is clear that δ_{yes} is maximized when $\kappa_1 = q_h$, which corresponds to a distinguisher who makes all q_h hash queries before making any of the q_s sign queries. Taking a loose bound $q_h + q_s$ for each κ_i , we obtain

$$\delta_{\text{yes}} < q_s \left(2 \text{coll}(\mathbf{A}, \mathbf{E}) \frac{\text{nwr}(\mathbf{A}, \mathbf{E})}{(1 - \text{nwr}(\mathbf{A}, \mathbf{E}))^2} + \varepsilon + \frac{(q_h + q_s)^2}{\varepsilon^2} \frac{\text{coll}(\mathbf{A}, \mathbf{E})}{1 - \text{nwr}(\mathbf{A}, \mathbf{E})} \right) \quad (82)$$

Finally, because the real and simulated oracles are δ_{yes} -close, it follows that

$$\Pr[\mathcal{S} \text{ output "yes"} \mid (\mathbf{A}, \mathbf{T}) \text{ yes-instance of LWE}] > \Pr[\text{forge}(\mathbf{A}, \mathbf{T})] - \delta_{\text{yes}} \quad (83)$$

as desired.

B.8 Appendix: Probability of Well-Roundedness

Let ϕ denote the probability that a random vector in \mathbb{Z}_q^m is not well-rounded:

$$\phi \stackrel{\text{def}}{=} \Pr_{x \in \mathbb{Z}_q^m} [x \text{ not well-rounded}] \leq m \left(\frac{2L}{2^d} + \frac{2L}{q} \right). \quad (84)$$

The quantity ϕ is a function of the TESLA parameters q, m, d, L . It is a constant that's noticeably smaller than 1.

Recall the definition of $\text{nwr}(\mathbf{A}, \mathbf{E})$: for TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) define $\text{nwr}(\mathbf{A}, \mathbf{E})$ as the probability over $(y, c) \in \mathbb{Y} \times \mathbb{H}$ that $\mathbf{A}y - \mathbf{E}c$ is not well-rounded:

$$\text{nwr}(\mathbf{A}, \mathbf{E}) \stackrel{\text{def}}{=} \Pr_{(y, c) \in \mathbb{Y} \times \mathbb{H}} [\mathbf{A}y - \mathbf{E}c \text{ not well-rounded}]. \quad (85)$$

We prove the following.

Lemma 3 (Probability of well-roundedness). *The following holds for all $K > 0$. With probability $1 - 1/K$ over the choice of TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) it holds that*

$$\text{nwr}(\mathbf{A}, \mathbf{E}) \leq \phi + \sqrt{\frac{K(q+1)}{\#\mathbb{Y}}}. \quad (86)$$

Proof. Our strategy is to bound the variance of $\text{nwr}(\mathbf{A}, \mathbf{E})$ over the choice of TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) and use Chebyshev's inequality. By definition,

$$\text{Var}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})] = \text{Ex}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})^2] - \text{Ex}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})]^2 \quad (87)$$

so it suffices to compute the expectation of $\text{nwr}(\mathbf{A}, \mathbf{E})$ and an upper bound on the expectation of $\text{nwr}(\mathbf{A}, \mathbf{E})^2$.

We begin by computing the expectation of $\text{nwr}(\mathbf{A}, \mathbf{E})$. We have

$$\mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})] \quad (88)$$

$$= \sum_{(\mathbf{A}, \mathbf{E})} \Pr[(\mathbf{A}, \mathbf{E})] \frac{1}{\#\mathbb{Y}\#\mathbb{H}} \sum_{(y, c)} \text{bool}[\mathbf{A}y - \mathbf{E}c \text{ not well-rounded}] \quad (89)$$

$$= \frac{1}{\#\mathbb{Y}\#\mathbb{H}} \sum_{(y, c)} \sum_{(\mathbf{A}, \mathbf{E})} \Pr[(\mathbf{A}, \mathbf{E})] \text{bool}[\mathbf{A}y - \mathbf{E}c \text{ not well-rounded}] \quad (90)$$

$$= \frac{1}{\#\mathbb{Y}\#\mathbb{H}} \sum_{(y, c)} \Pr_{(\mathbf{A}, \mathbf{E})} [\mathbf{A}y - \mathbf{E}c \text{ not well-rounded}]. \quad (91)$$

(Here we have used the notation $\text{bool}[s]$ for any statement s that can be either true or false to mean that $\text{bool}[s] = 1$ if the statement is true and $\text{bool}[s] = 0$ otherwise.) So we need to bound the probability

$$\Pr_{(\mathbf{A}, \mathbf{E})} [\mathbf{A}y - \mathbf{E}c \text{ not well-rounded}] \quad (92)$$

for each fixed choice of $(y, c) \in \mathbb{Y} \times \mathbb{H}$. There are two cases:

1. If $y \neq 0$ then $\mathbf{A}y$ is a uniformly random vector in \mathbb{Z}_q^m . So too is $\mathbf{A}y - \mathbf{E}c$, since c is fixed and \mathbf{E} is independent of \mathbf{A} . In this case, the probability (92) equals ϕ .
2. If $y = 0$ then the probability (92) equals 0, since $-\mathbf{E}c$ is well-rounded for all \mathbf{E}, c .

