

Submission to NIST's post-quantum project (2nd round):  
lattice-based digital signature scheme qTESLA

Name of the cryptosystem: qTESLA

Principal and auxiliary submitters:

**Nina Bindel**,  
(Principal submitter)                      University of Waterloo,  
QNC 4316, 200 University Ave West, Waterloo, ON N2L  
3G1, Canada  
Email: [nlbindel@uwaterloo.ca](mailto:nlbindel@uwaterloo.ca),  
Phone: +1 519-888-39072

Signature: 

**Sedat Akleylek**,                      Ondokuz Mayıs University, Turkey  
**Erdem Alkim**,                      Ondokuz Mayıs University, Turkey and Fraunhofer SIT,  
Germany  
**Paulo S. L. M. Barreto**,              University of Washington Tacoma, USA  
**Johannes Buchmann**,              Technische Universität Darmstadt, Germany  
**Edward Eaton**,                      ISARA Corporation and University of Waterloo, Canada  
**Gus Gutoski**,                      ISARA Corporation, Canada  
**Juliane Krämer**,                      Technische Universität Darmstadt, Germany  
**Patrick Longa**,                      Microsoft Research, USA  
**Harun Polat**,                      Technische Universität Darmstadt, Germany  
**Jefferson E. Ricardini**,              LG Electronics, USA  
**Gustavo Zanon**,                      University of São Paulo, Brazil

Inventors of the cryptosystem:

All the submitters by name based on a previous scheme by Shi Bai and Steven Galbraith and several other previous works, as explained in the body of this document.

Owners of the cryptosystem:

None (dedicated to the public domain).

# Changelog

This is the changelog of this document and the corresponding implementation of qTESLA.

Version	Date	Description of changes
1.0	11/30/2017	<ul style="list-style-type: none"><li>• Original submission to NIST (1st round).</li></ul>
2.0	06/14/2018	<ul style="list-style-type: none"><li>• qTESLA described generically using <math>k &gt; 1</math> R-LWE samples.</li><li>• Signing algorithm changed to probabilistic (instead of deterministic).</li><li>• New parameter sets proposed: three heuristic and two provably-secure parameter sets.</li><li>• Improved explanation of the realization of the different functions (Section 2.5).</li><li>• Minor changes and refinements throughout the document.</li><li>• C-only reference implementation corrected; e.g., to have proper protection against timing and cache attacks.</li><li>• C-only reference implementation improved; e.g., to have more resilience against certain fault attacks.</li></ul>
2.1	06/30/2018	<ul style="list-style-type: none"><li>• Corrected typo that assumed an exponent <math>d</math> or <math>d + 1</math> instead of <math>d - 1</math> or <math>d</math> (resp.) in some places.</li><li>• Small fix in the bounds of the signature rejection evaluation, line 18 of Algorithm 7. Updated KATs accordingly.</li><li>• Applied notation <math>\text{mod}^\pm</math> to denote the use of a centered representative in Algorithms 7 and 8.</li><li>• Updated correctness proof in Section 2.4.</li></ul>
2.2	08/27/2018	<ul style="list-style-type: none"><li>• Corrected typo in the definition of <math>\text{mod}^\pm</math>.</li><li>• Corrected typo in the signature verification algorithm, line 6 of Algorithm 8.</li><li>• Some corrections in Algorithm 10. Rearranged if-blocks to maximize use of cSHAKE128's output. Updated KATs.</li><li>• Corrected typos in Algorithm 14. Rearranged if-blocks to maximize use of cSHAKE128's output.</li><li>• Added rejection of value <math>B + 1</math> during sampling of <math>y</math>, Algorithm 12.</li></ul>
2.3	10/31/2018	<ul style="list-style-type: none"><li>• Introduced hash function <math>G : \{0, 1\}^* \rightarrow \{0, 1\}^{512}</math> that maps a message to a 512-bit string. See Algorithms 7 and 8.</li><li>• Corrected typo in the definition of the encoding function <math>\text{Enc}</math> in Section 2.3.</li><li>• Modified expression for hash function <math>H</math> in Algorithms 7 and 8 to match function definition in Algorithm 13.</li><li>• Corrected typo in line 1 of Algorithm 7. Counter is initialized to 1 instead of 0.</li><li>• Corrected typo in Section 2.5.1. Hashing in function <math>H</math> is instantiated with SHAKE, not cSHAKE.</li><li>• Corrected typo in line 5 of Algorithm 10.</li><li>• Rounded parameter <math>\xi</math> to the immediately smaller integer for parameter sets qTESLA-I and qTESLA-III-size. Updated KATs.</li></ul>

Version	Date	Description of changes
2.4	01/25/2019	<ul style="list-style-type: none"> <li>• Original Gaussian sampler based on the Bernoulli-based rejection sampling is replaced by a new portable and constant-time CDT-based Gaussian sampler that does not require floating-point arithmetic.</li> <li>• Added new AVX2-optimized implementations for the heuristic parameter sets.</li> </ul>
2.5	03/30/2019	<ul style="list-style-type: none"> <li>• Original submission to NIST (2nd round).</li> <li>• Revised conjecture used in the security reduction. Explained the usage of the conjecture in context and provided a heuristic argument for why it is true. Added a script to experimentally search for possible counterexamples.</li> <li>• Added parameter sets for levels II and V.</li> <li>• Added section about a qTESLA variant with smaller public keys and parameter sets <code>qTESLA-I-s</code>, <code>qTESLA-II-s</code>, <code>qTESLA-III-s</code>, <code>qTESLA-V-s</code>, and <code>qTESLA-V-size-s</code>, following [28].</li> <li>• Added a section to discuss a variant with <math>n</math> being a non-power-of-two.</li> </ul>
2.6	04/26/2019	<ul style="list-style-type: none"> <li>• Removed the section about a qTESLA variant with smaller public keys and parameter sets <code>qTESLA-I-s</code>, <code>qTESLA-II-s</code>, <code>qTESLA-III-s</code>, <code>qTESLA-V-s</code>, and <code>qTESLA-V-size-s</code>.</li> </ul>
2.7	08/19/2019	<ul style="list-style-type: none"> <li>• Removed heuristic parameter sets and the corresponding implementations.</li> <li>• Fixed minor bugs in the implementation.</li> </ul>
2.8	11/08/2019	<ul style="list-style-type: none"> <li>• Added the digest of <math>G(t_1, \dots, t_k)</math> during hashing of <math>c'</math>.</li> <li>• Hashing with <math>G</math> has been adjusted to produce a digest of 320 bits.</li> <li>• Fixed typos in the specifications.</li> </ul>
2.9	04/13/2020	<ul style="list-style-type: none"> <li>• Fixed typos in this specifications document.</li> <li>• Tightened the bound in Conjecture 5.</li> </ul>

# Contents

<b>1</b>	<b>Introduction</b>	<b>5</b>
1.1	qTESLA highlights . . . . .	5
1.2	Related work . . . . .	6
<b>2</b>	<b>Specification</b>	<b>8</b>
2.1	Notation . . . . .	8
2.2	Basic signature scheme . . . . .	9
2.3	Formal description of qTESLA . . . . .	10
2.4	Correctness of the scheme . . . . .	15
2.5	Realization of the required functions . . . . .	16
2.5.1	Hash and pseudorandom functions . . . . .	16
2.5.2	Pseudorandom bit generation of seeds, PRF <sub>1</sub> . . . . .	17
2.5.3	Generation of $\mathbf{a}_1, \dots, \mathbf{a}_k$ . . . . .	18
2.5.4	Gaussian sampling . . . . .	19
2.5.5	Sampling of $y$ . . . . .	20
2.5.6	Hash-based function $H$ . . . . .	22
2.5.7	Encoding function . . . . .	23
2.5.8	Polynomial multiplication and the number theoretic transform . . . . .	24
2.6	System parameters and parameter selection . . . . .	26
<b>3</b>	<b>Performance analysis</b>	<b>31</b>
3.1	Reference implementations . . . . .	31
3.2	Performance of qTESLA on x64 Intel . . . . .	32
3.3	AVX2 optimizations . . . . .	32
<b>4</b>	<b>Known answer values</b>	<b>33</b>
<b>5</b>	<b>Expected security strength</b>	<b>34</b>
5.1	Provable security in the quantum random oracle model . . . . .	34
5.2	Bit security of our proposed parameter sets . . . . .	35
5.2.1	Correspondence between security and hardness . . . . .	35
5.2.2	Estimation of the R-LWE hardness . . . . .	36
5.3	Resistance to implementation attacks . . . . .	38
5.3.1	Side-channel analysis . . . . .	38
5.3.2	Fault analysis . . . . .	39
<b>6</b>	<b>Advantages and limitations</b>	<b>39</b>
<b>A</b>	<b>Sketch of proof of Theorem 2</b>	<b>48</b>

# 1 Introduction

This document presents a detailed specification of **qTESLA**, a family of provably-secure post-quantum signature schemes based on the hardness of the decisional Ring Learning With Errors (R-LWE). **qTESLA** is an efficient variant of the Bai-Galbraith signature scheme — which in turn is based on the “Fiat-Shamir with Aborts” framework by Lyubashevsky — adapted to the setting of ideal lattices.

Concretely, this document proposes *two* parameter sets targeting *two* security levels:

- (1) **qTESLA-p-I**: NIST’s security category 1.
- (2) **qTESLA-p-III**: NIST’s security category 3.

The present document is organized as follows. In the remainder of this section, we summarize the main features of **qTESLA** and describe related previous work. In Section 2, we provide the specification details of the scheme, including a basic and a formal algorithmic description, the functions that are required for the implementation, and the proposed parameter sets. In Section 3, we analyze the performance of our implementations. Section 4 includes the details of the known answer values. Then, we discuss the (provable) security of our proposal in Section 5, including an analysis of the concrete security level and the security against implementation attacks. Section 6 ends this document with a summary of the advantages and limitations of **qTESLA**.

## 1.1 **qTESLA** highlights

**qTESLA**’s main features can be summarized as follows:

- **Simplicity.** **qTESLA** is simple and easy to implement, and its design makes possible the realization of compact and portable implementations that achieve good performance. In addition, the use of a simplified Gaussian sampler is limited to key generation.
- **Security foundation.** The underlying security of **provably-secure qTESLA** is based on the hardness of the decisional R-LWE problem, and comes accompanied by a tight security proof in the (quantum) random oracle model.
- **Practical security.** By design, **qTESLA** facilitates secure implementations. In particular, it supports *constant-time* implementations (i.e., implementations that are secure against timing and cache side-channel attacks since their execution time does not depend on secret values), and is inherently protected against certain simple yet powerful fault attacks. Moreover, it also comes with a built-in safeguard to protect

against Key Substitution (KS) attacks [20, 46] (a.k.a. Duplicate Signature Key Selection (DSKS) attacks) and, thus, improved security in the multi-user setting; see also [37].

- **Scalability and portability.** qTESLA’s simple design makes it straightforward to easily support more than one security level and parameter set with a single, efficient portable implementation.

**Security.** The security of qTESLA is proven using the reductionist approach, i.e., we construct an efficient reduction that turns any successful adversary against qTESLA into one that solves R-LWE.

Our qTESLA parameter sets then go one step further: That is, qTESLA instantiations are *provably* secure in the Quantum Random Oracle Model (QROM) since they are chosen taking into account the respective security reduction from R-LWE. Despite these security assurances, qTESLA achieves relatively good performance and offers relatively compact signatures.

## 1.2 Related work

The signature scheme proposed in this submission is the result of a long line of research. The first work in this line is the signature scheme proposed by Bai and Galbraith [12] which is based on the Fiat-Shamir construction of Lyubashevsky [44]. The scheme by Bai and Galbraith is constructed over standard lattices and comes with a (non-tight) security reduction from the LWE and the short integer solution (SIS) problems in the random oracle model. Dagdelen *et al.* [26] presented improvements and the first implementation of the Bai-Galbraith scheme. The scheme was subsequently studied under the name TESLA by Alkim, Bindel, Buchmann, Dagdelen, Eaton, Gutoski, Krämer, and Pawlega [8], who provided an alternative security reduction from the LWE problem in the quantum random oracle model.

A variant of TESLA over ideal lattices was derived under the name ring-TESLA [1] by Akleyek, Bindel, Buchmann, Krämer, and Marson. Since then, there have appeared subsequent works aimed at improving the efficiency of the scheme [14, 35]. Most notably, a scheme called TESLA# [14] by Barreto, Longa, Naehrig, Ricardini, and Zanon included several implementation improvements. Finally, several works [18, 19, 31] have focused on the analysis of ring-TESLA against side-channel and fault attacks.

In this document, we consolidate the most relevant features of the prior works with the goal of designing the quantum-secure signature scheme qTESLA.

## Acknowledgments

We are grateful to Michael Naehrig for his valuable feedback on an earlier version of this document. Thanks to Peter Pessl and Matthias Kannwischer for comments and suggestions for a previous version of our implementations, to Kevin Kane and François Gérard for pointing out various typos and implementation bugs, and for suggesting several improvements in a previous version of this document, to Wen Wang for pointing out an optimization in the NTT algorithm. Special thanks also to Vadim Lyubashevsky for pointing out a typo in an earlier version of the parameter generation script that invalidated the “provably-secure” property of the parameter sets originally submitted to NIST on November 2017 (version 1.0 of this document), and to Vadim Lyubashevsky and Peter Schwabe for pointing out a mistake in an earlier version of this document that contained a variant with smaller public keys and heuristic parameters that were lacking security estimates with respect to the short integer solution problem (SIS). We thank Fernando Virdia, Martin Albrecht and Shi Bai for fruitful discussions and helpful advice on the hardness estimation of SIS. Finally, we thank Joo Woo for pointing out that the bound in Conjecture 5 was too loose.

Sedat Akleylek and Erdem Alkim are partially supported by TÜBİTAK under grant no. EEEAG-116E279. Nina Bindel and Juliane Krämer are partially supported by the German Research Foundation (DFG) as part of project P1 within the CRC 1119 CROSSING. Edward Eaton is partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). Jefferson Ricardini is supported by the joint São Paulo Research Foundation (FAPESP)/Intel Research grant 2015/50520-6 “Efficient Post-Quantum Cryptography for Building Advanced Security Applications”.

## 2 Specification

In this section, we define basic notation and give an informal description of the basic scheme that is used to specify qTESLA. A formal specification of qTESLA’s key generation, signing, and verification algorithms follows in Section 2.3. The correctness of the scheme is discussed in Section 2.4. We describe the implementation of the functions required by qTESLA in Section 2.5. Finally, we explain all the system parameters and the proposed parameter sets in Section 2.6.

### 2.1 Notation

*Rings.* Let  $q$  be an odd prime throughout the document if not stated otherwise. Let  $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$  denote the quotient ring of integers modulo  $q$ , and let  $\mathcal{R}$  and  $\mathcal{R}_q$  denote the rings  $\mathbb{Z}[x]/\langle x^n + 1 \rangle$  and  $\mathbb{Z}_q[x]/\langle x^n + 1 \rangle$ , respectively. Given  $f = \sum_{i=0}^{n-1} f_i x^i \in \mathcal{R}$ , we define the reduction of  $f$  modulo  $q$  to be  $(f \bmod q) = \sum_{i=0}^{n-1} (f_i \bmod q) x^i \in \mathcal{R}_q$ . Let  $\mathbb{H}_{n,h} = \{\sum_{i=0}^{n-1} f_i x^i \in \mathcal{R} \mid f_i \in \{-1, 0, 1\}, \sum_{i=0}^{n-1} |f_i| = h\}$ , and  $\mathcal{R}_{[B]} = \{\sum_{i=0}^{n-1} f_i x^i \in \mathcal{R} \mid f_i \in [-B, B]\}$ .

*Rounding operators.* Let  $d \in \mathbb{N}$  and  $c \in \mathbb{Z}$ . For an even (odd) modulus  $m \in \mathbb{Z}_{\geq 0}$ , define  $c' = c \bmod^{\pm} m$  as the unique element  $c'$  such that  $-m/2 < c' \leq m/2$  (resp.  $- \lfloor m/2 \rfloor \leq c' \leq \lfloor m/2 \rfloor$ ) and  $c' = c \bmod m$ . We then define the functions  $[\cdot]_L : \mathbb{Z} \rightarrow \mathbb{Z}$ ,  $c \mapsto (c \bmod^{\pm} q) \bmod^{\pm} 2^d$ , and  $[\cdot]_M : \mathbb{Z} \rightarrow \mathbb{Z}$ ,  $c \mapsto (c \bmod^{\pm} q - [c]_L)/2^d$ . Hence,  $c \bmod^{\pm} q = 2^d \cdot [c]_M + [c]_L$  for  $c \in \mathbb{Z}$ . These definitions are extended to polynomials by applying the operators to each polynomial coefficient, i.e.,  $[f]_L = \sum_{i=0}^{n-1} [f_i]_L x^i$  and  $[f]_M = \sum_{i=0}^{n-1} [f_i]_M x^i$  for a given  $f = \sum_{i=0}^{n-1} f_i x^i \in \mathcal{R}$ .

*Infinity norm.* Given  $f \in \mathcal{R}$ , the function  $\max_k(f)$  returns the  $k$ -th largest absolute coefficient of  $f$ . That is, if the coefficients of  $f$  are reordered as to produce a polynomial  $g$  with coefficients ordered (without losing any generality) as  $|g_1| \geq |g_2| \geq \dots \geq |g_n|$ , then  $\max_k(f) = g_k$ . For an element  $c \in \mathbb{Z}_q$ , we have that  $\|c\|_{\infty} = |c \bmod^{\pm} q|$ , and we define the infinity norm for a polynomial  $f \in \mathcal{R}$  as  $\|f\|_{\infty} = \max_i \|f_i\|_{\infty}$ .

