

(One) failure is not an option
**Bootstrapping the search for failures in
lattice-based encryption schemes**

Jan-Pieter D’Anvers^{1*} and Mélissa Rossi^{2,3,4,5†} and Fernando Virdia^{6‡}

¹ imec-COSIC, KU Leuven, Belgium
janpieter.danvers@esat.kuleuven.be

² ANSSI, Paris, France

³ ENS Paris, CNRS, PSL University, Paris, France

⁴ Thales, Gennevilliers, France

⁵ INRIA, Paris, France

melissa.rossi@ens.fr

⁶ Information Security Group, Royal Holloway, University of London, UK
fernando.virdia.2016@rhul.ac.uk

Abstract. Lattice-based encryption schemes are often subject to the possibility of decryption failures, in which valid encryptions are decrypted incorrectly. Such failures, in large number, leak information about the secret key, enabling an attack strategy alternative to pure lattice reduction. Extending the “failure boosting” technique of D’Anvers et al. in PKC 2019, we propose an approach that we call “*directional* failure boosting” that uses previously found “failing ciphertexts” to accelerate the search for new ones. We analyse in detail the case where the lattice is defined over polynomial ring modules quotiented by $\langle X^N + 1 \rangle$ and demonstrate it on a simple Mod-LWE-based scheme parametrized *à la* Kyber768/Saber. We show that for a given secret key (single-target setting), the cost of searching for additional failing ciphertexts after one or more have already been found, can be sped up dramatically. We thus demonstrate that, in this single-target model, these schemes should be designed so that it is hard to even obtain one decryption failure. Besides, in a wider security model where there are many target secret keys (multi-target setting), our attack greatly improves over the state of the art.

Keywords: cryptanalysis, lattice-based cryptography, reaction attacks, decryption errors

*The research of D’Anvers was supported the European Commission through the Horizon 2020 research and innovation programme Cathedral ERC Advanced Grant 695305, by the CyberSecurity Research Flanders with reference number VR20192203 and by the Semiconductor Research Corporation (SRC), under task 2909.001.

†The research of Rossi was supported by the European Union’s H2020 Programme under PROMETHEUS project (grant 780701). It was also supported by the French Programme d’Investissement d’Avenir under national project RISQ P14158.

‡The research of Virdia was supported by the EPSRC and the UK government as part of the Centre for Doctoral Training in Cyber Security at Royal Holloway, University of London (EP/P009301/1).

1 Introduction

Algebraic lattices are a powerful tool in cryptography, enabling the many sophisticated constructions such as digital signatures [37,6], zero-knowledge proofs [39,43], FHE [26] and others. Applications of main interest are public-key encryptions (PKE) [44,38] and key encapsulation mechanisms (KEM).

The computational problems defined over lattices are believed to be hard to solve, even with access to large-scale quantum computers, and hence many of these constructions are considered to be quantum-safe. As industry starts to make steps forward into the concrete development of small quantum computers, the US National Institute of Standards and Technology (NIST) began an open standardization effort, with the aim of selecting quantum-safe schemes for public-key encryption and digital signatures [41]. At the time of writing, the process is in its second round, and 9 out of 17 candidates for PKE or KEM base their security on problems related to lattices, with or without special structure.

One commonly occurring characteristic of lattice-based PKE or KEM schemes is that of lacking perfect correctness. This means that sometimes, ciphertexts generated honestly using a valid public key may lead to decryption failures under the corresponding private key. Throughout this paper we'll refer to such ciphertexts as “failures”, “decryption failures”, or “failing ciphertexts”. While in practice, schemes are parametrised in such a way that decryption failures do not undermine overall performance, these can be leveraged as a vehicle for key recovery attacks against the key pair used to generate them. Such an attack was described by Jaulmes and Joux [31] against NTRU, after which it was extended in [30] and [25]. A similar attack on Ring-LWE based schemes was later presented by Fluhrer [22] and extended by B  etu et. al [5].

However, the aforementioned attacks all use specially crafted ciphertexts and can therefore be prevented with a transformation that achieves chosen ciphertext security. This can for example be obtained by means of an off-the-shelf compiler [23,29] that stops the adversary from being able to freely malleate honestly generated ciphertexts.

The NIST Post-Quantum Standardization Process candidate Kyber [8] noted that it was possible to search for ciphertexts with higher failure probability than average. D’Anvers et al. [16] extended this idea to an attack called “failure boosting”, where ciphertexts with higher failure probability are generated to speedup the search for decryption failures, and provided an analysis of the effectiveness of the attack on several NIST candidates. At the same time, Guo et al. [28] described an adaptive attack against the IND-CCA secure ss-ntru-pke variant of NTRUEncrypt [10], which used an adaptive search for decryption failures exploiting information from previously collected ciphertexts.

Our contributions. In this paper, we present a novel attack technique called “*directional failure boosting*”, aimed at enhancing the search for decryption failures in public-key encryption schemes based on the protocol by Lyubashevsky et al. [38], in the single-target setting. Our technique is an improvement of the “failure boosting” technique of D’Anvers et al. [16].

We consider a simple (but realistically parametrized) scheme based on the Mod-LWE problem as a case study and make some necessary orthogonality and independance assumptions that are reasonable in our range of parameters. We show that in this setting, the work and number of decryption queries needed to obtain multiple failing ciphertexts is only marginally larger than those necessary to obtain the first decryption failure. For example, obtaining 30 decryption failures requires only 25% more quantum work and only 58% more queries than obtaining one decryption failure. As previously shown in [16] and [28], we recall that having many decryption failures enables more efficient lattice reduction which leads to key recovery attacks. As a result, we conclude that when protecting against decryption failure attacks, in the single target setting, designers should make sure that an adversary can not feasibly obtain even a single decryption failure.

Our attack outperforms previously proposed attacks based on decryption failures. In particular, it improves over the multitarget attack of Guo et al. [28] on ss-ntru-pke, lowering the attack’s quantum complexity from $2^{139.5}$ to $2^{96.6}$.

Paper outline. In §2, we introduce some preliminaries about notation and structures. In §3, we describe the general idea of lattice-based encryption and how decryption failures are generated. In §4, we recall the original failure boosting technique from [12]. In §5, we describe our directional failure boosting technique. In §6, we show⁷ how this method impacts the total work and queries overhead. Finally in §7, we discuss the results by comparing them with the literature and conclude with possible future work.

2 Preliminaries

Let \mathbb{Z}_q be the ring of integers modulo q . For N a power of 2, we define R_q the ring $\mathbb{Z}_q[X]/(X^N + 1)$, and $R_q^{l_1 \times l_2}$ the ring of $l_1 \times l_2$ matrices over R_q . Vectors and polynomials will be indicated with bold lowercase letters, eg. \mathbf{v} , while matrices will be written in bold uppercase letters, eg. \mathbf{M} . Denote with $\lfloor \cdot \rfloor$ flooring to the nearest lower integer, and with $\lceil \cdot \rceil$ rounding to the nearest integer. These operations are extended coefficient-wise for vectors and polynomials. Throughout, we abuse notation and identify elements in \mathbb{Z}_q with their representatives in $[-q/2, q/2)$, and elements in R_q with their representatives of degree $< N$, with index i indicating the coefficient of X^i . This allows us to define the ℓ_2 -norm $\|\mathbf{x}\|_2$ of a polynomial $\mathbf{x} \in R_q$, so that $\|\mathbf{x}\|_2 = \sqrt{\sum_i \mathbf{x}_i^2}$ where $\mathbf{x}_i \in [-q/2, q/2)$, and extend this to vectors of polynomials $\mathbf{y} \in R_q^{l \times 1}$ as $\|\mathbf{y}\|_2 = \sqrt{\sum_i \|\mathbf{y}_i\|_2^2}$. Identically, we define and extend the ℓ_∞ -norm.

Let $x \leftarrow X$ denote sampling x according to the probability distribution X . We extend this notation for coefficient-wise sampling of a vector $\mathbf{x} \in R_q^{l \times 1}$ as $\mathbf{x} \leftarrow X(R_q^{l \times 1})$, and similarly for a matrix. We denote with $\mathbf{x} \leftarrow X(R_q^{l \times 1}; r)$ sampling $\mathbf{x} \in R_q^{l \times 1}$ pseudorandomly from the seed r with each coefficient following

⁷The software is available at: <https://github.com/KULeuven-COSIC/PQCRYPTO-decryption-failures>.

the distribution X . In algorithms, we also use $x \leftarrow \text{Alg}()$ to mean that the value x is assigned to be the output of a probabilistic algorithm Alg .

Let \mathcal{U} be the uniform distribution over \mathbb{Z}_q and let $\mathcal{N}_{\mu,\sigma}$ be the normal distribution with mean μ and standard deviation σ , so that the probability density function of $x \leftarrow \mathcal{N}_{\mu,\sigma}$ is defined as:

$$f_{\mathcal{N}_{\mu,\sigma}}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/2\sigma^2}. \quad (1)$$

The discrete Gaussian distribution $\mathcal{D}_{\mu,\sigma}$ is a discrete restriction to \mathbb{Z}_q of $\mathcal{N}_{\mu,\sigma}$, so that an integer x is sampled with a probability proportional to $e^{-(x-\mu)^2/2\sigma^2}$ and its remainder modulo q in $[-q/2, q/2)$ is returned.

For an event A we define $P[A]$ as its probability. For an element which does not correspond to an event, a ciphertext ct for example, we abusively write $P[ct]$ to denote the probability of the event $ct' = ct$ where ct' is drawn from a distribution which will be clear in the context. We will denote with $\mathbb{E}[A]$ the expected value of a variable drawn from a distribution A .

Security definitions. Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme, with message space \mathcal{M} , and let $K = (\text{KeyGen}, \text{Encaps}, \text{Decaps})$ be a key encapsulation mechanism (KEM). When a decapsulation or a decryption oracle is provided, we assume that the maximum number of ciphertexts that can be queried to it for each key pair is 2^K ; in practice, $K = 64$ is often considered [41, §4.A.2]. In this work, we keep the maximum number of queries as a parameter with no specific value, in order to provide a better granularity in the security assessment. Indeed, to mount an attack, the adversary trades off between number of queries and the work.

Definition 1 (IND-CPA $_{A,\Pi}(k)$ game [34]). Let A be an adversary and $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme. The experiment $\text{IND-CPA}_{A,\Pi}(1^k)$ runs as follows:

1. $(pk, sk) \leftarrow \text{KeyGen}(1^k)$
2. A is given pk . After evaluating $\text{Enc}(pk, \cdot)$ as desired, it outputs $(m_0, m_1) \in \mathcal{M} \times \mathcal{M}$.
3. A random bit $b \leftarrow_{\S} \{0, 1\}$ is sampled, and $c \leftarrow \text{Enc}(pk, m_b)$ is passed to A .
4. A keeps evaluating $\text{Enc}(pk, \cdot)$ as desired, until it returns a bit b' .
5. The experiment outputs 1 if $b = b'$ and 0 otherwise.

Definition 2 (IND-CCA $_{A,K}(k)$ game [34]). Let A be an adversary and $K = (\text{KeyGen}, \text{Encaps}, \text{Decaps})$ be a key encapsulation mechanism. The experiment $\text{IND-CCA}_{A,K}(1^k)$ runs as follows:

1. $(pk, sk) \leftarrow \text{KeyGen}(1^k)$
2. $(c, k) \leftarrow \text{Encaps}(pk)$
3. $b \leftarrow_{\S} \{0, 1\}$. If $b = 0$, set $\hat{k} = k$, else let $\hat{k} \leftarrow \{0, 1\}^n$.
4. A is given (pk, c, \hat{k}) , and access to a decapsulation oracle $\text{Decaps}(sk, \cdot)$. After evaluating $\text{Encaps}(pk, \cdot)$ and querying $\text{Decaps}(sk, \cdot)$ as desired (except for decapsulation queries on c), it returns $b' \in \{0, 1\}$.

5. The experiment outputs 1 if $b = b'$ and 0 otherwise.

Definition 3 (PKE and KEM security [24]). A public-key encryption scheme Π (resp. a key encapsulation mechanism K) is (t, ϵ) -GAME secure if for every t -time adversary A , we have that

$$\left| \Pr[\text{GAME}_{A,\Pi}(k) = 1] - \frac{1}{2} \right| \leq \epsilon \quad \left(\text{resp.} \quad \left| \Pr[\text{GAME}_{A,K}(k) = 1] - \frac{1}{2} \right| \leq \epsilon \right)$$

For a security parameter 1^k , we usually mean $t \approx \text{poly}(k)$ and $\epsilon \leq \text{negl}(k)$. If GAME is IND-CPA (resp. IND-CCA) we say that Π (resp. K) is (t, ϵ) -secure against chosen-plaintext attacks (resp. (t, ϵ) -secure against adaptive chosen-ciphertext attacks).