Case 2 occurs with probability $1/\#\mathbb{Y}$, from which it follows that

$$\mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})] = \left(1 - \frac{1}{\#\mathbb{Y}}\right) \phi. \quad (93)$$

Next, we compute an upper bound on the expectation of $\text{nwr}(\mathbf{A}, \mathbf{E})^2$. Similar to the above, we have

$$\mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})^2] = \frac{1}{(\#\mathbb{Y}\#\mathbb{H})^2} \sum_{(y, c), (y', c')} \Pr_{(\mathbf{A}, \mathbf{E})} [\mathbf{A}y - \mathbf{E}c, \mathbf{A}y' - \mathbf{E}c' \text{ not well rounded}] \quad (94)$$

and so we need to bound the probability

$$\Pr_{(\mathbf{A}, \mathbf{E})} [\mathbf{A}y - \mathbf{E}c, \mathbf{A}y' - \mathbf{E}c' \text{ not well-rounded}] \quad (95)$$

for each fixed choice of $(y, c), (y', c') \in \mathbb{Y} \times \mathbb{H}$. There are two cases:

1. If y, y' are nonzero and linearly independent then $\mathbf{A}y, \mathbf{A}y'$ are uniformly random vectors in \mathbb{Z}_q^m ; so too are $\mathbf{A}y - \mathbf{E}c, \mathbf{A}y' - \mathbf{E}c'$. In this case, the probability (95) equals ϕ^2 .
2. If y, y' are linearly dependent then the probability (95) is at most 1.

Case 2 occurs with probability at most $(q + 1)/\#\mathbb{Y}$, from which it follows that

$$\mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})^2] \leq \left(1 - \frac{q + 1}{\#\mathbb{Y}}\right) \phi^2 + \frac{q + 1}{\#\mathbb{Y}} \quad (96)$$

Combining these bounds on the expectation of $\text{nwr}(\mathbf{A}, \mathbf{E})$ and $\text{nwr}(\mathbf{A}, \mathbf{E})^2$ (and employing the inequality $1 - (q + 1)/\#\mathbb{Y} < (1 - 1/\#\mathbb{Y})^2$), we obtain the inequality

$$\text{Var}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})] \leq \frac{q + 1}{\#\mathbb{Y}}. \quad (97)$$

By Chebyshev's inequality it holds that

$$\Pr_{(\mathbf{A}, \mathbf{E})} \left[\left| \text{nwr}(\mathbf{A}, \mathbf{E}) - \mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{nwr}(\mathbf{A}, \mathbf{E})] \right| \geq \sqrt{\frac{K(q + 1)}{\#\mathbb{Y}}} \right] \leq \frac{1}{K}. \quad (98)$$

The lemma follows from the expression (93) for the expectation of $\text{nwr}(\mathbf{A}, \mathbf{E})$.

B.9 Appendix: Probability of Repetition

Let ψ denote the probability that a random vector $\mathbf{x} \in \mathbb{Z}_q^m$ is in $\Delta\mathbb{L}$:

$$\psi \stackrel{\text{def}}{=} \Pr_{\mathbf{x} \in \mathbb{Z}_q^m} [\mathbf{x} \in \Delta\mathbb{L}] \leq \left(\frac{2^{d+1}}{q}\right)^m. \quad (99)$$

The quantity ψ is a function of the TESLA parameters q, m, d . It is negligibly small.

Recall the definition of $\text{coll}(\mathbf{A}, \mathbf{E})$: for TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) define $\text{coll}(\mathbf{A}, \mathbf{E})$ as the maximum over all $\mathbf{w} \in \mathbb{W}$ of the probability over $(\mathbf{y}, \mathbf{c}) \in \mathbb{Y} \times \mathbb{H}$ that $[\mathbf{A}\mathbf{y} - \mathbf{E}\mathbf{c}] = \mathbf{w}$:

$$\text{coll}(\mathbf{A}, \mathbf{E}) \stackrel{\text{def}}{=} \max_{\mathbf{w} \in \mathbb{W}} \left\{ \Pr_{(\mathbf{y}, \mathbf{c}) \in \mathbb{Y} \times \mathbb{H}} [[\mathbf{A}\mathbf{y} - \mathbf{E}\mathbf{c}] = \mathbf{w}] \right\}. \quad (100)$$

We prove the following.