*Representation of polynomials and bit strings.* We write a given polynomial  $f \in \mathcal{R}_q$  as  $\sum_{i=0}^{n-1} f_i x^i$  or, in some instances, as the coefficient vector  $(f_0, f_1, \dots, f_{n-1}) \in \mathbb{Z}_q^n$ . When it is clear by the context, we represent some specific polynomials with a subscript (e.g., to represent polynomials  $a_1, \dots, a_k$ ). In these cases, we write  $a_j = \sum_{i=0}^{n-1} a_{j,i} x^i$ , and the corresponding vector representation is given by  $a_j = (a_{j,0}, a_{j,1}, \dots, a_{j,n-1}) \in \mathbb{Z}_q^n$  for  $j = 1, \dots, k$ . In the case of sparse polynomials  $c \in \mathbb{H}_{n,h}$ , these polynomials are encoded as the two arrays  $pos.list \in \{0, \dots, n-1\}^h$  and  $sign.list \in \{-1, 1\}^h$  representing the positions and signs of the nonzero coefficients of  $c$ , respectively. We denote this by  $c \triangleq \{pos.list, sign.list\}$ . In some cases,  $s$ -bit strings  $r \in \{0, 1\}^s$  are written as vectors over the set  $\{0, 1\}$ , in which



an element in the  $i$ -th position is represented by  $r_i$ . This applies analogously to other sets. Multiple instances of the same set are represented by appending an additional superscript. For example,  $\{0, 1\}^{s,t}$  corresponds to  $t$   $s$ -bit strings each defined over the set  $\{0, 1\}$ .

*Distributions.* The centered discrete Gaussian distribution with standard deviation  $\sigma$  is defined to be  $\mathcal{D}_\sigma = \rho_\sigma(c)/\rho_\sigma(\mathbb{Z})$  for  $c \in \mathbb{Z}$ , where  $\sigma > 0$ ,  $\rho_\sigma(c) = \exp(-\frac{c^2}{2\sigma^2})$ , and  $\rho_\sigma(\mathbb{Z}) = 1 + 2 \sum_{c=1}^{\infty} \rho_\sigma(c)$ . We write  $x \leftarrow_\sigma \mathbb{Z}$  to denote sampling of a value  $x$  with distribution  $\mathcal{D}_\sigma$ . For a polynomial  $f \in \mathcal{R}$ , we write  $f \leftarrow_\sigma \mathcal{R}$  to denote sampling each coefficient of  $f$  with distribution  $\mathcal{D}_\sigma$ . Moreover, for a finite set  $S$ , we denote sampling  $s$  uniformly from  $S$  with  $s \leftarrow_{\S} S$  or  $s \leftarrow \mathcal{U}(S)$ .

## 2.2 Basic signature scheme

Informal descriptions of the algorithms that give rise to the signature scheme qTESLA are shown in Algorithms 1, 2, and 3. These algorithms require two basic terms, namely, *B-short* and *well-rounded*, which are defined below.

Let  $q, L_E, L_S, E, S, B$ , and  $d$  be system parameters that denote the modulus, the bound constant for error polynomials, the bound constant for the secret polynomial, two rejection bounds used during signing and verification that are related to  $L_E$  and  $L_S$ , the bound for the random polynomial at signing, and the rounding value, respectively. An integer polynomial  $y$  is *B-short* if each coefficient is at most  $B$  in absolute value. We call an integer polynomial  $w$  *well-rounded* if  $w$  is  $(\lfloor q/2 \rfloor - E)$ -short and  $[w]_L$  is  $(2^{d-1} - E)$ -short.

In Algorithms 1–3, we assume for simplicity that the hash oracle  $H(\cdot)$  maps to  $\mathbb{H}$ , where  $\mathbb{H}$  denotes the set of polynomials  $c \in \mathcal{R}$  with coefficients in  $\{-1, 0, 1\}$  with exactly  $h$  nonzero entries, i.e., we ignore the encoding function  $\text{Enc}$  introduced in Section 2.3.

Because of the random generation of the polynomial  $y$  (see line 1 of Alg. 2), Algorithm 2 is described as a *non-deterministic* algorithm. This property implies that different randomness is required for each signature. For the formal specification of qTESLA we incorporate

---

### Algorithm 1 Informal description of the key generation

---

**Require:** -

**Ensure:** Secret key  $sk = (s, e_1, \dots, e_k, a_1, \dots, a_k)$ , and public key  $pk = (a_1, \dots, a_k, t_1, \dots, t_k)$

---

- 1:  $a_1, \dots, a_k \leftarrow \mathcal{R}_q$  ring elements.
  - 2: Choose  $s \in \mathcal{R}$  with entries from  $\mathcal{D}_\sigma$ . Repeat step if the  $h$  largest entries of  $s$  sum to at least  $L_S$ .
  - 3: For  $i = 1, \dots, k$ : Choose  $e_i \in \mathcal{R}$  with entries from  $\mathcal{D}_\sigma$ . Repeat step at iteration  $i$  if the  $h$  largest entries of  $e_i$  sum to at least  $L_E$ .
  - 4: For  $i = 1, \dots, k$ : Compute  $t_i \leftarrow a_i s + e_i \in \mathcal{R}_q$ .
  - 5: Return  $sk = (s, e_1, \dots, e_k, a_1, \dots, a_k)$  and  $pk = (a_1, \dots, a_k, t_1, \dots, t_k)$ .
-

---

**Algorithm 2** Informal description of the signature generation

---

**Require:** Message  $m$ , secret key  $sk = (s, e_1, \dots, e_k, a_1, \dots, a_k)$

**Ensure:** Signature  $(z, c)$

---

- 1: Choose  $y$  uniformly at random among  $B$ -short polynomials in  $\mathcal{R}_q$ .
  - 2:  $c \leftarrow \text{H}([a_1y]_M, \dots, [a_ky]_M, m)$ .
  - 3: Compute  $z \leftarrow y + sc$ .
  - 4: If  $z$  is not  $(B - S)$ -short then retry at step 1.
  - 5: For  $i = 1, \dots, k$ : If  $a_iy - e_ic$  is not well-rounded then retry at step 1.
  - 6: Return  $(z, c)$ .
- 

---

**Algorithm 3** Informal description of the signature verification

---

**Require:** Message  $m$ , public key  $pk = (a_1, \dots, a_k, t_1, \dots, t_k)$ , and signature  $(z, c)$

**Ensure:** “accept” or “reject” signature

---

- 1: If  $z$  is not  $(B - S)$ -short then return reject.
  - 2: For  $i = 1, \dots, k$ : Compute  $w_i \leftarrow a_iz - t_ic \in \mathcal{R}_q$ .
  - 3: If  $c \neq \text{H}([w_1]_M, \dots, [w_k]_M, m)$  then return reject.
  - 4: Return accept.
- 

an additional improvement: qTESLA requires a combination of fresh randomness and a fixed value for the generation of  $y$  (see Section 2.3). This design feature is added in order to prevent some implementation pitfalls and, at the same time, protect against some simple but devastating fault attacks. We discuss the advantages of our approach in Section 5.3.

### 2.3 Formal description of qTESLA

qTESLA is parameterized by  $\lambda, \kappa, n, k, q, \sigma, L_E, L_S, E, S, B, d, h$ , and  $b_{\text{GenA}}$ ; see Table 3 in Section 2.6 for a detailed description of all the system parameters. The following functions are required for the implementation of the scheme:

- The pseudorandom function  $\text{PRF}_1 : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{\kappa, k+3}$ , which takes as input a seed **pre-seed** that is  $\kappa$  bits long and maps it to  $(k + 3)$  seeds of  $\kappa$  bits each.
- The collision-resistant hash function  $\text{G} : \{0, 1\}^* \rightarrow \{0, 1\}^{320}$ , which maps a given input string to a 320-bit string.
- The pseudorandom function  $\text{PRF}_2 : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \times \{0, 1\}^{320} \rightarrow \{0, 1\}^\kappa$ , which takes as inputs  $\text{seed}_y$  and the random value  $r$ , each  $\kappa$  bits long, and the hash  $\text{G}$  of a message  $m$ , which is 320-bit long, and maps them to the  $\kappa$ -bit seed **rand**.
- The generation function of the public polynomials  $a_1, \dots, a_k$ ,  $\text{GenA} : \{0, 1\}^\kappa \rightarrow \mathcal{R}_q^k$ ,

which takes as input the  $\kappa$ -bit seed  $\text{seed}_a$  and maps it to  $k$  polynomials  $a_i \in \mathcal{R}_q$ .

- The Gaussian sampler function  $\text{GaussSampler} : \{0, 1\}^\kappa \times \mathbb{Z} \rightarrow \mathcal{R}$ , which takes as inputs a  $\kappa$ -bit seed  $\text{seed} \in \{\text{seed}_s, \text{seed}_{e_1}, \dots, \text{seed}_{e_k}\}$  and a nonce counter  $\in \mathbb{Z}_{>0}$ , and outputs a secret or error polynomial in  $\mathcal{R}$  sampled according to the Gaussian distribution  $\mathcal{D}_\sigma$ . To realize  $\text{GaussSampler}$ , we propose a simple yet efficient constant-time algorithm. This is described in Section 2.5.4.
- The encoding function  $\text{Enc} : \{0, 1\}^\kappa \rightarrow \{0, \dots, n-1\}^h \times \{-1, 1\}^h$ . This function encodes a  $\kappa$ -bit hash value  $c'$  as a polynomial  $c \in \mathbb{H}_{n,h}$ . The polynomial  $c$  is represented as the two arrays  $\text{pos\_list} \in \{0, \dots, n-1\}^h$  and  $\text{sign\_list} \in \{-1, 1\}^h$  containing the positions and signs of its nonzero coefficients, respectively.
- The sampling function  $\text{ySampler} : \{0, 1\}^\kappa \times \mathbb{Z} \rightarrow \mathcal{R}_{[B]}$  samples a polynomial  $y \in \mathcal{R}_{[B]}$  taking as inputs a  $\kappa$ -bit seed  $\text{rand}$  and a nonce counter  $\in \mathbb{Z}_{>0}$ .
- The hash-based function  $\text{H} : \mathcal{R}_q^k \times \{0, 1\}^{320} \times \{0, 1\}^{320} \rightarrow \{0, 1\}^\kappa$ . This function takes as inputs  $k$  polynomials  $v_1, \dots, v_k \in \mathcal{R}_q$  and first computes  $[v_1]_M, \dots, [v_k]_M$ . The result is then hashed together with the hash  $\text{G}(m)$  for a given message  $m$  and the hash  $\text{G}(t_1, \dots, t_k)$  to a string  $\kappa$  bits long.
- The correctness check function  $\text{checkE}$ , which gets an error polynomial  $e$  as input and rejects it if  $\sum_{k=1}^h \max_k(e)$  is greater than some bound  $L_E$ ; see Algorithm 5. The function  $\text{checkE}$  guarantees the correctness of the signature scheme by ensuring that  $\|e_i c\|_\infty \leq E$  for  $i = 1, \dots, k$  during key generation, as described in Section 2.4.
- The simplification check function  $\text{checkS}$ , which gets a secret polynomial  $s$  as input and rejects it if  $\sum_{k=1}^h \max_k(s)$  is greater than some bound  $L_S$ ; see Algorithm 4.  $\text{checkS}$  ensures that  $\|sc\|_\infty \leq S$ , which is used to simplify the security reduction.

We are now in position to describe  $\text{qTESLA}$ 's algorithms for key generation, signing, and verification, which are depicted in Algorithms 6, 7 and 8, respectively.

**Key generation.** First, the public polynomials  $a_1, \dots, a_k$  are generated uniformly at random over  $\mathcal{R}_q$  (lines 2–4) by expanding the seed  $\text{seed}_a$  using  $\text{PRF}_1$ . Then, a secret polynomial  $s$  is sampled with Gaussian distribution  $\mathcal{D}_\sigma$ . This polynomial must fulfill the requirement check in  $\text{checkS}$  (lines 5–8). A similar procedure to sample the secret error polynomials  $e_1, \dots, e_k$  follows. In this case, these polynomials must fulfill the correctness check in  $\text{checkE}$  (lines 10–13). To generate pseudorandom bit strings during the Gaussian sampling, the corresponding value from  $\{\text{seed}_s, \text{seed}_{e_1}, \dots, \text{seed}_{e_k}\}$  is used as seed, and a counter is used as nonce to provide domain separation between the different calls to the sampler. Accordingly, this counter is initialized at 1 and then increased by 1 after each invocation to the Gaussian sampler. Finally, the public key  $pk$  consists of  $\text{seed}_a$  and the polynomials  $t_i = a_i s + e_i \pmod q$  for  $i = 1, \dots, k$  and the secret key  $sk$  consists of  $s, e_1, \dots, e_k$ , the seeds  $\text{seed}_a$  and  $\text{seed}_y$ , and the

hash  $g = G(t_1, \dots, t_k)$ . All the seeds required during key generation are generated by expanding a pre-seed  $\text{pre-seed}$  using  $\text{PRF}_1$ .

**Signature generation.** To sign a message  $m$ , first a polynomial  $y \in \mathcal{R}_{[B]}$  is chosen uniformly at random (lines 1–4). To this end, a counter initialized at 1 is used as nonce, and a random string  $\text{rand}$  is used as seed. The random string  $\text{rand}$  is computed as  $\text{PRF}_2(\text{seed}_y, r, G(m))$  with  $\text{seed}_y$ , a random string  $r$ , and the digest  $G(m)$  of the message  $m$ . The counter is used to provide domain separation between the different calls to sample  $y$ . Accordingly, it is increased by 1 every time the algorithm restarts if any of the security or correctness tests fail to compute a valid signature (see below). Next,  $\text{seed}_a$  is expanded to generate the polynomials  $a_1, \dots, a_k$  (line 5) which are then used to compute the polynomials  $v_i = a_i y \bmod^{\pm} q$  for  $i = 1, \dots, k$  (lines 6–8). Afterwards, the hash-based function  $H$  computes  $[v_1]_M, \dots, [v_k]_M$  and hashes these together with the digests  $G(m)$  and  $g$  in order to generate  $c'$ . This value is then mapped deterministically to a pseudorandomly generated polynomial  $c \in \mathbb{H}_{n,h}$  which is encoded as the two arrays  $\text{pos\_list} \in \{0, \dots, n-1\}^h$  and  $\text{sign\_list} \in \{-1, 1\}^h$  representing the positions and signs of the nonzero coefficients of  $c$ , respectively. In order for the *potential* signature  $(z \leftarrow sc + y, c')$  at line 11 to be returned by the signing algorithm, it needs to pass a *security* and a *correctness* check, which are described next.

The security check (lines 12–15), also called the *rejection sampling*, is used to ensure that the signature does not leak any information about the secret  $s$ . It is realized by checking that  $z \notin \mathcal{R}_{[B-S]}$ . If the check fails, the algorithm discards the current pair  $(z, c')$  and repeats all the steps beginning with the sampling of  $y$ . Otherwise, the algorithm goes on with the correctness check.

The correctness check (lines 18–21) ensures the correctness of the signature scheme, i.e., it guarantees that every valid signature generated by the signing algorithm is accepted by the verification algorithm. It is realized by checking that  $\|[w_i]_L\|_{\infty} < 2^{d-1} - E$  and  $\|w_i\|_{\infty} < \lfloor q/2 \rfloor - E$ . If the check fails, the algorithm discards the current pair  $(z, c')$  and repeats all the steps beginning with the sampling of  $y$ . Otherwise, the algorithm returns the signature  $(z, c')$  on  $m$ .

**Verification.** The verification algorithm, upon input of a message  $m$ , a signature  $(z, c')$  and a public key  $pk$ , computes  $\{\text{pos\_list}, \text{sign\_list}\} \leftarrow \text{Enc}(c')$ , and then expands  $\text{seed}_a$  to generate  $a_1, \dots, a_k \in \mathcal{R}_q$  and computes  $w_i = a_i z - b_i c$  for  $i = 1, \dots, k$ . The hash-based function  $H$  computes  $[w_1]_M, \dots, [w_k]_M$  and hashes these together with the digests  $G(m)$  and  $G(t_1, \dots, t_k)$ . If the bit string resulting from the previous computation matches the signature bit string  $c'$ , and  $z \in \mathcal{R}_{[B-S]}$ , the signature is accepted; otherwise, it is rejected.