3 Lattice-based encryption

The Module-LWE (or Mod-LWE) problem [35] is a mathematical problem that can be used to build cryptographic primitives such as encryption [7,13], key exchange [13] and signatures [20]. It is a generalization of both the Learning With Errors (or LWE) problem [44], and the Ring-LWE problem [48,38].

Definition 4 (Mod-LWE [35]). Let n, q, k be positive integers, χ be a probability distribution on \mathbb{Z} and \mathbf{s} be a secret module element in R_q^k . We denote by \mathcal{L} the probability distribution on $R_q^k \times R_q$ obtained by choosing $\mathbf{a} \in R_q^k$ uniformly at random, choosing $e \in R$ by sampling each of its coefficients according to χ and considering it in R_q , and returning $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in R_q^k \times R_q$.

Decision-Mod-LWE is the problem of deciding whether pairs $(\mathbf{a}, c) \in R_q^k \times R_q$ are sampled according to \mathcal{L} or the uniform distribution on $R_q^k \times R_q$.

Search-Mod-LWE is the problem of recovering \mathbf{s} from $(\mathbf{a}, c) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in R_q^k \times R_q$ sampled according to \mathcal{L} .

3.1 Passively and actively secure encryption

Lyubashevsky et al. [38] introduced a simple protocol to build passively secure encryption from the Ring-LWE problem, inspired by Diffie-Hellman key exchange [19] and ElGamal public-key encryption [21]. Naturally, the protocol can also be adapted to work based on plain and Module LWE assumptions. A general extension of the protocol for all aforementioned assumptions is described in Algorithms 1, 2, and 3, where $r \in \mathcal{R} = \{0, 1\}^{256}$, and where the message space is defined as $\mathcal{M} = \{\text{polynomials in } R_q \text{ with coefficients in } \{0, 1\}\}$.

In order to obtain active security, designers usually use an off-the-shelf CCA compiler, usually a (post-quantum) variant [18,29,49,45,32] of the Fujisaki-Okamoto transform [23] (FO). These come with proofs of security in the (quantum) random oracle model, with explicit bounds about the loss of security caused by the transformation. We show such transformed KEM Decapsulation and Encapsulation in Algorithms 4 and 5.

Algorithm 1: PKE.KeyGen()	Algorithm 2: PKE.Enc($pk = (\mathbf{b}, \mathbf{A})$, $\mathbf{m} \in \mathcal{M}; r$)
<ol style="list-style-type: none"> 1 $\mathbf{A} \leftarrow \mathcal{U}(R_q^{l \times l})$ 2 $\mathbf{s}, \mathbf{e} \leftarrow \mathcal{D}_{0, \sigma_s}(R_q^{l \times 1}) \times \mathcal{D}_{0, \sigma_e}(R_q^{l \times 1})$ 3 $\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e}$ 4 return ($pk = (\mathbf{b}, \mathbf{A})$, $sk = \mathbf{s}$) 	<ol style="list-style-type: none"> 1 $\mathbf{s}', \mathbf{e}' \leftarrow \mathcal{D}_{0, \sigma_s}(R_q^{l \times 1}; r) \times \mathcal{D}_{0, \sigma_e}(R_q^{l \times 1}; r)$ 2 $\mathbf{e}'' \leftarrow \mathcal{D}_{0, \sigma_e}(R_q; r)$ 3 $\mathbf{b}' := \mathbf{A}^T \mathbf{s}' + \mathbf{e}'$ 4 $\mathbf{v}' := \mathbf{b}^T \mathbf{s}' + \mathbf{e}'' + \lfloor q/2 \rfloor \cdot \mathbf{m}$ 5 return $ct = (\mathbf{v}', \mathbf{b}')$
Algorithm 3: PKE.Dec($sk = \mathbf{s}$, $ct = (\mathbf{v}', \mathbf{b}')$)	
<ol style="list-style-type: none"> 1 $\mathbf{m}' := \lfloor \lfloor 2/q \rfloor (\mathbf{v}' - \mathbf{b}'^T \mathbf{s}) \rfloor$ 2 return \mathbf{m}' 	
Algorithm 4: KEM.Encaps(pk)	
<ol style="list-style-type: none"> 1 $m \leftarrow \mathcal{U}(\{0, 1\}^{256})$ 2 $(\bar{K}, r) := \mathcal{G}(pk, m)$ 3 $ct := \text{PKE.Enc}(pk, m, r)$ 4 $K := \mathcal{H}(\bar{K}, r)$ 5 return (ct, K) 	

In the case of FO for lattice-based schemes, the randomness used during the encryption is generated by submitting the message (and sometimes also the public key) to a random oracle. As this procedure is repeatable with knowledge of the message, one can check the validity of ciphertexts during decapsulation. Hence, an adversary wanting to generate custom ephemeral secrets \mathbf{s}' , \mathbf{e}' , \mathbf{e}'' in order to fabricate weak ciphertexts, would need to know a preimage of the appropriate random coins for the random oracle. Therefore, their only option is to mount a (Grover's) search by randomly generating ciphertexts corresponding to different messages \mathbf{m} , and testing if their predicted failure probability is above a certain threshold.

Remark 1. Several lattice-based candidates submitted to the NIST Post-Quantum Cryptography Standardization Process use a variant of the protocol by Lyubashevsky et al. [38]. Deviating from the original design, most candidates perform an additional rounding of the ciphertext \mathbf{v}' , in order to reduce bandwidth. The designers of New Hope [3] and LAC [36] choose to work directly over rings (or equivalently, they choose a module of rank $l = 1$) and add error correction on the encapsulated message, while the designers of Kyber [7] and Saber [13] choose a module of rank $l > 1$ and perform an additional rounding of \mathbf{b}' (and \mathbf{b} in case of Saber). We here focus on the basic version given in Algorithms 1 to 3 and leave the study of the effect of compression to further work.

We selected the parameters of the studied encryption scheme to ensure a similar failure probability and security to Kyber and Saber. These parameters can be found in Table 1. The security estimates are generated using the Core-

Algorithm 5: $\text{KEM.Decaps}(sk, pk, ct, K)$

```

1  $m' := \text{PKE.Dec}(sk, ct)$ 
2  $(\bar{K}, r') := \mathcal{G}(pk, m')$ 
3  $ct' := \text{PKE.Enc}(pk, m'; r')$ 
4 if  $ct = ct'$  then
5 |   return  $K := (\bar{K}, r')$ 
6 else
7 |   return  $K := \perp$  // Could return a pseudo-random string to implicitly
   |   reject

```

	l	N	q	σ_s	σ_e	$P[F]$	Classical	Quantum
Chosen parameters	3	256	8192	2.00	2.00	2^{-119}	2^{195}	2^{177}
Saber	3	256	8192	1.41	2.29	2^{-136}	2^{203}	2^{185}
Kyber 768	3	256	3329	1.00	$1.00/1.38^\dagger$	2^{-164}	2^{181}	2^{164}

[†] Standard deviation of the error term in the public key and ciphertext respectively

Table 1. Comparison between our target scheme and Saber and Kyber 768, as parametrised in Round 2 of the NIST PQC standardization process. The classical (resp. quantum) security is evaluated using the Core-SVP [3] methodology, assuming the cost of BKZ with block size β to be $2^{0.292\beta}$ (resp. $2^{0.265\beta}$).

SVP methodology [3] and the LWE estimator⁸[2], while the failure probability of Kyber and Saber is given as reported in their respective the NIST round 2 documentations [47,14]. The failure probability of our chosen parameters is determined by calculating the variance of the error term and assuming the distribution to be Gaussian.

Remark 2. We do not consider the case of “plain” LWE based schemes like FrodoKEM [40] or Round5 [4]. Nonetheless, we believe that the attack methodology would easily translate to the LWE setting as the failure condition and the failure probabilities are similar to the investigated case.

3.2 Decryption failures

Following the execution of the protocol, both messages \mathbf{m}' and \mathbf{m} are the same if the coefficients of the error term $\mathbf{e}^T \mathbf{s}' - \mathbf{s}^T \mathbf{e}' + \mathbf{e}''$ are small enough; more exactly if $\|\mathbf{e}^T \mathbf{s}' - \mathbf{s}^T \mathbf{e}' + \mathbf{e}''\|_\infty \leq q/4$. This expression can be simplified by defining the vector \mathbf{S} as the vertical concatenation of $-\mathbf{s}$ and \mathbf{e} , the vector \mathbf{C} as the concatenation of \mathbf{e}' and \mathbf{s}' , and by replacing \mathbf{e}'' with \mathbf{G} , as shown below:

$$\mathbf{S} = \begin{bmatrix} -\mathbf{s} \\ \mathbf{e} \end{bmatrix} \quad \mathbf{C} = \begin{bmatrix} \mathbf{e}' \\ \mathbf{s}' \end{bmatrix} \quad \mathbf{G} = \mathbf{e}'' . \quad (2)$$

⁸The estimator can be found at <https://bitbucket.org/malb/lwe-estimator/>.

Here, \mathbf{S} contains the secret elements of the secret key, and \mathbf{C} and \mathbf{G} consist of elements used to construct the ciphertexts⁹. Using these vectors, the error expression can be rewritten: a failure occurs when $\|\mathbf{S}^T \mathbf{C} + \mathbf{G}\|_\infty > q/4$.

The standard deviation of the terms in the polynomial $\mathbf{S}^T \mathbf{C}$ equals $\sqrt{2N}\sigma_s\sigma_e$, versus a standard deviation of σ_e for the terms of \mathbf{G} . Therefore, the influence of \mathbf{G} on the failure rate is limited, i.e. $\|\mathbf{S}^T \mathbf{C} + \mathbf{G}\|_\infty \approx \|\mathbf{S}^T \mathbf{C}\|_\infty$. Let $q_t := q/4$ denote the failure threshold, we will use

$$\|\mathbf{S}^T \mathbf{C}\|_\infty > q_t \quad (3)$$

as an approximation of the failure expression throughout our analysis. However, with some extra work, one can rewrite a more accurate Equation 3 as $\|\mathbf{S}^T \mathbf{C}\|_\infty > q_t - \|\mathbf{G}\|_\infty$, and instead of considering q_t to be fixed, taking the distribution of $q_t - \|\mathbf{G}\|_\infty$ as shown in [16]. For the ease of the implementation and due to the low influence of \mathbf{G} on the failure rate, we prefer to stick with Equation 3. We now introduce a more handy way of writing the failure condition (Equation 3) by only using vectors in \mathbb{Z}_q .

Definition 5 (Coefficient vector). For $\mathbf{S} \in R_q^{l \times 1}$, we denote by $\bar{\mathbf{S}} \in \mathbb{Z}_q^{lN \times 1}$, the representation of \mathbf{S} where each polynomial is decomposed as a list of its coefficients in¹⁰ \mathbb{Z}_q .

Definition 6 (Rotations). For $r \in \mathbb{Z}$ and $\mathbf{C} \in R_q^{l \times 1}$, we denote by $\mathbf{C}^{(r)} \in R_q^{l \times 1}$, the following vector of polynomials

$$\mathbf{C}^{(r)} := X^r \cdot \mathbf{C}(X^{-1}) \pmod{X^N + 1}.$$

Correspondingly, $\overline{\mathbf{C}^{(r)}} \in \mathbb{Z}_q^{lN \times 1}$ denotes its coefficient vector.

It is easy to show that $\overline{\mathbf{C}^{(r)}}$ is constructed as to ensure that for $r \in [0, \dots, N-1]$, the r^{th} coordinate of $\mathbf{S}^T \mathbf{C}$ is given by the scalar product $\bar{\mathbf{S}}^T \overline{\mathbf{C}^{(r)}}$. In other words, one is now able to decompose $\mathbf{S}^T \mathbf{C}$ as a sum of scalar products:

$$\mathbf{S}^T \mathbf{C} = \sum_{r \in [0, N-1]} \bar{\mathbf{S}}^T \overline{\mathbf{C}^{(r)}} \cdot X^r. \quad (4)$$

One can observe that this construction is only valid for the modulo $X^N + 1$ ring structure, but it could be adapted for other ring structures. Note that for any $r \in \mathbb{Z}$, $\mathbf{C}^{(r+N)} = -\mathbf{C}^{(r)}$ and $\mathbf{C}^{(r+2N)} = \mathbf{C}^{(r)}$. Besides, taking into account the extension of the norms to vectors of polynomials (defined in §2), one can make the following remark.