Lemma 4 (Probability of repetition). *The following holds for all $K > 0$. With probability $1 - 1/K$ over the choice of TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) it holds that*

$$\text{coll}(\mathbf{A}, \mathbf{E}) \leq K\psi. \quad (101)$$

Before proving Lemma 4 let us introduce some notation. Define the set

$$\mathbb{G}(\mathbf{A}, \mathbf{E}) \stackrel{\text{def}}{=} \{(\mathbf{y}, \mathbf{c}) \in \Delta\mathbb{Y} \times \Delta\mathbb{H} : \mathbf{A}\mathbf{y} - \mathbf{E}\mathbf{c} \in \Delta\mathbb{L}\}. \quad (102)$$

Some basic facts about the set $\mathbb{G}(\mathbf{A}, \mathbf{E})$ are listed below in Lemma 5.

Proof (Proof of Lemma 4). Let $\text{coll}'(\mathbf{A}, \mathbf{E})$ denote the probability over $(\mathbf{y}, \mathbf{c}) \in \Delta\mathbb{Y} \times \Delta\mathbb{H}$ that $(\mathbf{y}, \mathbf{c}) \in \mathbb{G}(\mathbf{A}, \mathbf{E})$:

$$\text{coll}'(\mathbf{A}, \mathbf{E}) \stackrel{\text{def}}{=} \Pr_{(\mathbf{y}, \mathbf{c}) \in \Delta\mathbb{Y} \times \Delta\mathbb{H}} [(\mathbf{y}, \mathbf{c}) \in \mathbb{G}(\mathbf{A}, \mathbf{E})]. \quad (103)$$

It follows from Lemma 5 that $\text{coll}'(\mathbf{A}, \mathbf{E}) \geq \text{coll}(\mathbf{A}, \mathbf{E})$. Thus, it suffices to prove the lemma with $\text{coll}'(\mathbf{A}, \mathbf{E})$ in place of $\text{coll}(\mathbf{A}, \mathbf{E})$.

Observe that $\text{coll}'(\mathbf{A}, \mathbf{E}) \geq 1/\#\Delta\mathbb{Y}$, which follows from the fact that $\#\mathbb{G}(\mathbf{A}, \mathbf{E}) \geq \#\Delta\mathbb{H}$ (Lemma 5). Our strategy is to bound the expectation of the positive random variable $\text{coll}'(\mathbf{A}, \mathbf{E}) - 1/\#\Delta\mathbb{Y}$ over the choice of TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) and use Markov's inequality. To this end let us compute the expectation of $\text{coll}'(\mathbf{A}, \mathbf{E})$:

$$\mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{coll}'(\mathbf{A}, \mathbf{E})] = \sum_{(\mathbf{A}, \mathbf{E})} \Pr[(\mathbf{A}, \mathbf{E})] \frac{1}{\#\Delta\mathbb{S}\#\Delta\mathbb{H}} \sum_{(\mathbf{y}, \mathbf{c})} \text{bool}[(\mathbf{y}, \mathbf{c}) \in \mathbb{G}(\mathbf{A}, \mathbf{E})] \quad (104)$$

$$= \frac{1}{\#\Delta\mathbb{S}\#\Delta\mathbb{H}} \sum_{(\mathbf{y}, \mathbf{c})} \sum_{(\mathbf{A}, \mathbf{E})} \Pr[(\mathbf{A}, \mathbf{E})] \text{bool}[(\mathbf{y}, \mathbf{c}) \in \mathbb{G}(\mathbf{A}, \mathbf{E})] \quad (105)$$

$$= \frac{1}{\#\Delta\mathbb{S}\#\Delta\mathbb{H}} \sum_{(\mathbf{y}, \mathbf{c})} \Pr_{(\mathbf{A}, \mathbf{E})} [(\mathbf{y}, \mathbf{c}) \in \mathbb{G}(\mathbf{A}, \mathbf{E})]. \quad (106)$$

So we need to bound the probability

$$\Pr_{(\mathbf{A}, \mathbf{E})} [(\mathbf{y}, \mathbf{c}) \in \mathbb{G}(\mathbf{A}, \mathbf{E})] \quad (107)$$

for each fixed choice of $(\mathbf{y}, \mathbf{c}) \in \Delta\mathbb{Y} \times \Delta\mathbb{H}$. There are two cases:

1. If $\mathbf{y} \neq 0$ then $\mathbf{A}\mathbf{y}$ is a uniformly random vector in \mathbb{Z}_q^m . So too is $\mathbf{A}\mathbf{y} - \mathbf{E}\mathbf{c}$, since \mathbf{c} is fixed and \mathbf{E} is independent of \mathbf{A} . In this case, the probability (107) is exactly ψ .
2. If $\mathbf{y} = 0$ then the probability (107) is exactly 1.