---

**Algorithm 4** checkS: simplifies the security reduction by ensuring that  $\|sc\|_\infty \leq S$ .

---

**Require:**  $s \in \mathcal{R}$

**Ensure:**  $\{0, 1\} \triangleright$  true, false

---

```

1: if  $\sum_{i=1}^h \max_i(s) > L_S$  then
2:   return 1
3: end if
4: return 0

```

---



---

**Algorithm 5** checkE: ensures correctness of the scheme by checking that  $\|ec\|_\infty \leq E$ .

---

**Require:**  $e \in \mathcal{R}$

**Ensure:**  $\{0, 1\} \triangleright$  true, false

---

```

1: if  $\sum_{i=1}^h \max_i(e) > L_E$  then
2:   return 1
3: end if
4: return 0

```

---



---

**Algorithm 6** qTESLA's key generation

---

**Require:** -

**Ensure:** secret key  $sk = (s, e_1, \dots, e_k, \text{seed}_a, \text{seed}_y, g)$ , and public key  $pk = (t_1, \dots, t_k, \text{seed}_a)$

---

```

1: counter  $\leftarrow$  1
2: pre-seed  $\leftarrow_{\S} \{0, 1\}^\kappa$ 
3: seeds, seede1, ..., seedek, seeda, seedy  $\leftarrow$  PRF1(pre-seed) [Algorithm 9]
4: a1, ..., ak  $\leftarrow$  GenA(seeda) [Algorithm 10]
5: do
6:   s  $\leftarrow_{\sigma} \mathcal{R}$  [GaussSampler(seeds, counter), Algorithm 11]
7:   counter  $\leftarrow$  counter + 1
8: while checkS(s)  $\neq$  0 [Algorithm 4]
9: for i = 1, ..., k do
10:  do
11:    ei  $\leftarrow_{\sigma} \mathcal{R}$  [GaussSampler(seedei, counter), Algorithm 11]
12:    counter  $\leftarrow$  counter + 1
13:  while checkE(ei)  $\neq$  0 [Algorithm 5]
14:  ti  $\leftarrow$  ais + ei mod q
15: end for
16: g  $\leftarrow$  G(t1, ..., tk)
17: sk  $\leftarrow$  (s, e1, ..., ek, seeda, seedy, g)
18: pk  $\leftarrow$  (t1, ..., tk, seeda)
19: return sk, pk

```

---

---

**Algorithm 7** qTESLA's signature generation

---

**Require:** message  $m$ , and secret key  $sk = (s, e_1, \dots, e_k, \text{seed}_a, \text{seed}_y, g)$

**Ensure:** signature  $(z, c')$

---

```
1: counter  $\leftarrow 1$ 
2:  $r \leftarrow_{\S} \{0, 1\}^{\kappa}$ 
3:  $\text{rand} \leftarrow \text{PRF}_2(\text{seed}_y, r, G(m))$ 
4:  $y \leftarrow \text{ySampler}(\text{rand}, \text{counter})$  [Algorithm 12]
5:  $a_1, \dots, a_k \leftarrow \text{GenA}(\text{seed}_a)$  [Algorithm 10]
6: for  $i = 1, \dots, k$  do
7:    $v_i = a_i y \bmod^{\pm} q$ 
8: end for
9:  $c' \leftarrow H(v_1, \dots, v_k, G(m), g)$  [Algorithm 13]
10:  $c \triangleq \{\text{pos\_list}, \text{sign\_list}\} \leftarrow \text{Enc}(c')$  [Algorithm 14]
11:  $z \leftarrow y + sc$ 
12: if  $z \notin \mathcal{R}_{[B-S]}$  then
13:   counter  $\leftarrow$  counter + 1
14:   Restart at step 4
15: end if
16: for  $i = 1, \dots, k$  do
17:    $w_i \leftarrow v_i - e_i c \bmod^{\pm} q$ 
18:   if  $\|[w_i]_L\|_{\infty} \geq 2^{d-1} - E \vee \|w_i\|_{\infty} \geq \lfloor q/2 \rfloor - E$  then
19:     counter  $\leftarrow$  counter + 1
20:     Restart at step 4
21:   end if
22: end for
23: return  $(z, c')$ 
```

---

---

**Algorithm 8** qTESLA's signature verification

---

**Require:** message  $m$ , signature  $(z, c')$ , and public key  $pk = (t_1, \dots, t_k, \text{seed}_a)$

**Ensure:**  $\{0, -1\} \triangleright$  accept, reject signature

---

```
1:  $c \triangleq \{\text{pos\_list}, \text{sign\_list}\} \leftarrow \text{Enc}(c')$  [Algorithm 14]
2:  $a_1, \dots, a_k \leftarrow \text{GenA}(\text{seed}_a)$  [Algorithm 10]
3: for  $i = 1, \dots, k$  do
4:    $w_i \leftarrow a_i z - t_i c \bmod^{\pm} q$ 
5: end for
6: if  $z \notin \mathcal{R}_{[B-S]} \vee c' \neq H(w_1, \dots, w_k, G(m), G(t_1, \dots, t_k))$  then
7:   return  $-1$ 
8: end if
9: return  $0$ 
```

---

## 2.4 Correctness of the scheme

In general, a signature scheme consisting of a tuple  $(\text{KeyGen}, \text{Sign}, \text{Verify})$  of algorithms is correct if, for every message  $m$  in the message space  $\mathcal{M}$ , we have that

$$\Pr [\text{Verify}(\text{pk}, m, \sigma) = 0 : (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(), \sigma \leftarrow \text{Sign}(\text{sk}, m) \text{ for } m \in \mathcal{M}] = 1,$$

where the probability is taken over the randomness of the probabilistic algorithms. To prove the correctness of qTESLA, we have to show that for every signature  $(z, c')$  of a message  $m$  generated by Algorithm 7 it holds that (i)  $z \in \mathcal{R}_{[B-S]}$  and (ii) the output of the hash-based function  $\mathbf{H}$  at signing (line 9 of Algorithm 7) is the same as the analogous output at verification (line 6 of Algorithm 8).

Requirement (i) is ensured by the security check during signing (line 12 of Algorithm 7). To ensure (ii), we need to prove that, for genuine signatures and for all  $i = 1, \dots, k$  it holds that  $[a_i y]_M = [a_i z - t_i c]_M = [a_i(y + sc) - (a_i s + e_i)c]_M = [a_i y + a_i sc - a_i sc - e_i c]_M = [a_i y - e_i c]_M$ . From the definition of  $[\cdot]_M$ , this means proving that  $(a_i y \bmod^{\pm} q - [a_i y]_L)/2^d = (a_i y - e_i c \bmod^{\pm} q - [a_i y - e_i c]_L)/2^d$ , or simply  $[a_i y]_L = e_i c + [a_i y - e_i c]_L$ .

The above equality must hold component-wise, so let us prove the corresponding property for individual integers.

Assume that for integers  $\alpha$  and  $\varepsilon$  it holds that  $|[\alpha - \varepsilon]_L| < 2^{d-1} - E$ ,  $|\varepsilon| \leq E < \lfloor q/2 \rfloor$ ,  $|\alpha - \varepsilon \bmod^{\pm} q| < \lfloor q/2 \rfloor - E$ , and  $-\lfloor q/2 \rfloor < \alpha \leq \lfloor q/2 \rfloor$  (i.e.,  $\alpha \bmod^{\pm} q = \alpha$ ). Then, we need to prove that

$$[\alpha]_L = \varepsilon + [\alpha - \varepsilon]_L. \quad (1)$$

*Proof.* To prove Equation (1), start by noticing that  $|\varepsilon| \leq E < 2^{d-1}$  implies  $[\varepsilon]_L = \varepsilon$ . Thus, from  $-2^{d-1} + E < [\alpha - \varepsilon]_L < 2^{d-1} - E$  and  $-E \leq [\varepsilon]_L \leq E$  it follows that

$$-2^{d-1} = -2^{d-1} + E - E < [\varepsilon]_L + [\alpha - \varepsilon]_L < 2^{d-1} - E + E = 2^{d-1},$$

and therefore

$$[[\varepsilon]_L + [\alpha - \varepsilon]_L]_L = [\varepsilon]_L + [\alpha - \varepsilon]_L = \varepsilon + [\alpha - \varepsilon]_L. \quad (2)$$

Next we prove that

$$[[\varepsilon]_L + [\alpha - \varepsilon]_L]_L = [\alpha]_L. \quad (3)$$

Since  $|\varepsilon| \leq E < \lfloor q/2 \rfloor$  and  $|\alpha - \varepsilon \bmod^{\pm} q| < \lfloor q/2 \rfloor$ , it holds further that

$$[[\varepsilon]_L + [\alpha - \varepsilon]_L]_L \quad (4)$$

$$= ((\varepsilon \bmod^{\pm} q) \bmod^{\pm} 2^d + (\alpha - \varepsilon \bmod^{\pm} q) \bmod^{\pm} 2^d) \bmod^{\pm} q \bmod^{\pm} 2^d \quad (5)$$

$$= (\varepsilon \bmod^{\pm} q + (\alpha - \varepsilon \bmod^{\pm} q)) \bmod^{\pm} 2^d. \quad (6)$$

Since  $|\varepsilon| \leq E$  and  $|\alpha - \varepsilon \bmod^\pm q| < \lfloor q/2 \rfloor - E$ , it holds that  $|\alpha - \varepsilon| + |\varepsilon| < (\lfloor q/2 \rfloor - E) + E = \lfloor q/2 \rfloor$ . Hence, Equation (6) is the same as

$$= (\varepsilon + \alpha - \varepsilon \bmod^\pm q) \bmod^\pm 2^d = (\alpha \bmod^\pm q) \bmod^\pm 2^d = [\alpha]_L.$$

By combining Equation (2) and Equation (3) we deduce that  $[\alpha]_L = \varepsilon + [\alpha - \varepsilon]_L$ , which is the equation we needed to prove.  $\square$

Now define  $\alpha := (a_i y)_j$  and  $\varepsilon := (e_i c)_j$  with  $i \in \{1, \dots, k\}$  and  $j \in \{0, \dots, n-1\}$ . From line 18 of Algorithm 7, we know that for  $i = 1, \dots, k$ ,  $\|[a_i y - e_i c]_L\|_\infty < 2^{d-1} - E$  and  $\|a_i y - e_i c\|_\infty < \lfloor q/2 \rfloor - E$  for an honestly generated signature, and that Algorithm 6 (line 13) guarantees  $\|e_i c\|_\infty \leq E$ . Likewise, by definition it holds that  $E < \lfloor q/2 \rfloor$ ; see Section 5. Finally,  $v_i = a_i y$  is reduced  $\bmod^\pm q$  in line 7 of Algorithm 7 and, hence,  $v_i$  is in the centered range  $-\lfloor q/2 \rfloor < a_i y \leq \lfloor q/2 \rfloor$ .

In conclusion, we get the desired condition for ring elements,  $[a_i y]_L = e_i c + [a_i y - e_i c]_L$ , which in turn means  $[a_i z - t_i c]_M = [a_i y]_M$  for  $i = 1, \dots, k$ .

## 2.5 Realization of the required functions

### 2.5.1 Hash and pseudorandom functions

In addition to the hash-based functions G and H and the pseudorandom functions PRF<sub>1</sub> and PRF<sub>2</sub>, several functions that are used for the implementation of qTESLA require pseudorandom bit generation. This functionality is provided by so-called extendable output functions (XOF).

For the remainder, the format that we use to call a XOF is given by XOF(X, L, D), where X is the input string, L specifies the output length in bytes, and D specifies an optional domain separator<sup>1</sup>.

Next, we summarize how XOFs are instantiated using SHAKE [29] and cSHAKE [39] in the different functions requiring hashing or pseudorandom bit generation.

- PRF<sub>1</sub>: the XOF is instantiated with SHAKE128 (resp. SHAKE256) for the Level-I parameter set (resp. for the Level-III parameter set); see Algorithm 9.
- PRF<sub>2</sub>: the same as PRF<sub>1</sub>.
- GenA: the XOF is instantiated with cSHAKE128 (see Algorithm 10).
- GaussSampler: the XOF is instantiated with cSHAKE128 (resp. cSHAKE256) for the Level-I parameter set (resp. for the Level-III parameter set); see Algorithm 11.

---

<sup>1</sup>The domain separator D is used with cSHAKE, but ignored when SHAKE is used.



- **Enc**: the XOF is instantiated with cSHAKE128 (see Algorithm 14).
- **ySampler**: the XOF is instantiated with cSHAKE128 (resp. cSHAKE256) for the Level-I parameter set (resp. for the Level-III parameter set); see Algorithm 12.
- **Hash G**: this function is instantiated with SHAKE128 (resp. SHAKE256) for the Level-I parameter set (resp. for the Level-III parameter set).
- **Hash-based function H**: the hashing in this function is instantiated with SHAKE128 (resp. SHAKE256) for the Level-I parameter set (resp. for the Level-III parameter set); see Algorithm 13.

In the cases of the functions GenA, Enc, G and H, implementations of qTESLA need to follow strictly the XOF specifications based on SHAKE/cSHAKE given above in order to be specification compliant. However, for the rest of the cases (i.e., PRF<sub>1</sub>, PRF<sub>2</sub>, ySampler and GaussSampler) users can opt for a different cryptographic PRF.

### 2.5.2 Pseudorandom bit generation of seeds, PRF<sub>1</sub>

qTESLA requires the generation of seeds during key generation; see line 3 of Algorithm 6. These seeds are then used to produce the polynomials  $s$ ,  $e_i$ ,  $a_i$  and  $y$ . Specifically, these seeds are:

- $\text{seed}_s$ , which is used to generate the polynomial  $s$ ,
- $\text{seed}_{e_i}$ , which are used to generate the polynomials  $e_i$  for  $i = 1, \dots, k$ ,
- $\text{seed}_a$ , which is used to generate the polynomials  $a_i$  for  $i = 1, \dots, k$ , and
- $\text{seed}_y$ , which is used to generate the polynomial  $y$ .

The size of each of these seeds is  $\kappa$  bits. In the accompanying implementations, the seeds are generated by first calling the system random number generator (RNG) to produce a pre-seed of size  $\kappa$  bits at line 2 of Algorithm 6, and then expanding this pre-seed through Algorithm 9. As explained in Section 2.5.1, in this case the XOF function is instantiated with SHAKE in our implementations.

---

#### Algorithm 9 Seed generation, PRF<sub>1</sub>

---

**Require:**  $\text{pre-seed} \in \{0, 1\}^\kappa$

**Ensure:**  $(\text{seed}_s, \text{seed}_{e_1}, \dots, \text{seed}_{e_k}, \text{seed}_a)$ , where each seed is  $\kappa$  bits long

---

1:  $\langle \text{seed}_s \rangle \| \langle \text{seed}_{e_1} \rangle \| \dots \| \langle \text{seed}_{e_k} \rangle \| \langle \text{seed}_a \rangle \| \langle \text{seed}_y \rangle \leftarrow \text{XOF}(\text{pre-seed}, \kappa \cdot (k + 3)/8)$ , where each  $\langle \text{seed} \rangle \in \{0, 1\}^\kappa$

2: **return**  $(\text{seed}_s, \text{seed}_{e_1}, \dots, \text{seed}_{e_k}, \text{seed}_a)$

---

### 2.5.3 Generation of $\mathbf{a}_1, \dots, \mathbf{a}_k$

In qTESLA, the polynomials  $a_1, \dots, a_k$  are freshly generated per secret/public keypair using the seed  $\text{seed}_a$  during key generation; see line 4 of Algorithm 6. This seed is then stored as part of both the private and public keys so that the signing and verification operations can regenerate  $a_1, \dots, a_k$ .

The approach above permits to save bandwidth since we only need  $\kappa$  bits to store  $\text{seed}_a$  instead of the  $k \cdot n \cdot \lceil \log_2 q \rceil$  bits that are required to represent the full polynomials. Moreover, the use of fresh  $a_1, \dots, a_k$  per keypair makes the introduction of backdoors more difficult and reduces drastically the scope of all-for-the-price-of-one attacks [9, 14].

The procedure depicted in Algorithm 10 to generate  $a_1, \dots, a_k$  is as follows. The seed  $\text{seed}_a$  obtained from Algorithm 9 is expanded to  $(\text{rate}_{\text{XOF}} \cdot b_{\text{GenA}})$  bytes using cSHAKE128, where  $\text{rate}_{\text{XOF}}$  is the SHAKE128 rate constant 168 [29] and  $b_{\text{GenA}}$  is a qTESLA parameter that represents the number of blocks requested in the first XOF call. Then, the algorithm proceeds to do rejection sampling over each  $8 \lceil \log_2 q \rceil$ -bit string of the cSHAKE output modulo  $2^{\lceil \log_2 q \rceil}$ , discarding every package that has a value equal or greater than the modulus  $q$ . Since there is a possibility that the cSHAKE output is exhausted before all the  $k \cdot n$  coefficients are filled out, the algorithm permits successive (and as many as necessary) calls to the function requesting  $\text{rate}_{\text{XOF}}$  bytes each time (lines 5–8). The first call to cSHAKE128 uses the value  $D = 0$  as domain separator. This value is incremented by one at each subsequent call.

The procedure above produces polynomials with uniformly random coefficients. Thus, following a standard practice, qTESLA assumes that the resulting polynomials  $a_1, \dots, a_k$  from Algorithm 10 are already in the NTT domain. This permits an important speedup of the polynomial operations without affecting security. We remark, however, that this assumption does affect the correctness and, hence, implementations should follow this design feature to be specification compliant.

Refer to Section 2.5.8 for details about the NTT computations.

---

**Algorithm 10** Generation of public polynomials  $a_i$ , GenA

---

**Require:**  $\text{seed}_a \in \{0, 1\}^\kappa$ . Set  $b = \lceil (\log_2 q)/8 \rceil$  and the SHAKE128 rate constant  $\text{rate}_{\text{XOF}} = 168$

**Ensure:**  $a_i \in \mathcal{R}_q$  for  $i = 1, \dots, k$

---

```
1:  $D \leftarrow 0, b' \leftarrow b_{\text{GenA}}$ 
2:  $\langle c_0 \rangle \| \langle c_1 \rangle \| \dots \| \langle c_T \rangle \leftarrow \text{cSHAKE128}(\text{seed}_a, \text{rate}_{\text{XOF}} \cdot b', D)$ , where each  $\langle c_t \rangle \in \{0, 1\}^{8b}$ 
3:  $i \leftarrow 0, \text{pos} \leftarrow 0$ 
4: while  $i < k \cdot n$  do
5:   if  $\text{pos} > \lfloor (\text{rate}_{\text{XOF}} \cdot b')/b \rfloor - 1$  then
6:      $D \leftarrow D + 1, \text{pos} \leftarrow 0, b' \leftarrow 1$ 
7:      $\langle c_0 \rangle \| \langle c_1 \rangle \| \dots \| \langle c_T \rangle \leftarrow \text{cSHAKE128}(\text{seed}_a, \text{rate}_{\text{XOF}} \cdot b', D)$ , where each  $\langle c_t \rangle \in \{0, 1\}^{8b}$ 
8:   end if
9:   if  $c_{\text{pos}} \bmod 2^{\lceil \log_2 q \rceil} < q$  then
10:     $a_{\lfloor i/n \rfloor + 1, i - n \cdot \lfloor i/n \rfloor} \leftarrow c_{\text{pos}} \bmod 2^{\lceil \log_2 q \rceil}$ , where a polynomial  $a_x$  is interpreted as a vector
    of coefficients  $(a_{x,0}, a_{x,1}, \dots, a_{x,n-1})$ 
11:     $i \leftarrow i + 1$ 
12:   end if
13:    $\text{pos} \leftarrow \text{pos} + 1$ 
14: end while
15: return  $(a_1, \dots, a_k)$ 
```

---

### 2.5.4 Gaussian sampling

One of the advantages of qTESLA is that discrete Gaussian sampling is only required during key generation to sample  $e_1, \dots, e_k$ , and  $s$  (see Alg. 6). Nevertheless, certain applications might still require an efficient and secure implementation of key generation and one that is, in particular, portable and protected against timing and cache side-channel attacks. Accordingly, we employ a *constant-time* discrete Gaussian sampler based on the well-established technique of cumulative distribution table (CDT) of the normal distribution, which consists of precomputing, to a given  $\beta$ -bit precision, a table  $\text{CDT}[i] := \lfloor 2^\beta \Pr[c \leq i \mid c \leftarrow_\sigma \mathbb{Z}] \rfloor$ , for  $i \in [-t + 1 \dots t - 1]$  with the smallest  $t$  such that  $\Pr[|c| \geq t \mid c \leftarrow_\sigma \mathbb{Z}] < 2^{-\beta}$ . To obtain a Gaussian sample, one picks a uniform sample  $u \leftarrow_\$ \mathbb{Z}/2^\beta \mathbb{Z}$ , looks it up in the table, and returns the value  $z$  such that  $\text{CDT}[z] \leq u < \text{CDT}[z + 1]$ .

