

A Hybrid Lattice Basis Reduction and Quantum Search Attack on LWE

Florian Göpfert², Christine van Vredendaal¹, and Thomas Wunderer²

¹ Department of Mathematics and Computer Science
Technische Universiteit Eindhoven, P.O. Box 513, 5600 MB Eindhoven, NL
`c.v.vredendaal@tue.nl`

² Fachbereich Informatik
Technische Universität Darmstadt, Hochschulstraße, 10, 64289 Darmstadt, DE
`fgoepfert@cdc.informatik.tu-darmstadt.de`,
`twunderer@cdc.informatik.tu-darmstadt.de`

Abstract. Recently, an increasing amount of papers proposing post-quantum schemes also provide concrete parameter sets aiming for concrete post-quantum security levels. Security evaluations of such schemes need to include all possible attacks, in particular those by quantum adversaries. In the case of lattice-based cryptography, currently existing quantum attacks are mainly classical attacks, carried out with quantum basis reduction as subroutine.

In this work, we propose a new quantum attack on the learning with errors (LWE) problem, whose hardness is the foundation for many modern lattice-based cryptographic constructions. Our quantum attack is based on Howgrave-Graham’s Classical Hybrid Attack and is suitable for LWE instances in recent cryptographic proposals. We analyze its runtime complexity and optimize it over all possible choices of the attack parameters. In addition, we analyze the concrete post-quantum security levels of the parameter sets proposed for the New Hope and Frodo key exchange schemes, as well as several instances of the Lindner-Peikert encryption scheme. Our results show that – depending on the assumed basis reduction costs – our Quantum Hybrid Attack either significantly outperforms, or is at least comparable to all other attacks covered by Albrecht–Player–Scott in their work “On the concrete hardness of Learning with Errors”. We further show that our Quantum Hybrid Attack improves upon the Classical Hybrid Attack in the case of LWE with binary error.

Keywords: Public-key encryption, lattice-based cryptography, LWE, quantum attack, hybrid attack.

1 Introduction

Over the past decade *lattice-based cryptography* [31] has proven to be one of the most promising candidates for post-quantum cryptography. One of the reasons for this is the seemingly strong resistance it has shown against quantum attacks.

On top of this lattice-based cryptography has shown a wealth of applications (e.g., [16,15,19,5,33,30,20]). The foundation for many recent lattice-based cryptographic constructions is the *Learning with Errors* (LWE) problem [33,30,32], which is provably as hard as worst-case lattice problems [33,8].

In order to evaluate the concrete post-quantum security levels of LWE-based schemes, cryptanalysts must evaluate the best known algorithms to solve the underlying LWE problem. This evaluation must not only consider classical attacks, but attacks by adversaries with quantum computing power. So far the only existing quantum attacks on LWE are classical attacks where the basis reduction subroutine is replaced by quantum basis reduction.

In this work we present a new quantum attack on LWE: the Quantum Hybrid Attack. The attack is based on Howgrave-Graham’s Classical Hybrid Attack [21], which combines lattice-based techniques such as basis reduction [24,14] with guessing techniques such as brute-force or meet-in-the-middle attack [6] (MitM). In its original form the Classical Hybrid Attack was designed to break the NTRU cryptosystem, but has recently been applied to instances of LWE with highly structured error distributions such as binary or trinary errors [11,37].

From a technical point of view our algorithm replaces the MitM-phase with a generalization of Grover’s quantum search algorithm by Brassard–Hoyer–Mosca–Tapp [9]. The idea to replace this phase by Grover’s search algorithm [17] was sketched in Schanck’s thesis [34], but is in its original form only practical for highly structured (e.g., uniform) NTRU keys. A straightforward application of this idea to LWE Hybrid Attacks would therefore only be practical for LWE with small (uniform) error such as binary or trinary errors. In contrast, our attack is applicable for LWE with arbitrary error distribution and is particularly suitable for LWE instances in recent cryptographic proposals. This is achieved by replacing Grover’s quantum search with the generalization by Brassard–Hoyer–Mosca–Tapp. For example, the time to recover r coefficients of a vector following the New Hope error distribution decreases from $2^{2.52r}$ with Grover’s algorithm to $2^{1.85r}$ with the variant used for the new attack.

We also give a detailed analysis of the Quantum Hybrid Attack and optimize the attack parameters selection. We apply the new attack to the LWE key-exchange schemes New Hope [3] and Frodo [7], and the R-BinLWEenc [10] and Lindner-Peikert [25] encryption schemes, and compare it to the runtimes estimations for existing attacks given by the LWE estimator [2,35]. Depending on the assumed basis reduction costs, our Quantum Hybrid Attack either significantly outperforms, or is at least comparable to all other attacks covered by [2]. We also show that our Quantum Hybrid Attack outperforms the Classical Hybrid Attack in the case of LWE with binary errors.

1.1 Structure of the paper

The remainder of this paper is organized as follows. In Section 2 we introduce lattice definitions and notations in order to explain the Classical Hybrid Attack in Section 3. Then in Section 4 we explain our improved Quantum Hybrid Attack for LWE instances with arbitrary secret and error distributions. We then analyze

the runtime complexity and optimize it over the choice of the attack parameters in Section 5. Lastly, in Section 6 we apply this runtime to concrete parameter sets and compare it to existing attacks using the LWE simulator on common key-exchange and encryption schemes.

Acknowledgement This work has been co-funded by the DFG as part of project P1 within the CRC 1119 CROSSING and supported by the Netherlands Organisation for Scientific Research (NWO) under grant 639.073.005.

2 Preliminaries

We denote vectors by bold lower case (e.g., $\mathbf{a} \in \mathbb{Z}^n$) letters, matrices by bold uppercase letters (e.g., $\mathbf{A} \in \mathbb{Z}^{m \times n}$) and probability distributions by upper case letters (e.g., D). We use the notation \mathbb{Z}_q for the quotient ring $\mathbb{Z}/q\mathbb{Z}$. By $\mathbf{a} \bmod q$ we indicate that each component of the vector is reduced modulo q to lie in the interval $[-\frac{q}{2}, \frac{q}{2})$.

For a probability distribution X , we write $x \stackrel{\$}{\leftarrow} X$ if an element x is sampled according to X . For every element a in the support of X , we write $x_a := \Pr[a = b|b \stackrel{\$}{\leftarrow} X]$. We will specifically refer to the discrete Gaussian distribution D_σ as the distribution such that

$$\forall y \in \mathbb{Z} : \Pr[x = y | x \stackrel{\$}{\leftarrow} D_\sigma] \sim \exp\left(-\frac{y^2}{2\sigma^2}\right).$$

For a probabilistic algorithm \mathcal{A} , $x \stackrel{\$}{\leftarrow} \mathcal{A}$ assigns the outcome of one (random) run of \mathcal{A} to x .