⁹When talking about ciphertexts throughout the paper, we will sometimes refer to their underlying elements \mathbf{C} and \mathbf{G} .

¹⁰Recall that, in this paper, all the elements in \mathbb{Z}_q are represented as integers belonging in $[-q/2, q/2]$.

Remark 3. Note that for any $r \in \mathbb{Z}$, $\|\overline{\mathbf{C}}^{(r)}\|_2 = \|\overline{\mathbf{C}}\|_2 = \|\mathbf{C}\|_2$ and $\|\overline{\mathbf{C}}^{(r)}\|_\infty = \|\overline{\mathbf{C}}\|_\infty = \|\mathbf{C}\|_\infty$.

The decomposition in Equation 4 will allow a geometric interpretation of the failures as it will be shown in the rest of the paper. First, let us introduce a brief example to illustrate Definitions 5 and 6.

Example 1. For a secret \mathbf{S} and a ciphertext \mathbf{C} in $\mathbb{Z}_q^{2 \times 1}[X]/(X^3 + 1)$:

$$\mathbf{S} = \begin{bmatrix} s_{0,0} + s_{0,1}X + s_{0,2}X^2 \\ s_{1,0} + s_{1,1}X + s_{1,2}X^2 \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} c_{0,0} + c_{0,1}X + c_{0,2}X^2 \\ c_{1,0} + c_{1,1}X + c_{1,2}X^2 \end{bmatrix} \quad (5)$$

we get the following vectors:

$$\overline{\mathbf{S}} = \begin{bmatrix} s_{0,0} \\ s_{0,1} \\ s_{0,2} \\ s_{1,0} \\ s_{1,1} \\ s_{1,2} \end{bmatrix}, \quad \overline{\mathbf{C}}^{(0)} = \begin{bmatrix} c_{0,0} \\ -c_{0,2} \\ -c_{0,1} \\ c_{1,0} \\ -c_{1,2} \\ -c_{1,1} \end{bmatrix}, \quad \overline{\mathbf{C}}^{(1)} = \begin{bmatrix} c_{0,1} \\ c_{0,0} \\ -c_{0,2} \\ c_{1,1} \\ c_{1,0} \\ -c_{1,2} \end{bmatrix}, \quad \overline{\mathbf{C}}^{(2)} = \begin{bmatrix} c_{0,2} \\ c_{0,1} \\ c_{0,0} \\ c_{1,2} \\ c_{1,1} \\ c_{1,0} \end{bmatrix}, \quad \overline{\mathbf{C}}^{(3)} = \begin{bmatrix} -c_{0,0} \\ c_{0,2} \\ c_{0,1} \\ -c_{1,0} \\ c_{1,2} \\ c_{1,1} \end{bmatrix} \dots$$

In case of a failure event, $\mathbf{S}^T \mathbf{C}$ satisfies Equation 3. Therefore, at least one element among all the coefficients

$$\overline{\mathbf{S}}^T \cdot \overline{\mathbf{C}}^{(0)}, \dots, \overline{\mathbf{S}}^T \cdot \overline{\mathbf{C}}^{(2N-1)}$$

is larger than q_t .

Definition 7 (Failure event). A failure event will be denoted with F , while we use S to indicate a successful decryption.

More precisely, for $r \in [0, 2N - 1]$, we denote by F_r the failure event where

$$\overline{\mathbf{S}}^T \cdot \overline{\mathbf{C}}^{(r)} > q_t.$$

The event F_r gives a twofold information: it provides the location of the failure in the $\mathbf{S}^T \mathbf{C}$ polynomial and it also provides the sign of the coefficient that caused the failure.

An assumption on the failing ciphertexts. In the rest of the paper, in order to predict the results of our attack, we will make the following orthogonality assumption.

Assumption 1 Let $n \ll 2Nl$, and $\mathbf{C}_0, \dots, \mathbf{C}_n$ be ciphertexts that lead to failure events F_{r_0}, \dots, F_{r_n} . The vectors $\overline{\mathbf{C}}_0^{(r_0)}, \dots, \overline{\mathbf{C}}_n^{(r_n)}$ are considered orthogonal when projected on the hyperplane orthogonal to $\overline{\mathbf{S}}$.

This assumption is an approximation that is supported by the fact that vectors in high dimensional space have a strong tendency towards orthogonality, as can be seen in Figure 2.

4 Failure Boosting attack technique

Failure boosting is a technique introduced in [16] to increase the failure rate of (Ring/Mod)-LWE/LWR based schemes by honestly generating ciphertexts and only querying weak ones, i.e. those that have a failure probability above a certain threshold $f_t > 0$. This technique is especially useful in combination with Grover’s algorithm [27], in which case the search for weak ciphertexts can be sped up quadratically. Failure boosting consists of two phases: a precomputation phase, and a phase where the decryption oracle is queried.

Precomputation phase. The adversary does an offline search for weak ciphertexts with the following procedure:

1. Generate a key encapsulation⁹ $ct = (\mathbf{C}, \mathbf{G})$.
2. If $P[F | ct] \geq f_t$, keep ct in a weak ciphertext list, otherwise go to Step 1.

In Step 2, $P[F | ct]$ is defined as the failure probability given a certain ciphertext ct . It is computed as follows.

$$P[F | ct] := \sum_{\mathbf{S}} P [\|\mathbf{S}^T \mathbf{C} + \mathbf{G}\|_{\infty} > q_t | \mathbf{S}] \cdot P[\mathbf{S}] \quad (6)$$

Given the probability of generating ciphertexts $P[ct] = P[\mathbf{C}, \mathbf{G}]$, the probability of finding such a weak ciphertext can be expressed as follows:

$$\alpha_{f_t} = \sum_{\forall ct: P[F|ct] > f_t} P[ct]. \quad (7)$$


An adversary thus needs to perform on average $\alpha_{f_t}^{-1}$ work to obtain one weak ciphertext, or $\sqrt{\alpha_{f_t}^{-1}}$ assuming Grover’s search achieves a full speed-up.

Decryption oracle query phase. After the precomputation phase, an adversary has a probability β_{f_t} that a weak ciphertext results in a failure, where β_{f_t} can be calculated as a weighted average of the failure probabilities of weak ciphertexts:

$$\beta_{f_t} = \frac{\sum_{\forall ct: P[F|ct] > f_t} P[ct] \cdot P[F|ct]}{\sum_{\forall ct: P[F|ct] > f_t} P[ct]}. \quad (8)$$

Thus to obtain one decryption failure with probability $1 - e^{-1}$, an adversary needs to perform approximately $\beta_{f_t}^{-1}$ queries and therefore $\alpha_{f_t}^{-1} \beta_{f_t}^{-1}$ work (or $\sqrt{\alpha_{f_t}^{-1} \beta_{f_t}^{-1}}$ using a quantum computer).

The better an adversary can predict $P[F|ct]$, the more efficient failure boosting will be. Having no information about the secret except its distribution, an adversary is bound to standard failure boosting, where the failure probability is estimated based on $\|\mathbf{C}\|_2$ and $\|\mathbf{G}\|_2$. For a graphical intuition, a two dimensional toy example is depicted in Figure 1a below, where the red arrow represents the



(a) Without directional information, as in [16], the weak ciphertexts (in blue) are defined as the ciphertexts with a probability higher than f_t .

(b) With directional information, the weak ciphertexts (in blue) are found according to a refined acceptance criterion, here represented as an ellipse.

Fig. 1. Simplified diagram trying to provide an intuition on the effect of directional failure boosting. The red arrow represents the secret vector $\bar{\mathbf{S}}$. Ciphertexts with $\bar{\mathbf{C}}$ that lie in the dashed area will provoke a failure as the inner product with $\bar{\mathbf{S}}$ will exceed the threshold q_t . Ciphertexts outside the blue circle are considered weak.

secret vector $\bar{\mathbf{S}}$. Ciphertexts with $\bar{\mathbf{C}}$ that lie in the dashed area will provoke a failure as the inner product with $\bar{\mathbf{S}}$ will exceed the threshold q_t . The blue circle is a circle of ciphertexts that have a certain failure probability f_t as estimated by an adversary who does not know the secret. During the failure boosting procedure, we will generate random ciphertexts, and only select the ciphertexts with a higher failure probability than f_t , i.e. that are outside the blue circle. One can graphically see in Figure 1a that these ciphertexts will have a higher failure probability and a higher norm. We refer to [16] for a full description of the failure boosting technique. Note that Figure 1a is an oversimplified 2-dimension example that does not take into account the polynomial structure and the high dimensionality of the space.

5 Directional Failure Boosting

Once $n \geq 1$ decryption failures $\mathbf{C}_0, \dots, \mathbf{C}_{n-1}$ are found, additional information about the secret key \mathbf{S} becomes available, and can be used to refine the failure estimation for new ciphertexts and thus speed up failure boosting. We now introduce an iterative two-step method to perform directional failure boosting.

- Step 1 An estimate, denoted $\bar{\mathbf{E}}$, of the ‘direction’ of the secret $\bar{\mathbf{S}}$ in \mathbb{Z}_q^{LN} is obtained from $\mathbf{C}_0, \dots, \mathbf{C}_{n-1}$.
- Step 2 The estimate $\bar{\mathbf{E}}$ is used to inform the search for weak ciphertexts and improve the failure probability prediction for a new ciphertext \mathbf{C}_n . One is able to refine the criterion $P[F | ct] \geq f_t$ with computing $P[F | ct, \bar{\mathbf{E}}] \geq f_t$ instead.

Once new failing ciphertexts are found in step 2, one can go back to step 1 and improve the estimate $\bar{\mathbf{E}}$ and thus bootstrap the search for new failures.

To give an intuition, a two dimensional toy representation can be found in Figure 1b. Like in the classical failure boosting technique, the red arrow depicts the secret $\bar{\mathbf{S}}$, while the additional blue arrow marks estimate $\bar{\mathbf{E}}$ (as calculated in step 1, see §5.2). Using this estimate, we can refine the acceptance criterion to the depicted ellipse to better reflect our knowledge about the secret (step 2, see §5.3). Ciphertexts outside this ellipse will be flagged as weak ciphertexts, and while the probability of finding such a ciphertext is the same, the failure probability of weak ciphertexts is now higher. As in, more of the blue zone lies in the dashed area.

5.1 Distributions

We now introduce some probability distributions that will be useful in following sections.

Scaled χ -distribution. The scaled χ -distribution $\chi_{n,\sigma}$ is the distribution of the ℓ_2 -norm of a vector with n coefficients, each following the normal distribution $\mathcal{N}_{0,\sigma}$. Denoting with Γ the gamma function, the probability density function of $\chi_{n,\sigma}$ is given by:

$$f_{\chi_{n,\sigma}}(x) = \frac{\left(\frac{x}{\sigma}\right)^{n-1} e^{-\frac{x^2}{2\sigma^2}}}{2^{\left(\frac{n}{2}-1\right)} \Gamma\left(\frac{n}{2}\right)} \quad \text{for } x \geq 0, \quad (9)$$

which has mean [33, §18.3] $\mathbb{E}_{\chi}[x] = \sqrt{2} \frac{\Gamma((n+1)/2)}{\Gamma(n/2)} \sigma \approx \sqrt{n} \sigma$.

We will approximate the probability distribution of $\|\mathbf{x}\|_2$ where $\mathbf{x} \leftarrow \mathcal{D}_{0,\sigma}(R_q^{l \times 1})$ with a discretized version of the $\chi_{(l \cdot N),\sigma}$ -distribution, which will be denoted with $\chi_{(l \cdot N),\sigma}^D$. Using this distribution, the probability density function of $\|\mathbf{x}\|_2$ is calculated as:

$$P[\|\mathbf{x}\|_2 = x] = C \cdot \left(\frac{x}{\sigma}\right)^{l \cdot N - 1} e^{-\frac{x^2}{2\sigma^2}} \quad \text{for } x \in \left\{0, \dots, \left\lfloor \frac{q}{2} \sqrt{lN} \right\rfloor\right\}, \quad (10)$$

with C a normalization constant.