A CDT-based approach has apparently first been considered for cryptographic purposes by Peikert [50] (in a somewhat more complex form). The approach was assessed and deemed mostly impractical by Ducas *et al.* [27], since it would take  $\beta t \sigma$  bits. Yet, they only considered a scenario where the standard deviation  $\sigma$  was at least 107, and as high as 271. As a result, table sizes around 78 Kbytes are reported (presumably for  $\sigma = 271$  with roughly 160-bit sampling precision). For the qTESLA parameter sets, however, the values of  $\sigma$  are much smaller, making the CDT approach feasible, as one can see in Table 1.

Table 1: CDT dimensions used in the accompanying qTESLA implementations (targeted precision  $\beta$  : implemented precision in bits : number of rows  $t$  : table size in bytes).

Parameter set	CDT parameters
qTESLA-p-I	64 : 63 : 78 : 624
qTESLA-p-III	128 : 125 : 111 : 1776

**Implementation details.** To implement the CDT-based Gaussian sampler in our implementations, we use an optimized constant-time version of Algorithm 11. This algorithm generates  $n$  Gaussian samples, doing a chunk of  $c \mid n$  samples at a time. The chunk size is fixed to  $c = 512$  when the dimension  $n$  is a multiple of 512. For the required pre-computed CDT tables, the targeted sampling precision  $\beta$  is conservatively set to a value much greater than  $\lambda/2$ , as can be seen in Table 1. The CDT tables, which only contain the right-hand-sided (i.e., positive) values of the cumulative distributed function to save memory, are generated using the script provided in the folder `\Supporting_Documentation\Script_for_Gaussian_sampler`. Note that the most significant bit in each row of the table is always set to 0. This is done to facilitate the sign generation that is used to recover the full distribution during sampling (see line 7 of Alg. 11). In addition, each column’s most significant bit is set to 0 to facilitate the multiprecision subtractions that are required to implement the comparison in line 10. Hence, in the implementation of the algorithm comparisons should take care of such top bits in this representation, as remarked in line 8.

In order to make the Gaussian sampler constant-time we make sure that basic operations such as comparisons are not implemented with conditional jumps that depend on secret data, and that lookup tables are always fully scanned at each pass generating samples with constant-time logical and arithmetic operations.

As stated in Section 2.5.1, for the pseudorandom bit generation required by Algorithm 11, we use cSHAKE as XOF using a seed `seed` produced by PRF<sub>1</sub> (see line 3 of Algorithm 6) as input string, and a nonce  $D$  (written as `counter` in Algorithm 6) as domain separator.

### 2.5.5 Sampling of $y$

The sampling of the polynomial  $y$  (line 4 of Algorithm 7) can be performed by generating  $n$  ( $\lceil \log_2 B \rceil + 1$ )-bit values uniformly at random, and then correcting each value to the range  $[-B, B + 1]$  with a subtraction by  $B$ . Since values need to be in the range  $[-B, B]$ , coefficients with value  $B + 1$  need to be rejected, which in turn might require the generation of additional random bits.

Algorithm 12 depicts the procedure used in our implementations. For the pseudorandom

---

**Algorithm 11** Constant-time CDT-based Gaussian sampling, GaussSampler

---

INPUT: seed  $\text{seed} \in \{0, 1\}^\kappa$  and nonce  $D \in \mathbb{Z}_{>0}$ .

OUTPUT: a sequence  $z$  of  $n$  Gaussian samples.

GLOBAL: dimension  $n$ ;  $\text{cdt.v}$ : the  $t$ -entry right-hand-sided,  $\beta$ -bit precision CDT;  $c$ : chunk size, s.t.  $c \mid n$ .

---

```
1:  $D' \leftarrow D \cdot 2^8$ 
2: for  $0 \leq i < n$  do
3:    $\langle r_0 \rangle \| \langle r_1 \rangle \| \dots \| \langle r_{c-1} \rangle \leftarrow \text{XOF}(\text{seed}, c \cdot \beta / 8, D')$ , where each  $\langle r_i \rangle \in \{0, 1\}^\beta$ 
4:    $D' \leftarrow D' + 1$ 
5:   for  $0 \leq j < c$  do
6:      $z_{i+j} \leftarrow 0$ 
7:      $\text{sign} \leftarrow r_j / 2^{\beta-1}$ 
8:      $\text{val} \leftarrow r_j$  after removing the bits corresponding to the top bits of the CDT
       columns
9:     for  $0 \leq k < t$  do
10:      if  $\text{val} \geq \text{cdt.v}_k$  then
11:         $z_{i+j} \leftarrow z_{i+j} + 1$ 
12:      end if
13:    end for
14:    if  $\text{sign} = 1$  then
15:       $z_{i+j} \leftarrow -z_{i+j}$ 
16:    end if
17:  end for
18:   $i \leftarrow i + c$ 
19: end for
20: return  $z$ 
```

---

bit generation, the seed  $\text{rand}$  produced by  $\text{PRF}_2$  (see line 3 of Algorithm 7) is used as input string to a XOF, while the nonce  $D$  (written as counter in Algorithm 7) is used for the computation of the values for the domain separation. The first call to the XOF function uses the value  $D' \leftarrow D \cdot 2^8$  as domain separator. Each subsequent call to the XOF increases  $D'$  by 1. Since  $D$  is initialized at 1 by the signing algorithm, and then increased by 1 at each subsequent call to sample  $y$ , the successive calls to the sampler use nonces  $D'$  initialized at  $2^8, 2 \cdot 2^8, 3 \cdot 2^8$ , and so on, providing proper domain separation between the different uses of the XOF in the signing algorithm.

Our implementations use cSHAKE as the XOF function.

---

### Algorithm 12 Sampling $y$ , ySampler

---

**Require:** seed  $\text{rand} \in \{0, 1\}^\kappa$  and nonce  $D \in \mathbb{Z}_{>0}$ . Set  $b = \lceil (\log_2 B + 1)/8 \rceil$

**Ensure:**  $y \in \mathcal{R}_{[B]}$

---

```

1:  $pos \leftarrow 0, n' \leftarrow n, D' \leftarrow D \cdot 2^8$ 
2:  $\langle c_0 \rangle \| \langle c_1 \rangle \| \dots \| \langle c_{n'-1} \rangle \leftarrow \text{XOF}(\text{rand}, b \cdot n', D')$ , where each  $\langle c_i \rangle \in \{0, 1\}^{8b}$ 
3: while  $i < n$  do
4:   if  $pos \geq n'$  then
5:      $D' \leftarrow D' + 1, pos \leftarrow 0, n' \leftarrow \lfloor \text{rate}_{\text{XOF}}/b \rfloor$ 
6:      $\langle c_0 \rangle \| \langle c_1 \rangle \| \dots \| \langle c_{n'-1} \rangle \leftarrow \text{XOF}(\text{rand}, \text{rate}_{\text{XOF}}, D')$ , where each  $\langle c_i \rangle \in \{0, 1\}^{8b}$ 
7:   end if
8:    $y_i \leftarrow c_{pos} \bmod 2^{\lceil \log_2 B \rceil + 1} - B$ 
9:   if  $y_i \neq B + 1$  then
10:     $i \leftarrow i + 1$ 
11:   end if
12:    $pos \leftarrow pos + 1$ 
13: end while
14: return  $y = (y_0, y_1, \dots, y_{n-1}) \in \mathcal{R}_{[B]}$ 

```

---

### 2.5.6 Hash-based function H

This function takes as inputs  $k$  polynomials  $v_1, \dots, v_k$  in  $\mathcal{R}_q$  and computes  $[v_1]_M, \dots, [v_k]_M$ . The result is hashed together with the hash  $G$  of a message  $m$  and the hash  $G(t_1, \dots, t_k)$  to a string  $c'$  that is  $\kappa$  bits long. The detailed procedure is as follows. Let each polynomial  $v_i$  be interpreted as a vector of coefficients  $(v_{i,0}, v_{i,1}, \dots, v_{i,n-1})$ , where  $v_{i,j} \in (-q/2, q/2]$ , i.e.,  $v_{i,j} = v_{i,j} \bmod^{\pm} q$ . We first compute  $[v_i]_L$  by reducing each coefficient modulo  $2^d$  and decreasing the result by  $2^d$  if it is greater than  $2^{d-1}$ . This guarantees a result in the range  $(-2^{d-1}, 2^{d-1}]$ , as required by the definition of  $[\cdot]_L$ . Next, we compute  $[v_i]_M$  as  $(v_i \bmod^{\pm} q - [v_i]_L)/2^d$ . Since each resulting coefficient is guaranteed to be very small it is stored as a byte, which in total makes up a string of  $k \cdot n$  bytes. Finally, SHAKE is used to hash the resulting  $(k \cdot n)$ -byte string together with the 40-byte digests  $G(m)$  and  $G(t_1, \dots, t_k)$  to the  $\kappa$ -bit string  $c'$ . This procedure is depicted in Algorithm 13.

---

**Algorithm 13** Hash-based function H

---

**Require:** polynomials  $v_1, \dots, v_k \in \mathcal{R}_q$ , where  $v_{i,j} \in (-q/2, q/2]$ , for  $i = 1, \dots, k$  and  $j = 0, \dots, n-1$ , and the digests  $G(m)$  and  $G(t_1, \dots, t_k)$ , each of length 40 bytes.

**Ensure:**  $c' \in \{0, 1\}^\kappa$

---

```
1: for  $i = 1, 2, \dots, k$  do
2:   for  $j = 0, 1, \dots, n-1$  do
3:      $\text{val} \leftarrow v_{i,j} \bmod 2^d$ 
4:     if  $\text{val} > 2^{d-1}$  then
5:        $\text{val} \leftarrow \text{val} - 2^d$ 
6:     end if
7:      $w_{(i-1) \cdot n + j} \leftarrow (v_{i,j} - \text{val}) / 2^d$ 
8:   end for
9: end for
10:  $\langle w_{k \cdot n} \rangle \| \langle w_{k \cdot n + 1} \rangle \| \dots \| \langle w_{k \cdot n + 39} \rangle \leftarrow G(m)$ , where each  $\langle w_i \rangle \in \{0, 1\}^8$ 
11:  $\langle w_{k \cdot n + 40} \rangle \| \langle w_{k \cdot n + 41} \rangle \| \dots \| \langle w_{k \cdot n + 79} \rangle \leftarrow G(t_1, \dots, t_k)$ , where each  $\langle w_i \rangle \in \{0, 1\}^8$ 
12:  $c' \leftarrow \text{SHAKE}(w, \kappa/8)$ , where  $w$  is the byte array  $\langle w_0 \rangle \| \langle w_1 \rangle \| \dots \| \langle w_{k \cdot n + 79} \rangle$ 
13: return  $c' \in \{0, 1\}^\kappa$ 
```

---

### 2.5.7 Encoding function

This function maps the bit string  $c'$  to a polynomial  $c \in \mathbb{H}_{n,h} \subset \mathcal{R}$  of degree  $n-1$  with coefficients in  $\{-1, 0, 1\}$  and weight  $h$ , i.e.,  $c$  has  $h$  coefficients that are either 1 or  $-1$ . For efficiency,  $c$  is encoded as two arrays *pos\_list* and *sign\_list* that contain the positions and signs of its nonzero coefficients, respectively.

For the implementation of the encoding function **Enc** we follow [1, 27]. Basically, the idea is to use a XOF to generate values uniformly at random that are interpreted as the positions and signs of the  $h$  nonzero entries of  $c$ . The outputs are stored as entries to the two arrays *pos\_list* and *sign\_list*.

The pseudocode of our implementation of this function is depicted in Algorithm 14. This works as follows. The algorithm first requests  $\text{rate}_{\text{XOF}}$  bytes from a XOF, and the output stream is interpreted as an array of 3-byte packets in little endian format. Each 3-byte packet is then processed as follows, beginning with the least significant packet. The  $\lceil \log_2 n \rceil$  least significant bits of the lowest two bytes in every packet are interpreted as an integer value in little endian representing the position *pos* of a nonzero coefficient of  $c$ . If such value already exists in the *pos\_list* array, the 3-byte packet is rejected and the next packet in line is processed; otherwise, the packet is accepted, the value is added to *pos\_list* as the position of a new coefficient, and then the third byte is used to determine the coefficient's sign as follows. If the least significant bit of the third byte is 0, the coefficient is assumed to be positive (+1), otherwise, it is taken as negative ( $-1$ ). In our implementations, *sign\_list* encodes positive and negative coefficients as 0's and 1's, respectively.

The procedure above is executed until  $pos\_list$  and  $sign\_list$  are filled out with  $h$  entries each. If the XOF output is exhausted before completing the task then additional calls are invoked, requesting  $rate_{\text{XOF}}$  bytes each time. qTESLA uses cSHAKE128 as the XOF function, with the value  $D = 0$  as domain separator for the first call.  $D$  is incremented by one at each subsequent call.

---

**Algorithm 14** Encoding function, Enc

---

**Require:**  $c' \in \{0, 1\}^\kappa$

**Ensure:** arrays  $pos\_list \in \{0, \dots, n-1\}^h$  and  $sign\_list \in \{-1, 1\}^h$  containing the positions and signs, resp., of the nonzero elements of  $c \in \mathbb{H}_{n,h}$

---

```

1:  $D \leftarrow 0, cnt \leftarrow 0$ 
2:  $\langle r_0 \rangle \| \langle r_1 \rangle \| \dots \| \langle r_T \rangle \leftarrow \text{cSHAKE128}(c', rate_{\text{XOF}}, D)$ , where each  $\langle r_t \rangle \in \{0, 1\}^8$ 
3:  $i \leftarrow 0$ 
4: Set all coefficients of  $c$  to 0
5: while  $i < h$  do
6:   if  $cnt > (rate_{\text{XOF}} - 3)$  then
7:      $D \leftarrow D + 1, cnt \leftarrow 0$ 
8:      $\langle r_0 \rangle \| \langle r_1 \rangle \| \dots \| \langle r_T \rangle \leftarrow \text{cSHAKE128}(c', rate_{\text{XOF}}, D)$ , where each  $\langle r_t \rangle \in \{0, 1\}^8$ 
9:   end if
10:   $pos \leftarrow (r_{cnt} \cdot 2^8 + r_{cnt+1}) \bmod n$ 
11:  if  $c_{pos} = 0$  then
12:    if  $r_{cnt+2} \bmod 2 = 1$  then
13:       $c_{pos} \leftarrow -1$ 
14:    else
15:       $c_{pos} \leftarrow 1$ 
16:    end if
17:     $pos\_list_i \leftarrow pos$ 
18:     $sign\_list_i \leftarrow c_{pos}$ 
19:     $i \leftarrow i + 1$ 
20:  end if
21:   $cnt \leftarrow cnt + 3$ 
22: end while
23: return  $\{pos\_list_0, \dots, pos\_list_{h-1}\}$  and  $\{sign\_list_0, \dots, sign\_list_{h-1}\}$ 

```

---

### 2.5.8 Polynomial multiplication and the number theoretic transform

Polynomial multiplication over a finite field is one of the fundamental operations in R-LWE based schemes such as qTESLA. In this setting, this operation can be efficiently carried out by satisfying the condition  $q \equiv 1 \pmod{2n}$  and, thus, enabling the use of the Number



Theoretic Transform (NTT).

Since qTESLA specifies the generation of the polynomials  $a_1, \dots, a_k$  directly in the NTT domain for efficiency purposes (see Section 2.5.3), we need to define polynomials in such a domain. For the remainder of this section we limit the discussion to the standard case of a power-of-two NTT.

Let  $\omega$  be a primitive  $n$ -th root of unity in  $\mathbb{Z}_q$ , i.e.,  $\omega^n \equiv 1 \pmod{q}$ , and let  $\phi$  be a primitive  $2n$ -th root of unity in  $\mathbb{Z}_q$  such that  $\phi^2 = \omega$ . Then, given a polynomial  $a = \sum_{i=0}^{n-1} a_i x^i$  the forward transform is defined as

$$\text{NTT} : \mathbb{Z}_q[x]/\langle x^n + 1 \rangle \rightarrow \mathbb{Z}_q^n, \quad a \mapsto \tilde{a} = \left( \sum_{j=0}^{n-1} a_j \phi^j \omega^{ij} \right)_{i=0, \dots, n-1},$$

where  $\tilde{a} = \text{NTT}(a)$  is said to be in *NTT domain*. Similarly, the inverse transformation of the vector  $\tilde{a}$  in the NTT domain is defined as

$$\text{NTT}^{-1} : \mathbb{Z}_q^n \rightarrow \mathbb{Z}_q[x]/\langle x^n + 1 \rangle, \quad \tilde{a} \mapsto a = \sum_{i=0}^{n-1} \left( n^{-1} \phi^{-i} \sum_{j=0}^{n-1} \tilde{a}_j \omega^{-ij} \right) x^i.$$

It then holds that  $\text{NTT}^{-1}(\text{NTT}(a)) = a$  for all polynomials  $a \in \mathcal{R}_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$ . The polynomial multiplication of  $a$  and  $b \in \mathcal{R}_q$  can be performed as  $a \cdot b = \text{NTT}^{-1}(\text{NTT}(a) \circ \text{NTT}(b))$ , where  $\cdot$  is the polynomial multiplication in  $\mathcal{R}_q$  and  $\circ$  is the coefficient wise multiplication in  $\mathbb{Z}_q^n$ .