The *Learning with Errors* (LWE) problem was introduced by Regev [33] and has been the foundation of many cryptographic constructions [33,30,32] since.

Definition 1. Let $n, m, q \in \mathbb{Z}$ be positive integers and let D_e be a distribution on \mathbb{Z}^m and D_s be a distribution on \mathbb{Z}^n . Let $\mathbf{s} \stackrel{\$}{\leftarrow} D_s$, \mathbf{A} be chosen uniformly at random from $\mathbb{Z}_q^{m \times n}$, $\mathbf{e} \stackrel{\$}{\leftarrow} D_e$, and $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$. The LWE problem is the problem of recovering \mathbf{s} , given (\mathbf{A}, \mathbf{b}) .

We now review some basic definitions for lattices (for a full survey see e.g., [26]). Throughout this paper, we only consider full-rank lattices. A set $\Lambda \subset \mathbb{R}^m$ is called a *lattice* Λ in \mathbb{R}^m if

$$\Lambda = \Lambda(\mathbf{B}) := \left\{ \mathbf{x} \in \mathbb{R}^m \mid \mathbf{x} = \sum_{i=1}^m \alpha_i \mathbf{b}_i, \text{ with } \alpha_i \in \mathbb{Z} \right\},$$

for some \mathbb{R} -linearly independent set $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_m\} \subset \mathbb{R}^m$. Such a set \mathbf{B} is called a basis of the lattice Λ .

The *determinant* $\det(\Lambda)$ of a lattice Λ is defined as $\det(\Lambda) = |\det(\mathbf{B})|$, where \mathbf{B} is some basis of Λ . This definition is independent of the choice of the basis.

The Hermite delta δ of a basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_m)$ is defined via $\|\mathbf{b}_1\| = \delta^m \det(\Lambda)^{1/m}$.

Lattice-based cryptography is based on a variety of lattice problems that are conjectured to be hard. In the following we list the problems relevant to this work. The *Shortest Vector Problem* (SVP) is to find a shortest non-zero lattice vector, given a basis of the lattice.

The *unique Shortest Vector Problem* (uSVP) is a variant of SVP with the additional promise that the shortest non-zero lattice vector \mathbf{y} is significantly shorter than all other lattice vectors that are not an integral multiple of \mathbf{y} .

The *Bounded Distance Decoding* (BDD) problem is the problem of given a basis of a lattice in \mathbb{R}^m and a target vector $\mathbf{t} \in \mathbb{R}^m$ that is close to a lattice vector \mathbf{v} , find the lattice vector \mathbf{v} . In this work we assume that the task is to find $\mathbf{t} - \mathbf{v}$ instead of \mathbf{v} , which is equivalently hard.

3 The Classical Hybrid Attack

In this section, we recap the approach of solving LWE problems with the Classical Hybrid Attack, see e.g., [11,37]. The first step is to transform the LWE problem instance $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \in \mathbb{Z}_q^m$, into a uSVP instance. Note that a “short” s is necessary for this transformation to be possible. This is however a requirement that is typical for LWE instances. The second step is to then solve the resulting uSVP instance with the Hybrid Attack.

We use the following common approach [27,6,2] to transform an LWE problem into uSVP. Consider the d -dimensional lattice

$$\Lambda = \{\mathbf{x} \in \mathbb{Z}^d : (\mathbf{A}|\mathbf{I}_m| - \mathbf{b})\mathbf{x} = \mathbf{0} \pmod{q}\},$$

where $d = n + m + 1$. With high probability, we have $\det(\Lambda) = q^m$ [27]. Since $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$, the vector $\mathbf{v} = (\mathbf{s}, \mathbf{e}, 1) \in \mathbb{Z}^d$ is a vector in the lattice Λ . Provided \mathbf{v} is sufficiently short (again, this is typical for LWE instances), this leads to a uSVP problem in the lattice Λ . In order to apply the Hybrid Attack to the uSVP instance we need to compute a basis \mathbf{B}' of Λ of the form

$$\mathbf{B}' = \begin{pmatrix} \mathbf{B} & \mathbf{C} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix} \in \mathbb{Z}^{d \times d},$$

for some $r \in \mathbb{N}$ with $r < m, n$. Wunderer [37] showed that with high probability such a basis of Λ exists and can be found efficiently.

The main idea of the Hybrid Attack is to split the short vector \mathbf{v} into two parts $\mathbf{v} = (\mathbf{v}_\ell, \mathbf{v}_g)$ with $\mathbf{v}_\ell \in \mathbb{Z}^{d-r}$ and $\mathbf{v}_g \in \mathbb{Z}^r$. With this notation it holds that

$$\mathbf{v} = \begin{pmatrix} \mathbf{v}_\ell \\ \mathbf{v}_g \end{pmatrix} = \mathbf{B}' \begin{pmatrix} \mathbf{x} \\ \mathbf{v}_g \end{pmatrix} = \begin{pmatrix} \mathbf{B}\mathbf{x} + \mathbf{C}\mathbf{v}_g \\ \mathbf{v}_g \end{pmatrix},$$

for some vector $\mathbf{x} \in \mathbb{Z}^{d-r}$, which implies $\mathbf{v}_\ell - \mathbf{B}\mathbf{x} = \mathbf{C}\mathbf{v}_g$. Note that $\mathbf{B}\mathbf{x}$ is a lattice vector in the lattice spanned by \mathbf{B} , and \mathbf{v}_ℓ is a short vector. Consequently,

if \mathbf{v}_g is known, we can recover \mathbf{v}_ℓ by solving BDD in the lattice spanned by \mathbf{B} with target vector $\mathbf{C}\mathbf{v}_g$. This idea results in the Hybrid Attack: loop through guesses for \mathbf{v}_g and check if a guess is correct (and if so recover \mathbf{v}_ℓ) by solving the corresponding BDD. To solve this problem Babai’s Nearest Plane algorithm [4] is used. Nearest Plane runs in polynomial time, but in order to achieve a high success probability it requires a lattice basis of sufficiently good quality, which in turn has to be generated by an exponential time precomputation step (basis reduction). This makes Nearest Plane suitable for the Hybrid Attack. For every guess of \mathbf{v}_g we have to solve one instance of BDD. Every such instance is a BDD in the same lattice spanned by \mathbf{B} , only with a different target vector. However, the most time-consuming step (basis reduction) of Nearest Plane is independent of the target vector. Consequently, we can precompute a good basis of the lattice spanned by \mathbf{B} *before* looping over the possible guesses of \mathbf{v}_g . Therefore, we can balance the time spent on basis reduction and on guessing values for \mathbf{v}_g to obtain the optimal trade-off.