Case 2 occurs with probability exactly $1/\#\Delta\mathbb{Y}$ over the choice of (\mathbf{y}, \mathbf{c}) . It follows that

$$\mathbb{E}_{(\mathbf{A}, \mathbf{E})} [\text{coll}'(\mathbf{A}, \mathbf{E})] = \left(1 - \frac{1}{\#\Delta\mathbb{Y}}\right) \psi + \frac{1}{\#\Delta\mathbb{Y}}. \quad (108)$$

Then by Markov's inequality we have

$$\Pr_{(\mathbf{A}, \mathbf{E})} \left[\text{coll}'(\mathbf{A}, \mathbf{E}) \geq K \left(1 - \frac{1}{\#\Delta\mathbb{Y}}\right) \psi \right] \leq \frac{1}{K}. \quad (109)$$

That is, with probability at least $1 - 1/K$ over the choice of TESLA keys (\mathbf{A}, \mathbf{T}) , (\mathbf{S}, \mathbf{E}) it holds that

$$\text{coll}'(\mathbf{A}, \mathbf{E}) \leq K \left(1 - \frac{1}{\#\Delta\mathbb{Y}}\right) \psi \quad (110)$$

from which the lemma follows.

Lemma 5. For all TESLA keys (A, T) , (S, E) and all $w \in \mathbb{W}$ it holds that

$$\#\mathbb{G}(A, E) \geq \#\Delta\mathbb{H} \quad (111)$$

$$\#\mathbb{G}(A, E) \geq \#\{(y, c) \in \mathbb{Y} \times \mathbb{H} : [Ay - Ec] = w\} \quad (112)$$

Proof. The inequality (111) is straightforward: For each $c, c' \in \mathbb{H}$ we have $[Ec] = [Ec'] = [0]$, from which it follows that $(0, c - c') \in \mathbb{G}(A, E)$.

It remains to prove the inequality (112). Let $(y, c), (y', c')$ be elements of $\mathbb{Y} \times \mathbb{H}$ with $[Ay - Ec] = [Ay' - Ec'] = w$. We claim that $(y - y', c - c')$ is in $\mathbb{G}(A, E)$. It is clear that $y - y' \in \Delta\mathbb{Y}$ and $c - c' \in \Delta\mathbb{H}$. It remains to verify $A(y - y') - E(c - c') \in \Delta\mathbb{L}$. We have

$$A(y - y') - E(c - c') = Ay - Ec - (Ay' - Ec'). \quad (113)$$

Since $Ay - Ec$ and $Ay' - Ec'$ have the same high bits, it must be that $A(y - y') - E(c - c')$ is the difference of two vectors from $[-(2^{d-1} - 1), 2^{d-1}]^m$, from which it follows that $A(y - y') - E(c - c') \in \Delta\mathbb{L}$.

If $(y_1, c_1), \dots, (y_k, c_k)$ are distinct elements of $\mathbb{Y} \times \mathbb{H}$ with $[Ay_i - Ec_i] = w$ for each $i = 1, \dots, k$ then $(0, 0), (y_1 - y_2, c_1 - c_2), \dots, (y_1 - y_k, c_1 - c_k)$ must be distinct elements of $\mathbb{G}(A, E)$ ¹³. We have thus listed k distinct elements of $\mathbb{G}(A, E)$, from which the lemma follows.

C No-Instances of LWE

In this section we prove that the probability

$$\Pr[S \text{ output "yes"} \mid (A, T) \text{ no-instance of LWE}] \quad (114)$$

is small. Our strategy is to identify a correspondence between valid message-signature pairs and “good” inputs to the hash oracle. We then argue that, with high probability over the choice of LWE no-instance (A, T) and hash oracle $H(\cdot)$, the number of good inputs is a very small fraction of the total number of inputs.

Moreover, whether a given input is good is determined solely by its corresponding output from the hash oracle, implying that the only way to discover good inputs is to perform a search through an unstructured space.

Thus, a computationally bounded forger cannot expect to find good input (and hence a valid forgery), even with quantum access to the random oracle. This argument establishes the claim that the LWE-solver S outputs “yes” only with small probability, as desired.

C.1 Correspondence Between Valid Signatures and Good Hash Inputs

Let $w \in \mathbb{W}$, and let μ be an arbitrary message. For any fixed choice of random oracle $H(\cdot)$ and LWE no-instance (A, T) the hash input (w, μ) is called *good for* $H(\cdot), A, T$ if there exists $z \in \mathbb{S}$ with

$$[Az - TH(w, \mu)] = w. \quad (115)$$

¹³ By contrast with no-instances, we cannot guarantee that the negations are distinct.

Proposition 8 (Correspondence between valid signatures and good hash inputs). *If $(\mu, (z, c))$ is a valid message-signature pair for public key (A, T) and hash oracle $H(\cdot)$ then $([Az - Tc], \mu)$ is good for $H(\cdot), A, T$.*

Proof. Write $w = [Az - Tc]$. Because (z, c) is a valid signature for μ we have

$$H(w, \mu) = H([Az - Tc], \mu) = c. \quad (116)$$

Then

$$[Az - TH(w, \mu)] = [Az - Tc] = w \quad (117)$$

as desired.

A corollary of Proposition 8 is that the ability to find a message-signature pair $(\mu, (z, c))$ that is valid for public key (A, T) using q_h classical or quantum queries to $H(\cdot)$ implies the ability to find a hash input (w, μ) that is good for $H(\cdot), A, T$ using the same number of classical or quantum queries to $H(\cdot)$.