As mentioned earlier, the outputs  $a_1, \dots, a_k$  of GenA are assumed to be in the NTT domain. In particular, let  $\tilde{a}_i$  be the output  $a_i$  in the NTT domain. Polynomial multiplications  $a_i \cdot b$  for some polynomial  $b \in \mathcal{R}_q$  can be efficiently realized as  $\text{NTT}^{-1}(\tilde{a}_i \circ \text{NTT}(b))$ .

To compute the NTT in our implementations, we adopt butterfly algorithms to compute the NTT that efficiently merge the powers of  $\phi$  and  $\phi^{-1}$  with the powers of  $\omega$ , and that at the same time avoid the need for a so-called bit-reversal operation which is required by some implementations [9, 54, 55]. Specifically, we use an algorithm that computes the forward NTT based on the Cooley-Tukey butterfly that absorbs the products of the root powers in bit-reversed ordering. This algorithm receives the inputs of a polynomial  $a$  in standard ordering and produces a result in bit-reversed ordering. Similarly, for the inverse NTT we use an algorithm based on the Gentleman-Sande butterfly that absorbs the inverses of the products of the root powers in bit-reversed ordering. The algorithm receives the inputs of a polynomial  $\tilde{a}$  in bit-reversed ordering and produces an output in standard ordering. Polished versions of these algorithms, which we follow for our implementations, can be found in [56, Algorithms 1 and 2].

**Sparse multiplication.** While standard polynomial multiplications can be efficiently carried out using the NTT as explained above, *sparse multiplications* with a polynomial  $c \in \mathbb{H}_{n,h}$ , which only contain  $h$  nonzero coefficients in  $\{-1, 1\}$ , can be realized more efficiently with a specialized algorithm that exploits the sparseness of the input. In our implementations we use Algorithm 15 to realize the multiplications in lines 11 and 17 of Algorithm 7 and in line 4 of Algorithm 8, which have as inputs a given polynomial  $g \in \mathcal{R}_q$  and a polynomial  $c \in \mathbb{H}_{n,h}$  encoded as the position and sign arrays  $pos\_list$  and  $sign\_list$  (as output by Enc, Algorithm 14).

---

**Algorithm 15** Sparse Polynomial Multiplication

---

**Require:**  $g = \sum_{i=0}^{n-1} g_i x^i \in \mathcal{R}_q$  with  $g_i \in \mathbb{Z}_q$ , and list arrays  $pos\_list \in \{0, \dots, n-1\}^h$  and  $sign\_list \in \{-1, 1\}^h$  containing the positions and signs, resp., of the nonzero elements of a polynomial  $c \in \mathbb{H}_{n,h}$

**Ensure:**  $f = g \cdot c \in \mathcal{R}_q$

---

```

1: Set all coefficients of  $f$  to 0
2: for  $i = 0, \dots, h - 1$  do
3:    $pos \leftarrow pos\_list_i$ 
4:   for  $j = 0, \dots, pos - 1$  do
5:      $f_j \leftarrow f_j - sign\_list_i \cdot g_{j+n-pos}$ 
6:   end for
7:   for  $j = pos, \dots, n - 1$  do
8:      $f_j \leftarrow f_j + sign\_list_i \cdot g_{j-pos}$ 
9:   end for
10: end for
11: return  $f$ 

```

---

## 2.6 System parameters and parameter selection

In this section, we describe qTESLA’s system parameters and our proposed parameter sets.

**Parameter sets.** Herein, we propose *two* parameter sets which were derived following a “provably-secure” parameter generation according to a security reduction. The proposed parameter sets are displayed in Table 2 together with their targeted security category, as defined by NIST in [48].

Our proposed parameter sets, namely qTESLA-p-I and qTESLA-p-III, were chosen according to the security reduction provided in Theorem 2, Section 5.1. This implies the following: by virtue of our security reduction, these parameters strictly correspond to an instance of

Table 2: Parameter sets and their targeted security.

qTESLA-p-I	NIST's category 1
qTESLA-p-III	NIST's category 3

the R-LWE problem. That is, the reduction provably guarantees that our scheme has the selected security level as long as the corresponding R-LWE instance is intractable. In other words, hardness statements for R-LWE instances have a provable consequence for the security levels of our scheme. Moreover, since the presented reduction is tight, the tightness gap of our reduction is equal to 1 for our choice of parameters and, hence, the concrete bit security of our signature scheme is essentially the same as the bit hardness of the underlying R-LWE instance.

Choosing parameters following the security statements, as described above, implies to follow specific security requirements and to take a reduction loss into account. This affects the performance and signature/key sizes of the scheme.

The sage script that was used to generate the various parameters is included in the submission package (see the file `parameterchoice.sage` found in the submission folder `\Supporting_Documentation\Script_to_choose_parameters`).

**System parameters.** qTESLA's system parameters and their corresponding bounds are summarized in Table 3. Concrete parameter values for each of the proposed parameter sets are compiled in Table 4.

Let  $\lambda$  be the security parameter, i.e., the targeted bit security of a given instantiation. In the standard R-LWE setting, we have  $\mathcal{R}_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$ , where the dimension  $n$  is a power of two, i.e.,  $n = 2^\ell$  for  $\ell \in \mathbb{N}$ . Let  $\sigma$  be the standard deviation of the centered discrete Gaussian distribution that is used to sample the coefficients of the secret and error polynomials. Let  $k \in \mathbb{Z}_{>0}$  be the number of public polynomials  $a_1, \dots, a_k$ . This also corresponds to the number of R-LWE samples. Choosing a larger/smaller  $k$  allows for efficiency trade-offs, e.g., in the selection of the size of the modulus  $q$  or the dimension  $n$ . Depending on the specific function, the parameter  $\kappa$  defines the input and/or output lengths of the hash-based and pseudorandom functions. This parameter is specified to be larger or equal to the security level  $\lambda$ . This is consistent with the use of the hash in a Fiat-Shamir style signature scheme such as qTESLA, for which preimage resistance is relevant while collision resistance is much less. Accordingly, we take the hash size to be enough to resist preimage attacks. The parameter  $h$  defines the number of nonzero elements in the output of the encoding function described in Section 2.5.7.

The parameter  $b_{\text{GenA}} \in \mathbb{Z}_{>0}$  represents the number of blocks requested in the first call to cSHAKE128 during the generation of the public polynomials  $a_1, \dots, a_k$  (see Algorithm 10).

Table 3: Description and bounds of all the system parameters.

Param.	Description	Requirement
$\lambda$	security parameter	-
$q_h, q_s$	number of hash and sign queries	-
$n$	dimension	power-of-two
$\sigma$	standard deviation of centered discrete Gaussian distribution	-
$k$	# public polynomials $a_1, \dots, a_k$	-
$q$	modulus	$q \equiv 1 \pmod{2n}, q > 2^{d+1} + 1$ $q^{nk} \geq  \Delta\mathbb{S}  \cdot  \Delta\mathbb{L}  \cdot  \Delta\mathbb{H} ,$ $q^{nk} \geq 2^{4\lambda+nkd} 4q_s^3 (q_s + q_h)^2$
$h$	# of nonzero entries of output elements of Enc	$2^h \cdot \binom{n}{h} \geq 2^{2\lambda}$
$\kappa$	output length of hash-based function H and input length of GenA, PRF <sub>1</sub> , PRF <sub>2</sub> , Enc and ySampler	$\kappa \geq \lambda$
$L_E, \eta_E$	bound in checkE	$\lceil \eta_E \cdot h \cdot \sigma \rceil$
$L_S, \eta_S$	bound in checkS	$\lceil \eta_S \cdot h \cdot \sigma \rceil$
$S, E$	rejection parameters	$= L_S, L_E$
$B$	determines interval the randomness is chosen from during signing	near a power-of-two, $B \geq \frac{\sqrt[3]{M}+2S-1}{2(1-\sqrt[3]{M})}$
$d$	number of rounded bits	$d > \log_2(B), d \geq \log_2\left(\frac{2E+1}{1-M^{\frac{1}{nk}}}\right)$
$b_{\text{GenA}}$	number of blocks requested to SHAKE128 for GenA	$b_{\text{GenA}} \in \mathbb{Z}_{>0}$
$ \Delta\mathbb{H} $		$\sum_{j=0}^h \sum_{i=0}^{h-j} \binom{kn}{2i} 2^{2i} \binom{kn-2i}{j} 2^j$
$ \Delta\mathbb{S} $	see definition in the text	$(4(B-S)+1)^n$
$ \Delta\mathbb{L} $		$(2^d+1)^{nk}$
$\delta_z$	acceptance probability of $z$ in line 12 during signing	determined experimentally
$\delta_w$	acceptance probability of $w$ in line 18 during signing	determined experimentally
$\delta_{\text{keygen}}$	acceptance probability of key pairs during key generation	determined experimentally
sig size	theoretical size of signature [bits]	$\kappa + n(\lceil \log_2(B-S) \rceil + 1)$
pk size	theoretical size of public key [bits]	$kn(\lceil \log_2(q) \rceil) + \kappa$
sk size	theoretical size of secret key [bits]	$n(k+1)(\lceil \log_2(t-1) \rceil + 1) + 2\kappa + 320$ with $t = 78$ or $111$

The values of  $b_{\text{GenA}}$  were chosen experimentally such that they maximize performance on the targeted x64 Intel platform; see Section 3.2.

**Bound parameters and acceptance probabilities.** The values  $L_S$  and  $L_E$  are used to bound the coefficients of the secret and error polynomials in the evaluation functions checkS and checkE, respectively. Bounding the size of those polynomials restricts the size of the key space; accordingly we compensate the security loss by choosing a larger bit hardness as explained in Section 5.2.1. Both bounds,  $L_S$  and  $L_E$  (and consequently  $S$  and  $E^2$ ), impact the rejection probability during the signature generation as follows. If one increases the

<sup>2</sup>In an earlier version of this document we needed to distinguish  $L_S/L_E$  and  $S/E$ . Although this is not necessary in this version, we keep all four values  $L_S, S, L_E, E$  for consistency reasons.

Table 4: Parameters for each of the proposed *provably-secure* parameter sets with  $q_h = \min\{2^\lambda, 2^{128}\}$  and  $q_s = \min\{2^{\lambda/2}, 2^{64}\}$ ; we choose  $M = 0.3$ .

Param.	qTESLA-p-I	qTESLA-p-III
$\lambda$	95	160
$\kappa$	256	256
$n$	1 024	2 048
$\sigma$	8.5	8.5
$k$	4	5
$q$	343 576 577 $\approx 2^{28}$	856 145 921 $\approx 2^{30}$
$h$	25	40
$L_E(= E), \eta_E$	554, 2.61	901, 2.65
$L_S(= S), \eta_S$	554, 2.61	901, 2.65
$B$	$2^{19} - 1$	$2^{21} - 1$
$d$	22	24
$b_{\text{GenA}}$	108	180
$ \Delta_{\mathbb{H}} $	$\approx 2^{435.8}$	$\approx 2^{750.9}$
$ \Delta_{\mathbb{S}} $	$\approx 2^{21502.4}$	$\approx 2^{47102.7}$
$ \Delta_{\mathbb{L}} $	$\approx 2^{94208.0}$	$\approx 2^{256000.0}$
$\delta_w$	0.37	0.33
$\delta_z$	0.34	0.42
$\delta_{\text{sign}}$	0.13	0.14
$\delta_{\text{keygen}}$	0.59	0.43
sig size [bytes]	2, 592	5, 664
pk size [bytes]	14, 880	38, 432
sk size [bytes]	5, 224	12, 392
classical bit hardness	150	304
quantum bit hardness	139	279

values of  $L_S$  and  $L_E$ , the acceptance probability during key generation, referred to as  $\delta_{\text{keygen}}$ , increases (see lines 8 and 13 in Alg. 6), while the acceptance probabilities of  $z$  and  $w$  during signature generation, referred to as  $\delta_z$  and  $\delta_w$  resp., decrease (see lines 12 and 18 in Alg. 7). We determine a good trade-off between the acceptance probabilities during key generation and signing experimentally. To this end, we start by choosing  $L_S = \eta_S \cdot h \cdot \sigma$  (resp.,  $L_E = \eta_E \cdot h \cdot \sigma$ ) with  $\eta_S = \eta_E = 2.8$  and compute the corresponding values for the parameters  $B$ ,  $d$  and  $q$  (which are chosen as explained later). We then carefully tune these parameters by trying different values for  $\eta_S$  and  $\eta_E$  in the range  $[2.0, \dots, 3.0]$  until we find a good trade-off between the different probabilities and, hence, runtimes.

The parameter  $B$  defines the interval of the random polynomial  $y$  (see line 4 of Alg. 7),

and it is determined by the parameters  $M$  and  $S$  as follows:

$$\left(\frac{2B - 2S + 1}{2B + 1}\right)^n \geq M \Leftrightarrow B \geq \frac{\sqrt[n]{M} + 2S - 1}{2(1 - \sqrt[n]{M})},$$

where  $M$  is a value of our choosing. Once  $B$  is chosen, we select the value  $d$  that determines the rounding functions  $[\cdot]_M$  and  $[\cdot]_L$  to be larger than  $\log_2(B)$ . Furthermore,  $d$  is chosen such that the acceptance probability of the check  $\|[w]_L\|_\infty \geq 2^{d-1} - E$  in line 18 of Algorithm 7 is lower bounded by  $M$ . This check determines the acceptance probability  $\delta_w$  during signature generation. The acceptance probability of  $z$ , namely  $\delta_z$ , is also related to the value of  $M$ . The final acceptance probabilities  $\delta_z$ ,  $\delta_w$  and  $\delta_{keygen}$  obtained experimentally following the procedure above are summarized in Table 4.

**The modulus  $q$ .** This parameter is chosen to fulfill several bounds and assumptions that are motivated by efficiency requirements and **qTESLA**'s security reduction. To enable the use of fast polynomial multiplication using the NTT,  $q$  must be a prime integer such that  $q \bmod 2n = 1$ . To choose parameters according to the security reduction, it is first convenient to simplify our security statement. To this end we ensure that  $q^{nk} \geq |\Delta\mathbb{S}| \cdot |\Delta\mathbb{L}| \cdot |\Delta\mathbb{H}|$  with the following definition of sets:  $\mathbb{S}$  is the set of polynomials  $z \in \mathcal{R}_{[B-S]}$  and  $\Delta\mathbb{S} = \{z - z' : z, z' \in \mathbb{S}\}$ ,  $\mathbb{H}$  is the set of polynomials  $c \in \mathcal{R}_{[1]}$  with exactly  $h$  nonzero coefficients and  $\Delta\mathbb{H} = \{c - c' : c, c' \in \mathbb{H}\}$ , and  $\Delta\mathbb{L} = \{x - x' : x, x' \in \mathcal{R} \text{ and } [x]_M = [x']_M\}$ . Then, the following equation (see Theorem 2 in Section 5.1) has to hold:

$$\frac{2^{3\lambda + nkd + 2} q_s^3 (q_s + q_h)^2}{q^{nk}} \leq 2^{-\lambda} \Leftrightarrow q \geq \left(2^{4\lambda + nkd + 2} q_s^3 (q_s + q_h)^2\right)^{1/nk}.$$

Following the NIST's call for proposals [48, Section 4.A.4], we choose the number of classical queries to the sign oracle to be  $q_s = \min\{2^{\lambda/2}, 2^{64}\}$  for all our parameter sets. Moreover, we choose the number of queries of a hash function to be  $q_h = \min\{2^\lambda, 2^{128}\}$ .

**Key and signature sizes.** The theoretical bitlengths of the signatures and public keys are given by  $\kappa + n \cdot (\lceil \log_2(B - S) \rceil + 1)$  and  $k \cdot n \cdot (\lceil \log_2(q) \rceil) + \kappa$ , respectively. To determine the size of the secret keys we first define  $t$  as the number of  $\beta$ -bit entries of the CDT tables (see Table 1) which corresponds to the maximum value that can be possibly sampled to generate the coefficients of secret polynomials  $s$ . Then, it follows that the theoretical size of the secret key is given by  $n(k + 1)(\lceil \log_2(t - 1) \rceil + 1) + 2\kappa + 320$  bits.

## 3 Performance analysis

### 3.1 Reference implementations

This document comes accompanied by simple yet efficient reference implementations written exclusively in portable C.

An important feature of `qTESLA` is that it enables efficient and compact implementations that can work for different security levels with minor changes. For example, our implementations of the parameter sets `qTESLA-p-I` and `qTESLA-p-III` share most of their codebase, and only differ in some packing functions and system constants that can be instantiated at compilation time. Moreover, the implementations are very compact and only consist of approximately 300 lines of C code each. This highlights the simplicity and scalability of software based on `qTESLA`.

All our implementations run in *constant-time*, i.e., they avoid the use of secret address accesses and secret branches and, hence, are protected against timing and cache side-channel attacks. The following functions are implemented securely via constant-time logical and arithmetic operations: `H`, `checkE`, `checkS`, the correctness test for rejection sampling, polynomial multiplication using the NTT, sparse multiplication, and all the polynomial operations requiring modular reductions or corrections. Some of the functions that perform some form of rejection sampling, such as the security test at signing, `GenA`, `ySampler`, and `Enc`, potentially leak the timing of the failure to some internal test, but this information is independent of the secret data. Table lookups performed in our implementation of the Gaussian sampler are done with linear passes over the full table and producing samples via constant-time logical and arithmetic operations.