As was already shown by Howgrave-Graham [21], the guessing part of the attack can be sped up using meet-in-the-middle techniques. However, this approach has three main drawbacks. First, it is only practical for highly structured LWE instances such as LWE with binary or ternary error distribution [37]. Second, its memory requirements are huge [36]. Third, the probability that collisions are actually recognized can be extremely small [37]. In this work, we show how the guessing part can be sped up by using quantum search algorithms. Our Quantum Hybrid Attack, outlined in the next section, eliminates all three drawbacks of the meet-in-the-middle approach and thus enables the Hybrid Attack to handle arbitrary error distributions of LWE.

4 The Quantum Hybrid Attack

We now introduce our new Quantum Hybrid Attack. The main idea is to use quantum search algorithms to speed up the guessing part of the classical Hybrid Attack. This section is structured as follows. We give a brief summary of Grover’s quantum search algorithm and its modified version developed by Brassard–Hoyer–Mosca–Tapp [9] in Section 4.1. In Section 4.2 we show how to use this quantum search algorithm inside the Hybrid Attack to obtain a new Quantum Hybrid Attack.

4.1 Amplitude Amplification

In 1996, Grover presented a quantum algorithm that can speed up the search in unstructured databases [17]. Given a function $f : S \rightarrow \{0, 1\}$ for some finite set S , we call $S_f := \{x \in \{0, 1\}^d \mid f(x) = 1\}$ the set of marked elements. Grover’s algorithm allows to find an element $x \in S_f$ in approximately $\frac{\pi}{4} \cdot \sqrt{|S|/|S_f|}$ evaluations of f (without any further knowledge about f), while classical algorithms require an average number of evaluations in the order of $|S|/|S_f|$.

The runtime of Grover’s search algorithm is independent of how the marked elements have been chosen. The drawback is that additional information about the choice of the marked elements is not used. A generalization of Grover’s search algorithm that can utilize the probability distribution on the search space was presented by Brassard–Hoyer–Mosca–Tapp [9]. Their generalization uses an additional algorithm \mathcal{A} sampling from some distribution on the search space S .

Theorem 1 ([9], Theorem 3). *There exists a quantum algorithm QSearch with the following property. Let \mathcal{A} be any quantum algorithm that uses no measurements (i.e., a unitary transformation), and let $f : S \rightarrow \{0, 1\}$ be any Boolean function. Let a denote the initial success probability of \mathcal{A} (i.e., $a = \Pr[f(x) = 1, x \xleftarrow{\$} \mathcal{A}]$). The algorithm QSearch finds a good solution using an expected number of applications of \mathcal{A} , \mathcal{A}^{-1} and f which are in $\Theta(1/\sqrt{a})$ if $a > 0$, and otherwise runs forever.*

The quantum algorithm \mathcal{A} can be constructed as follows: Given an arbitrary (efficient) probabilistic sampling algorithm \mathcal{S} , it can be easily transformed into a deterministic algorithm that gets random bits as input. This algorithm in turn can be transformed into a quantum algorithm. Instantiating this quantum algorithm with the uniform distribution as superposition for the input bits leads to the wanted algorithm \mathcal{A} .

Note that the complexity of the algorithm is only given asymptotically. This is only necessary because the probability a is unknown. In Appendix A, we show that the hidden constant is indeed small, and we can ignore the Landau notation in our runtime estimates.

4.2 The Attack

We now describe our new Quantum Hybrid Attack (Algorithm 2). We use the notation $\text{NP}_{\mathbf{B}}(\mathbf{t})$ to indicate that Nearest Plane is called on the target vector \mathbf{t} and input basis \mathbf{B} . Inputs for the Quantum Hybrid Attack are an LWE instance $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, the LWE error distribution D_e , and the attack parameters r, δ . The algorithm first transforms the LWE instance into a uSVP instance as described in Section 3 and then runs *QSearch* with the function defined by Algorithm 1.

As we show in Section 5, it is not optimal to use the error distribution for the sampling algorithm \mathcal{A} to find the solution. Instead we use the following transformed distribution.

Definition 2. *Let X be an arbitrary distribution with support S . We write $T(X)$ for the distribution defined by*

$$\forall a \in S : \Pr[a = b | b \xleftarrow{\$} T(X)] = \frac{x_a^{\frac{2}{\delta}}}{\sum_{c \in S} x_c^{\frac{2}{\delta}}}.$$

Our Quantum Hybrid Attack is presented in Algorithm 2. Recall that the attack parameter r indicates the guessing dimension and parameter δ is the Hermite delta used for basis reduction algorithms.

Algorithm 1: Function $f_{\mathbf{A},\mathbf{b},\mathbf{B},\mathbf{C}}(\mathbf{w}_g)$

```
1  $\mathbf{w}_\ell \leftarrow \text{NP}_{\mathbf{B}}(\mathbf{C}\mathbf{w}_g)$ ;  
2 Set  $(\mathbf{s}', \mathbf{e}', 1) = (\mathbf{w}_\ell, \mathbf{w}_g)$ ;  
3 if  $\mathbf{A}\mathbf{s}' + \mathbf{e}' = \mathbf{b}$  and  $\mathbf{s}', \mathbf{e}'$  are small then  
4    $\lfloor$  return 1;  
5 else  
6    $\lfloor$  return 0;
```

Algorithm 2: Quantum Hybrid Attack

Input: LWE instance $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, $\mathbf{b} \in \mathbb{Z}_q^m$, error distribution D_e on \mathbb{Z}^m , attack parameters $\delta \in \mathbb{R}_{>1}$, $r \in \mathbb{N}$, $r < m, n$

```
1 Let  $D$  be the distribution of the last  $r$  entries of the vector  $(\mathbf{x}, 1)$ , where  
    $\mathbf{x} \stackrel{\$}{\leftarrow} D_e$ ;  
2 Set  $\mathcal{A}$  to be a quantum algorithm without measuring for the distribution  $T(D)$ ;  
3 Calculate basis  $\mathbf{B}'$  of lattice  $\Lambda = \{\mathbf{x} \in \mathbb{Z}^{m+n+1} : (\mathbf{A}|\mathbf{I}_m| - \mathbf{b})\mathbf{x} = \mathbf{0} \pmod{q}\}$  of  
   form  $\mathbf{B}' = \begin{pmatrix} \mathbf{B} & \mathbf{C} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix}$ ;  
4 Perform basis reduction to reduce  $\mathbf{B}$  to Hermite delta  $\delta$ ;  
5 Let  $\mathbf{v}'_g$  be the result of  $QSearch$  (Theorem 1) with function  $f_{\mathbf{A},\mathbf{b},\mathbf{B},\mathbf{C}}$   
   (Algorithm 1) and quantum algorithm  $\mathcal{A}$ ;  
6 return  $(\text{NP}_{\mathbf{B}}(\mathbf{v}'_g), \mathbf{v}'_g)$ ;
```

5 Analysis

In this section, we analyze the runtime complexity of the Quantum Hybrid Attack and show how to minimize it over all choices of attack parameters.