Angle distribution. The distribution of angles between n -dimensional vectors in \mathbb{R}^n with coefficients drawn from a normal distribution $\mathcal{N}_{0,\sigma}$ can be modelled using the following probability density function [9]:

$$f_{\Theta_n}(\theta) = \sin^{n-2}(\theta) / \int_0^\pi \sin^{n-2}(t) dt, \quad \text{for } \theta \in [0, \pi]. \quad (11)$$

Due to the high dimensionality of the vector space used in this paper, vectors will have a very strong tendency towards orthogonality, i.e. θ is close to $\pi/2$, as can be seen in Figure 2.

For computational reasons, we will use a discretized version Θ_n^D of this distribution to model the distribution of the angles between discrete vectors, if no




Fig. 2. Probability density function (pdf) of the angle between two random vectors in 1536-dimensional space. As the dimension increases, the pdf tends to the Dirac delta function centered at $\frac{\pi}{2}$.

extra directional information is present. Given a uniformly spaced list of angles between 0 and π , we assign to each angle a probability

$$P[\theta] = C \sin^{n-2}(\theta) \quad (12)$$

with C a normalization constant. The higher the number of angles in this list, the better this distribution approximates the continuous distribution Θ_n .

Order statistics. The maximal order statistic of a distribution X in n dimensions, is the distribution of the maximum of n samples drawn from this distribution. We will denote this distribution with $M(X, n)$. For a discrete distribution X , the probability mass function of $M(X, n)$ can be computed as:

$$f_{M(X,n)}(x) = P[x \geq y | y \leftarrow X]^n - P[x > y | y \leftarrow X]^n \quad (13)$$

$$\approx n \cdot P[x = y | y \leftarrow X] \cdot P[x > y | y \leftarrow X]^{n-1}, \quad (14)$$

where the latter approximation gets better for smaller probabilities.

5.2 Step 1: Estimating the direction $\bar{\mathbf{E}}$

Informally, $\bar{\mathbf{E}}$ should be a vector that has approximately the same direction as $\bar{\mathbf{S}}$. Denoting the angle between $\bar{\mathbf{E}}$ and $\bar{\mathbf{S}}$ as θ_{ES} , the bigger $|\cos(\theta_{ES})|$, the closer our estimate is to $\pm\bar{\mathbf{S}}$ and the better our estimate of failure probability will be. Since we focus on estimating the direction of $\bar{\mathbf{S}}$, $\bar{\mathbf{E}}$ will always be normalized.

In this section, we derive an estimate $\bar{\mathbf{E}}$ of the direction of the secret $\bar{\mathbf{S}}$ given $n \geq 1$ ciphertexts $\mathbf{C}_0, \dots, \mathbf{C}_{n-1}$. Our goal is to find $\bar{\mathbf{E}}$ such that $|\cos(\theta_{ES})|$ is as big as possible. We will first discuss the case where the adversary has one ciphertext, then the case where she has two, followed by the more general case where she has n ciphertexts.

One ciphertext. Assume that a unique failing ciphertext \mathbf{C} is given. For a failure event F_r , $\bar{\mathbf{E}} = \overline{\mathbf{C}^{(r)}} / \|\overline{\mathbf{C}^{(r)}}\|_2$ is a reasonable choice as $\cos(\theta_{ES})$ is bigger than average. This can be seen as follows:

$$|\cos(\theta_{ES})| = \frac{|\bar{\mathbf{S}}^T \cdot \bar{\mathbf{E}}|}{\|\bar{\mathbf{S}}\|_2 \|\bar{\mathbf{E}}\|_2} = \frac{|\bar{\mathbf{S}}^T \cdot \overline{\mathbf{C}^{(r)}}|}{\|\bar{\mathbf{S}}\|_2 \|\overline{\mathbf{C}^{(r)}}\|_2} > \frac{q_t}{\|\bar{\mathbf{S}}\|_2 \|\overline{\mathbf{C}^{(r)}}\|_2}. \quad (15)$$

Keep in mind that the cosine of angles between random vectors strongly tend to zero in high dimensional space, so that even a relatively small value of $|\cos(\theta_{ES})|$ might be advantageous.

One can argue that it is not possible to compute $\overline{\mathbf{C}^{(r)}}$ without knowledge of r ; whereas in the general case, the failure location is unknown. However, $\overline{\mathbf{E}} = \overline{\mathbf{C}^{(0)}} / \|\overline{\mathbf{C}^{(0)}}\|_2$ is an equally good estimate regardless of the value of r . Indeed, $\overline{\mathbf{C}^{(0)}}$ approximates a rotation of the secret $\overline{\mathbf{S}'} := \overline{X^{-r} \cdot \mathbf{S}}$ instead of $\overline{\mathbf{S}}$, which can be seen using the equality $\overline{\mathbf{A}}^T \cdot \overline{\mathbf{B}} = \overline{X^i \mathbf{A}^T \cdot X^i \mathbf{B}}$:

$$\begin{aligned} \overline{\mathbf{S}}^T \cdot \overline{\mathbf{C}^{(r)}} &= \overline{X^{-r} \cdot \mathbf{S}^T \cdot X^{-r} X^r \mathbf{C}^{(0)}} \\ &= \overline{X^{-r} \cdot \mathbf{S}^T \cdot \mathbf{C}^{(0)}}. \end{aligned} \quad (16)$$

Furthermore, multiplying a polynomial in R_q with a power of X does not change its infinity norm, as the multiplication only results in the rotation or negation of coefficients. Thus, using an estimate of the direction of $\overline{X^{-r} \cdot \mathbf{S}}$ is as good as an estimate of the direction of $\overline{\mathbf{S}}$ when predicting the failure probability of ciphertexts, and we can use $\overline{\mathbf{E}} = \overline{\mathbf{C}^{(0)}} / \|\overline{\mathbf{C}^{(0)}}\|_2$.

Two ciphertexts. Now, assume that two linearly independent failing ciphertexts \mathbf{C}_0 and \mathbf{C}_1 , resulting from failure events F_{r_0} and F_{r_1} respectively, are given. Taking $\overline{\mathbf{E}}$ as the normalized version of an average $\overline{\mathbf{C}_{\text{av}}} = (\overline{\mathbf{C}_0^{(0)}} + \overline{\mathbf{C}_1^{(0)}}) / 2$ may not necessarily result in a good estimate. For example, if \mathbf{C}_0 comes from a failure event F_0 and \mathbf{C}_1 from a failure event F_N , the two directions cancel each other out as the ciphertexts $\overline{\mathbf{C}_0^{(0)}}$ and $\overline{\mathbf{C}_1^{(0)}}$ are in opposite directions.

Keeping the convention that $\overline{\mathbf{C}_0^{(0)}}$ approximates a rotation of the secret $\overline{\mathbf{S}'} = \overline{X^{-r_0} \cdot \mathbf{S}}$, we will compute the relative error position $\delta_{1,0} = r_1 - r_0$ and show that is enough to build a correct estimate $\overline{\mathbf{E}}$ as $\overline{\mathbf{E}} = \overline{\mathbf{C}_{\text{rav}}} / \|\overline{\mathbf{C}_{\text{rav}}}\|_2$ where:

$$\overline{\mathbf{C}_{\text{rav}}} := \left(\overline{\mathbf{C}_0^{(0)}} + \overline{\mathbf{C}_1^{(\delta_{1,0})}} \right) / 2. \quad (17)$$

The reason why such $\overline{\mathbf{E}}$ is a good estimator of $\overline{\mathbf{S}'}$ can be seen as follows:

$$\begin{aligned} \cos(\theta_{ES'}) &= \frac{1}{2 \|\overline{\mathbf{C}_{\text{rav}}}\|_2 \|\overline{\mathbf{S}'}\|_2} \cdot \left(\overline{X^{-r_0} \cdot \mathbf{S}^T \cdot \mathbf{C}_0^{(0)}} + \overline{X^{-r_0} \cdot \mathbf{S}^T \cdot X^{r_1 - r_0} \mathbf{C}_1^{(0)}} \right) \\ &= \frac{1}{2 \|\overline{\mathbf{C}_{\text{rav}}}\|_2 \|\overline{\mathbf{S}'}\|_2} \cdot \left(\overline{\mathbf{S}^T \cdot \mathbf{C}_0^{(r_0)}} + \overline{\mathbf{S}^T \cdot \mathbf{C}_1^{(r_1)}} \right) > \frac{q_t}{\|\overline{\mathbf{C}_{\text{rav}}}\|_2 \|\overline{\mathbf{S}'}\|_2}. \end{aligned}$$

Remark 4. In practice ciphertexts with smaller norm will on average be better aligned with the secret, as $\cos(\theta_{CS'}) > q_t / (\|\overline{\mathbf{C}}\|_2 \|\overline{\mathbf{S}'}\|_2)$. Therefore they carry more information than ciphertexts with larger norm. To compensate for this effect we will calculate $\overline{\mathbf{C}_{\text{rav}}}$ as $:= \left(\overline{\mathbf{C}_0^{(0)}} / \|\overline{\mathbf{C}_0^{(0)}}\|_2 + \overline{\mathbf{C}_1^{(\delta_{1,0})}} / \|\overline{\mathbf{C}_1^{(\delta_{1,0})}}\|_2 \right) / 2$. While it is possible to further refine the calculation of $\overline{\mathbf{E}}$ using extra directional information, this heuristic is good enough for our purposes.

Computation of the relative position $\delta_{1,0}$. One can use the fact that both $\overline{\mathbf{C}}_0^{(0)}$ and $\overline{\mathbf{C}}_1^{(\delta_{1,0})}$ are expected to be directionally close to $\overline{\mathbf{S}}$. Thus, the cosine of the angle between $\overline{\mathbf{C}}_0^{(0)}$ and $\overline{\mathbf{C}}_1^{(\delta_{1,0})}$ should be larger than usual. Therefore, $\delta_{1,0}$ can be estimated with the following distinguisher:

$$\delta'_{1,0} := \operatorname{argmax}_{r \in [0, 2N-1]} \mathcal{C}(r) \text{ where } \mathcal{C}(r) := \frac{\overline{\mathbf{C}}_0^{(0)T} \cdot \overline{\mathbf{C}}_1^{(r)}}{\|\overline{\mathbf{C}}_0^{(0)}\|_2 \|\overline{\mathbf{C}}_1^{(r)}\|_2}. \quad (18)$$

The next paragraph estimates the efficiency of using Equation 18 as a distinguisher for deriving $\delta_{1,0}$. We will show that, for Table 1 parameters, we expect

$$P[\delta'_{1,0} = \delta_{1,0}] \approx 89\%. \quad (19)$$

Experiments run by simulating the sampling 10^4 failing ciphertexts (refer to the full version of our paper [15] for the generation technique), and using Equation 18 for finding $\delta_{1,0}$ between pairs of them, return $P_{\text{Exp}}[\delta'_{1,0} = \delta_{1,0}] \approx 84.8\%$, in sufficiently good agreement.