C.2 The fraction of good hash inputs

We wish to bound the probability over hash oracles $H(\cdot)$ and LWE no-instances (A, T) that a non-negligible fraction of hash inputs (w, μ) are good. To this end, define the sets

$$\mathbb{M} \stackrel{\text{def}}{=} \{(w, \mu) : w \in \mathbb{W}, \mu \text{ is a message}\} \quad (118)$$

$$\mathbb{M}(H, A, T) \stackrel{\text{def}}{=} \{(w, \mu) \in \mathbb{M} : (w, \mu) \text{ is good for } H(\cdot), A, T\} \quad (119)$$

Discussion is somewhat complicated by the fact that there is an infinite number messages, and hence \mathbb{M} and $\mathbb{M}(H, A, T)$ are infinite sets. For ease of exposition we presume a fixed, large upper bound such as 2^{2^λ} on the size of \mathbb{M} . After all, no computationally bounded forger could possibly query the hash oracle on inputs whose bit length exceeds 2^λ . Under this presumption, \mathbb{M} is a finite set and so $\#\mathbb{M}$ is a positive integer. The ratio

$$\frac{\#\mathbb{M}(H, A, T)}{\#\mathbb{M}} \quad (120)$$

is the fraction of inputs that are good.

Our goal is to show that the ratio (120) is negligibly small with high probability over the choice of $H(\cdot), A, T$. To this end, for each message $(w, \mu) \in \mathbb{M}$ define the boolean random variable

$$X_{(w, \mu)} = \begin{cases} 1 & \text{if } (w, \mu) \text{ is good for } H(\cdot), A, T \\ 0 & \text{otherwise} \end{cases} \quad (121)$$

and observe

$$\frac{\#\mathbb{M}(H, A, T)}{\#\mathbb{M}} = \frac{1}{\#\mathbb{M}} \sum_{(w, \mu) \in \mathbb{M}} X_{(w, \mu)}, \quad (122)$$

which is an average over boolean random variables. Moreover, the random variables $X_{(w,\mu)}$ are independent and so we may apply Hoeffding bounds to obtain

$$\Pr_{\mathbb{H},(\mathbf{A},\mathbf{T})} \left[\frac{\#\mathbb{M}(\mathbb{H},\mathbf{A},\mathbf{T})}{\#\mathbb{M}} - \mathbb{E}_{\mathbb{H},(\mathbf{A},\mathbf{T})} \left[\frac{\#\mathbb{M}(\mathbb{H},\mathbf{A},\mathbf{T})}{\#\mathbb{M}} \right] \geq \delta \right] \leq \exp(-2\#\mathbb{M}\delta^2). \quad (123)$$

Because $\#\mathbb{M}$ is very large relative to other TESLA parameters, we may choose δ so small that it can safely be assumed to equal zero. For example, if $\#\mathbb{M} = 2^{2^\lambda}$ then the probability (123) is negligibly small even when δ is as small as $2^{-2^{\lambda-2}}$. Thus, the ratio (120) is almost certain to be very close to its expectation

$$\mathbb{E}_{\mathbb{H},(\mathbf{A},\mathbf{T})} \left[\frac{\#\mathbb{M}(\mathbb{H},\mathbf{A},\mathbf{T})}{\#\mathbb{M}} \right]. \quad (124)$$

This expectation equals

$$\frac{1}{\#\mathbb{M}} \sum_{(w,\mu) \in \mathbb{M}} \mathbb{E}_{\mathbb{H},(\mathbf{A},\mathbf{T})} [X_{(w,\mu)}] \quad (125)$$

and by definition,

$$\mathbb{E}_{\mathbb{H},(\mathbf{A},\mathbf{T})} [X_{(w,\mu)}] = \Pr_{\mathbb{H},(\mathbf{A},\mathbf{T})} [(w,\mu) \text{ is good for } \mathbb{H}(\cdot), \mathbf{A}, \mathbf{T}]. \quad (126)$$

It remains to bound this probability for each hash input (w,μ) .

C.3 Good Hash Inputs are Rare

For each choice of $w \in \mathbb{W}$ and LWE no-instance (\mathbf{A}, \mathbf{T}) we define the set $\mathbb{H}(w, \mathbf{A}, \mathbf{T}) \subset \mathbb{H}$ as

$$\mathbb{H}(w, \mathbf{A}, \mathbf{T}) \stackrel{\text{def}}{=} \{c \in \mathbb{H} \mid \exists z \in \mathbb{S} : [\mathbf{A}z - \mathbf{T}c] = w\}. \quad (127)$$

Observe that a hash input (w,μ) is good for $\mathbb{H}(\cdot), \mathbf{A}, \mathbf{T}$ if and only if $\mathbb{H}(w,\mu) \in \mathbb{H}(w, \mathbf{A}, \mathbf{T})$. Thus,

$$\Pr_{\mathbb{H},(\mathbf{A},\mathbf{T})} [(w,\mu) \text{ is good for } \mathbb{H}(\cdot), \mathbf{A}, \mathbf{T}] = \mathbb{E}_{(\mathbf{A},\mathbf{T})} \left[\frac{\#\mathbb{H}(w, \mathbf{A}, \mathbf{T})}{\#\mathbb{H}} \right] \quad (128)$$

We prove the following.