5.1 Success Probability and Number of Function Applications

In the following, we show our main result about the runtime of our Quantum Hybrid Attack.

Main Result. For an LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$, let the vectors $\mathbf{v}, \mathbf{v}_\ell, \mathbf{v}_g$, the matrices \mathbf{B}, \mathbf{B}' , the Distribution D , the algorithm \mathcal{A} , and the parameters n, m, q, d, r, δ be defined as in Sections 3 and 4.

The success probability p of the Quantum Hybrid Attack is approximately

$$p \approx \prod_{i=1}^{d-r} \left(1 - \frac{2}{B\left(\frac{(d-r)-1}{2}, \frac{1}{2}\right)} \int_{-1}^{\max(-r_i, -1)} (1-t^2)^{\frac{(d-r)-3}{2}} dt \right),$$

where $B(\cdot, \cdot)$ denotes the Euler beta function (see [29]),

$$r_i = \frac{R_i}{2\|\mathbf{v}_i\|} \quad \text{for all } i \in \{1, \dots, d-r\},$$

and R_1, \dots, R_{m-r} denote the lengths of the Gram-Schmidt basis vectors corresponding to the basis \mathbf{B} .

In case of success, the expected number of applications of f , \mathcal{A} , and \mathcal{A}^{-1} in Algorithm 2 is $\Theta(L)$, where

$$L = \left(\sum_{x \in \text{supp}(D)} d_x^{\frac{2}{3}} \right)^{\frac{3}{2}}.$$

Furthermore, the choice of the distribution for the sampling algorithm \mathcal{A} in Algorithm 4 is optimal.

We first determine the success probability of the attack. We then calculate and optimize the number of applications of f , \mathcal{A} , and \mathcal{A}^{-1} and compare our results with Grover's search algorithm.

Success Probability If $\text{NP}_{\mathbf{B}}(\mathbf{C}\mathbf{v}_g) = \mathbf{v}_\ell$, we have $f_{\mathbf{A}, \mathbf{b}, \mathbf{B}, \mathbf{C}}(\mathbf{v}_g) = 1$ with overwhelming probability and $QSearch$ recovers \mathbf{v}_g . An approximation of the probability that $\text{NP}_{\mathbf{B}}(\mathbf{C}\mathbf{v}_g) = \mathbf{v}_\ell$ is calculated in [11,37] and yields the success probability given in Theorem 5.1. If the components of the LWE error are distributed according to a discrete Gaussian distribution with standard deviation σ , this approximation can be replaced by the simpler (and more efficiently computable) formula

$$\Pr[\text{NP}_{\mathbf{B}}(\mathbf{C}\mathbf{v}_g) = \mathbf{v}_\ell] = \prod_{i=1}^{d-r} \text{erf}\left(\frac{R_i \sqrt{2}}{\sigma}\right) \quad (1)$$

given by Lindner and Peikert [25].

Number of Applications of f , \mathcal{A} , and \mathcal{A}^{-1} We now calculate the expected number of applications of f , \mathcal{A} and \mathcal{A}^{-1} (simply called loops in the following) in the Quantum Hybrid Attack in the case the attack is successful. We show how the choice of the sampling algorithm \mathcal{A} influences the number of loops, how to minimize this number over all possible choices of \mathcal{A} , and that our choice in Algorithm 2 is in fact optimal. In the following, let $S = \text{supp}(D)$ be a finite set. The support S is the search space of our quantum algorithm. Let \mathcal{A} be the initial sampling algorithm used in the Quantum Hybrid Attack and A be the distribution with support S corresponding to \mathcal{A} . According to Theorem 1, for a fixed target element $x \in S$ the expected number of loops in the Quantum Hybrid Attack is roughly $(\sqrt{a_x})^{-1}$. However, since the marked element (and its probability) is not known, we can only estimate the expected number of loops

$$L(A) = L((a_x)_{x \in S}) = \sum_{x \in S} \frac{d_x}{\sqrt{a_x}}. \quad (2)$$

In order to minimize the runtime of the quantum search we must determine the optimal distribution A that minimizes the number of loops $L(A)$. We emphasize

that minimizing the number of loops is of independent interest for any quantum search algorithm based on [9] applied in a similar way as in our attack.

Minimal number of loops. We first minimize the expected number of loops over all possible choices of A . Without loss of generality we assume $S = \{1, \dots, k\}$ for some $k \in \mathbb{N}$. We minimize the expected number of loops by minimizing the function

$$L : (0, 1)^k \rightarrow \mathbb{R}, \quad (a_1, \dots, a_k) \mapsto \sum_{i=1}^k \frac{d_i}{\sqrt{a_i}}, \quad (3)$$

in k variables $a_1, \dots, a_k \in (0, 1)$ under the constraint

$$a_1 + \dots + a_k = 1, \quad (4)$$

where $d_1, \dots, d_k \in (0, 1)$ are fixed. In order to minimize L under the constraints, we define the Lagrange function corresponding to L and Equation (4)

$$\mathcal{L}(\lambda, a_1, \dots, a_k) = \left(\sum_{i=1}^k \frac{d_i}{\sqrt{a_i}} \right) + \lambda \left(-1 + \sum_{i=1}^k a_i \right). \quad (5)$$