To obtain the value (19), the idea is to estimate the distribution of a correct guess $\mathcal{C}(\delta_{1,0})$ and an incorrect guess $\max_{r \neq \delta_{1,0}} \mathcal{C}(r)$ and quantify the discrepancy. First, we decompose the ciphertexts in a component parallel to $\overline{\mathbf{S}}$, denoted with \parallel , and a component orthogonal, denoted with \perp , we rewrite $\mathcal{C}(r)$ as follows:

$$\mathcal{C}(r) = \frac{\overline{\mathbf{C}}_{0,\parallel}^{(0)} \cdot \overline{\mathbf{C}}_{1,\parallel}^{(r)} + \overline{\mathbf{C}}_{0,\perp}^{(0)} \cdot \overline{\mathbf{C}}_{1,\perp}^{(r)}}{\|\overline{\mathbf{C}}_0^{(0)}\|_2 \|\overline{\mathbf{C}}_1^{(r)}\|_2} \quad (20)$$

In the first term, the scalar product of two parallel elements equals the product of their norms (up to their sign). For the second term, we apply the scalar product definition and introduce t as the angle between $\overline{\mathbf{C}}_{0,\perp}^{(0)}$ and $\overline{\mathbf{C}}_{1,\perp}^{(r)}$.


$$\mathcal{C}(r) = \pm \frac{\|\overline{\mathbf{C}}_{0,\parallel}^{(0)}\|_2}{\|\overline{\mathbf{C}}_0^{(0)}\|_2} \cdot \frac{\|\overline{\mathbf{C}}_{1,\parallel}^{(r)}\|_2}{\|\overline{\mathbf{C}}_1^{(r)}\|_2} \pm \frac{\|\overline{\mathbf{C}}_{0,\perp}^{(0)}\|_2}{\|\overline{\mathbf{C}}_0^{(0)}\|_2} \cdot \frac{\|\overline{\mathbf{C}}_{1,\perp}^{(r)}\|_2}{\|\overline{\mathbf{C}}_1^{(r)}\|_2} \cdot \cos(t) \quad (21)$$

$$= \cos(\theta_{S', \overline{\mathbf{C}}_0^{(0)}}) \cos(\theta_{S', \overline{\mathbf{C}}_1^{(r)}}) + \sin(\theta_{S', \overline{\mathbf{C}}_0^{(0)}}) \sin(\theta_{S', \overline{\mathbf{C}}_1^{(r)}}) \cos(t) \quad (22)$$

The vectors $\overline{\mathbf{C}}_{0,\perp}^{(0)}$ and $\overline{\mathbf{C}}_{1,\perp}^{(r)}$ are orthogonal to $\overline{\mathbf{S}}$. This means that they live in the $2Nl - 1$ dimensional space orthogonal to $\overline{\mathbf{S}}$. The high dimension of the space will strongly drive the vectors towards orthogonality as can be seen in Figure 2. Using Assumption 1, the angle t between $\overline{\mathbf{C}}_{0,\perp}^{(0)}$ and $\overline{\mathbf{C}}_{1,\perp}^{(r)}$ is then assumed to follow the distribution of random angles between vectors in a $2Nl - 1$ dimensional space (See Equation 11).

Now, let us study the distribution of $\mathcal{C}(r)$ depending of the value $r \in [0, 2N-1]$. One can refer to Figure 3 for a graphical interpretation based on the parameters of Table 1.

Fig. 3. Distributions used for finding $\delta_{1,0}$



- If $r = \delta_{1,0}$, the expected value of $\mathcal{C}(r)$ will be higher than average. Indeed, by definition of F_{r_1} and F_{r_0} the cosines forming the first term are positive. The distribution of $\mathcal{C}(r)$ can then be estimated using Equation 22 (blue curve).
- If $r = \delta_{1,0} + N \pmod{2N}$, the distribution of $\mathcal{C}(r)$ is equal to the distribution of $-\mathcal{C}(\delta_{1,0})$ and will be closer to -1 (orange curve).
- If $r \neq \delta_{1,0} \pmod{N}$, $\mathcal{C}(r)$ can be assumed to follow the distribution of random angles in a $2Nl$ dimensional space Θ_{2Nl} , as given in Equation 11 (green curve).
- The pdf of $\max_{r \neq \delta_{1,0}} \mathcal{C}(r)$ is then calculated as $M(\Theta_{2Nl}, 2N - 1)$ by definition of the maximal order statistic (red curve).

Figure 3 assembles the probability density functions of the above distributions in a plot. The probability of selecting the correct $\delta'_{1,0}$ using $\operatorname{argmax}_{r \in [0, 2N-1]} \mathcal{C}(r)$, can then be computed as:

$$P[\delta'_{1,0} = \delta_{1,0}] = P \left[\max_{r \neq \delta_{1,0}} \mathcal{C}(r) < \mathcal{C}(\delta_{1,0}) \right].$$

For our toy scheme’s parameters, this results in Equation 19.

Multiple ciphertexts. In this section, we assume that n linearly independent failing ciphertexts $\mathbf{C}_0, \dots, \mathbf{C}_{n-1}$, resulting from failure events $F_{r_0}, \dots, F_{r_{n-1}}$ respectively, are given. We introduce a generalized method to recover the relative positions $\delta_{1,0}, \dots, \delta_{n-1,0}$, based on “loopy belief propagation” [42]. Once these relative positions are found, they can be combined in an estimate $\overline{\mathbf{E}}$ with $\overline{\mathbf{E}} = \overline{\mathbf{C}}_{\text{rav}} / \|\overline{\mathbf{C}}_{\text{rav}}\|_2$ where

$$\overline{\mathbf{C}}_{\text{rav}} := \left(\overline{\mathbf{C}}_0^{(0)} / \|\overline{\mathbf{C}}_0^{(0)}\|_2 + \sum_{i \in [1, n-1]} \overline{\mathbf{C}}_i^{(\delta_{i,0})} / \|\overline{\mathbf{C}}_i^{(\delta_{i,0})}\|_2 \right) / n. \quad (23)$$

To find the correct rotations, we construct a weighted graph that models the probability of different relative rotations, and we will use loopy belief propagation to obtain the most probable set of these rotations:

- The nodes represent the obtained failing ciphertexts: $(\mathbf{C}_i)_{i \in [0, n-1]}$. In total, there are n nodes.
- Each node \mathbf{C}_i where $i \neq 0$ is associated a list with $2N$ probability values called *beliefs* and denoted $(b_i(0), \dots, b_i(2N-1))$ that together define a probability distribution over $[0, 2N-1]$. The r^{th} item in the list represents our belief that the correct relative position $\delta_{0,1}$ equals r . The correct rotation of the 0^{th} node will be fixed to 0 (i.e. $b_0(0) = 1$ and $b_0(i) = 0$ for all other i) as only the relative rotations of the ciphertexts is important. These node weights are initialized as follows:

$$b_i(r) := P[\delta_{i,0} = r] \quad (= P[F_r \text{ for } \mathbf{C}_i | F_0 \text{ for } \mathbf{C}_0])$$

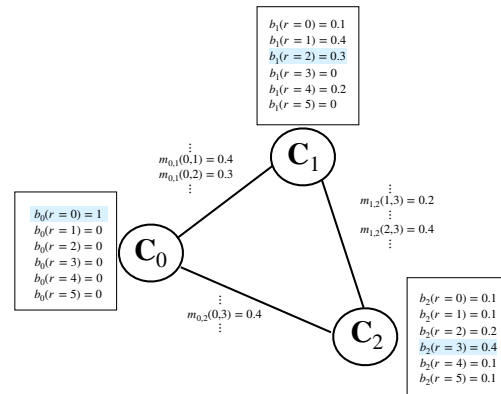
- For two nodes \mathbf{C}_i and \mathbf{C}_j , the value of the vertex called *message*, models the influence of the beliefs in the rotations s of node j towards the beliefs in rotation r of node i , which can be formalized as follows:

$$m_{i,j}(r, s) := P[\delta_{i,j} = r - s] \quad (= P[F_{r-s} \text{ for } \mathbf{C}_i | F_0 \text{ for } \mathbf{C}_j])$$

Loopy belief propagation tries to find the node values r for each node, so that the probabilities over the whole graph are maximized. This is done in an iterative fashion by updating the node beliefs according to the messages coming from all other nodes. Our goal is eventually to find $r = \delta_{i,0}$ for each node i .

Example 2. For example, with $N = 3$ and $n = 3$, the graph contains the nodes \mathbf{C}_0 , \mathbf{C}_1 , and \mathbf{C}_2 . In Figure 4, we represent how such a graph could look like where we arbitrarily instantiate the messages and beliefs. We can see that if one chooses the $r_i = \text{argmax}_r b_i(r)$ for each node, one would have chosen $r_1 = 1$ and $r_2 = 3$. Nevertheless, we notice that the influence of the other probabilities allows for a better choice (underlined in blue in the figure): $r_1 = 2, r_2 = 3$.

Fig. 4. Example of the graph for finding the relative rotations where $N = 3$ and $n = 3$. The beliefs are in the rectangles, the circles represent the nodes and some messages are represented between the nodes.



Vertex probabilities. As discussed, the edge between two nodes \mathbf{C}_i with rotation r and \mathbf{C}_j with rotation s is weighted with $P[\delta_{i,j} = r - s]$. This probability can be computed using a generalization of the distinguisher technique used for two ciphertexts as detailed in the full version of our paper [15].

Algorithm 6: GetRotation()

```

1 for  $i \in [1, n - 1]$  do // initialization
2   foreach  $r$  do
3      $b_i(r) := P[\delta_{0,i} = r]$ 
4 for # of iterations do // update phase
5   for  $i \in [1, n]$  do
6     for  $j \in [1, n]$  if  $i \neq j$  do
7       foreach  $r$  do
8          $\text{infl}_{ji}(r) := \sum_s m_{i,j}(r, s) \cdot b_j(s)$  // influence of node  $j$  on
           node  $i$ 
9          $\text{normalize}(\text{infl}_{ji})$ 
10      foreach  $r$  do
11         $b_i(r) := \prod_{j=0, j \neq i}^n \text{infl}_{ji}(r)$  // calculate new belief
12       $\text{normalize}(b_i)$ 
13 for  $i \in [1, n]$  do // finally
14    $r_i := \text{argmax}_{r \in [0, 2N-1]} b_i(r)$  // pick the  $r_i$  with highest belief
15 return  $(r_i)_{i \in [1, n-1]}$ 

```

	2 ciphertexts	3 ciphertexts	4 ciphertexts	5 ciphertexts
$P[r_i = \delta_{i,0} \forall i \in [1, n - 1]]$	84.0%	95.6%	> 99.0%	> 99.0%

Table 2. Probability of finding the correct relative rotations and thus building the correct estimate $\bar{\mathbf{E}}$ with the knowledge of 2, 3, 4 and 5 failing ciphertexts.

Loopy belief propagation. This technique looks for the best set (r_1, \dots, r_{n-1}) by iteratively correcting the beliefs using messages from other nodes. This procedure is detailed in Algorithm 6, where $\text{normalize}(f)$ normalizes the list $b()$ so that $\sum_{x \in \text{supp}(b)} b(x) = 1$. In each iteration, the belief of each node \mathbf{C}_i is updated according to the messages of the other nodes \mathbf{C}_j . For each i the belief is updated as follows:

$$b_i(r) = \prod_{j=0, j \neq i}^n \text{infl}_{ji}(r) \quad (24)$$

where $\text{infl}_{ji}(r)$ captures the influence of the value of the node \mathbf{C}_j to node \mathbf{C}_i . This influence can be calculated as $\text{infl}_{ji}(r) \leftarrow C \sum_x m_{i,j}(r, x) \cdot b_j(x)$, with C as normalizing constant.

Experimental verification. With Table 1 parameters, we obtained the correct values $r_i = \delta_{i,0}$ for all $i \in [1, n - 1]$ after 3 iterations with the probabilities as reported in Table 2, by generating 1000 times each number of ciphertexts and trying to find the correct values of the r_i .

Remark 5 (Consistency with the previous section). Note that this procedure also incorporates the setting where one has only 2 failing ciphertexts, which would yield exactly the same results as in the previous paragraph.

Finally, once all the rotations are found, recall that the estimate is obtained by $\overline{\mathbf{E}} = \overline{\mathbf{C}}_{\text{rav}} / \|\overline{\mathbf{C}}_{\text{rav}}\|_2$ where

$$\overline{\mathbf{C}}_{\text{rav}} = \left(\overline{\mathbf{C}}_0^{(0)} / \|\overline{\mathbf{C}}_0^{(0)}\|_2 + \sum_{i \in [1, n-1]} \overline{\mathbf{C}}_i^{(r_i)} / \|\overline{\mathbf{C}}_i^{(r_i)}\|_2 \right) / n. \quad (25)$$

5.3 Step 2: Finding weak ciphertexts

In this section, we are given an estimate $\overline{\mathbf{E}}$ and we refine the acceptance criterion. Instead of accepting if $P[F | ct] \geq f_t$, our condition is slightly changed.

1. Generate a key encapsulation $ct = (\mathbf{C}, \mathbf{G})$ with derived key K .
2. If $P[F | \overline{\mathbf{E}}, ct] \geq f_t$, keep ct in a weak ciphertext list, otherwise go to Step 1.

In Step 2, $P[F | \overline{\mathbf{E}}, ct]$ is defined as the failure probability, given a certain ciphertext ct and a certain estimate $\overline{\mathbf{E}}$. In the following, we explain a way to compute it.