Our polynomial arithmetic, which is dominated by polynomial multiplications based on the NTT, uses a signed 32-bit datatype to represent coefficients. Throughout polynomial computations, intermediate results are let to grow and are only reduced or corrected when there is a chance of exceeding 32 bits of length, after a multiplication, or when a result needs to be prepared for final packing (e.g., when outputting public keys). Accordingly, to avoid overflows the results of additions and subtractions are either corrected or reduced via Barrett reductions whenever necessary. We have performed a careful bound analysis for each of the proposed parameter sets in order to maximize the use of lazy reduction and cheap modular corrections in the polynomial arithmetic. In the case of multiplications, the results are reduced via Montgomery reductions. To minimize the cost of converting to/from Montgomery representation we use the following approach. First, the so-called “twiddle factors” in the NTT are scaled *offline* by multiplying with the Montgomery constant  $R = 2^{32} \bmod q$ . Similarly, the coefficients of the outputs  $a_i$  from `GenA` are scaled to remainders  $r' = rn^{-1}R \bmod q$  (and  $r' = rR \bmod q$  for the non-power-of-two case) by multiplying with the constant  $R^2 \cdot n^{-1}$ . This enables an efficient use of Montgomery reductions during the

NTT-based polynomial multiplication  $\text{NTT}^{-1}(\tilde{a} \circ \text{NTT}(b))$ , where  $\tilde{a} = \text{NTT}(a)$  is the output of GenA which is assumed to be in NTT domain. Multiplications with the twiddle factors during the computation of  $\text{NTT}(b)$  naturally cancel out the Montgomery constant. The same happens during the pointwise multiplication with  $\tilde{a}$ , and finally during the inverse NTT, which naturally outputs values in standard representation without the need for explicit conversions.

### 3.2 Performance of qTESLA on x64 Intel

We evaluated the performance of our implementations on a 3.4GHz Intel Core i7-6700 (Skylake) processor, running Ubuntu 16.04.3 LTS. As is standard practice, TurboBoost was disabled during the tests. For compilation we used gcc version 7.2.0 with the command `gcc -O3 -march=native -fomit-frame-pointer`.

The results for the reference implementations are summarized in Table 5.

Scheme	keygen	sign	verify	total (sign + verify)
qTESLA-p-I	2,358.6 (2,431.9)	2,299.0 (3,089.9)	814.3 (814.5)	3,113.3 (3,904.4)
qTESLA-p-III	13,151.4 (13,312.4)	5,212.3 (7,122.6)	2,102.3 (2,102.6)	7,314.6 (9,225.2)

Table 5: Performance (in thousands of cycles) of the reference implementations of qTESLA on a 3.4GHz Intel Core i7-6700 (Skylake) processor. Results for the median and average (in parenthesis) are rounded to the nearest  $10^2$  cycles. Signing is performed on a message of 59 bytes.

The combined (median) time of signing and verification on the Skylake platform is of approximately 0.92 and 2.15 milliseconds with qTESLA-p-I and qTESLA-p-III, respectively. This demonstrates that the speed of provably-secure qTESLA can be considered practical for most applications.

### 3.3 AVX2 optimizations

We optimized two functions with hand-written assembly exploiting AVX2 vector instructions, namely, polynomial multiplication and XOF expansion during sampling of  $y$ . Our polynomial multiplication follows the recent approach by Seiler [56], and the realization of the method has some similarities with the implementation from [28]. That is, our implementation processes 32 coefficients loaded in 8 AVX2 registers simultaneously, in such



a way that butterfly computations are carried out through multiple NTT levels without the need for storing and loading intermediate results, whenever possible.

One difference with [28, 56] is that our NTT coefficients are represented as 32-bit *signed* integers, which motivates a speedup in the butterfly computation by avoiding the extra additions that are required to make the result of subtractions positive when using an unsigned representation.

Our approach reduces the cost of the portable C polynomial multiplication from 76,300 to 18,400 cycles for  $n = 1024$ , and from 174,800 to 43,900 cycles for  $n = 2048$ .

Sampling of  $y$  is sped up by using the AVX2 implementation of SHAKE by Bertoni, Daemen, Hoffert, Peeters, Van Assche, and Van Keer [17], which allows us to sample up to 4 coefficients in parallel.

We note that it is possible to modify GenA to favor a vectorized computation of the XOF expansion inside this function. However, we avoid this optimization because it degrades the performance on smaller platforms with no vector instruction support.

Scheme	keygen	sign	verify	total (sign + verify)
qTESLA-p-I	2,212.4 (2,285.0)	1,370.4 (1,759.0)	678.4 (678.5)	2,048.8 (2,437.5)
qTESLA-p-III	12,791.0 (13,073.4)	3,081.9 (4,029.5)	1,745.3 (1,746.4)	4,827.2 (5,775.9)

Table 6: Performance (in thousands of cycles) of the AVX2 implementations of qTESLA on a 3.4GHz Intel Core i7-6700 (Skylake) processor. Results for the median and average (in parenthesis) are rounded to the nearest  $10^2$  cycles. Signing is performed on a message of 59 bytes.

## 4 Known answer values

The submission includes KAT values with tuples that contain message size (`mlen`), message (`msg`), public key (`pk`), secret key (`sk`), signature size (`slen`) and signature (`sm`) values for all the proposed parameter sets.

The KAT files for the reference implementations can be found in the media folder:

- qTESLA-p-I: `\KAT\ref\PQCsignKAT_qTesla-p-I.rsp`, and
- qTESLA-p-III: `\KAT\ref\PQCsignKAT_qTesla-p-III.rsp`.

## 5 Expected security strength

In this section we discuss the expected security strength of and possible attacks against qTESLA. This includes a statement about the theoretical security and the parameter choices depending on them. To this end we first define the decisional ring learning with errors (R-LWE) problem.

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

*Given  $k$  tuples  $(a_1, t_1), \dots, (a_k, t_k)$ , the decisional R-LWE problem R-LWE $_{n,k,q,\chi}$  is to distinguish whether  $(a_i, t_i) \leftarrow \mathcal{U}(\mathcal{R}_q \times \mathcal{R}_q)$  or  $(a_i, t_i) \leftarrow \mathcal{D}_{s,\chi}$  for all  $i$ . The R-LWE advantage is defined as*

$$\text{Adv}_{n,k,q,\chi}^{\text{R-LWE}}(\mathcal{A}) = \left| \Pr \left[ \mathcal{A}^{\mathcal{D}_{s,\chi}(\cdot)} = 1 \right] - \Pr \left[ \mathcal{A}^{\mathcal{U}(\mathcal{R}_q \times \mathcal{R}_q)(\cdot)} = 1 \right] \right|.$$

### 5.1 Provable security in the quantum random oracle model

The asymptotic security for qTESLA is expected to follow from [43], where it is shown how the security of the Fiat-Shamir transformation transfers over to the quantum setting.

In addition, the concrete security of qTESLA is supported by a security reduction (see Theorem 2) that gives a reduction from the R-LWE problem to the existential unforgeability under chosen-message attack (EUF-CMA) security of qTESLA in the quantum random oracle model. It is very similar to [8, Theorem 1]. Currently, Theorem 2 holds assuming a conjecture, as explained in Appendix A where we also sketch the proof. The formal security proof is not included in this document because it is very close to the original result [8].

**Theorem 2** (Security reduction from R-LWE). *Let the parameters be as in Table 3, in particular, let  $q^{nk} \geq 2^{4\lambda+nkd} 4q_s^3(q_s + q_h)^2$ . Assume that Conjecture 5 holds. Assume that there exists a quantum adversary  $\mathcal{A}$  that forges a qTESLA signature in time  $t_{\text{qTESLA}}$ , making at most  $q_h$  (quantum) queries to its quantum random oracle and  $q_s$  (classical) queries to its signing oracle. Then there exists a reduction  $\mathcal{S}$  that solves the R-LWE problem with*

$$\text{Adv}_{\text{qTESLA}}^{\text{EUF-CMA}}(\mathcal{A}) \leq \text{Adv}_{k,n,q,\sigma}^{\text{R-LWE}}(\mathcal{S}) + \frac{2^{3\lambda+nkd} \cdot 4 \cdot q_s^3(q_s + q_h)^2}{q^{nk}} + \frac{2(q_h + 1)}{\sqrt{2^h \binom{n}{h}}} \quad (7)$$

*and in time  $t_{\text{R-LWE}}$  which is about the same as  $t_{\text{qTESLA}}$  in addition to the time to simulate the quantum random oracle.*

**Remark 3.** (*Uniform sampling and PRF security*) Define the signature scheme  $\text{qTESLA}' = (\text{KeyGen}', \text{Sign}', \text{Verify}')$ .  $\text{KeyGen}'$  and  $\text{Verify}'$  are given in Algorithm 6 and 8, i.e., they are the same as the original  $\text{qTESLA}$  algorithms.  $\text{Sign}'$  differs from  $\text{qTESLA}$ 's  $\text{sign}$ , cf. Algorithm 7, as it chooses  $y \leftarrow_{\S} \mathcal{R}_{[B]}$  instead of computing it by  $\text{PRF}_2$  and  $\text{ySampler}$ . The two security reductions given in this section actually bound the advantage of forging a  $\text{qTESLA}'$  signature. The advantage of breaking the EUF-CMA security of  $\text{qTESLA}$  is, hence, upper bounded by the sum of the advantage of distinguishing  $\text{ySampler} \circ \text{PRF}_2$  from a truly random function and the advantage of breaking the EUF-CMA security of  $\text{qTESLA}'$ . In addition, the advantage against  $\text{qTESLA}$  is depending on the acceptance probability of  $\text{checkE}$  and  $\text{checkS}$ , since these functions reduce the key space.

## 5.2 Bit security of our proposed parameter sets

In the following, we describe how we estimate the concrete security of the proposed parameters described in Section 2.6. To this end, we first describe how the security of our scheme depends on the hardness of R-LWE and afterwards we describe how we derive the bit hardness of the underlying R-LWE instances.

### 5.2.1 Correspondence between security and hardness

The security reduction given by Theorem 2 in Section 5.1 provides a reduction from the hardness of the decisional ring learning with errors problem and bounds *explicitly* the forging probability with the success probability of the reduction. More formally, let  $\epsilon_{\Sigma}$  and  $t_{\text{qTESLA}}$  denote the success probability and the runtime (resp.) of a forger against our signature scheme, and let  $\epsilon_{LWE}$  and  $t_{\text{R-LWE}}$  denote analogous quantities for the reduction presented in the proof of Theorem 2. We say that R-LWE is  $\eta$ -bit hard if  $t_{\text{R-LWE}}/\epsilon_{LWE} \geq 2^{\eta}$ ; and we say that the signature scheme is  $\lambda$ -bit secure if  $t_{\text{qTESLA}}/\epsilon_{\Sigma} \geq 2^{\lambda}$ .

For our *provably-secure* parameter sets, we choose parameters such that  $\epsilon_{LWE} \approx \epsilon_{\Sigma}$  and  $t_{\text{qTESLA}} \approx t_{\text{R-LWE}}$ <sup>3</sup>, that is, the bit hardness of the R-LWE instance is *theoretically* the same as the bit security of our signature scheme, by virtue of the security reduction and its tightness. Hence, the security reduction provably guarantees that our scheme instantiated with the provably-secure parameter sets has the selected security level as long as the corresponding R-LWE instance gives the assumed hardness level. This approach provides a *stronger* security argument.

---

<sup>3</sup>To be precise, we assume that the time to simulate the (quantum) random oracle is smaller than the time to forge a signature. This assumption is commonly made in “provably secure” cryptography.

**Remark 4.** In practical instantiations of  $\text{qTESLA}$ , the bit security does not exactly match the bit hardness of R-LWE, see Table 4. This is because the bit security does not only depend on the bit hardness of R-LWE (as explained above), but also on the probability of rejected/accepted key pairs and on the security of other building blocks such as the encoding function  $\text{Enc}$ . First, in all our parameter sets, the key space is reduced by the rejection of polynomials  $s, e_1, \dots, e_k$  with large coefficients via  $\text{checkE}$  and  $\text{checkS}$ . In particular, depending on the instantiation, the size of the key space is decreased by  $\lceil \log_2(\delta_{\text{KeyGen}}) \rceil$  bits. We compensate this security loss by choosing an R-LWE instance of larger bit hardness. Hence, the corresponding R-LWE instances give at least  $\lambda + \lceil \log_2(\delta_{\text{KeyGen}}) \rceil$  bits of hardness against currently known (classical and quantum) attacks. Finally, we instantiate the encoding function  $\text{Enc}$  such that it is  $\lambda$ -bit secure.

Accordingly, we claim a bit security that is strictly smaller than the hardness of the corresponding R-LWE instance. For example, the hardness of the R-LWE instance corresponding to  $\text{qTESLA-p-I}$  is 139 bits but we claim a bit security of 95.

### 5.2.2 Estimation of the R-LWE hardness

Since the introduction of the learning with errors problem over rings [45], it has remained an open question to determine whether the R-LWE problem is as hard as the LWE problem. Several results exist that exploit the structure of some ideal lattices [22, 25, 30, 32]. However, up to now, these results do not seem to apply to R-LWE instances that are typically used in signature schemes and, therefore, do not apply to the proposed  $\text{qTESLA}$  instances. Consequently, we assume that the R-LWE problem is as hard as the LWE problem, and estimate the hardness of R-LWE using state-of-the-art attacks against LWE.

Albrecht, Player, and Scott [7] presented the *LWE-Estimator*, a software to estimate the hardness of LWE given the matrix dimension  $n$ , the modulus  $q$ , the relative error rate  $\alpha = \frac{\sqrt{2\pi}\sigma}{q}$ , and the number of given LWE samples. The LWE-Estimator estimates the hardness against the fastest LWE solvers currently known, i.e., it outputs an upper (conservative) bound on the number of operations an attack needs to break a given LWE instance. In particular, the following attacks are considered in the LWE-Estimator: the meet-in-the-middle exhaustive search, the coded Blum-Kalai-Wassermann algorithm [36], the dual lattice attacks recently published in [2], the enumeration approach by Lindner and Peikert [42], the primal attack described in [5, 13], the Arora-Ge algorithm [11] using Gröbner bases [3], and the latest analysis to compute the block sizes used in the lattice basis reduction BKZ, recently published by Albrecht *et al.* [6]. Moreover, quantum speedups for the sieving algorithm used in BKZ [40, 41] are also considered.

Arguably, the most important building block in most efficient attacks against the underlying R-LWE instance in  $\text{qTESLA}$  is BKZ. Hence, the model used to estimate the cost of

BKZ determines the overall hardness estimation of our instances. While many different cost models for BKZ exist [4], we decided to adopt the BKZ cost model of  $0.265\beta + 16.4 + \log_2(8d)$  for the hardness estimation of our parameters, where  $\beta$  is the BKZ block size and  $d$  is the lattice dimension. In the LWE-Estimator this corresponds to using the option `cost_model = BKZ.qsieve`. This cost model is very conservative in the following sense: it only takes into account the number of operations needed to solve a certain instance and it assumes that the attacker can handle huge amounts of quantum memory. At the same time it matches practical state-of-the-art attacks, where (classical) experiments [24, 52] show that during BKZ an SVP oracle is required to be called several times instead of only once as it is assumed in even more conservative models such as the cost model proposed in [9]. Still, to deal with potential future advances in cryptanalysis, we choose instances that currently provide a higher hardness level than the targeted security level. For example, the hardness of `qTESLA-p-I` is estimated to be 139 bits (see Table 7) while we target a security of 95 bits for security category 1. We compare our chosen hardness estimation with the BKZ model from [9] in Table 7 (in the LWE-Estimator this cost model corresponds to using the option `cost_model = partial(BKZ.ADPS16,mode="quantum")`). Furthermore, we display our hardness estimation according to the fastest classical algorithms, which is using classical sieving to implement the SVP oracle in BKZ. To this end, we use the cost model from [16] which corresponds to using the option `cost_model = BKZ.sieve`.

Table 7: Security estimation (bit hardness) under different BKZ cost models of the proposed parameter sets.

BKZ cost model	qTESLA-p-I	qTESLA-p-III
<code>BKZ.sieve</code>	150	304
<b><code>BKZ.qsieve</code></b>	<b>139</b>	<b>279</b>
<code>BKZ.ADPS16,mode="quantum"</code>	108	247
Targeted quantum bit security	95	160
Security category	I	III

We note that another recent quantum attack by Göpfert, van Vredendaal, and Wunderer [34], called quantum hybrid attack, is not considered in our analysis and the LWE-Estimator. This hybrid attack is most efficient on the learning with errors problem with very small secret and error, e.g., binary or ternary. Since the coefficients of the secret and error polynomials of `qTESLA` are chosen Gaussian distributed, the attack is not efficient for our instances.

The LWE-Estimator is the result of many different contributions and contributors. It is open source and hence easily checked and maintained by the community. Therefore, we find the LWE-Estimator to be a suitable tool to estimate the hardness of our chosen LWE instances. We integrated the LWE-Estimator with commit-id 3019847 on 2019-02-14 in

our sage script that is also included in this submission.

The most efficient LWE solvers for our instances are the decoding attack and the embedding approach. We refer to [51] for a description of these attacks.

### 5.3 Resistance to implementation attacks

Besides the theoretical security against computational attacks, such as lattice reduction, it is important for a cryptographic scheme to be secure against implementation attacks. These attacks come in two flavors: side-channel and fault analysis attacks.

#### 5.3.1 Side-channel analysis

These attacks exploit physical information such as timing or power consumption, electromagnetic emanation, etc., that is correlated to some secret information during the execution of a cryptographic scheme. Simple and differential side-channel attacks that rely on power and electromagnetic emanations are very powerful but typically require physical access (or close proximity) to the targeted device. Protecting lattice-based schemes against this class of attacks is a very active area of research.

In contrast, attacks that exploit timing leakage, such as timing and cache attacks, are easier to carry out remotely. Hence, these attacks represent a more immediate danger for most applications and, consequently, it has become a minimum security requirement for a cryptographic implementation to be secure against this class of attacks. One effective approach to provide such a protection is by guaranteeing so-called *constant-time* execution. In practice, this means that an implementation should avoid the use of secret address accesses and conditional branches based on secret information and that the execution time should be independent of secret data.

One of the main advantages of qTESLA is that the Gaussian sampler, likely the most complex part of the scheme, is restricted to key generation. This reduces drastically the attack surface to carry out a timing and cache attack against qTESLA. Moreover, we emphasize that qTESLA's Gaussian sampler is simple and can be implemented securely in a constant-time manner, as can be observed in the accompanying implementations. Other functions of qTESLA, such as polynomial arithmetic operations, are easy to implement in constant-time as well.