Proposition 9 (Good Hash Inputs are Rare). *For all $w \in \mathbb{W}$ it holds that*

$$\mathbb{E}_{(\mathbf{A},\mathbf{T})} \left[\max_{w \in \mathbb{W}} \left\{ \frac{\#\mathbb{H}(w, \mathbf{A}, \mathbf{T})}{\#\mathbb{H}} \right\} \right] \leq \frac{1}{2\#\mathbb{H}} \left(1 + \frac{\#\Delta\mathbb{H}\#\Delta\mathbb{S}\#\Delta\mathbb{L}}{q^m} \right). \quad (129)$$

Proof. Define the set

$$\mathbb{D}(\mathbf{A}, \mathbf{T}) \stackrel{\text{def}}{=} \{\mathbf{b} \in \Delta\mathbb{H} : \exists y \in \Delta\mathbb{S} \text{ with } \mathbf{A}y - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}\}. \quad (130)$$

In Lemma 6 below we prove

$$\#\mathbb{H}(\mathbf{w}, \mathbf{A}, \mathbf{T}) \leq \frac{\#\mathbb{D}(\mathbf{A}, \mathbf{T}) + 1}{2} \quad (131)$$

for all $\mathbf{w} \in \mathbb{W}$. Thus, it suffices to bound the expectation

$$\mathbb{E}_{(\mathbf{A}, \mathbf{T})} \left[\frac{\#\mathbb{D}(\mathbf{A}, \mathbf{T}) + 1}{2\#\mathbb{H}} \right] = \frac{1}{2\#\mathbb{H}} \left(1 + \mathbb{E}_{(\mathbf{A}, \mathbf{T})} [\#\mathbb{D}(\mathbf{A}, \mathbf{T})] \right) \quad (132)$$

We have

$$\mathbb{E}_{(\mathbf{A}, \mathbf{T})} [\#\mathbb{D}(\mathbf{A}, \mathbf{T})] = \frac{1}{\#(\mathbf{A}, \mathbf{T})} \sum_{(\mathbf{A}, \mathbf{T})} \#\{\mathbf{b} \in \Delta\mathbb{H} : \exists \mathbf{y} \in \Delta\mathbb{S} \text{ with } \mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}\} \quad (133)$$

$$= \frac{1}{\#(\mathbf{A}, \mathbf{T})} \sum_{(\mathbf{A}, \mathbf{T})} \sum_{\mathbf{b} \in \Delta\mathbb{H}} \text{bool}[\exists \mathbf{y} \in \Delta\mathbb{S} \text{ with } \mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}] \quad (134)$$

$$\leq \frac{1}{\#(\mathbf{A}, \mathbf{T})} \sum_{(\mathbf{A}, \mathbf{T})} \sum_{\mathbf{b} \in \Delta\mathbb{H}} \sum_{\mathbf{y} \in \Delta\mathbb{S}} \text{bool}[\mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}] \quad (135)$$

$$= \sum_{\mathbf{b} \in \Delta\mathbb{H}} \sum_{\mathbf{y} \in \Delta\mathbb{S}} \frac{1}{\#(\mathbf{A}, \mathbf{T})} \sum_{(\mathbf{A}, \mathbf{T})} \text{bool}[\mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}] \quad (136)$$

$$= \sum_{\mathbf{b} \in \Delta\mathbb{H}} \sum_{\mathbf{y} \in \Delta\mathbb{S}} \Pr_{(\mathbf{A}, \mathbf{T})} [\mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}] \quad (137)$$

For each fixed choice of $\mathbf{y} \in \Delta\mathbb{S}$, $\mathbf{b} \in \Delta\mathbb{H}$, if \mathbf{A}, \mathbf{T} are uniformly random matrices then $\mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b}$ is a uniformly random vector from \mathbb{Z}_q^m . Thus, the probability $\Pr_{(\mathbf{A}, \mathbf{T})} [\mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}]$ is simply the probability that a random vector lands in $\Delta\mathbb{L}$. That is,

$$\Pr_{(\mathbf{A}, \mathbf{T})} [\mathbf{A}\mathbf{y} - \mathbf{T}\mathbf{b} \in \Delta\mathbb{L}] = \frac{\#\Delta\mathbb{L}}{q^m}. \quad (138)$$

Thus, the expectation becomes

$$\mathbb{E}_{(\mathbf{A}, \mathbf{T})} [\#\mathbb{D}(\mathbf{A}, \mathbf{T})] \leq \sum_{\mathbf{b} \in \Delta\mathbb{H}} \sum_{\mathbf{y} \in \Delta\mathbb{S}} \frac{\#\Delta\mathbb{L}}{q^m} = \frac{\#\Delta\mathbb{H}\#\Delta\mathbb{S}\#\Delta\mathbb{L}}{q^m} \quad (139)$$

as desired.