First, for $r \in [0, 2N - 1]$, we will estimate the probability that a ciphertext leads to an error in the r^{th} location. Decomposing the vectors $\overline{\mathbf{S}}$ and $\overline{\mathbf{C}}$ in a component orthogonal to $\overline{\mathbf{E}}$, denoted with subscript \perp , and a component parallel to $\overline{\mathbf{E}}$, denoted with subscript \parallel , we obtain the failure expression:

$$\begin{aligned} P[F_r | \overline{\mathbf{E}}, \mathbf{C}] &= P[\overline{\mathbf{S}}^T \cdot \overline{\mathbf{C}}^{(r)} > q_t | \overline{\mathbf{E}}, \mathbf{C}] = P[\overline{\mathbf{S}}_{\parallel}^T \cdot \overline{\mathbf{C}}_{\parallel}^{(r)} + \overline{\mathbf{S}}_{\perp}^T \cdot \overline{\mathbf{C}}_{\perp}^{(r)} > q_t | \overline{\mathbf{E}}, \mathbf{C}] \\ &= P \left[\left(\frac{\|\overline{\mathbf{S}}\|_2 \|\overline{\mathbf{C}}^{(r)}\|_2 \cos(\theta_{SE}) \cos(\theta_{C^{(r)}E}) + \|\overline{\mathbf{S}}\|_2 \|\overline{\mathbf{C}}^{(r)}\|_2 \sin(\theta_{SE}) \sin(\theta_{C^{(r)}E}) \cos(t)}{\|\overline{\mathbf{S}}\|_2 \|\overline{\mathbf{C}}^{(r)}\|_2} \right) > q_t | \overline{\mathbf{E}}, \mathbf{C} \right] \\ &= P \left[\cos(t) > \frac{q_t - \|\overline{\mathbf{S}}\|_2 \|\overline{\mathbf{C}}^{(r)}\|_2 \cos(\theta_{SE}) \cos(\theta_{C^{(r)}E})}{\|\overline{\mathbf{S}}\|_2 \|\overline{\mathbf{C}}^{(r)}\|_2 \sin(\theta_{SE}) \sin(\theta_{C^{(r)}E})} | \overline{\mathbf{E}}, \mathbf{C} \right] \end{aligned}$$

where θ_{SE} and $\theta_{C^{(r)}E}$ are the angles of $\overline{\mathbf{S}}$ and $\overline{\mathbf{C}}^{(r)}$ with the estimate $\overline{\mathbf{E}}$ respectively, and where t is the angle between $\overline{\mathbf{S}}_{\perp}$ and $\overline{\mathbf{C}}_{\perp}^{(r)}$. We assume no other knowledge about the direction of the secret apart from the directional estimate $\overline{\mathbf{E}}$. In this case, using Assumption 1, t can be estimated as a uniform angle in a $2Nl - 1$ dimensional space. Then t is assumed to follow the probability distribution Θ_{2Nl-1} (defined in Equation 11).

The values $\overline{\mathbf{E}}$, $\|\mathbf{C}\|_2$ and $\cos(\theta_{C^{(r)}E})$ are known, meanwhile the values $\|\mathbf{S}\|$ and θ_{SE} can be modelled using their probability distribution. Thus, we can approximate $P[F_i | \overline{\mathbf{E}}, \mathbf{C}]$ with $P[F_i | \overline{\mathbf{E}}, \|\mathbf{C}\|_2, \cos(\theta_{C^{(r)}E})]$.

Assumption 2 *We assume that failures at different locations are independent.*

Assumption 2 is a valid assumption for schemes without error correcting codes, as discussed in [17]. We can then calculate the failure probability of a certain ciphertext as:

$$P[F | \bar{\mathbf{E}}, \mathbf{C}] = 1 - \prod_{r=0}^{2N} (1 - P[F_r | \bar{\mathbf{E}}, \|\bar{\mathbf{C}}\|_2, \cos(\theta_{C^{(r)}E})]) \quad (26)$$

As this expression gives us a better prediction of the failure probability of ciphertexts by using the information embedded in $\bar{\mathbf{E}}$, we can more accurately (Grover) search for weak ciphertexts and thus reduce the work to find the next decryption failure. Moreover, the better $\bar{\mathbf{E}}$ approximates the direction of $\bar{\mathbf{S}}$, the easier it becomes to find a new decryption failure.

5.4 Finalizing the attack with lattice reduction

Once multiple failures are found, the secret key can be recovered with lattice reduction techniques as presented in [17, §4] and in [28, Step 3 of the attack]. The following Section simply outlines how their technique transposes to our framework. As shown in §5, an estimate $\bar{\mathbf{E}}$ of the direction of a rotated version of $\bar{\mathbf{S}}' = \bar{X}^r \bar{\mathbf{S}}$ with an unknown value r is provided. Therefore, similarly to [28], an attacker can obtain an estimation of $\bar{\mathbf{S}}'$ (and not only its direction) by rescaling

$$\bar{\mathbf{E}}' := \bar{\mathbf{E}} \cdot nq_t \cdot \left(\left\| \bar{\mathbf{C}}_0^{(0)} + \sum_{i \in [1, n-1]} \bar{\mathbf{C}}_i^{(r_i)} \right\|_2 \right)^{-1},$$

using the approximation $\bar{\mathbf{E}}'^T \cdot 1/n \left(\bar{\mathbf{C}}_0^{(0)} + \sum_{i \in [1, n-1]} \bar{\mathbf{C}}_i^{(r_i)} \right) \approx q_t$.

Then, for each possible $r \in [0, 2N - 1]$, an attacker can perform lattice reduction and recover candidates for \mathbf{s}, \mathbf{e} that are accepted if they verify $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$. One caveat is that an attacker may have to run a lattice reduction up to $2N$ times. Since $\bar{\mathbf{E}}' - \bar{\mathbf{S}}'$ is small, the attacker can construct an appropriate lattice basis encoding $\bar{\mathbf{E}}' - \bar{\mathbf{S}}'$ as a unique shortest target vector, and solves the corresponding Unique-SVP problem with the BKZ algorithm [46, 11, 3, 1]. The block size of BKZ will depend on the accuracy of the estimate $\bar{\mathbf{E}}$. Indeed, the standard deviation of $\bar{\mathbf{E}}'_i - \bar{\mathbf{S}}'_i$ is of the order of $\sigma_s \cdot \sin(\theta_{S'E})$ (assuming that $\theta_{S'E}$ is small and $\|\bar{\mathbf{S}}'\|_2 \approx \|\bar{\mathbf{E}}'\|_2$). Thus, when many decryption failures are available, $\sin(\theta_{S'E})$ gets very small and the complexity of this step is dominated by the work required for constructing $\bar{\mathbf{E}}$. For example, in the case of our toy scheme, if $\cos(\theta_{S'E}) > 0.985$, using [2], the BKZ block size becomes lower than 363 which leads to less than 2^{100} quantum work (in the Core-SVP [3] 0.265β model). As we will see in §6.3, this is less than the work required to find the first failure.

Remark 6. One can think that the failures obtained by directional failure boosting will not be totally independent. It is true that the failing ciphertexts are roughly following the same direction. But applying our Assumption 1, in high dimensions, for a reasonable number n of failures ($n \ll 2lN$), the hypercone in which the failures belong is large enough that linear dependency will happen with very low probability.

6 Efficiency of Directional Failure Boosting

In this section, we experimentally verify the efficiency of the directional failure boosting technique. We first quantify the accuracy of the estimate $\overline{\mathbf{E}}$ computed according to §5.2. We then derive the necessary work required to run the directional failure boosting technique and the optimal number of queries. For the rest of the section, we focus on minimizing the total work for finding failures and we will assume there is no upper limit to the number of decryption queries.

Our key takeaway is that, for Table 1 parameters, the more failing ciphertexts have been found, the easier it becomes to obtain the next one, and that most of the effort is concentrated in finding the first failure. The final work and query overheads are stored in Table 4.

6.1 Accuracy of the estimate

Let $\mathbf{C}_0, \dots, \mathbf{C}_{n-1}$ be n previously found failing ciphertexts and we take the estimate defined according to Equation 25. Similarly to §5.2, we define $\overline{\mathbf{S}'} = \overline{X^{-r_0}} \cdot \overline{\mathbf{S}}$ as the secret vector for an unknown F_{r_0} . To estimate the accuracy of $\overline{\mathbf{E}}$, we compute $\cos(\theta_{S'E}) = \frac{\overline{\mathbf{S}'^T} \cdot \overline{\mathbf{E}}}{\|\overline{\mathbf{S}'}\|_2} = \frac{\overline{\mathbf{S}'^T} \cdot \overline{\mathbf{C}_{\text{rav}}}}{\|\overline{\mathbf{S}'}\|_2 \|\overline{\mathbf{C}_{\text{rav}}}\|_2}$ as

$$\cos(\theta_{S'E}) = \frac{\overline{\mathbf{S}'^T} \cdot \left(\frac{\overline{\mathbf{C}_0^{(0)}}}{\|\overline{\mathbf{C}_0^{(0)}}\|_2} + \sum_{i=1}^{n-1} \frac{\overline{\mathbf{C}_i^{(r_i)}}}{\|\overline{\mathbf{C}_i^{(r_i)}}\|_2} \right)}{n \|\overline{\mathbf{S}'}\|_2 \|\overline{\mathbf{C}_{\text{rav}}}\|_2} \quad (27)$$

$$= \frac{\cos(\theta_{\mathbf{C}_0^{(0)} S'}) + \sum_{i=1}^{n-1} \cos(\theta_{\mathbf{C}_i^{(r_i)} S'})}{\left\| \frac{\overline{\mathbf{C}_0^{(0)}}}{\|\overline{\mathbf{C}_0^{(0)}}\|_2} + \sum_{i=1}^{n-1} \frac{\overline{\mathbf{C}_i^{(r_i)}}}{\|\overline{\mathbf{C}_i^{(r_i)}}\|_2} \right\|_2} \quad (28)$$

First, we make the following approximation.

Approximation 1 We approximate the cosine with the secret $\overline{\mathbf{S}'}$ by its expected value denoted $\cos(\theta_{CS'}) := \mathbb{E} \left[\cos(\theta_{\mathbf{C}_i^{(r_i)} S'}) \right]$. In other words, for all $i \in [1, n-1]$ we assume

$$\cos(\theta_{CS'}) = \cos(\theta_{\mathbf{C}_i^{(r_i)} S'}) = \cos(\theta_{\mathbf{C}_0^{(0)} S'}).$$

Table 3. Accuracy of the estimate derived from several failures. Expected value of $\cos(\theta_{S'E})$ according to Equation 29. The closer to 1, the more accurate $\overline{\mathbf{E}}$ is.

n	1	2	3	5	10	20	30	50	100
theoretical	0.328	0.441	0.516	0.613	0.739	0.841	0.885	0.926	0.961
experimental	0.318	0.429	0.502	0.600	0.727	0.832	0.878	0.921	0.958

To estimate the denominator of Equation 28, we split the ciphertexts in a component parallel to the secret $\overline{\mathbf{C}}_{i,\parallel}^{(r_i)}$ and a component orthogonal $\overline{\mathbf{C}}_{i,\perp}^{(r_i)}$ to the secret. Following Assumption 1, we will assume orthogonality between the various $\overline{\mathbf{C}}_{i,\perp}^{(r_i)}$. As the norm of the sum of parallel vectors is the sum of their norm, and the norm of the sum of orthogonal vectors can be calculated using Pythagoras' theorem, we can approximate $\cos(\theta_{S'E})$ as follows:

$$\cos(\theta_{S'E}) \approx \frac{n \cos(\theta_{CS'})}{\sqrt{n^2 \cos(\theta_{CS'})^2 + n \sin(\theta_{CS'})^2}} = \frac{\cos(\theta_{CS'})}{\sqrt{\cos(\theta_{CS'})^2 + \sin(\theta_{CS'})^2/n}} \quad (29)$$

One can see from this equation that $\cos(\theta_{S'E})$ gets closer to 1 when n increases.

Experimental verification. The first line of Table 3 gives the expected values of $\cos(\theta_{S'E})$ for various n , according to Equation 29, with $\cos(\theta_{CS'})$ set to $q_t / \mathbb{E}[\|\overline{\mathbf{S}}\|] \mathbb{E}[\|\overline{\mathbf{C}}^{(0)}\|]$, which is a good approximation of $\cos(\theta_{CS'})$ as $\cos(\theta_{CS'}) > q_t / \|\overline{\mathbf{S}}\| \|\overline{\mathbf{C}}^{(0)}\|$ and because angles tend to orthogonality in high dimensional space.