Recently, the scheme ring-TESLA [1] was analyzed with respect to cache side channels with the software tool CacheAudit [19]. It was the first time that a post-quantum scheme was analyzed with program analysis. The authors found potential cache side channels, proposed countermeasures, and showed the effectiveness of their mitigations with CacheAudit. In

our implementations, we apply similar techniques to those proposed in [19] with some additional optimizations.

It is relevant to note that qTESLA includes *built-in* defenses against several attack scenarios, thanks to its probabilistic nature. Specifically, the seed used to generate the randomness  $y$  is produced by hashing the value  $\text{seed}_y$  that is part of the secret key, some fresh randomness  $r$ , and the digest  $G(m)$  of the message  $m$ . The random value  $r$  guarantees the use of a fresh  $y$  at each signing operation, which increases the difficulty to carry out side-channel attacks against the scheme. Moreover, this fresh  $y$  prevents some easy-to-implement but powerful fault attacks against deterministic signature schemes, as explained next.

### 5.3.2 Fault analysis

The use of  $\text{seed}_y$  makes qTESLA resilient to a catastrophic failure of the Random Number Generator (RNG) during generation of the fresh randomness, protecting against fixed-randomness attacks such as the one demonstrated against Sony’s Playstation 3 [23].

Recently, some studies have exposed the vulnerability of lattice-based schemes to fault attacks. We describe a simple yet powerful attack that falls in this category of attacks [21].

Assume that line 3 of Algorithm 7 is computed without the random value  $r$ , i.e., as  $\text{rand} \leftarrow \text{PRF}_2(\text{seed}_y, m)$ . Assume that a signature  $(z, c)$  is generated for a given message  $m$ . Afterwards, a signature is requested again for the same message  $m$ , but this time, a fault is injected on the computation of the hash value  $c$  yielding the value  $c_{\text{faulted}}$ . This second signature is  $(z_{\text{faulted}}, c_{\text{faulted}})$ . Computing  $z - z_{\text{faulted}} = sc - sc_{\text{faulted}} = s(c - c_{\text{faulted}})$ , reveals the secret  $s$  since  $c - c_{\text{faulted}}$  is known to the attacker. As stated in [53], this attack has broad implications since it is generically applicable to deterministic *Schnorr-like* signatures.

It is easy to see that, to prevent this (and other similar) fault attacks, every signing operation should be injected with fresh randomness, as precisely specified in line 3 of Algorithm 7. This makes qTESLA implicitly resilient to this line of attacks.

## 6 Advantages and limitations

In this section, we summarize some advantages and limitations of qTESLA.

**Security foundation.** qTESLA comes accompanied by an *explicit* and *tight* security reduction in the quantum random oracle model from R-LWE (Theorem 2), i.e., a quantum

adversary is allowed to ask the random oracle in superposition. This reduction is based on a variant of our scheme over standard lattices [8]. To port the reduction given in [8], we use a heuristic argument as explained in Section 5.1. Since our security reduction is *explicit*, we can explicitly give the relation between the success probabilities of solving the R-LWE problem and forging qTESLA signatures, enabling the selection of parameters according to this reduction. The tightness of the reduction enables smaller parameters and, thus, better performance of the provably-secure instantiations.

The bit security of the proposed parameters was estimated against known state-of-the-art classical and quantum algorithms that solve the LWE problem. The proposed parameters incorporate a very comfortable margin between the targeted and the estimated bit security in order to deal with future or unknown LWE solvers as well.

Finally, the parameter generation of qTESLA is easy to audit: the choice of parameters and their relation are clearly explained; and the generation procedure is simple and easy to follow. In order to facilitate the generation of new parameters (if needed), and also for transparency, we make our sage script for parameter generation available<sup>4</sup>.

**Ease of implementation and scalability.** qTESLA has a very compact and simple structure consisting of a few, easy-to-implement functions. The Gaussian sampler, arguably the most complex function in qTESLA, is only required during key generation. Therefore, even if the efficient Gaussian sampler included in this document is not used, most applications will not be impacted by the use of a slower Gaussian sampler.

qTESLA exhibits great scalability, i.e., it is easy to support different security levels with a common codebase. For instance, our reference implementations, written in portable C, consist of about 300 lines of code shared among all the parameter sets<sup>5</sup>.

**Security against implementation attacks.** qTESLA's simplicity facilitates its implementation in constant-time and, arguably, its protection against more powerful physical attacks such as differential power analysis. As stated before, Gaussian sampling is only required during key generation, which reduces significantly the attack surface over this function. Moreover, qTESLA requires a *simple* Gaussian sampler which further eases implementation and protection against side-channel attacks.

qTESLA comes equipped with built-in measures that provide a first layer of defense against some side-channel and fault attacks. Since qTESLA generates a fresh  $y$  per signing operation,

---

<sup>4</sup>The parameter generation script can be found at `\Supporting_Documentation\Script_to_choose_parameters\parameterchoice.sage`.

<sup>5</sup>This count excludes the parameter-specific packing functions, header files, NTT constants, and (c)SHAKE functions.



some simple side-channel attacks are more difficult to mount against the scheme. More importantly, this feature immediately renders some powerful and easy-to-carry out fault attacks unfeasible. At the same time, qTESLA is resilient to a catastrophic failure of the RNG during generation of the fresh randomness that is required to generate  $y$ .

In addition, qTESLA is protected against Key Substitution (KS) attacks [20, 46], providing improved security in the multi-user setting; see also [37].

**Cryptographic libraries and hybrid mode.** The reference implementation of qTESLA is integrated in several cryptographic libraries. In particular, the C reference implementation is integrated to the library *pqm4* [38] that targets microcontrollers, and the post-quantum cryptography library *liboqs* [10] as part of the *Open Quantum Safe* project [47]. Also, an implementation of qTESLA written in Java (not included in this submission) is integrated in the cryptographic library *BouncyCastle*.

Another contribution of the *Open Quantum Safe* project is an OpenSSL fork that integrates *liboqs* to OpenSSL and includes post-quantum and hybrid authentication in TLS [49]. Hybrid authentication enables to authenticate using classical and post-quantum signature schemes in order to preserve classical security, while adding post-quantum security. Additional implementations of hybrid certificates and hybrid authentication in TLS 1.3 can be found in [33] and [57] using the qTESLA integrations in *liboqs* and *BouncyCastle*, respectively. This shows the suitability of qTESLA used in hybrid schemes, an approach that seems to be favored by industry to enable a smooth and secure transition to post-quantum cryptography.

**Diverse applications.** qTESLA is present in several applications and projects from which we shortly mention two of them. qTESLA is used in a field study to protect medical health data through an experimental implementation of TLS. This collaboration between TU Darmstadt, the National Institute of Information and Communications Technology (NICT) and ISARA has been presented at the 6th ETSI/IQC Workshop on Quantum-Safe Cryptography.<sup>6</sup>

A variant of qTESLA was used as basis in the design of a new post-quantum certificate provisioning process for vehicle-to-everything (V2X) communications. This application allows provisioning of pseudonym certificates such that vehicles are able to engage in V2X communications in an environment with privacy by design, meaning that vehicles cannot be tracked by other entities in the system by the use of these pseudonym certificates [15].

---

<sup>6</sup>Contact persons: Atsushi Yamada (ISARA corporation), [Atsushi.Yamada@isara.com](mailto:Atsushi.Yamada@isara.com) and Masahide Sasaki (NICT), [psasaki@nict.go.jp](mailto:psasaki@nict.go.jp)

## References

- [1] Sedat Akleylek, Nina Bindel, Johannes A. Buchmann, Juliane Krämer, and Giorgia Azzurra Marson. An efficient lattice-based signature scheme with provably secure instantiation. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 2016*, volume 9646 of *LNCS*, pages 44–60. Springer, 2016.
- [2] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in HELib and SEAL. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part II*, volume 10211 of *Lecture Notes in Computer Science*, pages 103–129, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [3] Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, Robert Fitzpatrick, and Ludovic Perret. Algebraic algorithms for LWE problems. *ACM Comm. Computer Algebra*, 49(2):62, 2015.
- [4] Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the LWE, NTRU schemes! In Dario Catalano and Roberto De Prisco, editors, *SCN 18: 11th International Conference on Security in Communication Networks*, volume 11035 of *Lecture Notes in Computer Science*, pages 351–367, Amalfi, Italy, September 5–7, 2018. Springer, Heidelberg, Germany.
- [5] Martin R. Albrecht, Robert Fitzpatrick, and Florian Göpfert. On the efficacy of solving LWE by reduction to unique-SVP. In Hyang-Sook Lee and Dong-Guk Han, editors, *ICISC 13: 16th International Conference on Information Security and Cryptology*, volume 8565 of *Lecture Notes in Computer Science*, pages 293–310, Seoul, Korea, November 27–29, 2014. Springer, Heidelberg, Germany.
- [6] Martin R. Albrecht, Florian Göpfert, Fernando Virdia, and Thomas Wunderer. Revisiting the expected cost of solving uSVP and applications to LWE. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017, Part I*, volume 10624 of *Lecture Notes in Computer Science*, pages 297–322, Hong Kong, China, December 3–7, 2017. Springer, Heidelberg, Germany.
- [7] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
- [8] Erdem Alkim, Nina Bindel, Johannes A. Buchmann, Özgür Dagdelen, Edward Eaton, Gus Gutoski, Juliane Krämer, and Filip Pawlega. Revisiting TESLA in the quantum random oracle model. In Tanja Lange and Tsuyoshi Takagi, editors, *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017*, pages 143–162, Utrecht, The Netherlands, June 26–28 2017. Springer, Heidelberg, Germany.

- [9] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium, USENIX Security 16*, pages 327–343. USENIX Association, 2016.
- [10] Nicholas Allen, Maxime Anvari, Eric Crockett, Nir Drucker, Vlad Gheorghiu, Shay Gueron, Christian Paquin, Tancréde Lepoint, Shravan Mishra, and Douglas Stebila. liboqs – nist-branch: C library for quantum-resistant cryptographic algorithms, 2018. GitHub at <https://github.com/open-quantum-safe/liboqs>, commit-id: 86c6ab1.
- [11] Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP 2011: 38th International Colloquium on Automata, Languages and Programming, Part I*, volume 6755 of *Lecture Notes in Computer Science*, pages 403–415, Zurich, Switzerland, July 4–8, 2011. Springer, Heidelberg, Germany.
- [12] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *Topics in Cryptology – CT-RSA 2014*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47, San Francisco, CA, USA, February 25–28, 2014. Springer, Heidelberg, Germany.
- [13] Shi Bai and Steven D. Galbraith. Lattice decoding attacks on binary LWE. In Willy Susilo and Yi Mu, editors, *ACISP 14: 19th Australasian Conference on Information Security and Privacy*, volume 8544 of *Lecture Notes in Computer Science*, pages 322–337, Wollongong, NSW, Australia, July 7–9, 2014. Springer, Heidelberg, Germany.
- [14] 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. <http://eprint.iacr.org/2016/1026>.
- [15] Paulo S. L. M. Barreto, Jefferson E. Ricardini, Marcos A. Simplicio Jr., and Harsh Kupwade Patil. qscms: Post-quantum certificate provisioning process for v2x. Cryptology ePrint Archive, Report 2018/1247, 2018.
- [16] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In Robert Krauthgamer, editor, *27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 10–24, Arlington, VA, USA, January 10–12, 2016. ACM-SIAM.
- [17] Guido Bertoni, Joan Daemen, Seth Hoffert, Michaël Peeters, , Gilles Van Assche, and Ronny Van Keer. The eXtended Keccak Code Package (XKCP). <https://github.com/XKCP/XKCP>.
- [18] Nina Bindel, Johannes Buchmann, and Juliane Krämer. Lattice-based signature schemes and their sensitivity to fault attacks. In *2016 Workshop on Fault Diagnosis*

- and Tolerance in Cryptography, FDTC 2016*, pages 63–77. IEEE Computer Society, 2016.
- [19] Nina Bindel, Johannes Buchmann, Juliane Krämer, Heiko Mantel, Johannes Schickel, and Alexandra Weber. Bounding the cache-side-channel leakage of lattice-based signature schemes using program semantics. In *Proceedings of the 10th International Symposium on Foundations & Practice of Security (FPS)*, 2017. To appear.
- [20] Simon Blake-Wilson and Alfred Menezes. Unknown key-share attacks on the station-to-station (STS) protocol. In Hideki Imai and Yuliang Zheng, editors, *Public Key Cryptography (PKC'99)*, volume 1560 of *Lecture Notes in Computer Science*, pages 154–170. Springer, 1999.
- [21] Leon Groot Bruinderink and Peter Pessl. Differential fault attacks on deterministic lattice signatures. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(3):21–43, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/7267>.
- [22] Peter Campbell, Michael Groves, and Dan Shepherd. SOLILOQUY: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. [http://docbox.etsi.org/Workshop/2014/201410\\_CRYPT0/S07\\_Systems\\_and\\_Attacks/S07\\_Groves\\_Annex.pdf](http://docbox.etsi.org/Workshop/2014/201410_CRYPT0/S07_Systems_and_Attacks/S07_Groves_Annex.pdf).
- [23] H.M. Cantero, S. Peter, Bushing, and Segher. Console hacking 2010 – PS3 epic fail. 27th Chaos Communication Congress, 2010. [https://www.cs.cmu.edu/~dst/GeoHot/1780\\_27c3\\_console\\_hacking\\_2010.pdf](https://www.cs.cmu.edu/~dst/GeoHot/1780_27c3_console_hacking_2010.pdf).
- [24] Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. PhD thesis, Paris, France, 2013.
- [25] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 9666 of *LNCS*, pages 559–585. Springer, 2016.
- [26] 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 Diego F. Aranha and Alfred Menezes, editors, *Progress in Cryptology - LATINCRYPT 2014*, volume 8895 of *LNCS*, pages 84–103. Springer, 2015.
- [27] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer*

*Science*, pages 40–56, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

- [28] Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(1):238–268, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/839>.
- [29] Morris J. Dworkin. SHA-3 standard: Permutation-based hash and extendable-output functions. *Federal Inf. Process. Stds. (NIST FIPS) – 202*, 2015. Available at <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.202.pdf>.
- [30] Yara Elias, Kristin E. Lauter, Ekin Ozman, and Katherine E. Stange. Provably weak instances of ring-LWE. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference*, volume 9215 of *LNCS*, pages 63–92. Springer, 2015.
- [31] Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Loop-abort faults on lattice-based fiat-shamir and hash-and-sign signatures. In Roberto Avanzi and Howard M. Heys, editors, *Selected Areas in Cryptography - SAC 2016 - 23rd International Conference*, volume 10532 of *Lecture Notes in Computer Science*, pages 140–158. Springer, 2017.
- [32] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany.
- [33] Luca Gladiator and Tobias Stöckert. Fork of OQS-OpenSSL\_1.1.1-stable, 2019. GitHub at <https://github.com/CROSSINGTUD/openssl-hybrid-certificates>, commit-id: 6ea0607.
- [34] Florian Göpfert, Christine van Vredendaal, and Thomas Wunderer. A hybrid lattice basis reduction and quantum search attack on LWE. In Tanja Lange and Tsuyoshi Takagi, editors, *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017*, volume 10346 of *LNCS*, pages 184–202. Springer, 2017.
- [35] Shay Gueron and Fabian Schlieker. Optimized implementation of ring-TESLA. GitHub at <https://github.com/fschlieker/ring-TESLA>, 2016.
- [36] Qian Guo, Thomas Johansson, and Paul Stankovski. Coded-BKW: Solving LWE using lattice codes. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology – CRYPTO 2015*, volume 9215 of *LNCS*, pages 23–42. Springer, 2015.

- [37] Dennis Jackson, Cas Cremers, Katriel Cohn-Gordon, and Ralf Sasse. Seems legit: Automated analysis of subtle attacks on protocols that use signatures. Cryptology ePrint Archive, Report 2019/779 (to appear at ACM CCS'19), 2019. <https://eprint.iacr.org/2019/779>.
- [38] Matthias J. Kannwischer, Joost Rijneveld, Peter Schwabe, and Ko Stoffelen. pqm4–Post-quantum crypto library for the ARM Cortex-M4, 2018. GitHub at <https://github.com/mupq/pqm4>, commit-id: 133c0e8.
- [39] John Kelsey. SHA-3 derived functions: cSHAKE, KMAC, TupleHash, and Parallel-Hash. *NIST Special Publication*, 800:185, 2016. Available at <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-185.pdf>.
- [40] Thijs Laarhoven. *Search problems in cryptography*. PhD thesis, Eindhoven University of Technology, 2016.
- [41] Thijs Laarhoven, Michele Mosca, and Joop Pol. Solving the Shortest Vector Problem in Lattices Faster Using Quantum Search. In Philippe Gaborit, editor, *Post-Quantum Cryptography*, volume 7932 of *Lecture Notes in Computer Science*, pages 83–101. Springer Berlin Heidelberg, 2013.
- [42] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *Topics in Cryptology – CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339, San Francisco, CA, USA, February 14–18, 2011. Springer, Heidelberg, Germany.
- [43] Qipeng Liu and Mark Zhandry. Revisiting Post-Quantum Fiat-Shamir. Cryptology ePrint Archive, Report 2019/262, 2019. <https://eprint.iacr.org/2019/262>.
- [44] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.
- [45] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.
- [46] Alfred Menezes and Nigel P. Smart. Security of signature schemes in a multi-user setting. *Des. Codes Cryptogr.*, 33(3):261–274, 2004.
- [47] Michele Mosca and Douglas Stebila. Open quantum safe – software for prototyping quantum-resistant cryptography, 2018. <https://openquantumsafe.org/>. Accessed: 2018-07-26.