Lemma 6. *Let $\mathbb{D}(\mathbf{A}, \mathbf{T})$ be as defined in (130). For all LWE no-instances (\mathbf{A}, \mathbf{T}) and all $\mathbf{w} \in \mathbb{W}$ it holds that*

$$\#\mathbb{H}(\mathbf{w}, \mathbf{A}, \mathbf{T}) \leq \frac{\#\mathbb{D}(\mathbf{A}, \mathbf{T}) + 1}{2}. \quad (140)$$

Proof. Let $\mathbf{c}, \mathbf{c}' \in \mathbb{H}(\mathbf{w}, \mathbf{A}, \mathbf{T})$ as witnessed by $\mathbf{z}, \mathbf{z}' \in \mathbb{S}$, respectively. We claim that $\mathbf{c} - \mathbf{c}'$ is in $\mathbb{D}(\mathbf{A}, \mathbf{T})$. It is clear that $\mathbf{c} - \mathbf{c}' \in \Delta\mathbb{H}$ and $\mathbf{z} - \mathbf{z}' \in \Delta\mathbb{S}$. It remains to verify $\mathbf{A}(\mathbf{z} - \mathbf{z}') - \mathbf{T}(\mathbf{c} - \mathbf{c}') \in \Delta\mathbb{L}$. We have

$$\mathbf{A}(\mathbf{z} - \mathbf{z}') - \mathbf{T}(\mathbf{c} - \mathbf{c}') = \mathbf{A}\mathbf{z} - \mathbf{T}\mathbf{c} - (\mathbf{A}\mathbf{z}' - \mathbf{T}\mathbf{c}'). \quad (141)$$

Since $Az - Tc$ and $Az' - Tc'$ have the same high bits, it must be that $A(z - z') - T(c - c')$ is the difference of two vectors from $[-(2^{d-1} - 1), 2^{d-1}]^m$, from which it follows that $A(z - z') - T(c - c') \in \Delta\mathbb{L}$. A similar argument proves that the negation $c' - c \in \mathbb{D}(A, T)$.

If c_1, \dots, c_k are distinct elements of $\mathbb{H}(w, A, T)$ then $0, c_1 - c_2, \dots, c_1 - c_k$ must be distinct elements of $\mathbb{D}(A, T)$. Similarly, the negations $c_2 - c_1, \dots, c_k - c_1$ are also distinct elements of $\mathbb{D}(A, T)$. To see that $c_1 - c_2, \dots, c_1 - c_k$ are all distinct from their negations, observe that

$$c_1 - c_i = -(c_1 - c_j) \implies 2c_1 = c_i + c_j \implies c_i = c_j \quad (142)$$

where the final implication follows from the fact that the entries of c_1, c_i, c_j are all in $\{-1, 0, 1\}$. We have thus listed $2k - 1$ distinct elements of $\mathbb{D}(A, T)$, from which the lemma follows.

C.4 Forgers Cannot Forge on LWE No-Instances

Proposition 9 provides a bound on the fraction δ_{no} of hash inputs that are good. Moreover, since the goodness of a hash input (w, μ) depends solely on whether $H(w, \mu)$ is in $\mathbb{H}(w, A, T)$, the set of all good hash inputs is a randomly selected set. Thus, the only way to find a good hash input is via search through an unstructured space.

It then follows from lower bounds for quantum search [19] that any algorithm making no more than q_h quantum queries to $H(\cdot)$ finds a good hash input—and thus a valid TESLA forgery—with probability no larger than

$$2(q_h + 1)\sqrt{\delta_{\text{no}}}. \quad (143)$$

We therefore obtain

$$\Pr[\mathcal{S} \text{ output “yes”} \mid (A, T) \text{ no-instance of LWE}] \leq 2(q_h + 1)\sqrt{\delta_{\text{no}}}. \quad (144)$$

D Security: Putting it all Together

Assuming that no algorithm with run time comparable to that of \mathcal{S} can solve LWE with success bias exceeding ε , we have:

$$\varepsilon \geq \Pr[\mathcal{S} \text{ output “yes”} \mid (A, T) \text{ yes-instance of LWE}] \quad (145)$$

$$- \Pr[\mathcal{S} \text{ output “yes”} \mid (A, T) \text{ no-instance of LWE}]. \quad (146)$$

We know that

$$\Pr[\mathcal{S} \text{ output “yes”} \mid (A, T) \text{ yes-instance of LWE}] \geq \Pr[\text{forge}(A, T)] - \delta_{\text{yes}}. \quad (147)$$