Then, to verify the theory, we implemented a method to simulate the distribution of random failing ciphertexts. This technique is described in the full version of our paper [15]. Once the simulated failing ciphertexts are found, we combine them to build $\overline{\mathbf{E}}$ using their correct rotations, and we compute $\cos(\theta_{S'E})$. The latter experiment was repeated 100 times and the average values are reported in line two of Table 3.

6.2 Estimating α_{i,f_t} and β_{i,f_t}

To estimate the effectiveness of directional failure boosting given a certain number i of previously collected failing ciphertexts, we need to find the optimal weak ciphertext threshold f_t for each i . This corresponds to considering how much time to spend for one precalculation $\sqrt{\alpha_{i,f_t}^{-1}}$ and the average failure probability of weak ciphertexts β_{i,f_t} after the precalculation. Let us recall the definition of α_{n,f_t} and β_{n,f_t} , derived from Equations 7 and 8, where $\mathbf{C}_0, \dots, \mathbf{C}_{n-1}$ are the n previously found failing ciphertexts.

$$\alpha_{i,f_t} = \sum_{\forall ct: P[F|ct, \mathbf{C}_0, \dots, \mathbf{C}_{n-1}] > f_t} P[ct] \quad (30)$$

$$\beta_{i,f_t} = \frac{\sum_{\forall ct: P[F|ct, \mathbf{C}_0, \dots, \mathbf{C}_{n-1}] > f_t} P[ct] \cdot P[F|ct, \mathbf{C}_0, \dots, \mathbf{C}_{n-1}] > f_t}{\sum_{\forall ct: P[F|ct, \mathbf{C}_0, \dots, \mathbf{C}_{n-1}] > f_t} P[ct]} \quad (31)$$

To find the optimal values, we need to calculate Equation 30 and 31 as functions of f_t . This requires us to list the probability of all ciphertexts $ct := (\mathbf{C}, \mathbf{G})$, and their failure probability $P[F|ct, \mathbf{C}_0, \dots, \mathbf{C}_{t-1}]$. As discussed in [16], exhaustively computing both values is not practically feasible, and therefore we will make some assumptions to get an estimate.

A first simplification is to cluster ciphertexts that have similar $\|\mathbf{C}\|_2$ and $|\theta_{C^{(0)}E}| \cdots |\theta_{C^{(N-1)}E}|$ and thus a similar failure probability. To further reduce the list size, we only take into account the largest value of $|\cos(\theta_{C^{(i)}E})|$ denoted

$$\text{maxcos}(\theta_{CE}) := \max_i (|\cos(\theta_{C^{(i)}E})|),$$

which results in a slight underestimation of the effectiveness of the attack. In other words,

$$\begin{aligned} P[ct] &\text{ becomes } P[\|\mathbf{C}\|_2, \text{maxcos}(\theta_{CE})], \\ P[F|ct, \mathbf{C}_0, \dots, \mathbf{C}_{n-1}] &\text{ becomes } P[F | \|\mathbf{C}\|_2, \text{maxcos}(\theta_{CE})]. \end{aligned}$$

Assuming independence between the norm of \mathbf{C} and its angle with $\bar{\mathbf{E}}$, $P[\|\mathbf{C}\|_2, \text{maxcos}(\theta_{CE})]$ can be estimated using the distributions defined with Equations 10 and 13 as follows:

$$P[\|\mathbf{C}\|_2, \text{maxcos}(\theta_{CE})] = \underbrace{P[\|\mathbf{C}\|_2]}_{\chi_{Nl, \sigma}} \cdot \underbrace{P[\text{maxcos}(\theta_{CE})]}_{M(\Theta_{2Nl, 2N})}. \quad (32)$$

Denoting with r the position of the maximum angle, we can rewrite $P[F | \|\mathbf{C}\|_2, \text{maxcos}(\theta_{CE})]$ as follows:

$$\begin{aligned} P[F | \|\mathbf{C}\|_2, \text{maxcos}(\theta_{CE})] &= 1 - \prod_i \left(1 - P[F_i | \|\mathbf{C}\|_2, |\cos(\theta_{C^{(r)}E})|] \right), \quad (33) \\ &= 1 - \left(\left(1 - P[F_r | \|\mathbf{C}\|_2, |\cos(\theta_{C^{(r)}E})|] \right) \cdot \right. \\ &\quad \left. \prod_{i \neq r} \left(1 - P[F_i | \|\mathbf{C}\|_2, |\cos(\theta_{C^{(i)}E})| \leq |\cos(\theta_{C^{(r)}E})|] \right) \right), \quad (34) \end{aligned}$$

where $1 - P[F_r | \|\mathbf{C}\|_2, |\cos(\theta_{C^{(r)}E})|]$ can be estimated using Equation 26, and $P[F_i | \|\mathbf{C}\|_2, |\cos(\theta_{C^{(i)}E})| \leq |\cos(\theta_{C^{(r)}E})|]$ using an integral over Equation 26. The estimated failure probability of ciphertexts given $\|\mathbf{C}\|_2$ and $\cos(\theta_{CE})$ for the parameters listed in Table 1 is depicted in figure 5a.

Verification experiment. We verified these results experimentally by generating $5 \cdot 10^6$ failing ciphertexts and $5 \cdot 10^6$ successful ciphertexts, and calculating their norm and angle with 1000 estimates, or in this case other ciphertexts. The failing ciphertexts were produced using the methodology detailed in the full version of our paper [15]. Once they are generated, we estimate their failure probability with a procedure also detailed in the full version of our paper [15]. We combined

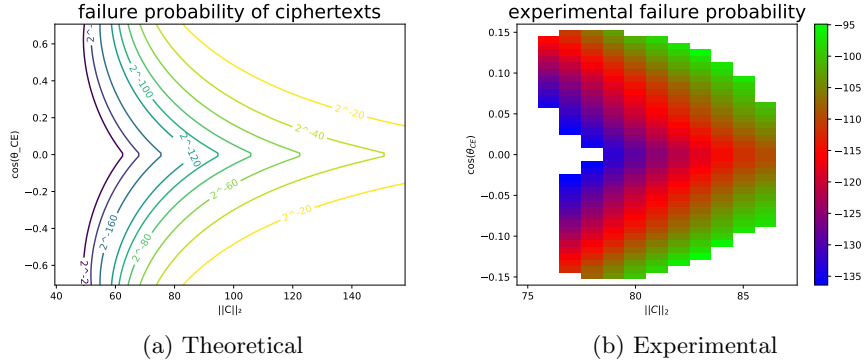


Fig. 5. Failure probability of ciphertexts as a function of $\|C\|_2$ and $\cos(\theta_{CE})$. A zoomed version of Figure 5a for easier comparison can be found the full version of our paper [15].

these results into Figure 5b. These experimental results confirm our theoretical estimates given in Figure 5a.

With the estimation of $P[F | \|C\|_2, \max \cos(\theta_{CE})]$ and $P[\|C\|_2, \max \cos(\theta_{CE})]$, α_{i, f_t} and β_{i, f_t} can be estimated as functions of i and f_t . Let us now define the optimal threshold f_t as a function of i as :

$$f_i := \operatorname{argmin}_{f_t} (\sqrt{\alpha_{i, f_t}} \cdot \beta_{i, f_t})^{-1}.$$

6.3 Total amount of work and queries

In this section, we will derive the optimal work and queries for an adversary to perform, in order to obtain n ciphertexts with probability $1 - e^{-1}$. We introduce the following notation: to find the $(i + 1)^{\text{th}}$ ciphertext, the adversary performs Q_i queries. Using a Poisson distribution, the success probability of finding the $(i + 1)^{\text{th}}$ ciphertext in Q_i queries is $1 - e^{-Q_i \beta_{i, f_i}}$. The probability of obtaining n failures can then be calculated as the product of the success probabilities of finding ciphertexts 0 to $n - 1$:

$$P_n = \prod_{i=0}^{n-1} (1 - e^{-Q_i \beta_{i, f_i}}). \quad (35)$$

This is a slight underestimation of the success probability of the attack, because if an adversary finds a failing ciphertext in less than Q_i samples, she can query more ciphertexts in the next stages $i + 1, \dots, n$. However, this effect is small due to the large value of Q_i .

The total amount of precomputation quantum work, and the total amount of queries to obtain the n failing ciphertexts by performing Q_i tries for each ciphertext, can be expressed as

ciphertexts n	1	2	3	5	10	20	30
$\log_2(\mathcal{W}_n^{\text{tot}})$	112.45	112.77	112.78	112.78	112.78	112.78	112.78
$\log_2(\mathcal{W}_n^{\text{tot}}/\mathcal{W}_1^{\text{tot}})$	—	0.32	0.33	0.33	0.33	0.33	0.33
$\log_2(\mathcal{Q}_n^{\text{tot}})$	102.21	102.86	102.87	102.87	102.87	102.87	102.87
$\log_2(\mathcal{Q}_n^{\text{tot}}/\mathcal{Q}_1^{\text{tot}})$	—	0.65	0.66	0.66	0.66	0.66	0.66

Table 4. Quantum work $\mathcal{W}_n^{\text{tot}}$ and queries $\mathcal{Q}_n^{\text{tot}}$ required to find n failing ciphertexts with probability $1 - e^{-1}$. Finding the first ciphertext requires the heaviest amount of computation. After the third failing ciphertext is found, the following ones are essentially for free.

$$\mathcal{W}_n^{\text{tot}} := \sum_{i=0}^{n-1} \frac{Q_i}{\underbrace{\sqrt{\alpha_{i,f_i}}}_{:=W_i}} \quad \mathcal{Q}_n^{\text{tot}} := \sum_{i=0}^{n-1} Q_i. \quad (36)$$

Recall that for now we assume there is no upper limit to the number of decryption queries that can be made, and we focus on minimizing the amount of work. The values of Q_i that minimizes the total quantum work $\mathcal{W}_n^{\text{tot}}$ can be found using the following Lagrange multiplier, minimizing the total amount of work to find n failures with probability $1 - e^{-1}$ using the above probability model:

$$L(Q_0, \dots, Q_{n-1}, \lambda) = \sum_{t=0}^{n-1} \frac{Q_i}{\sqrt{\alpha_{i,f_i}}} + \lambda \left((1 - e^{-1}) - \prod_{i=0}^{n-1} (1 - e^{-Q_i \beta_{i,f_i}}) \right) \quad (37)$$

By equating the partial derivative of L in Q_0, \dots, Q_{n-1} and λ to zero and solving the resulting system of equations, we obtain the optimal values of Q_0, \dots, Q_{n-1} to mount our attack.

The resulting total work and queries of obtaining n ciphertext using directional failure boosting are given in Table 4 and Figure 6. One can see that the majority of the work lies in obtaining the first ciphertext, and that obtaining more than one ciphertext can be done in less than double the work and queries, or less than one extra bit of complexity. For schemes with a lower failure probability, failing ciphertexts will be more correlated to the secret, so that the directional information is higher and directional failure boosting will be more effective.


In conclusion, the security of a scheme with low failure probability under a single target decryption failure attack can be approximated by the amount of work and queries that an adversary needs to do in order to obtain the first decryption failure. We emphasize the fact that obtaining many failures for a low overhead threatens the security of the scheme (See §5.4).

7 Discussion and variants

7.1 Comparison with D’Anvers et al. [16]

In figure 7, the total work and queries needed to obtain n ciphertexts with probability $1 - e^{-1}$ is plotted for both the traditional failure boosting, and

Fig. 6. Quantum work W_i and number of decryption queries Q_i required to find a new failing ciphertext, given the i failing ciphertexts found previously.



our directional failure boosting approach. For a fair comparison between both results, we adapted the method for estimating the total work and queries with success probability $1 - e^{-1}$ using the traditional failure boosting of [16]. For more information about our method, we refer to the full version of our paper [15].

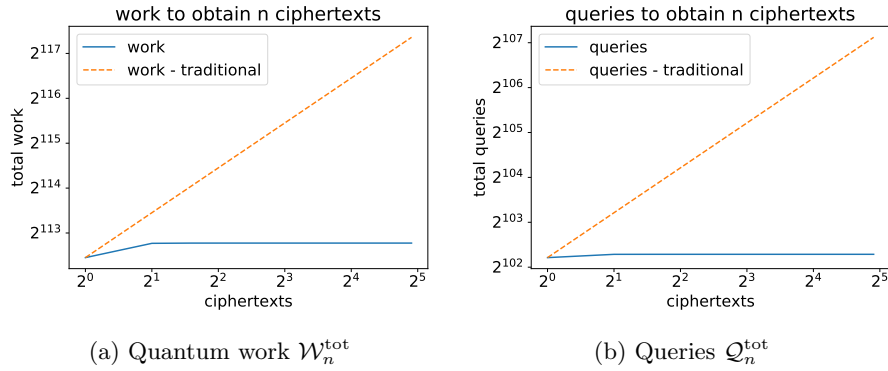



Fig. 7. Quantum work $\mathcal{W}_n^{\text{tot}}$ and number of decryption queries $\mathcal{Q}_n^{\text{tot}}$ required to obtain n failing ciphertexts with probability $1 - e^{-1}$, given the number of previously found failing ciphertexts.