To find the minimum of L we need to solve the following set of $k + 1$ equations

$$\begin{aligned} [\text{E}_i]_{i \in \{1, \dots, k\}} \quad & 0 = \mathcal{L}_{a_i}(\lambda, a_1, \dots, a_k) = -\frac{d_i}{2} a_i^{-\frac{3}{2}} + \lambda \\ [\text{E}_c] \quad & a_1 + \dots + a_k = 1, \end{aligned}$$

which gives

$$a_i = \frac{d_i^{\frac{2}{3}}}{\sum_{j=1}^k d_j^{\frac{2}{3}}}. \quad (6)$$

It remains to be shown that choosing the a_i according to Equation (6) leads in fact to a local *minimum* of L under the given constraints. If this is the case, this local minimum must indeed constitute the global minimum satisfying the constraints, since it is the only local minimum and L tends to infinity as one of the a_i approaches zero (hence the problem can be restricted to a compact domain). In order to show that the a_i constitute a local minimum, we compute the determinants of the leading principal minors of the bordered Hessian matrix evaluated in the a_i

$$H = \begin{pmatrix} 0 & 1 & 1 & \dots & 1 \\ 1 & x_1 & 0 & \dots & 0 \\ 1 & 0 & x_2 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 1 & 0 & \dots & 0 & x_k \end{pmatrix}, \quad \text{where } x_i = \frac{3d_i}{4a_i^{2.5}} > 0.$$

For $j \in \{1, \dots, k\}$ let

$$H_j = \begin{pmatrix} 0 & 1 & 1 & \dots & 1 \\ 1 & x_1 & 0 & \dots & 0 \\ 1 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 1 & 0 & \dots & 0 & x_j \end{pmatrix}$$

be the leading principal minors. Using Gaussian elimination we can see the determinants of all but the first principal minors of H are given by

$$\det(H_j) = \det \begin{pmatrix} x_0 & 1 & 1 & \dots & 1 \\ 0 & x_1 & 0 & \dots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ 0 & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \dots & 0 & x_j \end{pmatrix} \quad \text{where } x_0 = - \left(\sum_{i=0}^j \frac{1}{x_i} \right) < 0.$$

Hence all determinants of the leading principal minors of H (except the first one) are negative and thus choosing the a_i according to Equation (6) leads in fact to a local minimum of L under the given constraints. Inserting these a_i into Equation (3) yields the minimal number of loops

$$L_{\min} = \left(\sum_{x \in S} d_{x^{\frac{3}{2}}} \right)^{\frac{2}{3}}. \quad (7)$$

An important special case. While Equation (7) provides a simple formula for the minimal number of loops, evaluating it might be a computationally infeasible task for a large support S . In the following we consider the case that the support is of the form $S = S_0^r$ for some $r \in \mathbb{N}$ and smaller set S_0 and that $D = P^r$ for some distribution P on S_0 . Note that this is the case for most LWE-based cryptosystems, in particular for the ones we analyze in this work. We show how in this case Equation (7) can be evaluated by computing a sum of $|S_0|$ summands and raising it to the r -th power instead of computing a sum of $|S_0|^r$ summands.

This is true since Equation (7) can be rewritten and simplified to

$$\begin{aligned}
L_{\min} &= \left(\sum_{x \in S} d_x^{\frac{2}{3}} \right)^{\frac{3}{2}} = \left(\sum_{y_1 \in S_0} \dots \sum_{y_{r-1} \in S_0} \sum_{y_r \in S_0} \prod_{i=1}^r p_{y_i}^{\frac{2}{3}} \right)^{\frac{3}{2}} = \\
&= \left(\sum_{y_1 \in S_0} \dots \sum_{y_{r-1} \in S_0} \prod_{i=1}^{r-1} p_{y_i}^{\frac{2}{3}} \left(\sum_{y_r \in S_0} p_{y_r}^{\frac{2}{3}} \right) \right)^{\frac{3}{2}} = \\
&= \left(\sum_{y_1 \in S_0} \dots \sum_{y_{r-1} \in S_0} \prod_{i=1}^{r-1} p_{y_i}^{\frac{2}{3}} \left(\sum_{y \in S_0} p_y^{\frac{2}{3}} \right) \right)^{\frac{3}{2}} = \\
&= \dots = \left(\left(\sum_{y \in S_0} p_y^{\frac{2}{3}} \right)^r \right)^{\frac{3}{2}}, \tag{8}
\end{aligned}$$

since each of the d_x is exactly the product of r of the p_y .

Comparison with Grover's search algorithm. If in our Quantum Hybrid Attack the distribution D is the uniform distribution, then its complexity matches the one of Grover's search algorithm

$$L_{\min} = \left(\sum_{x \in S} d_x^{\frac{2}{3}} \right)^{\frac{3}{2}} = \left(\sum_{x \in S} \left(\frac{1}{|S|} \right)^{\frac{2}{3}} \right)^{\frac{3}{2}} = \left(|S| \frac{1}{|S|^{\frac{2}{3}}} \right)^{\frac{3}{2}} = \sqrt{|S|}.$$

For a structured search space, QSearch (see Theorem 1) gives a much better complexity. As an example we examine the distribution D on the set $S = \{-16, \dots, 16\}^r$ used in New Hope [3]. Then $|S| = 33^r$ and using Grover's search algorithm inside the Quantum Hybrid Attack would yield a complexity of

$$L_{\text{grover}} = \sqrt{33^r} \approx 2^{2.52r}.$$

In comparison, our Quantum Hybrid Attack only has complexity

$$L_{\text{our}} = \left(\left(\sum_{i=0}^{32} p_i^{\frac{2}{3}} \right)^r \right)^{\frac{3}{2}} \approx 2^{1.85r}, \quad \text{where } p_i = \binom{32}{i} \cdot 2^{-32}.$$

For $r = 200$ entries that are guessed during the Quantum Hybrid Attack this amounts to a speedup factor of 2^{134} of our approach over using Grover's algorithm inside the Hybrid Attack. This example showcases the significant improvement of our Quantum Hybrid Attack over one that is simply using Grover's search algorithm. It also demonstrates that our new Quantum Hybrid Attack opens the possibility to apply the Hybrid Attack to larger, non-uniform search spaces.

5.2 Total Runtime of the Quantum Hybrid Attack

In this section we estimate the total runtime of the Quantum Hybrid Attack by estimating the individual cost of one application of f , \mathcal{A} , and \mathcal{A}^{-1} , the pre-computation (i.e., basis reduction) cost, and combining the results with the ones of Section 5.1. The resulting runtime formula must then be optimized over all possible attack parameters.

Cost of f , \mathcal{A} , and \mathcal{A}^{-1} . The cost of the function f is dominated by the cost of one Nearest Plane call, which was experimentally found to be roughly $k^2/2^{1.06}$ bit operations, where k is the dimension of the lattice (in our case $k = d - r$), see [18].³ We assume that compared to this cost, the cost of the algorithm \mathcal{A} and \mathcal{A}^{-1} can be neglected.