- [48] National Institute of Standards and Technology (NIST). Submission requirements and evaluation criteria for the post-quantum cryptography standardization process, December, 2016. <https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf>. Accessed: 2018-07-23.
- [49] Christian Paquin and Douglas Stebila. OQS-OpenSSL\_1.1.1-stable, 2018. GitHub at <https://github.com/open-quantum-safe/openssl>, commit-id: 57a9724.
- [50] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 80–97, Santa Barbara (CA), 2010. Springer.
- [51] Rachel Player. *Parameter selection in lattice-based cryptography*. PhD thesis, Royal Holloway, University of London, 2017.
- [52] Rachel Player. *Parameter selection in lattice-based cryptography*. PhD thesis, Royal Holloway, University of London, London, United Kingdom, 2018.
- [53] Damian Poddebniak, Juraj Somorovsky, Sebastian Schinzel, Manfred Lochter, and Paul Rösler. Attacking deterministic signature schemes using fault attacks. *Cryptology ePrint Archive*, Report 2017/1014, 2017. <http://eprint.iacr.org/2017/1014>.
- [54] Thomas Pöppelmann and Tim Güneysu. Towards practical lattice-based public-key encryption on reconfigurable hardware. In Tanja Lange, Kristin E. Lauter, and Petr Lisonek, editors, *Selected Areas in Cryptography - SAC 2013*, volume 8282 of *Lecture Notes in Computer Science*, pages 68–85. Springer, 2014.
- [55] Sujoy Sinha Roy, Frederik Vercauteren, Nele Mentens, Donald Donglong Chen, and Ingrid Verbauwhede. Compact Ring-LWE cryptoprocessor. In Lejla Batina and Matthew Robshaw, editors, *Cryptographic Hardware and Embedded Systems - CHES 2014*, volume 8731 of *Lecture Notes in Computer Science*, pages 371–391. Springer, 2014.
- [56] Gregor Seiler. Faster AVX2 optimized NTT multiplication for Ring-LWE lattice cryptography. *Cryptology ePrint Archive*, Report 2018/039, 2018. <https://eprint.iacr.org/2018/039>.
- [57] Johannes Wirth. HybridCertificates, 2018. GitHub at <https://github.com/jojowi/HybridCertificate>, commit-id: 8518175.
- [58] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 758–775, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.



## A Sketch of proof of Theorem 2

**Theorem 2** (Security reduction from R-LWE). *Let the parameters be as in Table 3, in particular, let  $q^{nk} \geq 2^{4\lambda+nkd} 4q_s^3(q_s + q_h)^2$ . Assume that Conjecture 5 holds. Assume that there exists a quantum adversary  $\mathcal{A}$  that forges a qTESLA signature in time  $t_{\text{qTESLA}}$ , making at most  $q_h$  (quantum) queries to its quantum random oracle and  $q_s$  (classical) queries to its signing oracle. Then there exists a reduction  $\mathcal{S}$  that solves the R-LWE problem with*

$$\text{Adv}_{\text{qTESLA}}^{\text{EUF-CMA}}(\mathcal{A}) \leq \text{Adv}_{k,n,q,\sigma}^{\text{R-LWE}}(\mathcal{S}) + \frac{2^{3\lambda+nkd} \cdot 4 \cdot q_s^3(q_s + q_h)^2}{q^{nk}} + \frac{2(q_h + 1)}{\sqrt{2^h \binom{n}{h}}} \quad (7)$$

and in time  $t_{\text{R-LWE}}$  which is about the same as  $t_{\text{qTESLA}}$  in addition to the time to simulate the quantum random oracle.

The proof follows the approach proposed in [8], which shows the security of qTESLA’s predecessor TESLA. It uses the reductionist’s approach that assumes the existence of an adversary  $\mathcal{A}$  that forges a qTESLA signature. We then construct an algorithm  $\mathcal{S}$  that solves the (decisional) R-LWE problem in time  $t_{\text{R-LWE}} \approx t_{\text{qTESLA}}$  (plus the time to simulate the responses to  $\mathcal{A}$ ’s hash queries) and show that the advantage of solving R-LWE is about the same as the advantage of forging a qTESLA signature. Under the assumption that the R-LWE problem is computationally hard, it must follow that qTESLA is secure.

Specifically, the idea of the security reduction is as follows. Let  $\mathcal{A}$  be an algorithm that breaks qTESLA, i.e., given an “expanded” public key  $(t_1, \dots, t_k, a_1, \dots, a_k)$ , algorithm  $\mathcal{A}$  outputs  $(z, c', m)$  after some time  $t_{\text{qTESLA}}$ . Let  $\text{forge}(t_1, \dots, t_k, a_1, \dots, a_k)$  denote the event that the forger  $\mathcal{A}$  successfully produces a *valid* signature for  $(t_1, \dots, t_k, a_1, \dots, a_k)$ , i.e., with probability  $\Pr[\text{forge}(t_1, \dots, t_k, a_1, \dots, a_k)]$ ,  $(z, c')$  is a *valid* signature of message  $m$ . We model the hash-based function  $\text{H}$  as a quantum random oracle. In particular, algorithm  $\mathcal{A}$  is allowed to make (at most)  $q_h$  quantum queries to a quantum random oracle  $\text{H}$  and (at most)  $q_s$  classical queries to a qTESLA signing oracle. However, the message  $m$  that is returned by  $\mathcal{A}$  must not be queried to the signing oracle. We then build an algorithm  $\mathcal{S}$  that solves the decisional R-LWE problem with a runtime that is close to that of  $\mathcal{A}$  and with a probability of success close to  $\Pr[\text{forge}(t_1, \dots, t_k, a_1, \dots, a_k)]$ .

The reduction  $\mathcal{S}$  gets as input a tuple  $(t_1, \dots, t_k, a_1, \dots, a_k)$  and must decide whether it follows the R-LWE distribution (see Definition 1) or  $\mathcal{U}(\mathcal{R}_q^{2k})$ . It forwards  $(t_1, \dots, t_k, a_1, \dots, a_k)$  as the public key to  $\mathcal{A}$ . In the reduction,  $\mathcal{S}$  must simulate the responses to  $\mathcal{A}$ ’s hash queries and sign queries. It is then shown that if  $(t_1, \dots, t_k, a_1, \dots, a_k)$  follows the R-LWE distribution then the probability that  $\mathcal{S}$  answers *correctly* is close to  $\Pr[\text{forge}(t_1, \dots, t_k, a_1, \dots, a_k)]$ . Furthermore, if  $(t_1, \dots, t_k, a_1, \dots, a_k)$  follows the uniform distribution over  $\mathcal{R}_q^{2k}$  then  $\mathcal{S}$  returns the *wrong* answer with negligible probability.

The proof follows closely the approach proposed in [8], that shows the security of qTESLA’s predecessor TESLA, except for the computation of the two probabilities  $\text{coll}(\vec{a}, \vec{e})$  and



$\text{nwr}(\vec{a}, \vec{e})$  that we define and explain next. For simplicity, we assume that the randomness is sampled uniformly random in  $\mathcal{R}_{[B]}$ . We call a polynomial  $f$  *well-rounded* if  $f \in \mathcal{R}_{[\lfloor q/2 \rfloor - E]}$  and  $[f]_L \in \mathcal{R}_{[2^d - 1 - E]}$ . For our discussion we also define the following sets of polynomials:

$$\begin{aligned} \mathbb{Y} &= \{y \in \mathcal{R}_{[B]}\}, \\ \Delta\mathbb{Y} &= \{y - y' : y, y' \in \mathcal{R}_{[B]}\} = \mathcal{R}_{[2B]}, \\ \Delta\mathbb{L} &= \{x - x' : x, x' \in \mathcal{R} \text{ and } [x]_M = [x']_M\}, \\ \mathbb{H}_{n,h} &= \{c \in \mathcal{R}_{[1]} \mid \|c\|_2 = \sqrt{h}\}, \\ \Delta\mathbb{H}_{n,h} &= \{c - c' : c, c' \in \mathbb{H}_{n,h}\}, \\ \mathbb{W} &= \{[w]_M : w \in \mathcal{R}_q\}. \end{aligned}$$

Moreover, we denote  $\vec{a} = (a_1, \dots, a_k)$  and  $\vec{e} = (e_1, \dots, e_k)$  and define  $\text{nwr}(\vec{a}, \vec{e}) := \Pr_{(y,c) \in \mathcal{R}_{[B]} \times \mathbb{H}_{n,h}} [a_i y - e_i c \text{ not well-rounded for at least one } i \in \{1, \dots, k\}]$ . This quantity varies as a function of  $a_1, \dots, a_k, e_1, \dots, e_k$ . In contrast to [8], we cannot upper bound this in general in the ring setting. Instead, we check experimentally that our acceptance probability for  $w_i$  in line 18 of Algorithm 7 (signature generation) is at least 1/4 for our provably secure parameter sets (see Table 4). Hence,  $\text{nwr}(\vec{a}, \vec{e}) < 3/4$ .

Secondly, we need to bound the probability

$$\text{coll}(\vec{a}, \vec{e}) := \max_{(w_1, \dots, w_k) \in \mathbb{W}^k} \left\{ \Pr_{(y,c) \in \mathcal{R}_{[B]} \times \mathbb{H}_{n,h}} [[a_i y - e_i c]_M = w_i \text{ for } i = 1, \dots, k] \right\}.$$

In [8, Lemma 4] the corresponding probability  $\text{coll}(A, E)$  for standard lattices is upper bounded, given  $A \in \mathbb{Z}_q^{m \times n}$ ,  $E \in \mathbb{Z}_q^{m \times n'}$ , and  $n, m, n' \in \mathbb{Z}$ . Unfortunately, the proof does not carry over to the ring setting for the following reason. In the proof of [8, Lemma 4], it is used that if the randomness  $y \in [-B, B]^n$  is not equal to 0, the vector  $Ay \bmod q$  is uniformly random distributed over  $\mathbb{Z}_q$  and hence also  $Ay - Ec$  is uniformly random distributed over  $\mathbb{Z}_q$ . This does not necessarily hold if the *polynomial*  $y$  is chosen uniformly in  $\mathcal{R}_{[B]}$ . Moreover, in [8, Equation (99)],  $\psi$  denotes the probability that a random vector  $x \in \mathbb{Z}_q^m$  is in  $\Delta\mathbb{L}$ , i.e.,

$$\psi = \Pr_{x \in \mathbb{Z}_q^m} [x \in \Delta\mathbb{L}] \leq \left( \frac{2^d + 1}{q} \right)^m. \quad (8)$$

The quantity  $\psi$  is a function of the TESLA parameters  $q, m, d$ , and it is negligibly small. We cannot prove a similar statement for the signature scheme qTESLA over ideals. Instead, we need to *conjecture* the following.

**Conjecture 5.** *Let  $y$  be a random element of  $\Delta\mathbb{Y}$  and  $a_1, \dots, a_k$  be  $k$  random elements in the ring  $\mathcal{R}_q$ . Then with only negligible probability it will be the case that for each  $i$ , each coefficient of  $a_i y$  is in  $\{-2^d - 2E + 1, \dots, 2^d + 2E - 1\}$ . More formally,*

$$\Pr_{(\vec{a}, y) \leftarrow \mathcal{R}^k \times \mathcal{R}_{[2B]}} [\forall i \in \{1, \dots, k\} : a_i \cdot y \in \mathcal{R}_{[2^d + 2E - 1]}] \leq \frac{1}{2^{n+8\lambda}} \cdot \frac{|\mathbb{H}_{n,h}|}{|\Delta\mathbb{H}_{n,h}|}.$$

We briefly describe why this conjecture is needed in the security reduction for **qTESLA**, and why it should be expected to be true.

This conjecture is needed in bounding the value  $\text{coll}(\vec{a}, \vec{e})$  which corresponds to the maximum likelihood that the values  $a_i y - e_i c \pmod q$  round to some specific values. For example, if during key generation all of the  $a_i$ 's are set to be 0, then the rounding of any  $a_i y - e_i c$  is 0, which is a poor choice of a public key.

In the proof of **TESLA** [8], that such an event is unlikely to occur, the value  $\mathbb{G}(A, E)$  was defined, and a relation was shown between the values  $\text{coll}(A, E)$  and  $\mathbb{G}(A, E)$ . For **qTESLA**, we can define  $\mathbb{G}(\vec{a}, \vec{e})$  as

$$\mathbb{G}(\vec{a}, \vec{e}) = \{(y, c) \in \Delta\mathbb{Y} \times \Delta\mathbb{H}_{n,h} : \forall i, a_i y - e_i c \in \Delta\mathbb{L}\}. \quad (9)$$

A similar relation holds for **qTESLA** so that we can derive a bound on  $\text{coll}(\vec{a}, \vec{e})$  from a bound on  $\mathbb{G}(\vec{a}, \vec{e})$ . Specifically, following the same logic as [8, Lemma 5], we can see that

$$\text{coll}(\vec{a}, \vec{e}) \leq \frac{\mathbb{G}(\vec{a}, \vec{e})}{|\mathbb{Y}| \cdot |\mathbb{H}_{n,h}|}. \quad (10)$$

In fact, we can largely drop the  $\vec{e}$  part of the equation, and simply write  $\mathbb{G}(\vec{a})$ . Because each  $e_i$  is chosen so that  $e_i c$  is always quite small, we can see that the rounding  $[a_i y]_M$  of  $a_i y$  is almost always the same as  $[a_i y - e_i c]_M$ , and ignore the effects of  $e_i c$ . By considering the maximum difference between two elements that round to the same value, we can replace  $\Delta\mathbb{L}$  with  $\mathcal{R}_{[2^d-1]}$ . Then since each coefficient of  $e_i c \in \Delta\mathbb{H}_{n,h}$  is at most  $2E$ , we can see that  $a_i y - e_i c \in \Delta\mathbb{L}$  implies that  $a_i y \in \mathcal{R}_{[2^d+2E-1]}$ . This allows us to define the set  $\mathbb{G}(\vec{a}) = \{y \in \Delta\mathbb{Y} : \forall i, a_i y \in \mathcal{R}_{[2^d+2E-1]}\}$  and establish that  $|\mathbb{G}(\vec{a}, \vec{e})| \leq |\Delta\mathbb{H}_{n,h}| \cdot |\mathbb{G}(\vec{a})|$ .

To demonstrate a bound on the size of  $\mathbb{G}(\vec{a})$ , we calculate the expected value and employ Markov's inequality. Very similarly to the logic of [8, AppendixB.9], we determine that the expected size of  $\mathbb{G}(\vec{a})$  is equal to  $|\Delta\mathbb{Y}|$  times the probability that appears in Equation (5).

If this probability is lower than the bound in Equation (5), then we can employ Markov's inequality to establish that

$$\Pr_{\vec{a}, \vec{e}}[\text{coll}(\vec{a}, \vec{e}) \geq 2^{-\lambda}] \leq 2^{-7\lambda}. \quad (11)$$

The reason we require  $\text{coll}(\vec{a}, \vec{e})$  to be so low is because in the proof, it will be multiplied by factors such as the number of hash function queries squared. We refer to Equation (153) in [8] for how this quantity fits into the proof.

Here we sketch a brief argument as to why this conjecture should be expected to be true. The set  $\mathcal{R}_{q, [2^d+2E-1]}$  forms an incredibly tiny fraction of our ring  $R_q$ . That fraction is

$(2^{d+1} + 4E - 1)^n / q^n$ . For qTESLA-p-I and qTESLA-p-III, it is approximately  $1/2^{5,500}$  and  $1/2^{10,000}$ , respectively. So the closer picking a random  $\vec{a}$  and  $y$  and computing the products is to picking  $k$  uniform elements, the closer we get to this fraction.

For invertible  $y$ , it is easy to see that this corresponds exactly to picking out  $k$  uniform elements, and so the probability is much lower than we need. All that must be accounted for is the non-invertible  $y$ . For these, it should hold that the ideal generated by  $y$  still only has a negligible fraction that is in  $R_{q,[2^d+2E-1]}$ , and indeed it should be the case that only a small part of  $\Delta\mathbb{Y}$  is non-invertible.

To allow experimentation with our parameters, we wrote a script<sup>7</sup> that samples such a  $y$  and  $a$  and checks if their product is in  $\mathcal{R}_{[2^d+2E-1]}$ . After running the script over the parameter sets qTESLA-p-I and qTESLA-p-III 10,000 times each, we did not observe an instance in which a uniform element of  $\mathcal{R}_q$  and  $\mathcal{R}_{[2B]}$  was in  $\mathcal{R}_{[2^d+2E-1]}$ . This supports the claim that our conjecture holds true for the provably secure instantiations of qTESLA.

**Remark 6.** (*Expansion of public keys*) *The explanation above assumes an “expanded” public key  $(a_1, \dots, a_k, t_1, \dots, t_k)$ . In the description of qTESLA, however, the public polynomials  $a_1, \dots, a_k$  are generated from  $\text{seed}_a$  which is part of the secret and public key. This assumption can be justified by another reduction in the QROM: assume there exists an algorithm  $\mathcal{A}$  that breaks the original qTESLA scheme with public key  $(\text{seed}_a, t_1, \dots, t_k)$ . Then we can construct an algorithm  $\mathcal{P}$  that breaks a variant of qTESLA with “expanded” public key  $(a_1, \dots, a_k, t_1, \dots, t_k)$ . To this end, we model  $\text{GenA}(\cdot)$  as a (programmable) random oracle. The algorithm  $\mathcal{P}$  chooses first  $\text{seed}'_a \leftarrow_{\$} \{0, 1\}^{\kappa}$  and reprograms  $\text{GenA}(\text{seed}'_a) = (a_1, \dots, a_k)$ . Afterwards, it forwards  $(\text{seed}'_a, t_1, \dots, t_k)$  as the input tuple to  $\mathcal{A}$ . Quantum queries to  $\text{GenA}(\cdot)$  by  $\mathcal{A}$  can be simulated by  $\mathcal{P}$  according to the construction of Zhandry based on  $2q_h$ -wise independent functions [58]. Hence, the assumption above also holds in the QROM.*

---

<sup>7</sup>The script can be found at `/Supporting_Documentation/Script_for_conjecture/Script_for_experiments_supporting_the_security_conjecture.py` in the qTESLA submission package available at <https://qtesla.org>.