7.2 Minimizing the number of queries instead

In case there is a maximal number of decryption queries is imposed, say 2^K , the same attack strategy can be followed. However, to limit the number of queries $\mathcal{Q}_n^{\text{tot}}$ necessary in the attack, a stronger preprocessing $\sqrt{\alpha_{i,f_t}^{-1}}$ might be necessary to increase the failure probability β_{i,f_t} of weak ciphertexts over 2^{-K} . The only change to accomplish this is selecting the threshold f_t for each i appropriately. Note that for most practical schemes (e.g. Kyber, Saber, New Hope), increasing the failure probability β_{0,f_t} over 2^{-K} is not practically feasible or would require too much preprocessing $\sqrt{\alpha_{0,f_t}^{-1}}$.

Figure 8 depicts the amount of work $\sqrt{\alpha_{i,f_t}^{-1}\beta_{i,f_t}^{-1}}$ needed to increase the failure probability β_{i,f_t} to a certain failure probability (e.g. 2^{-K}) for the parameters given in Table 1. The various curves correspond to different numbers of available failing ciphertexts. From this figure, one can see that also in this case, the work is dominated by finding the first decryption failure. Another observation is that the attack gets much more expensive as the maximal number of decryption queries 2^K gets smaller.

Fig. 8. Quantum work $\mathcal{W}_n^{\text{tot}}$ required to find a new failing ciphertext, as a function of the decryption failure probability of a Mod-LWE scheme.



7.3 Application to ss-ntru-pke and improvement of Guo et al. [28]

In [28], an adaptive multitarget attack is proposed on the ss-ntru-pke version of NTRUEncrypt [10], a Ring-LWE based encryption scheme that claims security against chosen ciphertext attacks. The parameters of this scheme are given in Table 5. The attack performs at most 2^{64} queries on at most 2^{64} targets and has a classical cost of 2^{216} work, and a quantum cost of 2^{140} when speeding up the offline phase with Grover’s search. We adapt directional failure boosting to this attack model and propose both a single and multitarget attack.

For the single target attack, our proposed methodology in subsection 6.3 needs more than 2^{64} queries to obtain a ciphertext. To mitigate this, we increase the precomputational work $\sqrt{\alpha}^{-1}$ so that the failure probability of weak ciphertexts β increases over a certain f_t , which is chosen as 2^{-57} to make sure the total queries are below 2^{64} . The effect is a bigger overall computation, but a reduction in the number of necessary decryption queries. The rest of the attack proceeds as discussed in Subsection 6.3. The work or queries needed to obtain an extra ciphertexts with n ciphertexts can be seen in Figure 9a. The cost of this single target attack is $2^{139.6}$, which is close to the cost of their multitarget attack $2^{139.5}$, as can be seen in Table 6.

	l	N	q	σ_s	σ_e	$P[F]$	Claimed Security
ss-ntru-pke	1	1024	$2^{30} + 2^{13} + 1$	724	724	$> 2^{-80}$	2^{198}

Table 5. Parameters of the ss-ntru-pke [10] scheme.

scheme	claimed security	multitarget attack [28]	our single target attack	our multitarget attack
ss-ntru-pke	2^{198}	$2^{139.5}$	$2^{139.6}$	$2^{96.6}$

Table 6. Comparison of costs for different attacks against ss-ntru-pke [10].

In the multitarget case, we can use a maximum of $2^{64} \cdot 2^{64}$ queries to find the first failing ciphertext, after which we use the methodology of the single target attack to obtain further ciphertext with limited amount of queries. In practice we stay well below the query limit to find the first failure. In this case, the work is dominated by finding the second decryption failure, as we need to do this in under 2^{64} queries. The resulting work to obtain an extra ciphertext is depicted in Figure 9b. The cost of this attack is $2^{96.6}$, which is well below the cost of $2^{139.5}$ reported by Guo et al.

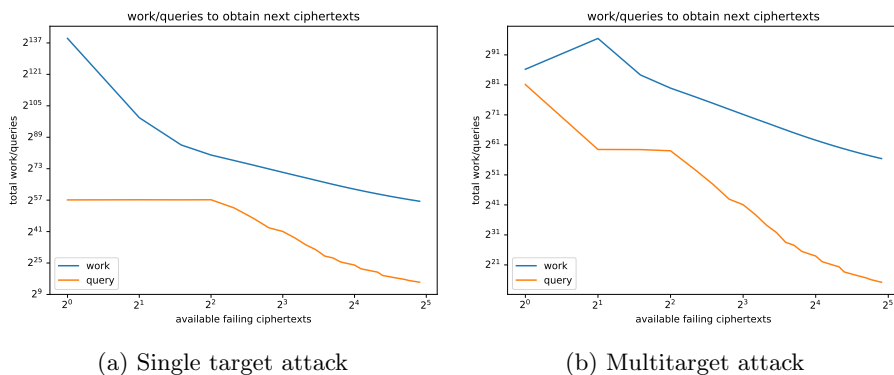


Fig. 9. Quantum work $\mathcal{W}_n^{\text{tot}}$ and number of decryption queries $\mathcal{Q}_n^{\text{tot}}$ required to find a new failing ciphertext for ss-ntru-pke, given the ones found previously.

Acknowledgements. We thank Henri Gilbert and Alessandro Budroni for the interesting discussions about decryption errors, and for providing advice during the writeup of this paper.

References

1. M. R. Albrecht, F. Göpfert, F. Virdia, and T. Wunderer. Revisiting the expected cost of solving uSVP and applications to LWE. pages 297–322.
2. M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of learning with errors. *JMC*, 2015.
3. E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange – a New Hope. In *USENIX Security 2016*, 2016.
4. H. Baan, S. Bhattacharya, O. Garcia-Morchon, R. Rietman, L. Tolhuizen, J.-L. Torre-Arce, and Z. Zhang. Round2: KEM and PKE based on GLWR. IACR ePrint 2017/1183.

5. C. B  etu, F. B. Durak, L. Hugu  enin-Dumittan, A. Talayhan, and S. Vaudenay. Misuse attacks on post-quantum cryptosystems. In *EUROCRYPT*. Springer, 2019.
6. S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*. Springer, 2014.
7. J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, and D. Stehl  . CRYSTALS – Kyber: a CCA-secure module-lattice-based KEM. IACR ePrint 2017/634.
8. J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, and D. Stehl  . CRYSTALS Kyber: a CCA-secure module-lattice-based KEM.
9. T. Cai, J. Fan, and T. Jiang. Distributions of Angles in Random Packing on Spheres. *Journal of machine learning research*, 14(1), jan 2013.
10. C. Chen, J. Hoffstein, W. Whyte, and Z. Zhang. NTRUEncrypt. Technical report, NIST, 2017.
11. Y. Chen and P. Q. Nguyen. Bkz 2.0: Better lattice security estimates. In *ASIACRYPT 2011*. Springer, 2011.
12. J.-P. D’Anvers, Q. Guo, T. Johansson, A. Nilsson, F. Vercauteren, and I. Verbauwhede. Decryption Failure Attacks on IND-CCA Secure Lattice-Based Schemes. In *PKC 2019*, 2019.
13. J.-P. D’Anvers, A. Karmakar, S. S. Roy, and F. Vercauteren. Saber: Module-LWR Based Key Exchange, CPA-Secure Encryption and CCA-Secure KEM. In *AFRICACRYPT 2018*, 2018.
14. J.-P. D’Anvers, A. Karmakar, S. S. Roy, and F. Vercauteren. SABER, Round 2 submission. Technical report, NIST, 2019.
15. J.-P. D’Anvers, M. Rossi, and F. Virdia. (one) failure is not an option: Bootstrapping the search for failures in lattice-based encryption schemes. Cryptology ePrint Archive, Report 2019/1399, 2019. <https://eprint.iacr.org/2019/1399>.
16. J.-P. D’Anvers, F. Vercauteren, and I. Verbauwhede. On the impact of decryption failures on the security of LWE/LWR based schemes. IACR ePrint 2018/1089.
17. J.-P. D’Anvers, F. Vercauteren, and I. Verbauwhede. The impact of error dependencies on Ring/Mod-LWE/LWR based schemes. IACR ePrint 2018/1172.
18. A. W. Dent. A designers guide to kems. In *IMACC*. Springer, 2003.
19. W. Diffie and M. Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 1976.
20. L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehl  . Crystals-dilithium: A lattice-based digital signature scheme. *TCHES*, 2018.
21. T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE transactions on information theory*, 31(4):469–472, 1985.
22. S. Fluhrer. Cryptanalysis of ring-LWE based key exchange with key share reuse. IACR ePrint 2016/085.
23. E. Fujisaki and T. Okamoto. Secure Integration of Asymmetric and Symmetric Encryption Schemes. *Journal of Cryptology*, 26(1), jan 2013.
24. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. 26(1):80–101, Jan. 2013.
25. N. Gama and P. Q. Nguyen. New Chosen-Ciphertext Attacks on NTRU. In T. Okamoto and X. Wang, editors, *Public Key Cryptography – PKC 2007*, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
26. C. Gentry and D. Boneh. *A fully homomorphic encryption scheme*. Stanford University Stanford, 2009.

27. L. K. Grover. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, New York, NY, USA, 1996. ACM.
28. Q. Guo, T. Johansson, and A. Nilsson. A Generic Attack on Lattice-based Schemes using Decryption Errors with Application to ss-ntru-pke. IACR ePrint 2019/043.
29. D. Hofheinz, K. Hövelmanns, and E. Kiltz. A Modular Analysis of the Fujisaki-Okamoto Transformation. 2017.
30. N. Howgrave-Graham, P. Q. Nguyen, D. Pointcheval, J. Proos, J. H. Silverman, A. Singer, and W. Whyte. The Impact of Decryption Failures on the Security of NTRU Encryption. In *CRYPTO 2003*, 2003.
31. É. Jaulmes and A. Joux. A Chosen-Ciphertext Attack against NTRU. In *CRYPTO 2000*, 2000.
32. H. Jiang, Z. Zhang, L. Chen, H. Wang, and Z. Ma. Post-quantum IND-CCA-secure KEM without Additional Hash. IACR ePrint 2017/1096.
33. N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions*. Houghton Mifflin Boston, 1970.
34. J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. Chapman & Hall/CRC, 2014.
35. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3), jun 2015.
36. X. Lu, Y. Liu, D. Jia, H. Xue, J. He, and Z. Zhang. LAC. Technical report, NIST, 2017.
37. V. Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT 2009*. Springer, 2009.
38. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. pages 1–23.
39. D. Micciancio and S. P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO 2003*. Springer, 2003.
40. M. Naehrig, E. Alkim, J. Bos, L. Ducas, K. Easterbrook, B. LaMacchia, P. Longa, I. Mironov, V. Nikolaenko, C. Peikert, A. Raghunathan, and D. Stebila. FrodoKEM. Technical report, NIST, 2017.
41. NIST. Submission requirements and evaluation criteria for the Post-Quantum Cryptography standardization process, 2016.
42. J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier, 2014.
43. C. Peikert and V. Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In *CRYPTO 2008*. Springer, 2008.
44. O. Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *STOC*. ACM, 2005.
45. T. Saito, K. Xagawa, and T. Yamakawa. Tightly-Secure Key-Encapsulation Mechanism in the Quantum Random Oracle Model. IACR ePrint 2017/1005.
46. C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 1994.
47. P. Schwabe, R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, G. Seiler, and D. Stehle. Crystals-Kyber, Round 2 submission. Technical report, NIST, Post-Quantum Standardization Process Round 2, 2019.
48. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. pages 617–635.
49. E. E. Targhi and D. Unruh. Post-Quantum Security of the Fujisaki-Okamoto and OAEP Transforms. In *TCC 2016*. Springer, 2016.