Basis Reduction Cost. In the following we examine the precomputation cost of basis reduction (BKZ) to achieve a BKZ-reduced basis of quality δ . According to [13,2], the minimal block size b needed by BKZ to achieve a certain Hermite delta δ can be determined via the relation

$$\delta = (((\pi b)^{1/b} b) / (2\pi e))^{1/(2(b-1))}.$$

We assume that the number of tours t with block size b is as given in the LWE estimator [2,35]. Then the precomputation cost of BKZ with block size b in the $(d - r)$ -dimensional lattice is roughly $T_{\text{red}} = t(d - r)T_{\text{SVP}_b}$, where T_{SVP_b} denotes the number of operations to solve SVP in dimension b .

The main two methods to solve SVP are *enumeration* and *sieving*. Asymptotically the runtime of sieving outperforms the one of enumeration, but the cross-over point is unknown (see e.g. the discussion in [22]). However, sieving algorithms require access to exponentially large memory (while enumeration only requires polynomial memory), which could turn out to be the limiting factor in high dimension, especially when it comes to quantum-sieving. In this work we compare the Quantum Hybrid Attack with existing attacks under two different basis reduction assumptions.

In the first model (called *quantum-sieving*) we assume that memory consumption is not a problem, sieving scales as predicted to higher dimensions and can be sped up with quantum computers as proposed by Laarhoven–Mosca–van de Pol [23] with a runtime complexity of

$$T_{\text{SVP}_b} = 2^{0.265b+16.4}.$$

³ In [18], Hirschhorn, Hoffstein, Howgrave-Graham and Whyte conservatively assume that if one has to perform multiple Nearest Plane calls with the same lattice basis (as it is the case in the Quantum Hybrid Attack), one can reduce this cost to $k/2^{1.06}$ bit operations using precomputation. However, since this speedup has not been confirmed in practice, we do not assume this linear cost for our runtime estimates. Note that assuming the linear cost instead of the quadratic one would lower the runtime of the Quantum Hybrid Attack.

From an attacker’s point of view this is an optimistic prediction, so the number derived in this model can be seen as a lower bound on the hardness of the LWE instances.

The second model (*enumeration*) assumes that sieving is not practical compared to enumeration for dimensions of cryptographic size and uses an interpolation of Albrecht–Player–Scott [2] based on runtimes of enumeration given by Chen and Nguyen [14] instead. The predicted number of operations necessary to solve SVP in dimension b is given by

$$T_{\text{SVP}_b} = 2^{0.27b \ln(b) - 1.019b + 16.10}.$$

This methodology leads to higher runtime estimations (for both the existing attacks and the Quantum Hybrid Attack).

Total Cost and Runtime Optimization. Using these estimates we obtain that the total runtime of the Quantum Hybrid Attack is given by

$$T_{\text{total}} = \frac{T_{\text{red}} + T_{\text{hyb}}}{p},$$

where

$$T_{\text{hyb}} = \left(\sum_{x \in S} d_x^{\frac{2}{3}} \right)^{\frac{3}{2}} \cdot (d - r)^2 / 2^{1.06},$$

T_{red} is the runtime of basis reduction, and p is the success probability as given in the Main Result.

The total runtime of the attack T_{total} depends on the attack parameters r and δ and must therefore be optimized over all such choices.

6 Results

In this section, we present concrete runtime estimates of our Quantum Hybrid Attack against the New Hope [3] and Frodo [7] key exchange schemes (Section 6.1) and the Lindner-Peikert [25] (Section 6.2) and R-BinLWEenc [10] (Section 6.3) encryption schemes. For the comparison, we always selected the maximal number of LWE samples available for the Quantum Hybrid Attack.

6.1 New Hope and Frodo

We analyze and optimize the runtime of the Quantum Hybrid Attack against the New Hope [3] and Frodo [7] key exchange schemes and compare our results to the security levels produced by the LWE estimator for LWE instances with limited number of samples [35]. Note that the LWE estimator handles LWE instances with Gaussian distribution, while the distributions of New Hope and Frodo are only approximations of such. Therefore, for the LWE estimator we use the Gaussian distributions that are approximated. To obtain a fair comparison,

we also use the approximated Gaussian distributions to determine the success probabilities according to Equation 1.

Table 1 shows that the Quantum Hybrid is significantly faster if enumeration is used as basis reduction subroutine. If we assume that quantum-sieving is practical and behaves as predicted [23], the Quantum Hybrid and the existing attacks are comparable (see Table 2).

Attack	New Hope	Frodo-592	Frodo-752	Frodo-864
Dual Embedding	1346	446	485	618
Decoding	833	—	—	—
Quantum Hybrid	725	254	310	377

Table 1: Quantum security estimates for New Hope and Frodo using *enumeration* as SVP oracle. Table shows the base-two logarithm of the expected runtimes.

Attack	New Hope	Frodo-592	Frodo-752	Frodo-864
Dual Embedding	389	173	184	219
Decoding	380	—	—	—
Quantum Hybrid	384	171	189	221

Table 2: Quantum security estimates for New Hope and Frodo using *quantum-sieving* as SVP oracle. Table shows the base-two logarithm of the expected runtimes.

Note that for both, the Quantum Hybrid Attack and the LWE estimator, the results differ substantially from the claimed security levels of the schemes [3,7]. This, is not surprising, since in the spirit of guaranteeing secure post-quantum parameter sets, [3] and [7] aim for highly conservative security estimates.

6.2 Lindner-Peikert

In 2011, Lindner and Peikert [25] introduced an LWE-based encryption scheme. The authors give four concrete parameter sets for various security levels. Later, Albrecht–Cabarcas–Fitzpatrick–Göpfert–Schneider [1] interpolated those sets to give an asymptotic instantiation.

We analyze such instances for dimensions ranging from 256 to 1024. Note that theoretically, the discrete Gaussian distributions used in the Lindner-Peikert encryption scheme have infinite support, while our analysis requires finite support. Using a standard tailbound argument [25] one can show that with overwhelming probability the absolute value of D_σ is bounded by 14σ . We therefore assume the distributions D_σ have finite support $\{-\lceil 14\sigma \rceil, \dots, \lceil 14\sigma \rceil\}$.

As Figure 1 shows, the Quantum Hybrid Attack outperforms all existing attacks for *enumeration* as SVP oracle. Again, the gap between the attacks nearly vanishes when *quantum-sieving* is used. However, the Quantum Hybrid Attack seems to benefit from a slightly better asymptotic complexity (see Figure 2). The exact hardness values are given in Appendix B.