Against a quantum forger, we have that

$$\Pr[\mathcal{S} \text{ output “yes”} \mid (A, T) \text{ no-instance of LWE}] \leq 2(q_h + 1)\sqrt{\delta_{\text{no}}}, \quad (148)$$

implying that

$$\Pr [\text{forge}(\mathbf{A}, \mathbf{T})] \leq \delta_{\text{yes}} + 2(q_h + 1)\sqrt{\delta_{\text{no}}} + \varepsilon. \quad (149)$$

Against a *classical* forger, we can remove the quadratic speedup on the lower bound query complexity, and the probability becomes

$$\Pr [\text{forge}(\mathbf{A}, \mathbf{T})] \leq \delta_{\text{yes}} + q_h \cdot \delta_{\text{no}} + \varepsilon. \quad (150)$$

We make some simplifying assumptions on the choice of TESLA parameters. These assumptions are not necessary in order to derive a negligibly small upper bound on the forger's success probability—they merely facilitate a simplified statement of the upper bound.

Definition 3 (Convenient TESLA Parameters). *TESLA parameters are convenient if the following bounds hold:*

$$\phi + \sqrt{\frac{2^\lambda(q+1)}{\#\mathbb{Y}}} < 1/2 \quad (151)$$

$$\#\Delta\mathbb{H}\#\Delta\mathbb{Y}\#\Delta\mathbb{L} < q^m. \quad (152)$$

All our proposed parameter sets for TESLA meet this condition.

We now incorporate our bounds on δ_{yes} and δ_{no} in order to derive an explicit upper bound on the forger's success probability.

By applying Lemma 3 with $K = 2^\lambda$, as well as equation 151 in definition 3, we can note that with probability $1 - 2^{-\lambda}$ over the choice of (\mathbf{A}, \mathbf{E}) , $\text{nwr}(\mathbf{A}, \mathbf{E}) \leq 1/2$. Incorporating this into equation 82, we see that

$$\delta_{\text{yes}} \leq q_s \gamma + 4q_s \text{coll}(\mathbf{A}, \mathbf{E}) \left(1 + \frac{(q_h + q_s)^2}{2\gamma^2} \right). \quad (153)$$

From lemma 4, we have a bound on $\text{coll}(\mathbf{A}, \mathbf{E})$ that holds with probability $1 - 1/K_{\text{coll}}$.

$$\delta_{\text{yes}} \leq q_s \gamma + 4q_s \left(\frac{2^{d+1}}{q} \right)^m K_{\text{coll}} \left(1 + \frac{(q_h + q_s)^2}{2\gamma^2} \right) \quad (154)$$

At this point, we note that our result on this bound holds for whatever γ we may choose. As we want the first term to be exponentially small, we will select $\gamma = \frac{1}{2^\lambda q_s}$. Then, using the simplification that $1 + \frac{(q_h + q_s)^2}{2\gamma^2} \approx \frac{(q_h + q_s)^2}{2\gamma^2}$ we get

$$\delta_{\text{yes}} \leq \frac{1}{2^\lambda} + \frac{2^{m(d+1)+2\lambda+1}}{q^m} (q_h + q_s)^2 q_s^3 K_{\text{coll}}. \quad (155)$$

From proposition 9 we have

$$\delta_{\text{no}} \leq \frac{1}{2\#\mathbb{H}} \left(1 + \frac{\#\Delta\mathbb{H}\#\Delta\mathbb{S}\#\Delta\mathbb{L}}{q^m} \right). \quad (156)$$

Using equation 152 of Definition 3 we can simplify this bound on δ_{no} to

$$\delta_{\text{no}} \leq \frac{1}{\#\mathbb{H}}. \quad (157)$$

Finally, we substitute this, and our bound for δ_{yes} into (149). We also note that $\#\mathbb{H} = 2^h \binom{n'}{h}$. We also must consider the probability with which our bounds do not hold. Doing this, we get that $\Pr[\text{forge}(\mathbf{A}, \mathbf{T})]$ is at most

$$\varepsilon + \frac{1}{2^\lambda} + \frac{2^{m(d+1)+2\lambda+1}}{q^m} (q_h + q_s)^2 q_s^3 K_{\text{coll}} + 2(q_h + 1) \sqrt{\frac{1}{2^h \binom{n'}{h}}} + \frac{1}{K_{\text{NWR}}} + \frac{1}{K_{\text{coll}}} \quad (158)$$

Then by choosing each K value to be 2^λ , we get that this is equal to

$$\varepsilon + \frac{3}{2^\lambda} + \frac{2^{m(d+1)+3\lambda+1}}{q^m} (q_h + q_s)^2 q_s^3 + 2(q_h + 1) \sqrt{\frac{1}{2^h \binom{n'}{h}}}. \quad (159)$$

Classically, we can similarly derive that the adversary's success is bounded by

$$\varepsilon + \frac{3}{2^\lambda} + \frac{2^{m(d+1)+3\lambda+1}}{q^m} (q_h + q_s)^2 q_s^3 + q_h \frac{1}{2^h \binom{n'}{h}} \quad (160)$$