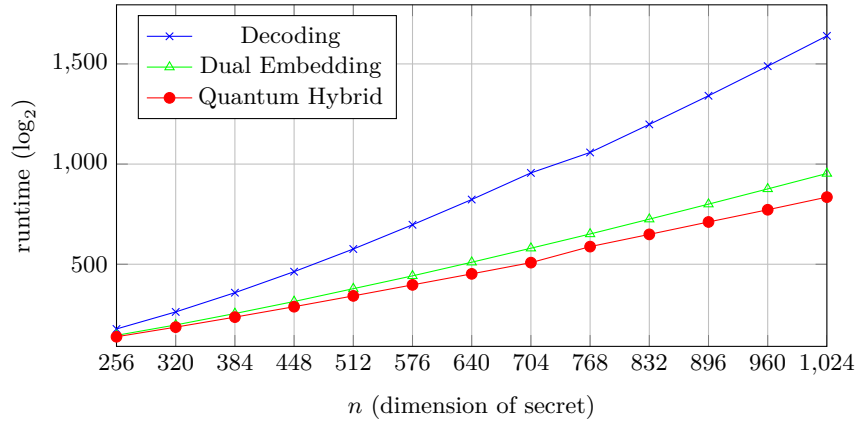


Fig. 1: Quantum security estimates for Lindner-Peikert parameter using *enumeration* as SVP oracle. Figure shows the base-two logarithm of the expected runtimes.




Fig. 2: Quantum security estimates for Lindner-Peikert parameter using *quantum-sieving* as SVP oracle. Figure shows the base-two logarithm of the expected runtimes.

6.3 R-BinLWEenc

So far, all instances considered either use Gaussian errors or approximations of such. However, the Classical Hybrid Attack is most efficient on LWE with binary error [11,37]. In order to compare the Classical and the Quantum Hybrid Attack, we investigate the hardness of LWE instances with binary error as used by Buchmann–Göpfert–Player–Wunderer [10] for their Lindner–Peikert-like encryption scheme.

The runtime of the Classical Hybrid Attack on binary LWE instances was estimated by Wunderer [37]. The author provides security over- and underestimates of the attack. For our comparison we use the security overestimates, since their underlying assumptions match the ones in this work.

Attack	Set-I	Set-II	Set-III
Classical Hybrid estimates (<i>enumeration</i> SVP oracle)	99	90	197
Quantum Hybrid (<i>enumeration</i> SVP oracle)	82	75	167
Quantum Hybrid (<i>quantum-sieving</i> SVP oracle)	79	73	140

Table 3: Quantum security estimates for R-BinLWEenc. Table shows the base-two logarithm of the expected runtimes.

References

1. Martin R. Albrecht, Daniel Cabarcas, Robert Fitzpatrick, Florian Göpfert, and Michael Schneider. A generator for lwe and ring-lwe instances. IACR archive, 2013. <https://www.iacr.org/news/files/2013-04-29lwe-generator.pdf>.
2. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *J. Mathematical Cryptology*, 9(3):169–203, 2015.
3. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. NEWHOPE without reconciliation, 2016. <http://cryptojedi.org/papers/#newhopesimple>.
4. László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
5. 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 - The Cryptographer’s Track at the RSA Conference 2014, San Francisco, CA, USA, February 25-28, 2014. Proceedings*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47. Springer, 2014.
6. Shi Bai and Steven D. Galbraith. Lattice decoding attacks on binary LWE. In Willy Susilo and Yi Mu, editors, *Information Security and Privacy - 19th Australasian Conference, ACISP 2014, Wollongong, NSW, Australia, July 7-9, 2014. Proceedings*, volume 8544 of *Lecture Notes in Computer Science*, pages 322–337. Springer, 2014.
7. Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring!

- practical, quantum-secure key exchange from LWE. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 1006–1018. ACM, 2016.
8. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584. ACM, 2013.
 9. Gilles Brassard, P. Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information: A Millennium Volume*, volume 305 of *AMS Contemporary Mathematics Series*, pages 53–74. American Mathematical Society, 2002. Earlier version in arxiv:quant-ph/0005055.
 10. Johannes A. Buchmann, Florian Göpfert, Tim Güneysu, Tobias Oder, and Thomas Pöppelmann. High-performance and lightweight lattice-based public-key encryption. In Richard Chow and Gökay Saldamli, editors, *Proceedings of the 2nd ACM International Workshop on IoT Privacy, Trust, and Security, CPSSAsiaCCS, Xi'an, China, May 30 - June 3, 2016*, pages 2–9. ACM, 2016.
 11. Johannes A. Buchmann, Florian Göpfert, Rachel Player, and Thomas Wunderer. On the hardness of LWE with binary error: Revisiting the hybrid lattice-reduction and meet-in-the-middle attack. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *Progress in Cryptology - AFRICACRYPT 2016 - 8th International Conference on Cryptology in Africa, Fes, Morocco, April 13-15, 2016, Proceedings*, volume 9646 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2016.
 12. Ran Canetti and Juan A. Garay, editors. *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*. Springer, 2013.
 13. Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complet-homomorphe*. PhD thesis, ENS-Lyon, France, 2013.
 14. Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, volume 7073 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2011.
 15. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In Canetti and Garay [12], pages 40–56.
 16. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Mitzenmacher [28], pages 169–178.
 17. Lov 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*, pages 212–219, New York, NY, USA, 1996. ACM.
 18. Philip S. Hirschhorn, Jeffrey Hoffstein, Nick Howgrave-Graham, and William Whyte. Choosing NTRUEncrypt parameters in light of combined lattice reduction and MITM approaches. In Michel Abdalla, David Pointcheval, Pierre-Alain Fouque, and Damien Vergnaud, editors, *Applied Cryptography and Network Security, 7th International Conference, ACNS 2009, Paris-Rocquencourt, France, June*

- 2-5, 2009. *Proceedings*, volume 5536 of *Lecture Notes in Computer Science*, pages 437–455, 2009.
19. Jeffrey Hoffstein, Jill Pipher, John M. Schanck, Joseph H. Silverman, and William Whyte. Practical signatures from the partial fourier recovery problem. In Ioana Boureanu, Philippe Owesarski, and Serge Vaudenay, editors, *Applied Cryptography and Network Security - 12th International Conference, ACNS 2014, Lausanne, Switzerland, June 10-13, 2014. Proceedings*, volume 8479 of *Lecture Notes in Computer Science*, pages 476–493. Springer, 2014.
 20. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In Joe Buhler, editor, *Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.
 21. Nick Howgrave-Graham. A hybrid lattice-reduction and meet-in-the-middle attack against NTRU. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, volume 4622 of *Lecture Notes in Computer Science*, pages 150–169. Springer, 2007.
 22. Thijs Laarhoven. *Sieving for Shortest Vectors in Lattices Using Angular Locality-Sensitive Hashing*, pages 3–22. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.
 23. Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Finding shortest lattice vectors faster using quantum search. *Designs, Codes and Cryptography*, 77(2):375–400, 2015.
 24. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.
 25. Richard Lindner and Chris Peikert. Better key sizes (and attacks) for lwe-based encryption. In Aggelos Kiayias, editor, *Topics in Cryptology - CT-RSA 2011 - The Cryptographers' Track at the RSA Conference 2011, San Francisco, CA, USA, February 14-18, 2011. Proceedings*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.
 26. Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, March 2002.
 27. Daniele Micciancio and Oded Regev. Lattice-based cryptography. In Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen, editors, *Post-Quantum Cryptography*, pages 147–191. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
 28. Michael Mitzenmacher, editor. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*. ACM, 2009.
 29. Frank WJ Olver. *NIST handbook of mathematical functions*. Cambridge University Press, 2010.
 30. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Mitzenmacher [28], pages 333–342.
 31. Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016.
 32. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. *SIAM J. Comput.*, 40(6):1803–1844, 2011.

33. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.
34. John M. Schanck. *Practical Lattice Cryptosystems: NTRUEncrypt and NTRUMLS*. PhD thesis, University of Waterloo, 2015.
35. Markus Schmidt and Nina Bindel. Estimation of the hardness of the learning with errors problem with a restricted number of samples. Cryptology ePrint Archive, Report 2017/140, 2017. <http://eprint.iacr.org/2017/140>.
36. Christine van Vredendaal. Reduced memory meet-in-the-middle attack against the NTRU private key. *IACR Cryptology ePrint Archive*, 2016:177, 2016.
37. Thomas Wunderer. Revisiting the hybrid attack: Improved analysis and refined security estimates. *IACR Cryptology ePrint Archive*, 2016:733, 2016.

A About the constant in Theorem 1

Brassard–Hoyer–Mosca–Tapp [9] give two different results about amplitude amplification: one for known probability a , and one if a is unknown. One disadvantage of the result about amplification with unknown a is that it is an asymptotic result (see Theorem 1). Such results give a way to group algorithms into complexity classes, but are of limited value for runtime estimations on concrete instances, since the constant factor is unknown. In this section, we show that the hidden constant factor of Theorem 1 is small.

In the analysis of their algorithm with known a , Brassard–Hoyer–Mosca–Tapp show that the success probability of their quantum amplification algorithm after m rounds is given by $p = \sin^2((2m + 1)\theta_a)$ with θ_a such that $\sin^2(\theta_a) = a$.

Our goal in this section is to produce an algorithm that succeeds at least with $p = 1/2$. This leads to

$$\begin{aligned}
 p \geq \frac{1}{2} &\Leftrightarrow \sin((2m + 1)\theta_a) \geq \frac{1}{\sqrt{2}} \\
 &\Leftrightarrow \frac{1}{4}\pi \leq (2m + 1)\theta_a \leq \frac{3}{4}\pi \\
 &\Leftrightarrow \frac{\pi}{4(2m + 1)} \leq \theta_a \leq \frac{3\pi}{4(2m + 1)} \\
 &\Leftrightarrow \sin^2\left(\frac{\pi}{4(2m + 1)}\right) \leq a \leq \sin^2\left(\frac{3\pi}{4(2m + 1)}\right)
 \end{aligned}$$

Since m is big in our applications, we can approximate the bounds by

$$a \in \left[\frac{\pi^2}{64m^2}, \frac{9\pi^2}{64m^2} \right] \tag{9}$$

Assume we know that $a \in [b_{min}, b_{max}]$. In the following, we find a sequence of rounds m_0, \dots, m_k such that $[b_{min}, b_{max}] \subseteq \bigcup_i \left[\frac{\pi^2}{64m_i^2}, \frac{9\pi^2}{64m_i^2} \right]$. Given this sequence, we can find a solution as follows. We start with running the algorithm

for m_0 rounds. If this succeeds, we found a solution. If not, we run the algorithm for m_1 rounds, and so on. After the last run (with m_k rounds) at least one of the algorithm calls had a success probability of at least $1/2$, so the overall success probability is at least $1/2$.

To find the sequence of m_i , we start with selecting m_0 such that $\frac{9\pi^2}{64m_0^2} = b_{max}$, which is equivalent to $m_0 = \frac{3\pi}{8\sqrt{b_{max}}}$. The other m_i are then defined iteratively by selecting m_{i+1} such that $\frac{9\pi^2}{64m_{i+1}^2} = \frac{\pi^2}{64m_i^2}$, which is equivalent to $m_{i+1} = 3m_i$, which in turn leads directly to $m_i = 3^{i+1} \frac{\pi}{8\sqrt{b_{max}}}$. The second condition of our sequence is that $\frac{\pi^2}{64m_k^2} \leq b_{min}$. A simple calculation shows that this is equivalent to $3^{2k+2} \geq \frac{b_{max}}{b_{min}}$. Finally, we take a look at the special when a is distributed according to a Gaussian distribution. By the definition of the Gaussian distribution, we have

$$\Pr[D_\sigma = x] = c \exp\left(-\frac{x^2}{2\sigma^2}\right),$$

which leads directly to $b_{min} = c$. It is common knowledge that with overwhelming probability, only elements smaller than 14σ get sampled, so we set

$$b_{max} = c \exp\left(-\frac{(14\sigma)^2}{2\sigma^2}\right) = c \exp(-98).$$

Consequently, we require $3^{2k+2} \geq \frac{c}{c \exp(-98)} = \exp(98)$, which is satisfied for $k \geq 45$.

B Hardness Tables for Lindner/Peikert LWE

n	256	320	384	448	512	576	640	704	768	832	896	960	1024
Dual Embedding	177	262	358	463	576	697	823	956	1058	1198	1341	1489	1640
Decoding	145	197	254	314	378	442	510	580	651	725	800	876	953
Quantum Hybrid	138	186	236	288	342	397	452	508	588	649	711	772	835

Table 4: Quantum security estimates for Lindner-Peikert parameter using *enumeration* as SVP oracle. Table shows the base-two logarithm of the expected runtimes.

n	256	320	384	448	512	576	640	704	768	832	896	960	1024
Dual Embedding	92	120	149	178	208	238	269	301	325	356	388	420	453
Decoding	101	127	155	182	210	239	266	295	324	353	383	412	442
Quantum Hybrid	108	134	161	187	214	242	269	297	326	355	383	412	441

Table 5: Quantum security estimates for Lindner-Peikert parameter using *quantum-sieving* as SVP oracle. Table shows the base-two logarithm of the expected runtimes.